MITSUBISHI ELECTRIC RESEARCH LABORATORIES
https://www.merl.com

Abort-safe spacecraft rendezvous under stochastic actuation and navigation uncertainty

Vinod, Abraham P.; Weiss, Avishai; Di Cairano, Stefano

TR2021-148 January 11, 2022

Abstract

We propose a tractable approach to generate abort-safe trajectories for safe spacecraft rendezvous that guarantees safety (the spacecraft does not enter a keep-out set defined around the rendezvous target), despite process and measurement noise, and the possibility of partial propulsion failure. We use a combination of stochastic reachability, computational geometry, and optimization to synthesize a nominal rendezvous trajectory and its associated controller. The designed trajectory is such that safe recovery is also guaranteed with high likelihood in the event of a partial propulsion failure. The recovery controllers can be computed only when needed using offline pre-computation, thereby reducing the online computational effort. Numerical experiments show the efficacy of the proposed approach.

IEEE Annual Conference on Decision and Control (CDC) 2021

Mitsubishi Electric Research Laboratories, Inc.
201 Broadway, Cambridge, Massachusetts 02139

Abort-safe spacecraft rendezvous under stochastic actuation and navigation uncertainty

Abraham P. Vinod*, Avishai Weiss, and Stefano Di Cairano

***Abstract*— We propose a tractable approach to generate abort-safe trajectories for spacecraft rendezvous that guarantees safety, i.e., the spacecraft does not enter a keep-out set defined around the rendezvous target. We guarantee safety of the rendezvous trajectory even in the event of propulsion failure and in the presence of stochastic uncertainty in actuation and navigation. We use a combination of stochastic reachability, computational geometry, and optimization to synthesize a nominal rendezvous trajectory and its associated controller. The designed trajectory is such that safe recovery, in the event of a propulsion failure, is guaranteed with pre-specified, sufficiently high probability. The recovery controllers are available when needed via an offline pre-computation, which significantly reduces the online computational effort. Numerical experiments show the efficacy of the proposed approach.**

I. Introduction

We consider spacecraft rendezvous maneuvers where the approaching spacecraft, also known as the deputy, must reach the target spacecraft, also known as the chief. We require that the deputy can guarantee safety (stay outside a keep out set defined around the chief) in the event of actuation failure (Figure 1), *despite the stochastic uncertainties arising from propulsion mismatch as well as navigational uncertainty*. Our prior effort considered the problem of generating such abort-safe rendezvous trajectories in the absence of uncertainties [1]. In this paper, we combine stochastic reachability and optimization to generate a nominal rendezvous trajectory with an associated controller that has a high likelihood of safety, in the presence of stochastic uncertainties. We also guarantee that safe abort (recovery) is possible at every time step with high likelihood using an off-nominal controller that accommodates uncertainty.

Traditionally, a deputy spacecraft performs a predetermined active *collision avoidance maneuver* if it deviates significantly from its nominal rendezvous approach and its current trajectory is no longer safe in proximity to the chief [2]. However, such a maneuver may not always be possible depending upon the current trajectory and the available actuation. The authors in [3] tackle this issue by generating nominal and abort sequences concurrently at every iteration, which is computationally expensive. Alternatively, researchers have used continuous-time and discrete-time backward reachability to verify desired behaviors in the rendezvous trajectories when no uncertainty is present [1], [4]–[6]. In the presence of uncertainty, researchers have turned to stochastic viability to ensure that line-of-sight constraints are satisfied with high likelihood [7]. However, the computation of the stochastic reachable set is in general computationally expensive and does not scale well to high dimensional systems. Moreover, existing work in stochastic reachability typically assumes access to accurate state information [7], [8]. To enforce the safe abort constraint under uncertain state measurements, we present a tractable inner-approximation of a *stochastic viability set with state measurement uncertainty*, which is the set of state measurements from which safe recovery of the deputy is guaranteed in the event of actuation failure with high likelihood despite uncertainty. These sets also enable the synthesis of the safe recovery controller on-demand.

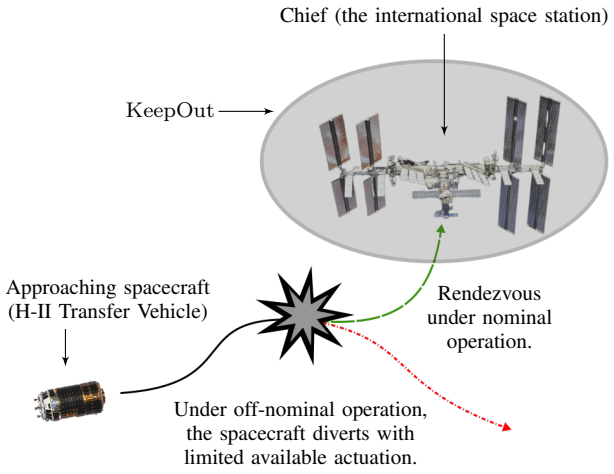

Fig. 1. Illustration of the abort-safe rendezvous problem. We design rendezvous trajectories that are safe, both in nominal and off-nominal operation, despite stochastic uncertainty in actuation uncertainty and navigation. In off-nominal operation, we only have access to a limited actuation, possibly from a redundant thruster. Picture courtesy: NASA and JAXA.

The main contribution of this paper is a *tractable approach for abort-safe spacecraft rendezvous trajectory optimization under uncertainty in actuation and navigation using stochastic reachability*. The proposed approach synthesizes a rendezvous trajectory that approaches the chief, while staying outside a pre-specified keep-out set with high likelihood, including the scenarios of propulsion failures. We utilize offline stochastic reachability computations to synthesize a set of safe state measurements from which recovery under off-nominal operation is possible, despite the uncertainties and the resulting limited actuation. We cast the original stochastic trajectory optimization with reachability-based constraints as

A. P. Vinod, A. Weiss, and S. Di Cairano are with Mitsubishi Electric Research Laboratories, Cambridge, MA 02139, USA. Emails: {vinod,weiss,dicairano}@merl.com
*Corresponding author.

a mixed-integer optimization problem, which we solve in a receding horizon framework in practice. In addition, we discuss a convexification of the optimization problem to enable real-time control at the expense of additional conservativeness. We demonstrate the efficacy of the proposed approach using numerical simulations.

II. Preliminaries and problem formulation

We employ the following notation throughout the paper: The interval $\mathbb{N}_{[a,b]}$ enumerates all natural numbers between and including $a, b \in \mathbb{N}$. Random vectors are denoted in bold, $\|\cdot\|$ denotes Euclidean distance, and $\mathrm{dist}(x, \mathcal{S})$ denotes the Euclidean distance between a point $x \in \mathbb{R}^n$ and a set $\mathcal{S}$.

A. Problem formulation

Consider a chief in a circular Keplerian orbit around Earth with mean motion ω, and a deputy in proximity to the chief. The deputy is assumed to be a rigid body such that all control forces act on its center of mass, the chief is uncontrolled, and orbital perturbations are neglected. We consider in-plane motion of the deputy relative to the chief in Hill's frame with radial, x, and along-track, y, components. The linearized dynamics of the deputy relative to chief are given by the Clohessy-Wiltshire (CW) equations [2]

$$\ddot{x} - 3\omega^2 x - 2\omega\dot{y} = \frac{F_x}{m_\mathrm{d}}, \quad \ddot{y} + 2\omega\dot{x} = \frac{F_y}{m_\mathrm{d}}, \tag{1}$$

where $[x, y]^\top \in \mathbb{R}^2$ is the relative position of the deputy resolved in Hill's frame, m_d is the mass of the deputy, and $[F_x, F_y]^\top \in \mathbb{R}^2$ is the control force applied to the deputy resolved in Hill's frame. We discretize (1) in time with sampling period Δt and model actuation mismatch as a stochastic uncertainty to obtain a discrete-time LTI system

$$\boldsymbol{x}_{t+1} = A\boldsymbol{x}_t + B(u_t + \boldsymbol{w}_t), \tag{2}$$

where $\boldsymbol{x}_t = [x, y, \dot{x}, \dot{y}]^\top \in \mathcal{X} = \mathbb{R}^4$ is the deputy's state, $u_t = [F_x, F_y]^\top \in \mathcal{U} \subset \mathbb{R}^2$ is the input, $\boldsymbol{w}_t \in \mathcal{W} = \mathbb{R}^2$ is the error due to actuation mismatch, and A and B are appropriately defined matrices. We assume that the input set $\mathcal{U}$ is convex and compact, and the process noise $\boldsymbol{w}_t$ is an independent and identically distributed Gaussian vector $\boldsymbol{w}_t \sim \mathcal{N}(\mu_w, \Sigma_w)$ with $\mu_w \in \mathbb{R}^2$ and $\Sigma_w \in \mathbb{R}^{2\times 2}$.

Measurement model: We model navigational uncertainty as noisy state measurements $y_k \in \mathcal{Y} = \mathbb{R}^4$. We assume y_k are samples of a random vector $\boldsymbol{y}_k$ at every time step k,

$$\boldsymbol{y}_k = x_k + \boldsymbol{\gamma}_k. \tag{3}$$

Here, x_k is the unknown true state of the deputy relative to the chief, and $\boldsymbol{\gamma}_k \sim \mathcal{N}(\mu_\gamma, \Sigma_\gamma)$ is an independent and identically distributed Gaussian noise that models the state measurement uncertainty with known mean vector $\mu_\gamma \in \mathbb{R}^4$ and known covariance matrix $\Sigma_\gamma \in \mathbb{R}^{4\times 4}$.

Effect of actuation: Next, we consider the effect of executing an open-loop control sequence $U \in \mathcal{U}^N$ over a *planning horizon* $N \in \mathbb{N}$. Due to the linearity of the system (2) and the Gaussianity of the stochastic process and measurement noises, the future state $\boldsymbol{x}_{t|k}$ is a Gaussian random vector for all $t \in \mathbb{N}_{[k,k+N]}$,

$$\boldsymbol{x}_{t|k} \sim \mathcal{N}(\mu_{t|k}, \Sigma_{t|k}), \tag{4a}$$
$$\mu_{t|k} = A^{t-k}\mu_{k|k} + \mathscr{C}_u(t,k)(U + \mu_W), \tag{4b}$$
$$\Sigma_{t|k} = A^{t-k}\Sigma_{k|k}\left(A^{t-k}\right)^\top + \mathscr{C}_u(t,k)\Sigma_W(\mathscr{C}_u(t,k))^\top \tag{4c}$$
$$\boldsymbol{y}_{t|k} \sim \mathcal{N}(\nu_{t|k}, \Gamma_{t|k}),\ \nu_{t|k} = \mu_{t|k} + \mu_\gamma,\ \Gamma_{t|k} = \Sigma_{t|k} + \Sigma_\gamma \tag{4d}$$

Here, $\boldsymbol{W} \sim \mathcal{N}(\mu_W, \Sigma_W)$ is the concatenated disturbance random vector for the process noise, and $\mathscr{C}_u(t,k)$ is the controllability matrix for (2) of appropriate dimensions.

Safety definition: We define $\mathrm{KeepOut}$ as a convex and compact set in (x, y)-coordinates around the chief located at the origin. The safety of the mission requires the following:

1) *Safety under nominal operation*: We require the state $\boldsymbol{x}_{t|k} \notin \mathrm{KeepOut}$ at all time steps $t \in \mathbb{N}_{[k,k+N]}$.
2) *Safety under off-nominal operation*: In the event of actuation failure at some failure time $T \in \mathbb{N}_{[k,k+N]}$, the input space is restricted to a convex and compact subset $\mathcal{V} \subset \mathcal{U}$ for all future time $t \geq T$. We require the state $\boldsymbol{x}_{t|k} \notin \mathrm{KeepOut}$ at all time steps $t \in \mathbb{N}_{[T,T+M]}$ for a longer *safety horizon* $M \in \mathbb{N}$, with $M > N$.

Due to the stochastic nature of the dynamics (2), we seek probabilistic safety guarantees. For some user-specified safety probability $\alpha \in (0, 1]$, we require that the probability of violating either of the safety requirements is no larger than $1-\alpha$. We refer to the off-nominal safety requirement with an empty set $\mathcal{V}$ as *passive safety*, and refer to the off-nominal safety requirement as *active safety* otherwise [1].

Cost definition: For a planning horizon N, we seek an open-loop controller that approaches the set $\mathrm{KeepOut}$ safely with minimal control effort. We denote the mission cost $J : \mathcal{X}^N \times \mathcal{U}^N \to \mathbb{R}$ as a quadratic function,

$$J(\mu_{k+1|k}, \ldots, \mu_{k+N|k}, U) = \sum_{t\in\mathbb{N}_{[k+1,k+N]}} \mathrm{dist}(\mu_{t|k}, \mathrm{KeepOut})^2 + \lambda\|U\|_2^2. \tag{5}$$

Here, $\lambda \geq 0$ trades-off the control effort with the proximity to $\mathrm{KeepOut}$.

Optimal control problem: At each step k, given the current state measurement y_k, solve

$$\min.\ \text{Cost } J \text{ as defined in (5)} \tag{6a}$$
$$\text{s. t. Dynamics (4) for the state } \boldsymbol{x}_{t|k} \text{ with } U \in \mathcal{U}^N, \tag{6b}$$
$$\mathbb{P}\left\{\forall t \in \mathbb{N}_{[k+1,k+N]}, \boldsymbol{x}_{t|k} \notin \mathrm{KeepOut}\right\} \geq \alpha_\mathrm{nom}, \tag{6c}$$
$$\begin{cases} \forall \text{ potential failure time } T \in \mathbb{N}_{[k,k+N-1]}, \\ \mathbb{P}\left\{\forall t \in \mathbb{N}_{[T,T+M]},\ \boldsymbol{x}^\mathrm{fail}_{t|T} \notin \mathrm{KeepOut}\right\} \geq \alpha_\mathrm{off} \\ \text{with state } \boldsymbol{x}^\mathrm{fail}_{t|T} \text{ for dynamics (4) under limited} \\ \text{actuation } \mathcal{V} \subseteq \mathcal{U} \text{ with } \boldsymbol{x}^\mathrm{fail}_{T|T} \text{ initialized to } \boldsymbol{x}_{T|k}. \end{cases} \tag{6d}$$

Here, $\alpha_\mathrm{nom}, \alpha_\mathrm{off} \in (0, 1]$ are probability thresholds selected to ensure that the safety objectives are met with a likelihood no smaller than α. In practice, we solve (6) in the receding horizon control framework, which motivates the need for computationally efficient solutions.

Problem 1. *Design a tractable chance constrained optimization formulation of* (6) *for* α*-safe rendezvous, in the presence of stochastic uncertainty in actuation and navigation, and even in presence of actuation failure.*

Motivated by our previous work [1], we approach Problem 1 using stochastic reachability theory to account for stochastic uncertainty in actuation and navigation.

Problem 1.a. *Use stochastic reachability to enforce* (6d) *without an explicit computation of the controllers needed for the recovery from a potential actuation failure.*

We generalize recent results in stochastic viability set computation [7] to the computation of a *stochastic viability set with state measurement uncertainty*, the set of safe state measurements from which recovery using noisy state measurements and limited actuation is possible with sufficiently high probability. These sets provide a tractable approach to enforce (6d).

III. Main results

The key component of the proposed approach is the construction of *stochastic viability sets under state measurement uncertainty*. We use these sets to cast (6) as a chance constrained optimization problem, and to avoid the need for the synthesis of off-nominal state-measurement-feedback controllers online. We solve the resulting optimization problem via mixed-integer quadratic programming, propose a convexification approach that yields a feasible (possibly suboptimal) solution to (6) via a convex quadratic program, and discuss the conservativeness introduced in the proposed approach for the sake of tractability.

A. Safety under actuation failure and state measurement uncertainty via stochastic viability sets

We first define the *stochastic viability set under state measurement uncertainty* and discuss a tractable inner-approximation. The proposed definition extends existing definition of stochastic viability sets (see [8]) to problems where state information is available *with measurement uncertainty*.

Let $\rho : \mathcal{Y} \to \mathcal{V}$ denote any Borel-measurable state-measurement-feedback control law. For a safety horizon M, let $\pi \triangleq (\rho_0, \rho_1, \ldots, \rho_{M-1}) \in \mathcal{M}$ denote a (possibly time-varying) state-measurement-feedback control policy, and let $\{\boldsymbol{x}_k^{\pi,y_0}\}_{k=0}^M$ denote the random process corresponding to the stochastic model (2), (3) under the influence of π and an initial state measurement y_0.

Definition 1 (Stochastic viability set with state measurement uncertainty)**.** *We define the stochastic viability set with state measurement uncertainty* $\mathcal{L}(\mathcal{S},\beta)$ *for the stochastic dynamics* (2) *and* (3)*, a safe set* $\mathcal{S} \subseteq \mathbb{R}^n$*, a safety horizon* $M \in \mathbb{N}$*, and a safety probability* $\beta \in (0,1]$ *as the set of all initial state measurements* y_0 *that satisfies*

$$\mathcal{L}(\mathcal{S},\beta) = \{y_0 | \exists \pi \in \mathcal{M},\ \mathbb{P}\{\forall k \in \mathbb{N}_{[0,M]}, \boldsymbol{x}_k^{\pi,y_0} \in \mathcal{S}\} \geq \beta\}$$

Intuitively, $\mathcal{L}(\mathcal{S},\beta)$ are the set of initial measurements y_0 from which an state-measurement-feedback control policy π exists such that the underlying stochastic state process $\{\boldsymbol{x}_k^{\pi,y_0}\}_{k=0}^M$ stays within $\mathcal{S}$ with a likelihood no smaller than β. However, the exact computation of $\mathcal{L}(\mathcal{S},\beta)$ is expected to be hard, since even for the case of perfect state knowledge, the state-of-the-art relies on grid-based dynamic programming [8]. Such methods suffer from the curse of dimensionality, and can not reliably solve systems with dimension $n \geq 3$ without incurring significant approximations [9].

We can utilize Definition 1, (4d), and a sufficiently high $\alpha_{\text{out}} \in (0,1]$ to cast (6d) as a chance constraint,

$$\forall t \in \mathbb{N}_{[k,k+N]},\ \mathbb{P}\left\{\boldsymbol{y}_{t|k} \in \mathcal{L}(\mathcal{X} \setminus \text{KeepOut}, \alpha_{\text{off}})\right\} \geq \alpha_{\text{out}}. \quad (7)$$

Consequently, we require an inner-approximation of $\mathcal{L}(\mathcal{S},\beta)$ to remain conservative in the right direction.

Next, we introduce the notion of robust control invariance for the non-stochastic dynamics with state measurement uncertainty to propose an inner-approximation of $\mathcal{L}(\mathcal{S},\beta)$. In contrast to Definition 1, robust control invariance seeks safety guarantees, despite any realization of bounded disturbances.

Definition 2 (Robust control invariance under state measurement uncertainty)**.** *For some bounded process noise set* $\mathcal{E}_{\text{act}}$ *and measurement noise set* $\mathcal{E}_{\text{meas}}$*, consider*

$$x_{t+1} = Ax_t + B(u_t + w_t),\ y_t = x_t + \gamma_t, \quad (8)$$

where no stochastic information is available on $w_t \in \mathcal{E}_{\text{act}} \subset \mathbb{R}^m$ *and* $\gamma_t \in \mathcal{E}_{\text{meas}} \subset \mathbb{R}^n$*. Let* $M \in \mathbb{N}$ *be the safety horizon, and* $\mathcal{S} \subset \mathcal{X}$ *be the safe set of interest. The robust control invariant set with state measurement uncertainty* $\mathcal{O}_k(\mathcal{E}_{\text{act}}, \mathcal{E}_{\text{meas}}, \mathcal{S})$ *is defined recursively for* $k \in \mathbb{N}_{[0,M-1]}$*,*

1) safety of the current state*: for every measurement noise* $\gamma_k \in \mathcal{E}_{\text{meas}}$ *and every state measurement* $y_k \in \mathcal{O}_k(\mathcal{E}_{\text{act}}, \mathcal{E}_{\text{meas}}, \mathcal{S})$*, the associated state is safe, i.e.,* $x_k = y_k - \gamma_k \in \mathcal{S}$*, and*
2) safety of the next state and the state measurement*: for every state measurement* $y_k \in \mathcal{O}_k(\mathcal{E}_{\text{act}}, \mathcal{E}_{\text{meas}}, \mathcal{S})$*, there exists an input* $u_k \in \mathcal{V}$ *such that for every process noise* $w_k \in \mathcal{E}_{\text{act}}$ *and every measurement noise* $\gamma_k \in \mathcal{E}_{\text{meas}}$*, the resulting next state is safe, i.e.,* $x_{k+1} = A(y_k - \gamma_k) + B(u_k + w_k) \in \mathcal{S}$*. Additionally, the system can be rendered safe from every possible next state measurement,* $y_{k+1} = x_{k+1} + \gamma_{k+1} \in \mathcal{O}_{k+1}(\mathcal{E}_{\text{act}}, \mathcal{E}_{\text{meas}}, \mathcal{S})$ *for every measurement noise* $\gamma_{k+1} \in \mathcal{E}_{\text{meas}}$*.*

We define $\mathcal{O}_M(\mathcal{E}_{\text{act}}, \mathcal{E}_{\text{meas}}, \mathcal{S}) = \mathcal{S} \ominus \mathcal{E}_{\text{meas}}$ *for the safety of terminal state* x_M*, since no control occurs after* $k = M$*.*

By Definition 2, a system that starts from a robust control invariant set with state measurement uncertainty $y_k \in \mathcal{O}_k(\mathcal{E}_{\text{act}}, \mathcal{E}_{\text{meas}}, \mathcal{S})$ at time k can be driven safely, $x_{t|k} \in \mathcal{S}$ for every $t \in \mathbb{N}_{[k,k+M]}$, using a state-measurement-based feedback policy, despite the bounded disturbance values. Lemma 1 follows from standard computational geometry arguments, similar to robust control invariant sets [10].

Lemma 1 (Set-based dynamic programming recursion)**.** *For any bounded sets* $\mathcal{E}_{\text{act}}$ *and* $\mathcal{E}_{\text{meas}}$*, the*

sets $\{\mathcal{O}_k(\mathcal{E}_{\rm act}, \mathcal{E}_{\rm meas}, \mathcal{S})\}_{k=0}^{M}$ *can be obtained from the following recursion for* $k \in \mathbb{N}_{[0,M-1]}$*, initialized with* $\mathcal{O}_M(\mathcal{E}_{\rm act}, \mathcal{E}_{\rm meas}, \mathcal{S}) = \mathcal{S} \ominus \mathcal{E}_{\rm meas}$*,*

$$\mathcal{P}_k(\mathcal{E}_{\rm act}, \mathcal{E}_{\rm meas}, \mathcal{S}) = \mathcal{O}_{k+1}(\mathcal{E}_{\rm act}, \mathcal{E}_{\rm meas}, \mathcal{S}) \ominus (-\mathcal{E}_{\rm meas}) \ominus (B\mathcal{E}_{\rm act}) \ominus (-A\mathcal{E}_{\rm meas}) \quad (9a)$$

$$\mathcal{O}_k(\mathcal{E}_{\rm act}, \mathcal{E}_{\rm meas}, \mathcal{S}) = A^{-1}(\mathcal{P}_k(\mathcal{E}_{\rm act}, \mathcal{E}_{\rm meas}, \mathcal{S}) \oplus (-B\mathcal{V})) \cap (\mathcal{S} \ominus (-\mathcal{E}_{\rm meas})). \quad (9b)$$

Theorem 1 (INNER-APPROXIMATION OF STOCHASTIC VIABILITY SETS WITH STATE MEASUREMENT UNCERTAINTY)**.** *For any* $\alpha_{\rm act}, \alpha_{\rm meas} \in (0,1]$*, let* $\mathcal{E}_{\rm act}$ *and* $\mathcal{E}_{\rm meas}$ *be bounded sets such that*

$$\mathbb{P}\{\boldsymbol{w}_t \in \mathcal{E}_{\rm act}\} \geq \alpha_{\rm act}^{\frac{1}{M}}, \text{ and } \mathbb{P}\{\boldsymbol{\gamma}_t \in \mathcal{E}_{\rm meas}\} \geq \alpha_{\rm meas}^{\frac{1}{M}}. \quad (10)$$

Let the current time be $k \in \mathbb{N}$*. Then, for all failure times* $T \in \mathbb{N}_{[k,k+N]}$ *and the corresponding state measurement* $y_{T|k} \in \mathcal{O}_0(\mathcal{E}_{\rm act}, \mathcal{E}_{\rm meas}, \mathcal{X} \setminus \text{KeepOut})$*,*

$$\mathbb{P}\left\{\forall t \in \mathbb{N}_{[T,T+M]},\ \boldsymbol{x}_{t|T}^{\rm fail} \notin \text{KeepOut}\right\} \geq \alpha_{\rm act}\alpha_{\rm meas}, \quad (11)$$

where $\boldsymbol{x}_{t|T}^{\rm fail}$ *corresponds to the state at time* $t \geq T$ *under limited actuation* $\mathcal{V} \subseteq \mathcal{U}$ *and dynamics* (2)*. In other words,* $\mathcal{O}_0(\mathcal{E}_{\rm act}, \mathcal{E}_{\rm meas}, \mathcal{X} \setminus \text{KeepOut}) \subseteq \mathcal{L}(\mathcal{X} \setminus \text{KeepOut}, \alpha_{\rm act}\alpha_{\rm meas})$.

We omit the proof of Theorem 1 due to space constraints. The key insight used in the proof is the law of total probability and the robust state constraint satisfaction offered by Definition 2. Based on Theorem 1, $\mathcal{O}_0(\mathcal{E}_{\rm act}, \mathcal{E}_{\rm meas}, \mathcal{S})$ is the set of state measurements from which safety can be guaranteed with high likelihood for the stochastic dynamical system (2) and measurement noise model (3). For any choice of $\alpha_{\rm act}$ and $\alpha_{\rm meas}$, the safety guarantee holds with a likelihood no smaller than $\alpha_{\rm act}\alpha_{\rm meas}$ over the safety horizon, despite actuation failure (future control is limited to $u_t \in \mathcal{V} \subseteq \mathcal{U}$), stochastic actuation uncertainty $\boldsymbol{w}_t$, and stochastic measurement uncertainty $\boldsymbol{\gamma}_t$.

In light of Theorem 1, every feasible solution of

$$\begin{aligned}
&\text{min. Cost } J \text{ as defined in (5)}\\
&\text{s. t. Dynamics (4) for the state } \boldsymbol{x}_{t|k} \text{ and measurement } \boldsymbol{y}_{t|k} \text{ under } U \in \mathcal{U}^N, && (12a)\\
&\mathbb{P}\left\{\forall t \in \mathbb{N}_{[k+1,k+N]},\ \boldsymbol{x}_{t|k} \notin \text{KeepOut}\right\} \geq \alpha_{\rm nom}, && (12b)\\
&\mathbb{P}\left\{\boldsymbol{y}_{t|k} \in \mathcal{O}_0(\mathcal{E}_{\rm act}, \mathcal{E}_{\rm meas}, \mathcal{X} \setminus \text{KeepOut})\right\} \geq \alpha_{\rm out}, \quad \forall t \in \mathbb{N}_{[k+1,k+N]}. && (12c)
\end{aligned}$$

satisfies (6), where (6d) is relaxed to (7). Thus, (12) addresses Problem 1.a, for sufficiently high $\alpha_{\rm out}$.

Off-nominal recovery controller synthesis: If a failure indeed occurs at some time $T \in \mathbb{N}_{[k,k+N]}$, then the execution of a control action, drawn from the following set at every subsequent time step $T+t$ with $t \in \mathbb{N}_{[0,M]}$, ensures the desired level of probabilistic safety, $\mathcal{V}_{t,\rm eff} = \{u \in \mathcal{V} | Bu \in \mathcal{P}_{t+1}(\mathcal{E}_{\rm act}, \mathcal{E}_{\rm meas}, \mathcal{X} \setminus \text{KeepOut}) \oplus \{-Ay_{t+T|T}\}\}$. The effective control set $\mathcal{V}_{t,\rm eff}$, defined for $t \in \mathbb{N}_{[0,M]}$, is parameterized by the current state measurement $y_{t+T|T}$, the dynamics (2), the limited control authority $\mathcal{V}$, the sets $\mathcal{E}_{\rm act}$ and $\mathcal{E}_{\rm meas}$, and the robust control invariant sets with state measurement uncertainty $\{\mathcal{O}_t(\mathcal{E}_{\rm act}, \mathcal{E}_{\rm meas}, \mathcal{X} \setminus \text{KeepOut})\}_{t=0}^{M}$. Thus, the use of (12c) instead of (6d) avoids the need for synthesizing off-nominal controllers, while ensuring that a safe recovery controller can be synthesized on-demand.

B. Tractable abort-safe motion planning under uncertainty

Next, we discuss tractable approaches to enforce the chance constraints (12b) and (12c) for polytopic keep-out sets. Motivated by existing work in stochastic obstacle avoidance [11], we assume that the set KeepOut is a polytope, with $\text{KeepOut} = \cap_{i \in \mathbb{N}_{[1,|\text{KeepOut}|]}} \text{KeepOut}_i = \cap_{i \in \mathbb{N}_{[1,|\text{KeepOut}|]}} \{x | a_i^\top x \leq b_i\}$. Consequently, we utilize Boole's inequality, quantile reformulation, and disjunctive constraint enforcement [12] to propose a conservative but tractable enforcement of the chance constraints (12b)–(12c).

1) Chance constraints for nominal operation (12b)*:* From Boole's inequality, a sufficient condition for (12b) is

$$\mathbb{P}\{\boldsymbol{x}_{t|k} \in \text{KeepOut}\} \leq \frac{1-\alpha_{\rm nom}}{N}, \ \forall t \in \mathbb{N}_{[k,k+N]}. \quad (13)$$

Since $\mathbb{P}\{\cap \mathcal{A}_i\} \leq \min_i \mathbb{P}\{\mathcal{A}_i\}$ for any finite collection of sets $\mathcal{A}_i$, the following collection of disjunctive chance constraints is sufficient to satisfy (13) (hence (12b)),

$$\forall t \in \mathbb{N}_{[k,k+N]},\ \exists i \in \mathbb{N}_{[1,|\text{KeepOut}|]},\ \mathbb{P}\{\boldsymbol{x}_{t|k} \in \text{KeepOut}_i\} \leq \frac{1-\alpha_{\rm nom}}{N}. \quad (14)$$

Since $\boldsymbol{x}_{t|k}$ is Gaussian, $\mathbb{P}\{\boldsymbol{x}_{t|k} \in \text{KeepOut}_i\} = \Phi\left(\frac{b_i - a_i^\top \mu_{t|k}}{\sqrt{a_i^\top \Sigma_{t|k} a_i}}\right)$, where $\mu_{t|k}$ and $\Sigma_{t|k}$ are given by (4), and Φ is the standard normal cumulative distribution. Using the quantile reformulation [13] and disjunctive constraint enforcement [12], we arrive at the following collection of mixed-integer linear constraints that conservatively enforces (12b) using auxillary binary variables $\delta_{i,t}^{\rm nom}$ and sufficiently large constants $\kappa_{i,t}$,

$$a_i^\top \mu_{t|k} + \sqrt{a_i^\top \Sigma_{t|k} a_i}\,\Phi^{-1}\left(\frac{1-\alpha_{\rm nom}}{N}\right) \geq b_i - (1-\delta_{i,t}^{\rm nom})\kappa_{i,t}, \quad (15a)$$

$$\sum\nolimits_{i \in \mathbb{N}_{[1,|\text{KeepOut}|]}} \delta_{i,t}^{\rm nom} \geq 1, \quad (15b)$$

for every $t \in \mathbb{N}_{[k,k+N]}$ and every $i \in \mathbb{N}_{[1,|\text{KeepOut}|]}$.

2) Chance constraints for off-nominal operation (12c)*:* For safety under the off-nominal operation, we simplify (12c) into a tractable collection of disjunctive chance constraints, which we enforce similarly to (15).

Proposition 1. *For a polytopic* KeepOut*,* $\mathbb{P}\left\{\boldsymbol{y}_{t|k} \in \mathcal{O}_0(\mathcal{E}_{\rm act}, \mathcal{E}_{\rm meas}, \mathcal{X} \setminus \text{KeepOut})\right\} \geq \max_{i \in \mathbb{N}_{[1,|\text{KeepOut}|]}} \mathbb{P}\left\{\boldsymbol{y}_{t|k} \in \mathcal{O}_0(\mathcal{E}_{\rm act}, \mathcal{E}_{\rm meas}, \mathcal{X} \setminus \text{KeepOut}_i)\right\}$.

Proposition 1 follows from the observation that $\mathbb{P}\{\cup_i \mathcal{A}_i\} \geq \max_i \mathbb{P}\{\mathcal{A}_i\}$ for any finite collection of sets $\mathcal{A}_i$. Then, a nominal open-loop control sequence U

satisfies (12b), if

$$\forall t \in \mathbb{N}_{[k,k+N]},\ \exists i \in \mathbb{N}_{[1,|\mathrm{KeepOut}|]},$$
$$\mathbb{P}\left\{\boldsymbol{y}_{t|k} \in \mathcal{O}_0(\mathcal{E}_{\mathrm{act}}, \mathcal{E}_{\mathrm{meas}}, \mathcal{X} \setminus \mathrm{KeepOut}_i)\right\} \geq \alpha_{\mathrm{out}}. \quad (16)$$

In contrast to (12c), (16) simplifies the necessary computation of the robust control invariant sets with state measurement uncertainty required to solve (12). The sets $\mathcal{O}_0(\mathcal{E}_{\mathrm{act}}, \mathcal{E}_{\mathrm{meas}}, \mathcal{X} \setminus \mathrm{KeepOut}_i)$ in (16) are easier-to-compute via (9) using existing computational geometry tools [14]. In contrast, the set $\mathcal{O}_0(\mathcal{E}_{\mathrm{act}}, \mathcal{E}_{\mathrm{meas}}, \mathcal{X} \setminus \mathrm{KeepOut})$ in (12c) is hard-to-compute via (9), since the set $\mathcal{X} \setminus \mathrm{KeepOut}$ is non-convex.

For each hyperplane (a_i, b_i) of the KeepOut with $i \in \mathbb{N}_{[1,|\mathrm{KeepOut}|]}$, it is easy to show that the sets $\mathcal{O}_0(\mathcal{E}_{\mathrm{act}}, \mathcal{E}_{\mathrm{meas}}, \mathcal{X} \setminus \mathrm{KeepOut}_i)$ are polytopes for a polytopic input set $\mathcal{V}$. We denote the set $\mathcal{O}_0(\mathcal{E}_{\mathrm{act}}, \mathcal{E}_{\mathrm{meas}}, \mathcal{X} \setminus \mathrm{KeepOut}_i) \triangleq \cap_{j \in \mathbb{N}_{[1,L_i]}} \{p_{ij}^\top y \leq q_{ij}\}$, where $L_i = |\mathcal{O}_0(\mathcal{E}_{\mathrm{act}}, \mathcal{E}_{\mathrm{meas}}, \mathcal{X} \setminus \mathrm{KeepOut}_i)| \in \mathbb{N}$, $p_{ij} \in \mathbb{R}^n$, and $q_{ij} \in \mathbb{R}$. Using Boole's inequality, the following disjunctive constraints are sufficient to satisfy (16) (hence (12c)),

$$\forall t \in \mathbb{N}_{[k,k+N]},\ \exists i \in \mathbb{N}_{[1,|\mathrm{KeepOut}|]},\ \forall j \in \mathbb{N}_{[1,L_i]},$$
$$\mathbb{P}\{p_{ij}^\top \boldsymbol{y}_{t|k} > q_{ij}\} \leq \frac{1-\alpha_{\mathrm{out}}}{L_i}, \quad (17)$$

where $\boldsymbol{y}_{t|k} \sim \mathcal{N}(\nu_{t|k}, \Gamma_{t|k})$ by (4d). Similarly to (15), we use mixed-integer linear constraints with auxiliary binary variables $\delta_{i,t}^{\mathrm{off}}$ to conservatively enforce (12c),

$$p_{ij}^\top \nu_{t|k} + \sqrt{p_{ij}^\top \Gamma_{t|k} p_{ij}}\,\Phi^{-1}\left(1 - \frac{1-\alpha_{\mathrm{out}}}{N}\right) \leq q_{ij} + (1-\delta_{i,t}^{\mathrm{off}})\kappa'_{i,t}, \quad (18a)$$
$$\sum\nolimits_{i \in \mathbb{N}_{[1,|\mathrm{KeepOut}|]}} \delta_{i,t}^{\mathrm{off}} \geq 1, \quad (18b)$$

for every $t \in \mathbb{N}_{[k,k+N]}$, every $i \in \mathbb{N}_{[1,|\mathrm{KeepOut}|]}$, and every $j \in \mathbb{N}_{[1,L_i]}$. Here, $\kappa'_{i,t} > 0$ are sufficiently large constants.

We describe the complete mixed-integer optimization problem as follows,

$$\begin{aligned} \min_{U, \delta_{i,t}^{\mathrm{nom}}, \delta_{i,t}^{\mathrm{off}}} \quad & \text{Cost } J \text{ as defined in (5)} \\ \text{s. t.} \quad & \delta_{i,t}^{\mathrm{off}}, \delta_{i,t}^{\mathrm{nom}} \in \{0,1\},\ \mu_{t|k}, \nu_{t|k} \text{ via affine} \\ & \text{transformation (4b) and (4d) of } U \in \mathcal{U}^N, \\ & \text{Mixed-integer linear constraints (15)} \\ & \text{and (18) in } U,\ \delta_{i,t}^{\mathrm{nom}}, \text{ and } \delta_{i,t}^{\mathrm{off}}. \end{aligned} \quad (19)$$

for every $t \in \mathbb{N}_{[k,k+N]}$ and $i \in \mathbb{N}_{[1,|\mathrm{KeepOut}|]}$. We select the parameters $\alpha_{\mathrm{nom}}, \alpha_{\mathrm{off}}, \alpha_{\mathrm{out}}, \alpha_{\mathrm{act}}, \alpha_{\mathrm{meas}} \in (0,1]$ such that the overall safety probability is no smaller than the user-specified threshold $\alpha \in (0,1]$ as follows. Recall that Theorem 1 requires $\alpha_{\mathrm{act}}\alpha_{\mathrm{meas}} \geq \alpha_{\mathrm{off}}$.

Proposition 2. *For any user-specified safety probability threshold $\alpha \in (0,1]$, any nominal open-loop controller that solves* (19) *with $\alpha_{\mathrm{nom}} \in [\alpha, 1]$, $\alpha_{\mathrm{out}} \in (0,1]$, $\alpha_{\mathrm{off}} \in (0,1]$, and $\alpha_{\mathrm{nom}} + \alpha_{\mathrm{out}}\alpha_{\mathrm{off}} \geq 1 - \alpha$, guarantees that all safety objectives are met with probability no smaller than α.*

Proposition 2 follows from the law of total probability, Theorem 1 and (12b). Proposition 2 describes a balance of safety violation risks between the nominal and off-nominal scenarios. By requiring high α_{nom}, we obtain more conservative nominal trajectories, while admitting more risky maneuvers during the potential off-nominal operation. On the other hand, reducing α_{nom} increases the lower bound on $\alpha_{\mathrm{off}}\alpha_{\mathrm{out}}$ in Proposition 2, which, in turn, restricts the regions of the state measurement space $\mathcal{Y}$ from which safe recovery is possible in the event of actuation failure. For a user-specified safety probability $\alpha = 0.9$, we can use $\alpha_{\mathrm{nom}} = 0.95$, $\alpha_{\mathrm{off}} = 0.96$, $\alpha_{\mathrm{out}} = 0.99$, and $\alpha_{\mathrm{act}} = \alpha_{\mathrm{meas}} = 0.98$.

Convexification via scheduling: While commercial solvers have brought in tremendous improvements in solving mixed-integer programs, we may need to compute a suboptimal solution for (19) to satisfy real-time requirements, or increase the control horizon to accommodate the computation time for solving mixed-integer programs online. Alternatively, we may satisfy real-time requirements by convexifying (19). Specifically, we select a feasible assignment of binary variables $\delta_{i,t}^{\mathrm{nom}}$ and $\delta_{i,t}^{\mathrm{off}}$, which renders the trajectory optimization problem (19) convex. Intuitively, the assignment of the binary variables *schedules* the particular halfspace of KeepOut that we must avoid at every time step. A simple strategy is to expand the KeepOut set to one of the hyperplanes $\mathrm{KeepOut}_{i^\dagger}$ for some user-specified $i^\dagger \in \mathbb{N}_{[1,|\mathrm{KeepOut}|]}$, resulting in a convex quadratic problem,

$$\begin{aligned} \underset{U}{\text{minimize}} \quad & \sum_{t \in \mathbb{N}_{[k+1,k+N]}} \mathrm{dist}(\mu_{t|k}, \mathrm{KeepOut}_i)^2 + \lambda \|U\|_2^2 \\ \text{subject to} \quad & \mu_{t|k}, \nu_{t|k} \text{ via affine transformation of } U \in \mathcal{U}^N, \\ \forall t \in \mathbb{N}_{[k,k+N]}, \quad & a_{i^\dagger}^\top \mu_{t|k} + \sqrt{a_{i^\dagger}^\top \Sigma_{t|k} a_{i^\dagger}}\,\Phi^{-1}\left(\tfrac{1-\alpha_{\mathrm{nom}}}{N}\right) \geq b_{i^\dagger}, \\ \forall t \in \mathbb{N}_{[k,k+N]}, \ \forall j \in \mathbb{N}_{[1,L_{i^\dagger}]} \quad & p_{i^\dagger j}^\top \nu_{t|k} + \sqrt{p_{i^\dagger j}^\top \Gamma_{t|k} p_{i^\dagger j}}\,\Phi^{-1}\left(1 - \tfrac{1-\alpha_{\mathrm{out}}}{N}\right) \\ & \leq q_{i^\dagger j}. \end{aligned} \quad (20)$$

C. Computation of $\mathcal{E}_{\mathrm{act}}$ and $\mathcal{E}_{\mathrm{meas}}$

From Theorem 1, we must choose $\mathcal{E}_{\mathrm{act}}$ and $\mathcal{E}_{\mathrm{meas}}$ such that (10) is satisfied. Additionally, we require that $\mathcal{E}_{\mathrm{act}}$ and $\mathcal{E}_{\mathrm{meas}}$ are convex and compact for the computation of $\mathcal{O}_0(\mathcal{E}_{\mathrm{act}}, \mathcal{E}_{\mathrm{meas}}, \mathcal{X} \setminus \mathrm{KeepOut}_i)$. Since $\boldsymbol{w}_t$ and $\boldsymbol{\gamma}_t$ are Gaussian random vectors, we choose $\mathcal{E}_{\mathrm{act}}$ and $\mathcal{E}_{\mathrm{meas}}$ as ellipsoids, $\mathcal{E} = \{\Sigma^{\frac{1}{2}} z + \mu \mid \|z\| \leq r\}$ where μ and Σ are the mean and covariance matrices of $\boldsymbol{w}$ and $\boldsymbol{\gamma}$ respectively, and $r > 0$ is obtained from the chi-distribution's quantile function. See [7] for more details.

D. How conservative is the proposed approach?

We conclude this section by briefly discussing the various points in the algorithm design, where we introduced conservativeness for the sake of tractability. A main source of conservativeness is the use of an inner-approximation to the stochastic viability sets with state measurement uncertainty (Theorem 1 and Proposition 1) for the sake of tractability. Another source of conservativeness is the use of Boole's inequality to enforce joint chance constraints (12b) and (12c). The use of piecewise-affine approximation of the quantile function instead of fixed risk allocation can

TABLE I

RESULTS OF ABORT-SAFE RENDEZVOUS TRAJECTORY OPTIMIZATION

	Mixed-integer QP (19)	Convex QP (20)
Time to stopping criterion	25 minutes	30 minutes
Cost J	27714.93	28211.07
Maneuver ΔV	4.84 m/s	4.97 m/s
Computation of $\mathcal{O}$ (offline)	1263.74 minutes	0.12 minutes
Computation of U (online)	112.29 seconds	0.57 seconds

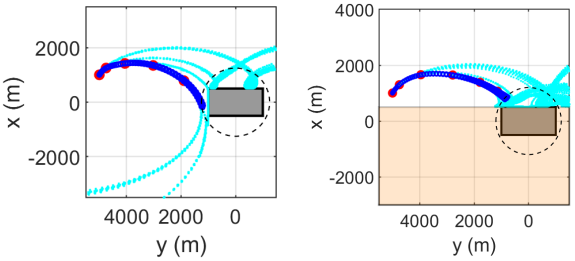

Fig. 2. Monte-Carlo validation (showing 20 out of 500 simulations) of the abort-safe rendezvous trajectories computed via (19) (left) and (20) (right). We mark the nominal rendezvous trajectory (blue), the off-nominal trajectory (cyan) after propulsion failure at time $T \in \{0, 10, 20, 30, 40, 50\}$ (red dots) after which actuation is limited $u_t \in \mathcal{V} \subset \mathcal{U}$ for $t \geq T$, the set KeepOut in grey, the halfspace selected for convexification $\text{KeepOut}_{i^\dagger}$ in orange, and the dotted circle marks the stopping criterion.

reduce the conservativeness from Boole's inequality [15]. When convexification (20) is used in place of the mixed-integer program (19), we incur additional conservativeness in the nominal trajectory since the scheduling of the keep-out constraints is now fixed. However, the convexification significantly reduces the computational effort.

IV. NUMERICAL EVALUATION

We apply the proposed approach to achieve active safety on a rendezvous example with the keep-out set $\text{KeepOut} = [-500, 500] \times [1000, 1000] \times \mathbb{R}^2$ (units in meter, unconstrained in velocity space). The discrete-time CW dynamics (2) have mean motion $\omega = \sqrt{\mu_\text{e}/R_\text{c}}$ with $R_\text{c} = 7228.1$ km, deputy's mass $m_\text{d} = 300$ kg, a sampling time of $\Delta t = 30$ seconds, a nominal actuation set $\mathcal{U} = [-1, 1]^2$ N, and an off-nominal actuation set $\mathcal{V} = \{0\} \times [-1, 1]$ N. We used a planning horizon of $N = 60$ (30 minutes), a safety horizon $M = 120$ (1 hour), trade-off parameter $\lambda = 100$, and set the desired rendezvous safety probability $\alpha = 0.9$. We set $\boldsymbol{w}$ and $\boldsymbol{\gamma}$ as zero-mean Gaussian random vectors with covariance matrices $\Sigma_w = \text{diag}(0.001, 0.001)$, $\Sigma_\gamma = \text{diag}(0.003, 0.003, 0, 0)$, and the initial state measurement $y_0 = [1000, 5000, 0, 0]$.

We used an Intel i7-4790K CPU with 4 GHz clock rate and 32 GB RAM running MATLAB 2020a for the computations. We used SReachTools [16] and YALMIP [17] to setup the stochastic optimal control problems (19) and (20), GUROBI [18] and ECOS [19] as the backend solvers, and MPT3 [14] to implement the recursion (9).

Figure 2 shows the nominal and off-nominal trajectories from six failure times $T \in \{0, 10, \ldots, 50\}$ via Monte-Carlo simulation. We stopped the simulation whenever the nominal trajectory is within a ball of radius 1250 m around the origin. We found the desired probability specifications were satisfied in every simulation. We also notice that the solution to (19) effectively utilizes its ability to schedule the keep-out halfspaces to come closer to the keep-out set within the planning horizon. Table I reports the time to stopping criterion, the resulting costs, the maneuver $\Delta V = (\Delta t/m_\text{d}) \sum_{k=0}^{N-1} \|u_k\|$, and the computation times, and illustrates the benefits of offline computation and convexification for Problem 1.

V. CONCLUSION

We present an abort-safe rendezvous trajectory optimization method that accounts for actuation uncertainty, measurement uncertainty, and the possibility of actuation failure. We generate a safe nominal rendezvous trajectory with high likelihood guarantees of safety, while ensuring that the recovery is possible using limited available actuation in the event of actuation failure. We utilize stochastic reachability to bypass the need for the synthesis of off-nominal, state-measurement-based recovery controllers.

REFERENCES

[1] D. Marsillach, S. Di Cairano, and A. Weiss, "Abort-safe spacecraft rendezvous in case of partial thrust failure," in *Proc. Conf. Dec. & Ctrl.*, 2020, pp. 1490–1495.

[2] W. Fehse, *Automated rendezvous and docking of spacecraft*. Cambridge Univ. Press, 2003, vol. 16.

[3] L. Breger and J. How, "Safe trajectories for autonomous rendezvous of spacecraft," *J. Guid. Ctrl. & Dyn.*, vol. 31, pp. 1478–1489, 2008.

[4] B. HomChaudhuri, M. Oishi, M. Shubert, M. Baldwin, and S. Erwin, "Computing reach-avoid sets for space vehicle docking under continuous thrust," in *Proc. Conf. Dec. & Ctrl.*, 2016, pp. 3312–3318.

[5] C. Zagaris and M. Romano, "Reachability analysis of planar spacecraft docking with rotating body in close proximity," *J. Guid. Ctrl. & Dyn.*, vol. 41, no. 6, pp. 1416–1422, 2018.

[6] D. Marsillach, S. Di Cairano, and A. Weiss, "Fail-safe rendezvous control on elliptic orbits using reachable sets," in *Proc. Amer. Ctrl. Conf.*, 2020, pp. 4920–4925.

[7] J. Gleason, A. Vinod, and M. Oishi, "Lagrangian approximations for stochastic reachability of a target tube," *Automatica*, vol. 125, 2021.

[8] A. Abate, M. Prandini, J. Lygeros, and S. Sastry, "Probabilistic reachability and safety for controlled discrete time stochastic hybrid systems," *Automatica*, vol. 44, no. 11, pp. 2724–2734, 2008.

[9] A. Abate *et al.*, "ARCH-COMP20 Category Report: Stochastic Models," *EPiC Series in Computing*, vol. 74, pp. 76–106, 2020.

[10] F. Borrelli, A. Bemporad, and M. Morari, *Predictive control for linear and hybrid systems*. Cambridge Univ. Press, 2017.

[11] L. Blackmore, M. Ono, and B. Williams, "Chance-constrained optimal path planning with obstacles," *IEEE Trans. Robot.*, vol. 27, no. 6, pp. 1080–1094, 2011.

[12] M. Conforti, G. Cornuéjols, and G. Zambelli, *Integer programming*. Springer, 2014.

[13] F. Oldewurtel, C. Jones, A. Parisio, and M. Morari, "Stochastic model predictive control for building climate control," *IEEE. Trans. Ctrl. Syst. Tech.*, vol. 22, no. 3, pp. 1198–1205, 2013.

[14] M. Herceg, M. Kvasnica, C. Jones, and M. Morari, "Multi-Parametric Toolbox 3.0," in *Proc. Euro. Ctrl. Conf.*, 2013, pp. 502–510.

[15] A. Vinod, V. Sivaramakrishnan, and M. Oishi, "Piecewise-affine approximation-based stochastic optimal control with Gaussian joint chance constraints," in *Proc. Amer. Ctrl. Conf.*, 2019, pp. 2942–2949.

[16] A. Vinod, J. Gleason, and M. Oishi, "SReachTools: A MATLAB stochastic reachability toolbox," in *Proc. Hybrid Syst.: Comp. & Ctrl.*, 2019, pp. 33–38.

[17] J. Löfberg, "YALMIP : A toolbox for modeling and optimization in MATLAB," in *Proc. CACSD Conf.*, 2004.

[18] L. Gurobi Optimization, "Gurobi optimizer reference manual," 2021.

[19] A. Domahidi, E. Chu, and S. Boyd, "ECOS: An SOCP solver for embedded systems," in *Proc. Euro. Ctrl. Conf.*, 2013, pp. 3071–3076.