Chance Constrained Model Predictive Control for Multi-Agent Systems with Coupling Constraints

Daniel Lyons, Jan-P. Calliess, and Uwe D. Hanebeck

Abstract— **We consider stochastic model predictive control of a multi-agent systems with constraints on the probabilities of inter-agent collisions. First, we discuss a method based on sample average approximation of the collision probabilities to make the stochastic control problem computationally tractable. Empirical results indicate that the complexity of the resulting optimization problem can be too high to be solved under real-time requirements. To reduce the computational burden we propose a second approach. It employs probabilistic bounds to determine regions of increased probability of presence for each agent and introduce constraints for the control problem prohibiting overlap of these regions. We prove that the resulting problem is conservative for the original problem, i.e., every control strategy that is feasible under our new constraints will automatically be feasible for the true original problem. Furthermore, we present simulations demonstrating improved run-time performance of our second approach and compare our stochastic method to robust control.**

I. Introduction

In this work, we propose two approaches to formulate stochastic model predictive control for multi-agent systems (MA-MPC) with chance constraints on the probability of inter-agent collisions as a mixed integer linear program (MILP). Chance constraints are user-defined upper bounds on the probability that the uncertain state of system violates state constraints and their application in model predictive control has received great interest in the recent past (for a brief literature review see Sec. I-A). MILPs are well-understood optimization problems that can be solved efficiently for moderate problem sizes and find many applications in robot planning, flight control, and receding horizon control (see for example [1], [2]).

Constraints on the probability that system states violate state bounds are in all but the simplest cases given by integrals of multivariate state distributions over complex non-convex regions [3]. The main challenge in chance constrained control is therefore to make the chance constraints computationally tractable while keeping the assumptions on state distributions and state constraints as general as possible.

In order to achieve this goal for the control of multi-agent systems, we first study a sample average approximation of the inter-agent collision probabilities. This approximation converges for fixed control inputs as the number of samples goes to infinity. However, the complexity of the resulting optimization problem is so high that this approach proves unsuitable for control under real-time requirements.

To alleviate the computational, burden we propose a second, more efficient approach. It uses probabilistic bounds to determine regions of increased probability of presence (RIPP) for each agent. We then formulate constraints for the control problem that guarantee that these regions do not overlap. We prove that the resulting problem is conservative for the original problem with probabilistic constraints. That is, every control strategy feasible under our new constraints will also satisfy the original constraints.

The main contribution of this work is that we find a deterministic and efficient constraint formulation that guarantees feasibility for the chance constrained problem without the need to evaluate complicated inter-agent collision probabilities. Since we employ a sample representation of the agents' uncertain positions and the probabilistic bounds we use for the RIPP regions hold for arbitrary uncertain states, we do not have to make any assumptions (such as being Gaussian) on the nature of the occurring noise or disturbances. Also in our approach we do not have to assume that the chance constraints are given by linear inequalities only and thus, are able to model the inherently non-convex coupling constraints on the states of the agents.

We show in simulations in a UAV path planning scenario that our proposed RIPP approach grants significantly better run-time performance compared to a controller with the sample average approximation of collision probabilities, with only a small degree of sub-optimality resulting from the conservativeness of our new approach. We also compare the probabilistic control methods with robust control for multi-UAV collision avoidance. We provide empirical evidence that the former are better suited for settings in which disturbance can be modelled stochastically, since it allows the user to precisely specify an upper bound on the probability of collisions for the UAVs.

To the best of our knowledge this is the first time a practical approach for the control of a multi-agent system with chance constraints on the probability of a collision of the agents is proposed. Existing work on chance constrained control either deals with planning for single agent systems or for systems without coupling chance constraints on joint states of different agents. The only other work that also considers coupling state constraints in a slightly different setting is our own work [4].

The paper is structured as follows. In Section II, we define the general problem of model predictive control for a multi-agent system with chance constraints. In Section III, we

D. Lyons and U. D. Hanebeck are with the Intelligent Sensor-Actuator-Systems Laboratory (ISAS), Karlsruhe Institute of Technology (KIT), Karlsruhe, Germany. lyons@kit.edu, uwe.hanebeck@ieee.org.

J. Calliess is with the Robotics Research Group, University of Oxford, Oxford, OX1 3PJ, UK. jan@robots.ox.ac.uk.

describe the system dynamics of the agents and the sample approximation of the state distributions. The contribution of this paper is described in Sections IV and V. In IV, we derive a sample average approximation of the collision avoidance constraints. In Section V, we detail the novel approach employing regions of increased probability of agent presence and prove its conservativeness. In Section VI, we illustrate several properties of our approach in a realistic UAV path planning scenario with non-Gaussian wind turbulence models and compare it with robust control techniques. Section VII concludes the paper.

A. Related Work

The body of work on interactions of agents in a multi-agent system is immense, so we focus here on recent results in chance constrained control of single-agent and multi-agent systems.

Most recent work on chance constrained model predictive control (MPC) for linear systems can roughly be classified into three parts: conservative approximations [5], conservative control of systems with Gaussian disturbance [6], [7], and sample-based control [8].

The authors of [5] approximate chance constraints by *conservative constraints* ensuring that confidence ellipsoids are completely contained in the feasible region. Computational comparisons in [6] and [7] indicate that, albeit being fast, this approach introduces a high degree of sub-optimality through its conservativeness.

For systems where the disturbance distributions are known to be *Gaussian*, [6], [7] propose that instead of enforcing that all state constraints are satisfied with a certain probability at the same time, each linear state constraint should be considered separately and the probability of violating this constraint be enforced separately. The authors of [9] extend these approaches to the control of multi-agent systems. Instead of considering coupling constraints on the states of the agents this work's emphasis is on the optimal distribution of risk among the agents.

In [8], the authors propose to approximate the probabilities for collisions with obstacles for the single agent case by *sample-based Monte-Carlo methods*. This approximation has the advantage that almost arbitrary state and noise distributions can be treated. They transform the search for an optimal solution to the stochastic single agent control problem under chance constraints to finding a solution of a mixed integer linear program (MILP). However, the considerations in [8] are restricted to chance constrained control of a single agent.

In [4], we have considered a decentralized auction-based mechanism to address the MA-MPC problem. As an application we have considered similar probabilistic inequalitiy-based techniques for multi-agent collision avoidance as we present in this paper.

B. Notation and Conventions

Bold face letters $\underline{\boldsymbol{x}}$ denote random variables, where the underline indicates that the random variable is multivariate. Bold face letters $\mathbf{A}$, $\mathbf{B}$, or $\mathbf{C}$ denote matrices. A superscript $\underline{x}^T$ denotes the transpose of $\underline{x}$. Deterministic quantities such as the system input $\underline{u}_t^i$ will be in normal type, where the underline indicates that the variable is a vector. The first superscript of a variable identifies the agent this variable refers to. The first subscript identifies the time instance in the planning horizon.

II. General Problem Formulation

The general problem we aim to solve is as follows: For M agents $i = 1, \dots, M$, with discrete-time stochastic state space model, we plan over a horizon of length H in order to minimize the sum of the agents' cost functions. This minimization is subject to the constraints that the probabilities of inter-agent collisions and the probability of agents leaving the feasible region are kept below certain user-defined thresholds. The formal formulation of this problem is

$$\underset{\underline{u}^i_{1:H}, i=1,\dots,M}{\text{minimize}} \sum_{i=1}^{M} \mathbb{E}_{\underline{\boldsymbol{x}}^i_{0:H}} \left\{ h^i(\underline{x}^i_{0:H}, \underline{u}^i_{1:H}) \right\} \quad (1)$$

$$\text{s.t.} \quad \text{for all agents } i, j = 1, \dots, M, \ i \neq j \quad (2)$$

$$\underline{u}^i_{1:H} \in F^i_u \quad (3)$$

$$\underline{\boldsymbol{x}}^i_t = f^i_t(\underline{\boldsymbol{x}}^i_0, \underline{u}^i_{1:t}, \underline{\boldsymbol{\nu}}^i_{1:t}), \forall t \in \{1, \dots, H\} \quad (4)$$

$$\Pr(\underline{\boldsymbol{x}}^i_{1:H} \notin F^i) \leq \delta^i \quad (5)$$

$$\Pr((\underline{\boldsymbol{x}}^i_t, \underline{\boldsymbol{x}}^j_t) \notin F^{i,j}_t) \leq \delta^{i,j}, \forall t \in \{1, \dots, H\} \ . \quad (6)$$

The decision variables $\underline{u}^i_{1:H} = [(\underline{u}^i_1)^T, (\underline{u}^i_2)^T, \dots, (\underline{u}^i_H)^T]^T$ are the (deterministic) control inputs to agent i, confined to lie in the polytopal feasible region F^i_u. The functions h^i are the control objectives and rate how desirable certain states of the agents are. They depend on the control inputs and the system states of the agents. The system state of agent i over the planning horizon is modeled as a random vector and is denoted by $\underline{\boldsymbol{x}}^i_{0:H} = [(\underline{\boldsymbol{x}}^i_0)^T, (\underline{\boldsymbol{x}}^i_1)^T, \dots, (\underline{\boldsymbol{x}}^i_H)^T]^T$. Reasons for modelling states as random vectors are for example uncertain localization of robotic systems caused by noisy sensor measurements. We assume these random vectors to be stochastically independent for different agents $i_0 \neq i_1$. Since the states are modeled as random vectors, we take the expectation of $h^i(\underline{x}^i_{0:H}, \underline{u}^i_{1:H})$ w.r.t. the agents' state distributions in (1).

The mapping f^i_t describes the model of the dynamics of agent i. The state of agent i at time step t depends on the initial state $\underline{\boldsymbol{x}}^i_0$, the control inputs $\underline{u}^i_1, \dots, \underline{u}^i_t$ up to time step t, and stochastic noise $\underline{\boldsymbol{\nu}}^i_1, \dots, \underline{\boldsymbol{\nu}}^i_t$. The stochastic noise terms are used to account for possible errors in the dynamic model or exogenous disturbances that act upon the agents, such as wind turbulence on UAVs.

We assume here that the second moments of the noise terms $\underline{\boldsymbol{\nu}}^i_t$ and the prior state distributions $\underline{\boldsymbol{x}}^i_0$ are known and that the disturbance is independent from the control inputs.

F^i is the feasible region for the state of agent i and $\Pr(\underline{\boldsymbol{x}}^i_{1:H} \notin F^i)$ is the probability that the state leaves the feasible region. The feasible region can for example model an area the agents are not supposed to leave or obstacles the agents have to avoid. $\Pr((\underline{\boldsymbol{x}}_t^i, \underline{\boldsymbol{x}}_t^j) \notin F_t^{i,j})$ specifies the probability that agent i and agent j do not meet the coupling constraints at time step t defined by the feasible region $F_t^{i,j}$ that controls the interaction among agents. For simplicity consider constraints consistent for all agents and all time steps and set $F^c := F_t^{i,j}$. We interpret these constraints as collision avoidance constraints, i.e., constraints on the expected distance of agent i and j.

The bounds $0 \leq \delta^i, \delta^{i,j} \leq 1$ on the probability of the agents leaving their own feasible region F^i or the joint feasible regions F^c characterize the chance constraints on the general failure of planning [10], [11], [3]. These chance constraint bounds are specified by the user and can be used to adjust how cautious the agents' plans will become.

III. Sample-based Methods for Systems with Linear Dynamics

For arbitrarily distributed initial states $\underline{\boldsymbol{x}}_0^i$ and for arbitrarily (non-Gaussian) distributed system noise $\underline{\boldsymbol{\nu}}_t^i$ there is in general no closed-form representation (or one with a finite number of parameters) of the densities of the random vectors $\underline{\boldsymbol{x}}_t^i$. We therefore make use of sample-based methods to represent the uncertain states of the agents and describe how to do so in this section.

We assume that the stochastic discrete-time dynamic state space model of each agent $i = 1, \ldots, M$ is given by the linear system equation

$$\underline{\boldsymbol{x}}_{t+1}^i = \mathbf{A}^i \underline{\boldsymbol{x}}_t^i + \mathbf{B}^i \underline{u}_t^i + \underline{\boldsymbol{\nu}}_t^i,\ t = 1, \ldots, H-1, \qquad (7)$$

where $\underline{\boldsymbol{x}}_t^i$ is the uncertain state of agent i modeled as a random variable, $\underline{u}_t^i$ is the deterministic control input and $\underline{\boldsymbol{\nu}}_t^i$ is the stochastic system noise. For the linear system dynamics (7) the mapping $\underline{\boldsymbol{x}}_t^i = f_t^i(\underline{\boldsymbol{x}}_0^i, \underline{u}_{1:t}^i, \underline{\boldsymbol{\nu}}_{1:t}^i)$ that determines how the agent's state depends on the initial state $\underline{\boldsymbol{x}}_0^i$ and behaves under control inputs $\underline{u}_{1:t}^i$, is linear in the control inputs and given by

$$\underline{\boldsymbol{x}}_{t+1}^i = (\mathbf{A}^i)^t \underline{\boldsymbol{x}}_0^i + \sum_{s=1}^{t} (\mathbf{A}^i)^{t-s-1} (\mathbf{B} \underline{u}_s^i + \underline{\boldsymbol{\nu}}_s^i) \ . \qquad (8)$$

We assume that for each agent $i = 1, \ldots, M$ we can draw N independent samples from the distribution of $\underline{\boldsymbol{x}}_0^i$ at time step $t = 0$. We will denote these samples by $\{\underline{x}_{0,j}^i\}_{j=1}^N$ and assume, for notational convenience only, that we draw an equal number N of samples from each agent's prior distribution. Furthermore for each agent i we draw N noise samples $(\underline{\nu}_{1,j}^i, \underline{\nu}_{2,j}^i, \ldots, \underline{\nu}_{H,j}^i)$, $j = 1, \ldots, N$ from the distribution of the system noise $\underline{\boldsymbol{\nu}}_1^i, \ldots, \underline{\boldsymbol{\nu}}_H^i$ that affects agent i over the planning horizon of length H.

The model of the dynamics (7) of each agent i allows us to generate N sample trajectories over time. These trajectories are obtained by propagating each initial sample in combination with a noise sample through the system equation (8). Each sample trajectory is of length $H + 1$ and consists of samples

$$\underline{x}_{0:H,j}^i := [(\underline{x}_{0,j}^i)^T, (\underline{x}_{1,j}^i)^T, \ldots, (\underline{x}_{H,j}^i)^T]^T \ . \qquad (9)$$

The most important property of this construction is that the sample trajectories depend deterministically and linearly on the control inputs. This will enable us to formulate a sample-based approximation of the stochastic control problem as a deterministic optimization problem.

With this construction we have obtained a representation of agent i's uncertain states. For instance, at time step t_0 the samples $\{\underline{x}_{t_0,j}^i\}_{j=1}^N$ can be utilized to estimate important statistics of the distribution of the random variable $\underline{\boldsymbol{x}}_{t_0}^i$ like its mean or the probability of violating state constraints as we will highlight in the next section. [8] showed that under certain assumptions sample-based chance-constrained control (1)-(5) for single agent systems can be formulated as a mixed integer linear optimization problem. We will extend these methods to control in multi-agent systems and focus our attention on the inter-agent collision avoidance chance constraints (6).

IV. Multi-Agent Collision Avoidance

This section is concerned with the approximation of the probabilities of inter-agent collisions and the derivation of constraints that keep these probabilities below user-defined thresholds. We will first formally define our notion of an inter-agent collision ("the agents come too close") and then study how from this definition, the true probability of an inter-agent collision can be derived. Because the probability of a collision of two agents will depend on *both* states of the agents its evaluation is even more computationally expensive than the computing single agent constraint violation probabilities. In a first step we use the sample approximations of the agents' uncertain states to derive a sample approximation of the probability of an inter-agent collision. The complexity of the resulting optimization problem can be so high that this approach proves problematic for control under real-time requirements.

For simplicity of exposition and for notational convenience, we consider path planning in the two-dimensional plane and we will from now on only consider two agents, denoted by a superscript 1 and 2. Since we defined the constraints on the probability of a collision in (6) as pairwise constraints on the states of two agents, the following considerations can be extended verbatim for every other possible combination of two agents.

A. Definition of Inter-Agent Collision Probability

Let $\epsilon > 0$ be the prespecified minimum distance between two agents for collision avoidance, for instance we could set ϵ to be twice the diameter of a robot or the twice the wingspan of a fixed-wing UAV.

Let $\underline{x}_t^1$ and $\underline{x}_t^2$ denote the two-dimensional positions of agents 1 and 2 in the plane at a time step t. We define the event of a *collision* as $\|\underline{x}_t^1 - \underline{x}_t^2\|_2 < \epsilon$, i.e., the two agents are closer than the minimum clearance ϵ. F^c, defined by the inter-agent constraints, is the set of joint agent positions that have a distance greater or equal than ϵ:

$$F^c := \{(\underline{x}^1, \underline{x}^2) \colon \|\underline{x}^1 - \underline{x}^2\|_2 \geq \epsilon\} \ . \qquad (10)$$

The probability of a collision of two agents is

$$\begin{aligned}\Pr((\underline{\boldsymbol{x}}_t^1, \underline{\boldsymbol{x}}_t^2) \notin F^c) &= \Pr(\|\underline{\boldsymbol{x}}_t^1 - \underline{\boldsymbol{x}}_t^2\|_2 < \epsilon) && (11)\\ &= \mathbb{E}_{\underline{\boldsymbol{x}}_t^1,\ \underline{\boldsymbol{x}}_t^2}\{\chi_{CF^c}(\underline{x}_t^1, \underline{x}_t^2)\} && (12)\\ &= \int\int \chi_{CF^c}(\underline{x}_t^1, \underline{x}_t^2) g_{\underline{\boldsymbol{x}}_t^1}(\underline{x}_t^1) g_{\underline{\boldsymbol{x}}_t^2}(\underline{x}_t^2) d\underline{x}_t^1 d\underline{x}_t^2\ , && (13)\end{aligned}$$

where

$$\chi_{CF^c}(\underline{x}_t^1, \underline{x}_t^2) = \begin{cases} 1, & \text{if } \|\underline{x}_t^1 - \underline{x}_t^2\|_2 < \epsilon \\ 0, & \text{otherwise} \end{cases} \tag{14}$$

is the indicator function of the complement of F^c and $g_{\underline{\boldsymbol{x}}_t^1}(\underline{x}_t^1)$ and $g_{\underline{\boldsymbol{x}}_t^2}(\underline{x}_t^2)$ are the probability density functions of the position estimates of the agents.

The integral (13) is the integral of multivariate density functions multiplied by the indicator function χ_{CF^c}. For general density functions this integral will have no closed-form solution because of the possibly complex structure of the densities. Even for multivariate Gaussian distributions it will become difficult to evaluate the integral since the indicator function is a non-convex and nonlinear function and there are no known formulas to evaluate this integral.

B. Sample Average Approximation of Inter-agent Collision probability

In order to render the probability of an inter-agent collision computationally tractable, we replace it by the sample average approximation (SAA) [12]. We do this by replacing the continuous distributions $g_{\underline{\boldsymbol{x}}_t^1}(\underline{x}_t^1)$ and $g_{\underline{\boldsymbol{x}}_t^2}(\underline{x}_t^2)$ by their corresponding sample approximations, given by the two sets of samples $\{\underline{x}_{t,j}^1\}_{j=1}^N$ and $\{\underline{x}_{t,l}^2\}_{l=1}^N$. When the continuous distributions are replaced by their sample counterparts the double integrals in Equation (13) are replaced by a nested sum and, hence, the approximated chance constraint on the probability of a collision is

$$\frac{1}{N^2}\sum_{j=1}^{N}\sum_{l=1}^{N}\chi_{CF^c}(\underline{x}_{t,j}^1, \underline{x}_{t,l}^2) \leq \delta_t^{1,2}\ . \tag{15}$$

In (15) the indicator function has to be evaluated N^2 times. Each evaluation can be formulated as binary linear constraints and the control problem can be formulated as MILP and solved by standard commercial MILP solvers like CPLEX [13]. However, each of these N^2 evaluations requires a fixed number of binary constraints in the MILP formulation. Since one usually desires to employ as many samples as possible to guarantee good approximation properties and since binary variables have a crucial impact on the complexity and thus the run-time of the program, a quadratic dependence of the number of binaries on the number of samples renders this formulation problematic to solve under real-time requirements.

Due to space restrictions we refer the reader to the technical report accompanying this submission [14] for a detailled formulation of (15) as integer linear constraints and a complexity analysis.

V. Efficient Collision Avoidance Approximation

In this section, we will propose a more efficient formulation of inter-agent collision avoidance constraints based on *regions of increased probability of presence (RIPP)* of agents. Not only does the RIPP formulation allow us to generate controls under real-time requirements with only a small degree of sub-optimality compared to the optimal solution of the chance constrained problem, but we can also formally prove that controls found with the RIPP algorithm are feasible to the original problem with chance constraints on the true probabilities of agent collisions.

The derivation of the RIPP algorithm proceeds in two steps: In the first step, we define a region around the mean of the uncertain position $\underline{\boldsymbol{x}}^i$ of an agent i and study the probability that the agent is outside of this region. The intuition is that the larger this region is, the less probable it is that the position of the agent lies outside the region. This intuition can be quantified by a form of the Chebychev inequality that gives an upper bound on the probability that the position of the agent lies outside of the RIPP region.

In the second step we introduce constraints to the control problem that ensure that for different agents their respective RIPPs do not overlap. We will prove that if these RIPPs have the adequate size and they do not overlap, then we can control the complex collision probabilities in such a way that they do not exceed the predefined bounds.

As before in Section IV, we will only consider two agents 1 and 2 and their two-dimensional positions in the plane. We will omit the subscript t denoting the time step in this section for notational convenience.

A. Regions of Increased Probability of Presence (RIPP)

Let $\underline{\mu}^i := [\mu_x^i, \mu_y^i]^T \in \mathbb{R}^2$ be the mean of the uncertain position $\underline{\boldsymbol{x}}^i$ of agent i, where the subscript x and y denote the x and y coordinates in the plane. *We define a rectangular region of increased probability of presence (RIPP) around the mean through*

$$R^i := \{[x^i, y^i]^T : |x^i - \mu_x^i| \leq \alpha_x^i, |y^i - \mu_y^i| \leq \alpha_y^i\} \subset \mathbb{R}^2\ .$$

The RIPP R^i describes the set of points in $\mathbb{R}^2$ for which both coordinates deviate at most some distance from the mean of the uncertain state of the agent. Its position depends on the mean of the uncertain state of the agent and its size depends on the two parameters α_x^i and α_y^i.

For each agent i, we define the probability that the uncertain position of agent i lies outside of the RIPP R^i as

$$P^i := \Pr(\underline{\boldsymbol{x}}^i \notin R^i) = 1 - \Pr(\underline{\boldsymbol{x}}^i \in R^i)\ . \tag{16}$$

For larger RIPPs R^i, the probability P^i becomes smaller as is quantified by the following theorem.

Theorem 1 (P. Whittle [15]): Let $\underline{\boldsymbol{x}}^i = [\boldsymbol{x}_x^i, \boldsymbol{x}_y^i]^T$ be a zero mean bivariate random vector with covariance matrix $(C_{kl})_{k,l=1,2}$ and probability P^i for $\underline{\boldsymbol{x}}^i$ defined as in (16), then P^i can be bounded by

$$P^i \leq \frac{C_{11}(\alpha_y^i)^2 + C_{22}(\alpha_x^i)^2}{2(\alpha_x^i)^2(\alpha_y^i)^2} + \frac{\sqrt{[C_{11}(\alpha_y^i)^2 + C_{22}(\alpha_x^i)^2]^2 - 4C_{12}^2(\alpha_x^i)^2(\alpha_y^i)^2}}{2(\alpha_x^i)^2(\alpha_y^i)^2} \ . \ \Diamond \quad (17)$$

For later reference we will denote the right-hand side of (17) as $\mathrm{C}(\underline{\boldsymbol{x}}^i, R^i)$.

In the rest of this section, we want to investigate the following question: Given an amount γ^i of probability mass, how do we determine the size of the region such that at most γ^i of the probability mass of the uncertain state $\underline{\boldsymbol{x}}^i$ lies outside of the region? Since in the next section, we will introduce constraints that prohibit an overlap of RIPP regions and these constraints are the more conservative the bigger the regions become, resulting in a loss of optimality of the motion plans, we aim to find RIPP of minimal size such that (18) still holds. That is, we desire to determine R^i of minimal size such that

$$\Pr(\underline{\boldsymbol{x}}^i \notin R^i) \leq \gamma^i \ . \quad (18)$$

The size of the region R^i is determined by the parameters α_x^i, which is the region's extent in the direction of the x coordinate, and α_y^i, which is the extent in the y direction. If we find α_x^i and α_y^i so that $\mathrm{C}(\underline{\boldsymbol{x}}^i, R^i) = \gamma^i$ holds, then by Thm. 1 we can guarantee that $\Pr(\underline{\boldsymbol{x}}^i \notin R^i) \leq \mathrm{C}(\underline{\boldsymbol{x}}^i, R^i) = \gamma^i$, the probability mass of the uncertain position $\underline{\boldsymbol{x}}^i$ outside of the region R^i is at most γ^i.

In order to bound the probability mass of the uncertain position of agent i outside of the region R^i we have to determine the parameters α_x^i and α_y^i such that the equality $\mathrm{C}(\underline{\boldsymbol{x}}^i, R^i) = \gamma^i$ holds. To achieve this, we have to add another equation, since the equation $\mathrm{C}(\underline{\boldsymbol{x}}^i, R^i) = \gamma^i$ is only one equation for the two unknowns α_x^i and α_y^i and hence it is under-determined. We propose to choose the parameters α_x^i and α_y^i such that additionally

$$\frac{\alpha_x^i}{\alpha_y^i} = \sqrt{\frac{C_{11}^i}{C_{22}^i}} \quad (19)$$

holds. This choice is motivated by the intuition that for an uncertain position with axis-aligned Gaussian distribution (i.e., $C_{12}^i = 0$ in the covariance), the diagonal of the covariance matrix is given by C_{11}^i and C_{22}^i, quantifying the extent of the covariance ellipsoid in x-direction and y-direction. If the ratio of α_x^i and α_y^i equals the ratio of $\sqrt{C_{11}^i}$ and $\sqrt{C_{22}^i}$, the shape and extent of R^i follows the shape and extent of the covariance ellipsoid. So if there is considerable uncertainty in one of the coordinate directions, indicated by a covariance ellipsoid with strong extent in this direction, the region will also have a stronger spread in this direction to account for this increased uncertainty. Other relations could be introduced and their influence on the region R^i and the resulting constraints will be subject to future research.

When we insert the equation for α_y^i that results from (19)

$$\alpha_y^i = \sqrt{\frac{C_{22}^i}{C_{11}^i}} \alpha_x^i \quad (20)$$

into the equation $C(\underline{\boldsymbol{x}}^i, R^i) = \gamma^i$ the latter is an equation with only one remaining unknown, namely α_x^i. The equation for the remaining unknown is a polynomial of degree four with two real solutions that can be determined analytically. The positive real solution of the equation is

$$\alpha_x^i = \sqrt{\frac{C_{11}^i}{\gamma^i} + \frac{\sqrt{C_{11}^i C_{22}^i (C_{11}^i C_{22}^i - (C_{12}^i)^2)(\gamma^i)^2}}{C_{22}^i (\gamma^i)^2}} \ . \quad (21)$$

Given an uncertain position $\underline{\boldsymbol{x}}^i$ together with the covariance matrix of this uncertain position and given a level γ^i, Eqs. (21) and (19) allow us to determine the RIPP R^i such that $C(\underline{\boldsymbol{x}}^i, R^i) = \gamma^i$ holds. Together with Whittle's Chebychev inequality, we can then guarantee that the probability mass of $\underline{\boldsymbol{x}}^i$ outside the RIPP R^i is at most γ^i, i.e., $\Pr(\underline{\boldsymbol{x}}^i \notin R^i) \leq \gamma^i$.

B. Collision Avoidance Based on Non-overlapping RIPPs

In this section, we will apply the results of the previous section to deriving our RIPP formulation of collision avoidance constraints. We consider two agents 1 and 2 with uncertain positions $\underline{\boldsymbol{x}}_t^1$ and $\underline{\boldsymbol{x}}_t^2$ together with an upper bound $0 \leq \delta_t^{1,2} \leq 1$ on the probability of a collision of these agents.

The RIPP constraints are constructed in two steps: In the first step, we determine the RIPPs R_t^1 and R_t^2 for agents 1 and 2 such that the probability mass outside of the RIPPs is at most some γ_t^i for both $i = 1, 2$. In the second step, a constraint is added to the optimization problem prohibiting that the two RIPPs overlap at time step t.

We assume that the system noise $\underline{\boldsymbol{\nu}}_t^i$ for all agents and for all time steps is zero-mean. [1] Then the mean $\underline{\mu}_t^i$ of agent i's uncertain state follows the recursive rule

$$\underline{\mu}_t^i = \mathbf{A}^i \underline{\mu}_{t-1}^i + \mathbf{B}^i \underline{u}_{t-1}^i \ . \quad (22)$$

and always *depends linearly* on the control inputs.

The covariances of the uncertain states $\underline{\boldsymbol{x}}_t^i$ do not depend on any control inputs, but only on the covariances of the prior distribution $\underline{\boldsymbol{x}}_0^i$ and the noise terms $\underline{\boldsymbol{\nu}}_{1:t}^i$. The recursive formula for the evolution of the covariances is

$$\begin{aligned}\mathrm{Cov}(\underline{\boldsymbol{x}}_t^i) &= \mathbf{A}^i \mathrm{Cov}(\underline{\boldsymbol{x}}_{t-1}^i)(\mathbf{A}^i)^T + \mathrm{Cov}(\underline{\boldsymbol{\nu}}_{t-1}^i) \\ &\quad + \mathbf{A}^i \mathrm{Cov}(\underline{\boldsymbol{x}}_{t-1}^i, \underline{\boldsymbol{\nu}}_{t-1}^i) + \mathrm{Cov}(\underline{\boldsymbol{\nu}}_{t-1}^i, \underline{\boldsymbol{x}}_{t-1}^i)(\mathbf{A}^i)^T \ .\end{aligned} \quad (23)$$

This property can be derived from basic matrix manipulations and covariance matrix properties. We assumed that covariances of the prior distributions $\underline{\boldsymbol{x}}_0^i$ and the noise terms $\underline{\boldsymbol{\nu}}_t^i$ are known, so the agents can recursively compute the covariance of their uncertain state at time step t.

We will split the chance constraint bound $\delta_t^{1,2}$ into two parts γ_t^1 and γ_t^2 and distribute these parts onto the agents to construct a RIPP for each agent. We propose that agents 1 and 2 split the chance constraint bound $\delta_t^{1,2}$ into parts according to

$$\gamma_t^1 = \frac{1}{d}\delta_t^{1,2} \text{ and } \gamma_t^2 = \frac{d-1}{d}\delta_t^{1,2} \ , \quad (24)$$

[1] If the noise had a non-zero mean, we could subtract this mean from the system equations as deterministic disturbance and would have reduced this noise to zero-mean noise again.

with free parameter $d > 1$. In our simulations in Section VI, we used an even split at $d = 2$. One could also employ a negotiation algorithm to find a split that is optimal for the agents and this will be addressed in future work. Equations $C(\underline{\boldsymbol{x}}_t^1, R_t^1) = \gamma_t^1$ and $C(\underline{\boldsymbol{x}}_t^2, R_t^2) = \gamma_t^2$ then uniquely determine the α-size parameters for the RIPP regions. We are in a position to define the constraints preventing overlap of the RIPPs.

Definition 1 (Constraint $\mathcal{C}_t$): The expected values $\underline{\mu}_t^1$ and $\underline{\mu}_t^2$ have a distance of more than $\alpha_{t,x}^1 + \alpha_{t,x}^2 + \epsilon$ in the x-direction, i.e., $|\mu_{t,x}^1 - \mu_{t,x}^2| > \alpha_{t,x}^1 + \alpha_{t,x}^2 + \epsilon$ *or* a distance of more than $\alpha_{t,y}^1 + \alpha_{t,y}^2 + \epsilon$ in the y-direction, i.e., $|\mu_{t,y}^1 - \mu_{t,y}^2| > \alpha_{t,y}^1 + \alpha_{t,y}^2 + \epsilon$.

The following theorem shows that controls for agents 1 and 2 that satisfy constraint $\mathcal{C}_t$, satisfy that the probability of a collision between agent 1 and 2 at time step t is less or equal than the bound $\delta_t^{1,2}$, i.e. they are feasible for the problem with bounds on inter-agent collision probabilities.

Theorem 2: Let F^c be the feasible region for inter-agent collision avoidance and $\underline{\boldsymbol{x}}_t^1$, $\underline{\boldsymbol{x}}_t^2$ the uncertain positions of agents 1 and 2 at time step t. Furthermore let γ_t^1 and γ_t^2 be such that $\gamma_t^1 + \gamma_t^2 = \delta_t^{1,2}$ and let R_t^1 and R_t^2 be so that $C(\underline{\boldsymbol{x}}_t^i, R_t^i) = \gamma_t^i$. Then for any control sequences $\underline{u}_{1:H}^1$ and $\underline{u}_{1:H}^2$ inducing states satisfying constraint $\mathcal{C}_t$ at time step $t \leq H$, the probability of a collision of agents 1 and 2 is below $\delta_t^{1,2}$. That is,

$$\Pr((\underline{\boldsymbol{x}}_t^1, \underline{\boldsymbol{x}}_t^2) \notin F^c) \leq \delta_t^{1,2} \ . \tag{25}$$

Proof: (Sketch) Due to space restriction we omit the complete proof and refer the reader to the technical report [14]. The proof procedes in two steps. In the first step we show that the true probability of a collision can be bounded from above by the sum of the probabilities that each agent is outside of its respective RIPP region, if constraint $\mathcal{C}_t$ holds, i.e., $\Pr((\underline{\boldsymbol{x}}_t^1, \underline{\boldsymbol{x}}_t^2) \notin F^c | \mathcal{C}_t) \leq \Pr(\underline{\boldsymbol{x}}_t^1 \notin R_t^1 | \mathcal{C}_t) + \Pr(\underline{\boldsymbol{x}}_t^2 \notin R_t^2 | \mathcal{C}_t)$. This is done by marginalizing the true collision probability and then showing that the event that the agents are within their resp. RIPPs and a collision occurs $\{\underline{x}_t^1 \in R_t^1, \underline{x}_t^2 \in R_t^2, \|\underline{x}_t^1 - \underline{x}_t^2\| < \epsilon\}$ has zero probability measure. In the second step, Whittle's Chebychev inequality and the results of the previous section are used to bound the probabilities of agents being outside of their RIPP. ■

The theorem guarantees that if we solve the chance constrained MA-MPC problem with RIPP constraints of the form $\mathcal{C}_t$, $t = 1, \ldots, H$, then the obtained controls are automatically feasible for the chance constrained MA-MPC problem with full constraints on inter-agent collision probabilities. Hence, the novel RIPP constraints allow us to find solutions to the MA-MPC problem including the complicated probabilistic coupling constraints without any knowledge about the agents' uncertain states besides the covariance and also without ever actually computing or evaluating the probability of a collision of two agents.

VI. Simulations

In this section, we consider path planning for multiple UAVs whose movements are affected by wind disturbances [16]. We assume that the UAVs all fly at the same fixed height and therefore consider collision avoidance in the two-dimensional plane. The task of the UAVs is to reach a certain goal point on as quickly as possible in order to save time. Bounds on the control inputs are given by bounds on the maximum acceleration and bounds on the maximum speed of the UAVs.

A. Model Parameters

We assume that the UAVs all have the same linear motion model given by the double integrator model

$$\underline{\boldsymbol{x}}_t^i = \mathbf{A}\underline{\boldsymbol{x}}_{t-1}^i + \mathbf{B}\underline{u}_{t-1}^i + \underline{\boldsymbol{\nu}}_{t-1}^i \tag{26}$$

with $\underline{\boldsymbol{x}}_t^i = [\boldsymbol{x}_t^i,\ \boldsymbol{y}_t^i,\ \dot{\boldsymbol{x}}_t^i,\ \dot{\boldsymbol{y}}_t^i]^T$, $\underline{u}_t^i = [\ddot{x}_t^i, \ddot{y}_t^i]^T$, $\|\underline{u}_t^i\|_\infty \leq 12$ and

$$\mathbf{A} = \begin{bmatrix} \mathbf{I}_2 & \mathbf{I}_2 \\ \mathbf{0}_2 & \mathbf{I}_2 \end{bmatrix} \ , \ \mathbf{B} = \begin{bmatrix} \mathbf{0}_2 \\ \mathbf{I}_2 \end{bmatrix} \ ,$$

where $\mathbf{I}_2$ is the two-dimensional identity matrix and $\mathbf{0}_2$ the two-dimensional zero matrix. We assume that the initial uncertain states of the agents have Gaussian distribution with covariance $\underline{\boldsymbol{x}}_0^i \sim \mathcal{N}(\underline{\mu}^i, \mathbf{C}_0^i)$, $\mathbf{C}_0^i = \text{diag}[10^{-3},\ 10^{-3},\ 10^{-5},\ 10^{-5}]$. We assume that the target way points are given as two-dimensional positions $\underline{Z}^i = [Z_1^i,\ Z_2^i]^T$ for each agent and that the objective function is the distance of the position to the target way point (normalized for better comparability with $(HMN)^{-1}$): $h^i(\underline{x}_t^i) = (HMN)^{-1}\left(|x_t^i - Z_1^i| + |y_t^i - Z_2^i|\right)$.

The disturbance samples $\underline{\nu}_{t,j}^i$ affecting the UAVs are drawn from the discrete Dryden low-altitude model to simulate wind turbulence acting on the UAVs [16]. The UAVs are assumed to fly with a maximum speed of 45 feet per s at a fixed altitude of 200 feet through a turbulence field with light turbulence with wind speed of 15 knots at 20 feet height. The minimum distance between the UAVs is set to $\epsilon = 5$ feet. The mixed integer linear solver we used is CPLEX [13].

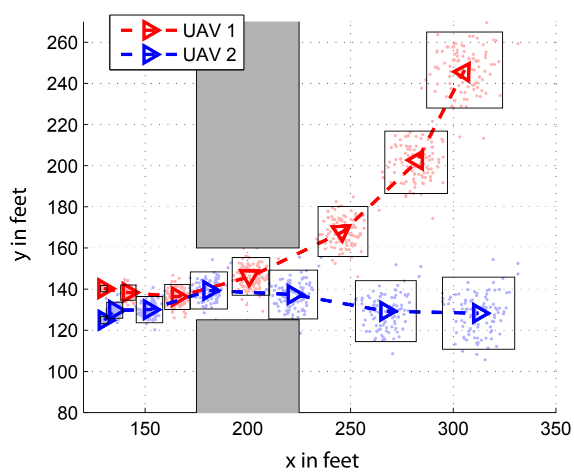

Fig. 1. The plot shows the trajectories of two UAVs passing a gap between obstacles (depicted as solid blue boxes) computed with our novel RIPP constraints. The boxes around the agents' position depict RIPPs.

B. Example Scenario

Figure 1 depicts the trajectories of two UAVs, computed with the RIPP constraints. The planning horizon for this example was set to $H = 7$, the starting parameters were $\underline{\mu}_1 = [130, 140]^T$, $\underline{\mu}_2 = [130, 125]^T$, $\underline{Z}_1 = [300, 250]^T$ and $\underline{Z}_2 = [320, 130]^T$ with $N = 100$ samples for each UAV's uncertain state. Obstacle avoidance was conducted according to [8]. The algorithm with our RIPP constraints solved the MILP to a value of the control objective of $h_{min} = 146.8437$ while the optimal solution with sample average approximation (SAA) of inter-agent collision probabilities achieved only a slightly better value of $h_{min} = 141.2349$. The run-time of the SAA was almost 50 times longer than with the RIPP approach, while the optimal value of the control objective is only 4% better. Here and in the next simulation, we consider SAA as the reference, since its approximation of collision probabilities will converge as the number of samples goes to infinity. Hence, we understand plans found with SAA as being close to optimal plans.

From the trajectories in Fig. 1 it can be seen that the lower UAV lets the upper UAV pass through the gap first, since this behavior grants a better overall performance for the system consisting of both UAVs. Also it can be seen that our RIPP constraints did not force a strict separation of the samples as some samples are allowed to fall below the minimum distance. The theoretical results guarantee that the probability of a collision of the two UAVs did not exceed the predefined threshold.

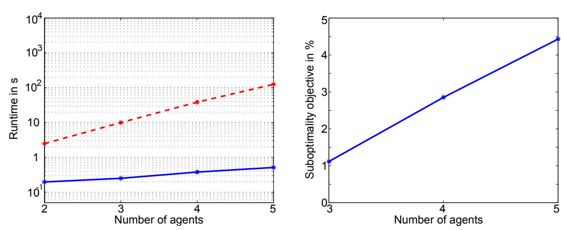

(a) Run-time comparison in seconds. (b) Suboptimality in percent.

Fig. 2. In (a) the run-times in seconds of the controller with sample average approximation of inter-agent collision probabilities (SAA) are depicted as red dashed line. The blue continuous line depicts the run-times of a controller with RIPP constraints. In (b) we plot the sub-optimality of controls found with our novel RIPP constraints in terms of percent of the optimal objective.

C. Quantitative Results

In this section, we evaluate how the run-times of problems with sample approximation constraints compare with the run-time of problems with RIPP constraints and how high the degree of suboptimality was for the more conservative RIPP constraints.

In Fig. 2(a), we plotted the run-time of a controller with the sample average approximation (SAA) and of a controller with the RIPP constraints for collision avoidance, both as a function of the number of agents that are controlled. The UAVs' starting positions and their target way points were randomly drawn to lie within certain areas. All results were averaged over 50 Monte-Carlo runs with each UAV's uncertain state approximated by 30 samples. The run-times of the MA-MPC algorithm with RIPP constraints for the same Monte-Carlo runs can also be found in Table I. It can be seen that the controller SAA had high run-time even for a small number of agents. The RIPP constraints yielded significantly faster solutions suggesting this approach is better applicable under real-time requirements.

	Run-time RIPP constraints [s]		
Agents	min	max	average
2	0.14	1.12	0.19
3	0.22	0.29	0.25
4	0.28	1.09	0.37
5	0.39	0.73	0.51

TABLE I

RUN-TIMES OF PLANNING WITH RIPP CONSTRAINTS.

Fig. 2(b) shows the sub-optimality of plans based on RIPP constraints in percent of the optimal objective achieved by the SAA controller. This plot is read in the way that for example for three agents the degree of sub-optimality in the optimal objective of the RIPP controller is a little above 1% in comparison to the objective of the SAA controller. The slight increase in sub-optimality for higher agent numbers stems from the fact that the RIPP constraints introduce some conservatism and the degree of conservatism adds up for more agents in the system.

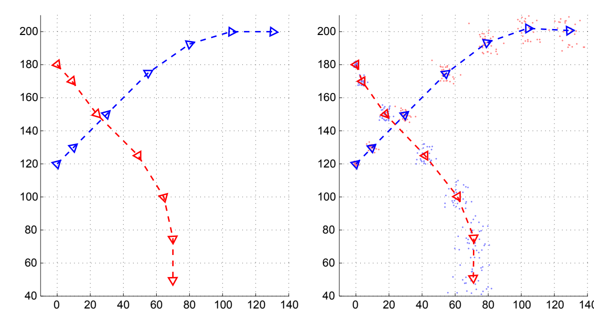

(a) Robust coupling constraints. (b) Probabilistic coupling constraints.

Fig. 3. Plans for two UAVs with robust and probabilistic coupling constraints for collision avoidance. The robust contraints allow the UAVs to come closer in the third time step than the probabilistic constraints.

D. Comparison with Robust Control

To the best of our knowledge, there are no other approaches to model chance constraints on the probability of inter-agent collisions, so we compare our approach to robust control approaches, namely to constraints from the robust MPC literature [17]. The work [17] assumes that the UAVs are affected by exogenous disturbances but does not model them stochastically. Instead, the authors assume them to be unknown but confined to lie in a bounded set. The aim of the robust controller is to guarantee robust constraint satisfaction, i.e., constraints on the states of the UAVs have to be satisfied under the strongest possible disturbance.

For the simulations, the model of the UAVs dynamics was as in Eq. 26 without stochastic disturbance $\underline{\nu}_t^i$ but with unknown but bounded disturbance with $\|\underline{\nu}_t^i\|_\infty \leq a_{\max}$. The objective remains the same and the bounds on the control inputs are set to be bounded by $\|\underline{u}^i\|_\infty \leq 10$. As in [17], we modeled the disturbances for robust control to be bounded by up to 10% of the control input. Hence, the disturbances acting on the states were bounded by $a_{\max} = 1$. The constraints on the positions of UAV 1 and UAV 2 resulting from the constraint tightening in [17] are

$$\|\underline{x}_t^1 - \underline{x}_t^2\|_\infty > \epsilon + 2\alpha(t) \ , \tag{27}$$

where $\alpha(1) = 0,\ \alpha(2) = 0,\ \alpha(t) = \|[1\ 0\ 0\ 0]\mathbf{LB}\|a_{\max}$ for $t > 2$ and $\mathbf{L} = \mathbf{A} + \mathbf{BK}$. The matrix $\mathbf{K}$ is a two-step nilpotent controller $\mathbf{K}$ for the system $(\mathbf{A}, \mathbf{B})$ and it can be checked that $\mathbf{K} = -\begin{bmatrix}\mathbf{I}_2 & 2\cdot\mathbf{I}_2\end{bmatrix}$ fulfills the requirements that $\mathbf{L}^2 = \mathbf{0}$.

In Figure 3 we plotted the open loop plans for two UAVs computed with robust constraints and the probabilistic constraints for collision avoidance. From the plot it becomes apparent that the robust contraints allowed the UAVs to come closer in the third time step than the probabilistic constraints did. This fact is also reflected in the collision probabilities for this time step: the robust plan had a probability of a collision of more than 28% in the third time step, the probabilistic plan of 0.7%. These collision probabilities were determined with 10^6 Monte-Carlo samples. The chance constraint bound for collision probabilities was set to $\delta_t^{1,2} = 1\%$ at each time step t.

It can be seen that the plans with the robust constraints allowed the UAVs to take higher risks, while the probabilistic constraints allowed to precisely control how cautious the UAVs behave or how much risk they take. It could be possible to adjust the robust constraints in such a way that the UAVs behave more cautiously for example by raising the bound $a_{\max}$ on the unknown but bounded disturbances. But in a setting with stochastic disturbance the question remains on how to adjust this bound such that a pre-specified collision probability can be guaranteed?

VII. CONCLUSIONS

In this work we proposed two methods to formulate chance constraints on the probability of inter-agent collisions in order to render them computationally tractable. The first formulation is a sample average approximation of the probability of collisions between agents. In the second formulation, we constructed a region of increased probability of presence (RIPP) for the uncertain positions of the agents and introduced constraints that these RIPPs do not overlap for differing agents. Because the construction of the RIPPs is based on a probabilistic inequality we were able to prove that controls that satisfy the RIPP constraints are automatically feasible for MA-MPC problem with chance constraints on inter-agent collision probabilities. The theoretical statements made on the feasiblity of the RIPP approach hold for uncertain states with arbitrary state distribution. Therefore, our approach is generally applicable and not limited to Gaussian state distributions. We compared both the sample average approximation and the RIPP constraints in simulations and showed that run-time of the sample average approximation can become unfeasible under real-time requirements while thr RIPP constraints proved to be fast. Furthermore, we provided empirical evidence that probabilistic constraints can be better suited than robust collision avoidance constraints in situations in which there is a stochastic model of the system disturbance.

ACKNOWLEDGEMENTS

This work was partially supported by the German Research Foundation (DFG) within the Research Training Group GRK 1194 “Self-organizing Sensor-Actuator-Networks”. Jan Calliess is grateful for funds via the UK EPSRC ”Orchid” project EP/I011587/1.

REFERENCES

[1] T. Schouwenaars, J. How, and E. Feron, “Receding horizon path planning with implicit safety guarantees,” in *Proceedings of the 2004 American Control Conference*. IEEE, 2004.
[2] N. Yilmaz, C. Evangelinos, P. Lermusiaux, and N. Patrikalakis, “Path planning of autonomous underwater vehicles for adaptive sampling using mixed integer linear programming,” *IEEE Journal of Oceanic Engineering*, vol. 33, no. 4, pp. 522–537, 2008.
[3] A. Prékopa, *Stochastic programming*. Springer, 1995.
[4] J. Calliess, D. Lyons, and U. Hanebeck, “Lazy auctions for multi-robot collision avoidance and motion control under uncertainty,” Dept. of Engineering Science, University of Oxford, Tech. Rep. PARG-01-11, 2011.
[5] D. Van Hessem, C. Scherer, and O. Bosgra, “LMI-based closed-loop economic optimization of stochastic process operation under state and input constraints,” in *Proceedings of the 40th IEEE Conference on Decision and Control*, 2002.
[6] M. Ono and B. C. Williams, “Iterative risk allocation: A new approach to robust model predictive control with a joint chance constraint,” in *Proceedings of 47th IEEE Conference on Decision and Control*, 2008.
[7] L. Blackmore and M. Ono, “Convex chance constrained predictive control without sampling,” in *Proceedings of the AIAA Guidance, Navigation and Control Conference*, 2009.
[8] L. Blackmore, M. Ono, A. Bektassov, and B. Williams, “A probabilistic particle approach to optimal, robust predictive control,” *IEEE Transactions on Robotics*, vol. 26, no. 5, 2010.
[9] M. Ono and B. C. Williams, “Decentralized chance-constrained finite-horizon optimal control for multi-agent systems,” in *Proceedings of 49th IEEE Conference on Decision and Control*, 2010.
[10] A. Schwarm and M. Nikolaou, “Chance-constrained model predictive control,” *AIChE Journal*, vol. 45, no. 8, pp. 1743–1752, 1999.
[11] J. R. Birge and F. V. Louveaux, *Introduction to stochastic programming*, ser. Springer series in operations research. New York: Springer, 1997.
[12] B. Pagnoncelli, S. Ahmed, and A. Shapiro, “Sample average approximation method for chance constrained programming: theory and applications,” *Journal of optimization theory and applications*, vol. 142, no. 2, pp. 399–416, 2009.
[13] CPLEX, “ILOG CPLEX User’s Manual.”
[14] D. Lyons, J. Calliess, and U. Hanebeck, “Chance-constrained model predictive control for multi-agent systems,” *Arxiv preprint arXiv:1104.5384*, 2011.
[15] P. Whittle, “A multivariate generalization of Tchebichev’s inequality,” *The Quarterly Journal of Mathematics*, vol. 9, no. 1, p. 232, 1958.
[16] “MIL-F-8785C, Military Specification,” Tech. Rep., 1980.
[17] A. G. Richards, “Robust constrained model predictive control,” Ph.D. dissertation, Massachusetts Institute of Technology, Department of Aeronautics and Astronautics, Cambridge MA, February 2005.