

Distributionally Robust RRT with Risk Allocation

Kajsa Ekenberg, Venkatraman Renganathan, and Björn Olofsson

Abstract—An integration of distributionally robust risk allocation into sampling-based motion planning algorithms for robots operating in uncertain environments is proposed. We perform non-uniform risk allocation by decomposing the distributionally robust joint risk constraints defined over the entire planning horizon into individual risk constraints given the total risk budget. Specifically, the deterministic tightening defined using the individual risk constraints is leveraged to define our proposed exact risk allocation procedure. Embedding the risk allocation technique into sampling-based motion planning algorithms realises guaranteed conservative, yet increasingly more risk-feasible trajectories for efficient state-space exploration.

I. INTRODUCTION

Motion planning under uncertainty becomes challenging when only limited information about the system uncertainty is known. Such lack of information adds complexity to the existing path planning problem formulation for finding guarantees on the safety of the path generated by algorithms that aim to address such uncertainty. For instance, critical robotics operations such as Mars Rover and rescue robot missions cannot afford the risk of obstacle collision given their high stakes and environmental uncertainty. Often assumptions (such as Gaussian uncertainties) are made in the name of tractability as in [1], [2] and they may lead to significant miscalculation of risk. Recently, this shortcoming due to non-Gaussian stochastic uncertainties in motion planning was considered in [3]–[7] using approaches like distributionally robust optimization (DRO) [8], [9] and conditional value-at-risk [10].

Though many risk-bounded path planning techniques work with stochastic uncertainties characterised by their moments, they suffer from unwanted conservatism as a result of the uniform risk allocation (URA) being used. That is, given a total risk budget for safety violation, it is a common practice to distribute it uniformly across all the obstacles and the planning horizon. The conservatism drawback of URA was identified in [11] and approached using a two-stage optimization method based iterative risk allocation strategy. This strategy has yielded promising, less conservative and guaranteed results for covariance steering problems, for example in spacecraft maneuvering as in [12], [13].

The conservatism that arises in motion planning because of the lack of exact information about the stochastic uncertainties should not restrict the ability of sampling-based

algorithms like RRT [14] to efficiently explore the state-space too much. To overcome this shortcoming, one can use non-uniform risk allocation technique as in [15]–[17], so that they do not exceed the risks allocated uniformly.

Contributions: We extend the DR-RRT algorithm presented in [3] by embedding our proposed risk allocation technique into it. Our main contributions are as follows:

- 1) We propose a new distributionally robust risk allocation technique called Exact Risk Allocation (ERA) for sampling-based motion planning algorithms that allocates as minimum risks as possible while respecting a given total risk budget (See Theorem 1).
- 2) We prove that all feasible paths with the uniform risk allocation of length $T_{\text{path}} \in \mathbb{N}_{\geq 1}$ and total risk budget $\Delta_{\text{path}} \in (0, 0.5]$ is also feasible with the ERA but the vice-versa is not necessarily true (See Theorem 2).
- 3) We demonstrate our proposed technique using simulation results and show that by switching from uniform risk allocation to ERA, it is possible to give the same risk guarantees for sampling-based motion planning algorithms while maintaining a reduced conservatism.

The rest of the paper is organized as follows: The main problem statement of risk-bounded motion planning with risk allocation is presented in §II. Then, the proposed Distributionally Robust Exact Risk Allocation (DR-ERA) algorithm is presented in §III. Subsequently, the embedding of DR-ERA into the sampling-based motion planning algorithm RRT is discussed in §IV. Then, the proposed idea is demonstrated using simulation results in §V. Finally, the paper is closed in §VI along with the directions for future research.

NOTATIONS & PRELIMINARIES

The set of real numbers and natural numbers are denoted by \mathbb{R} and \mathbb{N} , respectively. The subset of natural numbers between and including a and b with $a < b$ and beyond b with b included are denoted by $[a : b]$ and $\mathbb{N}_{\geq b}$, respectively. The operators \oplus, \setminus , and $|\cdot|$ denote the set translation, set subtraction and set cardinality, respectively. The transpose of a matrix A is denoted by A^T . An identity matrix of dimension n is denoted by I_n . For a non-zero vector $x \in \mathbb{R}^n$ and a matrix $P \in \mathbb{S}_{++}^n$ (here, \mathbb{S}_{++}^n denotes the set of positive definite matrices), let $\|x\|_P = \sqrt{x^T P x}$. A binary condition being true and false is denoted by \top and \perp , respectively.

II. PROBLEM FORMULATION

A. Robot & Environment Model

Our problem formulation follows the problem setup given in [3]. Consider a robot operating in an uncertain environment, $\mathcal{X} \subseteq \mathbb{R}^n$, with dynamic obstacles. The set of obstacles

This project has received funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation program under grant agreement No 834142 (Scalable Control). K. Ekenberg was a Master’s Thesis Student of the Department of Automatic Control LTH, Lund University, Sweden. V. Renganathan and B. Olofsson are with the Department of Automatic Control LTH, Lund University, Lund, Sweden. E-mail: ka2560ek-s@student.lu.se, (venkat, bjorn.olofsson)@control.lth.se.

is denoted as \mathcal{B} with $|\mathcal{B}| = N$. The robot model is given by a stochastic discrete-time linear time invariant system

$$x_{k+1} = Ax_k + Bu_k + w_k, \quad (1)$$

where $x_k \in \mathbb{R}^n$ and $u_k \in \mathbb{R}^m$ are the system state and input at time step k , respectively. The matrices A and B denote the dynamics matrix and the input matrix, respectively. The process noise $w_k \in \mathbb{R}^m$ is a zero-mean random vector that is independent and identically distributed over time. The distribution of w_k , namely \mathbb{P}_{w_k} , is unknown but belongs to a moment-based ambiguity set of distributions,

$$\mathcal{P}^w = \left\{ \mathbb{P}_{w_k} \mid \mathbb{E}[w_k] = 0, \mathbb{E}[w_k w_k^\top] = \Sigma_w \right\}. \quad (2)$$

The initial state x_0 is subject to a similar uncertainty model as the process noise, with its distribution belonging to a moment-based ambiguity set, $\mathbb{P}_{x_0} \in \mathcal{P}^{x_0}$, given by

$$\mathcal{P}^{x_0} = \left\{ \mathbb{P}_{x_0} \mid \mathbb{E}[x_0] = \hat{x}_0, \mathbb{E}[(x_0 - \hat{x}_0)(x_0 - \hat{x}_0)^\top] = \Sigma_{x_0} \right\}. \quad (3)$$

We assume the obstacles to perform a random walk around their initial position. That is,

$$\mathcal{O}_{ik} = \mathcal{O}_i^0 \oplus \hat{c}_{ik} \oplus c_{ik}, \quad \forall i \in \mathcal{B}, \quad (4)$$

where \mathcal{O}_{ik} denotes the position of the obstacle $i \in \mathcal{B}$ at time step k . The known shape of the obstacle is represented by $\mathcal{O}_i^0 \subset \mathbb{R}^n$, while \hat{c}_{ik} represents a known nominal translation. The location uncertainty and unpredictable motion of obstacle $i \in \mathcal{B}$ is represented by $c_{ik} \in \mathbb{R}^n$, which is a random vector with unknown distribution $\mathbb{P}_{c_{ik}}^c \in \mathcal{P}_{c_{ik}}^c$. The robot is expected to be in the free space at all time steps k . That is,

$$x_k \in \mathcal{X}_k^{\text{free}} := \mathcal{X} \setminus \bigcup_{i \in \mathcal{B}} \mathcal{O}_{ik}, \quad (5)$$

and the input of the robot is subject to the constraint $u_k \in \mathcal{U}$. Here, \mathcal{U} , \mathcal{X} and \mathcal{O}_{ik} are all assumed to be convex polytopes that can be represented by a conjunction of linear inequalities

$$\mathcal{U} = \{u_k \mid A_u u_k \leq b_u\}, \quad (6)$$

$$\mathcal{X} = \{x_k \mid A_x x_k \leq b_x\}, \quad (7)$$

$$\mathcal{O}_{ik} = \{x_k \mid A_{ik} x_k \leq b_{ik}\}. \quad (8)$$

B. Distributionally Robust Path Planning Problem

Problem 1. We seek to approximately solve the distributionally robust risk-constrained path planning problem. Given an uncertain initial state $x_0 \sim \mathbb{P}_{x_0}$ and a set of goal locations $\mathcal{X}_{\text{goal}} \subset \mathbb{R}^n$, we seek to find a feedback control policy $\pi = \{\pi_k\}_{k=0}^{T-1}$ such that applying the control inputs $u_k = \pi_k(x_k), k = [0 : T-1]$ yields a probabilistically feasible path from the initial state to the goal that minimises a finite-horizon cost function. That is,

$$\underset{\pi}{\text{minimize}} \quad \sum_{k=0}^{T-1} \ell_t(\hat{x}_k, \mathcal{X}_{\text{goal}}, u_t) + \ell_T(\hat{x}_T, \mathcal{X}_{\text{goal}}) \quad (9a)$$

$$\text{subject to} \quad (1), x_0 \sim \mathbb{P}_{x_0} \in \mathcal{P}^{x_0}, w_k \sim \mathbb{P}_w \in \mathcal{P}^w, \quad (9b)$$

$$(4), (5), u_k \in \mathcal{U}, c_{ik} \sim \mathbb{P}_{c_{ik}}^c \in \mathcal{P}_{c_{ik}}^c, \quad (9c)$$

$$\sup_{\mathbb{P}_{x_k} \in \mathcal{P}^{x_k}} \mathbb{P}_{x_k} \left[\bigwedge_{k=1}^T x_k \notin \mathcal{X}_k^{\text{free}} \right] \leq \Delta. \quad (9d)$$

Problem 1 is just a reformulated version of the problem in [3]. Here, $\ell_t(\cdot)$ is the stage cost function quantifying the distance to the goal set and actuator effort and it is expressed in terms of the robot mean state, \hat{x}_k . As (9d) is an infinite dimensional DR risk constraint, solving (9) exactly is practically hard and so we resort to approximate solutions using sampling-based motion planning algorithms. The constant $\Delta \in (0, 0.5]$ represents the user-prescribed total risk budget for the entire planning horizon, such that the worst-case probability of colliding with any of the N obstacles or being outside \mathcal{X} over the planning horizon should be at most Δ . As in [3], an LQR fixed affine feedback control policy given by $u_k = K_k x_k + g_k$ is used, and the state mean \hat{x}_k and covariance matrix Σ_{x_k} evolve as

$$\hat{x}_{k+1} = (A + BK_k)\hat{x}_k + Bg_k, \quad (10)$$

$$\Sigma_{x_{k+1}} = (A + BK_k)\Sigma_{x_k}(A + BK_k)^\top + \Sigma_w. \quad (11)$$

Note that (9d) can be decomposed into individual chance constraints for each obstacle and the state constraint \mathcal{X} at each time step. The individual risk bound for each obstacle $i \in \mathcal{B}$ and the constraints $j = 1, \dots, n_e$ defining \mathcal{X} at time step k , denoted by δ_{ik} and κ_{jk} respectively, should respect

$$\sum_{k=1}^T \sum_{i=1}^N \delta_{ik} + \sum_{k=1}^T \sum_{j=1}^{n_e} \kappa_{jk} \leq \Delta. \quad (12)$$

The following lemma is an adaptation of Theorem 1 in [3] with inclusion of time horizon from $t = 1, \dots, T$.

Lemma 1. If (12) holds true, then (9d) holds true as well if the worst-case probability of colliding with obstacle i and the worst-case probability of violating any one of $j = 1, \dots, n_e$ constraints defining \mathcal{X} at time step $k \in [1 : T]$ are

$$\sup_{\mathbb{P}_{x_k} \in \mathcal{P}^{x_k}} \mathbb{P}_{x_k}(x_k \in \mathcal{O}_{ik}) \leq \delta_{ik}, \quad \text{and}, \quad (13)$$

$$\sup_{\mathbb{P}_{x_k} \in \mathcal{P}^x} \mathbb{P}_{x_k}(a_{xj}^\top x_k \geq a_{xj}^\top c_{xj}) \leq \kappa_{jk}. \quad (14)$$

Proof. We know that $x_k \notin \mathcal{X}_k^{\text{free}} \iff \left\{ x_k \in \bigcup_{i=1}^N \mathcal{O}_{ik} \right\} \cup \{x_k \notin \mathcal{X}\}$. We denote the event of colliding with obstacle i at time step k as $C_{ik} := x_k \in \mathcal{O}_{ik}$ and similarly let the event of violating the j^{th} constraint defining the state constraint set \mathcal{X} at time step k be $D_{jk} := \{a_{xj}^\top x_k \geq a_{xj}^\top c_{xj}\}$. Then, the left-hand side of (9d) can be equivalently written as

$$\begin{aligned} & \sup_{\mathbb{P}_{x_k} \in \mathcal{P}^{x_k}} \mathbb{P}_{x_k} \left[\bigvee_{k=1}^T \left[\left\{ x_k \in \bigcup_{i=1}^N \mathcal{O}_{ik} \right\} \cup \{x_k \notin \mathcal{X}\} \right] \right] \\ & \leq \sup_{\mathbb{P}_{x_k} \in \mathcal{P}^{x_k}} \mathbb{P}_{x_k} \left[\bigvee_{k=1}^T \bigvee_{i=1}^N C_{ik} \right] + \sup_{\mathbb{P}_{x_k} \in \mathcal{P}^{x_k}} \mathbb{P}_{x_k} \left[\bigvee_{k=1}^T \bigvee_{j=1}^{n_e} D_{jk} \right] \\ & \leq \sum_{k=1}^T \sum_{i=1}^N \sup_{\mathbb{P}_{x_k} \in \mathcal{P}^{x_k}} \mathbb{P}_{x_k}[C_{ik}] + \sum_{k=1}^T \sum_{j=1}^{n_e} \sup_{\mathbb{P}_{x_k} \in \mathcal{P}^x} \mathbb{P}_{x_k}[D_{jk}] \\ & \leq \sum_{k=1}^T \sum_{i=1}^N \delta_{ik} + \sum_{k=1}^T \sum_{j=1}^{n_e} \kappa_{jk} \\ & \leq \Delta. \end{aligned}$$

Here, we applied the Boole's inequality [18] to get the second inequality, (13) and (14) to get the third inequality and (12) for the fourth inequality to obtain the desired result. \square

Future research will seek to reduce the conservatism resulting from the Boole's inequality by using sharper bounds such as the Kwerel's, Kounias' or Hunter's bounds [19], [20]. We now reformulate Problem 1 with individual risk bounds.

Problem 2. *We seek to approximately solve the following distributionally robust path planning problem with individual risk bounds as follows:*

$$\text{minimize}_{\pi, \delta} \quad \sum_{k=0}^{T-1} \ell_t(\hat{x}_k, \mathcal{X}_{\text{goal}}, u_t) + \ell_T(\hat{x}_T, \mathcal{X}_{\text{goal}}) \quad (15a)$$

$$\text{subject to} \quad (1), x_0 \sim \mathbb{P}_{x_0} \in \mathcal{P}^x, w_k \sim \mathbb{P}_w \in \mathcal{P}^w, \quad (15b)$$

$$(4), (5), u_k \in \mathcal{U}, c_{ik} \sim \mathbb{P}_{c_{ik}}^c \in \mathcal{P}_{c_{ik}}^c, \quad (15c)$$

$$\sup_{\mathbb{P}_{x_k} \in \mathcal{P}^x} \mathbb{P}_{x_k}(x_k \in \mathcal{O}_{ik}) \leq \delta_{ik}, \quad \forall i \in \mathcal{B}, \forall k \in [1:T], \quad (15d)$$

$$\sup_{\mathbb{P}_{x_k} \in \mathcal{P}^x} \mathbb{P}_{x_k}(a_{xj}^\top x_k \geq a_{xj}^\top c_{xj}) \leq \kappa_{jk}, \quad \forall j \in [1:n_e], \forall k \in [1:T], \quad (15e)$$

$$\sum_{k=1}^T \sum_{i=1}^N \delta_{ik} + \sum_{k=1}^T \sum_{j=1}^{n_e} \kappa_{jk} \leq \Delta. \quad (15f)$$

The only difference between Problems 1 and 2 is that Problem 2 is expressed with individual risk constraints and the allocated individual risks satisfy the total risk budget Δ .

III. DISTRIBUTIONALLY ROBUST RISK ALLOCATION

Allocating the individual risks in a non-uniform way while still enforcing the DR risk constraint (9d) can minimise the conservatism of the resulting path from source to the goal. Let us define the vector of all individual risk bounds as

$$\delta := [\delta_{11} \quad \dots \quad \delta_{NT}]^\top, \kappa := [\kappa_{11} \quad \dots \quad \kappa_{n_e T}]^\top. \quad (16)$$

A. Risk Treatment: Polytopic Obstacles & State Constraints

Since the obstacle $\mathcal{O}_{ik}, \forall i \in \mathcal{B}$ is a convex polytope, it can be represented by n_i hyperplanes. Collision with obstacle $i \in \mathcal{B}$ at time step k occurs if the position of the robot lies inside the obstacle, $x_k \in \mathcal{O}_{ik}$. This can be expressed as a conjunction of n_i linear constraints on the robot's position,

$$x_k \in \mathcal{O}_{ik} \iff \bigwedge_{j=0}^{n_i} a_{ij}^\top x_k < b_{ikj}. \quad (17)$$

The individual chance constraints given by (13) encode the fact that the worst-case probability of colliding with obstacle i at time step k should be at most δ_{ik} . That is,

$$\sup_{\mathbb{P}_{x_k} \in \mathcal{P}^x} \mathbb{P}_{x_k} \left[\bigwedge_{j=1}^{n_i} a_{ij}^\top x_k < a_{ij}^\top c_{ikj} \right] \leq \delta_{ik}, \quad (18)$$

where $c_{ikj} = \hat{c}_{ikj} + c_{ik}$ is a point on the j th constraint of obstacle \mathcal{O}_{ik} , with its first and second moments being \hat{c}_{ikj} and $\Sigma_{c_{jk}}$ respectively. The distributionally robust individual risk constraint in (18) can be handled through linear constraints

on the state mean \hat{x}_k defined using deterministic constraint tightening as in [3], [21]. That is,

$$a_{ikj}^\top \hat{x}_k \geq a_{ikj}^\top \hat{c}_{ikj} + \gamma_{ikj}^j, \quad (19)$$

$$\gamma_{ikj}^j(\delta_{ik}) := \sqrt{\frac{1 - \delta_{ik}}{\delta_{ik}}} \left\| (\Sigma_{x_k} + \Sigma_{c_{jk}})^{\frac{1}{2}} a_{ikj} \right\|_2, \quad (20)$$

where, γ_{ikj}^j is the deterministic constraint tightening of the j th constraint of obstacle $i \in \mathcal{B}$ at time k . To this end, we define Boolean quantities \mathbf{h}_{ik}^j and \mathbf{h}_{ik} that represent the mean state being outside the tightened j th constraint of \mathcal{O}_{ik} and outside the tightened obstacle \mathcal{O}_{ik} , respectively:

$$\mathbf{h}_{ik}^j = \begin{cases} \top, & (19) \text{ is satisfied} \\ \perp, & \text{otherwise,} \end{cases} \quad (21)$$

$$\mathbf{h}_{ik} = \begin{cases} \top, & \bigvee_{j=1}^{n_i} \mathbf{h}_{ik}^j = \top \\ \perp, & \text{otherwise.} \end{cases} \quad (22)$$

Here, (19) encodes the condition that the mean position of the robot should lie outside the tightened j th constraint of obstacle $i \in \mathcal{B}$ at time k to fulfill $\mathbf{h}_{ik}^j = \top$. A similar approach can be taken for treating the state constraints. The distributionally robust individual risk constraint in (15e) can be handled by linear constraints on the state mean \hat{x}_k defined using deterministic constraint tightening:

$$a_{xj}^\top \hat{x}_k \leq a_{xj}^\top c_{xj} - \underbrace{\sqrt{\frac{1 - \kappa_{jk}}{\kappa_{jk}}} \left\| \Sigma_{\hat{x}_k}^{\frac{1}{2}} a_{xj} \right\|_2}_{:= \eta_k^j(\kappa_{jk})}. \quad (23)$$

Similarly, we define Boolean quantities \mathbf{g}_k^j representing the mean state being inside the tightened j th constraint of \mathcal{X} :

$$\mathbf{g}_k^j = \begin{cases} \top, & (23) \text{ is satisfied} \\ \perp, & \text{otherwise.} \end{cases} \quad (24)$$

Here, η_k^j is the deterministic constraint tightening of the j th constraint of \mathcal{X} at time k and (24) encodes the condition that the mean position of the robot should lie inside the tightened state constraint set \mathcal{X} in order to fulfill $\mathbf{g}_k^j = \top$.

B. Exact Risk Allocation (ERA) Algorithm

The aim of ERA is to allocate as little risks δ_{ik} and κ_{jk} as possible for all obstacles $i \in \mathcal{B}$ and the constraints defining the state constraint set \mathcal{X} at all time steps k that fulfill the DR risk constraint in (15d) and (15e) respectively. Note that ERA cannot be done if the mean state \hat{x}_k is either inside the obstacle or outside \mathcal{X} as such paths will be deemed as non-feasible. Hence, we define the ERA problem with Boolean conditions for each constraints defining the obstacle $i \in \mathcal{B}$ and the constraints defining the state constraint set \mathcal{X} .

Problem 3. *Find the minimum risk δ_{ik} for $i \in \mathcal{B}$ at all time steps $k \in [1:T]$ such that $\mathbf{h}_{ik} = \top$, and the minimum risk κ_{jk} for which $\mathbf{g}_k^l = \top$ for $l = [1:n_e]$.*

The following theorem tells us how to obtain the required minimum risks δ_{ik} and κ_{jk} from (19) and (23), respectively.

Theorem 1. *The minimum risk for obstacle $i \in \mathcal{B}$ satisfying $\mathbf{h}_{ik} = \top$, at all time steps $k = 1, \dots, T$ is obtained by setting $a_{ikj}^\top \hat{x}_k = a_{ikj}^\top \hat{c}_{ikj} + \gamma_{ik}^j$ and is given by*

$$\delta_{ik}^* = \left(1 + \left(\frac{a_{ikj}^\top \hat{x}_k - a_{ikj}^\top \hat{c}_{ikj}}{\left\| (\Sigma_{x_k} + \Sigma_{c_{jk}})^{\frac{1}{2}} a_{ikj} \right\|_2} \right)^2 \right)^{-1}, \quad (25)$$

and the minimum risk for the j^{th} constraint defining \mathcal{X} satisfying $\mathbf{g}_k^j = \top, j = 1, \dots, n_e$ at all time steps $k = 1, \dots, T$ is obtained by setting $a_{xj}^\top \hat{x}_k = a_{xj}^\top c_{xj} - \eta_k^j$ and

$$\kappa_{jk}^* = \left(1 + \left(\frac{a_{xj}^\top c_{xj} - a_{xj}^\top \hat{x}_k}{\left\| \Sigma_{x_k}^{\frac{1}{2}} a_{xj} \right\|_2} \right)^2 \right)^{-1}. \quad (26)$$

Proof. Since Σ_k and $\Sigma_{c_{jk}}$ are known constants and $\sqrt{\frac{1-\delta_{ik}}{\delta_{ik}}}$ is a decreasing function of δ_{ik} , rearranging (21) for the case of $\mathbf{h}_{ik}^j = \top, j = 1, \dots, n_i$ leads to the below risk bound,

$$\delta_{ik} \geq \underbrace{\left(1 + \left(\frac{a_{ikj}^\top \hat{x}_k - a_{ikj}^\top \hat{c}_{ikj}}{\left\| (\Sigma_{x_k} + \Sigma_{c_{jk}})^{\frac{1}{2}} a_{ikj} \right\|_2} \right)^2 \right)^{-1}}_{:=\delta_{ik}^*}. \quad (27)$$

Similarly, rearranging (24) for the case of $\mathbf{g}_k^j = \top, j = 1, \dots, n_e$ leads to the following individual risk lower bound,

$$\kappa_{jk} \geq \underbrace{\left(1 + \left(\frac{a_{xj}^\top c_{xj} - a_{xj}^\top \hat{x}_k}{\left\| \Sigma_{x_k}^{\frac{1}{2}} a_{xj} \right\|_2} \right)^2 \right)^{-1}}_{:=\kappa_{jk}^*}. \quad (28)$$

□

IV. DISTRIBUTIONALLY ROBUST RRT WITH EXACT RISK ALLOCATION

In this section, we extend the sampling-based Distributionally Robust RRT (DR-RRT) algorithm in [3] which grows trees of state distributions while enforcing DR risk constraints, using the proposed ERA algorithm. Usually, DR-RRT employs the URA as it trivially satisfies (12), where each obstacle and time step are first assigned the same risk $\delta_{ik} = \frac{\Delta}{TN}$, and the assigned risks are then used to check the probabilistic feasibility of the generated path according to the constraint in (21). With ERA, the problem is tackled in the opposite way by *first* assigning risks δ_{ik} that fulfill the DR risk constraints in (21) and *then* checking if (12) holds.

A. Tree Expansion

Algorithm 1 outlines the DR-RRT tree expansion with Exact Risk Allocation incorporated and the readers are referred to [3] for information on DR-RRT tree expansion. Note that the trajectory generated from the LQR finite horizon steering function does not depend on the risk allocations δ_{ik} . In the next step, Exact Risk Allocation is applied to the generated trajectory, as outlined in Algorithm 2. The

ERA-function returns risk allocations δ_{ik} and κ_{jk} for all obstacles $i \in \mathcal{B}$ and all the constraints $j = 1, \dots, n_e$ defining the state constraint set \mathcal{X} at all time steps k along the trajectory. The risk allocation is done so that (19) and (23) are fulfilled and the total risk leading up to each time step is obtained by summing up all risk allocations δ_{ik} up to a certain time step, denoted as k^* . The path from N_{near} up to time step T_{steer} is then checked for distributionally robust feasibility, as outlined in Algorithm 3. If the path is feasible, the total cost J and the residual risk δ_{res} are calculated and used to assign a score to the path from the near node N_{near} as $\text{score}(N_{\text{near}}) = (\theta_J/J) + \theta_{\text{res}}\delta_{\text{res}}$, where $\theta_J, \theta_{\text{res}} \in [0, 1], \theta_J + \theta_{\text{res}} = 1$ are left to the user's choice to emphasize the cost and the residual risk appropriately. When paths from all near nodes that are DR-feasible have been assigned a score, the path with the best score is chosen and a new node and edge is added to the tree. The residual risk δ_{res} is also added to the node, which can in turn be re-allocated as described in subsection IV-B when steering from this node to a new sample. Feasible portions of the trajectories are also added to the tree in the same manner.

B. Feasibility Check

The feasibility check is based on the total risk allocated up to time step k , denoted by $\delta_{\text{tot}}(k)$. The risk constraints (15d)–(15f) have to hold for the entire planning horizon T and not just over the steering horizons $T_{\text{steer}} \in \mathbb{N}_{\geq 1}$. To assure this is the case, we begin by distributing the total risk budget Δ uniformly over all steering horizons according to $\Delta_{\text{steer}} = \frac{\Delta \cdot T_{\text{steer}}}{T}$, where Δ_{steer} is the risk budget over each steering horizon T_{steer} . An entire trajectory from a near node to the sample is deemed to be feasible, provided the total risk allocated over the steering horizon, $\delta_{\text{tot}}(T_{\text{steer}}) \leq \Delta_{\text{steer}}$. A similar reasoning can be applied to assure the feasibility of a portion of the steered path, from a near node up to a certain time step k . Then, the total risk allocated up to that time step, $\delta_{\text{tot}}(k)$, has to fulfill $\delta_{\text{tot}}(k) \leq \Delta_k$, where $\Delta_k := \frac{k \cdot \Delta_{\text{steer}}}{T_{\text{steer}}}$ is the uniformly allocated risk budget up to time step k . This means that a trajectory, or a portion of it, is considered feasible only when the total allocated risk (using ERA) does not exceed the corresponding total uniformly allocated risk. While this method has less conservatism than URA, there are still a lot of conservatism present from allocating the total risk budget uniformly over all steering horizons. This conservatism can be mitigated by reallocating residual risk of a horizon to the subsequent steering. If the entire risk budget Δ_{steer} or Δ_k is not used, such that $\delta_{\text{tot}}(T_{\text{steer}}) < \Delta_{\text{steer}}$ or $\delta_{\text{tot}}(k) < \Delta_k$, a residual for the newly generated node at time step k or T_{steer} can be created according to

$$\delta_{\text{res}} = \Delta_{\text{steer}} - \delta_{\text{tot}}(T_{\text{steer}}) \quad \text{or} \quad (29)$$

$$\delta_{\text{res}} = \Delta_k - \delta_{\text{tot}}(k). \quad (30)$$

These residual risks can then be reallocated to new trajectories generated from this node. When a new point x_s is sampled, the residual risk of the near node $\delta_{\text{res}}[N_{\text{near}}]$ can be allocated to the trajectory generated by steering from N_{near} to x_s . The total risk budget for the new trajectory or

its portion is then $\Delta_{\text{steer}} + \delta_{\text{res}}[N_{\text{near}}]$ or $\Delta_k + \delta_{\text{res}}[N_{\text{near}}]$, respectively. Then, the feasibility of the trajectory generated from N_{near} depends upon the relaxed risk budget constraints

$$\delta_{\text{tot}}(\mathbf{T}_{\text{steer}}) \leq \Delta_{\text{steer}} + \delta_{\text{res}}[N_{\text{near}}] \quad \text{or} \quad (31)$$

$$\delta_{\text{tot}}(k) \leq \Delta_k + \delta_{\text{res}}[N_{\text{near}}], \quad (32)$$

and the residual of N_{near} is added to the residual of newly created nodes originating from N_{near} . That is,

$$\delta_{\text{res}} = \Delta_{\text{steer}} + \delta_{\text{res}}[N_{\text{near}}] - \delta_{\text{tot}}(\mathbf{T}_{\text{steer}}) \quad \text{or} \quad (33)$$

$$\delta_{\text{res}} = \Delta_k + \delta_{\text{res}}[N_{\text{near}}] - \delta_{\text{tot}}(k). \quad (34)$$

Remarks: Note that the above risk allocation procedure still has some conservatism. An inevitable conservatism stems from the usage of Boole's inequality to decompose the joint risk constraint in (9d) into individual risk constraints. Though some trajectories are deemed to be infeasible and dismissed as they fail to satisfy (31), they could be potentially stored with the hope that they become feasible when they are connected with new trajectories such that the *combination* of the trajectories becomes risk-feasible. Albeit, such an effort would come at the expense of increased computational burden and memory storage along with the book-keeping to correctly identify feasible branches as near nodes to a random sample during the RRT tree expansion. For the ease of exposition, we decided not to implement the above mentioned aspects and only reallocate risk to future horizons. Interested readers are referred to [17] for additional details.

Theorem 2. *All DR-RRT paths feasible with the ERA of length T_{path} and total risk budget $\Delta_{\text{path}} \in (0, 0.5]$ is also feasible with the ERA but the opposite is not necessarily true.*

Proof. Without loss of generality, we present the proof assuming that the environmental borders given by (7) are not treated probabilistically. Consider a path that is feasible with ERA meaning that all risk allocations are assigned the same value $\delta_{\text{uni}} = \frac{\Delta_{\text{path}}}{N \cdot T_{\text{path}}}$ and they satisfy (12) and

$$\forall_{j=1}^{n_i} (a_{ij}^T \hat{x}_k - a_{ij}^T \hat{c}_{ikj} \geq \gamma_{ik}^j(\delta_{\text{uni}})), \quad \forall_{i \in [1:N]}, \quad \forall_{k \in [1:T_{\text{path}}]}.$$

On the other hand, ERA sets risk allocations δ_{ik} such that

$$\forall_{j=1}^{n_i} (a_{ij}^T \hat{x}_k - a_{ij}^T \hat{c}_{ikj} \geq \gamma_{ik}^j(\delta_{ik})), \quad \forall_{i \in [1:N]}, \quad \forall_{k \in [1:T_{\text{path}}]}.$$

Then, $\forall i \in \mathcal{B}, \forall k = 1, \dots, T_{\text{path}}$, and $j = 1, \dots, n_i$

$$\begin{aligned} \gamma_{ik}^j(\delta_{\text{uni}}) \leq \gamma_{ik}^j(\delta_{ik}) &\iff \sqrt{\frac{1 - \delta_{\text{uni}}}{\delta_{\text{uni}}}} \leq \sqrt{\frac{1 - \delta_{ik}}{\delta_{ik}}} \\ &\iff \delta_{ik} \leq \delta_{\text{uni}}. \end{aligned}$$

Further, the sum of all exact risk allocations satisfies

$$\sum_{k=1}^{T_{\text{path}}} \sum_{i=1}^N \delta_{ik} \leq \sum_{k=1}^{T_{\text{path}}} \sum_{i=1}^N \delta_{\text{uni}} = \sum_{k=1}^{T_{\text{path}}} \sum_{i=1}^N \frac{\Delta_{\text{path}}}{N \cdot T_{\text{path}}} = \Delta_{\text{path}}.$$

Thus, the path is also feasible with the ERA. Conversely, we just need to prove that there exists a path which is feasible with ERA but not with URA. Suppose that at time step $k, \exists i \in \mathcal{B}$ such that the given path with risk allocation $\delta_{ik} < \delta_{\text{uni}}$ is feasible with ERA meaning that $\sum_{k=1}^{T_{\text{path}}} \sum_{i=1}^N \delta_{ik} =$

$\bar{\Delta}_{\text{path}} < \Delta_{\text{path}}$. Since, $\sum_{k=1}^{T_{\text{path}}} \sum_{i=1}^N \delta_{\text{uni}} = \Delta_{\text{path}} > \bar{\Delta}_{\text{path}}$, we can conclude that $\exists j = 1, \dots, n_i$ for which $\gamma_{ik}^j(\delta_{\text{uni}}) < \gamma_{ik}^j(\delta_{ik})$, and the path will be deemed infeasible with URA as both (22) with $\gamma_{ik}^j(\delta_{\text{uni}})$ and the risk summation condition that $\sum_{k=1}^{T_{\text{path}}} \sum_{i=1}^N \delta_{\text{uni}} \leq \Delta_{\text{path}}$ cannot hold true simultaneously. Hence, there exist paths that are feasible with ERA but not with URA. \square

Algorithm 1 DR-RRT With ERA: Tree Expansion

Inputs: Tree \mathcal{T} , time k , $\mathbf{T}_{\text{steer}}, \theta_J, \theta_{\text{res}} \in [0, 1]$

$x_s = \text{sample}(\mathcal{X}_k^{\text{free}})$

$N_{\text{near}} = \text{NearestMNodes}(x_s, \mathcal{T}, M)$

for all N_{near} **do**

$(\hat{x}_{\text{path}}, \Sigma_{\text{path}}) = \text{Steer}(N_{\text{near}}, x_s, \mathbf{T}_{\text{steer}})$

$\delta, \kappa = \text{ExactRiskAllocation}(\hat{x}_{\text{path}}, \Sigma_{\text{path}}, \mathbf{T}_{\text{steer}})$

$\delta_{\text{tot}}(k^*) = \sum_{k=1}^{k^*} (\sum_{i=1}^N \delta_{ik} + \sum_{j=1}^{n_e} \kappa_{jk}), \forall k^* \leq T_{\text{steer}}$

if DRFeasible($\delta_{\text{tot}}(\mathbf{T}_{\text{steer}}), \delta_{\text{res}}[N_{\text{near}}]$) **then**

$J = J[N_{\text{near}}] + J(\hat{x}_{\text{path}}, \Sigma_{\text{path}})$

$\delta_{\text{res}} = \delta_{\text{res}}[N_{\text{near}}] + \Delta_{\text{steer}} - \delta_{\text{tot}}(\mathbf{T}_{\text{steer}})$

score(N_{near}) = $(\theta_J/J) + \theta_{\text{res}}\delta_{\text{res}}$

Select path $(\hat{x}_{\text{path}}, \Sigma_{\text{path}})$ from N_{near} with best score

$\mathcal{T}.\text{AddNode}(\hat{x}_{\text{path}}(\mathbf{T}_{\text{steer}}), \Sigma_{\text{path}}(\mathbf{T}_{\text{steer}}), \delta_{\text{res}})$

$\mathcal{T}.\text{AddEdge}(N_{\text{near}}, \hat{x}_{\text{path}}(\mathbf{T}_{\text{steer}}))$

for $k = 1 : T_{\text{steer}} - 1$ **do**

if DRFeasible($\delta_{\text{tot}}(k), \delta_{\text{res}}[N_{\text{near}}]$) **then**

$\delta_{\text{res}} = \delta_{\text{res}}[N_{\text{near}}] + \Delta_k - \delta_{\text{tot}}(k)$

$\mathcal{T}.\text{AddNode}(\hat{x}_{\text{path}}(k), \Sigma_{\text{path}}(k), \delta_{\text{res}})$

$\mathcal{T}.\text{AddEdge}(N_{\text{near}}, \hat{x}_{\text{path}}(k))$

Algorithm 2 ExactRiskAllocation

Inputs: Path $\hat{x}_{\text{path}}, \Sigma_{\text{path}}, \mathbf{T}_{\text{steer}}$

Output: Risk matrices $\delta \in \mathbb{R}^{N \times T_{\text{steer}}}, \kappa \in \mathbb{R}^{n_e \times T_{\text{steer}}}$

for $k = 1 : T_{\text{steer}}$ **do**

for $i = 1 : N$ **do**

Assign δ_{ik} satisfying (25).

for $j = 1 : n_e$ **do**

Assign κ_{jk} satisfying (26).

return δ, κ

Algorithm 3 DRFeasible

Inputs: total risk $\delta_{\text{tot}}(k)$, residual of $N_{\text{near}}, \delta_{\text{res}}[N_{\text{near}}]$

Output: true if DR-feasible, otherwise false

if $\delta_{\text{tot}}(k)$ satisfies (32) **then**

return true

else

return false

V. NUMERICAL SIMULATIONS

The environment for the simulation is a square area $[0, 50]^2 \text{ m}^2$ where $N = 10$ rectangular obstacles are randomly placed. As for the simulations of DR-RRT in [3], a




Fig. 1. DR-RRT tree in \mathbb{R}^2 with URA using $\Delta = 0.1$ along with zoomed-in covariances.




Fig. 2. DR-RRT tree in \mathbb{R}^2 with ERA using $\Delta = 0.1$.




Fig. 3. DR-RRT tree in \mathbb{R}^2 with ERA using $\Delta = 0.02$.

robot with discrete-time stochastic double-integrator dynamics having a mass of 1 kg is considered. The initial position is $[0, 0]$ m and the initial velocity is zero. The dynamics and input matrices are

$$A = \begin{bmatrix} I_2 & dtI_2 \\ 0_{2 \times 2} & I_2 \end{bmatrix}, \quad B = \begin{bmatrix} \frac{dt^2}{2}I_2 \\ dtI_2 \end{bmatrix}, \quad (35)$$

with $dt = 0.1$ s. The robot state is the position and velocity along each axis, with the corresponding force as inputs. The covariance matrices of the initial state and the disturbance are chosen as in [3],

$$\Sigma_{x_0} = 10^{-3} \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix}, \quad \Sigma_w = 10^{-3} \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 2 & 1 \\ 0 & 0 & 1 & 2 \end{bmatrix}.$$

All obstacles are static and treated as deterministic, so that all uncertainty comes from the unknown state of the robot. The robot is treated as a point mass and the bounds on the environment are not treated probabilistically. As in [3], the steering from a near node to a sample x_s is done by solving a discrete-time linear quadratic optimal control problem to compute the affine state feedback policy that minimises

$$\sum_{k=0}^{T_s-1} \hat{e}_k^T Q \hat{e}_k + u_k^T R u_k + \hat{e}_{T_s}^T Q \hat{e}_{T_s}, \quad (36)$$

where $\hat{e}_k = \hat{x}_k - x_s$ and $T_s = T_{\text{steer}}$, $Q = 40I_4$ and $R = 0.1$. The quadratic optimal cost-to-go function is also used as the distance metric in the selection of the nearest tree nodes. In all simulations, the trajectories of the mean state \hat{x}_k is represented by lines and the uncertainty is represented by ellipses of one standard deviation, derived from the covariance Σ_{x_k} . Note that in all the simulations, the ellipses are in the range of 0.01 m (visibly too small). The planning horizon is $T = 1000$ and the steering horizon is $T_{\text{steer}} = 10$. The risk budget for the entire planning horizon T is denoted as Δ . Three DR-RRT trees with 1000 samplings are simulated, namely: 1) Using URA and risk budget $\Delta = 0.1$, 2) Using ERA and risk budget $\Delta = 0.1$, and 3) Using ERA and risk budget $\Delta = 0.02$. The risk allocation of ERA was done using results from Theorem 1. Besides from the risk allocation and risk budget, everything in the trees and environment are exactly the same, including the random sampling points. This is to get a fair comparison of the different trees. With URA, the same risk is allocated for all obstacles and time steps,

such that $\delta_{ik} = \frac{\Delta}{T \cdot N} = \frac{0.1}{1000 \cdot 10} = 10^{-5}$. With ERA, the risk budget for a steering horizon is $\Delta_{\text{steer}} + \delta_{\text{res}}$, where δ_{res} is the residual of the node from which steering is done and

$$\Delta_{\text{steer}} = \frac{\Delta \cdot T_{\text{steer}}}{T} = \begin{cases} 10^{-3}, & \text{when } \Delta = 0.1 \\ 2 \times 10^{-4}, & \text{when } \Delta = 0.02. \end{cases}$$

TABLE I

AVERAGE RESULTS OF 1000 INDEPENDENT SIMULATIONS EACH WITH 1000 ITERATIONS

Methodology	Δ	# Nodes ($ \mathcal{T} $)
URA	0.10	3101
ERA	0.10	8175
ERA	0.02	3348

Discussion: Figs. 1 and 2 show that DR-RRT with ERA generates less conservative paths than DR-RRT with uniform risk allocation when the same risk budget $\Delta = 0.1$ is used. With the same Δ , DR-RRT with ERA explores the state-space more efficiently than with URA (consequence of Theorem 2), still by having the same risk guarantees. Fig. 3 illustrates how DR-RRT with ERA can be used with lower risk budget $\Delta = 0.02$ and still generate paths with comparably a similar degree of conservatism as DR-RRT with a URA using a higher risk budget of $\Delta = 0.1$. The comparison in Table I is a good indication of the above observation. In general, the selection of best Δ is not straightforward as it depends upon T , Σ_{x_0} , Σ_w and the steering law being used. Overall, ERA gives the same risk guarantees for sampling-based motion planning algorithms, while maintaining a reduced conservatism and with almost no additional computational complexity resulting from the risk-allocation procedure.

VI. CONCLUSION

An extension of the sampling-based probabilistically complete DR-RRT motion planning algorithm in [3] with an optimal risk allocation was presented. We proved that our risk allocation based embedding technique realises guaranteed conservative, yet increasingly more risk feasible trajectories for efficient state-space exploration. That is, all DR-RRT paths feasible with the URA are feasible with the ERA but not vice-versa. Future research will aim to design a slightly more involved risk allocation based embedding into the DR-RRT* algorithm from [4], [22] to guarantee both risk-bounded and asymptotically optimal trajectories.

REFERENCES

- [1] B. Luders, M. Kothari, and J. How, "Chance constrained RRT for probabilistic robustness to environmental uncertainty," in *AIAA Guidance, Navigation, and Control Conference*, 2010, pp. 8160–8181.
- [2] B. D. Luders, S. Karaman, and J. P. How, "Robust sampling-based motion planning with asymptotic optimality guarantees," in *AIAA Guidance, Navigation, and Control (GNC) Conference*, 2013, pp. 5097–5122.
- [3] T. Summers, "Distributionally robust sampling-based motion planning under uncertainty," in *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2018, pp. 6518–6523.
- [4] V. Renganathan, I. Shames, and T. H. Summers, "Towards integrated perception and motion planning with distributionally robust risk constraints," *IFAC-PapersOnLine*, vol. 53, no. 2, pp. 15 530–15 536, 2020.
- [5] W. Han, A. Jasour, and B. Williams, "Non-gaussian risk bounded trajectory optimization for stochastic nonlinear systems in uncertain environments," in *2022 International Conference on Robotics and Automation (ICRA)*, 2022, pp. 11 044–11 050.
- [6] P. Lathrop, B. Boardman, and S. Martinez, "Distributionally safe path planning: Wasserstein safe RRT," *IEEE Robotics and Automation Letters*, vol. 7, no. 1, pp. 430–437, 2021.
- [7] B. Luders and J. How, "Probabilistic feasibility for nonlinear systems with non-gaussian uncertainty using RRT," in *Infotech@ Aerospace*. AIAA, 2011, pp. 1589–1602.
- [8] P. Mohajerin Esfahani and D. Kuhn, "Data-driven distributionally robust optimization using the wasserstein metric: Performance guarantees and tractable reformulations," *Mathematical Programming*, vol. 171, no. 1-2, pp. 115–166, 2018.
- [9] A. R. Hota, A. Cherukuri, and J. Lygeros, "Data-driven chance constrained optimization under Wasserstein ambiguity sets," in *2019 American Control Conference (ACC)*. IEEE, 2019, pp. 1501–1506.
- [10] A. Hakobyan and I. Yang, "Wasserstein distributionally robust motion planning and control with safety constraints using conditional value-at-risk," in *2020 IEEE International Conference on Robotics and Automation (ICRA)*, 2020, pp. 490–496.
- [11] M. Ono and B. C. Williams, "Iterative risk allocation: A new approach to robust model predictive control with a joint chance constraint," in *IEEE Conference on Decision and Control*, 2008, pp. 3427–3432.
- [12] J. Pilipovsky and P. Tsiotras, "Covariance steering with optimal risk allocation," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 57, no. 6, pp. 3719–3733, 2021.
- [13] V. Renganathan, J. Pilipovsky, and P. Tsiotras, "Distributionally robust covariance steering with optimal risk allocation," *arXiv preprint arXiv:2210.00050*, 2022.
- [14] S. M. LaValle and J. J. Kuffner Jr, "Randomized kinodynamic planning," *The International Journal of Robotics Research*, vol. 20, no. 5, pp. 378–400, 2001.
- [15] M. Ono and B. C. Williams, "An efficient motion planning algorithm for stochastic dynamic systems with constraints on probability of failure," in *AIAA Guidance, Navigation, and Control Conference*, 2008, pp. 1376–1382.
- [16] M. P. Vitus and C. J. Tomlin, "On feedback design and risk allocation in chance constrained control," in *Conference on Decision and Control and European Control Conference*. IEEE, 2011, pp. 734–739.
- [17] K. Ekenberg, "Distributionally robust risk-bounded path planning through exact spatio-temporal risk allocation," *MS Thesis, TFRT-6174, Lund University, Sweden*, 2022.
- [18] D. Hunter, "An upper bound for the probability of a union," *Journal of Applied Probability*, vol. 13, no. 3, pp. 597–603, 1976.
- [19] S. M. Kwerel, "Bounds on the probability of the union and intersection of m events," *Advances in Applied Probability*, vol. 7, no. 2, pp. 431–448, 1975.
- [20] A. Patil and T. Tanaka, "Upper and lower bounds for end-to-end risks in stochastic robot navigation," *arXiv preprint arXiv:2110.15879*, 2021.
- [21] G. C. Calafiore and L. E. Ghaoui, "On distributionally robust chance-constrained linear programs," *Journal of Optimization Theory and Applications*, vol. 130, no. 1, pp. 1–22, 2006.
- [22] V. Renganathan, S. Safaoui, A. Kothari, B. Gravell, I. Shames, and T. Summers, "Risk bounded nonlinear robot motion planning with integrated perception & control," *Artificial Intelligence*, vol. 314, p. 103812, 2023.