2026. Volume 10, Number 1 Pages 146-160

UDC: 007.52:519.71
DOI: https://doi.org/10.30546/09090.2026.001.20.2056

RISK-BOUNDED REAL-TIME COLLISION AVOIDANCE UNDER PERCEPTION UNCERTAINTY

ELMIR ADILZADE

[1]Azerbaijan Oil and Industry University, Baku, Azerbaijan

elmira21620@sabah.edu.az

ARTICLE INFO

Article history:
Received:2025-05-26
Received in revised form:2025-06-20
Accepted:2025-07-26
Available online:2026-06-23

Keywords:
Chance-constrained MPC;
Perception uncertainty;
Multi-robot navigation;
Kalman filtering;
2010 Mathematics Subject Classifications: 93E20, 90C15, 68T40, 93C95

ABSTRACT

The uncertainty of perception makes it difficult to avoid collision in real-time for mobile robots because the states of obstacles will appear displaced, as well as future interactions with other agents. The main purpose of this study is to establish a perception of risk by framing multi-robot navigation as a risk-aware optimization problem where safety values are approximated through a risk-inflected deterministic surrogate that is parameterized by an α that users can set. Also constructed is a stochastic receding horizon controller by propagating Gaussian-belief-state representations for moving obstacles and transforming probabilistic considerations of separation into deterministic robustness constraints through confidence-region inflation. This formulation thus provides a practical approximation of chance-based safety; However, it does not provide a formally certified probabilistic assurance based on all of the complicated non-linear system dynamics.

The framework has been evaluated in a two-dimensional multi-agent simulation comprised of multiple robots, with moving obstacles and noise in the measurements used to infer the location of the obstacles; the locations of the obstacles were filtered as they were tracked online. The empirical results indicate that increasing the risk coefficient creates more conservative trajectories, which increases the amount of clearance between agents and the obstacle, while at the same time, maintaining real-time task completion and the times it takes to solve the problem were within acceptable ranges for short replanning intervals. This framework provides a systematic means for tuning the tradeoff of safety versus efficiency given uncertain perception and therefore is applicable to autonomous systems that must adjust their motion planning to accommodate risk in real-time.

1. INTRODUCTION

Autonomous robots need to avoid collisions to operate properly in populated, dynamic environments like warehouses, hospitals, and cities. Many of the problems associated with this difficulty have been made worse by the increasing numbers of agents, the presence of moving obstacles, and the requirement to replan quickly with limited computational resources. With the increasing number of deployments, there have also been increasing expectations of safety because of the potential for physical damage caused by collision events and the potential for non-compliance with regulations, while too much safety leads to lower throughput and usability than desired. Because of this, there are increasing expectations that method for avoiding collisions will develop systematic, adjustable methods for ensuring safety while still maintaining real-time performance.

The greatest source of uncertainty during navigation in the real world is perception uncertainty. The presence of sensor noise, occlusions and imperfect object tracking algorithms will generate increasingly large deviations from the actual state of the obstacle, which will be most evident when extrapolated into the future for a number of seconds. In probabilistic robotics, the estimate of the state of a system is sometimes represented as a belief distribution that allows for uncertainty-based reasoning to be done through Bayesian filtering and through Kalman filtering when linear/Gauss assumptions are used [4] [7]. Nevertheless, there continue to be challenges associated with incorporating belief uncertainty into the real-time motion planning process because the safety constraints are generated probabilistically and are therefore fraught with difficulty in relation to the execution of an appropriate enforcement method for a given optimisation problem.

Collision avoidance systems can be roughly classified into two groups: reactive geometric methods and optimization-based predictive controllers. Reactive methods like velocity-obstacle formulations help facilitate quick, local decisions to prevent imminent crashes between many agents [8]. While these methods provide effective means of utilizing high update rate data, they face significant obstacles in capturing long-horizon objectives, explicitly modeling uncertainty in moving obstacles, and structuring risk control.

In contrast to reactive systems, optimization-based approaches such as model predictive control (MPC) provide a structured means to balance task progress against constraints by employing constrained optimization techniques each time a control cycle executes. Stochastic MPC extends these techniques to uncertain systems with disturbances making it possible to reason probabilistically about constraint violation [5]. The chance-constrained formulation in robotic planning approximates probabilistic obstacle avoidance through the use of deterministic constraints to embed confidence intervals for predicted states [1] [6]. However, translating these approaches into multi-robot, real-time environments with perception uncertainty is a still unsolved methodological issue and using traditional probabilistic guarantees would require impractically high levels of computational power.

An important outstanding question is how estimation uncertainty and multi-agent coupling interact when navigating in real-time. When there are multiple robots operating in the same environment, the possible actions of each robot depend on uncertain predictions of how obstacles will move/where the other robots will go. Many of the current approaches to this problem ignore uncertainty completely, treat other robots as fixed, or assume the "worst" case regarding the motion of other robots. Because of this these approaches may produce solutions that are less efficient than they could be. Thus, there is a gap in the ability to design controllers that (i) use estimates of uncertainty when navigating, (ii) control the amount of risk associated with collisions, and (iii) compute their solutions in a timeframes compatible with re-planning frequently.

This study addresses this gap by developing a real-time collision avoidance controller that is risk-aware and considers uncertainty in the perception of objects in the environment. The objectives of the study are as follows: First, develop a method for modeling moving obstacles using a Gaussian belief state based on Kalman filtering techniques and updating that state in real-time. Second, create a multi-agent receding horizon optimization problem that finds solutions for avoiding colliding with moving obstacles via using a risk-inflated parameterized deterministic surrogate model (using a user defined α as a parameter to the surrogate). Finally, test the safety efficiency trade-of within a multi-agent simulation with multiple moving obstacles.

Our primary contribution is to combine belief-state estimation and risk-modulated receding horizon planning for multi-agent collision avoidance into a single system. As part of this proposed framework, we introduce a practical method for estimating belief states by utilizing a known deterministic surrogate function within the context of standard nonlinear optimization to achieve real-time solutions. Although we have been inspired by principles of chance-constrained control, the way we implemented our formulation does not mathematically guarantee probabilistic performance for any full-nonlinear dynamics of the underlying system; our formulation generally yields a conservative estimate of behaved errors; conservative between conservative and risk modulated depending on α specified in our optimization model. We also provide empirical evidence on how changing risk coefficients affect margin of clearance and computational complexity within a simulated environment.

Therefore, research that examines risk-modulated collision avoidance in dynamic environments where perception uncertainty is a factor has become very timely and critical. The structured and modifiable methodology for controlling risk offers a solid foundation for deploying autonomous robotic systems in environments characterized by high levels of uncertainty, while avoiding designing systems with overly- conservative worst-case scenarios based on static assumptions during the execution of online operations.

2. LITERATURE REVIEW AND PROBLEM STATEMENT

2.1 Change-constrained obstacle avoidance

The use of chance constraints allows for the formal quantification of the probable rate of unsafe events occurring due to stochastic uncertainties. Ono and Williams [6] describes an efficient approach to motion planning for stochastic dynamic systems where the probability of failure along planned trajectories can be bounded, thus allowing for probabilistically safe specifications to be incorporated into an optimization framework. This work is significant in that it shows how to include probability-of-failure constraints in trajectory generation without having to enumerate all stochastic outcomes.

Blackmore [1] extended chance-constrained planning to obstacle avoidance by planning over the future distribution of the vehicle state and ensuring that the probability of intersecting with obstacles is no greater than a specified threshold. This work is important in that it offers a practical transformation of probabilistic conditions for collision into deterministic constraints that can be solved as single-variate Gaussian random variables. However, this work has primarily focused on single vehicle applications and has not addressed the coupling of multiple vehicles working together simultaneously.

2.2 Stochastic model predictive control and risk management

Chance constraints are formal methods for establishing certainty bounds on the likelihood of unsafe occurrences in environments with random chance. Ono and Williams [6] propose a viable strategy for planning motion through a stochastic dynamic system by modelling the expected positions of objects within a planned path while guaranteeing that the expected failure probabilities of those objects will be <50%. The authors point out that this provides an example of how to include a probabilistic constraint on failure into a planning process without explicitly listing the random outcome space.

Blackmore [1] takes the concept of chance-constrained planning further, to include isolation from expected collisions with obstacles, through future expected states of the planner and residual states of the planner, or by ensuring that the expected collision probability for an obstacle is <50%. This research provides a means to convert probabilistic criteria to deterministic constraints for the single-agent case; however, it does not adequately consider the coupling effect of having multiple agents moving at the same time.

2.3 Perception uncertainty and belief-state estimation

Perceptual noise is a major source of risk during real worldwide travel due to sensory noise, displaced markers, and inaccurate location tracking. Errors introduced by these factors will compound over time, particularly when predictions of an object's predicted position extend for several seconds into the future. In order to minimize this risk the approach adopted in this research will be to create a probabilistic model of state estimation that provides a distribution of probability for the object's current position rather than a deterministic coordinate.

Kalman [1960] [4] first provided a mathematical description for linear filter state estimation systems that uses Gaussian noise to estimate position, and since then this model has become authoritative for use in the establishment of continuous position tracking through robotics. In the linear state space model each object being estimated is assigned a fixed velocity and has an associated state vector $s = [x, y, v_x, v_y]^T$. Therefore, at each time step k the sensor will measure a noise-corrupted position z_k . The measurement error is assumed to be Gaussian with mean 0, variance σ_{meas}= 0.25 m, consistent with the traditional linear filtering model. These sensor measurements will be passed through the Kalman filter, which incorporates the sensor measurement noise into the belief state and alters the belief state in order to provide an updated position estimate (mean) and updated variance (covariance).

Thrun outlined how probabilistic methods can be used to represent beliefs and methods for integrating information from multiple sensors, estimation processes, and decision-making [7]. This research is relevant because a controller needs to have corresponding uncertainty estimates that will depend on the accuracy of those estimates, and those estimates can be calculated through multiple means. In order to incorporate prediction uncertainty into the planning task, the Kalman filter state is propagated through the motion model in a forward manner, providing a series of predicted means and covariances associated with the location of a potential obstacle. In this way, the predicted location may be used to convert the requirements for probabilistic separation into certain constraints by increasing the effective radius of that obstacle based on the covariance and risk factor provided by the user.

2.4 Multi-agent collision avoidance and remaining gap

The Reciprocal Velocity Obstacle framework offers an efficient geometric means of providing reciprocal collision avoidance among multiple agents. It is notable because of its ability to provide scalability to many agents in real time. However, it does not provide explicit probabilistic bounds of collision risk and assumes at least locally (based on time) deterministic states of agents along with short time period interaction rules.

Overall, much of the literature suggests that chance-constrained planning provides a principled approach to risk management with uncertain conditions. At the same time, SMPC provides a natural online optimization framework for addressing on-line collision avoidance. Yet a significant unresolved problem exists for the development of a real-time controlled multi-robot collision avoidance system using chance-constrained optimization: A controller must function in conjunction with an online belief state estimation method and a computationally efficient chance-constrained optimization algorithm that allows collision risk bounds to be explicitly placed in the multi-agent coupling of predicted agent interaction trajectories. This supports the need for further research into a risk-bounded real-time collision avoidance system with uncertain perceptions.

3. THEORETICAL FRAMEWORK AND RESEARCH METHODOLOGY

3.1 Object of research, hypothesis, and assumptions

The study of the multi-robot navigation system takes place on two-dimensional environments that have dynamic objects that can be moving or partially known. The system allows multiple robots to reach their respective goals while keeping a safe distance away from one another and avoiding collisions with moving or dynamic objects in an uncertain environment. The primary focus of the research is to provide a way for real-time trajectory generation to have probabilistic safety assurances while the robots are navigating through an environment where their perception could be less than perfect.

The main hypothesis of the research is that by including uncertainty in perception into the planning process using probabilistic state estimation and risk-based safe distance constraints, the controllers of the robots will perform better at avoiding collisions than if the controllers used deterministic planning methods without considering uncertainties in their environment. It is expected that the robot controller should adjust the amount of safety margin it uses when passing other robots in relation to the estimated risk of a collision by considering obstacle uncertainties in its model. This will allow the robot to be able to trade safety for motion efficiency without reducing the real-time computational feasibility of the control method.

In order to provide the necessary accuracy for real-time predictive optimization, several general modeling assumptions have been adopted in this research project.

Robot motion has been modeled as first-order discrete-time kinematic with bounded planar velocities to simplify the complicated higher-order non-linearities associated with holonomically mobile robots without losing a great deal of accuracy with regards to their major translational dynamics (i.e. the position and velocity would be the only motion parameters taken into account). Actuator performance was also considered by creating control input limits that are representative of actuator capabilities and will result in physically possible trajectories.

Moving obstacles have been modeled on a linear constant velocity state-space form with added Gaussian process noise, while the obstacle's acceleration is considered an unobservable state variable that is implied by the process noise applied (which also represents unknown or partially known disturbances incurred during the motion of the obstacle). The constant-velocity assumption is a good compromise between giving the ability to predict an obstacle's future position too far into the future (greater than the length of time the obstacle will be in motion) while providing sufficient computational simplicity to allow for a short amount of time in predicting an obstacle's future position.

Measurement error is assumed to be random with respect to actual measurements and will be treated as Gaussian distributed to conform with conventional linear filtering theory [4], [7]. The obstacle state for every obstacle will consequently be recursively estimated using Kalman filtering methods under the usual linear-Gaussian measurement and dynamical assumptions. Thus, at every time t, the uncertainty for each obstacle's location is represented as a multivariate Gaussian with a mean and covariance.

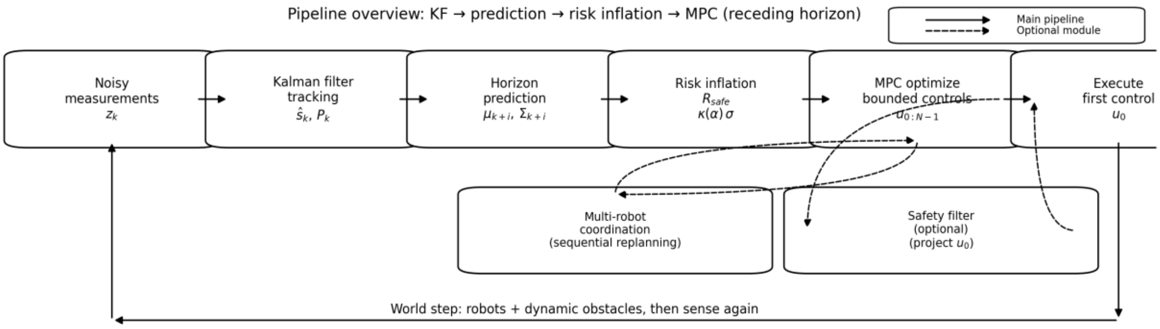

Fig 1. *System pipeline for real-time collision avoidance under perception uncertainty.*

The approach to filtering occurs in real time by using filtering techniques to obtain estimates and uncertainty for every obstacle in addition to propagating those values across the planning horizon to produce predicted estimates and uncertainty (i.e., mean and covariance). These predictions assist with calculation of an inflation margin associated with the uncertainty that will define the safety constraint/penalty in a receding-horizon MPC optimization. The first control action is executed and then the loop reinitiates at the next time step; optional coordination among robots and a safety filter may occur during execution. (Fig 1.)

This Gaussian approach allows compact representation of uncertainty during propagation along the planning horizon and serves as a basis for establishing risk-bounded safety constraints in the optimization component of the trajectory optimization component.

The use of a Gaussian representation facilitates the compact propagation of uncertainty throughout the planning horizon while creating a structured basis for trajectory optimization layer risk-bounded safety constraints. The uncertainty is expressed in terms of mean and covariance, preserving the probabilistic information while retaining computational tractability that is suitable for real-time application of model predictive control (MPC).

3.2 Simulation environment and robot dynamics

A continuous 2D workspace, or bounded rectangle $W \subset R^2$, is used for the evaluation, with several robots and moving dynamic obstacles. All entities are modeled as circular discs to minimize computational effort required for collision checking with regard to the robots (Rrobot = 0.25 m) and dynamic obstacles (Robs = 0.35 m).

Each robot and obstacle has an initial position that is sampled to ensure that there will not be an overlap during the execution of the robots' trajectories, therefore any collisions that occur will be due to dynamic interaction between the robots and not because of a failure to provide a feasible solution before the start of executing the robots' path.

The initial configuration of the multi-robot system is shown in Figure 2 according to a chance-constrained MPC formulation with a risk level of 0.95; there are three robots and two dynamically moving obstacles in the scenario. The colour of the markers represents the current position of a robot, while the cross markers represent the robot's assigned goal location. The grey circles represent nominal static position(s) of dynamic obstacle(s) while the large black circles represent the uncertainty-inflated safety region (from the chance constraints).

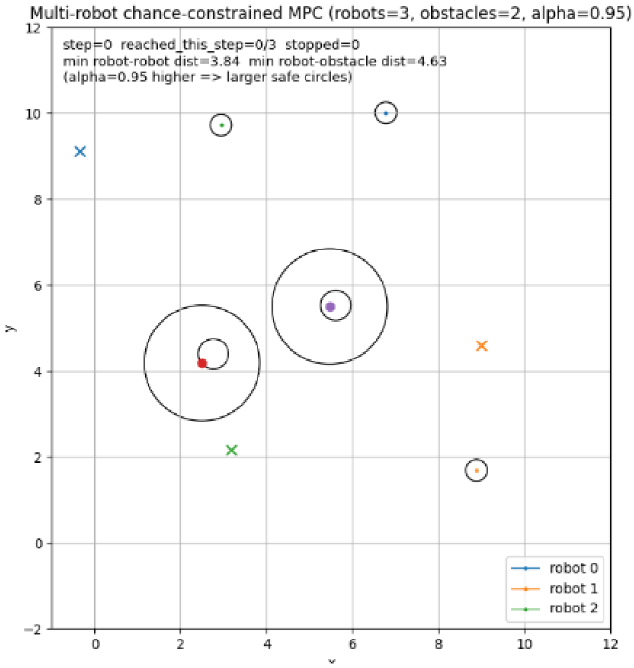

Fig 2. Simulation environment (3 Robots and 2 obstacles)

The trajectory optimization procedure is implemented in a receding-horizon manner with a fixed replanning interval (Δt=0.2 s), and a total number of planning steps for optimizing will consist of 8 steps (total time prediction of 8xΔt=1.6 s). Thus, there is a good balance between cognitive prediction behaviour and adequate computational capacity in the design phase of a robot's control path.

The discrete time model that is used to represent how the robot moves is a first-order integrator model (i.e. single integrable continuous kinematic motion) as represented by (1)

$$p_{k+1} = p_k + \Delta t * u_k \qquad (1)$$

This type of linear dynamical model is used very frequently in creating the model predictive control (MPC) for holonomic mobile robots because of its linear nature and ability to support real-time optimization capabilities for many of the tasks required to be accomplished throughout an entire robot's operation.

The physical constraint on the control inputs is defined in (2):

$$|u_{k,x}| \leq v_{max}, |u_{k,y}| \leq v_{max}, \text{ where } v_{max} = 1.0\ m/s \qquad (2)$$

These bounds shows us the activator limitations and ensures dynamic motion. The single-integrator assumes velocity tracking and shows high-order dynamics such as speed limits and constraints. This simplification is suitable for studying collision avoidance behavior and risk-aware path planning strategies, as it reduces the effect of uncertainity modelling from lower-level controlling.

3.3 Obstacle state estimation under perception uncertainty

Each of the moving obstacles in space is contained in four-dimensional state space with constant velocity, i.e., $s = [x, y, v_x, v_y]^T$. At every time step, Hikaru produces a noisy position measurement $z_k = [x_{meas}, y_{meas}]^T$ from the real position of an obstacle in space. A Kalman filter then looks at the previous belief state and produces an updated belief state with the resulting estimated mean and covariance of the updated position measurement. The measurement noise for each axis has a standard deviation of σ_{meas} = 0.25 m, while the process noise features diagonal terms with respect to the position and velocity components. The resulting belief state provides an estimate of the current position of the obstacle and the uncertainty associated with predicting that obstacle's future position over a N-step planning horizon.

When planning with predictive uncertainty, the Kalman filter state is propagated into the future using the obstacle motion model, which generates a sequence of N position mean and covariance predictions ($\mu k+i$ and $\Sigma(k+i)$) for all $i = 1, \ldots, N$, of all obstacles. This propagation of the newly generated belief states is consistent with the principles of planning within the belief space and provides the probabilistic parameters necessary for use in chance-constrained collision avoidance methods.

3.4 Risk-bounded collision model and deterministic approximation

Safety is defined as having minimal distance between every robot and obstacles, and between pairs of robots, for the prediction horizon. There is uncertainty in the position of the obstacles. The amount of separation between robots and obstacles is determined by the risk-inflation coefficient (α) ∈ (0, 1), where a higher value of α means a more conservative confidence region will be created. Instead of providing a strictly probabilistic collision-prevention constraint, the following formulation provides a deterministic proxy for the collision-prevention constraint that was created by inflating the confidence region for an obstacle based on the obstacle's position distribution (assumed to be Gaussian) and double-precision floating-point chi-square distribution.

Assuming the obstacle has a position distribution that is Gaussian, a conservative, deterministic approximation for the effective radius of the obstacle is generated by inflating the effective radius of the obstacle based on the chi-square distribution with two degrees of freedom. The confidence inflation factor ($\kappa(\alpha)$) is defined as follows, where the notation $\chi^2_2(.)$ is the chi-square quantile for the chi-square distribution with two degrees of freedom, and is the area under the two-dimensional chi-square distribution in the positive quadrant and between two values.

$$\kappa(\alpha) = \sqrt{(\chi^2_2(\alpha))} \qquad (2)$$

The obstacles' covariance Σ is represented by their largest principal standard deviation $\sigma_{max} = \sqrt{(\lambda_{max}(\Sigma))}$ in the implemented planner. The risk-inflated safety radius is given by $R_{safe} = R_{robot} + R_{obs} + b_{obs} + \kappa(\alpha)\ \sigma_{max}$, where $b_{obs} = 0.15$m serves as an additional buffer to account for any unmodeled effects. For robot-robot interactions, a deterministic separation radius $R_{rr} = 2R_{robot} + b_{rr}$ with $b_{rr} = 0.10$ m, this reflects the assumption that self-localization uncertainty of robots is negligible compared to the perception uncertainty of the obstacles in the current simulation.

The result is a deterministically approximated safety margin that is adjustable: increasing α results in an increase in R_{safe} and therefore produces a more conservative trajectory with respect to the clearance of obstacles. Although based on concepts similar to chance constrained formulations, this structure does not impose a validated probability bound on collision events using complete system dynamics; rather, it provides a conservative and computationally tractable substitute for creating trajectories using real time planning, while providing an easy to interpret risk adjustment.

3.5 Receding-horizon optimization and multi-robot coordination

Each robot solves a receding-horizon optimization problem at every control cycle, solving for a set of velocity commands {u0, …, uN−1}. This optimization problem includes three objectives in relation to the prediction horizon: reaching a goal, controlling effort and avoiding collisions. Collisions are avoided by applying a smooth soft penalty when predicted distances are less than the risk-inflated safety distance contained in Section 3.4. This formulation promotes separation, while maintaining the differentiability of the formulation, and will thus work with gradient-based optimization subject to box constraints on control inputs.

Since avoidance of collisions is created as a soft penalty instead of a hard constraint, strict feasibility of the separation constraint cannot be assured mathematically at all times during the optimization process. Rather safety is encouraged through penalized weighting, which when coupled with conservative inflation of the safety radius will ensure that robots behave practically collision-free in simulation while maintaining numerical stability and guaranteeing real-time solvability.

Multi-robot collaboratively controlling themselves is done via sequential planning by optimally controlling all robots in a predetermined order at each timestep. During its optimization, a robot must consider the previously predicted trajectory for the other robots as dynamic and deterministic bounds on its optimization. The above design guarantees scalability while including near-term interaction relations, without using centralized, joint optimization techniques. After an optimization, the first control input is executed for each robot, and the process of determining the optimization for all robots continues at the next timestep, as is consistent with using model predictive control techniques.

The nonlinear optimization problem is solved using the limited-memory BFGS (L-BFGS) algorithm with bound constraints (L-BFGS-B). Each optimization problem is limited to a maximum of 18 iterations per solve and to component-wise bounds $u_k \in [-v_{max}, v_{max}]^2$ for all components. A warm-start condition is created by moving the previous optimal solution forward in time, which helps to improve convergence rate and is capable of executing in real time throughout the given replanning interval.

3.6 Evaluation protocol and metrics

To evaluate performance, we conduct multiple simulations while setting fixed random seeds to guarantee identical initial conditions and movements of the obstacles. The main outcome of interest will be whether the robots successfully complete their assigned tasks within a specified number of time steps. In addition, we will quantify safety considerations as the minimum distance between two robots or a robot and an obstacle at their disk centers during each run. Finally, we will quantify the computational feasibility of this process as the optimization time (for each robot) to generate a control sequence for each of the robots for each time step using both mean solve time and 95th percentile solve time.

The main variable in an experiment will be the risk bound (α), allowing for an evaluation of the trade-offs between safety and efficiency. In all experiments unless otherwise indicated a total of three robots and two moving obstacles will be used with dynamic and uncertainty parameters specified in Sections 3.2–3.4.

4. RESULTS

4.1 Real-time optimization performance

The controller executed a replanning interval of $\Delta t = 0.2$ s and a horizon length of $N = 8$. The nonlinear optimization algorithm was able to produce successful control sequences for each active robot at each time step throughout all of the evaluated risk levels. Tables summarizes the average and 95th percentile of the time taken to find solutions to each per-robot planning call during each run. Each of the columns in this table represents an attribute of each planning run:

1. **α** - The α is the risk inflation coefficient used in the deterministic surrogate model. This coefficient regulates how much bigger the confidence region is enlarged based on the model prediction of where the obstacles will be. A larger value would indicate that the robot requires a larger margin of separation to behave conservatively. This should not be viewed as a proven probability of no collisions occurring.
2. **Success** - Refers to whether the robots that were active during the simulation reached their goal locations without experiencing a solver failure or an observed collision.
3. **Time to finish (s)** - is the total simulated time that was required for all robots to reach their goals, and is an indication of overall task performance based on the identified risk level.
4. **Min d(R,R) (m)** - Is the smallest distance between any two robots that occurred at any time during the entire simulation period. This value provides an indication of how close the robots actually came to each other at the closest point, and is used as an indicator of safe distance from other robots.
5. **Min d(R,O) (m)** - Minimum distance recorded between a mobile robot and obstacles throughout the experiment. This metric indicated the shortest distance achieved between any of the mobile robots and moving obstacles when evaluating performance of the robots to avoid dynamic obstacles while they are working.
6. **Mean solve (ms)** - Average amount of time (in milliseconds) to complete a single MPC optimization for one robot at every timestep during the period of time being tested. This metric indicates how much processing overhead was typical for all robots.
7. **P95 solve (ms)** - If we take the total time of all solves and rank them from lowest to highest, 95% of all solves will fall under this number. This metric can indicate the worst-case processing time periodically and is critical to determining if real time solutions can be provided given specific timing constraints.
8. **MPC solves** - Total number of calls to the MPC optimization were made during the simulation for this test period. The number of solves depends entirely on the duration of the mission and the number of robots in use and represents the total workload for the mission planner.

Table 1. *Performance summary across risk levels in a 3-robot, 2-obstacle scenario (seed = 7).*

α	Success	Time to finish (s)	Min d(R,R) (m)	Min d(R,O) (m)	Mean solve (ms)	P95 solve (ms)	MPC solves
0.90	Yes	8.2	2.83	2.41	58.7	150.56	41
0.95	Yes	8.2	2.96	2.65	59.26	154.02	41
0.99	Yes	8.6	3.12	3.24	69.25	142.27	43

Table 2. *Performance summary across risk levels in a 5-robot, 2-obstacle scenario (seed = 7).*

α	Success	Time to finish (s)	Min d(R,R) (m)	Min d(R,O) (m)	Mean solve (ms)	P95 solve (ms)	MPC solves
0.90	Yes	12.6	0.594	1.994	144.86	318.59	63
0.95	Yes	13.6	0.588	2.292	143.56	291.97	68
0.99	Yes	20.2	0.581	1.056	94.94	184.13	101

Table 3. *Performance summary across risk levels in a 5-robot, 4-obstacle scenario (seed = 7).*

α	Success	Time to finish (s)	Min d(R,R) (m)	Min d(R,O) (m)	Mean solve (ms)	P95 solve (ms)	MPC solves
0.90	Yes	16.2	0.908	1.598	264.84	427.47	81
0.95	Yes	16.8	0.570	2.037	344.51	580.90	84
0.99	Yes	36.0	0.563	0.501	241.59	241.59	180

4.2 Clearance distances under varying risk bounds

For both robot-to-robot and robot-to-obstacle interactions, the entire execution time horizon was recorded for minimum separation distances. The separation distance is defined as the smallest center-to-center distance recorded at any simulation step.

For the three robot simulation (Table 1), increasing the confidence level from a confidence level of 0.90 to $\alpha = 0.99$ resulted in a corresponding increase in all of the separation distances. The robot to obstacle minimum distance increased from 2.41 meters at a confidence level of 0.90 to 3.24 meters at a confidence level of 0.99; similarly, the minimum robot to robot separation distance increased from 2.83 meters at a confidence level of 0.90 to 3.12 meters at a confidence level of 0.99. The expected behavior is consistent with the theoretical properties of chance-constrained formulations, where increasing the confidence level increases the uncertainty associated with the safety regions and encourages more conservative strategies for avoiding issues.

On the other hand, the five robot configurations demonstrate non-linear, density-dependent behavior. For example, when using two obstacles, the minimum robot to obstacle distance was increased from 1.99 meters ($\alpha = 0.90$) to 2.29 meters ($\alpha = 0.95$); however, it was ultimately reduced to 1.06 meters ($\alpha = 0.99$). A similar trend was observed when using four obstacles; for example, the minimum robot to obstacle distance was increased from 1.60 meters ($\alpha = 0.90$) to 2.04 meters ($\alpha = 0.95$); however, it was reduced substantially to 0.50 meters ($\alpha = 0.99$). The inter-robot separation also demonstrated similar non-monotonic behavior.

Congestion-created coordination limitations explain this resulting behavior. As the value of α increases, the area occupied by obstacles' area will expand, thus increasing safety margins. In open environments, this expanding area translates directly into a larger clearance, as indicated by the spacing; however, in multi-robot systems with many other robots and obstacles nearing one another, the large areas of inflated obstacles will create less available space for the robots to maneuver, and the level of interaction between the two systems will be heightened. This will force the robots to execute their tasks in close proximity to each other during the coordination stages of their operation and will result in lower than normal minimum distances between the center of one robot and the center of another, even when there are high levels of probabilistic assurance.

The distance measurements were made by means of the total distances from the centre of one object (robot or obstacle) to the centre of the next object in the same group. To determine effective distance clearance, the combined radius of the entities being measured should be subtracted from the overall distance between centres; that is, using the minimum recorded distance ($\alpha = 0.99$ with 5 robots and 4 obstacles) all measured distances were above the distance referenced by the collision threshold, so there was no collision.

The results show that the ε parameter (risk) adjusts the degree of spatial conservatism in a predictable manner in sparse configurations, while producing highly dependent spatial buffer zones in less sparse configurations (density) of multiple agents.

4.3 Task completion under risk constraints

To assess the effectiveness of these robot navigation strategies, measured time was used for completion when robots reached their final locations. Measuring completion times provided insight into how well the robots will use the navigation methods as the level of risk for not meeting their goals increases; therefore, measuring completion times provides a detailed picture of the efficiency of the robots' navigation strategy under all three levels of risk and at increasing levels of environmental complexity.

In the three robot case (as summarized in Table 1), the completion times for the robots were stable across all levels of risk, with the robots taking 8.2 seconds to complete their work at levels of risk of α = 0.90 or α = 0.95, while robots took 8.6 seconds to complete their work when at α = 0.99. Thus, the maximum difference in completion times for the three robots is 0.4 seconds across all levels of risk, thus indicating that only minor degradation in the overall efficiency of the navigation strategy can result from strong safety guarantees in low-density environments.

Conversely, the five robot case demonstrated considerable variability in sensitivity to changes in the level of risk. Changes in completion times were also greater for robots navigating around two obstacles; at α = 0.90, the robots took 12.6 seconds, while at α = 0.95 they took 13.6 seconds and at α = 0.99 took 20.2 seconds to complete their work. In contrast, when four obstacles were added, the coverage times of the robots took substantially longer; at α = 0.90, 16.2 seconds were needed to navigate effectively compared to α = 0.95, which took 16.8 seconds, and at α = 0.99, the robots took 36.0 seconds to navigate effectively. This difference in coverage times illustrates the non-linear interaction of risk proliferation and the complexity of multi-agent coordinating tasks, resulting in longer planning time and times for resolving conflicts in high-density environments.

Risk-bounded planning behavior has been shown through an increase in task duration with increasing values of α. As you increase your confidence level, you will increase the amount of obstacle uncertainty by a larger factor which means you will have to reduce the amount of available navigation space. In situations with low density, the increase in the amount of obstacle uncertainty causes detours and slightly longer paths. When you consider a situation with a high density, the increase of obstacle uncertainty will increase the amount of interaction between all of the robots. Therefore the amount of time taken to reach globally spatially consistent behaviour will be increased at least fivefold, even though the cost of optimizing the trajectory is bounded per step.

All robots were successfully able to reach their assigned destination and there were no collisions. In each of the runs, the robot minimum separation distances (d) were still above the geometric collision threshold and therefore provided probabilistically safe guarantees despite the increase in completion time due to high-density environments.

The results of this study indicate that the risk parameter α is a controllable method to balance efficiency and probabilistically safe execution of tasks by multiple autonomous robots. While the trade-off between efficiency and probabilistic safety in environments with low to medium density can be predicted and considered moderate, in environments of high density where multiple robots are operating simultaneously, you must take nonlinear effects of scale into account when determining the risk boundary for a real-world deployment.

5. DISCUSSION

The results show that through comprising real-time risk-bounded collision avoidance by using a receding-horizon optimizer that combines both an online belief-state estimation with an inflated-risk safety model, you can achieve this goal. In three robot, two obstacle collision avoidance, there is an increase in confidence level alpha results in an increase in minimum separation distances based on the theoretical interpretations of chance constrained formulations stating that a higher requirement of confidence will increase the uncertainty dependent feasible set [1] [5]. Importantly, increasing conservatism did not prevent goal achievements on these occassions.

In addition to that, during five-robots under numerous obstacles, we did not retain the monotonic relationship between alpha and spatial clearance we had previously seen at lower densities. In multiple robot scenarios, specifically where there were four or more obstacles, the clearance behavior exhibits a non-linear characteristic. Moderate increases in alpha resulted in increased minimum clearances however; at greater confidence levels, we see a reduction in minimum clearances due to congestion. The increased size of the inflated obstacle regions reduced the maneuverability option and increased co-operative coupling of the agents and led to a need for tighter coordination. This behavior indicates that the effects of probabilistic conservatism and interaction density do interact with one another, and that the stronger risk bounds do not necessarily yield linear increases in minimum actual clearances in high density environments.

The data was collected and analyzed based on task efficiency and there is very little variation between completion times (8.2 s to 8.6 s) of low density configurations across risk levels. In contrast, the five robots demonstrated an inconsistent increase in time across risk levels, with two obstacles causing completion time to increase from 12.6 s at $\alpha = 0.90$ to 20.2 s at $\alpha = 0.99$ and four obstacles causing completion times to increase from 16.2 s to 36.0 s. Thus, with the increased risk of inflation there is greater coordination overhead resulting in an increase in the replanning cycle time proportional to the increase in number of agents in complex environments, taking longer to resolve the conflict.

Interestingly, there was no increase in the mean solve time per replanning step with respect to alpha. In some of the dense scenarios, greater levels of confidence caused a longer mission total, but produced similar or lower average solve times per replanning step when compared to the lower confidence level scenarios. Therefore, the main source of performance degradation is through the increased complexity of coordination and the addition of planning steps to those agents as opposed to the cost associated with optimizing at each iteration. Regardless of configuration, the average solve time remained within the bounds of real-time feasibility for a 0.2 s replanning period but increased dramatically with the number of agents and obstacles present.

Reciprocal Velocity Obstacles [8] are considered a reactive approach and provide high update rates and scalability at the cost of having no explicit way to guarantee probabilistic risk. The presented framework provides a means to directly tune the safety margin using α, and explicitly includes obstacle covariance through online estimation for the use of obstacle covariance in risk assessment. This explicit coupling of the uncertainty to the obstacle will cause the safety area to automatically grow based on the added uncertainty to the perception system. Although computational costs remain reasonable for risk-adaptive systems, there are some scalability challenges when working with density-coupled systems.

The limits of this research include several assumptions made for the systems presented. The obstacle uncertainty is assumed to be Gaussian distributed and is assumed to only account for principal components of uncertainty in the summary. Therefore, the assumption of using principal standard deviation inflation may not represent an accurate Gaussian distribution for many types of objects with strong anisotropic covariance. There is no explicit way to model the uncertainty in the robot state; thus the risk guarantees are based only on the uncertainty of object perception which cannot fully take into account the entire closed loop system uncertainty. The sequential planning of the adaptive risk creates order dependencies on the risk adaptive planning and while the ability to estimate and predict cleared obstacles with respect to the robot will show some level of partial coupling, there will not be a global optimally or symmetry for the adaptive risk planning.

An additional methodological limitation is the use of a soft penalty approach to avoid collisions. Soft penalty constraints provide efficient gradient-based optimization but can create strict feasibility issues based on penalty weightings. For scenarios with densely-coupled systems when minimum clearances between objects become small, using explicit equality constraints or barrier-function based approaches may provide better theoretical guarantees than that provided by a penalty function approach [2] [3].

Future exploration must include robot localization uncertainty, develop the framework to include more complex dynamic models than single integrator models; and develop decentralized or distributed optimization techniques that address order effects and promote scalability. To better assess how the analysis of Gaussian uncertainty relates to real-world robot deployments, additional experimental proof of concept must be made using heterogeneous robotic platforms with heterogeneous sensor modalities. In addition, confirmation through practical experimentation, can facilitate a measurement method for calibrating probabilistic risk metrics to operational safety benchmarks.

6. CONCLUSION

The study used an optimization-based approach to evaluate multi-robot collision avoidance using online Belief State Estimation and uncertainty-based safety modeling in a receding horizon framework as a basis for creating a probabilistic measure of effective risk through the use of probabilistic risk modulation via the α parameter. Results indicate that probabilistic risk modulation can be used to modulate safety constraints while being able to accomplish tasks reliably at real-time.

As density conditions decrease, the behavior of robots is fairly predictable with regard to the ability to maintain minimum separation distances and increase task completion time (i.e., three robots at 8.2–8.6 s) for increases in α. The preceding hypothesis is in line with previous research which has demonstrated that confidence-region based deterministic surrogate approaches exhibit this same type of behavior; higher confidence requirement leads to larger feasible safety sets and leads to more conservative motion strategies; therefore this result is likely to be true across all density conditions.

However, the results of the experiments conducted in the higher density scenarios showed a nonlinear scaling effect in terms of time to complete a task for five robots in multiple obstacles; the majority of the trials with density showed an increase in task duration of up to 36.0 s for the densest density trial (i.e., with five robots). In addition, the results exhibited an increase in spacing between robots not only due to increased congestion leading to increased coordination due to coordination constraints, but also resulted in non-monotonic behavior with respect to minimum separation distance. These results also demonstrate that the effect of probabilistic conservative motion strategies does not translate uniformly into larger actual clearances at higher density levels, but rather acts to modify the coordination dynamics.

From the computation perspective, the per-step optimization times across all density conditions remained well within the real-time limits associated with a 0.2 s replanning interval. The primary source of degradation from the performance in the higher density trials was primarily due to the increased level of coordination (or multiple replanning steps) and not directly due to the integrated increase in optimization costs for each individual iteration.

Unlike Reciprocal Velocity Obstacles [8], which achieve high (100Hz) update rates at the expense of structured probabilistic risk guarantees, the system presented here has tunable and explicit safety modulations via alpha and includes obstacle covariance from online estimation in an explicit fashion. This coupling of uncertainties will allow for adaptive expansion of safety regions as the uncertainty from perception increases.

There are several limitations to the framework presented above. The obstacles are modeled with Gaussian uncertainty and conservatively summed, the robot's state uncertainty is not considered, and there is an order dependency with respect to sequential planning that does not guarantee a globally optimal solution. In addition, while collision avoidance is achieved via soft penalties, there is no formal certification of a chance constraint with respect to a full nonlinear dynamic model, i.e., via strict inequality constraints.

Future work should explore extending the presented framework to include uncertainty in localization, richer dynamic models, and developing distributed optimization techniques to enhance scalability in dense multi-agent settings. Experimental validation using physical robot platforms would also lend credibility to the calibration of the probabilistic risk parameters in relation to operationally established safety standards.

In conclusion, the experimental results provide evidence that it is possible to perform uncertainty aware model predictive control with tunable probabilistic risk bounds in real time and provide effective safety modulations while demonstrating that the density of interactions is an important factor in determining scalability and emergent spatial behavior.

REFERENCE

1. Blackmore, L., Ono, M., & Williams, B. C. (2011). Chance-constrained optimal path planning with obstacles. IEEE Transactions on Robotics, 27(6), 1080–1094. https://doi.org/10.1109/TRO.2011.2161160
2. Calafiore, G. C., & Campi, M. C. (2006). The scenario approach to robust control design. IEEE Transactions on Automatic Control, 51(5), 742–753. https://doi.org/10.1109/TAC.2006.875041
3. Campi, M. C., & Garatti, S. (2008). The exact feasibility of randomized solutions of uncertain convex programs. SIAM Journal on Optimization, 19(3), 1211–1230. https://doi.org/10.1137/07069821X
4. Kalman, R. E. (1960). A new approach to linear filtering and prediction problems. Journal of Basic Engineering, 82(1), 35–45. https://doi.org/10.1115/1.3662552
5. Mesbah, A. (2016). Stochastic model predictive control: An overview and perspectives for future research. IEEE Control Systems Magazine, 36(6), 30–44. https://doi.org/10.1109/MCS.2016.2602087
6. Ono, M., & Williams, B. C. (2008). An efficient motion planning algorithm for stochastic dynamic systems with constraints on probability of failure. In Proceedings of the Twenty-Third AAAI Conference on Artificial Intelligence (AAAI-08) (pp. 1376–1382). AAAI Press.
7. Thrun, S., Burgard, W., & Fox, D. (2005). Probabilistic robotics. MIT Press.
8. van den Berg, J., Lin, M., & Manocha, D. (2008). Reciprocal velocity obstacles for real-time multi-agent navigation. In Proceedings of the IEEE International Conference on Robotics and Automation (ICRA) (pp. 1928–1935). https://doi.org/10.1109/ROBOT.2008.4543489