

Risk-Aware Non-Myopic Motion Planner for Large-Scale Robotic Swarm Using CVaR Constraints

Xuru Yang¹, Yunze Hu¹, Han Gao¹, Kang Ding¹, Zhaoyang Li², Pingping Zhu³, Ying Sun⁴ and Chang Liu¹

Abstract—Swarm robotics has garnered significant attention due to its ability to accomplish elaborate and synchronized tasks. Existing methodologies for motion planning of swarm robotic systems mainly encounter difficulties in scalability and safety guarantee. To address these limitations, we propose a Risk-aware swarm motion planner using conditional Value-at-Risk (ROVER) that systematically navigates large-scale swarms through cluttered environments while ensuring safety. ROVER formulates a finite-time model predictive control (FTMPC) problem predicated upon the macroscopic state of the robot swarm represented by a Gaussian Mixture Model (GMM) and integrates conditional value-at-risk (CVaR) to ensure collision avoidance. The key component of ROVER is imposing a CVaR constraint on the distribution of the Signed Distance Function between the swarm GMM and obstacles in the FTMPC to enforce collision avoidance. Utilizing the analytical expression of CVaR of a GMM derived in this work, we develop a computationally efficient solution to solve the non-linear constrained FTMPC through sequential linear programming. Simulations and comparisons with representative benchmark approaches demonstrate the effectiveness of ROVER in flexibility, scalability, and safety guarantee.

I. INTRODUCTION

Large-scale swarm robotic systems comprised of numerous autonomous and interacting robots are witnessing a surge in popularity due to their superior robustness and flexibility in applications such as target detection [1], cooperative object transport [2], and search and rescue [3]. In recent years, there has been a growing interest in developing motion planning techniques for large-scale swarm robots [4, 5].

Present large-scale swarm motion planning approaches can be categorized into *microscopic* methods and *macroscopic* methods. Microscopic methods consider agent-wise interaction and coordination and directly generate individual control inputs for each robot. Despite the satisfying performance in

*This work was sponsored by Beijing Nova Program (20220484056) and the National Natural Science Foundation of China (62203018).

All correspondences should be sent to Chang Liu.

¹Xuru Yang, Yunze Hu, Han Gao, Kang Ding, and Chang Liu are with the Department of Advanced Manufacturing and Robotics, College of Engineering, Peking University, Beijing 100871, China (xuru.yang@stu.pku.edu.cn; hu_yun_ze@stu.pku.edu.cn; hangao-coe@pku.edu.cn; kangding@stu.pku.edu.cn; changliucoe@pku.edu.cn).

²Zhaoyang Li is with the Department of Automation, Tsinghua University, Beijing 100080, China (lizhaoya21@mails.tsinghua.edu.cn)

³Pingping Zhu is with the Department of Computer Sciences and Electrical Engineering (CSEE), Marshall University, Huntington, WV 25755, USA (zhup@marshall.edu).

⁴Ying Sun is with the School of Electrical Engineering and Computer Science, The Pennsylvania State University, State College, PA 16802, USA (ybs5190@psu.edu).




Fig. 1. Figure (a) illustrates the task of large-scale robotic swarm motion planning through a cluttered environment, with dashed lines indicating swarm transport trajectories and white dots denoting individual robots. Obstacles are represented as gray polygons. Figure (b) illustrates the first-person view in the red circular sector in Figure (a).

small and medium-scale robot swarms [6, 7], these methods always suffer from dramatic increases in computational overhead as the swarm size increases, hindering their applications to large-scale robotic systems.

Macroscopic approaches overcome the scalability issue by treating the swarm as an entity instead of focusing on individual robots. One representative category of macroscopic approaches is based on the mean-field theory [8, 9], which represents the swarm state as the average of kinodynamics and cost of all robots. Another typical category is density control [5], which models the swarm as a probability density function (PDF) and formulates the optimal control problem over the space of PDFs to transform the distribution into the target one. These methods achieve improved scalability as the state space is independent of the swarm size. However, these works usually assume obstacle-free environments and oversimplified agent kinematics, and cannot be implemented in real-time due to high computational demands, making them inapplicable to practical situations.

Recently, hierarchical approaches that plan the reference trajectory of the entirety of the swarm at the upper level and generate individual control for each robot to track the reference trajectory at the lower level are receiving increased attention as they integrate the advantages of both macroscopic and microscopic methods. For example, Alonso et al. [10] and Mao et al. [11] proposed to generate local goals or reference trajectories for the swarm at the upper level,

and developed distributed lower-level tracking controllers for each robot to track the reference trajectory. An evident drawback of these approaches lies in the lack of flexibility as the whole swarm is restricted to certain manually designed formations, which prevents the swarm from conducting adaptive behaviours such as splitting and merging when navigating through obstacle-cluttered environments.

To tackle this challenge, adaptive distributed optimal control (ADOC) [4] proposed to model the swarm state as a Gaussian Mixture Model (GMM), and developed a macroscopic planner to navigate the GMM in a cluttered environment. The individual agents are then controlled by an artificial potential field (APF) method to track the GMM at the microscopic level. ADOC is able to navigate hundreds of agents in a cluttered environment. However, ADOC uses myopic planning and lacks a systematic way of enforcing collision avoidance, thus leading to risky robot trajectories that are close to obstacles and cannot be easily adjusted.

A promising remedy is the conditional value-at-risk (CVaR), which has attracted significant interest as a measure of risk. CVaR measures the expected risk that would occur beyond a certain risk level. Compared to the traditional chance constraint risk measure, CVaR takes into account the long-tail distribution of risks and can better discern rare events. Due to these desirable properties, CVaR has been used in motion planning of both single [12] and small-scale [13] robotic systems. However, there is few work adopting CVaR for large-scale robotic systems due to the high computational overhead.

To address the existing research gaps, we propose a Risk-aware mOtion planner for large-scale swarm robotics using conditional ValuE-at-Risk (ROVER). In particular, ROVER adopts a hierarchical motion planning framework, where the macroscopic state of the swarm is represented as a GMM, and the macroscopic planning is formulated as a finite-time model predictive control (FTMPC) problem integrating CVaR as the risk measure. To the best of the authors' knowledge, ROVER is the first online motion planner for large-scale swarms that analytically integrates CVaR of a GMM (GMM-CVaR) as risk constraints. The specific contributions of this work can be summarized as follows:

- We propose to compute the distribution of Signed Distance Function (SDF) between swarm GMM and obstacles to evaluate collision between robots and obstacles. We further prove that the SDF distribution can be approximated as a GMM via Taylor expansion.
- We derive an analytical expression of CVaR for GMM, which is then used as risk constraints to prevent collision between the swarm and obstacles.
- We present an online solution of the proposed FTMPC by Sequential Linear Programming (SLP). Extensive simulations and comparisons with benchmark approaches verify the outstanding performance of ROVER in terms of flexibility, scalability, and risk control ability.

II. BACKGROUND AND PROBLEM FORMULATION

This section provides an overview of the important preliminaries followed by the problem formulation. For simplicity, we define a set notation \underline{N} to represent the set $\{1, 2, \dots, N\}$ for any $N \in \mathbb{N}_+$.

A. Wasserstein Metric

The Wasserstein metric measures the distance between two probability distributions on a given metric space, and is widely used in the optimal transport (OT) problems [14].

We first consider two Gaussian PDFs denoted by $g_1 = \mathcal{N}(\boldsymbol{\mu}_1, \boldsymbol{\Sigma}_1) \in \mathcal{P}(\mathbb{R}^d)$ and $g_2 = \mathcal{N}(\boldsymbol{\mu}_2, \boldsymbol{\Sigma}_2) \in \mathcal{P}(\mathbb{R}^d)$, where $\mathcal{P}(\Omega)$ denotes the set of all probability distributions defined on the sample space Ω . The Wasserstein metric W_2 between g_1 and g_2 can be analytically calculated by [14]:

$$W_2(g_1, g_2) = \left\{ \|\boldsymbol{\mu}_1 - \boldsymbol{\mu}_2\|^2 + \text{tr} \left[\boldsymbol{\Sigma}_1 + \boldsymbol{\Sigma}_2 - 2 \left(\boldsymbol{\Sigma}_1^{1/2} \boldsymbol{\Sigma}_2 \boldsymbol{\Sigma}_1^{1/2} \right)^{1/2} \right] \right\}^{1/2}, \quad (1)$$

where $\text{tr}(\cdot)$ indicates the trace operator, and $\|\cdot\|$ denotes the Euclidean norm. Furthermore, the Wasserstein metric of two GMMs can be approximated utilizing eq. (1) [14]. Specifically, consider two GMMs with N_1 and N_2 Gaussian components (GCs) denoted by $\varphi_1 = \sum_{i=1}^{N_1} \omega_1^i g_1^i, \varphi_2 = \sum_{j=1}^{N_2} \omega_2^j g_2^j$, where g_1^i, g_2^j denote the i th GC and the j th GC of φ_1 and φ_2 , respectively, and ω_1^i, ω_2^j denote their corresponding weights that satisfy $\sum_{i=1}^{N_1} \omega_1^i = \sum_{j=1}^{N_2} \omega_2^j = 1$. The Wasserstein metric $W_2(\varphi_1, \varphi_2)$ can be approximated as:

$$W_2(\varphi_1, \varphi_2) \triangleq \left\{ \min_{\pi \in \Pi(\boldsymbol{\omega}_1, \boldsymbol{\omega}_2)} \sum_{i=1}^{N_1} \sum_{j=1}^{N_2} [W_2(g_1^i, g_2^j)]^2 \pi(i, j) \right\}^{1/2}. \quad (2)$$

where $\boldsymbol{\omega}_1 = [\omega_1^1, \dots, \omega_1^{N_1}]$ and $\boldsymbol{\omega}_2 = [\omega_2^1, \dots, \omega_2^{N_2}]$ denote the weight vector, and $\Pi(\boldsymbol{\omega}_1, \boldsymbol{\omega}_2)$ denotes the space of joint distributions between $\boldsymbol{\omega}_1$ and $\boldsymbol{\omega}_2$. From an intuitive perspective, $\pi(i, j)$ represents the weight transported from the i th GC in φ_1 to the j th GC in φ_2 .

B. Signed Distance Function (SDF)

The SDF measures the orthogonal distance between a point and the boundary of a set in a metric space with the sign indicating whether or not the point is in the interior of the set. In particular, the SDF between a point $\mathbf{p} \in \mathbb{R}^2$ and an obstacle $\mathcal{O}_i \subset \mathbb{R}^2$ can be calculated as follows,

$$sd(\mathbf{p}, \mathcal{O}_i) = \begin{cases} -d(\mathbf{p}, \partial\mathcal{O}_i), & \mathbf{p} \in \mathcal{O}_i \\ d(\mathbf{p}, \partial\mathcal{O}_i), & \mathbf{p} \notin \mathcal{O}_i \end{cases}, \quad (3)$$

where $d(\mathbf{p}, \partial\mathcal{O}_i)$ is the minimum distance from \mathbf{p} to the boundary of \mathcal{O}_i , noted as $\partial\mathcal{O}_i$. The closest point to \mathbf{p} on $\partial\mathcal{O}_i$ is denoted as \mathbf{o}_i^* , and the normal vector along the SDF direction can be calculated as follows:

$$\mathbf{n} = \text{sgn}(sd(\mathbf{p}, \mathcal{O}_i)) \cdot (\mathbf{p} - \mathbf{o}_i^*) / \|\mathbf{p} - \mathbf{o}_i^*\|, \quad (4)$$

where $\text{sgn}(\cdot)$ is the sign function. The SDF and its corresponding normal vector \mathbf{n} can be efficiently calculated

using Gilbert–Johnson–Keerthi (GJK) algorithm [15] and Expanding Polytope Algorithm (EPA) [16].

C. Value-at-Risk and Conditional Value-at-Risk

The value-at-risk (VaR) represents the minimum possible value of risk that can be achieved given a risk tolerance level $\alpha \in (0, 1]$. Specifically, the VaR of a random variable ζ under α is defined as

$$VaR_\alpha(\zeta) = \min \{z | Pr(\zeta \leq z) \geq 1 - \alpha\}, \quad (5)$$

where $Pr(\cdot)$ denotes the probability. The CVaR of a random variable ζ with a continuous distribution is defined as

$$CVaR_\alpha(\zeta) = \min_{z \in \mathbb{R}} \mathbb{E}[z + (\zeta - z)^+ / \alpha] = \mathbb{E}[\zeta | \zeta \geq VaR_\alpha(\zeta)], \quad (6)$$

where $(\cdot)^+ = \max(\cdot, 0)$, and \mathbb{E} is the expectation operator.

Let $\phi(\cdot)$ and $\Phi(\cdot)$ denote the PDF and the cumulative distribution function (CDF) of a standard normal distribution, respectively. The CVaR of a Gaussian random variable $\zeta \sim \mathcal{N}(\mu, \sigma^2)$ has the following closed-form expression [17]

$$CVaR_\alpha(\zeta) = \mu + \sigma \phi(\Phi^{-1}(1 - \alpha)) / \alpha. \quad (7)$$

D. Problem Formulation

Consider a workspace $\mathcal{W} \subset \mathbb{R}^2$ that contains a robot swarm consisting of N_r robots and N_o obstacles $\mathcal{O}_i \subset \mathcal{W}, \forall i \in \underline{N}_o$. Let $\mathcal{O} = \bigcup_{i=1}^{N_o} \mathcal{O}_i$ represents the union of all obstacles. Obstacles are static and known a priori. For simplicity, \mathcal{O}_i is assumed to be convex. For non-convex obstacles, the convex hull or convex decomposition can be applied to obtain convex obstacles.

The proposed swarm motion planning problem aims to safely transport robots from an initial area to a target region in obstacle-cluttered environments (Fig. 1). We propose to address this task through a hierarchical strategy involving a macroscopic planning stage and a microscopic control stage.

At the macroscopic level, the swarm state can be represented by a time-varying random variable $\mathbf{X}(k)$ with k denoting the time step. The PDF of $\mathbf{X}(k)$ indicates the density distribution of robots across \mathcal{W} .

Considering the universal approximation property of Gaussian Mixture Models (GMMs), we choose a GMM $\varphi_k = \sum_{j=1}^{N_k} \omega_k^j g_k^j$ with $\sum_{j=1}^{N_k} \omega_k^j = 1$ from the GMM distribution space $\mathcal{G}(\mathcal{W})$ to represent the macroscopic state of the robot swarm, i.e., $\mathbf{X}(k) \sim \varphi_k$. Here N_k denotes the number of GCs in the GMM and g_k^j is the j th Gaussian distribution with mean μ_k^j , covariance matrix Σ_k^j and weight $\omega_k^j \geq 0$.

The macroscopic planning stage is formulated as a swarm PDF OT problem (P0) from an initial distribution $\varphi_0 \in \mathcal{G}(\mathcal{W})$ to a target distribution $\varphi_{target} \in \mathcal{G}(\mathcal{W})$ while minimizing the total transport cost and satisfying a set of collision avoidance constraints, as follows:

$$\min_{\varphi_1, \varphi_2, \dots, \varphi_{T_f}} J \triangleq \sum_{k=0}^{T_f-1} \lambda_k W_2(\varphi_k, \varphi_{k+1}) \quad (8a)$$

$$s.t. \quad CVaR_\alpha(-sd(\mathbf{X}(p), \mathcal{O}_i)) < \epsilon, \quad (8b)$$

$$f_i(\varphi_p) < \epsilon, \quad (8c)$$

$$\forall p \in \underline{T}_f, \forall i \in \underline{N}_o, \quad (8d)$$

where the optimization objective J is the sum of the Wasserstein distance between consecutive PDFs from the

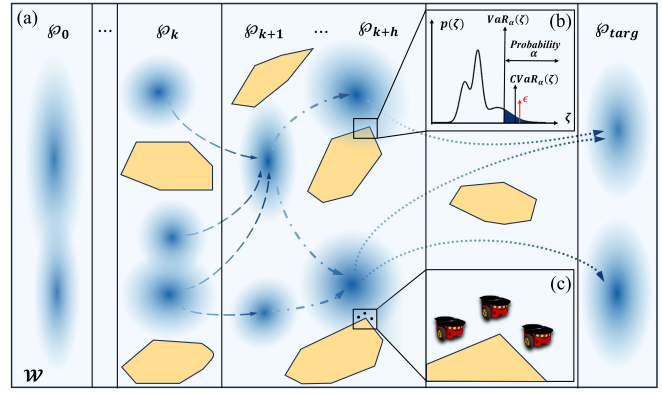


Fig. 2. Illustration of ROVER transporting the robot swarm from an initial area in the left column to the target area in the right column in a cluttered workspace $\mathcal{W} \subset \mathbb{R}^2$. Static obstacles are denoted by pale yellow polygons. Columns separated by verticle black lines correspond to different time steps. (a) **Macroscopic planning in ROVER.** Blue clouds represent the GMM distributions φ of the swarm state. Blue dashed lines illustrate the GMM transformation from φ_k to φ_{k+1} , while the dash-dotted lines represent $(h-1)$ subsequent transformations from φ_{k+1} to φ_{k+h} . The dotted lines represent the transformation from φ_{k+h} to φ_{target} . (b) **Collision avoidance between GMM and obstacles using CVaR.** The random variable ζ represents the distribution of distance between a GMM and obstacles. The shaded area denotes the $\alpha\%$ of the area under $p(\zeta)$. The collision avoidance under a risk acceptance level $\alpha \in (0, 1]$ is to constrain the $CVaR_\alpha(\zeta)$, the expected value of ζ under the shaded area, below a user-defined threshold ϵ . (c) **Microscopic control in ROVER.** Robots track the GMM trajectories given from the macroscopic level while avoiding collision.

initial time step to the terminal time step T_f when $\varphi_{T_f} = \varphi_{target}$, with corresponding weight coefficients $\lambda_k > 0$. The constraint eq. (8b) limits the expectation of SDF between the swarm PDF and obstacles. Taking the opposite sign of SDF is performed to accommodate the definition of CVaR. Details regarding $sd(\cdot, \cdot)$ will be provided in Sec. III-B. The constraint eq. (8c) with a linear function $f_i(\cdot)$ and a probability bound $\epsilon \in (0, 1)$ ensures that the robots do not come too close to each other. A more detailed elaboration of (P0) will be presented in Sec. III-A.

From the microscopic aspect, the position of every single robot $\mathbf{x}_i(k)$ can be treated as a sample from the swarm GMM, with a motion model formulated as:

$$\begin{aligned} \mathbf{x}_i(k+1) &= f(\mathbf{x}_i(k), \mathbf{u}_i(k)), \\ k &= 0, 1, \dots, T_f - 1, \forall i \in \underline{N}_r, \end{aligned} \quad (9)$$

where $f(\cdot, \cdot)$ denotes the robot motion function, $\mathbf{x}_i(k)$ and $\mathbf{u}_i(k)$ denote the state and control input of every single robot at time step k . At the microscopic control stage, each robot computes its own control inputs by the APF method to keep track of the swarm PDF obtained from the macroscopic planning level and avoid collisions with static obstacles and other robots. The microscopic control falls beyond the scope of this work, and interested readers are encouraged to refer to [18] for details.

III. METHODOLOGY

This section provides a comprehensive outline of the methodology, including the detailed formulation of macroscopic planning (Sec. III-A), the SDF distribution between




Fig. 3. An example of determining \mathcal{GC}_k and N_k , $\forall k \in T_f$, based on assumption 1 and assumption 2. Figures (a), (b), and (c) represent GMMs at time steps k , $k+1$, and $k+2$, respectively. Light blue circles depict GCs in \mathcal{C} and pentagrams with numbers represent the GCs at specific time step and their corresponding indices. The dotted lines give an example of weight transported between consecutive time steps within a transport range denoted by light orange circles.

swarm GMM and obstacles (Sec. III-B), the formulation of GMM-CVaR collision avoidance constraints (Sec. III-C), followed by an online implementation with SLP framework (Sec. III-D). A schematic overview is provided in Fig. 2.

A. Macroscopic Planning Formulation

The proposed (P0) is an infinite-dimensional optimization problem, which is computationally prohibitive. Thus, we proposed two approximations to simplify the problem.

In the first approximation, we uniformly discretize the planning space and establish a set \mathcal{C} of GCs whose mean positions are constrained on the discrete points and the covariance matrices are uniform. We then propose the following assumptions:

Assumption 1. All GCs of the GMM in the planning stage can only be chosen from \mathcal{C} , i.e., $g_k^j \in \mathcal{C}$, $\forall j \in \underline{N_k}, \forall k \in T_f$.

Assumption 2. Between consecutive planning steps, GCs can be only transported to the GCs within a predefined transport radius measured by the Wasserstein Metric.

Leveraging assumption 1 and assumption 2, a set of all GCs at time step k denoted as \mathcal{GC}_k and $N_k = |\mathcal{GC}_k|$ can be predetermined, $\forall k \in T_f$. Only weights of the GMMs remain to be optimized. A schematic explanation of these assumptions can be found in Fig. 3.

In the second approximation, we propose to reduce the dimension of (P0) utilizing an FTMPC framework. The transport cost is composed of stage cost and terminal cost. The stage cost represents the transport cost within the FTMPC planning horizon, and terminal cost is defined as the cost from the last step in the planning horizon to T_f . To calculate the terminal cost, we make the following assumption:

Assumption 3. The GCs of the GMM from the last step in the planning horizon to $T_f - 1$ do not split or merge. Transformation into the target distribution is conducted in the terminal time step T_f .

With the help of these two approximations, (P0) is reformulated into an FTMPC problem (P) as follows:

$$\min_{\pi} J \triangleq \sum_{p=1}^h \sum_{i=1}^{N_{k+p-1}} \sum_{j=1}^{N_{k+p}} \lambda_{k+p-1} W_2^2(g_{k+p-1}^i, g_{k+p}^j) \pi_{k+p-1}(i, j) + \sum_{m=1}^{N_{k+h}} \sum_{n=1}^{N_{target}} \lambda_{k+h} \mathcal{Q}_{k+h}(m, n) \pi_{k+h}(m, n) \quad (10a)$$

$$s.t. \quad \omega_{k+p-1}^i = \sum_{j=1}^{N_{k+p}} \pi_{k+p-1}(i, j), \quad \forall i \in \underline{N_{k+p-1}} \quad (10b)$$

$$\omega_{k+p}^j = \sum_{i=1}^{N_{k+p-1}} \pi_{k+p-1}(i, j), \quad \forall j \in \underline{N_{k+p}} \quad (10c)$$

$$\omega_{k+h}^m = \sum_{n=1}^{N_{target}} \pi_{k+h}(m, n), \quad \forall m \in \underline{N_{k+h}} \quad (10d)$$

$$\omega_{target}^n = \sum_{m=1}^{N_{k+h}} \pi_{k+h}(m, n), \quad \forall n \in \underline{N_{target}} \quad (10e)$$

$$\sum_{m=1}^{N_{k+p}} \omega_{k+p}^m = \sum_{n=1}^{N_{target}} \omega_{target}^n = 1 \quad (10f)$$

$$\omega_{k+p} \geq \mathbf{0}, \omega_{target} \geq \mathbf{0} \quad (10g)$$

$$\sum_{i=1}^{N_{k+p-1}} \sum_{j=1}^{N_{k+p}} f_i(g_{k+p}^j) \pi_{k+p-1}(i, j) < \epsilon \quad (10h)$$

$$CVaR_{\alpha}(-sd(\mathbf{X}(k+p), \mathcal{O}_i)) < \epsilon, \quad (10i)$$

$$\forall p \in \underline{h}, \forall i \in \underline{N_o} \quad (10j)$$

where h represents the planning horizon. The transport policy $\pi = [\pi_k(\cdot, \cdot), \pi_{k+1}(\cdot, \cdot), \dots, \pi_{k+h}(\cdot, \cdot)]$ is a stack of joint Probability Mass Functions (PMFs) whose marginal PMFs are denoted by eqs. (10b) to (10e). As an example, $\pi_k(i, j)$ represents the weight transported from g_k^i to g_{k+1}^j . The optimal swarm PDF at the next time step is obtained by

$$\varphi_{k+1}^* = \sum_{j=1}^{N_{k+1}} \sum_{i=1}^{N_k} \pi_k^*(i, j) g_{k+1}^j. \quad (11)$$

The former half of eq. (10a) represents the stage cost with weights $\lambda_{k:k+h-1}$, while the latter half represents the terminal cost with a weight λ_{k+h} . Under assumption 1 and assumption 2, the squared terms of $W_2(\cdot, \cdot)$ in the stage cost can be calculated offline. For the terminal cost, we adopt the transport policy as in assumption 3 and establish a directed graph \mathbf{G} whose vertices are GCs from \mathcal{C} and weights are determined by the transport distance and the collision risk between two vertices. The cost term $\mathcal{Q}_{k+h}(m, n)$ in the terminal cost can be computed by applying a shortest-path-planning algorithm on \mathbf{G} in advance. Therefore, the cost function eq. (10a) is linear w.r.t. π .

Similarly, the constraint eq. (10h) is linear w.r.t. π . But the formulation of constraint eq. (10i) remains unclear, which will be elaborated in Sec. III-B and Sec. III-C.

B. SDF under GMM Representation

The collision avoidance constraint eq. (10i) involves SDF between the swarm GMM and obstacles. We consider a GMM random variable $\mathbf{X} \sim \sum_{j=1}^N \omega^j \mathcal{N}(\boldsymbol{\mu}^j, \boldsymbol{\Sigma}^j)$ and derive the SDF distribution between \mathbf{X} and an obstacle noted as $sd(\mathbf{X}, \mathcal{O}_i)$.

For simplicity, we first consider the j th GC of the GMM denoted as $\mathbf{X}^j \sim \mathcal{N}(\boldsymbol{\mu}^j, \boldsymbol{\Sigma}^j)$ and its corresponding $sd(\mathbf{X}^j, \mathcal{O}_i)$. Utilizing the EPA or GJK algorithm and eq. (4), the deterministic SDF between the mean position of \mathbf{X}^j and an obstacle \mathcal{O}_i , denoted by $sd(\boldsymbol{\mu}^j, \mathcal{O}_i)$, and the corresponding normal vector \mathbf{n}^j , can be calculated.

The $sd(\mathbf{X}^j, \mathcal{O}_i)$ can be subsequently approximated by the first-order Taylor expansion as follows,

$$sd(\mathbf{X}^j, \mathcal{O}_i) \approx sd(\boldsymbol{\mu}^j, \mathcal{O}_i) + \nabla sd(\mathbf{X}^j, \mathcal{O}_i)|_{\mathbf{X}^j=\boldsymbol{\mu}^j}(\mathbf{X}^j - \boldsymbol{\mu}^j), \quad (12)$$

where $\nabla sd(\mathbf{X}^j, \mathcal{O}_i)|_{\mathbf{X}^j=\boldsymbol{\mu}^j} = (\mathbf{n}^j)^T$. The reasonableness of this approximation stems from the assumption that when calculating $sd(\mathbf{X}^j, \mathcal{O}_i)$, the closest point \mathbf{o}_i^* and the normal vector \mathbf{n}^j remain unchanged for any sample of $\mathbf{X}^j \sim \mathcal{N}(\boldsymbol{\mu}^j, \boldsymbol{\Sigma}^j)$ [19]. The approximation error is acceptable when samples of \mathbf{X}^j lie in the neighbourhood of $\boldsymbol{\mu}^j$. Notice that $sd(\mathbf{X}^j, \mathcal{O}_i)$ is a Gaussian random variable as a result of the linear transformation of \mathbf{X}^j , i.e.,

$$sd(\mathbf{X}^j, \mathcal{O}_i) \sim \mathcal{N}(sd(\boldsymbol{\mu}^j, \mathcal{O}_i), (\mathbf{n}^j)^T \boldsymbol{\Sigma}^j \mathbf{n}^j). \quad (13)$$

We further derive the distribution of $sd(\mathbf{X}, \mathcal{O}_i)$ based on eq. (13) and the following proposition:

Proposition 1 (SDF of GMM¹). *The distribution $\sum_{j=1}^N \omega^j Pr(sd(\mathbf{X}^j, \mathcal{O}_i))$ equals distribution $Pr(sd(\mathbf{X}, \mathcal{O}_i))$.*

Proposition 1 yields that $sd(\mathbf{X}, \mathcal{O}_i)$ follows a GMM distribution as follows:

$$sd(\mathbf{X}, \mathcal{O}_i) \sim \sum_{j=1}^N \omega^j \mathcal{N}(sd(\boldsymbol{\mu}^j, \mathcal{O}_i), (\mathbf{n}^j)^T \boldsymbol{\Sigma}^j \mathbf{n}^j). \quad (14)$$

Therefore, the $sd(\cdot, \cdot)$ in eq. (10i) can be calculated by $sd(\mathbf{X}, \mathcal{O}_i) = \sum_{j=1}^N \omega^j sd(\mathbf{X}^j, \mathcal{O}_i)$ utilizing eq. (12).

C. Analytical Expression of GMM-CVaR

The collision avoidance constraints eq. (10i) involve GMM-CVaR, for which there is currently no analytical expression available. Therefore, we first propose a proposition concerning GMM-CVaR.

Proposition 2 (CVaR of GMM²). *The CVaR of a GMM random variable Y at risk acceptance level α can be represented by the summation of CVaRs of the corresponding Gaussian component Y^j of Y at risk acceptance level α_j :*

$$CVaR_\alpha(Y) = \frac{1}{\alpha} \sum_{j=1}^N \omega^j \alpha_j CVaR_{\alpha_j}(Y^j), \quad (15)$$

where ω^j denotes the weight of Y^j , and α_j is the tail probability of the Y^j distribution at the $VaR_\alpha(Y)$ quantile, i.e., $\alpha_j = \int_{VaR_\alpha(Y)}^{+\infty} p(Y^j) dy$. Besides, there exists relationship between α and α_j denoted as $\alpha = \sum_{j=1}^N \omega_j \alpha_j$.

Thus, the constraint eq. (10i) can be explicitly expressed by substituting Y in proposition 2 with $Y(k+p) = -sd(\mathbf{X}(k+p), \mathcal{O}_i)$ and Y^j in proposition 2 with $Y^j(k+p) = -sd(\mathbf{X}^j(k+p), \mathcal{O}_i)$, $\forall i \in \underline{N}_o$:

¹Due to space constraint, the proof of Proposition 1 is provided in [20].

²Due to space constraint, the proof of Proposition 2 is provided in [20].

$$CVaR_\alpha(Y(k+p)) = \frac{1}{\alpha} \sum_{j=1}^{N_{k+p}} \omega_{k+p}^j \alpha_j CVaR_{\alpha_j}(Y^j(k+p)) < \epsilon, \quad (16a)$$

$$\forall p \in \underline{h}. \quad (16a)$$

Although $CVaR_{\alpha_j}(Y^j(k+p))$ can be explicitly expressed by eq. (7), α_j relating with $VaR_\alpha(Y(k+p))$ is an implicit function of the weights of GCs, which renders eq. (10i) a nonlinear constraint. In the next section, we propose a computationally efficient algorithm customized to (P) by leveraging the structure of eq. (16a).

D. Online implementation via Sequential Linear Programming (SLP)

Since the constraint eq. (10i) turns out to be the only nonlinear term in (P), we propose to linearize the constraint eq. (10i) to achieve online implementation of ROVER. Specifically, we adopt an SLP framework [21] and seek to solve the optimization problem (P) by solving a series of linear programming (LP) subproblems (P^v). For each subproblem (P^v), we linearize the $CVaR_\alpha(Y(k+p))$ constraint using the first-order Taylor expansion at a feasible solution of (P^v) $\boldsymbol{\pi}^v$ as

$$CVaR_\alpha(Y(k+p)) = CVaR_\alpha(Y_v(k+p)) + \nabla_{\boldsymbol{\pi}}^T CVaR_\alpha(Y(k+p))|_{\boldsymbol{\pi}=\boldsymbol{\pi}^v}(\boldsymbol{\pi} - \boldsymbol{\pi}^v), \quad (17a)$$

where $Y_v(k+p)$ denotes the SDF random variable $Y(k+p)$ under a given $\boldsymbol{\pi}^v$, and its corresponding $CVaR_\alpha(Y_v(k+p))$ can be calculated with eq. (16a). It should be noted that the $\boldsymbol{\pi}$ used in this section represents the flattened vector form of the matrix form $\boldsymbol{\pi}$ in Sec. III-A.

For each iteration, the subproblem (P^v) is as follows:

$$\min_{\boldsymbol{\pi}} J \quad (18a)$$

$$s.t. \quad (10b) - (10h), (10j), \quad (18b)$$

$$(17a) < \epsilon, \quad (18c)$$

$$-s^v < \boldsymbol{\pi} - \boldsymbol{\pi}^v < s^v, \quad (18d)$$

where eq. (18d) is an element-wise inequality constraint and each dimension of s^v is a step bound parameter $s^v > 0$. This constraint ensures that the feasible domain of (P^v) remains in the neighbourhood of $\boldsymbol{\pi}^v$, where linearization occurs. Every iteration involves first the linearization of $CVaR_\alpha(Y(k+p))$ and solving (P^v) to find the optimal iterative solution $\boldsymbol{\pi}^{v*}$. Then the $\boldsymbol{\pi}^v$ and s^v are updated for the next iteration. This algorithm is sequentially executed until $\boldsymbol{\pi}^v$ meets the convergence condition.

IV. SIMULATION AND RESULTS

In this section, we evaluate the performance of ROVER via simulations. The workspace \mathcal{W} is a $[0, 200] \times [0, 160]m^2$ area with static obstacles. The initial distribution of the swarm consists of four GCs with a shared covariance matrix $\boldsymbol{\Sigma} = 100\mathbf{I}_2$, where \mathbf{I}_2 denotes a 2×2 identity matrix. Other parameters are $\boldsymbol{\mu}_1 = [25, 35]$, $\omega_1 = 0.25$, $\boldsymbol{\mu}_2 = [25, 55]$, $\omega_2 = 0.375$, $\boldsymbol{\mu}_3 = [25, 115]$, $\omega_3 = 0.1875$, $\boldsymbol{\mu}_4 = [25, 135]$, $\omega_4 = 0.1875$. The target distribution is composed of three GCs with parameters $\boldsymbol{\mu}_1 = [175, 120]$, $\omega_1 = 0.25$, $\boldsymbol{\mu}_2 =$

[175, 60], $\omega_2 = 0.375$, $\mu_3 = [175, 40]$, $\omega_3 = 0.375$ and an identical covariance matrix $\Sigma = 100I_2$. The GC set \mathcal{C} is predefined by taking the mean of each GC on fixed grids, where the X coordinates range from 5 to 195 and the Y coordinates range from 5 to 155, with a discretization interval of 10. Consequently, the set \mathcal{C} comprises a total of $20 \times 16 = 320$ GCs, each with a same covariance matrix $\Sigma_c = 50I_2$. We set the discretization time interval Δt to 0.1s. Robots are assumed to adopt omnidirectional models with a maximal speed of 1.5m/s and the radius of each robot is set to 0.12m. To specify the optimization problem, α and ϵ are set to 0.05 and -0.2 respectively. In the objective function, we set h to 2, λ_{k+h} to 3, and λ_{k+p-1} to 1, $\forall p \in \underline{h}$. In the subproblem (P^v), we adopt a fixed step bound as $s^v = 0.1, \forall v \in \mathbb{N}_+$. The η in the convergence condition is 10^{-5} . All simulations are run on a desktop (13th Intel(R) i7 CPU@2.10GHz) and LP is solved using the “mosek” solver with interior-point algorithm in MATLAB.

We propose five metrics to quantify the performance of a motion planner:

- Total runtime t excluding the offline data preparation phase. t is set to be infinite when the swarm transport task fails to be completed within $T_f^{max} = 3000$.
- Average runtime per time step $t_f = t/T_f$.
- Average trajectory length \bar{D} until reaching the goal.
- Minimum inter-robot distance $\min(d_{ij})$ and minimum robot-obstacle distance $\min(d_{io})$ maintained by all robots over the whole trajectory.

A. Comparison with Benchmark Approaches

We compared the performance of our method ROVER with several state-of-art approaches in a more complex scenario characterized by a host of fragmented polygonal obstacles (Fig. 4). Approaches highly pertinent to our research include multi-robot Formation Control (FC) [10], Predictive Control (PC) [6, 7], and Optimal Tube swarm planning and control (OT) [11]. However, OT involving a single virtual tube planning cannot accomplish our swarm transport task, where our swarm is partitioned into distinct subgroups for various goal areas concerning variation in swarm density. Therefore, the following section presents a comparative analysis of ROVER, FC and PC. Notice that FC is a sampling-based method, the numerical results of FC are mean values obtained from five random simulations. Furthermore, PC utilizes soft constraints for collision avoidance, allowing robots to continue moving towards target positions even after collisions occur. We approximate polygonal obstacles using circular obstacles in the simulation of PC to accommodate the collision avoidance mechanism of PC.

1) *Flexibility Analysis of Different Approaches:* A comparison with benchmark approaches with a fixed swarm size $N_r = 500$ is presented in this subsection. The trajectories of robots are illustrated in Fig. 4.

Fig. 4 showcases the flexibility of different methods. Through planning and control at the individual level, the microscopic approach PC exhibits the utmost flexibility. As hierarchical approaches, ROVER demonstrates greater

flexibility compared to FC. ROVER enables the swarm to split and merge when encountering obstacles, while FC constrains the swarm to travel as a group, compromising its flexibility. The flexibility of ROVER benefits from the GMM representation of a swarm at the macroscopic level. The adaptive change of GCs and their corresponding weights in the GMMs facilitates the swarm’s splitting and merging behaviour, which results in likewise flexibility as provided by the microscopic method PC.

2) *Scalability in Various Swarm Size:* We conduct a comparison with varying swarm sizes and the results in Tab. I are presented in the manner of ROVER / FC / PC, respectively. The time results in the parentheses denote the aggregated macroscopic planning duration of ROVER. FC fails to satisfy the inter-robot distance constraints when $N_r = 1000$, leading to the NaN results.

In terms of computational efficiency (t and t_f), PC outperforms the other two approaches at smaller swarm scales ($N_r = 25, 50$). However, as the swarm size increases, the advantages of hierarchical methods (ROVER and FC) become increasingly significant, while the computational cost of PC rises sharply. Additionally, ROVER exhibits superior performance than FC across various swarm sizes. It is worth noticing that ROVER only requires a runtime of around 10min when $N_r = 1000$ and can reach the 1Hz control frequency, enabling online implementations. The remarkable scalability of ROVER is attributed to the complete decoupling of the macroscopic planning level from the microscopic control phase. Robots in our work can be viewed as samples from swarm PDFs. Thus, increasing the swarm size does not affect macroscopic PDF planning, which corresponds to the stable planning duration within the parentheses.

Regarding the average trajectory length (\bar{D}), ROVER exhibits a decreasing trend as the swarm scale increases. This can be explained by the fact that the PDF paths provided by ROVER’s macroscopic planning are similar across various swarm sizes, and the growth in robot density forces peripheral robots in the swarm to choose trajectories closer to obstacles but with shorter distances. Conversely, the trajectory length of FC is the utmost and does not follow a discernible pattern since FC only samples a feasible formation sequence without guaranteeing optimality in travel distance. PC maintains a relatively stable and shortest path across all swarm sizes because PC optimizes every robot’s path individually with fixed start and goal positions.

As for collision avoidance, ROVER exhibits a decrease in the distance to obstacles and among robots as the robot density surges, while still guaranteeing collision-free trajectories. FC demonstrates the most conservative obstacle avoidance performance but reduces the inter-robot distance ($\min(d_{ij} = 0)$ in $N_r = 250, 500$). PC encounters collisions (negative $\min(d_{ij})$ and $\min(d_{io})$) except for small swarm scales ($N_r = 25, 50$) due to soft collision avoidance constraints.

V. CONCLUSION

In this work, we propose ROVER for risk-aware motion planning of large-scale robotic swarms. The GMM represen-




Fig. 4. Figure (a)-(c) shows the trajectories of 500 robots generated by ROVER, FC, and PC, respectively. The initial positions of robots are demonstrated by circles, while their corresponding final positions are highlighted by diamonds. The grey polygons denote static obstacles.

TABLE I

N_r	$t(\text{min})$	$t_f(\text{s})$	$\bar{D}(\text{m})$	$\min(d_{ij})(\text{m})$	$\min(d_{io})(\text{m})$
25	1.85 (0.66) / 2.35 / 1.01	0.21 / 0.05 / 0.33	223.00 / 345.40 / 165.83	1.03 / 0 / 0.29	0.38 / 16.63 / 0.09
50	2.10 (0.67) / 3.55 / 1.71	0.22 / 0.09 / 0.55	210.87 / 305.19 / 158.35	0.84 / 0.44 / 0.06	0.37 / 11.35 / 0.05
250	3.46 (0.67) / 24.12 / 22.10	0.36 / 0.58 / 6.20	202.81 / 323.31 / 160.60	0.38 / 0 / -0.2	0.23 / 8.95 / -2.71
500	6.05 (0.67) / 58.58 / 120.78	0.64 / 1.28 / 31.51	202.26 / 364.16 / 159.036	0.34 / 0 / -0.21	0.22 / 6.68 / -3.08
1000	10.89 (0.66) / NaN / 564.05	1.09 / NaN / 120.01	204.37 / NaN / 159.42	0.32 / NaN / -0.22	0.17 / NaN / -4.10

tation of the swarm brings great flexibility to ROVER, allowing the robot swarm to split and merge to travel around obstacles. The SLP framework allows ROVER to conduct online motion planning for large-scale swarm robotic systems. The incorporation of CVaR enables effective risk management between the swarm and obstacles, ensuring safe navigation through cluttered environments. Extensive comparisons with benchmarks illustrate the effectiveness of ROVER in terms of flexibility, scalability, and safety guarantee. Future work will extend ROVER to unknown environments and evaluate the approach using real-world experiments.

REFERENCES

- [1] M. Brambilla, E. Ferrante, M. Birattari, and M. Dorigo, "Swarm robotics: a review from the swarm engineering perspective," *Swarm Intelligence*, vol. 7, pp. 1–41, 2013.
- [2] E. Tuci, M. H. Alkilabi, and O. Akanyeti, "Cooperative object transport in multi-robot systems: A review of the state-of-the-art," *Frontiers in Robotics and AI*, vol. 5, p. 59, 2018.
- [3] G. Bevacqua, J. Cacace, A. Finzi, and V. Lippiello, "Mixed-initiative planning and execution for multiple drones in search and rescue missions," in *Proceedings of the International Conference on Automated Planning and Scheduling*, vol. 25, pp. 315–323, 2015.
- [4] P. Zhu, C. Liu, and S. Ferrari, "Adaptive online distributed optimal control of very-large-scale robotic systems," *IEEE Transactions on Control of Network Systems*, vol. 8, no. 2, pp. 678–689, 2021.
- [5] Y. Chen, "Density control of interacting agent systems," *IEEE Transactions on Automatic Control*, 2023.
- [6] E. Soria, F. Schiano, and D. Floreano, "Predictive control of aerial swarms in cluttered environments," *Nature Machine Intelligence*, vol. 3, no. 6, pp. 545–554, 2021.
- [7] E. Soria, F. Schiano, and D. Floreano, "Distributed predictive drone swarms in cluttered environments," *IEEE Robotics and Automation Letters*, vol. 7, no. 1, pp. 73–80, 2021.
- [8] K. Elamvazhuthi and S. Berman, "Mean-field models in swarm robotics: A survey," *Bioinspiration & Biomimetics*, vol. 15, no. 1, p. 015001, 2019.
- [9] T. Zheng, Q. Han, and H. Lin, "Transporting robotic swarms via mean-field feedback control," *IEEE Transactions on Automatic Control*, vol. 67, no. 8, pp. 4170–4177, 2021.
- [10] J. Alonso-Mora, S. Baker, and D. Rus, "Multi-robot formation control and object transport in dynamic environments via constrained optimization," *The International Journal of Robotics Research*, vol. 36, no. 9, pp. 1000–1021, 2017.
- [11] P. Mao, R. Fu, and Q. Quan, "Optimal virtual tube planning and control for swarm robotics," *The International Journal of Robotics Research*, p. 02783649231210012, 2022.
- [12] A. Hakobyan, G. C. Kim, and I. Yang, "Risk-aware motion planning and control using cvar-constrained optimization," *IEEE Robotics and Automation Letters*, vol. 4, no. 4, pp. 3924–3931, 2019.
- [13] A. Navsalkar and A. R. Hota, "Data-driven risk-sensitive model predictive control for safe navigation in multi-robot systems," in *2023 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 1442–1448, IEEE, 2023.
- [14] Y. Chen, T. T. Georgiou, and A. Tannenbaum, "Optimal transport for gaussian mixture models," *IEEE Access*, vol. 7, pp. 6269–6278, 2018.
- [15] E. G. Gilbert, D. W. Johnson, and S. S. Keerthi, "A fast procedure for computing the distance between complex objects in three-dimensional space," *IEEE Journal on Robotics and Automation*, vol. 4, no. 2, pp. 193–203, 1988.
- [16] G. Van Den Bergen, "Proximity queries and penetration depth computation on 3d game objects," in *Game Developers Conference*, vol. 170, 2001.
- [17] M. Norton, V. Khokhlov, and S. Uryasev, "Calculating cvar and bpoe for common probability distributions with application to portfolio optimization and density estimation," *Annals of Operations Research*, vol. 299, pp. 1281–1315, 2021.
- [18] P. Zhu, S. Ferrari, J. Morelli, R. Linares, and B. Doerr, "Scalable gas sensing, mapping, and path planning via decentralized hilbert maps," *Sensors*, vol. 19, no. 7, p. 1524, 2019.
- [19] X. Yang, H. Gao, P. Zhu, and C. Liu, "Risk-aware motion planning for very-large-scale robotics systems using conditional value-at-risk," in *International Conference on Intelligent Robotics and Applications*, pp. 513–525, Springer, 2023.
- [20] X. Yang, Y. Hu, H. Gao, K. Ding, P. Zhu, Y. Sun, and C. Liu, "Rover: Risk-aware swarm robotics motion planner using conditional value at risk," *arXiv preprint arXiv:2402.16690*, 2024.
- [21] R. E. Griffith and R. Stewart, "A nonlinear programming technique for the optimization of continuous processing systems," *Management science*, vol. 7, no. 4, pp. 379–392, 1961.