Continuously Optimizing Radar Placement with Model Predictive Path Integrals

Michael Potter, Student Member, IEEE
Northeastern University, Boston, MA 02115, USA

Shuo Tang, Student Member, IEEE
Northeastern University, Boston, MA 02115, USA

Paul Ghanem, Student Member, IEEE
Northeastern University, Boston, MA 02115, USA

Milica Stojanovic, Fellow, IEEE
Northeastern University, Boston, MA 02115, USA

Pau Closas, Senior Member, IEEE
Northeastern University, Boston, MA 02115, USA

Murat Akcakaya, Senior Member, IEEE
University of Pittsburgh, Pittsburgh, PA 15260, USA

Ben Wright
Kostas Research Institute, Burlington, MA, USA

Marius Necsoiu, Member, IEEE
DEVCOM ARL, San Antonio, TX 78204, USA

Deniz Erdoğmuş, Senior Member, IEEE
Northeastern University, Boston, MA 02115, USA

Michael Everett*, Member, IEEE
Northeastern University, Boston, MA 02115, USA

Tales Imbiriba*, Member, IEEE
University of Massachusetts Boston, MA 02125, USA

***Abstract*— Continuously optimizing sensor placement is essential for precise target localization in various military and civilian applications. While information theory has shown promise in optimizing sensor placement, many studies oversimplify sensor measurement models or neglect dynamic constraints of mobile sensors. To address these challenges, we employ a range measurement model that incorporates radar parameters and radar-target distance, coupled with Model Predictive Path Integral (MPPI) control to manage complex environmental obstacles and dynamic constraints. We compare the proposed approach against stationary radars or simplified range measurement models based on the Root Mean Squared Error (RMSE) of the Cubature Kalman Filter (CKF) estimator for the targets' state. Additionally, we visualize the evolving geometry of radars and targets over time, highlighting areas of highest measurement information gain, demonstrating the strengths of the approach. The proposed strategy outperforms stationary radars and simplified range measurement models in target localization, achieving a 38-74% reduction in mean RMSE and a 33-79% reduction in the upper tail of the 90% Highest Density Interval (HDI) over 500 Monte Carlo (MC) trials across all time steps. Code available at https://github.com/mlpotter/OptimalTrackingAndControl.**

***Index Terms*— Model Predictive Path Integral, Model Predictive Control, Radar, Cubature Kalman Filter**

I. INTRODUCTION

Deploying a fleet of autonomous mobile sensors to estimate the positions and velocities of targets over time is an important research problem, with applications in transportation, navigation, defense, surveillance, and emergency rescue [1]–[4], exemplified in Figure 1. This problem involves multiple challenges, including predicting expected sensor information gain without exact target or measurement data, adapting mobile sensors planned trajectories to dynamic environmental and target changes, and selecting suitable sensor models and modalities. We will discuss these challenges in the following paragraphs.

The multi-target tracking and sensor placement literature has explored various sensing modalities, including cameras [5], ultrasonic sensors [6], [7], Light Detection and Ranging (LIDAR) [8], acoustic range sensors [9], and radar-based techniques dating back to World War II. This paper focuses on tracking with radar in the Ultra High Frequency (UHF) (particularly "cognitive radars" [10]), due to the long-range detection capability, ability to penetrate atmospheric obstacles, and versatility in capturing detailed target features [11].

A key challenge in radar-based target tracking is in incorporating realistic sensor models into the information metrics used for controlling the sensors. In particular, most existing control strategies for this problem use the Posterior Fisher Information Matrix (PFIM) [12] (inverse of the Posterior Cramer Rao Lower Bound (PCRLB)) to quantify the relationship between radar-target geometry and target localization performance [13]–[15]. However, studies optimizing sensor placement often oversimplify by using a constant covariance matrix for range measurements, neglecting signal path attenuation effects [4]. The approach in [4] focuses on optimizing the angle subtended between the radars and targets but does not encourage the mobile radars to move closer to the target. In [16], the authors provide the error covariance matrix of a Best Linear Unbiased Estimator (BLUE) estimator for the range measurement using a first-order Taylor Series approximation of a distance dependent range (DDR) measurement model, which yields the Fisher Information Matrix (FIM). Similarly, we calculate the FIM for the DDR measurement model from [16], where the covariance matrix of the range measurement is directly proportional to the received power at the radar after signal path attenuation, but we do not apply a first-order Taylor Series approximation to the DDR measurement model. This strategy encourages mobile radars to move closer to the targets, thereby increasing the Signal Noise Ratio (SNR) and improving radar range measurements.

Another major challenge in scenarios with moving targets and movable radars is that the optimal radar-target

Manuscript received May 29, 2024; revised December 18, 2024; accepted December 30, 2024.
* denotes equal contribution of authors. Research was sponsored by the Army Research Laboratory and was accomplished under Cooperative Agreement Number W911NF-23-2-0014. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Army Research Laboratory or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation herein.

(a) Search and rescue in forest fire (b) Enemy drone localization and tracking

Fig. 1: Subfigure (a) shows the civilian application of a search and rescue mission where mobile radars navigate through a dense forest (on fire) to quickly localize and track lost victims. Subfigure (b) shows the military application of enemy drone localization and tracking in a city environment during a cloudy day.

geometry may change substantially over time. Consequently, many studies use Model Predictive Controller (MPC) [17] for choosing the control inputs of mobile sensors over a finite horizon in response to uncertain target state estimates [18]–[22]. However, natural choices for the MPC objective in this problem, such as minimizing/maximizing the D-optimality of PCRLB/PFIM, lead to challenging nonlinear optimization problems that are computationally prohibitive to solve exactly. To address this issue, we instead utilize Model Predictive Path Integral (MPPI) to approximate the optimal control input distribution for the mobile radars. MPPI enables the handling of discontinuous, non-smooth objectives through highly parallelized computation for real-time control [23], [24]. However, as the number of agents, planning horizon, and the actuation space increases, parameter estimation, particularly for the covariance, requires an exponentially growing number of samples. To mitigate this issue and reduce the number of proposal distribution updates and trajectory samples needed in the high-dimensional control space, we integrate advanced Adaptive Importance Sampling (AIS) techniques, such as cross-entropy weighting with covariance shrinkage [25], to prevent overfitting of the sample covariance matrix.

This paper introduces a improved pipeline for continuously optimizing radar placement, aiming to overcome the previously mentioned challenges. We focus on range-only radars, deriving objective functions using FIMs that incorporate signal path attenuation in the range measurement model. These functions are optimized with respect to radar control inputs using MPPI control. Crucially, MPPI control, which is rapidly gaining popularity in control research, offers enhanced flexibility for real-time optimization of nonlinear, discontinuous, and non-differentiable objectives. These advantages enable MPPI to handle future extensions of the proposed objective, such as incorporating more complex signal measurement models or adding more constraints through occupancy grid maps.

The paper is structured as follows: Section II describes the current literature; Section III describes the simulation assumptions and notation; Section IV describes the radar signal model and DDR measurement model; Section V formulates the targets' state-space model and the radars' kinematic model; Section VI derives the FIM which accounts for the radar signal attenuation; Section VII formulates our objective function based on the derived FIMs and outlines our MPPI controller, and Section VIII discusses the results of the proposed approach for three simulation scenarios.

II. Related Work

We explore the literature on optimizing sensor trajectories using the FIM of the likelihood function of the range measurements, where the likelihood function parameters correspond to the target positions. Various forms of the FIM establish a fundamental estimation lower bound limit for the mean squared error covariance matrix [26]. Therefore, many authors construct an objective as the D/A-optimality [27] of the PFIM, parameterized by sensor and estimated target positions [28]. We divide these works into two categories: those that consider static-sensors versus those that consider dynamic-sensors.

A. Static Sensors

Static sensors lack dynamic models or control inputs. Optimizing their placement involves determining where they would be positioned if placed instantaneously. Many studies on optimal radar placement use range measurements assumed to follow Gaussian distributions to construct the FIM, where the covariance may or may not vary with sensor-target distance [13], [14], [16], [29]–[31].

Several studies (e.g., [13], [29]) suggest that when the range measurement covariance is distance-independent, optimal sensor placement depends only on the angle between the sensor and the target. However, this approach overlooks the impact of signal path attenuation, which ideally results in higher information gain when the sensor is closer to the target.

To address this, [14], [32] adopt a heuristic assumption concerning covariance, combining a constant diagonal matrix with another diagonal matrix which is a function of the range between sensors and targets. However, signal processing fundamentals provide more realistic range measurement distributions. We take a similar approach to [16], [31], where the covariance, which scales with the return radar power, is determined by optimizing the Taylor series expansion of a match filter over time delay for a Narrowband return signal subject to attenuation and passing through an Additive White Gaussian Noise (AWGN) channel.

B. Dynamic Sensors

Practically, radars cannot simply "appear" at the optimal location; they require constrained dynamics and path planning. Moreover, dynamic sensors are essential to track dynamic targets, as the optimal radar configuration at one time step may not remain optimal at the next time step.

Historically, two main approaches have been used for optimal radar placement in target localization of dynamic targets: 1) moving the radars to track dynamic targets while maintaining an optimal radar-target localization geometry [4], [33], and 2) "spraying" radars over a large area to ensure high spatial coverage and multiplicity, then selecting which radars to employ for sensing [15]. This paper primarily focuses on the first approach, which is more appropriate given a limited number of radars.

Since the objective is a function of a finite horizon of time steps, many studies utilize an MPC [17] objective for optimizing the control inputs of mobile sensors [18]–[22]. Consequently, rather than directly optimizing the radar's position at a specific time step, as in the static-sensor formulation, the optimal radar trajectory is indirectly determined by optimizing the control inputs to the radar kinematic model using the method of direct shooting [34].

The finite horizon MPC objective, formulated in terms the PFIM, may be efficiently computed through recursive evaluation using discrete-time nonlinear filtering [12]. Several studies on cooperative range-based underwater target localization with multiple autonomous surface vehicles adopt this approach, while also considering constraints such as distances between targets and sensors [9], [21], [22]. However, these methods often assume a constant diagonal covariance matrix for range measurements, irrespective of range, typically involve slow-moving sensors and targets, assume the target remains at a constant depth, and may assume a zero covariance for the target dynamic model. Moreover, solving the resulting nonlinear optimization problems remains challenging, with several recent approaches leveraging `CasADi` [35].

He et al [36] employs a more complex measurement model which integrates Multi-Target Tracking (MTT) characteristics such as data association, track insertion, and track deletion using the Joint Probabilistic Data Association Filter (JPDAF). However, this work assumes a constant diagonal covariance matrix for range and bearing measurements. Furthermore, the trajectory optimization only optimizes the sensor's heading angle one time step ahead due to the highly nonlinear objective function, potentially leading to myopic and suboptimal trajectories.

Similar to information-based MPC objectives like the D/A optimality of the FIMs, [37] minimizes the mutual information gain between the target state distributions before and after the update step of a linear Kalman Filter, while parameterizing the measurement noise covariance based on the robot's state. However, the measurement covariance only scales linearly with the sensor-target distance and is clipped at an extremely small variance of 0.01 meters, restricting experiments to small square arenas of 40 to 60 meters.

Hence, the proposed approach aims to enhance target localization through continuously optimizing radar placement by integrating an improved range measurement model into the objective function. Additionally, we leverage MPPI control to increase flexibility in solving the MPC objective over non-smooth and highly nonlinear objectives, enabling real-time control capabilities.

III. Preliminaries

We consider N radars and M targets. The state for radar n at time step k is denoted as $\boldsymbol{\chi}_k^{R_n} = [x, y, z, \theta, v, \omega]$, while for target m, it is $\boldsymbol{\chi}_k^{T_m} = [x, y, z, \dot{x}, \dot{y}, \dot{z}]$, where θ is the radar's heading angle with respect to the x-axis of the radar coordinate frame, v is the heading velocity, and ω is the angular velocity. The positions for radar n and target m are defined as $\boldsymbol{\chi}_{xyz}^{R_n} = [x, y, z]$ and $\boldsymbol{\chi}_{xyz}^{T_m} = [x, y, z]$, respectively. The time step k has a duration of Δt. We denote the Hadamard product (elementwise multiplication) as $\odot$, and $\otimes$ is the kronecker product. Additionally, we use the notation $\Delta_x^y = \nabla_x \nabla_y^T$ for the Hessian operator.

To highlight the focus of our paper, we make the following simplifying assumptions to our analysis:

1) *Perfect data association and target detection*. We assume that the process of selecting which measurement(s) to incorporate into the target state estimator, is perfect, and the radars constantly track the targets.
2) *Doppler shift can be neglected*. For slow relative velocity ν between the radar and target, when the Doppler shift is much smaller than the effective bandwidth of the return signal, $\frac{f_d}{\beta} \ll 1$, the Doppler shift $f_d = 2 \times \frac{f_c}{c} \left(\frac{(\boldsymbol{\chi}_{xyz}^T - \boldsymbol{\chi}_{xyz}^R)^T}{|\boldsymbol{\chi}_{xyz}^T - \boldsymbol{\chi}_{xyz}^R|} \right) \nu$ [38] may be disregarded in the radar return signal. f_c is the carrier frequency and c is the speed of light.
3) *Radars have zero altitude*. The targets move in R^3, but the radars move on a plane at $z = 0$. This is a common assumption that radars are attached to ground vehicles.
4) *We assume targets follow a constant velocity model with acceleration process noise and therefore do not account for drastic target maneuvers in our model.*

With these assumptions, the following sections detail each computation block needed to continuously optimize radar placement (flow diagram depicted in Figure 2). We first discuss the more realistic range measurement model based on [16] compared to previous works such as [9], [13], [22].

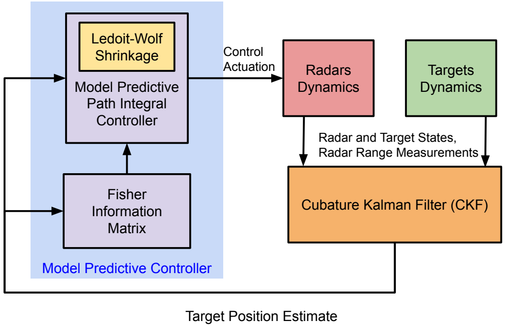

Fig. 2: Block Diagram of proposed approach.

IV. Radar Signal Model

Following the approach in [16], we assume there are N radars transmitting orthogonal narrowband signals to track M targets. In this setup, the received signal at radar n for target m is given by:

$$r^{nm}(t) = \sum_{b=1}^{N} \zeta^{nm} e^{-j2\pi f_c \tau^{nm}} s^b(t - \tau^{nm}) + \epsilon_a^n(t), \quad (1)$$

where the time delay is:

$$\tau^{nm} = \frac{\|\boldsymbol{\chi}_{xyz}^{R_n} - \boldsymbol{\chi}_{xyz}^{T_m}\| + \|\boldsymbol{\chi}_{xyz}^{R_b} - \boldsymbol{\chi}_{xyz}^{T_m}\|}{c}, \quad (2)$$

and the set of orthonormal lowpass equivalent signals $s(t)$ such that there is no signal interference when multiple radars transmit simultaneously. The return signal noise $\epsilon_a(t)$ is circularly symmetric, zero-mean, complex Gaussian noise with autocorrelation function $\sigma_a^2\delta(\tau)$. ζ is the amplitude of the return signal, where $|\zeta|^2$ is the radar power return specified by the radar equation [39]:

$$|\zeta|^2 = P_r = \frac{P_t \Lambda_t \Lambda_r \lambda^2 \Xi}{(4\pi)^3 \|\boldsymbol{\chi}_{xyz}^{R} - \boldsymbol{\chi}_{xyz}^{T}\|^4 L} \quad (3)$$

The radar parameters specifying the received power P_r $[W]$ at the radar include P_t as the transmit power $[W]$, Λ_t and Λ_r as the transmit and receive gains, L as the general loss factor, Ξ as the radar cross section $[m^2]$, and λ as the carrier frequency wavelength $[m]$.

It has been shown that the range measurement between radar n and target m approximately follows a Gaussian Probability Density Function (PDF) with variance parameterized by the radar-target geometry [13].

$$z^{nm} \sim \mathcal{N}\left(2\|\boldsymbol{\chi}_{xyz}^{R_n} - \boldsymbol{\chi}_{xyz}^{T_m}\|\ ,\ \frac{c^2\sigma_a^2}{8\pi^2 f_c^2 |\zeta^{nm}|^2}\right) \quad (4)$$

We denote all the radar parameters, signal parameters, and noise parameters as a constant $\Gamma = \frac{\sigma_a^2 \pi L}{2P_t\Lambda_t\Lambda_r\Xi}$. For brevity, we define shorthand notation as $\mu_{z^{nm}} = 2\|\boldsymbol{\chi}_{xyz}^{R_n} - \boldsymbol{\chi}_{xyz}^{T_m}\|$ and $\sigma_{z^{nm}} = \Gamma\|\boldsymbol{\chi}_{xyz}^{R_n} - \boldsymbol{\chi}_{xyz}^{T_m}\|^4$.

Existing research has explored optimal radar placement when the covariance matrix of the range measurement depends on the radar-target geometry [14], [31], [40], primarily focusing on stationary targets and radars. Given that the range measurement model relies on the positions of both the radars and targets, the next section describes the target and radar kinematic models. This will facilitate the definition of the targets' conditional transition probability in the targets' state space model.

V. Process Model

We describe the kinematic model of a target to finalize the probabilistic state space model for the concatenated state of all the targets, encompassing both a conditional transition probability and the conditional measurement probability described in IV.

A. Target Kinematics: Constant Velocity Model

The target kinematic model for a single target follows constant velocity motion with acceleration noise [41], [42], which specifies the single-target transition model:

$$\boldsymbol{\chi}_{k+1}^{T} = A_{\text{single}}\boldsymbol{\chi}_k^{T} + \epsilon_w \quad (5)$$

where

$$A_{\text{single}} = \begin{bmatrix} \boldsymbol{I}_{3\times3} & \Delta t\ \boldsymbol{I}_{3\times3} \\ \boldsymbol{0}_{3\times3} & \boldsymbol{I}_{3\times3} \end{bmatrix} \quad (6)$$

and ϵ_w is the acceleration noise which follows a zero mean Gaussian PDF with covariance

$$\boldsymbol{W}_{\text{single}} = \boldsymbol{W}_{\Delta t}\boldsymbol{\Sigma}_w\boldsymbol{W}_{\Delta t}^T \quad (7)$$

where

$$\boldsymbol{W}_{\Delta t} = \begin{bmatrix} \frac{1}{2}\Delta t^2\ \boldsymbol{I}_{3\times3} \\ \Delta t\ \ \boldsymbol{I}_{3\times3} \end{bmatrix}\ ,\ \Sigma_w = \boldsymbol{I}_{3x3}\ \sigma_W^2 \quad (8)$$

and σ_W determines the degree of deviation of the dynamics from constant velocity via the strength of the acceleration variance. Thus, the targets' transition model (for all targets concatenated as a column vector) is:

$$\boldsymbol{\chi}_{k+1}^{T_{1:M}} = A\boldsymbol{\chi}_k^{T_{1:M}} + \epsilon_W \quad (9)$$

$$A = \boldsymbol{I}_M \otimes A_{\text{single}} \quad W = \boldsymbol{I}_M \otimes W_{\text{single}} \quad (10)$$

where $I_M \in \Re^{M\times M}$ is the identity matrix. Similarly, the measurements of all radar-target pairs may be concatenated to measurement measurement model of the targets' state-space equations (based on Equation (4)):

$$\mathbf{z}^{T_{1:M}} = \begin{bmatrix} \mu_{z^{11}} & \dots & \mu_{z^{1M}} \\ \mu_{z^{21}} & \dots & \mu_{z^{2M}} \\ \vdots & \ddots & \vdots \\ \mu_{z^{N1}} & \dots & \mu_{z^{NM}} \end{bmatrix} + \begin{bmatrix} \sigma_{z^{11}} & \dots & \sigma_{z^{1M}} \\ \sigma_{z^{21}} & \dots & \sigma_{z^{2M}} \\ \vdots & \ddots & \vdots \\ \sigma_{z^{N1}} & \dots & \sigma_{z^{NM}} \end{bmatrix} \odot \epsilon_{\mathbf{z}} \quad (11)$$

$$= \mathbf{H}(\boldsymbol{\chi}_{\mathbf{xyz}}^{\mathbf{T_{1:M}}}) + \Omega \odot \epsilon_{\mathbf{z}} \quad (12)$$

where $\sigma_{z^{nm}}$ is the standard deviation of the range measurement for radar n and target m, as defined by the variance in Equation (4), while $\epsilon_{\mathbf{z}}$ is standard normal Gaussian PDF. Therefore the final state-space equations for the targets is shown in Equations (13) and (14):

$$\boldsymbol{\chi}_{k+1}^{T_{1:M}} = A\boldsymbol{\chi}_k^{T_{1:M}} + \epsilon_W \quad (13)$$

$$\mathbf{z}_{k+1}^{T_{1:M}} = \mathbf{H}\left((\boldsymbol{\chi}_{xyz}^{T_{1:M}})_{k+1}\right) + \Omega \odot \epsilon_{\mathbf{z}} \quad (14)$$

Given our focus on dynamic targets and dynamic radars, it's essential to outline a basic kinematic model where the radar position is determined by control inputs and adheres to specific control and dynamic constraints. Thus, we next define the second-order unicycle kinematic model for the radars.

B. Radar Kinematics: Second-Order Unicycle Model

Each radar follows the second-order unicycle model. We use the Euler discretization of the unicycle continuous dynamic model with heading acceleration and angular acceleration control inputs $u = [u_a, u_{\dot{\omega}}]^T$ to define the discrete time kinematic model:

$$\boldsymbol{\chi}^R_{k+1} = \boldsymbol{\chi}^R_k + G_k(\boldsymbol{\chi}^R_k, u_k) \tag{15}$$

with control, velocity, and angular velocity limits,

$$\begin{aligned} \underline{u}_a \leq u_{a_k} \leq \bar{u}_a \ \forall k \quad & \underline{v} \leq v_k \leq \bar{v} \ \forall k \\ \underline{u}_{\dot{\omega}} \leq u_{\dot{\omega}_k} \leq \bar{u}_{\dot{\omega}} \ \forall k \quad & \underline{\omega} \leq \omega_k \leq \bar{\omega} \ \forall k. \end{aligned} \tag{16}$$

With the radars' positions defined relative to control inputs, the following section outlines the derivation of the FIMs with respect to the targets' positions. This is based on the radar-target measurement model outlined in Equations (4) and (11), and the radars' positions parameterized by Equations (15) and (16).

VI. Fisher Information Matrix

The Standard Fisher Information Matrix (SFIM) measures the amount of information that an observable random variable z (range measurements) carries about an unknown parameter θ (targets' states) [12], [43]:

$$\boldsymbol{J}_D(\theta) = E\left[\left(\nabla_\theta \log f\left(z|\theta\right)\right)\left(\nabla_\theta \log f\left(z|\theta\right)\right)^T\right] \tag{17}$$

Under the assumption of independence, the range measurements likelihood in Equation (11) is represented as:

$$f\left(\mathbf{z}^{T_{1:M}}|\boldsymbol{\chi}^{T_{1:M}}_{xyz}; \boldsymbol{\chi}^{R_{1:N}}_{xyz}\right) = \prod_{n=1}^{N}\prod_{m=1}^{M} f(z^{nm}|\boldsymbol{\chi}^{T_m}_{xyz}; \boldsymbol{\chi}^{R_n}_{xyz}) \tag{18}$$

Given that the measurement is a Gaussian PDF, with the mean and covariance parameterized by the targets' positions, and that information of measurements is additive by the independence assumption, we use the derivation of general Gaussian FIM [44] to calculate the SFIM for a multiple targets' positions with multiple radars:

$$\begin{aligned} \boldsymbol{J}(\boldsymbol{\chi}^{T_m}_{xyz}; \boldsymbol{\chi}^{R_n}_{xyz}) = (\boldsymbol{\chi}^{T_m}_{xyz} - \boldsymbol{\chi}^{R_n}_{xyz})(\boldsymbol{\chi}^{T_m}_{xyz} - \boldsymbol{\chi}^{R_n}_{xyz})^T \times \\ \left(\frac{4}{\Gamma\|\boldsymbol{\chi}^R_{xyz} - \boldsymbol{\chi}^T_{xyz}\|^6} + \frac{8}{\|\boldsymbol{\chi}^R_{xyz} - \boldsymbol{\chi}^T_{xyz}\|^4}\right) \end{aligned} \tag{19}$$

$$\boldsymbol{J}(\boldsymbol{\chi}^{T_m}_{xyz}; \boldsymbol{\chi}^{R_{1:N}}_{xyz}) = \sum_{n=1}^{N} \boldsymbol{J}(\boldsymbol{\chi}^{T_m}_{xyz}; \boldsymbol{\chi}^{R_n}_{xyz}) \tag{20}$$

$$\begin{aligned} &\boldsymbol{J}(\boldsymbol{\chi}^{T_{1:m}}_{xyz}; \boldsymbol{\chi}^{R_{1:n}}_{xyz}) = \\ &\begin{bmatrix} \boldsymbol{J}(\boldsymbol{\chi}^{T_1}_{xyz}; \boldsymbol{\chi}^{R_{1:N}}_{xyz}) & 0 & \dots & 0 \\ 0 & \boldsymbol{J}(\boldsymbol{\chi}^{T_2}_{xyz}; \boldsymbol{\chi}^{R_{1:N}}_{xyz}) & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & \boldsymbol{J}(\boldsymbol{\chi}^{T_M}_{xyz}; \boldsymbol{\chi}^{R_{1:N}}_{xyz}) \end{bmatrix} \end{aligned} \tag{21}$$

Equation Equation (19) is the SFIM of a single target position parameterized by a single radar and Equation (20) is the SFIM of a single target positions parameterized by multiple radars.

When there is prior knowledge about the dynamics of the target state (Equation (5)), then this information may be incorporated into the SFIM to form the PFIM:

$$\boldsymbol{J} = \boldsymbol{J}_D + \boldsymbol{J}_P \tag{22}$$

$$\boldsymbol{J}_P = E[\Delta^{\boldsymbol{\chi}^T}_{\boldsymbol{\chi}^T} \log f(\boldsymbol{\chi}^T)] \tag{23}$$

Here, $\boldsymbol{J}_D$ represents the information about the unknown target state obtained from the data, while $\boldsymbol{J}_P$ reflects the information about the unknown target state based on prior assumptions. Now, considering that the targets' state includes velocity, we update Equation (21) to be $\boldsymbol{J}_D$ for the PFIM, with the block matrices as:

$$\boldsymbol{J}(\boldsymbol{\chi}^{T_n}; \boldsymbol{\chi}^{R_{1:N}}_{xyz}) = \begin{bmatrix} \boldsymbol{J}(\boldsymbol{\chi}^{T_n}_{xyz}; \boldsymbol{\chi}^{R_{1:N}}_{xyz}) & 0 \\ 0 & 0 \end{bmatrix} \tag{24}$$

For discrete time filtering systems, such as Equations (13) and (14), the Equation (23) is recursively computed at time $k+1$ from time step k using the famous recursion updates in [12]. The recursion equations are written as

$$\boldsymbol{J}_{k+1}(\boldsymbol{\chi}^{T_{1:M}}_{k+1}) = \boldsymbol{D}^{22}_k - \boldsymbol{D}^{21}_k(\boldsymbol{J}_k + \boldsymbol{D}^{11}_k)^{-1}\boldsymbol{D}^{12}_k \tag{25}$$

where

$$\boldsymbol{D}^{11}_k = E[-\Delta^{\boldsymbol{\chi}^{T_{1:M}}_k}_{\boldsymbol{\chi}^{T_{1:M}}_k} \log p(\boldsymbol{\chi}^{T_{1:M}}_{k+1}|\boldsymbol{\chi}^{T_{1:M}}_k)] \tag{26}$$

$$\boldsymbol{D}^{12}_k = E[-\Delta^{\boldsymbol{\chi}^{T_{1:M}}_{k+1}}_{\boldsymbol{\chi}^{T_{1:M}}_k} \log p(\boldsymbol{\chi}_{k+1}|\boldsymbol{\chi}_k)] = [\boldsymbol{D}^{21}_k]^T \tag{27}$$

$$\begin{aligned} \boldsymbol{D}^{22}_k = E\left[-\Delta^{\boldsymbol{\chi}^{T_{1:M}}_{k+1}}_{\boldsymbol{\chi}^{T_{1:M}}_{k+1}} \log p(\boldsymbol{\chi}^{T_{1:M}}_{k+1}|\boldsymbol{\chi}^{T_{1:M}}_k)\right] + \\ E\left[-\Delta^{\boldsymbol{\chi}^{T_{1:M}}_{k+1}}_{\boldsymbol{\chi}^{T_{1:M}}_{k+1}} \log p(z_{k+1}|\boldsymbol{\chi}^{T_{1:M}}_{k+1})\right] \end{aligned} \tag{28}$$

Each D term in equation Equation (25) is (PDF regularity assumptions) as

$$\boldsymbol{D}^{11}_k = \boldsymbol{A}^T\boldsymbol{W}^{-1}\boldsymbol{A} \tag{29}$$

$$\boldsymbol{D}^{12}_k = -\boldsymbol{A}^T\boldsymbol{W}^{-1} \tag{30}$$

$$\boldsymbol{D}^{22}_k = \boldsymbol{W}^{-1} + E_{f(\boldsymbol{\chi}^{T_{1:M}}_{1:k+1})}\left[\boldsymbol{J}_{\boldsymbol{D}_k}(\boldsymbol{\chi}^{T_{1:M}}_{k+1}))\right] \tag{31}$$

where due to the Markov assumption on the transition model in the targets' state space equation (Equation (9)), the joint PDF $f(\boldsymbol{\chi}^{T_{1:M}}_{1:k})$ is simplified as

$$f(\boldsymbol{\chi}^{T_{1:M}}_{1:k}) = f(\boldsymbol{\chi}^{T_{1:M}}_0)\prod_{t=1}^{k+1} f(\boldsymbol{\chi}^{T_{1:M}}_t|\boldsymbol{\chi}^{T_{1:M}}_{t-1}) \tag{32}$$

After substituting equations Equations (29) to (31) into Equation (25), and leveraging the Woodbury matrix identity [26], we arrive at [45]:

$$\begin{aligned} \boldsymbol{J}_{k+1}(\boldsymbol{\chi}^{T_{1:M}}_{k+1}) = (\boldsymbol{W} + \boldsymbol{A}^T\boldsymbol{J}_k^{-1}\boldsymbol{A})^{-1} + \\ E_{f(\boldsymbol{\chi}^{T_{1:M}}_{1:k+1})}\left[\boldsymbol{J}_{\boldsymbol{D}_k}(\boldsymbol{\chi}^{T_{1:M}}_{k+1}))\right] \end{aligned} \tag{33}$$

Equation (33) is preferred over Equation (25) for numerical stability. This preference arises because it avoids the inversion of the W matrix, which can be close to singular in cases where the conditional transition PDF is described by constant velocity models.

To compute the expectation over the target state for Equation (31), as well as objective terms involving Equation (21) or any other objective term in the MPC objective that involves the targets' states, we employ the Cubature Kalman Filter (CKF) to describe Equation (32).

CKFs efficiently leverage third-order spherical cubature integration into both the prediction and correction steps of an additive Gaussian noise Kalman filter [46]. To compute the expectations described in the previous paragraph, we propagate the sigma points of the CKF multiple time steps into the future, employing a number of sigma points equal to twice the dimension of the state variable. Compared to a Particle Filter (PF) [46], the CKF has lower computational complexity, as the number of evaluation points scales linearly with the state size, while the number of particles in a PF may substantially increase.

With our current formulations in Sections IV to VI, we formulate a discrete-time finite-horizon MPC objective using dynamic, control, and mission-aware constraints.

VII. Model Predictive Controller Objective

MPC is an optimal control technique that determines the control actuation minimizing a cost function for a constrained dynamical system over a finite, receding horizon [47], [48]. At each time step of the receding horizon objective, we enforce minimum distances between radars and targets to prevent potential destruction by adversaries using soft penalties. Similarly, we ensure safe separation distances between mobile radars to prevent collisions. Additionally, we incorporate the radar's acceleration, angular acceleration, angular speed, and heading velocity constraints as described in Equation (16) into the dynamic equation using clipping functions. Finally, we directly integrate the radar's dynamic equation (Equation (16)) into the objective using direct shooting [34].

As a result, converting hard constraints to soft/implicit constraints results in an unconstrained, discontinuous, and non-differentiable optimization formulation. Lastly, we incorporate the uncertainty of the targets' state by taking expectation of the MPC objective function with respect to the CKF estimator of the targets' state:

$$\begin{aligned}S(\boldsymbol{u}_{k:k+K+1};\boldsymbol{\chi}^{T_{1:M}}_{k:k+K+1}) = E_{\tilde{f}}\big[&S_{traj}(\boldsymbol{u}_{k:k+K+1};\boldsymbol{\chi}^{T_{1:M}}_{k:k+K+1})\\ &+\alpha_1 S_{R2R}(\boldsymbol{u}_{k:k+K})\\ &+\alpha_2 S_{R2T}(\boldsymbol{u}_{k:k+K};\boldsymbol{\chi}^{T_{1:M}}_{k:k+K})\big]\end{aligned} \tag{34}$$

where

$$\begin{aligned}&S_{traj}\left(\boldsymbol{\chi}^{R_{1:N}}_{k:k+K+1};\boldsymbol{\chi}^{T_{1:M}}_{k:k+K+1}\right) =\\ &\quad\sum_{t=k}^{k+K+1}\gamma^{t-k}\log\det\left(J(\boldsymbol{\chi}^{T_{1:M}}_t;\boldsymbol{\chi}^{R_{1:M}}_t\right)\\ &S_{R2T}(\boldsymbol{u}_{k:k+K+1}) =\\ &\quad\sum_{t=k}^{k+K+1}\gamma^{t-k}\sum_{m=1}^{M}\sum_{n=1}^{N}\mathbf{1}\left(\|\boldsymbol{\chi}^{T_n}_t-\boldsymbol{\chi}^{R_m}_t(\boldsymbol{u}_t)\|\le R_{R2T}\right)\\ &S_{R2R}(\boldsymbol{u}_{k:k+K+1}) =\\ &\quad\sum_{t=k}^{k+K+1}\gamma^{t-k}\sum_{\{i<j:i,j\in M\}}\mathbf{1}\left(\|\boldsymbol{\chi}^{R_i}_t(\boldsymbol{u}_t)-\boldsymbol{\chi}^{R_j}_t(\boldsymbol{u}_t)\|\le R_{R2R}\right)\end{aligned} \tag{35}$$

where α_1, α_2 are hyper parameters that determine how much to penalize the radars for violating the soft constraints S_{R2R}, S_{R2T} respectively, γ is the discount factor, and K is the horizon length. The CKF estimator PDF on the targets' state over the horizon is

$$\tilde{f}(\boldsymbol{\chi}^{T_{1:M}}_{k:k+K+1}) = f_{ckf}(\boldsymbol{\chi}^{T_{1:M}}_k|\mathbf{z}_{1:k-1})\prod_{t=k+1}^{k+K+1} f(\boldsymbol{\chi}^{T_{1:M}}_t|\boldsymbol{\chi}^{T_{1:M}}_{t-1}) \tag{36}$$

This MPC problem requires an efficient solver that can handle highly nonlinear, discontinuous, and non-differentiable functions. Therefore, we employ a sampling-based strategy, MPPI [24], as discussed in the following section.

A. Model Predictive Path Integral

MPPI is a sampling-based MPC algorithm designed for optimizing nonlinear stochastic systems subject to complex cost functions in real time [49]. The assumption of MPPI is the control input is random and distributed as Gaussian random vector $\boldsymbol{v} \sim N(\boldsymbol{u}, \boldsymbol{\Sigma})$. Due to the stochasticity of the control input, rather than solving for a deterministic optimal control input we solve for an optimal control PDF Q^* [50]:

$$Q^* = \frac{1}{\eta}\exp\left(-\frac{1}{b}S(\boldsymbol{v}_{k:k+K+1})\right)P(\boldsymbol{v}_{k:k+K+1}) \tag{37}$$

where $P(\cdot)$ is the PDF of an control input sequence in an uncontrolled system ($\boldsymbol{u}_{k:k+K+1} = 0$), b is the temperature of the distribution, and η is the evidence. However, the optimal control PDF likely does not possess an analytical form. Therefore, by minimizing the Kullback-Leibler Divergence (KL-Divergence) between Q^* and a variational Gaussian PDF Q parameterized by u, Σ [24] with respect to u gives the optimal control at k as [24]

$$\boldsymbol{u}^*_k = \int q^*(\boldsymbol{V})\boldsymbol{v}_k d\boldsymbol{V} \tag{38}$$

Since we cannot express Q^* analytically in Equation (38), we leverage AIS to approximate the expectation as follows [24], [25]:

1) sample control trajectories $\boldsymbol{v}^{(i)}_{k:k+K+1}$ from distribution Q with mean $\boldsymbol{u}_{k:k+K+1}$ and covariance Σ
2) evaluate the objective function for each control trajectory sample $S(\boldsymbol{v}^{(i)}_{k:k+K+1})$
3) compute the importance sampling corrective weights for each sample $\tilde{w}^{(i)} \propto S(\boldsymbol{v}^{(i)})$
4) adapt the mean and covariance of the PDF Q
5) repeat steps 1-4

Due to the high dimensionality of the targets' control input over a horizon and the limited number of samples in MPPI, the sample covariance matrix for control inputs can overfit, leading to inaccurate variances [51]. This often results in poor control inputs from the proposal distribution of MPPI. To address this, we adopt the Ledoit-Wolf (LW) for AIS, which approximates Oracle Approximating

Shrinkage (OAS) shrinkage and helps mitigate overfitting, following the method described in [25] [1].

```
Algorithm 1: Continuous Optimal Radar Placement
Data: $\chi_0^{R_{1:N}}$
Result: $\hat{\chi}_{1:T}^{T_{1:M}}, \chi_{1:T}^{R_{1:N}}$
for $k = 1 : T$ do
    $\chi_{k+1}^{T_{1:M}}$ = TransitionFn$(\chi_k^{T_{1:M}})$
    $\hat{\chi}_{k+1}^{T_{1:M}}$ = CKF.predict$(\hat{\chi}_k^{T_{1:M}})$
    if $k\%T_{control} = 0$ then
        $\hat{\chi}_{k+1:k+K+1}^{T_{1:M}} \leftarrow$ CKF.propogate$(\hat{\chi}_{k+1}^{T_{1:M}})$;
        $\mathbf{u}_{k:k+K+1} \leftarrow$
          MPPI$(\hat{\chi}_{k+1:k+K+1}^{T_{1:M}}, \chi_k^{R_{1:M}},)$;
    end
    $\chi_{k+1}^{R_{1:N}} \leftarrow$ Actuate$(\chi_k^{R_{1:N}}, \mathbf{u}_k)$ ;
    $\mathbf{z}_{k+1}^{R_{1:N}}$ = MeasureFn$(\chi_{k+1}^{T_{1:M}}, \chi_{k+1}^{R_{1:N}})$ ;
    $\hat{\chi}^{T_{1:M}}$ = CKF.update$(\chi_{k+1}^{R_{1:N}}, \mathbf{z}_{k+1}^{R_{1:N}})$ ;
end
```

The integration of the algorithms outlined in Sections IV to VII forms the enhanced continuous optimal radar placement pipeline, described in Algorithm 1. This pipeline includes two key enhancements: an improved range measurement model, which significantly impacts the FIM, and a flexible and computationally efficient MPC objective optimizer using MPPI.

VIII. EXPERIMENTS

We examine three scenarios: 1. an "underdetermined" system, with fewer radars than targets, 2. an "overdetermined" system, with more radars than targets, and 3. an "full" system, with equal amount of radar and targets. We highlight not only the poorer performance of a misspecified measurement model (DDR versus constant covariance range (CCR) measurement model) but also the pathologies in radar movement caused by the controller using a FIM-based objective specified from the CCR measurement model. The proposed approach outperforms methods employing stationary radars and simplified range measurement models in target localization, achieving a reduction in mean RMSE and the upper tail of the 90% HDI by 38-74% and 33-79% respectively, across 500 Monte Carlo (MC) trials at all time steps.

A. Simulation

The hyperparameters for each experiment for the MPPI control, the AIS, the radar configurations , and the MPC objective is found in Tables I to III.

Radar Parameter	Value
Carrier Frequency	1×10^8
Radar Loss	1
Transmitter Gain	200
Receiver Gain	200
Radius (for SNR)	500 [m]
Radar Cross Section	1 [m^2]
Power Transmit	1000 [W]
SNR	-20 [dB]

TABLE I: The Radar Equation configuration settings for each experiment

MPPI Parameters	Value
u_a std	25 [m/s^2]
$u_{\dot{w}}$ std	45°
Horizon	15
Number of Sample Trajectories	200
Number of MPPI Sub-iterations	5
Temperature	0.1
Elite Threshold	0.9
AIS Method	Cross Entropy
Covariance Shrinkage Method	Ledoit-Wolf Shrinkage
Radar Cross Section	1

TABLE II: The MPPI settings for each experiment

MPC Parameter	Value
γ	0.95
$R2T$	125 [m]
$R2R$	10 [m]
α_1	500
α_2	1000

TABLE III: The MPC configuration settings for each experiment

We assume that the radars move on the ground plane (z=0), while the drones move in free space by the constant velocity kinematic model (Equation (5)) with acceleration noise constant $\sigma_W = \sqrt{10}$. For each experiment , we conduct 500 MC trials, where every MC trial is 600 time steps (1 minute in real time). The initial heading velocity and acceleration of the radars are both set to 0, and their initial positions in the XY-plane are uniformly sampled within a square with edge sizes 800 [m], centered at the origin (0 [m], 0 [m]). Although the targets may start in the same position for each MC trial, the random acceleration noise in the transition model (Equation (9)) leads to significant variations in the trajectories of the targets across trials. The observed range measurements at each time step are simulated using the more realistic range measurement model presented in Equation (4).

We compare our approach quantitatively and qualitatively with stationary radars. Additionally, using the MPPI controller for mobile radars, we compare the improved measurement model [16] to the CCR measurement model used in [9], [13], [22], [29]. This comparison not only empirically demonstrates the suboptimality of the constant covariance model as range measurements become noisier with increasing distance between the radar and the target

---

[1] The authors in [25] do not explicitly discuss Oracle Approximating Shrinkage (OAS) or Ledoit-Wolfe, but their publicly available code demonstrates methods for estimating the sample covariance, including Ledoit & Wolf, Schaffer & Strimmer, Rao-Blackwell, and OAS. A reference in the publicly available code is https://github.com/sisl/MPOPIS/wiki/MPOPIS-Details.

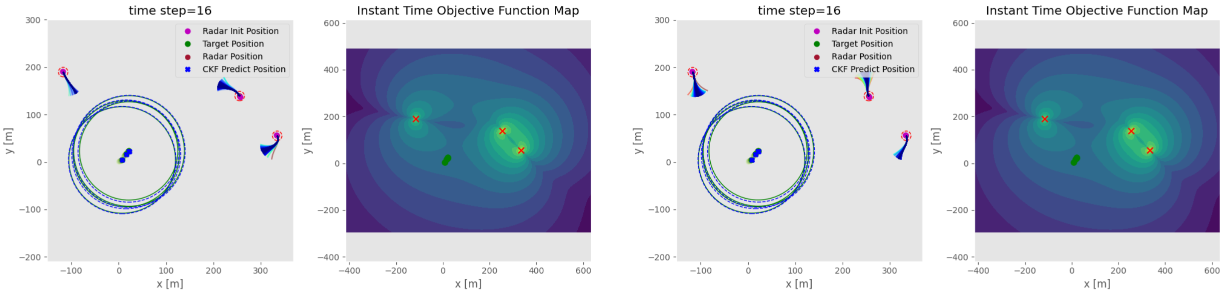

(a) Proposed approach: Each radar moves towards the targets.

(b) [9], [13], [22] range model: The upper right radar moves away from targets

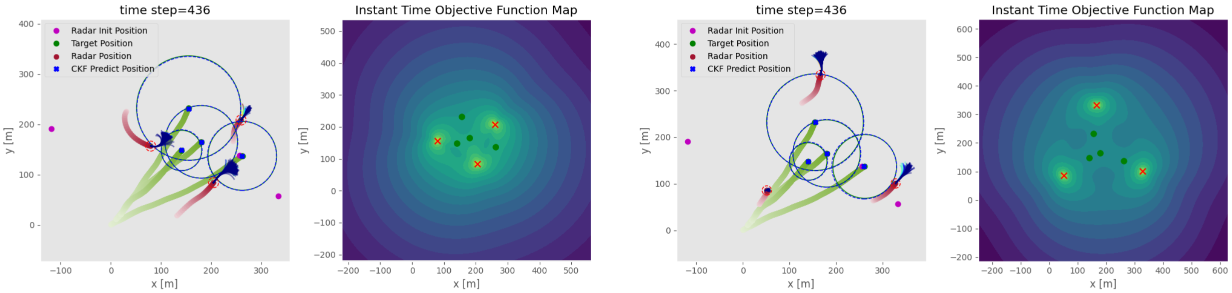

(c) Radars move toward targets as closely as feasible.

(d) Radars encircle targets, but the bottom most radar remains distant.

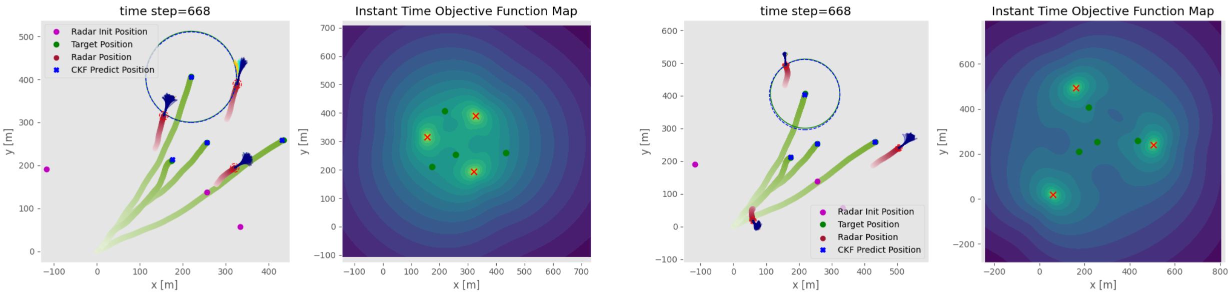

(e) Radars maintain circular formation with targets circle formation.

(f) Radars encircle the targets, but bottom most radar moves away from targets.

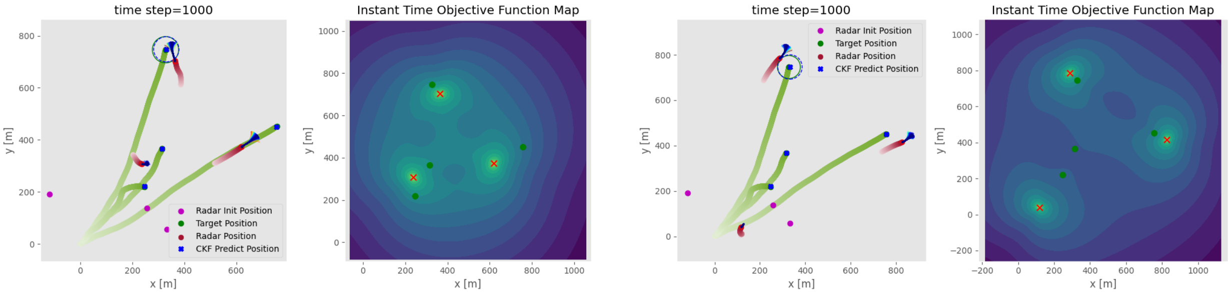

(g) Radars maintain a circular formation within the circle of targets.

(h) Radars maintain encirclement formation around targets cricle

Fig. 3: The left column shows snapshots of target localization using a range measurement model with covariance dependent on radar-target distance [16], while the right column depicts tracking with a range measurement model employing constant diagonal covariance [13]. In the first three rows, each subfigure's left panel shows the 3D projection onto 2D of continuously optimizing radar placement at a specific time. The colored lines emanating from the radars depict the weighted MPPI-planned trajectories, with blue indicating better paths and red indicating worse ones. Dashed red and green circles mark the radar's (10 [m]) and the target's (125 [m]) collision avoidance boundaries, respectively, while blue dashed circles represent the targets' collision avoidance boundary based on the CKF position estimate. The radar's trajectory is the last 25 time steps. The right panel displays the log determinant of the FIM (for the ddr measurement model) concerning the 2D target position for the specific time step, meshed over the free space area, and evaluated at each mesh point.

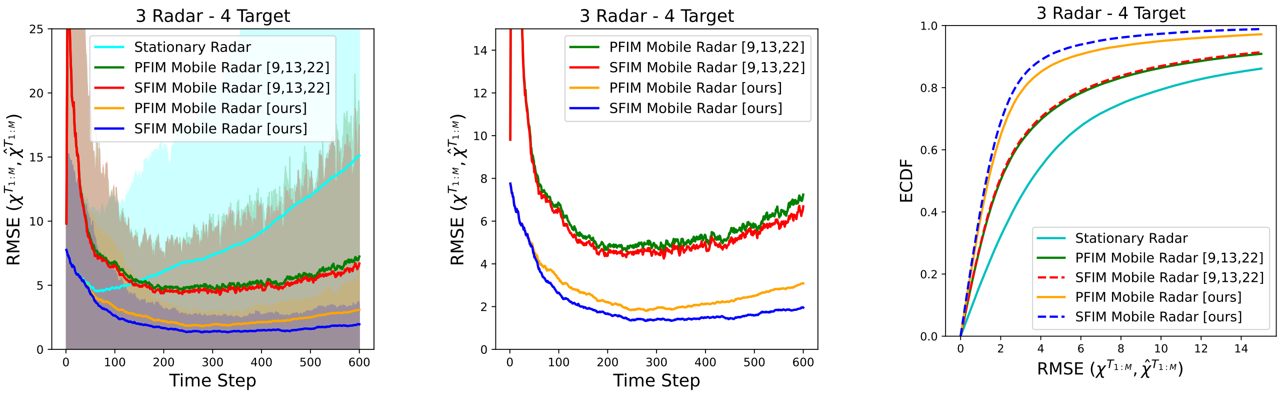

(a) 3 Radar - 4 Target: RMSE with 90% HDI

(b) 3 Radar - 4 Target: RMSE without stationary radar or HDI

(c) 3 Radar - 4 Target: RMSE ECDF

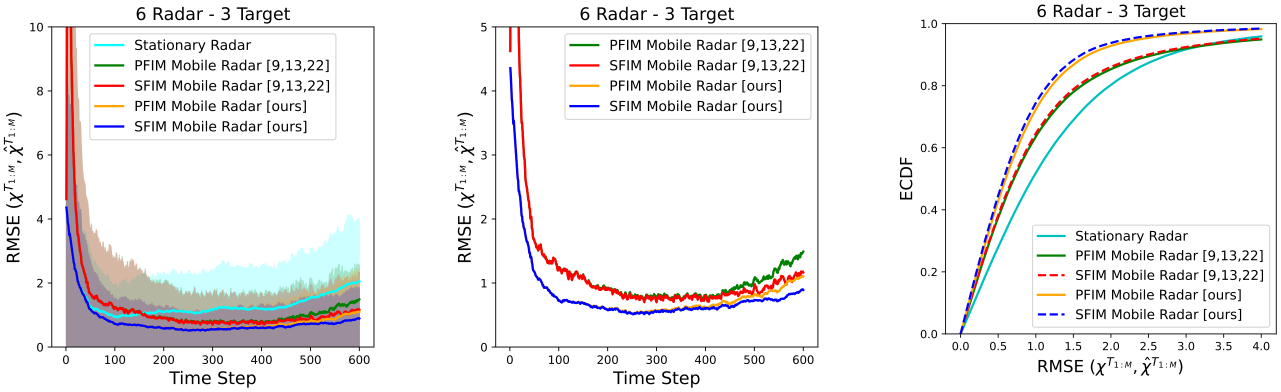

(d) 6 Radar - 3 Target: RMSE with 90% HDI

(e) 6 Radar - 3 Target: RMSE without stationary radar or HDI

(f) 6 Radar - 3 Target: RMSE ECDF

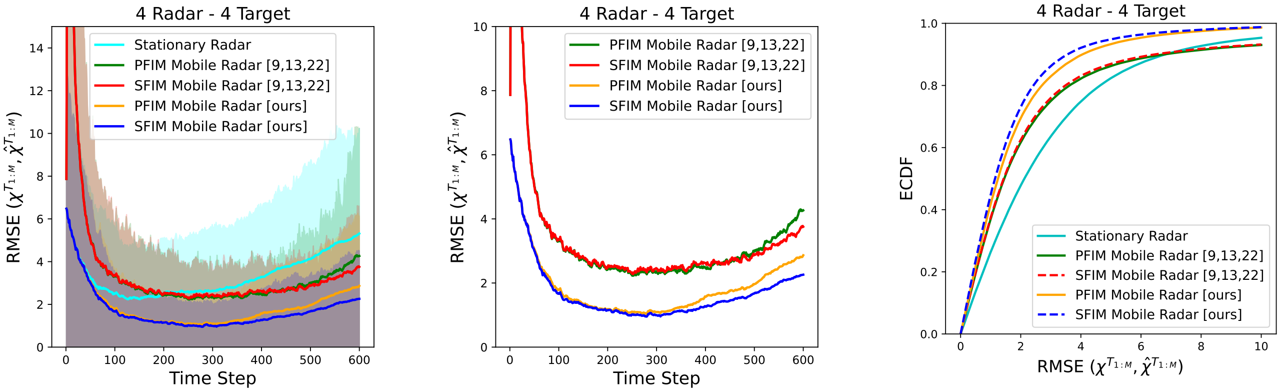

(g) 4 Radar - 4 Target: RMSE with 90% HDI

(h) 4 Radar - 4 Target: RMSE without stationary radar or HDI

(i) 4 Radar - 4 Target: RMSE ECDF

Fig. 4: The difference in target localization errors between constant covariance and radar-based covariance for 3 radars aand 4 targets, 6 radars aand 3 targets, and 4 radars and 4 targets for the SFIM and PFIM based objective via Root Mean Squared Error (RMSE) plot with 90% HDI and RMSE Emperical Cumulative Distribution Function (ECDF) plots. The left column shows the mean RMSE over the targets' true state versus the CKF predicted state over 500 MC trials. The lightly shaded region is the 90% highest density interval (HDI) of the RMSE over 500 MC trials. The middle column shows a zoom in of the left column subfigurese without stationary radar and 90% HDI RMSEs. The right column is the empirical ECDF of the CKF predicted state RMSE over time steps and MC trials. Higher ECDF is better at a given RMSE.

(due to model mispecification), but also highlights the pitfalls of a controller using an objective based on a CCR measurement model. In the CCR measurement model, we use an "optimistic" diagonal covariance matrix with elements $R2T^4 * \Gamma$ meters, where $R2T$ is the radar-target distance constraint.

For experimental purposes, we maintain constant noise power for the AWGN in Equation (1). We establish the desired SNR level at a specified distance using the formula:

$$\sigma_a^2 = \frac{M\tilde{P}_r}{10^{SNR/10}} \tag{39}$$

Here, $\tilde{P}_r$ represents the radar received power at the specified distance 500 [m]. For all experiments we set SNR = -20 [dB].

We note that the MPPI controller was written in JAX, which on a NVIDIA GeForce RTX 2080 Ti runs at 26-27 Hz (realtime), despite not fully optimizing our code.

B. Improved Target Localization

In over 500 MC trials, we empirically show that our approach exhibits significantly lower RMSE for all time steps (see Figure 4 (a,b,d,e,g,h)) compared to both the stationary radar approach and the mobile radar approach with constant range measurement. Additionally, the 90% HDI interval in localization performance for the CCR measurement model is significantly larger compared to the DDR measurement model presented in [16] (Figure 4 (c,f,i)). The PFIM-based objective for both range measurement models has slightly higher targets' state estimate RMSE compared to the SFIM-based objective.

However, Bayesian methods are notorious for overfitting, attributed to ensembling with model mispecification and the assumption of a single true model generating the data [52], [53]. We further examine the performance of the proposed approach inn the three simulation scenarios in the next subsections.

1. 3 Mobile Radars and 4 Targets

We initialize the targets such that all target directional velocities point towards the upper right (northeast) in the world coordinate frame, and begin around the same position (see Figure 3):

$$\chi_0^{T_{1:M}} = \begin{bmatrix} 0 & 0 & 55 & 20 & 10 & 0 \\ 15.4 & 15.32 & 70 & 15 & 20 & 0 \\ 10 & 10 & 55 & 17 & 19 & 0 \\ 20 & 20 & 45 & 6 & 8 & 0 \end{bmatrix}$$

where the first 3 columns are the 3D xyz coordinates, and the last 3 columns are the directional velocities.

Therefore, with fewer radars than targets, we anticipate a circular formation of radars near the targets when continuously optimizing radar placement, aiming to maintain nearest feasible proximity to the targets.

To illustrate the shortcomings of a controller using a CCR measurement model-based objective for radar placement, we showcase four time steps of a MC simulation realization, as depicted in Figure 3. For our proposed approach, at the beginning of the simulation, radars move towards targets, forming a circular arrangement within the nearest feasible proximity while adhering to radar-target constraints. As targets disperse, radars maintain this formation, but increased range measurement noise prompts a tighter circle, resulting in most targets being slightly outside the radar formation (left subfigure in Figure 3). In contrast, approaches with distance-independent noise maintain optimal angular spacing regardless of radar-target distance, often moving very far from the targets (Figure 3 (b,d,f,h)) and causing high localization errors (Figure 4 (a-i)). Counterintuitively, to achieve the optimal angle, mobile radars using the CCR measurement model move away from the targets, resulting in larger localization errors (Figure 3(b,f)). The target localization error over 1000 time steps for the extended realized MC simulation in Figure 3 is depicted in Figure 5.

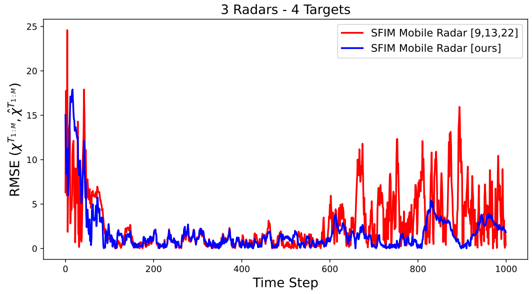

Fig. 5: RMSE for single simulation realization in Figure 3

2. 6 Mobile Radars and 3 Targets

We initialize each target's velocity direction to be different, with target 1 heading towards the upper right, target 2 towards the lower right, and target 3 towards the lower left:

$$\chi_0^{T_{1:M}} = \begin{bmatrix} 0 & 0 & 70 & 25 & 20 & 0 \\ -100.4 & -30.32 & 45 & 20 & -10 & 0 \\ 30 & 30 & 80 & -10 & -10 & 0 \end{bmatrix}$$

where the first 3 columns are the 3D xyz coordinates, and the last 3 columns are the directional velocities.

The initial positions of the targets are depicted in Appendix Figure 1. Since the targets' are spreading out over an open space, we hypothesize that the radars will send atleast 1 radar to follow each target, while keeping a reasonable distance to all other targets.

As the targets disperse in free space, the 6 radars coordinate to position 2 radars on either side of each target's direction of motion. This strategy intuitively avoids maintaining a circular formation, which would place the radars too far from the targets, leading to noisy measurements and inaccurate target localization (Appendix Figure 1 (a-h)). Optimal radar placement remains relatively similar regardless of whether a distance-independent or distance-dependent range noise model is used, though the radars are positioned farther from the targets in the former case which leads to slightly higher localization error and 90% HDI interval (Appendix Figure 1 (c,d,g,h)). In particular, the upper-most radar in Appendix Figure 1(h) is quite far away from the upper right-most target. The extended MC simulation over 1000 time steps is shown in Appendix Figure 1 with corresponding target localization error depicted in Figure 6.

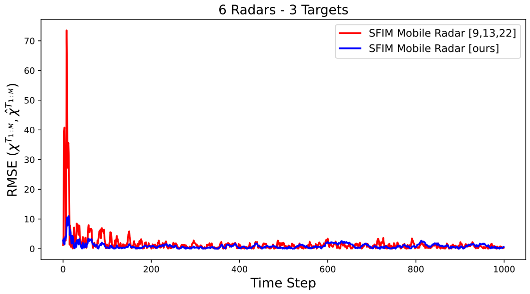

Fig. 6: RMSE for single simulation realization in Appendix Figure 1

3. *4 Mobile Radars and 4 Targets*

We initialize two targets to move towards the upper right and two targets to move towards the lower left. The initial positions of the targets are:

$$\chi_0^{T_{1:M}} = \begin{bmatrix} 0 & 15 & 70 & 15 & 15 & 0 \\ 40.4 & 15 & 70 & 15 & 15 & 0 \\ -30 & -15 & 45 & -10 & -10 & 0 \\ 20 & -15 & 45 & -10 & -10 & 0 \end{bmatrix}$$

where the first 3 columns are the 3D xyz coordinates, and the last 3 columns are the directional velocities.

We hypothesize that the radars will pair off, with two radars following each group of two targets.

As the targets disperse in free space, the 4 radars coordinate to position 2 radars on either side of each pair of targets' direction of motion. This strategy also avoids maintaining a circular formation, which would place the radars too far from the targets, leading to noisy measurements and inaccurate target localization (Appendix Figure 2 (a-h)). Optimal radar placement for the CCR measurement model based controller is quite different from the DDR measurement model based controller, where the former positions one of the radars extremely far away from all targets in the simulation (Appendix Figure 2 (b,d,f,i)) leading to high localization error. In particular, the upper-left most radar in Appendix Figure 2(h) is quite far all targets. The extended MC simulation over 1000 time steps is shown in Appendix Figure 2 with corresponding target localization error depicted in Figure 7.

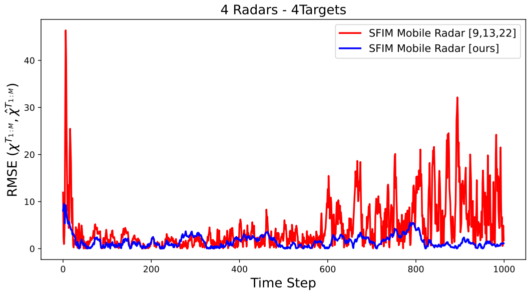

Fig. 7: RMSE for single simulation realization in Appendix Figure 2

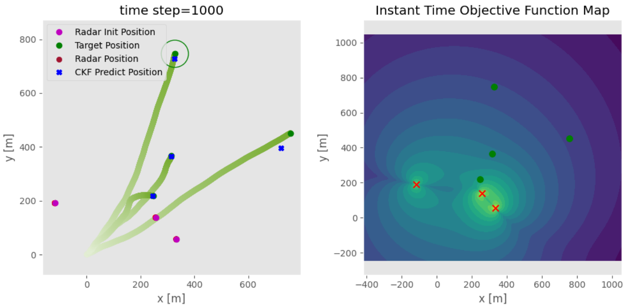

(a) 3 radars and 4 targets. The uppermost target shows significant height error, while the leftmost target exhibits substantial position error in the xy plane.

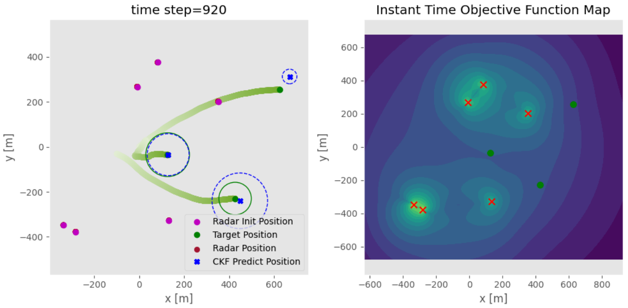

(b) 6 radars and 3 targets. The bottom rightmost target exhibits significant height error, while the top rightmost target displays a position error in the xy plane.

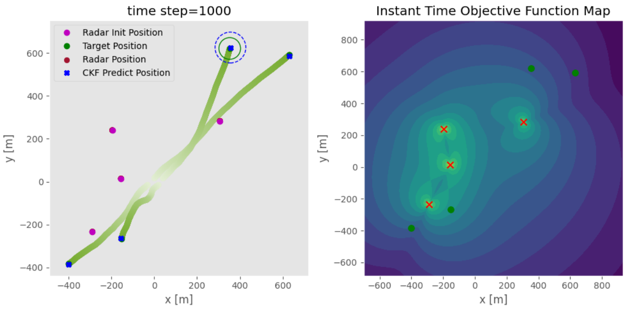

(c) 4 radars and 4 targets. The right uppermost target exhibits significant height error.

Fig. 8: Snapshots of each simulation scenario for stationary radars highlight the poor geometric layout of stationary radars for target localization.

4. *Stationary Radars*

Mobile radar methods, whether employing constant covariance or distance-dependent covariance in the range measurement model, outperform stationary radars (Figure 4 (a-i)). We present a qualitative example for each simulation scenario, illustrating the failure of stationary radar to track targets (Figure 8 (a-c)) and showing how radar positioning affects the information gain of measurements. As targets move farther away from the stationary radars, the noise in the range measurements increases drastically, resulting in higher localization errors.

IX. CONCLUSION

We implemented two improvements to optimize radar placement for multi-target localization: a more realistic radar range measurement model and an efficient stochastic optimal controller (Model Predictive Path Integral (MPPI)) for handling non-smooth, non-convex objectives in real time. Our approach significantly improves the Cubature Kalman Filter (CKF) target state estimate Root Mean Squared Error (RMSE) compared to stationary radars. Additionally, we show the pitfalls of the constant covariance range (CCR) measurement model, which may increase target localization error by moving radars in the opposite direction of target movement to optimize angles rather than the radar-target distance. Overall, the proposed strategy outperforms stationary radars and simplified range measurement models in target localization, achieving a 38-74% reduction in mean RMSE and an 33-79% reduction in the upper tail of the 90% Highest Density Interval (HDI) over 500 Monte Carlo (MC) trials across all time steps. For future work, we plan to integrate occupancy grid maps into the collision cost function to account for static and dynamic obstacles, such as trees and buildings, and further improve the radar range measurement model with traditional Multi-Target Tracking (MTT) characteristics.

REFERENCES

[1] R. Ding, Z. Wang, L. Jiang, and S. Zheng, "Radar target localization with multipath exploitation in dense clutter environments," *Applied Sciences*, vol. 13, no. 4, p. 2032, 2023.

[2] S. Venkatasubramanian, S. Gogineni, B. Kang, A. Pezeshki, M. Rangaswamy, and V. Tarokh, "Data-driven target localization using adaptive radar processing and convolutional neural networks," 2023.

[3] R. Grant, "The radar game," *Mitchell Institute*, 2010.

[4] M. Kumar and S. Mondal, "Recent developments on target tracking problems: A review," *Ocean Engineering*, vol. 236, p. 109558, 2021.

[5] Y. Yoon, S. Gruber, L. Krakow, and D. Pack, "Autonomous target detection and localization using cooperative unmanned aerial vehicles," in *Optimization and Cooperative Control Strategies: Proceedings of the 8th International Conference on Cooperative Control and Optimization*. Springer, 2009, pp. 195–205.

[6] C.-W. Juan and J.-S. Hu, "Object localization and tracking system using multiple ultrasonic sensors with newton–raphson optimization and kalman filtering techniques," *Applied Sciences*, vol. 11, no. 23, p. 11243, 2021.

[7] P. K. Ray and A. Mahajan, "A genetic algorithm-based approach to calculate the optimal configuration of ultrasonic sensors in a 3d position estimation system," *Robotics and Autonomous Systems*, vol. 41, no. 4, pp. 165–177, 2002.

[8] X. Wang, C. Fu, Z. Li, Y. Lai, and J. He, "Deepfusionmot: A 3d multi-object tracking framework based on camera-lidar fusion with deep association," *IEEE Robotics and Automation Letters*, vol. 7, no. 3, pp. 8260–8267, 2022.

[9] N. Crasta, D. Moreno-Salinas, A. M. Pascoal, and J. Aranda, "Multiple autonomous surface vehicle motion planning for cooperative range-based underwater target localization," *Annual Reviews in Control*, vol. 46, pp. 326–342, 2018.

[10] S. Z. Gurbuz, H. D. Griffiths, A. Charlish, M. Rangaswamy, M. S. Greco, and K. Bell, "An overview of cognitive radar: Past, present, and future," *IEEE Aerospace and Electronic Systems Magazine*, vol. 34, no. 12, pp. 6–18, 2019.

[11] M. A. Richards, J. Scheer, W. Holm, and W. Melvin, "Principles of modern radar, raleigh, nc," 2010.

[12] P. Tichavsky, C. H. Muravchik, and A. Nehorai, "Posterior cramér-rao bounds for discrete-time nonlinear filtering," *IEEE Transactions on signal processing*, vol. 46, no. 5, pp. 1386–1396, 1998.

[13] A. N. Bishop, B. Fidan, B. D. Anderson, K. Doğançay, and P. N. Pathirana, "Optimality analysis of sensor-target localization geometries," *Automatica*, vol. 46, no. 3, pp. 479–492, 2010.

[14] D. Moreno-Salinas, A. Pascoal, and J. Aranda, "Optimal sensor placement for multiple underwater target localization with acoustic range measurements," *IFAC Proceedings Volumes*, vol. 44, no. 1, pp. 12 825–12 832, 2011.

[15] F. Zhao, J. Shin, and J. Reich, "Information-driven dynamic sensor collaboration," *IEEE Signal processing magazine*, vol. 19, no. 2, pp. 61–72, 2002.

[16] H. Godrich, A. M. Haimovich, and R. S. Blum, "Target localization accuracy gain in mimo radar-based systems," *IEEE Transactions on Information Theory*, vol. 56, no. 6, pp. 2783–2803, 2010.

[17] N. Khaled and B. Pattel, *Practical design and application of model predictive control: MPC for MATLAB® and Simulink® Users*. Butterworth-Heinemann, 2018.

[18] S. Khan, J. Guivant, and X. Li, "Design and experimental validation of a robust model predictive control for the optimal trajectory tracking of a small-scale autonomous bulldozer," *Robotics and Autonomous Systems*, vol. 147, p. 103903, 2022.

[19] S. Uluskan, "Noncausal trajectory optimization for real-time range-only target localization by multiple uavs," *Aerospace Science and Technology*, vol. 99, p. 105558, 2020.

[20] X. Chu, Z. Liang, and Y. Li, "Trajectory optimization for rendezvous with bearing-only tracking," *Acta Astronautica*, vol. 171, pp. 311–322, 2020.

[21] Ø. Sture, P. Norgren, and M. Ludvigsen, "Trajectory planning for navigation aiding of autonomous underwater vehicles," *IEEE Access*, vol. 8, pp. 116 586–116 604, 2020.

[22] N. T. Hung, N. Crasta, D. Moreno-Salinas, A. M. Pascoal, and T. A. Johansen, "Range-based target localization and pursuit with autonomous vehicles: An approach using posterior crlb and model predictive control," *Robotics and Autonomous Systems*, vol. 132, p. 103608, 2020.

[23] J. Reuter and H. Homburger. Tutorial on model predictive path integral (mppi) control. [Online]. Available: https://www.syscop.de/event/tutorial-model-predictive-path-integral-mppi-control

[24] G. Williams, P. Drews, B. Goldfain, J. M. Rehg, and E. A. Theodorou, "Information-theoretic model predictive control: Theory and applications to autonomous driving," *IEEE Transactions on Robotics*, vol. 34, no. 6, pp. 1603–1622, 2018.

[25] D. M. Asmar, R. Senanayake, S. Manuel, and M. J. Kochenderfer, "Model predictive optimized path integral strategies," 2023.

[26] C. M. Bishop, "Pattern recognition and machine learning," *Springer google schola*, vol. 2, pp. 5–43, 2006.

[27] B. Jones, K. Allen-Moyer, and P. Goos, "A-optimal versus d-optimal design of screening experiments," *Journal of Quality Technology*, vol. 53, no. 4, pp. 369–382, 2021.

[28] F. Gustafsson and F. Gunnarsson, "Mobile positioning using wireless networks: possibilities and fundamental limitations based on available wireless network measurements," *IEEE Signal Processing Magazine*, vol. 22, no. 4, pp. 41–53, 2005.

[29] B. Jiang, A. N. Bishop, B. D. Anderson, and S. P. Drake, "Optimal path planning and sensor placement for mobile target detection," *Automatica*, vol. 60, pp. 127–139, 2015.

[30] S. Hu, F. Rusek, and O. Edfors, "Cramér-rao lower bounds for positioning with large intelligent surfaces," in *2017 IEEE 86th Vehicular Technology Conference (VTC-Fall)*. IEEE, 2017, pp. 1–6.

[31] J. Yan, H. Liu, W. Pu, and Z. Bao, "Exact fisher information matrix with state-dependent probability of detection," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 53, no. 3, pp. 1555–1561, 2017.

[32] D. B. Jourdan and N. Roy, “Optimal sensor placement for agent localization,” *ACM Transactions on Sensor Networks (TOSN)*, vol. 4, no. 3, pp. 1–40, 2008.

[33] S. Martínez and F. Bullo, “Optimal sensor placement and motion coordination for target tracking,” *Automatica*, vol. 42, no. 4, pp. 661–668, 2006.

[34] K. Hauser. Robotic systems. [Online]. Available: https://motion.cs.illinois.edu/RoboticSystems/OptimalControl.html

[35] J. A. Andersson, J. Gillis, G. Horn, J. B. Rawlings, and M. Diehl, “Casadi: a software framework for nonlinear optimization and optimal control,” *Mathematical Programming Computation*, vol. 11, pp. 1–36, 2019.

[36] S. He, H.-S. Shin, and A. Tsourdos, “Trajectory optimization for multitarget tracking using joint probabilistic data association filter,” *Journal of Guidance, Control, and Dynamics*, vol. 43, no. 1, pp. 170–178, 2020.

[37] X. Cai, B. Schlotfeldt, K. Khosoussi, N. Atanasov, G. J. Pappas, and J. P. How, “Non-monotone energy-aware information gathering for heterogeneous robot teams,” in *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2021, pp. 8859–8865.

[38] I. Shames, A. N. Bishop, M. Smith, and B. D. Anderson, “Doppler shift target localization,” *IEEE Transactions on Aerospace and Electronic Systems*, vol. 49, no. 1, pp. 266–276, 2013.

[39] D. K. Barton, *Radar equations for modern radar*. Artech House, 2013.

[40] Q. He, J. Hu, R. S. Blum, and Y. Wu, “Generalized cramér–rao bound for joint estimation of target position and velocity for active and passive radar networks,” *IEEE Transactions on Signal Processing*, vol. 64, no. 8, pp. 2078–2089, 2015.

[41] N. L. Baisa, “Derivation of a constant velocity motion model for visual tracking,” *arXiv preprint arXiv:2005.00844*, 2020.

[42] S. Blackman and R. Popoli, “Design and analysis of modern tracking systems(book),” *Norwood, MA: Artech House, 1999.*, 1999.

[43] S. M. Kay, *Fundamentals of statistical signal processing: estimation theory*. Prentice-Hall, Inc., 1993.

[44] ——, “Statistical signal processing: estimation theory,” *Prentice Hall*, vol. 1, pp. Chapter–3, 1993.

[45] L. Arienzo and M. Longo, “Posterior cramér-rao bound for range-based target tracking in sensor networks,” in *2009 IEEE/SP 15th Workshop on Statistical Signal Processing*, 2009, pp. 541–544.

[46] S. Särkkä and L. Svensson, *Bayesian filtering and smoothing*. Cambridge university press, 2023, vol. 17.

[47] M. Morari and J. H. Lee, “Model predictive control: past, present and future,” *Computers & chemical engineering*, vol. 23, no. 4-5, pp. 667–682, 1999.

[48] N. Khaled and B. Pattel, *Practical design and application of model predictive control: MPC for MATLAB® and Simulink® Users*. Butterworth-Heinemann, 2018.

[49] G. Williams, P. Drews, B. Goldfain, J. M. Rehg, and E. A. Theodorou, “Aggressive driving with model predictive path integral control,” in *2016 IEEE International Conference on Robotics and Automation (ICRA)*, 2016, pp. 1433–1440.

[50] E. Theodorou, J. Buchli, and S. Schaal, “A generalized path integral control approach to reinforcement learning,” *The Journal of Machine Learning Research*, vol. 11, pp. 3137–3181, 2010.

[51] J. Theiler, “The incredible shrinking covariance estimator,” in *Automatic target recognition XXII*, vol. 8391. SPIE, 2012, pp. 225–236.

[52] K. Monteith, J. L. Carroll, K. Seppi, and T. Martinez, “Turning bayesian model averaging into bayesian model combination,” in *The 2011 International Joint Conference on Neural Networks*, 2011, pp. 2657–2663.

[53] P. Domingos, “Bayesian averaging of classifiers and the overfitting problem,” in *ICML*, vol. 747, 2000, pp. 223–230.

Michael Potter is a Ph.D. student at Northeastern University (NEU) advised by Deniz Erdoğmuş of the Cognitive Systems Laboratory (CSL). He received his B.S, M.S., and M.S. degrees in Electrical and Computer Engineering from NEU and University of California Los Angeles (UCLA) in 2020, 2020, and 2022 respectively. His research interests are Bayesian Neural Networks, uncertainty quantification, and dynamics based manifold learning.

Shuo Tang received the BS degree in mechanical engineering from China Agricultural University, China and the MS degree in mechanical engineering from Northeastern University, Boston, MA, in 2018 and 2020, respectively. He is currently a PhD candidate in electrical and computer engineering at Northeastern University. His research interests include GNSS signal processing, sensor fusion and computational statistics.

Paul Ghanem received the B.E. and M.S. degrees in Electrical Engineering and Robotics in 2016 from the Lebanese American University and in 2019 from the University of Maryland, College Park. He is currently a PhD student at Northeastern University. He is a research assistant at the cognitive systems Lab at NEU. His research interests are in the area of machine learning, robotics and automatic Control.

Milica Stojanovic (Fellow, IEEE), graduated from the University of Belgrade, Serbia, in 1988, and received M.S. (’91) and Ph.D. (’93) degrees in electrical engineering from Northeastern University, Boston, Massachusetts. She was a Principal Scientist at the Massachusetts Institute of Technology, and in 2008 joined Northeastern University, where she is currently a Professor of electrical and computer engineering. She is also a Guest Investigator at the Woods Hole Oceanographic Institution. She is an Associate Editor for the IEEE Journal of Oceanic Engineering, a past Associate Editor for the IEEE Transactions on Signal Processing, IEEE Transactions on Vehicular Technology, and a Senior Editorial Board Member of the IEEE Signal Processing Magazine. She chairs the IEEE Ocean Engineering Society’s (OES) Technical Committee for Underwater Communication, Navigation and Positioning, and is an IEEE OES Distinguished Lecturer. Milica is the recipient of the 2015 IEEE/OES Distinguished Technical Achievement Award, the 2019 IEEE WICE Outstanding Achievement Award, and the 2023 IEEE Communications Society’s Women Stars in Computer Networking and Communications Award. In 2022, she was awarded an honorary doctorate from the Aarhus University in Denmark, and was elected to the Academy of Engineering Sciences of Serbia.

Pau Closas (Senior Member, IEEE), is an Associate Professor in Electrical and Computer Engineering at Northeastern University, Boston MA. He received the MS and PhD in Electrical Engineering from UPC in 2003 and 2009, respectively. He also holds a MS in Advanced Maths and Mathematical Engineering from UPC since 2014. He is the recipient of the EURASIP Best PhD Thesis Award 2014, the 9^{th} Duran Farell Award for Technology Research, the 2016 ION’s Early Achievements Award, 2019 NSF CAREER Award, and the IEEE AESS Harry Rowe Mimno Award in 2022. His primary areas of interest include statistical signal processing, stochastic filtering, robust filtering, and machine learning, with applications to positioning and localization systems. He is the EiC of the IEEE AESS Magazine, volunteered in other editorial roles (e.g. NAVIGATION, Proc. IEEE, IEEE Trans. Veh. Tech., and IEEE Sig. Process. Mag.), and has been actively involved in organizing committees of a number of conference such as EUSIPCO (2011, 2019, 2021, 2022), IEEE SSP'16, IEEE/ION PLANS (2020, 2023), or IEEE ICASSP'20.

Murat Akcakaya (Senior Member, IEEE) received his Ph.D. degree in Electrical Engineering from the Washington University in St. Louis, MO, USA, in December 2010. He is currently an Associate Professor in the Electrical and Computer Engineering Department of the University of Pittsburgh. His research interests are in the areas of statistical signal processing and machine learning.

Marius Necsoiu (Member, IEEE) received his PhD in Environmental Science (Remote Sensing) from the University of North Texas (UNT), in 2000. He has broad experience and expertise in remote sensing systems, radar data modeling, and analysis to characterize electromagnetic environment behavior and geophysical deformation. As part of the DEVCOM ARL he leads research in cognitive radars and explores new paradigms in AI/ML science that are applicable in radar/EW research.

Deniz Erdoğmuş (Sr Member, IEEE), received BS in EE and Mathematics (1997), and MS in EE (1999) from the Middle East Technical University, PhD in ECE (2002) from the University of Florida, where he was a postdoc until 2004. He was with CSEE and BME Departments at OHSU (2004-2008). Since 2008, he has been with the ECE Department at Northeastern University. His research focuses on statistical signal processing and machine learning with applications data analysis, human-cyber-physical systems, sensor fusion and intent inference for autonomy. He has served as associate editor and technical committee member for multiple IEEE societies.

Ben Wright is a Research Engineer at Kostas Research Institute. He received his Ph.D. in Computer Science from New Mexico State University in 2018. He was a Postdoctoral Researcher in the Navy Center for Applied Research in AI at the US Naval Research Laboratory from 2019-2021. His research interests involve Multi-agent Reasoning, Computational Logic, and Human-Robot Interaction.

Michael Everett received the S.B., S.M., and Ph.D. degrees in mechanical engineering from the Massachusetts Institute of Technology (MIT), Cambridge, MA, USA, in 2015, 2017, and 2020, respectively. He was a Post-Doctoral Associate and Research Scientist in the Department of Aeronautics and Astronautics at MIT. He was a Visiting Faculty Researcher at Google Research. He joined Northeastern University in 2023, where he is currently an Assistant Professor in the Department of Electrical & Computer Engineering and Khoury College of Computer Sciences at Northeastern University, Boston, MA, USA. His research lies at the intersection of machine learning, robotics, and control theory, with specific interests in the theory and application of safe and robust neural feedback loops. Dr. Everett's work has been recognized with numerous awards, including the Best Paper Award in Cognitive Robotics at IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS) 2019.

Tales Imbiriba (Member, IEEE) is an Assistant Professor in Computer Science at University of Massachusetts Boston (UMB). Before joining UMB, he served as an Assistant Research Professor at the ECE dept. and Senior Research Scientist at the Institute for Experiential AI, both at Northeastern University (NU), Boston, MA, USA. He received his Doctorate degree from the Department of Electrical Engineering (DEE) of the Federal University of Santa Catarina (UFSC), Florianópolis, Brazil, in 2016. He served as a Postdoctoral Researcher at the DEE–UFSC (2017–2019) and at the ECE dept. of the NU (2019–2021). His research interests include audio and image processing, pattern recognition, Bayesian inference, online learning, and physics-guided machine learning.

Appendix

A. Single Simulation Realization

The extended MC simulation over 1000 time steps for 6 radars and 3 targets, and 4 radars and 4 targets are shown in Appendix Figure 1 and Appendix Figure 2, respectively. GIFs for the undetermined, overdetermined, and full system are included in supplemental materials.

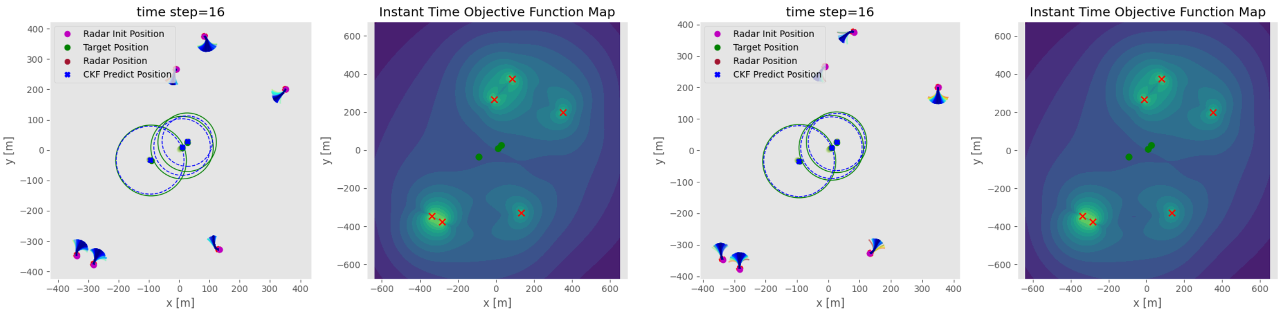

(a) Proposed approach: Each radar moves towards target

(b) [9], [13], [22] range model: Radars move to encircle targets

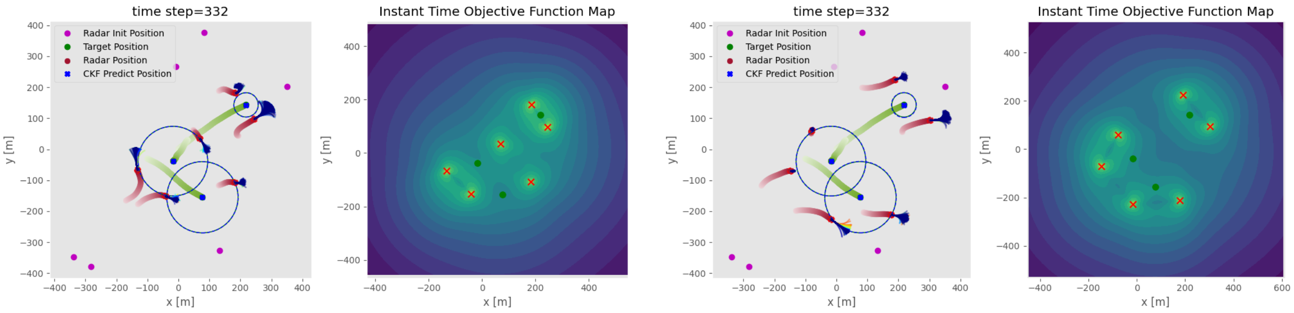

(c) Radars move toward targets as closely as feasible

(d) Radars encircle targets

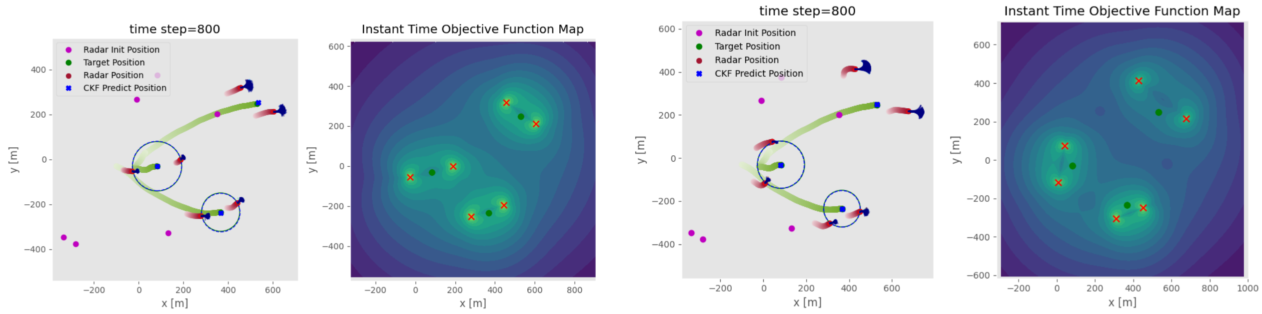

(e) Two radars follow each target individually

(f) Radars encircle the targets, but the upper most two radars move far away from the upper most target.

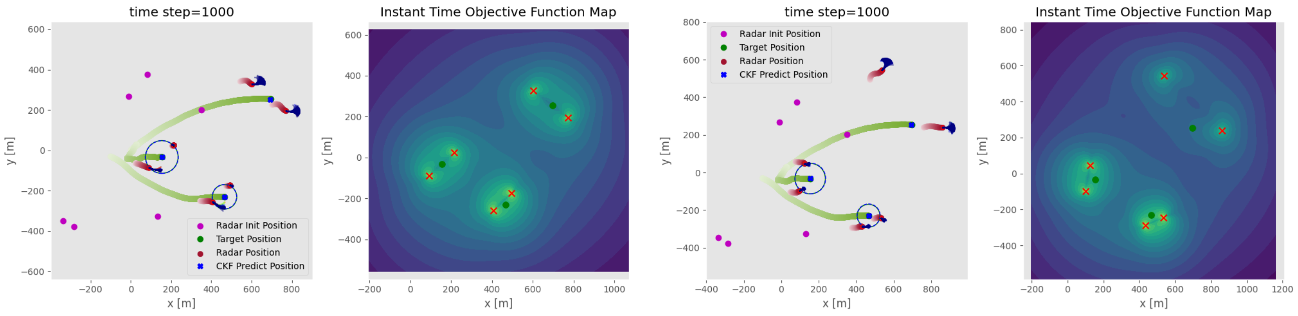

(g) Two radars track each target individually by moving side-by-side parallel to the target's trajectory.

(h) Radars maintain encirclement formation around target, but the uppermost radar is far from the uppermost target

Appendix Figure 1: 6 radars and 3 targets. The left column shows snapshots of target localization using a range measurement model with covariance dependent on radar-target distance [16], while the right column depicts tracking with a range measurement model employing constant diagonal covariance [13]. In the first three rows, each subfigure's left panel shows a 3D projection onto 2D of the radar placement at a specific time. The colored lines emanating from the radars depict the weighted MPPI-planned trajectories, with blue indicating better paths and red indicating worse ones. Dashed red and green circles mark the radar's (10 [m]) and the target's (125 [m]) collision avoidance boundaries, respectively, while blue dashed circles represent the targets' collision avoidance boundary based on the CKF position estimate. The radar's trajectory is the last 25 time steps. The right panel displays the log determinant of the FIM (for the ddr measurement model) concerning the 2D target position for the specific time step, meshed over the free space area, and evaluated at each mesh point.

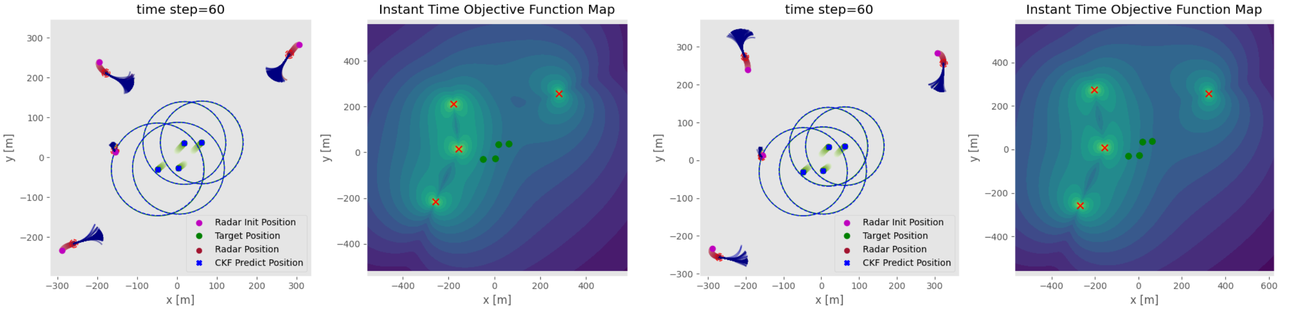

(a) Proposed approach: Pairs of radars move towards pairs of targets

(b) [9], [13], [22] range model: the upperleft radar moves away from target

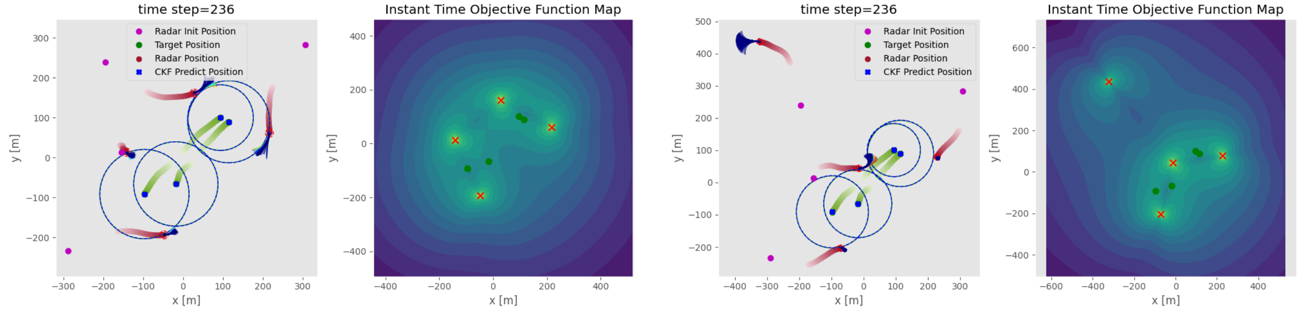

(c) Radars move toward targets as closely as feasible, with pairs of radars positioned side by side with pairs of targets.

(d) Radars encircle the targets, with the upper left radar positioned farther away from the targets.

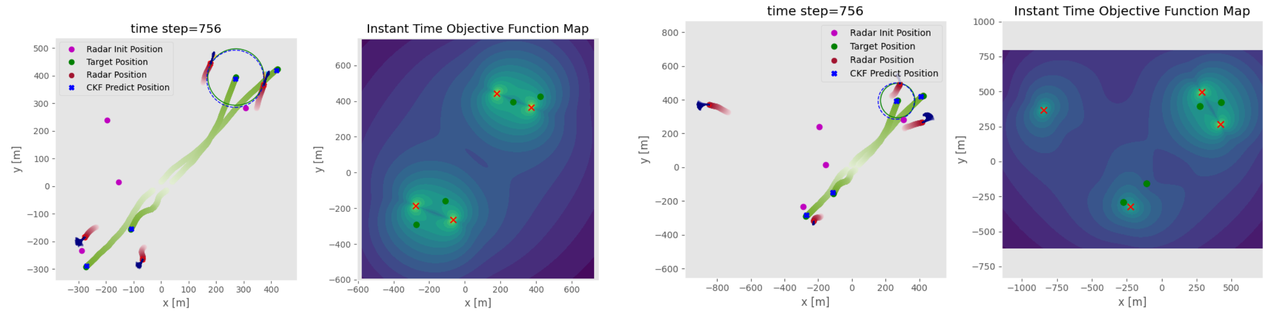

(e) The pairs of targets are kept in the maximum information gain locations of the radar pairs.

(f) The bottom left pair of targets only has one radar following, leading to worse target localization.

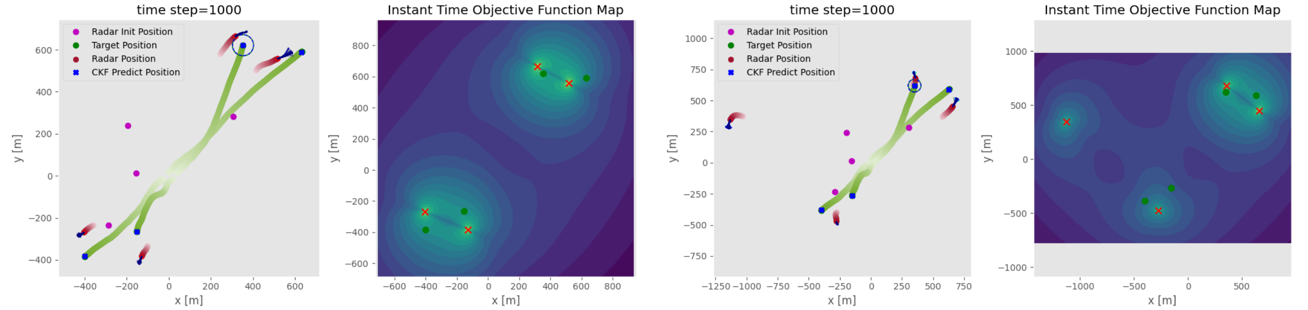

(g) Pairs of radars track each pair of targets by moving side-by-side parallel to the pair of targets' trajectory.

(h) Radars maintain encirclement formation around target, but the bottom left radar is far from the bottom left target

Appendix Figure 2: 4 radars and 4 targets. The left column shows snapshots of target localization using a range measurement model with covariance dependent on radar-target distance [16], while the right column depicts tracking with a range measurement model employing constant diagonal covariance [13]. In the first three rows, each subfigure's left panel shows a 3D projection onto 2D of the radar placement at a specific time. The colored lines emanating from the radars depict the weighted MPPI-planned trajectories, with blue indicating better paths and red indicating worse ones. Dashed red and green circles mark the radar's (10 [m]) and the target's (125 [m]) collision avoidance boundaries, respectively, while blue dashed circles represent the targets' collision avoidance boundary based on the CKF position estimate. The radar's trajectory is the last 25 time steps. The right panel displays the log determinant of the FIM (for the ddr measurement model) concerning the 2D target position for the specific time step, meshed over the free space area, and evaluated at each mesh point.