# Distributed Multi-Robot Ergodic Coverage Control for Estimating Time-Varying Spatial Processes

Mattia Mantovani, Mattia Catellani and Lorenzo Sabattini

***Abstract*— This paper presents a fully distributed approach for multi-robot systems to explore and cover time-varying spatial processes in complex, non-convex environments. Building on the heat-equation-driven adaptive coverage (HEDAC) framework, our method enables autonomous navigation, process reconstruction, and dynamic balancing of exploration and coverage in a collision-free manner. A temporal decay component regulates the trade-off between revisiting known regions and exploring new ones, ensuring adaptive and efficient monitoring. Simulations and real-world experiments validate the approach's effectiveness and robustness.**

## I. Introduction

Coverage control in multi-robot systems allocates sensing effort according to a spatial density. Classical Voronoi-based methods [1], [2] optimize static coverage, while ergodic coverage aligns a robot's time-averaged trajectory with a distribution of interest. The Spectral Multiscale Coverage (SMC) framework [3], later extended with time-optimal [4], [5] and receding-horizon [6] methods, assumes convexity, lacks collision avoidance, and emphasizes global coverage. HEDAC [7] addresses some of these issues via a stationary heat equation embedding.

Prior work assumes full knowledge of static processes, whereas real applications require online monitoring of unknown, time-varying phenomena. Gaussian Processes (GPs) are standard for modeling [8]–[10] but face computational and scalability limits [11], [12]. No existing method fully combines distributed ergodic control, time-varying learning, and collision-aware motion; [13] addresses decentralized ergodic control but not dynamic process estimation.

We present a fully distributed multi-robot framework that leverages HEDAC and time-varying GPs to collaboratively explore and track dynamic environments, adapting to changes in real time.

## II. Distributed Ergodic Coverage Control

### A. Distributed Gaussian Process

In this work, we adopt time-varying GPs to enable robots to explore, reconstruct, and persistently monitor dynamic spatial phenomena. Each robot computes its own $\text{GP}_i$ online, collaborating with neighbors via dataset sharing. Following [12], we estimate the underlying process by applying spatio-temporal decay, so that older or spatially distant samples contribute less to predictions defined as $(\phi^*|\mathfrak{D}_t) \sim \mathcal{N}(\mu(\mathbf{X}^*), \Sigma(\mathbf{X}^*))$, where

$$
\begin{aligned}
\mu(\mathbf{X}^*) &= \tilde{\kappa}(\mathbf{X}^*, \mathbf{X}_t)^\intercal [\tilde{\kappa}(\mathbf{X}_t, \mathbf{X}_t) + \sigma_n^2 \mathbf{I}]^{-1} \mathbf{y}_t, \\
\Sigma(\mathbf{X}^*) &= \kappa(\mathbf{X}^*, \mathbf{X}^*) - \tilde{\kappa}(\mathbf{X}^*, \mathbf{X}_t)^\intercal \\
&\quad \cdot [\tilde{\kappa}(\mathbf{X}_t, \mathbf{X}_t) + \sigma_n^2 \mathbf{I}]^{-1} \tilde{\kappa}(\mathbf{X}^*, \mathbf{X}_t).
\end{aligned} \tag{1}
$$

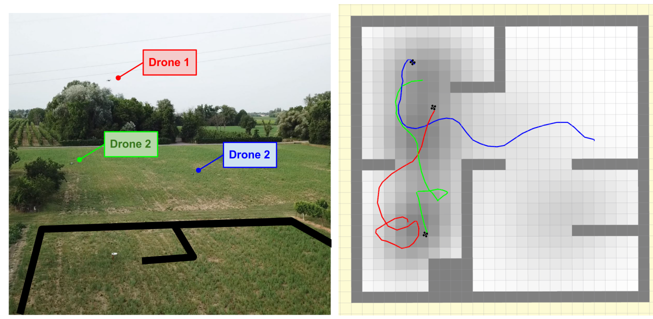


Fig. 1. Snapshot of a real-world experiment with three Uvify IFO-S drones running the algorithm onboard. Drones collaborate via minimal communication. RViz visualization shows the environment, trajectories, and spatial process estimation.

The decayed covariance matrices are

$$
\begin{aligned}
\tilde{\kappa}(\mathbf{X}_t, \mathbf{X}_t) &= \kappa(\mathbf{X}_t, \mathbf{X}_t) \odot \mathbf{D}_t \odot \mathbf{T}^D, \\
\tilde{\kappa}(\mathbf{X}^*, \mathbf{X}_t) &= \kappa(\mathbf{X}^*, \mathbf{X}_t) \odot \mathbf{d}_t^\intercal \odot \mathbf{T}^d,
\end{aligned} \tag{2}
$$

with $\odot$ denoting element-wise multiplication. The temporal decay components $\mathbf{T}^D$ and $\mathbf{T}^d$ are defined following [12]. Unlike [12], which bases spatial decay on sample index, we define it using true Euclidean distance to avoid treating distant samples as nearby

$$
\begin{aligned}
[\mathbf{D}_t]_{ij} &= \exp(-\|\mathbf{x}_i - \mathbf{x}_j\|_2^2 / \lambda_s), \\
[\mathbf{d}_t]_i &= \exp(-\|\mathbf{x}_i - \mathbf{z}_i\|_2^2 / \lambda_s),
\end{aligned} \tag{3}
$$

where $\mathbf{x}_i, \mathbf{x}_j \in \mathbf{X}_t$ are data samples, $\mathbf{z}_i$ is the current robot position, and $\lambda_s$ sets the spatial decay rate.

To mitigate the $\mathcal{O}(N^3)$ GP complexity, we adopt threshold-based filtering similar to [11], [12], retaining samples that reduce epistemic uncertainty and discarding outdated ones. Specifically, a sample $\mathbf{x}$ is included if $\sigma_{\text{in}}(\mathbf{x}) \geq e_{\text{in}}/z$ and removed if $\sigma_{\text{out}}(\mathbf{x}) \geq e_{\text{out}}/z$, where $e_{\text{in}}$ and $e_{\text{out}}$ are inclusion and removal tolerances, and $z$ sets the confidence level.

### B. Dynamic Goal Density

We define a dynamic goal density that fuses each robot's GP mean $\mu_i(\mathbf{x})$ and uncertainty $\sigma_{u,i}(\mathbf{x}) = \sqrt{\Sigma_i(\mathbf{x})}$ into a unified spatial distribution

$$
\tilde{\mathrm{F}}_i(\mathbf{x}) = (\exp(\mu_i(\mathbf{x})) - 1) + (\exp(\sigma_{u,i}(\mathbf{x})) - 1) \tag{4}
$$

with the normalized unified spatial distribution defined as $\mathrm{F}_i(\mathbf{x}) = \tilde{\mathrm{F}}_i(\mathbf{x}) / \int_\Omega \tilde{\mathrm{F}}_i(\mathbf{x}) d\mathbf{x}$. This distribution prioritizes exploration in regions of high uncertainty and monitoring in high-mean areas, with both contributions normalized. The robot dynamically updates its local goal density using a smoothed inner product:

$$g_i(\mathbf{x}) = \langle \varphi_\gamma, \mathrm{F}_i \rangle_\Omega = \int_\Omega \varphi_\gamma(\mathbf{x}) \mathrm{F}_i(\mathbf{x}) d\mathbf{x}. \tag{5}$$

where $\varphi_\gamma$ is an RBF smoothing kernel [7]. Updates occur continuously as robots explore or share data, enabling adaptive coverage of time-varying spatial processes.

### C. Distributed Ergodic Control

To enable scalable, distributed collaboration, robots exchange only local coverage information within a communication range $r$, avoiding unnecessary global sharing. Each robot maintains a local coverage density

$$c_i(\mathbf{x}, t) = \frac{1}{t} \int_0^t p_i(\mathbf{r}_i^\tau) d\tau, \quad p_i(\mathbf{r}_i^t) = \exp(-(\epsilon \mathbf{r}_i^t)^2), \tag{6}$$

where $\mathbf{r}_i^t$ is the sensor-relative position, $\epsilon > 0$ shapes the RBF footprint, and time-dependence captures dynamic sensing. Neighbor contributions are aggregated via

$$\tilde{c}_i(\mathbf{x}, t) = c_i(\mathbf{x}, t) + \sum_{j \in N_{i,r}} w_r(\mathbf{x}, \mathbf{z}_i(t)) c_j(\mathbf{x}, t), \tag{7}$$

with $w_r$ enforcing locality and $N_{i,r}$ denoting neighbors within range $r$. The resulting distributed error field is

$$\tilde{e}_i(\mathbf{x}, t) = g_i(\mathbf{x}) - \tilde{c}_i(\mathbf{x}, t), \quad \mathcal{E}_i = \|\tilde{e}_i(\mathbf{x}, t)\|_2, \tag{8}$$

where $g_i(\mathbf{x})$ is the dynamic goal density (5). Each robot defines a local source for the heat equation as

$$s_i(\mathbf{x}, t) = \begin{cases} \tilde{e}_i(\mathbf{x}, t)^2 & \tilde{e}_i(\mathbf{x}, t) \geq 0, \\ 0 & \text{otherwise}, \end{cases} \tag{9}$$

and solves the distributed heat equation

$$\begin{aligned} \rho \Delta u(\mathbf{x}, t) &= \beta u(\mathbf{x}, t) - s_i(\mathbf{x}, t), \\ \frac{\partial u}{\partial \mathbf{n}} &= 0 \quad \text{on } \delta\Omega, \end{aligned} \tag{10}$$

where $\rho$ controls spatial diffusion, $\beta$ ensures scale-invariant decay, and $\mathbf{n}$ is the outward normal. Collision avoidance [14] maintains a minimum safety distance $r_s$, enabling safe, distributed ergodic coverage.

## III. Experimental Validation

### A. Time-varying Process

We evaluate our distributed strategy in a scenario where the spatial process abruptly changes. In Fig. 2, four robots explore, learn, and monitor using local information and dynamic goal densities. After the shift at step 5000, our method quickly adapts, redirecting coverage and learning the updated process, while centralized HEDAC remains focused on outdated regions. Fig. 3 illustrates one robot's trajectory and ergodic metric before and after the change, showing effective adaptation from the initial to the new key area.

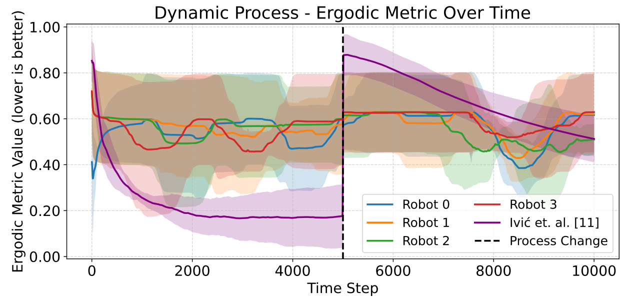


Fig. 2. Ergodic metric over time for four robots monitoring a dynamic spatial process using our distributed strategy (colored lines) versus centralized HEDAC [7] (purple). At step 5000 (dashed), the process abruptly shifts. Our method adapts via time-varying GPs and dynamic goal densities, while the centralized baseline shows lower initial metrics and slow responsiveness.

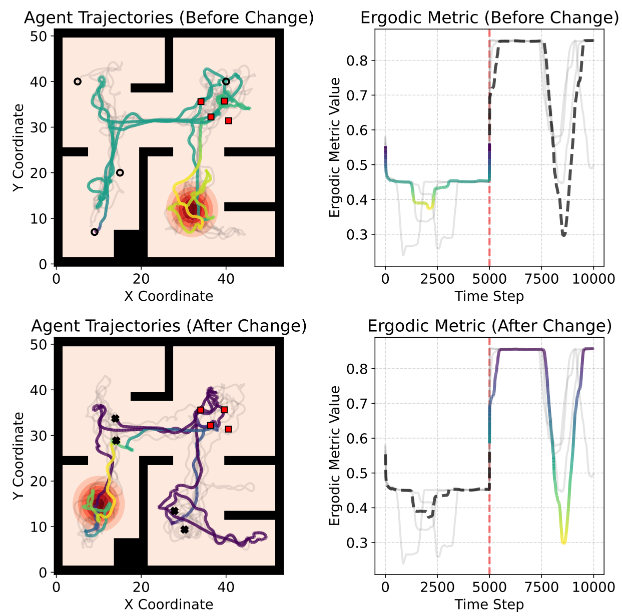


Fig. 3. Simulation of one robot before and after a sudden spatial process change at step 5000. Trajectory is color-coded by ergodic metric (brighter = focused coverage). Black dots: initial positions; red squares: positions at the change; black crosses: final positions. Other robots shown in gray.

### B. Real-world experiments

We tested our distributed strategy with three Uvify IFO-S drones, each running the algorithm onboard on a Jetson Nano alongside a PX4 controller. Robots share minimal information, position and dataset samples, via the base station with neighbors only. Fig. 1 shows a snapshot and RViz visualization.

## IV. Conclusion

This work presents a fully distributed multi-robot system for adaptive monitoring of time-varying spatial processes. Each robot leverages its own GP, dynamic goal and coverage strategies, and the HEDAC ergodic potential-based coordination mechanism to explore efficiently and safely. Simulations and real-world experiments demonstrate the approach's robustness and effectiveness.

# References


[1] J. Cortes, S. Martinez, and F. Bullo, "Spatially-distributed coverage optimization and control with limited-range interactions," *ESAIM: Control, Optimisation and Calculus of Variations*, vol. 11, no. 4, pp. 691–719, 2005.

[2] F. Pratissoli, B. Capelli, and L. Sabattini, "On coverage control for limited range multi-robot systems," in *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2022, pp. 9957–9963.

[3] G. Mathew and I. Mezić, "Metrics for ergodicity and design of ergodic dynamics for multi-agent systems," *Physica D: Nonlinear Phenomena*, vol. 240, no. 4-5, pp. 432–442, 2011.

[4] L. M. Miller and T. D. Murphey, "Trajectory optimization for continuous ergodic exploration," in *2013 American Control Conference*. IEEE, 2013, pp. 4196–4201.

[5] M. Sun, A. Gaggar, P. Trautman, and T. Murphey, "Fast ergodic search with kernel functions," 2024. [Online]. Available: https://arxiv.org/abs/2403.01536

[6] D. Dong, H. Berger, and I. Abraham, "Time optimal ergodic search," *arXiv preprint arXiv:2305.11643*, 2023.

[7] S. Ivić, B. Crnković, and I. Mezić, "Ergodicity-based cooperative multiagent area coverage via a potential field," *IEEE Transactions on Cybernetics*, vol. 47, no. 8, pp. 1983–1993, 2017.

[8] M. Santos, U. Madhushani, A. Benevento, and N. E. Leonard, "Multi-robot learning and coverage of unknown spatial fields," in *2021 International Symposium on Multi-Robot and Multi-Agent Systems (MRS)*. IEEE, 2021, pp. 137–145.

[9] K.-C. Ma, L. Liu, and G. S. Sukhatme, "Informative planning and online learning with sparse gaussian processes," in *2017 IEEE International Conference on Robotics and Automation (ICRA)*, 2017, pp. 4292–4298.

[10] L. Wei, A. McDonald, and V. Srivastava, "Multi-robot gaussian process estimation and coverage: Deterministic sequencing algorithm and regret analysis," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*, 2021, pp. 9080–9085.

[11] M. Mantovani, F. Pratissoli, and L. Sabattini, "Distributed coverage control for spatial processes estimation with noisy observations," *IEEE Robotics and Automation Letters*, vol. 9, no. 5, pp. 4431–4438, 2024.

[12] F. Pratissoli, M. Mantovani, A. Prorok, and L. Sabattini, "Distributed coverage control for time-varying spatial processes," *IEEE Transactions on Robotics*, vol. 41, pp. 1602–1617, 2025.

[13] I. Abraham and T. D. Murphey, "Decentralized ergodic control: distribution-driven sensing and exploration for multiagent systems," *IEEE Robotics and Automation Letters*, vol. 3, no. 4, pp. 2987–2994, 2018.

[14] I. I. Hussein and D. M. Stipanovic, "Effective coverage control for mobile sensor networks with guaranteed collision avoidance," *IEEE Transactions on Control Systems Technology*, vol. 15, no. 4, pp. 642–657, 2007.