

Resilient Decentralized Ergodic Coverage for Scalable Multi-Robot Systems in Unknown Time-Varying Environments

Maria G. Mendoza¹ Victoria Tuck² Chinmay Maheshwari³ Shankar Sastry¹
University of California-Berkeley¹, University of Pennsylvania², Johns Hopkins University³



MOTIVATION

In disaster response, mission success depends on **real-time situational awareness**.


Drone-as-First-Responder programs

- provide fast response times,
- put fewer humans in dangerous situations, and
- offer extended capabilities such as aerial viewing.

However, **deployment remains largely human-operated 1-to-1**, diverting critical personnel from direct response.

We design an **adaptive data collection coverage protocol** that

1. operates multiple drones decentralized,
2. allows the team to track evolving priorities,
3. scales to 200+ robots and remains resilient under degraded conditions: intermittent or delayed communication, limited onboard memory, and robot drop-out.



PROBLEM

Prior work lacks *persistence* in long run decentralized sensing in unknown, dynamic environments. This motivates the central question of our work:

Can a team of robots maintain persistent long-term coverage of the most critical regions in unknown, time-varying environments?


Our goal is captured as **ergodic coverage**. **Ergodic methods** shape long-run visitation so that the time-averaged spatial statistics of a trajectory converge to a target distribution.

We desire that time spent in each region captured by the empirical distribution

$$\hat{\rho}_K = \frac{1}{M} \sum_{m \in \mathcal{M}} \frac{1}{K+1} \sum_{k=0}^K 1(x_k^m)$$

tracks the time-varying, true map distribution ρ_k^* . We wish to achieve **no-regret** asymptotically:

$$\text{Regret}(K) = \frac{1}{K} \sum_{k=1}^K \mathbb{E}[\|\hat{\rho}_k - \rho_k^*\|_1]$$



ERGODIC COVERAGE FOR TIME-VARYING ENVIRONMENTS





Fig. The environment is abstracted into a grid representation. Robots travel between cells using a lower-level planner and control.

Each robot m follows locally from the first epoch $e = 0$:

1. **Observe**. Noisy readings from its own sensor and its neighbors'.
2. **Learn**. If $\text{mod}(K, \tau_p) = 0$. A map weighting posterior $\tilde{\phi}_e^m | D_e^m \sim GP(\mu_e^m, \kappa_e^m)$ of a Gaussian Process (GP) on the dataset D_e^m . Determine $\text{UCB} \phi_{\text{UCB}}^m = \mu_e^m + \beta \sigma_e^m$. Set belief distribution $\hat{\rho}_e^m = \frac{\phi_{\text{UCB}}^m}{\sum_{i \in \mathcal{M}} \phi_{\text{UCB}}^i}$. Use Metropolis-Hastings to determine transition P_e^m from $\hat{\rho}_e^m$.
3. **Move**. Sample the next region using P_e^m



REGRET ANALYSIS

The system is not overly sensitive to precise tuning of τ_p (frequency of K). Regret is reduced over time with environment changes.




Fig. Performance under different update τ_p . The environment changes at $\mathcal{K} = \{500, 1000, 2000\}$.

EMERGENT DISTRIBUTED ERGODICITY AT SCALE

200 robots are tasked to explore a large environment and on average, proportionally divide time amongst regions according to their importance/interest level.

Parameters:
200 robots
Large environment (50x50 grid)
40,000 steps of simulation
4,000-observation memory cap
Communication reaching only ~3% of the grid

Key Result: The team empirical visitation matches the desired true map (top) **without any agent needing to correctly map the entire environment (bottom)**

Ergodic coverage is an emergent team property!




Fig. The average visit time over the fleet, the "empirical distribution" (right), closely mimics the desired true distribution (left) over regions.





Fig. Individual agents only correctly map a **portion** of the environment (red boxes). However, collectively they achieve ergodicity!


ROBUSTNESS TO SYSTEM FAILURES

Robustness to Belief Staleness



Constrained communication does not impact time to find critical regions. For full map coverage, local communication is better because beliefs remain heterogeneous, delaying premature agreement².

Robustness to Comm. Delays & Robot Loss



Communication delays do not impact ROI discovery;
2/3 of the team can be lost and ROIs are still found in under 500 steps.


FASTER REMAPPING AND DISCOVERY VS GREEDY APPROACH

Our approach shows

- 100% ROI discovery,
- faster remapping after every change,
- occupancy proportional to importance

The greedy baseline (MAC-DT)

- misses 9% of ROIs,
- skew towards top-reward regions



Additional data in paper Table I & II.

Multi-Agent Coverage with Doubling Trick (MAC-DT)³ greedy-based centralized planner fails to min regret.

Ongoing Work

LLM-based persistent orchestration of heterogeneous agents for semantic missions with delegation to ergodic coverage in time-varying environments to provide persistence and improve scalability in time and fleet size (Under Review, ICRA 2027).

Theoretical no-regret guarantees with high probability for synchronous and asynchronous multi-agent ergodic coverage (In Preparation).

References

1. D. A. Levin and Y. Peres, Markov Chains and Mixing Times, 2nd ed. American Mathematical Society, 2017.
2. M. S. Talamali, A. Saha, J. A. R. Marshall, and A. Reina, "When less is more: Robot swarms adapt better to changes with constrained communication," Science Robotics, 2021.
3. R. Zhang, H. Ma, and N. Li, "Multi-agent coverage control with transient behavior consideration," in Proceedings of the 6th Annual Learning for Dynamics and Control Conference, ser. Proceedings of Machine Learning Research. PMLR, 2024.