

Strategic Sacrifice: Self-Organized Robot Swarm Localization for Inspection Productivity

Sneha Ramshanker, Hungtang Ko, and Radhika Nagpal

Princeton University, Princeton NJ 08544, USA,
s.ramshanker@princeton.edu, hk1581@princeton.edu, rn1627@princeton.edu

Abstract. Robot swarms offer significant potential for inspecting diverse infrastructure, ranging from bridges to space stations. However, effective inspection requires accurate robot localization, which demands substantial computational resources and limits productivity. Inspired by biological systems, we introduce a novel cooperative localization mechanism that minimizes collective computation expenditure through self-organized sacrifice. Here, a few agents bear the computational burden of localization; through local interactions, they improve the inspection productivity of the swarm. Our approach adaptively maximizes inspection productivity for unconstrained trajectories in dynamic interaction and environmental settings. We demonstrate the optimality and robustness using mean-field analytical models, multi-agent simulations, and hardware experiments with metal climbing robots inspecting a 3D cylinder.

Keywords: Collaborative localization, mean-field models, task-allocation

1 Introduction

Imagine a future where small robotic teams roam our infrastructure—bridges, pipelines, buildings, and satellites—detecting problems promptly, such as leaks and cracks (Figure 1(b)) [1–4]. Teams of robots offer many advantages for inspection, including high parallelization, resilience to failure, and potentially low unit cost. However, for effective inspection, robots need to know where they are, a problem known as localization. In most indoor and remote locations, external localization mechanisms like GPS are unavailable or unreliable [5]. Robots must autonomously perform localization, overcoming difficulties such as sensor noise, limited landmarks, and slippage. While the literature primarily addresses the localization problem from a single-robot perspective, with approaches such as SLAM achieving good accuracy [6], these methods are computationally intensive. Computation is a finite and valuable resource; if most of an agent’s computation is used for localization, little remains for productive inspection tasks like finding cracks, measuring vibration, and monitoring rust accumulation.

We explore if collaboration among robots can simplify this localization problem, allocating more computation for inspection. We draw inspiration from nature, where individuals sacrifice for the group’s benefit. In anti-predator vigilance, some members watch for predators, allowing others to eat safely [7, 8].




Fig. 1. Rovables [14] (a) on a 3D metal cylinder, (b) on a piping system, and (c) in the palm of a human hand. (d) Rovables with markers for Vicon Motion Capture.

Army ants form bridges out of their bodies, facilitating cargo transport across gaps and cracks [9, 10]. The vigilant animals forego eating, and the bridge ants do not carry cargo; instead, by sacrificing individual productivity, they enhance group success. What is especially interesting about these biological cases is that the fraction of sacrificers is self-organized and adapts to environmental demands.

We introduce a novel cooperative localization mechanism for robot swarms that leverages this idea of self-organized sacrifice for the group’s benefit. While previous studies have used cooperation to enhance localization accuracy and enable localization in unknown environments [11–13], these methods often increase computational requirements, constrain robot trajectories, and lack adaptability to changing environments. In our approach, individuals become dedicated localizers or inspectors, with this distribution self-organized based on local interactions. We demonstrate that this decentralized mechanism optimizes collective productivity in dynamic conditions, validated through theoretical models, numerical simulations, and hardware experiments. By deriving mean-field models (Sec 4), we prove that sacrificing agents as dedicated localizers improves swarm productivity. We also show that the swarm can reconfigure based on local interactions to always maximize productivity. Using agent-based numerical simulations (Sec 5), we demonstrate that group productivity increases further through smarter collaboration. Hardware experiments (Sec 6) conducted with a swarm of 10 metal climbing robots inspecting a 3D metal cylinder (Figure 1) demonstrate the effectiveness of this approach in navigating complex physical environments. Moreover, these experiments show that the emergent behavior optimizes productivity amidst dynamically changing interactions.

2 Related Work

Single-robot localization is extensively studied in robotics literature [6, 15]. The simplest approach is dead reckoning, where the robot uses internal sensor data like wheel encoders and IMU to estimate its position. While computationally inexpensive, dead reckoning suffers from error accumulation and requires regular external feedback for accuracy [6]. Probabilistic methods like particle filters and SLAM offer higher accuracy by incorporating external landmark and feature detection but demand heavy computational resources [15]. In this paper, we will not explore single-agent localization methods in detail. Instead, we will assume robots can achieve accurate localization at a high computational cost. An analysis of this accuracy-computation trade-off can be found in [16].

In multi-robot settings, various approaches improve localization using collaboration [17]. A common strategy to boost accuracy is *sharing personal estimates* and relative measurements [11, 18–20]. For instance, in [11], each agent employs a particle filter (PF) combining personal dead reckoning with relative position estimates from neighboring agents. While these methods improve localization accuracy, they also come with an increase in computational costs, as all robots must process not only their individual measurements but also the estimates of their neighbors. Other collaborative localization (CL) approaches enhance accuracy by *constraining swarm motion* [12, 13, 21–23]. For example, in [4], robots disperse to form a landmark tree for navigation and in [12], robots move in coordinated steps, with subgroups acting as stationary landmarks. However, the constrained motion can conflict with the motion required for inspection, increasing the complexity of inspection algorithms. By contrast, our approach is compatible with any unconstrained trajectory, offering the flexibility to execute any inspection algorithm. *Heterogeneous CL* methods involve some robots performing more computation and sensing to localize others in the group [13, 24–26]. For example, in [13], observer robots with cameras localize beacon robots, sacrificing their computation for the group. However, in such heterogeneous approaches, the fraction of localizing agents is fixed a priori. Thus, the robot swarm cannot reconfigure itself if the environment or the needs of the group change. In our method, the fraction of sacrificing agents constantly adapts to maximize efficiency.

Our approach is also related to self-organized task allocation in biological and robot swarms [27–29], building upon mean-field techniques commonly used in this area [30]. It maps to a task allocation problem where one task is computationally intensive localization, and the other is inspection. Unlike most studies, these tasks are tightly coupled in our approach, allowing the swarm to collectively achieve more than the sum of individual efforts via optimal task distribution.

3 Problem Formulation

In our **collective localization** mechanism, agents operate in one of two modes:

- **Dead Reckoner (DR) Mode:** Dead reckoners integrate IMU data to determine position and orientation. While computationally inexpensive, this




Fig. 2. *Top:* Summary of state properties. *Bottom:* Transitions between $DR_{NotLost}$, DR_{Lost} , PL^\dagger , PL for (a) fixed mode, (b) individual mode-switching, (c) collaborative mode-switching. Dashed arrows indicate instantaneous transitions while solid arrows indicate transitions at a finite rate.

approach suffers from error accumulation over time. Dead reckoners are productive when not lost but are increasingly prone to getting lost.

- **Perfect Localizer (PL) Mode:** Perfect localizers allocate all computational power to achieve perfect localization, sacrificing all inspection productivity. This assumption aligns with previous SLAM implementations on resource-constrained systems [16]. We introduce the PL^\dagger state as a transitional phase, representing the startup cost for achieving perfect localization. In this state, localization is incomplete, and the agents cannot assist others. While undesirable, this is a necessary state for achieving perfect localization.

We represent the four states using the notation: $DR_{NotLost}$, DR_{Lost} , PL^\dagger , PL . A dead reckoner’s lost status is characterized using a disorientation function, $\gamma = 1 - \exp\left(-\frac{|\delta\mathbf{p}|}{\delta p_0}\right)$. Here, $\delta\mathbf{p}$ quantifies the localization error and δp_0 represents the characteristic localization error. Disorientation asymptotically approaches 1 as $|\delta\mathbf{p}|$ increases. We set a disorientation threshold, γ_{thresh} , to determine if an agent is lost: if $\gamma > \gamma_{thresh}$, the agent is considered lost. Designers can adjust the system’s sensitivity to localization errors by changing $(\delta p_0, \gamma_{thresh})$. We assume agents can accurately estimate their disorientation, a strong assumption we aim to relax in the future. We use a step function to map dead reckoner disorientation to productivity, Π . If $\gamma < \gamma_{thresh}$, $\Pi = 1$; otherwise, $\Pi = 0$. For perfect localizers, $\Pi = 0$ and $\gamma = 0$. The table in Figure 2 summarizes these states.

While perfect localizers are not directly productive, they play a crucial role in enhancing overall productivity by correcting the localization errors of dead reckoners through pairwise interactions. During such an interaction, the agents measure the relative position \mathbf{p}_{rel} between them. The perfect localizer then shares its position, \mathbf{p}_{PL} , enabling the dead reckoner to calculate its true position as $\mathbf{p}_{DR} = \mathbf{p}_{PL} + \mathbf{p}_{rel}$. For now, we assume no uncertainty in \mathbf{p}_{rel} or \mathbf{p}_{PL} , causing the dead reckoner’s localization error to collapse to zero after the interaction.

In the next section, we analyze these dynamics using analytical models, demonstrating the benefits of strategic sacrifice.

4 Analytical Model using Mean Field

To model swarm dynamics and the effect of collaboration on productivity, we use mean-field ordinary differential equations (ODEs). Mean-field models simplify system dynamics by focusing on the evolution of mean values. This approach has been applied to various physical, social, and robotic systems [30, 31].

The swarm's occupancy configuration n denotes the number of agents in each state: $n = \{n_{DR}^{NotLost}, n_{DR}^{Lost}, n_{PL^\dagger}, n_{PL}\}$, where $n_{DR}^{NotLost} + n_{DR}^{Lost} + n_{PL^\dagger} + n_{PL} = N$ (total agents). We express the change in mean occupancy over time as $\frac{d\langle n \rangle}{dt}$ using ODEs. By solving for the steady-state configuration $\langle n \rangle_{t \rightarrow \infty}$ where $\frac{d\langle n \rangle_{t \rightarrow \infty}}{dt} = 0$, we estimate mean productivity per agent using eq 1. The form of eq 1 reflects that only non-lost dead reckoners are productive. We omit the $t \rightarrow \infty$ notation for clarity. $\frac{\langle \Pi \rangle}{N}$ is normalized by experiment duration; e.g., if an agent is productive for half the experiment, $\frac{\langle \Pi \rangle}{N} = 0.5$.

$$\frac{\langle \Pi \rangle}{N} = \frac{\langle n_{DR}^{NotLost} \rangle}{N} \quad (1)$$

We next find the steady-state dynamics for (1) Fixed Modes (2) Individual Mode-Switching, and (3) Collaborative Mode-Switching (Figure 2 a-c).

4.1 Fixed Modes, with Collaboration

We first assume a designer fixes the number of perfect localizers a priori (N_{PL}). The remaining agents are dead reckoners (N_{DR}). As shown in Figure 2(a), dead reckoners transition to the lost state at a fixed rate r_L , which accounts for factors such as environmental slippage, observable landmarks, and sensor noise. Lost dead reckoners can return to the not-lost state via interactions with perfect localizers. At this stage, we assume a well-mixed system where every agent has an equal probability of interacting with any other agent in the swarm; the interaction rate between any two arbitrary agents is denoted as r_{int} . We want to determine the optimal number of perfect localizers that maximize productivity. The system's occupancy configuration is expressed by eq 2. The right-hand side (RHS) terms are the rates of dead reckoners becoming lost and re-localizing via interactions, respectively. $n_{PL} = N_{PL}$, $n_{DR}^{NotLost} + n_{DR}^{Lost} = N_{DR}$, and $N_{DR} + N_{PL} = N$. Solving eq 2 for the steady-state gives eq 3. By computing the jacobian of eq 2, we can show that the steady-state is a stable equilibrium.

$$\frac{d\langle n_{DR}^{NotLost} \rangle}{dt} = -r_L \langle n_{DR}^{NotLost} \rangle + \frac{2r_{int}}{N(N-1)} \langle n_{PL} \rangle \langle n_{DR}^{Lost} \rangle \quad (2)$$

$$\frac{\langle \Pi \rangle}{N} = \frac{1 - \left(\frac{N_{PL}}{N}\right)}{1 + \frac{r_L(N-1)}{2r_{int}} \left(\frac{N}{N_{PL}}\right)} \quad (3)$$

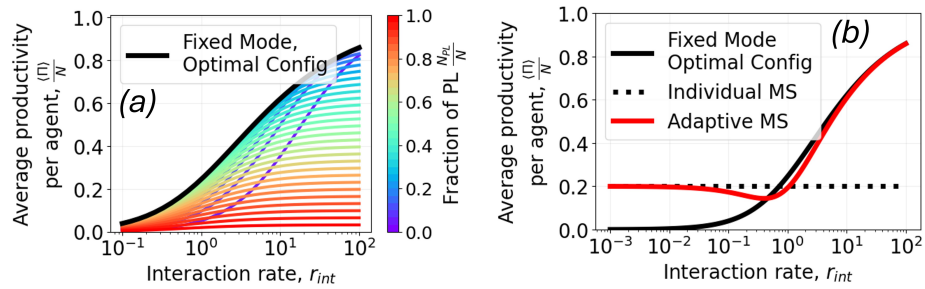


Fig. 3. Mean-field predictions for (a) Fixed Mode (eq 3), and (b) Mode-Switching. Solid black line is the fixed-mode productivity if initialized with optimal fraction of perfect localizers (eq. 4). Dotted line is the mode-switching productivity without collaboration (eq 6). Red curve is the adaptive mode-switching strategy, where $r_{MS} = \frac{0.01}{r_{int}}$ (eq 8).

Differentiating eq 3 by the fraction of perfect localizers, $\frac{N_{PL}}{N}$, we find the $\frac{N_{PL}}{N}$ that maximizes the mean productivity per agent is eq 4.

$$\left(\frac{N_{PL}}{N}\right)^* = \frac{r_L(N-1)}{2r_{int}} \left(-1 + \sqrt{1 + \frac{2r_{int}}{(N-1)r_L}}\right) \quad (4)$$

Figure 3(a) shows $\frac{\langle \Pi \rangle}{N}$ against the interaction rate, r_{int} , for different N_{PL} , where $N = 30$ agents and $r_L = 0.04$. Increasing r_{int} enhances productivity per agent across all $\frac{N_{PL}}{N}$. Moreover, the optimal fraction of perfect localizers shifts with r_{int} : more perfect localizers are needed for maximum productivity at lower r_{int} and fewer are required at higher r_{int} .

4.2 Individual Mode-Switching, with No Collaboration

With fixed modes, average productivity per agent hinges on the initial mode configuration, i.e. the ratio between perfect localizers and dead reckoners (Figure 3(a)). Sub-optimal initial configurations can significantly reduce productivity. If the mode configurations are flexible, can swarms reconfigure to more optimal configurations by switching their modes?

We assume agents have the flexibility to switch modes independently but cannot interact with each other. Each agent optimizes productivity through periodic mode-switching: starting as a dead reckoner and transitioning to perfect localizer mode upon becoming lost. However, this transition incurs a cost as the agent must spend computation re-localizing in the PL^\dagger state. This cost is captured by a fixed time-penalty τ_p , which varies based on factors such as environmental stability and the number of observable features and landmarks. Once localized, the agent reverts to dead reckoning and resumes productivity until it gets lost again. Based on Figure 2(b), the mean field is expressed by eq 5.

$$\frac{d\langle n_{DR}^{\text{Not Lost}} \rangle}{dt} = r_p \langle n_{PL^\dagger} \rangle - r_L \langle n_{DR}^{\text{Not Lost}} \rangle \quad (5)$$

The RHS terms are the rates of mode-switching of perfect localizers to dead reckoners and vice versa, assuming instantaneous transitions. We approximate r_p as $\frac{1}{\tau_p}$. Using the conservation of agent number, eq 6 gives the mean productivity per agent at stable equilibrium.

$$\left\langle \frac{\Pi}{N} \right\rangle = \frac{r_p}{r_p + r_L} \quad (6)$$

Mode-switching allows lost agents to re-localize, regaining productivity for the collective. However, in this individualistic setup, productivity remains constant and independent of the interaction rate (Eq. 6 and Figure 3(b)). The system thus overlooks potential productivity gains from interactions, especially in scenarios where re-localization costs are high.

4.3 Collaborative Mode-Switching

We combine mode-switching (Sec 4.2) with collaboration from our fixed modes scenario (Sec 4.1). This allows self-organization: in low interaction regimes, individuals mode-switch to be most productive, whereas, in high interaction regimes, fixed roles emerge. We examine whether this leads to optimal behavior.

Agents operate according to Figure 2(c), a superposition of Figure 2(a) and 2(b). We represent mode-switching using a rate, r_{MS} . The mean-field ODEs are

$$\frac{d\langle n_{DR}^{\text{Lost}} \rangle}{dt} = r_L \langle n_{DR}^{\text{NotLost}} \rangle - r_{MS} \langle n_{DR}^{\text{Lost}} \rangle - \frac{2r_{int}}{N(N-1)} \langle n_{PL} \rangle \langle n_{DR}^{\text{Lost}} \rangle \quad (7a)$$

$$\frac{d\langle n_{PL^\dagger} \rangle}{dt} = r_{MS} \langle n_{DR}^{\text{Lost}} \rangle - r_p \langle n_{PL^\dagger} \rangle \quad (7b)$$

$$\frac{d\langle n_{PL} \rangle}{dt} = r_p \langle n_{PL^\dagger} \rangle - r_{MS} \langle n_{PL} \rangle \quad (7c)$$

In eq 7(a), the first and last RHS terms mirror eq 2. The middle term in eq 7(a) and the first term in eq 7(b) denote the rate of mode-switching of lost dead reckoners to PL^\dagger . The final term in eq 7(b) and the first term in eq 7(c) represent the rate at which PL^\dagger transition to PL . In eq 7(c), the last term is the rate at which perfect localizers mode-switch to dead reckoners. Using conservation of the number of agents and solving for the steady state we get

$$\frac{\langle \Pi \rangle}{N} = 1 - \left(2 + \frac{r_{MS}}{r_p} \right) \frac{\langle n_{PL} \rangle_{t \rightarrow \infty}}{N} \quad (8)$$

$\langle n_{PL} \rangle_{t \rightarrow \infty}$ is the steady-state number of perfect localizers. Jacobian stability analysis and the Routh-Hurwitz Criteria confirm that this steady state is stable. If $r_{MS} \gg r_p, r_L, r_{int}$, the system self-organizes to individual mode-switching,

as seen by eq 8 approaching eq 6. Conversely, if $r_{MS} \ll r_p, r_L, r_{int}$, the system maintains fixed modes with an optimal number of perfect localizers, as evidenced by eq 8 approaching eq 3 with $N_{PL} = N_{PL}^*$ (eq 4).

What is the optimal r_{MS} ? In low r_{int} scenarios, where perfect localizers provide minimal support, maintaining a high r_{MS} prevents unnecessary loss in productivity that occurs by retaining dedicated perfect localizers. In high r_{int} scenarios, where interactions alone can maintain the productivity of dead reckoners, a low r_{MS} is preferable to avoid unnecessary mode-switching costs. Thus, if r_{MS} is inversely correlated with r_{int} , the system can maximize productivity by adapting its mode-switching rate based on the agents' interaction rate.

We call this form of collaboration **adaptive mode-switching**. We set $r_{MS} = \frac{\alpha}{r_{int}}$, where the proportionality constant, α , determines the sensitivity of this dependence. Figure 3(b) illustrates the performance of adaptive mode-switching with $\alpha = 0.01$, $\tau_p = 100$, and $r_L = 0.04$. It achieves near-optimal productivity for all interaction rates by converging to the optimal fixed mode configuration (Sec 4.1) when interaction rates are high and clipping to the individual mode-switching strategy (Sec 4.2) when interaction rates are low.

5 Simulation: Multi-Agent Well-Mixed Scenarios

In this section, we conduct agent-based simulations to validate the accuracy of the mean-field analytical model (Sec 4) and explore the impact of initial conditions and non-steady state transients. Additionally, we introduce a *smart collaboration* strategy to enhance productivity further.

We initialize $N = 30$ agents, each transitioning to the lost state (DR_{Lost}) after an interval of $\tau_L = 3.46$ s in dead reckoning mode. Every τ_{int} seconds, two agents in the swarm randomly interact in a well-mixed manner. If a lost dead reckoner (DR_{Lost}) interacts with a perfect localizer (PL), the dead reckoner corrects its localization error to 0. For each simulation run, we collect γ and Π for every agent at each timestep. The normalized average productivity per agent, $\langle \frac{\Pi}{N} \rangle$, is calculated as the cumulative productivity achieved by all agents, divided by N and the total run duration, $\tau_{total} = 200$ s.

We first consider the fixed modes scenario (Sec 4.1). Across 900 simulation runs, we vary the interaction rate $r_{int} = \frac{1}{\tau_{int}}$ and fraction of perfect localizers $\frac{N_{PL}}{N}$. The results are shown in Figure 4(a), where each data point represents a single simulation run. We see close alignment between simulations and the analytical model, with minor variations due to transient behavior.

Furthermore, smarter collaboration strategies significantly improve productivity. In Sec 4, only lost dead reckoners correct their localization estimates through interactions with perfect localizers. However, not-lost dead reckoners can also benefit from these interactions, correcting errors and maintaining productivity longer. Although challenging to model analytically due to non-linearity, numerical simulations confirm this effect. Figure 4(b) shows higher productivity in fixed mode configurations with this smart collaboration strategy, which we incorporate into subsequent simulations and experiments.

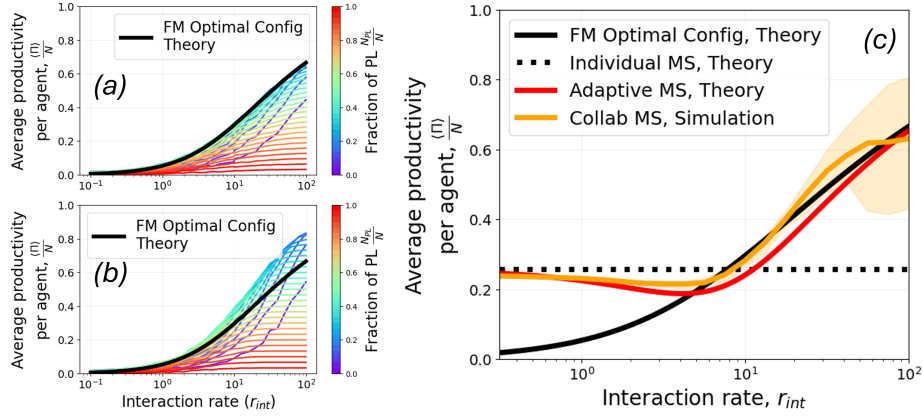


Fig. 4. Agent-based Simulation Results. (a) Fixed Mode and Basic Collaboration, (b) Fixed Mode and Smart Collaboration, (c) Collaborative Mode-Switching

A second difference between the analytical model and multi-agent simulations occurs for collaborative mode-switching. In the analytical model (Sec 4.3), we demonstrate that the ideal mode-switching rate is inversely related to the interaction rate r_{int} . However, in the multi-agent simulations, agents lack knowledge of r_{int} and must estimate it themselves. Each agent measures its local interaction rate, \hat{r}_{int} , using a simple counter and adjusts its mode-switching rate accordingly: $r_{MS} = \frac{1}{\hat{r}_{int}}$. If a dead reckoner mode-switches, it remains in the PL^\dagger state for $\tau_p = 10$ s before transitioning to a perfect localizer. Interactions between agents in the PL^\dagger and dead reckoner states have no impact on the swarm. Like before, we conduct 900 simulation runs, exploring the interaction rate parameter space and initial mode configuration space. Unlike the fixed mode case, agents can change their modes if the initial configuration is sub-optimal. This setup allows us to observe the effect of initial conditions on performance variability.

Figure 4(c) presents the simulation results for collaborative mode-switching, which mostly outperforms the analytical adaptive mode-switching scheme with $\alpha = 1$ across various r_{int} values. The orange-shaded region represents the range of productivity across all initial mode configurations. Mode-switching reduces the sensitivity of average productivity to the initial mode configuration compared to fixed modes. However, at higher interaction rates, average productivity still exhibits some dependence on the initial configuration, although this effect is weaker. This is because swarms that start in optimal configurations for high interaction rates achieve high productivity from the outset. In contrast, swarms beginning in sub-optimal configurations must undergo transient mode-switching, temporarily reducing total productivity.

The well-mixed simulations above showcase the flexibility and effectiveness of collaborative mode-switching. However, real-world swarms have spatial dependencies and non-uniform localization errors not captured by these simulations. We address these complexities in the next section by testing with real robots.

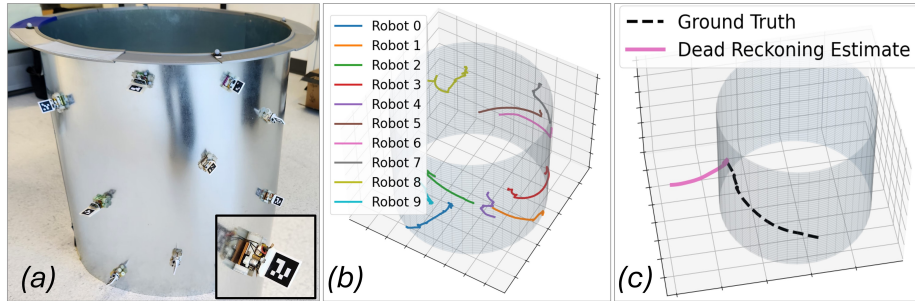


Fig. 5. (a) 10 Rovables on a 3D cylinder and zoom in of a single Rovable, (b) Ground truth trajectories of the 10 Rovables, (c) Dead reckoning estimate of Robot 6.

6 Hardware Experiments: 3D Inspection Swarm

Our group focuses on future applications of robot swarms for space structure inspection, aiming to reduce space debris and astronaut risks [2, 32]. We use a laboratory testbed of metal climbing Rovable robots [14] inspecting a 3D cylinder structure (Figure 5(a)). These robots are equipped with magnetized wheels for movement on metal surfaces and various onboard sensors, including radio communication and IMU for dead reckoning [14]. Future Rovables will include cameras for detecting relative 6D pose using fiducial tags [33]. Currently, we employ Vicon Motion Capture for line-of-sight vision over the curved 3D cylinder surface and to obtain ground-truth localization data.

We placed ten Rovables on a 3D metallic cylinder to simulate an inspection mission (depicted in Figure 5(a)). This setup allows us to explore two significant factors. (1) We investigate the transition from well-mixed to spatial settings, where robots are more likely to interact with nearby neighbors, leading to dynamic interaction rates and interaction networks that are influenced by inspection trajectories. (2) We consider the variable rate at which robots become lost, influenced by physical factors like slippage, robot maneuvers, and sensor noise.

To assess the impact of our collective localization method on swarm productivity, we let robots perform random walks on the cylinder. IMU measurements were recorded, and ground truth positions were tracked using Vicon (owing to Vicon limitations, trajectories were collected asynchronously). The trajectories of the ten robots are illustrated in Figure 5(b). We tested various scenarios offline using this trajectory dataset, allowing us to compare the performance of different algorithmic strategies given the same experimental conditions.

Robots are initially classified as perfect localizers or dead reckoners, with perfect localizers utilizing ground truth position data from the Vicon system, and dead reckoners relying solely on onboard IMU data for localization. The disorientation observed in our experiments is influenced by physical factors such as slippage, sensor noise, and gravitational effects. Additionally, the rate at which robots become lost varies, with certain maneuvers and orientations leading to

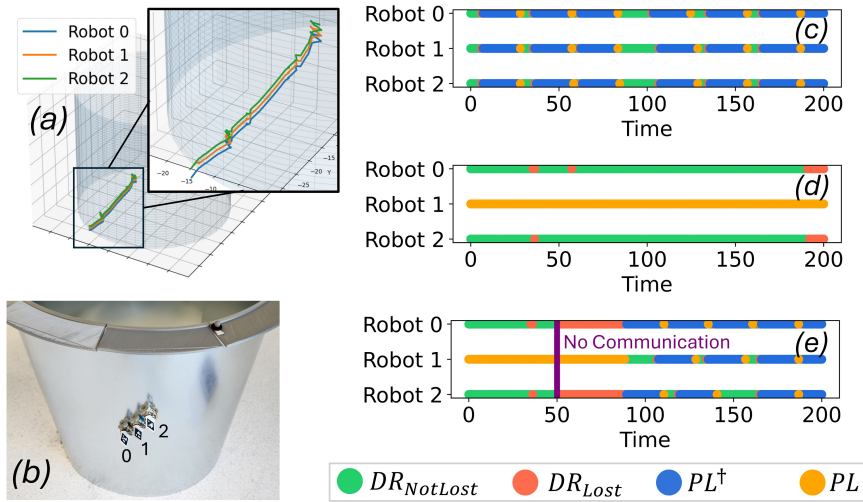


Fig. 6. Collective localization of three Rovables in formation. (a) Ground truth trajectories of robots in formation, (b) 3 Rovables in formation on the physical cylinder. Robot mode time series for (c) individual mode-switching, (d) collaborative mode-switching, and (e) communication failure.

increased localization errors. Dead reckoning alone proves highly unreliable in the 3D curved setting, as evidenced by Figure 5(c). We maintain $\gamma_{thresh} = 0.4$ as before and set the localization penalty as $\tau_p = 20$ s; this parameter will eventually be time-varied and dependent on the perfect localizers' localization mechanism (e.g. particle filtering or SLAM). For now, we do not consider the productivity effects of different τ_p . We also allow robots to know their localization error, so dead reckoners can detect when they are lost.

Interactions occur when robots establish line-of-sight vision of one another. During an interaction between a dead reckoner and a perfect localizer, the dead reckoner corrects its position using absolute position \mathbf{p}_{PL} and relative position \mathbf{p}_{rel} from the perfect localizer, reducing its localization error to zero (see Section 3). Currently, we do not account for errors in this process. We define the effective interaction rate as the frequency at which $DR_{NotLost}$ and DR_{Lost} interact with PL and PL^\dagger robots. Robots determine this effective interaction rate by averaging their effective interaction count over a time window τ_{window} . Averaging over τ_{window} allows for memory retention, enabling robots to make decisions based on recent observations while gradually forgetting older ones. The window size τ_{window} determines the sensitivity to interaction fluctuations. We set $\tau_{window} = 2\tau_p$. Mode-switching rate is inversely related to the average effective interaction rate. Unlike the well-mixed scenario where we could use the absolute interaction rate for mode-switching, in this context, we rely on the effective interaction rate to prevent the swarm from being trapped in local minima.

6.1 Collective Localization of Three Robots in Formation

We analyze $N = 3$ robots moving in formation to explore collective localization in a high interaction rate regime. However, formation control may not be the optimal strategy for inspection, as all robots cover the same subspace of the cylinder, limiting inspection coverage. Achieving formation control without vision is challenging, so we simulate it by shifting the Vicon data collected for a single robot (Robot 0) along the z-axis by 5 mm and 10 mm. The same IMU data is utilized for all three robots. To break symmetry, we vary the δp_0 parameter: [1.0, 1.3, 1.5]. In the spatial setting, the choice of formation also influences productivity. For the formation used in this case study, the middle robot interacts with two edge robots, but the edge robots do not see each other due to occlusion. Figure 6(a) and (b) display the trajectories and provide a visual representation of formation control on a physical cylinder. We examine three distinct cases:

- **Individual Mode-Switching** This scenario models individual robot localization without any collaboration. As depicted in Figure 6(c), the blue-yellow-green snake-like pattern demonstrates the robots alternating between productive phases and incurring time penalties in the PL^\dagger state. This behavior aligns with the periodic mode-switching pattern discussed in section 4.2, resulting in an average productivity per agent of 0.22.
- **Collaborative Mode-Switching:** Robots collaborate and are initialized as [DR, PL, DR]. Robot 1 interacts with Robot 0 and Robot 2 throughout the simulation. As illustrated in Figure 6(d), the swarm organizes to designate Robot 1 as the perfect localizer, while Robots 0 and 2 remain dedicated dead reckoners. Robots 0 and 2 use Robot 1 like a GPS, continuously obtaining position information through interactions. The total productivity per agent reaches 0.63, nearly three times higher than without collaboration.
- **Communication Breaks** We show the system’s robustness to interaction rate changes by implementing a scenario where robots can interact for the first 50 seconds, after which communication is disabled. As shown in Figure 6(e), initially, the swarm designates Robot 1 as the dedicated perfect localizer. However, after communication is disrupted, it dynamically self-organizes into a mode-switching behavior. Note, this behavior occurs purely via local interactions. Between 50 and 80 seconds, despite being unproductive, the robots persist in dead reckoning and perfect localizer mode. This hysteresis behavior is influenced by the memory window τ_{window} . The effect of the hysteresis is more profound for larger τ_{window} , but the swarm is also more resistant to fluctuation in interaction rate. Adjusting the memory window size allows system designers to modify swarm behavior.

This simple case study highlights several intriguing properties. First, in scenarios where robots do not collaborate, we illustrate how mode-switching enables robots to sustain productivity. Second, within the high interaction rate regime, the swarm realizes a substantial productivity increase by allocating one robot to support the maximum productivity of the remaining two robots. Lastly, the swarm demonstrates resilience to failures, such as signal jamming, reconfiguring to mode-switching when the interaction rate declines.

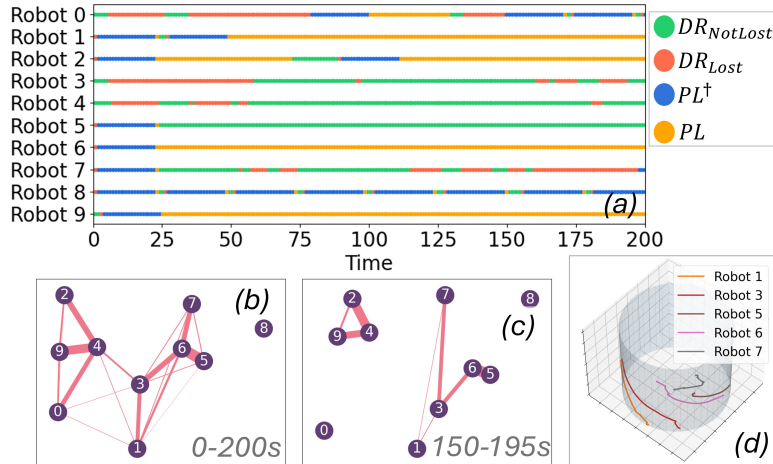


Fig. 7. 10 Random-Walkers Results. (a) Robot mode time series. Interaction network (b) from 0-200 s (full experiment duration) and (c) 150-195s. Edges represent interactions between robots within the specified time frames, with edge weight proportional to the number of interactions. (d) Ground truth trajectory of Robots 1, 3, 5, 6, 7.

6.2 Coverage by Ten Random-Walking Robots

We explore a complex scenario where all 10 robots follow unconstrained trajectories, randomly traversing the cylinder for coverage. We generate 30 runs by randomly rotating and translating the 10 original trajectories (Figure 5(b)) around the cylinder’s axis of symmetry. This approach produces complex interaction profiles, leading to dynamic interaction networks and varying interaction rates within and across experimental runs. We then analyze the productivity under Fixed Modes, Individual Mode-Switching, and Collaborative Mode-Switching.

All robots are initialized as dead reckoners, with $\delta p_0 = 1$ and $\tau_p = 20$ s. For Fixed Mode and Individual Mode-Switching, results are consistent across all runs as these modes are unaffected by interaction profiles. Under Fixed Mode, the productivity per agent is low at 0.02, as robots quickly become lost. Individual Mode-Switching yields slightly improved productivity, reaching 0.138 per agent. However, Collaborative Mode-Switching, which is sensitive to interaction profiles, performs the best for all 30 runs, achieving an average productivity per agent of 0.278 ± 0.046 .

To gain insight into the underlying dynamics and understand how interactions improve productivity, we examine the Robot Mode Time Series and Interaction Networks (Figure 7) of a single run. The productivity per agent for this run under collaborative mode-switching is 0.304.

- **Robot 8 is isolated from the rest of the swarm** throughout the run, as seen in Figure 7(b). Operating on the opposite side of the cylinder, it lacks line-of-sight connectivity with other robots. Consequently, as depicted

in Figure 7(a), Robot 8 exhibits mode-switching behavior similar to that shown in Figure 6(c). Robot 0 encounters similar isolation around $t = 150$ s, as illustrated in Figure 7(c). In Figure 7(a), we observe Robot 0 transitioning to pure mode-switching following a period of hysteresis.

- **Robot 6 interacts with Robots 7, 5, 3, and 1.** From Figure 7(a), the swarm self-organizes to sacrifice Robot 6 as the perfect localizer, allowing Robots 3, 5, and 7 to reap the productivity benefits (similar to Figure 6(d)). The behavior of Robot 1 is complex as it interacts with several other robots.
- **Robot 4 exhibits an interesting individual behavior**, consistently remaining a dead reckoner while benefiting from other perfect localizers (Figure 7(a)). This strategy, while reasonable on the individual level, is sub-optimal for the group. Given Robot 4’s central location (Figure 5(b) and 7(b)), it could assist more robots as a perfect localizer. If we rerun the trial with Robot 4 starting as a perfect localizer, productivity per agent increases to 0.38, with Robot 4 maintaining this role throughout the experiment. This highlights that the swarm can settle into sub-optimal mode configurations in spatial settings but still achieve high productivity.

We demonstrate that our collective localization approach operates in complex physical settings involving sensor noise, slippage, dynamic interaction profiles, and unconstrained trajectories. We showcase an array of interesting emergent sub-behavior that depends on local interactions: isolated robots individually mode-switch; robots with many interactions designate perfect localizers; robots adapt to changing interaction rates. These emergent behaviors collectively boost the productivity of the swarm. As collective localization is compatible with unconstrained trajectories, it can be applied to numerous inspection scenarios.

7 Conclusion and Future Work

Inspired by nature, we introduce a self-organized approach to collaborative localization that enhances inspection productivity of a robot swarm. This is achieved by strategically sacrificing some individuals to optimize computational efficiency. We showcase the advantages of this approach using theoretical models, numerical simulations, and hardware experiments conducted with a swarm of 10 metal-climbing robots. Our analytical mean-field model demonstrates that sacrificing agents as perfect localizers can improve the productivity of the swarm and that the swarm can dynamically reconfigure itself to maximize productivity for any interaction regime. This collaborative strategy depends only on local interactions and scales for any number of agents. We can further boost productivity using smart collaboration, as validated using numerical simulations. Our hardware experiments demonstrate the applicability and robustness of this collective localization mechanism in complex inspection-like scenarios. Future work involves running hardware experiments using more sophisticated metal climbing robots with vision and SLAM capabilities. We also plan to generalize this approach to scenarios beyond localization, such as vigilance and obstacle avoidance.

References

1. Halder, S., Afsari, K.: Robots in inspection and monitoring of buildings and infrastructure: A systematic review. *Applied Sciences* **13**(4), 2304 (2023)
2. Haghghat, B., Boghaert, J., Minsky-Primus, Z., Ebert, J., Liu, F., Nisser, M., Ekblaw, A., Nagpal, R.: An approach based on particle swarm optimization for inspection of spacecraft hulls by a swarm of miniaturized robots. In: *International Conference on Swarm Intelligence*, pp. 14–27. Springer (2022)
3. Correll, N., Martinoli, A.: Collective inspection of regular structures using a swarm of miniature robots. In: *Experimental robotics IX: The 9th international symposium on experimental robotics*, pp. 375–386. Springer (2006)
4. McLurkin, J., Smith, J., Frankel, J., Sotkowitz, D., Blau, D., Schmidt, B.: Speaking swarmish: Human-robot interface design for large swarms of autonomous mobile robots. In: *AAAI spring symposium: to boldly go where no human-robot team has gone before*, pp. 72–75. Palo Alto, CA (2006)
5. Mautz, R.: Overview of current indoor positioning systems. *Geodezija ir kartografija* **35**(1), 18–22 (2009)
6. Cadena, C., Carlone, L., Carrillo, H., Latif, Y., Scaramuzza, D., Neira, J., Reid, I., Leonard, J.J.: Past, present, and future of simultaneous localization and mapping: Toward the robust-perception age. *IEEE Transactions on robotics* **32**(6), 1309–1332 (2016)
7. Brügger, R., Willems, E., Burkart, J.: Looking out for each other: Coordination and turn taking in common marmoset vigilance. *Animal Behaviour* **196**, 183–199 (2023)
8. Beauchamp, G.: *Animal vigilance: monitoring predators and competitors*. Academic Press (2015)
9. Malley, M., Haghghat, B., Houel, L., Nagpal, R.: Eciton robotica: Design and algorithms for an adaptive self-assembling soft robot collective. In: *2020 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 4565–4571. IEEE (2020)
10. McCreery, H.F., Gemayel, G., Pais, A.I., Garnier, S., Nagpal, R.: Hysteresis stabilizes dynamic control of self-assembled army ant constructions. *Nature communications* **13**(1), 1160 (2022)
11. Prorok, A., Bahr, A., Martinoli, A.: Low-cost collaborative localization for large-scale multi-robot systems. In: *2012 IEEE International Conference on Robotics and Automation*, pp. 4236–4241. Ieee (2012)
12. Kurazume, R., Nagata, S., Hirose, S.: Cooperative positioning with multiple robots. In: *Proceedings of the 1994 IEEE International Conference on Robotics and Automation*, pp. 1250–1257. IEEE (1994)
13. Nemsick, B.E., Buchan, A.D., Nagabandi, A., Fearing, R.S., Zakhor, A.: Cooperative inchworm localization with a low cost team. In: *2017 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 6323–6330. IEEE (2017)
14. Dementyev, A., Kao, H.L., Choi, I., Ajilo, D., Xu, M., Paradiso, J.A., Schmandt, C., Follmer, S.: Rovables: Miniature on-body robots as mobile wearables. In: *Proceedings of the 29th Annual Symposium on User Interface Software and Technology*, pp. 111–120 (2016)
15. Thrun, S.: Probabilistic robotics. *Communications of the ACM* **45**(3), 52–57 (2002)
16. Eyvazpour, R., Shoaran, M., Karimian, G.: Hardware implementation of slam algorithms: a survey on implementation approaches and platforms. *Artificial Intelligence Review* **56**(7), 6187–6239 (2023)

17. Sullivan, N., Grainger, S., Cazzolato, B.: Analysis of cooperative localisation performance under varying sensor qualities and communication rates. *Robotics and Autonomous Systems* **110**, 73–84 (2018)
18. Lajoie, P.Y., Ramtoula, B., Wu, F., Beltrame, G.: Towards collaborative simultaneous localization and mapping: a survey of the current research landscape. arXiv preprint arXiv:2108.08325 (2021)
19. Luft, L., Schubert, T., Roumeliotis, S.I., Burgard, W.: Recursive decentralized localization for multi-robot systems with asynchronous pairwise communication. *The International Journal of Robotics Research* **37**(10), 1152–1167 (2018)
20. Klingner, J., Ahmed, N., Correll, N.: Fault-tolerant covariance intersection for localizing robot swarms. *Robotics and Autonomous Systems* **122**, 103,306 (2019)
21. Pires, A.G., Rezeck, P.A., Chaves, R.A., Macharet, D.G., Chaimowicz, L.: Cooperative localization and mapping with robotic swarms. *Journal of Intelligent & Robotic Systems* **102**(2), 47 (2021)
22. Li, Q., De Rosa, M., Rus, D.: Distributed algorithms for guiding navigation across a sensor network. In: *Proceedings of the 9th annual international conference on Mobile computing and networking*, pp. 313–325 (2003)
23. Tully, S., Kantor, G., Choset, H.: Leap-frog path design for multi-robot cooperative localization. In: *Field and Service Robotics: Results of the 7th International Conference*, pp. 307–317. Springer (2010)
24. Wanasinghe, T.R., Mann, G.K., Gosine, R.G.: Distributed leader-assistive localization method for a heterogeneous multirobotic system. *IEEE Transactions on Automation Science and Engineering* **12**(3), 795–809 (2015)
25. Allotta, B., Costanzi, R., Meli, E., Pugi, L., Ridolfi, A., Vettori, G.: Cooperative localization of a team of auvs by a tetrahedral configuration. *Robotics and Autonomous Systems* **62**(8), 1228–1237 (2014)
26. Haldane, D.W., Fankhauser, P., Siegwart, R., Fearing, R.S.: Detection of slippery terrain with a heterogeneous team of legged robots. In: *2014 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 4576–4581. IEEE (2014)
27. Pagliara, R., Gordon, D.M., Leonard, N.E.: Regulation of harvester ant foraging as a closed-loop excitable system. *PLoS computational biology* **14**(12), e1006,200 (2018)
28. Albani, D., Manoni, T., Nardi, D., Trianni, V.: Dynamic uav swarm deployment for non-uniform coverage. In: *Proceedings of the 17th international conference on autonomous agents and multiagent systems*, pp. 523–531 (2018)
29. Mayya, S., Wilson, S., Egerstedt, M.: Closed-loop task allocation in robot swarms using inter-robot encounters. *Swarm Intelligence* **13**, 115–143 (2019)
30. Elamvazhuthi, K., Berman, S.: Mean-field models in swarm robotics: A survey. *Bioinspiration & Biomimetics* **15**(1), 015,001 (2019)
31. Helbing, D.: *Quantitative sociodynamics: stochastic methods and models of social interaction processes*. Springer Science & Business Media (2010)
32. Chiu, D., Nagpal, R., Haghghat, B.: Optimization and evaluation of multi robot surface inspection through particle swarm optimization. In: *2017 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE (2024)
33. Kalaitzakis, M., Cain, B., Carroll, S., Ambrosi, A., Whitehead, C., Vitzilaios, N.: Fiducial markers for pose estimation: Overview, applications and experimental comparison of the artag, apriltag, aruco and stag markers. *Journal of Intelligent & Robotic Systems* **101**, 1–26 (2021)
34. Zurn, P., Bassett, D.S., Rust, N.C.: The citation diversity statement: a practice of transparency, a way of life. *Trends in Cognitive Sciences* **24**(9), 669–672 (2020)

8 Acknowledgments

This work is supported by NSF CMMI-2036359 and Amazon Robotics Research Award. Hungtang Ko is supported by the James S. McDonnell Foundation's Postdoctoral Fellowship. We thank Dr. Merihan Alhafnawi, Darren Chiu, Thiemen Siemensma, and Prof. Bahar Haghighat for assistance with the hardware testbed.

Inspired by a recent initiative to increase awareness and mitigate citation bias [34], we include a gender citation diversity statement using manually compiled data. Our references contain 12.1% woman (first author)/woman (last author), 18.2% man/woman, 6.1% woman/man, and 63.6 % man/man. The statistics we report fail to include intersex, non-binary, and transgender people. We look forward to future work that could help us better understand how to support equitable practices in science.