

Contact-Driven Localization in a Freeform Robotic Self-Assembled Structure

Mohammadali Rashidioun¹, Michael Sosa^{2,*}, Petras Swissler³

Abstract—Accurate localization remains a key challenge in swarm robotics, particularly for self-reconfigurable systems that must identify relative positions to form diverse structures. Most existing approaches rely on external tracking infrastructure or high-cost sensors, which limit scalability and deployment in unstructured environments. In this paper, we propose a novel contact-driven localization method for modular robots that leverages only local communication through binary contact information (whether two robots are physically connected or not). To exploit these contact cues, we introduce a virtual-force framework in which robots iteratively refine their poses—attracting toward dock-connected neighbors and repelling from non-connected ones. The method requires no external infrastructure and relies only on minimal onboard sensing. Simulations show effective localization during the assembly of towers and cantilevers, enabling accurate, scalable, free-form self-assembly.

I. INTRODUCTION

A. Background

A fundamental requirement for coordinated multi-robot behavior is *localization* [1]. Swarm robotics—multi-robot systems inspired by the collective behaviors of social animals—uses large groups of simple robots that cooperate to perform tasks beyond the capability of a single agent [2], [3]. In such swarms, agents must estimate their relative positions to maintain formations, avoid collisions, and achieve precise objectives in applications such as environmental monitoring [4], disaster response [5], and autonomous self-assembly—even in environments where external infrastructure is unavailable [6].

Absolute positioning using Global Navigation Satellite System (GPS, BeiDou, GLONASS, Galileo) [7] is mostly impractical in densely packed swarms or obstructed environments (underground, mountainous terrain) [8]. Accordingly, alternative localization strategies, each with their own set of limitations, have been proposed across different platforms to address these challenges. These strategies often fall into two categories, either being centralized through an external controller or observer or decentralized within the swarm. Several swarm platforms achieve localization by relying on overhead or environment-mounted systems. Coachbots [9] used fiducial tags mounted along arena walls, with cameras onboard estimating pose relative to detected tags. Later, in its second

¹ Industrial and Mechanical Engineering Department, New Jersey Institute of Technology. Mr943@njit.edu

² Ying Wu College of Computing, New Jersey Institute of Technology. ms3383@njit.edu

³ Industrial and Mechanical Engineering Department, New Jersey Institute of Technology. Petras.Swissler@njit.edu

* Work done prior to current employment

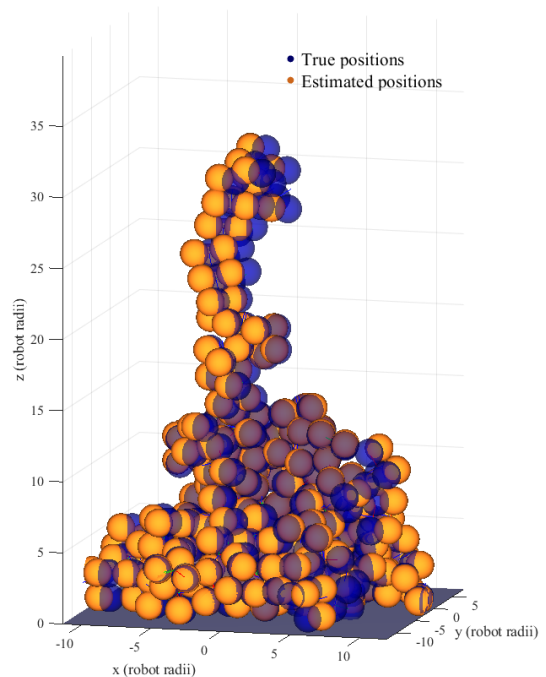


Fig. 1. Contact-driven localization in a ReactiveBuild-assembled tower. Blue spheres indicate ground-truth positions; orange spheres indicate estimated positions.

version [10], VR-based localization using infrared timing signals replaced the fiducial tags, which were decoded by onboard receivers for pose estimation. Zooids [11] employed projector-based optical patterns read by photodiodes on each robot, while SwarmBang [12] tracked ceiling beacons using upward-facing cameras. All of the mentioned methods are dependent on external installations and are largely confined to flat arenas, and thus cannot be easily generalized to unknown or outdoor environments.

In contrast to methods that use external installations, decentralized methods rely only on onboard sensing. Among decentralized 2D approaches, Kilobots [13] use range-only infrared sensing ‘seed’ robots to initialize a shared frame that then propagates hop-by-hop via distance gradients. This removes the need for external observers, but remains planar; errors accumulate with hop count, and the absence of bearing measurements limits orientation accuracy.

Bridging to 3D, aerial swarms routinely achieve precise

TABLE I
SENSOR COST PER MODULE AND MAIN LIMITATIONS.

System	Approximate Cost	Description
Coachbot VI [17]	\$20	2D localization using a camera and external ArUco tags.
Coachbot [18]	\$20 sensor + \$150 base	2D localization using IR signals via VR lighthouse.
Zoids [11]	\$1 sensor + \$1000 projector	High-precision 2D localization using high-speed structured light.
Kilobot [13]	Negligible	Decentralized 2D localization via range sensing and communication.
Free-SN [16]	\$150	Freeform 3D localization via 42 magnetometers + IMU; centralized.
M-Blocks [15]	\$50	3D Lattice-based localization using face tags.
ATRON [19]	Negligible	3D Lattice-based localization using IR communication
SwarmBang [12]	\$50	2D localization using line-of-sight to moving beacon
Our method	\$20 for FireAntV3 [?]	Freeform 3D contact-driven localization.

3D formations under external motion-capture (e.g., Vicon), but that centralizes estimation and ties deployment to lab infrastructure (not onboard sensing) [14].

Some modular platforms adopt lattice architectures. M-Blocks [15] use magnetic fiducial tags and absolute magnetic encoders to identify neighbors and connection orientation [15]. Although compact and passive, accuracy depends on precise face alignment and manufacturing tolerances. Because sensing requires face-to-face alignment, localization is naturally tied to axis-aligned lattice configurations.

FreeSN enables 3D freeform self-assembly through configuration identification, which combines dense magnetic sensing, inertial measurements, and odometry. A central computer aggregates these outputs to reconstruct the topology [16]. Despite its capability, this approach is hardware-intensive and depends on centralized processing.

Table I summarizes the localization methods discussed above. As you can see, most of the centralized 2D approaches rely on external infrastructure and line-of-sight; decentralized 2D approaches suffer error accumulation and do not readily extend to 3D; and freeform 3D systems offer expressiveness at substantial hardware and computational cost.

This paper presents an approach to robot localization in a 3D freeform self-assembled structure using modest sensing requirements. The major contributions of the presented work are as follows:

- 1) A method to enable freeform modular self-assembling robots to localize within a self-assembled structure.
- 2) Insight regarding the importance of robot physical characteristics in the context of localization.
- 3) Insight into sensing requirements for freeform self-assembling robotic systems.

The remainder of the paper is organized as follows: First, we present the unique characteristics of a FireAnt robot that provide additional localization information inherent in the physical design. Next, we present the methodology for using robot contacts to localize robots within a self-assembled structure. Finally, we examine several cases for sensor availability and present the localization quality for each.

II. FIREANT DESIGN FEATURES FOR LOCALIZATION

In most modular systems, localization is minimally constrained and high-dimensional; FireAntV3’s geometry and motion simplify the problem considerably. Each dock is fixed relative to the robot’s center, limiting possible connection configurations, and its discrete flipping gait provides predictable rotation angles and step heights. These constraints transform what is typically a difficult 3D localization problem into one that can be estimated fairly accurately using only local contact sensing and proprioception. Figure 2 illustrates how FireAntV3’s geometry and sensing capabilities progressively constrain the localization problem.

Initially, proprioceptive sensing from the IMU limits each robot’s orientation to a small set of candidate configurations. Identifying which previously added robot is connected can provide relational context that further narrows the feasible positions. Each FireAntV3 module consists of three spherical docks at fixed offsets from the robot’s center, and the robot can detect which of its docks is engaged with a neighbor. Recognizing specific dock-to-dock contacts imposes geometric constraints, as the distance between connected docks is fixed by the sphere diameter of the robot. Additional connections (e.g., Dock 1 of robot A to Dock 1 of robot B) further reduce the solution space. Finally, because FireAntV3 moves in discrete flips with known rotation angles and unit axes, the robot’s height (z) can be calculated directly from its motion history without external measurement. Through this layered process; combining orientation from the IMU, dock identity, and the geometric tangency and flipping history; the initially large configuration space collapses into an almost unique relative placement.

III. METHODOLOGY

A. Target platform

Our contact-based localization approach targets the simulated models of FireAnt-style robots (shown in Fig. 3), specifically **FireAntV3**. FireAntV3 is the latest generation of the FireAnt series, a set of modular, self-climbing robots designed for decentralized self-assembly in contact-rich environments [20]. Each module has a *three-sphere design* and employs a *flipping locomotion mechanism*, enabling traversal over arbitrarily arranged peer robots. Their *continuous docking surfaces* support strong, orientation-independent connec-

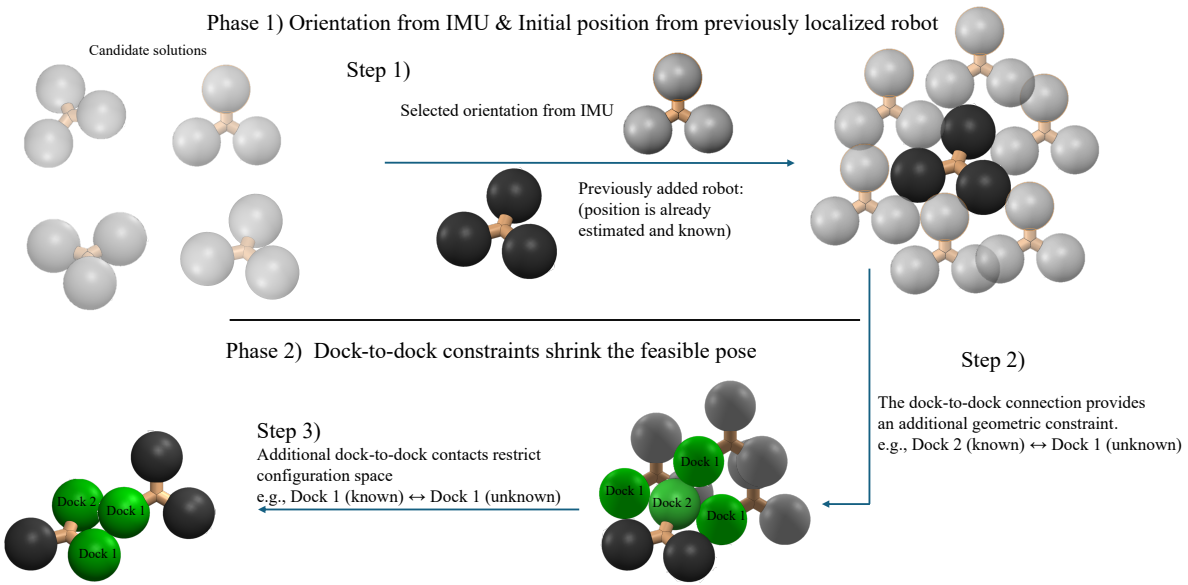


Fig. 2. **Contact-driven localization constraints with FireAntV3.** **Phase 1 (top):** Proprioceptive sensing (IMU) constrains the unknown robot’s orientation to a finite set of candidate poses. In addition, its connection to a previously placed robot provides an initial positional reference that further restricts feasible locations. **Phase 2 (bottom):** Dock-to-dock identities impose additional geometric constraints. For example, if Dock 2 of the known robot connects to Dock 1 of the unknown robot, the relative position is bounded by the fixed geometry of tangential spheres. Each additional dock-to-dock contact further constrains the pose, yielding a near-unique localization.




Fig. 3. (Left) The FireAnt V3 robot [20]; (Center) Orientation of the axes in a simulated FireAnt robot; (Right) A tower constructed by the ReactiveBuild algorithm [21], [22].

tions, allowing attachment at any point or angle without requiring precise alignment. FireAntV3 uses vibration-based neighbor identification and detects vibrations by measuring changes in consecutive 3D acceleration vectors with its IMUs. This allows it to distinguish peer-to-peer contacts from environmental noise and to identify which dock is engaged. The same IMU data also provides orientation, enabling FireAntV3 to localize both contact and pose.

FireAntV3 relies only on onboard IMUs for sensing, making it far simpler and cheaper than conventional localization approaches. External fiducials, motion capture systems, or GNSS are impractical for FireAnt and other self-assembly robots: line-of-sight is often blocked within dense 3D towers, GPS is unavailable indoors or underground, and lattice-constrained algorithms do not extend to the amorphous structures. Beyond individual sensing, FireAntV3 has also provided the foundation for an algorithmic study of self-assembly. In particular, the *ReactiveBuild* algorithm was developed and validated in simulation using FireAnt-style robots, showing how local contact and stress cues can drive the growth of amorphous, environment-adaptive structures.

ReactiveBuild enables modular robots like FireAntV3 to self-assemble structures in complex environments without needing a pre-defined shape or external planning [?]. Rather than building a rigid template, robots react to local conditions and other—growing structures, such as towers, chains, cantilevers, or bridges, based on where stress is high or where structural support is needed. Robots decide where to go based on goal direction and contact information. We extend this idea of using contact information in order to effectively calculate each robots position. This turns FireAntV3’s contact interaction into both construction and localization.

B. Proposed Localization Method

Our proposed contact-driven localization method builds on the sensing and interaction primitives of FireAntV3, and integrates naturally with the ReactiveBuild self-assembly framework. As new robots join a growing structure, the algorithm incrementally estimates their poses using only local contact sensing and proprioception. The first robot is designated as an anchor, and subsequent robots are localized relative to the anchor and FireAnts that have joined the structure (and have had their own positions calculated). In this way, each estimate builds upon earlier estimates. Although each FireAnts calculated location is relative to the anchor, as the number of robots increases, the given information restricts the possible configurations of robots within the structure; additionally, robots in contact with the ground provide a known z -height, offering a partial but independent vertical constraint that further bounds positional estimates. improving the overall accuracy.

C. Algorithm Overview

1) Initialization: When a new robot joins the structure, it first identifies previously localized neighbors with which it has direct physical contact. The centers of these neighbors are averaged to compute an initial position estimate. To resolve overlap and maintain spatial separation, the estimate is then projected outward along the robots local z -axis (See Fig. 3) by a fixed scalar offset referred to as the *scale factor*.

2) Localization refinement using virtual forces:

2.a) Attraction and repulsion from neighbors: The initial estimate is iteratively refined over T steps using virtual forces derived from dock-to-dock interactions. For each contacted neighbor, the Euclidean distance d and unit direction $\hat{\mathbf{d}}$ between connected dock pairs are computed. The virtual force applied to the robot's center is given by:

$$\mathbf{F} = \begin{cases} A \cdot \alpha_t \cdot (d - D) \cdot \hat{\mathbf{d}}, & d > D, \\ -B \cdot (1 - \alpha_t) \cdot (d - D)^4 \cdot \hat{\mathbf{d}}, & d < D, \end{cases} \quad (1)$$

where D is the desired dock spacing, A and B are attraction/repulsion gains (constants obtained by optimizing sum of squared error across all configurations), and α_t is a ramping factor. Early iterations prioritize repulsion to prevent overlap, followed by attraction to converge toward equilibrium spacing. This case reflects settings in which communication is limited to local neighbors.

2.b) Global disconnected repulsion: In a purely local-communication setting, a robot only knows which of its own docks are in physical contact; it has no information about nearby docks that are not contacting. The *undocked-repulsion* term therefore assumes a minimal global channel: robots periodically broadcast an estimated pose proxy (e.g., center and dock position). Using these messages, any pair of docks (d_i, d_j) that are not connected but lie within a collision threshold $d_{ij} < D$ exert a repulsive force to discourage crowding

$$\mathbf{F}_{ij}^{\text{undock}} = \begin{cases} -B(1 - \alpha_t)(d_{ij} - D)^4 \hat{\mathbf{d}}_{ij} & \text{if } d_{ij} < D \\ \mathbf{0} & \text{otherwise} \end{cases} \quad (2)$$

where d_{ij} is the Euclidean distance between the docks and $\hat{\mathbf{d}}_{ij}$ is the unit vector from dock j to dock i . This mirrors the connected-dock repulsion but applies it to nearby *unconnected* docks when the minimal global channel is available.

2.c) Z-Axis Correction: Each robot maintains a discrete record of its flipping history, including the rotation angles and axes of successive moves. Because FireAntV3 modules are initialized from a known base height, the z -position of each robot can be inferred from the cumulative vertical displacement implied by its flips. To exploit this, we introduce an additional virtual force directed from the current estimated position toward the z -coordinate derived from flipping history:

$$\mathbf{F}_z = k_z(z_{\text{flip}} - z_{\text{est}}), \quad (3)$$

Algorithm 1 Virtual Force-Based Localization Algorithm

Input : Swarm structure S , base robot b with known pose

Output: Estimated positions P for all robots

foreach new robot r joining the structure **do**

$P_r \leftarrow \text{average}(P_{N_r}) + s\hat{z}$; // Initialization

repeat

foreach neighbor $n \in N_r$; // Refinement
do

foreach dock pair (d_r, d_n) **do**

Compute distance $d_{r,n}$ and direction $\hat{u}_{r,n}$

if docks are connected **then**

if $d_{r,n} > D$ **then**

Apply attraction: $F \propto (d_{r,n} - D)$

else if $d_{r,n} < D$ **then**

Apply repulsion: $F \propto -(d_{r,n} - D)^4$

else

if $d_{r,n} < D$ **then**

Apply repulsion to avoid overlap

Update P_r using cumulative forces

if flipping history is available **then**

Apply attraction: $F \propto (z_{\text{flip}} - z_{\text{est}})$

until error < threshold or max iterations reached;

return P

where k_z is a correction gain and z_{flip} denotes the height predicted from flip records. This correction is especially effective in tower-like assemblies, where vertical stacking dominates, as it reinforces consistent alignment along the z -axis across robots. .

IV. SIMULATION-BASED VALIDATION

To validate the proposed localization algorithm, we use a virtual swarm of FireAntV3 robots in the ReactiveBuild simulation environment. Robots move toward a target region by flipping around their spherical docks and attaching sequentially; each robot localizes itself relative to previously placed neighbors. The simulator waits for each attachment to complete before introducing the next robot. We consider two self-assembly scenarios: *towers* (vertical stacking) and *cantilevers* (horizontal extension from a base). For each configuration, we generate and analyze 25 independent structures with 100 robots to quantify localization accuracy.

A. Evaluation Metrics

Let $\mathbf{p}_i \in \mathbb{R}^3$ denote the *true* center of robot i and $\hat{\mathbf{p}}_i \in \mathbb{R}^3$ its *estimated* center.

SSE and RMSE: We calculate the absolute localization error by

$$\text{SSE} = \sum_{i=1}^N \|\hat{\mathbf{p}}_i - \mathbf{p}_i\|^2, \quad (4)$$

$$\text{RMSE} = \sqrt{\frac{1}{N} \sum_{i=1}^N \|\hat{\mathbf{p}}_i - \mathbf{p}_i\|^2}, \quad (5)$$

where RMSE (Root mean square error), reported in robot radii, reflects the typical per-robot localization error.

Procrustes disparity (shape similarity): RMSE measures accuracy in the world frame, so global translation or rotation, like anchor drift, inflates error even if the relative arrangement of robots is correct. Following geometric morphometrics in biology (where shapes are compared after removing pose), we also report *Procrustes shape disparity* d , computed after optimally aligning the estimated configuration to ground truth with a rigid transform (rotation + translation; no scaling, no reflection). We first solve the optimization problem

$$\min_{\mathbf{R}, \mathbf{t}} \sum_{i=1}^N \|\mathbf{p}_i - (\mathbf{R} \hat{\mathbf{p}}_i + \mathbf{t})\|^2 \quad (6)$$

for \mathbf{R} and \mathbf{t} . Let $\mathbf{z}_i = \mathbf{R} \hat{\mathbf{p}}_i + \mathbf{t}$ be the aligned estimates and $\bar{\mathbf{p}} = \frac{1}{N} \sum_{i=1}^N \mathbf{p}_i$ the true centroid. We report the **Procrustes disparity**

$$d = \frac{\sum_{i=1}^N \|\mathbf{p}_i - \mathbf{z}_i\|^2}{\sum_{i=1}^N \|\mathbf{p}_i - \bar{\mathbf{p}}\|^2} \in [0, 1], \quad (7)$$

where $d = 0$ indicates identical configurations after alignment and $d = 1$ indicates no geometric similarity [23]. This makes d sensitive to structural fidelity (who is where relative to whom) while being invariant to global pose. As a simple example, consider two 100-robot towers with identical stacking order—one perfectly vertical, the other inclined by 10° . The inclined tower can show a large RMSE because every robot is displaced in the world frame, yet its shape disparity d remains small once aligned, indicating that the assemblies are geometrically similar. We therefore present both metrics: RMSE for deployment-relevant absolute accuracy and d to quantify how well our estimate matches the true structure.

For validation, we first run the algorithm with the full set of available features, including *Global disconnected repulsion*, which applies virtual repulsive forces between nearby but unconnected docks to prevent overlap in dense regions, and *Z-estimation*. We then progressively remove these features to assess their individual impact on localization accuracy.

B. Tower

To evaluate the algorithm, we simulated a scenario where the swarm constructs a vertical tower to a target height of 65 robot radii, with new robots added sequentially from random initial poses.

1) *Baseline*: The algorithm’s performance in reducing the RMSE error is shown in Fig. 6 as the red curve. The initial error from the projection step is rapidly minimized, with the system converging to a final RMSE of 1.64 robot radii (approximately 54 mm). Performance was observed to be dependent on the final structure geometry, with certain tower configurations achieving an even lower RMSE of 0.7 radii (approximately 23 mm). Figure 4 overlays the estimated and ground-truth robot positions for a sample configuration, demonstrating a close match. Another important observation




Fig. 4. Estimated vs True robot positions for a tower configuration. Each sphere marks a robot center.




Fig. 5. Per-robot localization error versus true height for 25 tower assemblies. As the tower grows taller, the pyramid-like geometry narrows the feasible region for each robot, reducing localization ambiguity and yielding lower error variability at greater heights. Robot joining order does not correlate with localization error. structural geometry is the dominant factor governing accuracy.

is that, as the tower grows, robots at greater heights exhibit lower error variation. This arises because their feasible placement region is progressively constrained by the hierarchical tower geometry: the cross-section narrows with height, reducing the number of admissible solutions. This trend is illustrated in Fig. 5, where localization error (Euclidean distance between the true position and the estimated one) decreases in variability at higher levels of the tower.

We next ablate additional sensory/interaction cues from the swarm system.

2) *Removing Z information (flipping history)*: We disable the flipping-history cue used to estimate each robot’s outward Z direction. Without this orientation prior, vertical mirror ambiguities remain, which slows convergence and increases localization error (blue vs. red in Fig. 6). The Procrustes disparity increases proportionally from $d = 0.045$ to $d = 0.051$ (+13%), consistent with the +12% increase in RMSE.

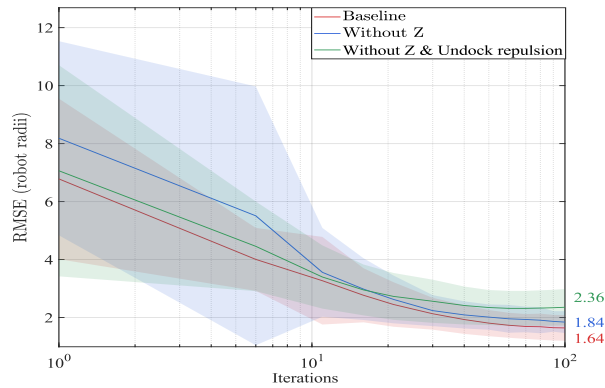


Fig. 6. Convergence of localization RMSE versus iterations for the tower configuration. Curves show the mean across 25 runs and shaded bands denote ± 1 STD. Baseline uses Z estimation from flipping history and global disconnected repulsion; Without Z removes the flipping-history cue; Without Z & undock repulsion removes both cues. Ablations slow convergence and increase steady-state error (final mean RMSE = 1.64, 1.84, and 2.36 robot radii, respectively).

This proportional degradation indicates that removing the Z cue uniformly affects both absolute accuracy and structural fidelity. In other words, the error is not simply a global vertical shift of the whole assembly, but a genuine distortion of relative robot placement. Nevertheless, the moderate magnitude of both increases suggests that the algorithm degrades gracefully without the flipping-history cue, preserving the overall tower shape reasonably well.

3) *Removing global disconnected repulsion*: We further simplify the algorithm by removing the undock-repulsion term. This allows non-contacting docks to collapse toward each other, biasing pose estimates, and yielding the largest error among conditions (green in Fig. 6). The Procrustes disparity increases from $d = 0.045$ to $d = 0.066$ (+47%), proportional to the +44% rise in RMSE. The magnitude of degradation is substantially larger than the Z -ablation case, indicating that global disconnected repulsion is more critical for maintaining structural fidelity in towers. Without it, nearby unconnected docks collapse inward, introducing genuine misplacement of robots relative to their true neighbors.

C. Cantilever

A cantilever is a structure that projects horizontally. To evaluate ReactiveBuild’s ability to produce such a form, we simulate robots growing a cantilever toward a goal located 45 units horizontally outward from the platform edge and 10 units(robot radii) vertically above it. As illustrated in Fig. 7, the horizontal plane terminates at a cylindrical boundary at its edge. Robots are introduced one at a time behind the current leading module, with randomly sampled initial offsets and orientations relative to the furthest placed robot, and then attach sequentially to extend the cantilever.

1) *Baseline*: The red curve in Fig. 9 is the *baseline* (all sensory cues enabled: flipping-history z cue and undocked-repulsion). Its steady-state error is higher than the

tower case (Fig. 6) for three reasons. First, the cantilever grows horizontally and does not develop the pyramid-like narrowing of a tower, so the solution space does not shrink as much based on height, and the z derived from flipping history is less informative. Secondly, robots are localized relative to previously estimated neighbors. In towers, robots 2 and 3 attach immediately to the anchor ($N=1$), creating a rigid base and limiting error propagation. In cantilevers, the first robot that actually contacts the anchor often appears much later in the join order (e.g., $N \approx 13$), so robots with indices 2–12 are placed relative to an unanchored chain, and their errors accumulate. This contrast is also visible in the scatter plots: in towers (Fig. 5, early-joining robots (dark blue, index ≤ 10) cluster among the lowest-error points, whereas in cantilevers (Fig. 8) the dark-blue points are not concentrated at low-error regions.

Despite these challenges, the baseline converges quickly from the initial projection step and reaches a final RMSE of 3.91 robot radii (≈ 92 mm), i.e., < 10 cm—acceptable for our module dimensions.

2) *Removing the Z cue (flipping history)*: In cantilevers, the vertical cue is weaker than in towers, but it is still helpful: it fixes the sign of vertical drift and reduces pose ambiguity near the plate. Removing it increases both error and uncertainty, slowing early convergence and raising the steady-state RMSE from 3.91 to 5.26 robot radii (Fig. 9). The wider shaded band reflects greater ambiguity when the structure grows horizontally without the flipping-history constraint. The Procrustes disparity increases from $d = 0.166$ to $d = 0.219$ (+32%), proportional to the +35% rise in RMSE. This consistency indicates that removing the Z cue degrades both absolute accuracy and structural fidelity uniformly, rather than introducing a simple global offset.

3) *Removing global disconnected repulsion (in addition to flipping history)*: Turning off global disconnected repulsion on top of the no- Z condition changes performance only in terms of final RMSE = 5.12 vs. 5.26 radii, and the shaded bands largely overlap in Fig. 9. This small effect is expected for cantilevers: robots extend into free space, so near but unconnected docks are rare, and spacing is already constrained by connected contacts and the plate geometry. However, the Procrustes disparity increases from $d = 0.219$ to $d = 0.251$ (+14%), indicating that the relative arrangement of robots within the structure worsens.

V. DISCUSSION

Table II summarizes RMSE and Procrustes disparity across all conditions for both tower and cantilever. RMSE values are expressed in Robot radii units; the baseline tower accuracy of 1.64 radii corresponds to approximately 54 mm error, while cantilevers reach about 92 mm. In both cases, robots localize themselves with accuracy on the order of one dock diameter, which is sufficient for reliable attachment in self-assembly. Localization quality depends strongly on structure geometry and join order. In towers, robots 2 and 3 attach directly to the anchor, forming a rigid base that locks pose early; errors stay low and do not accumulate (baseline

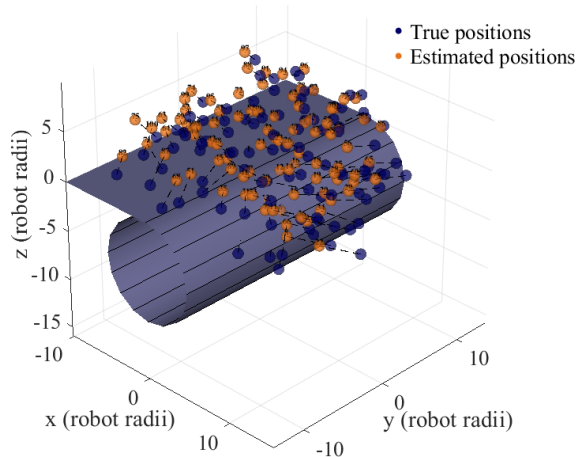


Fig. 7. Estimated vs. true robot positions for a cantilever configuration. Each sphere marks a robot center.




Fig. 8. Per-robot localization error versus true longitude for 25 Cantilever assemblies.

RMSE = 1.64 radii; some runs near 0.7). Cantilevers, by contrast, exhibit a large join-index gap before the anchor’s *first* physical neighbor appears (e.g., the 13th robot), so early placements drift and error accumulates (baseline RMSE = 3.91 radii). The global communication that enables undocked repulsion is clearly beneficial for *towers* since it reduces both RMSE and shape disparity. The reason is that towers are compact, so there are many near-but-unconnected docks, and the repulsion term helps prevent local collapse and bias. For *cantilevers*, the story is different: robots extend into free space, so undocked pairs are rarer; the term can lower d slightly but has little impact on RMSE. The Z cue (from flipping history) is a robust, low-cost constraint and should be enabled by default for both configurations.

VI. CONCLUSIONS

Through simulation with FireAntV3 modules, we demonstrated a contact-driven, infrastructure-free localization method that operates during freeform, environment-adaptive self-assembly. Because it relies primarily on IMU-based proprioception, which is a common capability of many self-assembling platforms, the method is not specific to FireAntV3 and can transfer with minimal changes to other modular systems. Motivated by the incremental nature of estimation and the join-index gap observed in cantilever as-




Fig. 9. Convergence of localization RMSE versus iterations for the cantilever configuration. Curves show the mean across 25 runs and shaded bands denote ± 1 STD. Baseline uses the flipping-history z cue and Global disconnected repulsion; Without Z removes the flipping-history cue; Without Z & undock repulsion removes both cues. The baseline converges to a final RMSE of 3.91 robot radii (≈ 92 mm).

TABLE II
RMSE AND PROCRUSTES ERROR ACROSS CONDITIONS

Condition		RMSE	Procrustes disparity
Tower	Baseline	1.64	0.045
	No Z Estimation	1.84 (+12%)	0.051 (+13%)
	No global comm. or Z est.	2.36 (+44%)	0.066 (+47%)
Cantilever	Baseline	3.91	0.166
	No Z Estimation	5.26 (+35%)	0.219 (+32%)
	No global comm. or Z est.	5.12 (+31%)	0.251 (+51%)

semblies, we propose exploring a local neighbor-to-neighbor *back-propagation* pass as part of future work to overcome error accumulation. In this approach, a back-propagation step would be triggered after every ten joins to revisit recent placements, allowing error corrections to propagate backward. This direction of investigation could help identify when and how frequently back-propagation should be applied to improve overall localization.

We also plan to reduce reliance on magnetometer yaw, which can be unreliable in practice due to disturbances and vibration, by introducing a virtual torque that aligns yaw using dock geometry and neighbor constraints, paired with IMU-centric filtering. Finally, we will deploy this decentralized localization on Simplified version of Fireant2d hardware to evaluate this method in experiments.

ACKNOWLEDGMENT

The authors are grateful to Roberto Torres for his help with the literature review.

REFERENCES

- [1] S. Wang, Y. Wang, D. Li, and Q. Zhao, “Distributed Relative Localization Algorithms for Multi-Robot Networks: A Survey,” *Sensors*, vol. 23, no. 5, p. 2399, Jan. 2023, number: 5 Publisher: Multidisciplinary Digital Publishing Institute.
- [2] M. Brambilla, E. Ferrante, M. Birattari, and M. Dorigo, “Swarm robotics: a review from the swarm engineering perspective,” *Swarm Intelligence*, vol. 7, no. 1, pp. 1–41, Mar. 2013.
- [3] Y. Tan and Z. yang Zheng, “Research advance in swarm robotics,” *Defence Technology*, vol. 9, no. 1, pp. 18–39, 2013.
- [4] B. Khaldi, F. Harrou, and Y. Sun, “Collaborative swarm robotics for sustainable environment monitoring and exploration: Emerging trends and research progress,” *Energy Nexus*, vol. 17, p. 100365, 2025.

- [5] A. Hafid, R. Hocine, and L. Guezouli, "Analyzing swarm robotics approaches in natural disaster scenarios: A comparative study," in *2024 1st International Conference on Innovative and Intelligent Information Technologies (IC3IT)*, 2024, pp. 1–6.
- [6] S. Chen, D. Yin, and Y. Niu, "A survey of robot swarms' relative localization method," *Sensors*, vol. 22, no. 12, 2022.
- [7] S. Jin, Q. Wang, and G. Dardanelli, "A review on multi-gnss for earth observation and emerging applications," *Remote Sensing*, vol. 14, no. 16, 2022.
- [8] S. Wang, Y. Wang, D. Li, and Q. Zhao, "Distributed relative localization algorithms for multi-robot networks: A survey," *Sensors*, vol. 23, no. 5, 2023.
- [9] G. Espinosa and M. Rubenstein, "Using hardware specialization and hierarchy to simplify robotic swarms," in *2018 IEEE International Conference on Robotics and Automation (ICRA)*, 2018, pp. 7667–7673.
- [10] H. Wang and M. Rubenstein, "Shape formation in homogeneous swarms using local task swapping," *IEEE Transactions on Robotics*, vol. 36, no. 3, pp. 597–612, 2020.
- [11] M. Le Goc, L. H. Kim, A. Parsaei, J.-D. Fekete, P. Dragicevic, and S. Follmer, "Zoooids: Building blocks for swarm user interfaces," in *Proceedings of the 29th annual symposium on user interface software and technology*, 2016, pp. 97–109.
- [12] M. Duan, X. Lei, Z. Duan, and Z. Zheng, "A minimalist self-localization approach for swarm robots based on active beacon in indoor environments," *Sensors*, vol. 23, no. 10, p. 4926, 2023.
- [13] M. Rubenstein, A. Cornejo, and R. Nagpal, "Programmable self-assembly in a thousand-robot swarm," *Science*, vol. 345, no. 6198, pp. 795–799, 2014.
- [14] A. Kushleyev, D. Mellinger, C. Powers, and V. Kumar, "Towards a swarm of agile micro quadrotors," *Auton. Robots*, vol. 35, no. 4, p. 287–300, Nov. 2013.
- [15] J. W. Romanishin, J. Mamish, and D. Rus, "Decentralized control for 3d m-blocks for path following, line formation, and light gradient aggregation," in *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2019, pp. 4862–4868.
- [16] Y. Tu, G. Liang, and T. L. Lam, "Freesn: A freeform strut-node structured modular self-reconfigurable robot-design and implementation," in *2022 International Conference on Robotics and Automation (ICRA)*. IEEE, 2022, pp. 4239–4245.
- [17] G. Espinosa and M. Rubenstein, "Using hardware specialization and hierarchy to simplify robotic swarms," in *2018 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2018, pp. 7667–7673.
- [18] V. Dornadula, L. Liu, M. Vejnovic, M. Zhang, and M. Rubenstein, "Coachbot swarm testbed: A 100-robot automated and remotely-accessible platform," in *International Symposium Distributed Autonomous Robotic Systems*, 2024.
- [19] D. Brandt, D. J. Christensen, and H. H. Lund, "Atron robots: versatility from self-reconfigurable modules," in *2007 International Conference on Mechatronics and Automation*. IEEE, 2007, pp. 26–32.
- [20] P. Swissler and M. Rubenstein, "Fireantv3: a modular self-reconfigurable robot toward free-form self-assembly using anywhere continuous docks," *IEEE Robotics and Automation Letters*, vol. 8, no. 8, 2023.
- [21] —, "Reactivebuild: Environment-adaptive self-assembly of amorphous structures," in *Distributed Autonomous Robotic Systems - 15th International Symposium, 2022*, ser. Springer Proceedings in Advanced Robotics, F. Matsuno, S. ichi Azuma, and M. Yamamoto, Eds. United States: Springer Nature, 2022, pp. 363–375, publisher Copyright: © 2022, The Author(s), under exclusive license to Springer Nature Switzerland AG.; 15th International Symposium on Distributed Autonomous Robotic Systems, DARS 2021 and 4th International Symposium on Swarm Behavior and Bio-Inspired Robotics, SWARM 2021 ; Conference date: 01-06-2021 Through 04-06-2021.
- [22] P. Swissler, "Behavioral heterogeneity enhances self-assembly: Exploring variants of the reactivebuild algorithm," in *International Symposium on Distributed Autonomous Robotic Systems*. Springer, 2024, pp. 100–114.
- [23] S. Zhang, H. Zhang, Y. Zhan, X. Wei, and Y. Liu, "Cluster networking and cooperative localization based on biogeography optimization and improved super-multidimensional scaling for multi-unmanned aerial vehicles," *Sensors*, vol. 25, p. 2887, 05 2025.