Available online at www.sciencedirect.com
ScienceDirect
Procedia CIRP 00 (2025) 000–000
www.elsevier.com/locate/procedia

CIRP Conference on Manufacturing Systems (CIRP CMS)

A Minimum-Risk Online Planner for Collision-Free Multi-Robot Additive Manufacturing

Ronnie F. P. Stone[a], Rachel Reynolds[b], Zhenghui Sha[a,*]

[a] *Center for Additive Manufacturing and Design Innovation, Walker Department of Mechanical Engineering, The University of Texas at Austin, TX 78712*
[b] *Department of Mechanical Engineering, Massachusetts Institute of Technology, Cambridge, MA 02139*

* Corresponding author. Tel.: +1-512-471-1818. *E-mail address:* zsha@austin.utexas.edu

---

Abstract

Cooperative 3D printing has recently emerged as a scalable extension of traditional additive manufacturing, enabling substantial gains in both makespan and build volume through the use of multiple robots. However, coordinating several printers within a shared workspace introduces the critical challenge of preventing inter-robot collisions. While prior work has largely focused on computationally intensive methods, such as optimal partitioning and coordinated toolpath generation, these approaches can be difficult to deploy in real-time. We introduce a lightweight, threat field–based method for generating minimum-risk trajectories when two printers require repositioning (e.g., homing to avoid collisions). The proposed algorithm integrates naturally with existing constraint-based strategies that command robots to retreat to safe configurations during hazardous interactions, providing a rapid, maximum-clearance motion solution. This paper offers a formal analysis of the underlying framework and numerical characteristics of such online planners, with particular emphasis on minimum-risk motion sequences for SCARA-based systems. Our numerical study examines how different distance metrics and workspace configurations affect computation time, threat-field reconstruction, and the accuracy of the resulting trajectories, as well as a simulated success rates.

Keywords: Cooperative 3D Printing; Multi-Robot Additive Manufacturing; SCARA Robots; Swarm Manufacturing; Robot Motion Planning

---

1. Introduction

Recent years have seen growing interest in multi-robot additive manufacturing (AM) [1, 22], also known as cooperative 3D printing (C3DP). The primary motivation behind these efforts is to develop systems that can scale traditional AM processes (e.g., FDM) in size, speed, and process capability, thereby enabling broader industrial adoption (see Fig. 1). In C3DP, multiple printing agents operate within a shared workspace, often collaborating on the fabrication of a single large-volume structure. Successful examples of C3DP systems include multi-laser powder bed fusion (LPBF) [29], multi-robot direct energy deposition (DED) [2], and multi-nozzle micro cold-spray (MCS) [27]. While the benefits of this paradigm are evident, introducing multiple autonomous manufacturing units also brings significant challenges. In particular, C3DP is inherently a multi-stage process that hinges on three core components: task division, task scheduling, and toolpath-level motion planning [20]. Optimizing these steps simultaneously requires a hierarchical or multi-level approach, which is computationally demanding, especially as the geometry of each printed layer evolves during the process. Additionally, offline optimization methods are sensitive to process uncertainties, including robot speed variations and printing tolerances [18, 23].

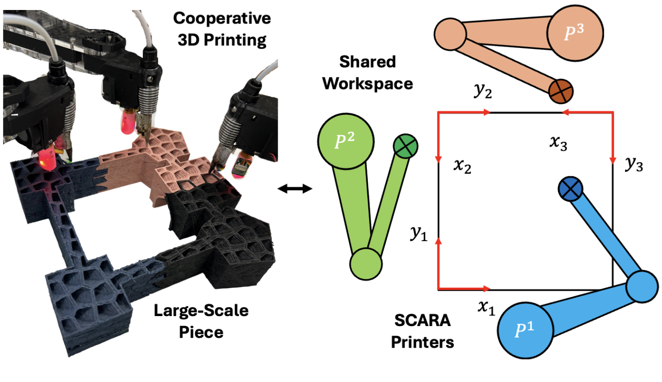

Fig. 1. An example of 3 SCARA printers engaging in a cooperative process [6].

A lightweight alternative to fully optimized C3DP pipelines is to adopt a constraint-based strategy in which a printer temporarily moves to a safe configuration when at risk of colliding [28]. This reactive approach does not guarantee a globally optimal makespan, but it is straightforward to deploy across a wide range of scenarios and preserves the originally designed toolpath, an important characteristic when specific mechanical are desired, given the inherent anisotropy of additively manufactured parts. However, current literature offers limited guidance on how to execute such homing and resuming (i.e., repositioning) maneuvers under general operating conditions, how to minimize collision risk during motion, and how to determine the desired configurations dynamically. In this context, the contributions of this paper are:

- A general computational framework for leveraging C3DP with SCARA robots.
- An optimal motion sequence algorithm that minimizes risk of collision in multi-robot arm manufacturing during repositioning maneuvers (e.g., homing).
- A sensitivity analysis of the methodology, including discretization resolution and failure modes.

As such, this paper stands as a unique addition to C3DP literature. The remainder of the paper is organized as follows. We summarize the existing literature on C3DP motion planning in Section 2. In Section 3, we present the conceptual framework supporting the investigation, including the SCARA kinematics model, threat fields, and safe regions. The methodology for the optimal homing sequence algorithm and the associated evaluation metrics are outlined in Section 4. Finally, a summary of the results is presented in Section 5, followed by a discussion of the limitations and future work in Section 6.

2. Related Work

Multi-robot motion planning and collision avoidance is not a topic exclusive to C3DP. In fact, there is a wealth of planners available for mobile robotics that involve sampling-based [14, 11], potential field [12], exact geometry [4], and, more recently, neural network-based approaches [21]. However, C3DP has a number of peculiarities in the problem definition that distinguish the methodological framework from traditional motion planning. As C3DP involves a multi-step approach: task division, task scheduling, and motion planning [20], any global optimization considerations must involve a multi-level approach.

One approach is to solve the collision problem purely from the task division and scheduling angle. In this context, specific regions are assigned and scheduled to printers in such a way that they are never working in adjacent areas [25, 26, 5, 13]. Alternatively, one can optimize at the toolpath level, by building traditional mathematical formulations such as linear programming or optimal control problems [9, 10, 24]. In any case, the assignment of the tasks or the toolpath can then be optimized according to some metric (e.g., makespan), however the process is computationally demanding and is generally done in a layer-wise manner, requiring re-optimization when the layer geometry changes.

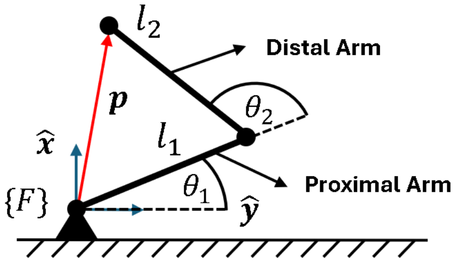

Fig. 2. Right-handed (elbow down) SCARA configuration used.

Finally, in the context of field-base methods, which are leveraged in this work, a significant family of algorithms come from potential field approaches [12]. The idea is to construct a field that can guide the robot's motion based on collision constraints. We leverage this approach in a hybridized sense, by incorporating the distance field of the printing robots to provide local minimum-risk motion trajectories, resulting in a balance between purely constraint-based C3DP approaches [28] and expensive, offline-optimal ones.

3. Conceptual Framework

In this section, we present the conceptual framework underlying the proposed online collision-avoidance strategy. We focus on discussing the SCARA robot kinematic model used, how the shared workspace is represented in the C3DP process, and the key notions of *safe regions* [28, 25, 26] and *threat fields*. However, the framework is general and can be applied to other robot kinematics.

3.1. SCARA Robot Representation

In this paper we follow the well-established manipulator theory for open-chain robots [19]. Without loss of generality, we assume that the robots are right-handed, have a fixed base denoted $\mathbf{b} \in \mathbb{R}^2$ in a two-dimensional Euclidean space, and distal and proximal link lengths l_1 and l_2. Thus, given the local frame of reference of the printer $\{F\}$ and any desired end-effector position $\mathbf{p} \in \mathbb{R}^2$, the forward kinematics will follow the system of equations

$$\begin{cases} x = l_1 \cos\theta_1 + l_2 \cos(\theta_1 + \theta_2), \\ y = l_1 \sin\theta_1 + l_2 \sin(\theta_1 + \theta_2), \end{cases} \tag{1}$$

where θ_1 is the joint angle between the proximal arm and the x-axis of $\{F\}$ and θ_2 is the joint angle between between the proximal and distal arms (see Fig. 2).

3.2. Shared Workspace in C3DP

We denote the printers in the C3DP process as P^i, where the superscript $i \in \mathbf{Z}^+$ indexes the printer number. The total number of printers in the process is given by $N \in \mathbf{Z}^+$ (i.e., $i \leq N$). Each printer P^i has a corresponding local frame of reference denoted

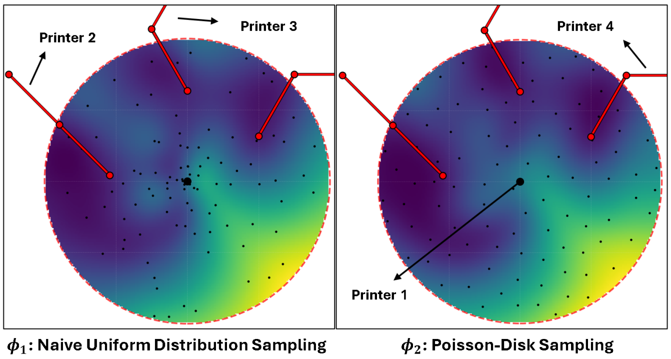

Fig. 3. Construction of two threat fields ϕ_1 and ϕ_2 based on multiquadric radial basis function interpolation. Using Poisson-disk sampling provides a better coverage of the SCARA workspace, resulting in a more accurate scalar field.

by $\{F_i\}$. Conversely, the shared workspace has a global frame of reference $\{G\}$ where the bases of the robots are given by $\mathbf{b}_i$. As such, all end-effector positions $\mathbf{p}^i$, expressed in their local frames of reference $\{F_i\}$, can be represented in the global frame via the rigid-body transformation

$$ {}^{G}\mathbf{p}^i = \mathbf{a} + R({}^{F_i}\mathbf{p}^i), \tag{2}$$

where vector $\mathbf{a}$ and the rotation matrix R represent the displacement and orientation of frame $\{F_i\}$ with respect to frame $\{G\}$, respectively. Therefore, given a set of printers, their local frames of reference, and their current end effector positions, the kinematic chains can be projected onto the global shared workspace via Eq. 2 and collision can be detected using various strategies.

3.3. Threat Fields

Given any two rigid bodies in a two-dimensional Euclidean workspace, a collision occurs if and only if the intersection of their occupied regions is non-empty. In practice, this is typically checked by discretizing the bodies into polygon-shaped approximations and applying standard computational geometry collision algorithms [16]. In our application, however, we are not solely interested in detecting whether two SCARA robots are colliding, but also in quantifying how close they are under a suitable distance metric $d : \mathbb{R}^2 \to \mathbb{R}$. By examining how the kinematic chain evolves as the end-effector position $\mathbf{p}^i$ varies, we construct, for each printer P^i, a scalar field $\phi_i(\mathbf{p}^i) = d(\mathbf{p}^i)$ that encodes the distance to collision with other printers at each configuration $\mathbf{p}^i$.

As deriving a closed-form analytical expression for ϕ_i is, except in very simple geometries, intractable [17, 15], we discretize the local workspace in polar coordinates centered at $\mathbf{b}_i$ (e.g., the base of the ith printer). Polar coordinates are preferable to Cartesian coordinates because the planar workspace of a SCARA robot is annular [7]. We employ a Poisson-disk sampling strategy, which, compared to a naive uniform sampling scheme, provides better spatial coverage and consequently a more accurate scalar field ϕ for the same computational cost [3]. The resulting samples are then interpolated using multiquadric radial basis functions (RBFs) [8] to generate the final threat fields ϕ_1 (naive) and ϕ_2 (proposed) (see Fig. 3).

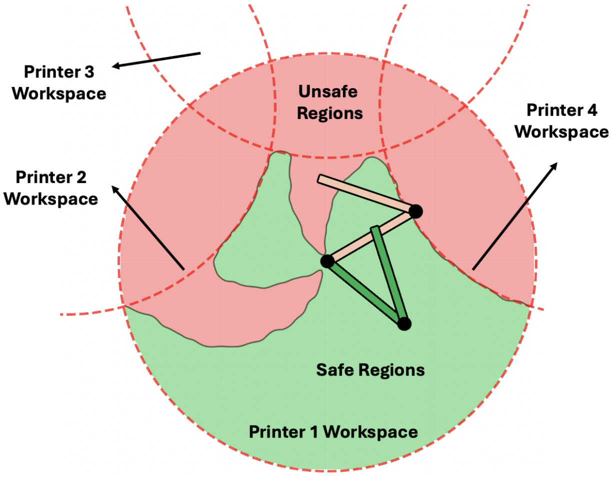

Fig. 4. Construction of the safe regions of P^1 based on the workspace of the other SCARA robots. Sample safe and unsafe configurations are shown.

3.4. Safe Regions

The threat fields ϕ in C3DP applications are not static, but rather time-varying as the robots move in the shared workspace during the manufacturing process. As such, the threat field can be represented as a time-dependent function $\phi_i(x, y, t)$, which has to be recalculated at every discretized time interval Δt. In this context, we define the *safe region* of printer P^i as the set of end-effector positions for which $\phi_i(x, y, t)$ is non-zero for all values of t. Mathematically, we can define this set as

$$ S_i = \{(x, y) \mid \phi_i(x, y, t) > 0, \forall t \leq T\}, \tag{3}$$

where T is the total duration of the manufacturing process. Intuitively, if the SCARA robot has its end-effector position in the safe region (i.e., $\mathbf{p}^i \in S$), then it will not collide with any other robot j regardless of their configurations. Conservatively, the safe zone of robot i can be constructed geometrically by defining the union of the circles C_j, centered at the base of the each robot j and with radius $l_1^j + l_2^j$, as the worst case obstacle space for printer P^i. If we let $A(\mathbf{p}^i)$ be the area occupied by robot i at any end-effector position $\mathbf{p}^i$, the safe region is then defined as:

$$ S_i = \{\mathbf{p}^i \mid A(\mathbf{p}^i) \cap C_j, \forall j \neq i\}, \tag{4}$$

which can be pre-computed and used throughout the C3DP process in real-time for collision avoidance applications (Fig. 4).

4. Methodology

In this section, we present the methodology, based on the proposed conceptual framework, to achieve:

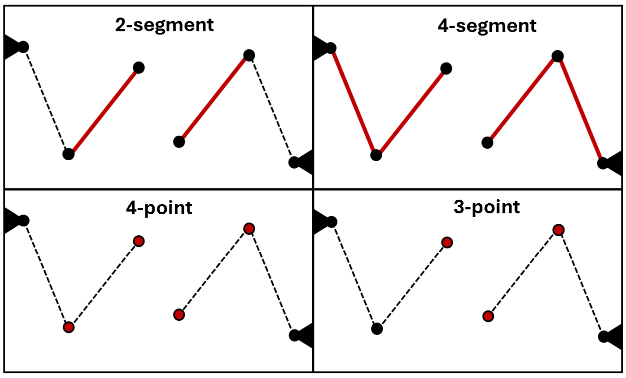

Fig. 5. Different types of discretization used in the sensitivity analysis.

- An algorithm for finding the minimum-risk trajectory to the nearest safe region for a given SCARA robot in case of a potential collision.
- A sensitivity analysis of the proposed algorithm based on different distance metrics, number of robots, and failure modes.

4.1. Minimum-Risk Homing Sequence

In a shared workspace it is essential to be able to move one or multiple arms away in the event of a predicted collision. Our goal is to develop a method for finding an optimal homing sequence that minimizes the risk of collision while remaining time-efficient. Thus, the optimal retreat sequence must maximize the distance from all colliders along the path while minimizing travel time. The reason for maximizing the distance is three-fold. First, avoiding collisions takes priority in the manufacturing process over small gains in makespan. Second, the end-effector position may be subject to uncertainties during the motion process (e.g., calibration). Finally, due to the dynamic nature of the problem (i.e., the threat field is time-varying), the homing sequence will be recomputed at every discretized time step Δt. Thus, a path with a higher clearance (i.e., higher ϕ) will allow the robots to react to moving obstacles in sufficient time.

To achieve this, we propose leveraging *gradient descent* on the static scalar field $-\phi(x, y)$. The original field $\phi(x, y)$ is multiplied by negative one to ensure that the final path (e.g, the homing sequence) goes toward the safe regions and away from the other robots. As such, the optimal homing sequence for the ith printer, in the static threat field, is sequential set of end-effector positions $\Pi = (\mathbf{p}^i_1, \mathbf{p}^i_2, ...\mathbf{p}^i_k)$ is given by the update rule

$$\mathbf{p}_{n+1} = \mathbf{p}_n - \eta\nabla\phi(\mathbf{p}_n), \tag{5}$$

where $\nabla\phi(\mathbf{p}_n)$ is the gradient of the threat field at end-effector position $\mathbf{p}_n$, and η is the step size (i.e., learning rate) of the descent. Note that, in general, if two printers are in danger of colliding, only one of them will follow the optimal homing sequence, as the other can continue the manufacturing process. By following the gradient $\nabla\phi$ we can guarantee that the trajectory locally minimizes the collision risk by moving towards the neighboring region with maximal clearance (i.e., higher ϕ.)

The stop condition of the algorithm is given when the end-effector position $\mathbf{p}_{n+1}$ is contained within the safe region of the printer (i.e., $\mathbf{p}_{n+1} \in A$), resulting in the final sequence starting at $\mathbf{p}_1$ and ending at $\mathbf{p}_{n+1}$. This condition allows for a faster computation of the homing sequence, while guaranteeing that the printer is collision-free irrespective of the movement of nearby printers. Finally, the calculation of the gradient at any point is done by first computing the gradient via central differences on the sampled points that generated $\phi(x, y)$ and then using bilinear interpolation for non-sampled points.

4.2. Sensitivity Analysis and Evaluation Methods

To evaluate the algorithm and its feasibility in physical scenarios, four evaluation metrics were used: computation time, threat field reconstruction accuracy, homing sequence accuracy, and success rate. These metrics change depending on how accurately the threat field is constructed, which is dependent on the distance metric $d(\mathbf{p}^i)$ used, and the physical characteristics of the environment (e.g, location and number of robots).

As the ground truth for comparison, the Euclidean distance using all robot links $(l^1_1, l^1_2, ...l^N_1, l^N_2)$, here denoted *4-segments*, was used. This metric represents the distance from the entire kinematic chain (i.e., both distal and proximal arms) of the printer of interest, to the entire kinematic chain of all other printers. While the most accurate, it can be computationally expensive and not able to find optimal homing solutions in real-time. For comparison, we considered three other metrics, in order descending order of more accurate to less accurate (see Figure 5):

- 2-segment, which considers only the distal arms of all printers.
- 4-points, which considers the end-effectors, positioned at $\mathbf{p}_i$, and the elbow joints that connect both arms of all printers.
- 3-points, which considers the end-effector position, positioned at $\mathbf{p}_i$, and the elbow joint of only the printer of interest (i.e., not all printers).

Time to compute was determined by checking the computers internal clock at the start of the program call and comparing it to the time value once the output was returned. This calculation included only the time time to compute $\phi(x, y)$, $\nabla\phi(x, y)$, and the iterative calculation of the gradient descent path.

To evaluate the accuracy D_m of the optimal homing sequence, the distance between the final end-effector position, based on the baseline 4-segment case, and the final end-effector position of the other cases was compared. Given n-runs in different scenarios, each final homing sequence end-effector position $\mathbf{p}^i_k$ was compared to the initial position $\mathbf{p}^i_1$ via

$$D_m = \frac{1}{n}\sum_{n=1}^{n} || \mathbf{p}^i_k - \mathbf{p}^i_1|| \tag{6}$$

where n is the number of runs, and $||.||$ is the Euclidean norm. Additionally, in order to demonstrate the effect of outliers and the skewness of the distance deviation average across samples, we also recorded the percent of samples that were within a certain tolerance (e.g., a small radius τ) of the baseline value. The value of τ was considered to be the minimum distance between points accepted during the Poisson-disk sampling (see Section 3.3).

The threat field reconstruction accuracy was determined by taking the point-difference between the baseline and the other methods. The values were normalized before subtraction. This is due to each method having different magnitudes of distances for the same point, despite having similar shapes. This allows comparison of relative shapes rather than raw distance values. The average of those differences was then computed and used as the reconstruction accuracy value. The numerical value of the distance accuracy does not reflect the distance offset, due to its normalization, but rather the relative difference between values.

Finally, the success rate of the method was estimated using a Monte Carlo simulation in which random scenarios were generated with various number of SCARA robots and configurations. For each scenario, successes were registered if the printer requiring repositioning ultimately reached the safe zone without collisions. Otherwise, a failure was added to the tally, even if the printers did not collide.

5. Results

In this section we summarize the numerical results derived from the evaluation of the proposed algorithm.

5.1. Simulation Environment & Computation Time

Simulations used 2–6 SCARA robots arranged around a shared print area with dimensions 550 × 250 *mm*. Without loss of generality, all SCARA had link lengths of 1 (i.e., $l_1^i = l_2^i = 1$). The values for grid size, number of arms, and the choice of distance metric d(.) varied according to Section 4.2. Each configuration was run 100 times with identical randomized end-effector positions. From there, gradient descent was performed with a maximum number of iterations of 500 and a η value of 1.

The 4-segment method was approximately 3.2× slower than 2-segment, and both point-based methods completed in less than $0.1s$ for most runs. Increasing arm count increased times for all methods due to more collision checks. The slowest case (4-segment, grid size 110, 6 arms) required $1.94s$. Time ranking was consistent: 4-segment, 2-segment, 4-point, and then 3-point (see Fig. 6, top left).

5.2. Path & Threat Field Reconstruction Accuracy

The path accuracy follows the same pattern as the previous two, with three-point being the worst and two-segment being the best. Even the largest deviation (75 *mm* for 3-point) is small relative to the build plate size, making all methods safe though not equally optimal (see Fig. 6, top right). Considering the mean error based on percentage of solutions outside of tolerance τ, it is important to note that the increase of the grid size decreases the tolerance, accentuating the decrease in accuracy (see Fig. 6, bottom left).

The mean error of the threat field reconstruction stays relatively constant over larger sampling sizes. This implies that when given the same resolution, the other three methods will continue to be off by the same amount. The grid mean error did increase slightly as the number of arms increased (see Fig 6, bottom right). The two-segment metric was the best, compared to the baseline, with significantly lower than the other two point-based methods. While an informative metric, the scalar field reconstruction did not provide significant insights into the difference between the distance metrics, as opposed to the path accuracy and the computational time.

5.3. Success Rate

After running a 1000 randomized scenarios in which the proposed method was applied, the success rate was of 77.1%. Failures could be reported in two cases, either the printer stalled in regions of the field ϕ in which the gradient was near zero (e.g., a saddle point or a minima), corresponding to the remaining 22.9%, or it collided with a neighboring printer, which never occurred (i.e., 0%). While the success rate is not ideal, we posit that most of the randomized scenarios were unrealistic, setting up the printer of interest in the middle of, or too close to the others. Such scenarios create highly non-convex fields ϕ in which gradient descent struggles, however, they would not be commonplace in C3DP applications. Finally, one could alleviate the observed failure modes by introducing small stochastic perturbations to the end-effector position when $\nabla\phi \approx 0$.

6. Conclusion, Limitations and Future Work

Overall, we find that the two-segment method creates a balanced approach between accuracy and computation time, especially on lower sampling resolutions. For applications requiring high precision, the 4-segment is best and should be used as the baseline value, albeit it is likely to be less reactive in real-time manufacturing scenarios. The point methods are close in performance, so we recommend considering all four points entirely, and use it in applications requiring fast frequency of computations (e.g., high-speed printing). There is also the possibility of employing a hybrid approach, in which the initial calculation can be performed with the 4-point method to quickly give the printer an initial sequence, and then, time-permitting, future sequence steps can use more accurate methods to find more optimal routes. Since the sampling number of the grid seems to have relatively little effect on accuracy, the authors recommend a choosing a value that provides a lightweight computation to match the frequency of update needed on a case-by-case basis.

The use of threat field- and gradient-based methods to find optimal motion sequences for SCARA arms extends beyond reactive collision-avoidance. For instance, the length of the final

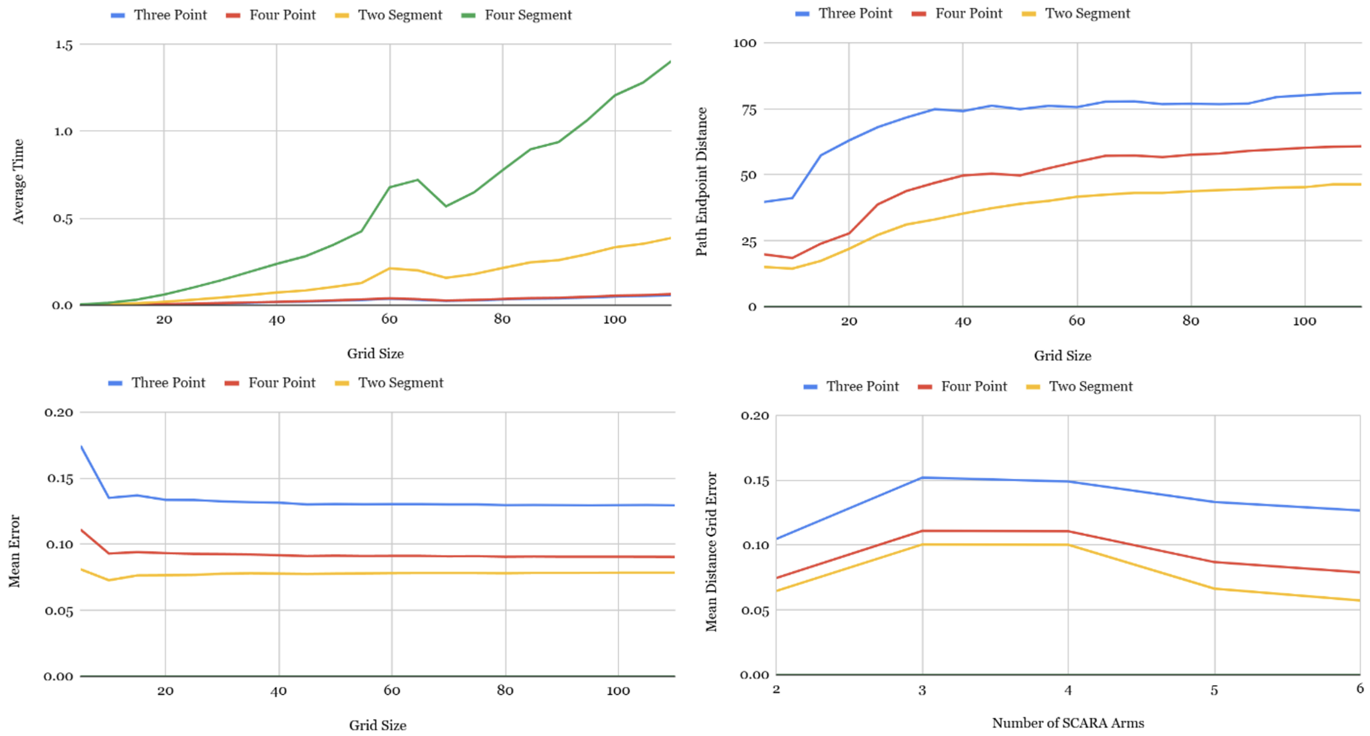

Fig. 6. Summary of numerical results. Top left: discretization computation time comparison in seconds. Top right: path accuracy based on final end-effector distance to baseline in millimeters. Bottom left: mean error of percentage of off-tolerance solutions. Bottom right: mean error of threat field reconstruction via normalized point-wise difference.

homing path can be used as a metric for the ease in which a SCARA arm can move in a specific region of the shared workspace. This can be used to create a kinematic-induced partition, improving existing safe region-based methods. Finally, the framework can be implemented in a variety of applications, and is extendable to other types of manipulators (e.g., 6-DOF arms). Future work will investigate more robust gradient-based methods such as medial axis transform and potential field approaches, which can alleviate premature convergence at local minima. Physical experiments will also be carried out in a swarm manufacturing testbed using multiple mobile SCARA robots.

References

[1] Alhijaily, A., Kilic, Z.M., Bartolo, A.P., 2023. Teams of robots in additive manufacturing: a review. Virtual and Physical Prototyping 18, e2162929.

[2] Arbogast, A., Nycz, A., Noakes, M.W., Wang, P., Masuo, C., Vaughan, J., Love, L., Lind, R., Carter, W., Meyer, L., et al., 2024. Strategies for a scalable multi-robot large scale wire arc additive manufacturing system. Additive Manufacturing Letters 8, 100183.

[3] Bridson, R., 2007. Fast poisson disk sampling in arbitrary dimensions. SIGGRAPH sketches 10, 1.

[4] Canny, J., 1988. The complexity of robot motion planning. MIT press.

[5] Ebert, M., Stone, R., Akleman, E., Sha, Z., Krishnamurthy, V., 2025a. Traveling cellsman: Partition-cluster co-parameterization for multi-robot cooperative 3d printing. Additive Manufacturing , 104987.

[6] Ebert, M., Stone, R.F., Koithan, J., Zhou, W., Pharr, M., Estrin, Y., Akleman, E., Sha, Z., Krishnamurthy, V., 2025b. Noodleprint: Cooperative multi-robot additive manufacturing with helically interlocked tiles. Journal of Manufacturing Science and Engineering 147, 061002.

[7] Gosselin, C.M., Guillot, M., 1991. The synthesis of manipulators with prescribed workspace. Journal of Mechanical Design 113, 451–455. doi:`10.1115/1.2912804`.

[8] Hardy, R.L., 1971. Multiquadric equations of topography and other irregular surfaces. Journal of geophysical research 76, 1905–1915.

[9] Jiang, Z., Wang, H., Sun, Y., 2021. Improved co-scheduling of multi-layer printing path scanning for collaborative additive manufacturing. IISE Transactions 53, 960–973.

[10] Jin, Y., Pierson, H.A., Liao, H., 2019. Toolpath allocation and scheduling for concurrent fused filament fabrication with multiple extruders. IISE Transactions 51, 192–208.

[11] Kavraki, L.E., Svestka, P., Latombe, J.C., Overmars, M.H., 2002. Probabilistic roadmaps for path planning in high-dimensional configuration spaces. IEEE transactions on Robotics and Automation 12, 566–580.

[12] Khatib, O., 1986. Real-time obstacle avoidance for manipulators and mobile robots. The international journal of robotics research 5, 90–98.

[13] Krishnamurthy, V., Poudel, L., Ebert, M., Weber, D.H., Wu, R., Zhou, W., Akleman, E., Sha, Z., 2022. Layerlock: Layer-wise collision-free multi-robot additive manufacturing using topologically interlocked space-filling shapes. Computer-Aided Design 152, 103392.

[14] Kuffner, J.J., LaValle, S.M., 2000. Rrt-connect: An efficient approach to single-query path planning, in: Proceedings 2000 ICRA. Millennium conference. IEEE international conference on robotics and automation. Symposia proceedings (Cat. No. 00CH37065), IEEE. pp. 995–1001.

[15] Latombe, J.C., 2012. Robot motion planning. volume 124. Springer Science & Business Media.

[16] Lin, M., Gottschalk, S., 1998. Collision detection between geometric models: A survey, in: Proc. of IMA conference on mathematics of surfaces, pp. 602–608.

[17] Lozano-Perez, T., 1983. Spatial planning: A configuration space approach. IEEE transactions on computers 32, 108–120.

[18] Mensch, C., Swaminathan, A., Sha, Z., 2024. A real-time monitoring framework for cooperative 3d printing .

[19] Murray, R.M., Li, Z., Sastry, S.S., 2017. A mathematical introduction to

robotic manipulation. CRC press.

[20] Poudel, L., Marques, L.G., Williams, R.A., Hyden, Z., Guerra, P., Fowler, O.L., Sha, Z., Zhou, W., 2022. Toward swarm manufacturing: architecting a cooperative 3d printing system. Journal of Manufacturing Science and Engineering 144, 081004.

[21] Qureshi, A.H., Simeonov, A., Bency, M.J., Yip, M.C., 2019. Motion planning networks, in: 2019 International Conference on Robotics and Automation (ICRA), IEEE. pp. 2118–2124.

[22] Rescsanski, S., Hebert, R., Haghighi, A., Tang, J., Imani, F., 2025. Towards intelligent cooperative robotics in additive manufacturing: Past, present, and future. Robotics and Computer-Integrated Manufacturing 93, 102925.

[23] Sanam, H., Swaminathan, A., Sha, Z., 2025. Computer vision-based in-situ monitoring of cooperative 3d printing in a closed-loop system, in: International Design Engineering Technical Conferences and Computers and Information in Engineering Conference, American Society of Mechanical Engineers. p. V02BT02A041.

[24] Stamatopoulos, M.N., Banerjee, A., Nikolakopoulos, G., 2024. Conflict-free optimal motion planning for parallel aerial 3d printing using multiple uavs. Expert systems with applications 246, 123201.

[25] Stone, R.F., Ebert, M., Zhou, W., Akleman, E., Krishnamurthy, V., Sha, Z., 2024a. Safezone: A topologically-aware voronoi-based framework for fast collision-free cooperative 3d printing, in: International Design Engineering Technical Conferences and Computers and Information in Engineering Conference, American Society of Mechanical Engineers. p. V02AT02A038.

[26] Stone, R.F., Ebert, M., Zhou, W., Akleman, E., Krishnamurthy, V., Sha, Z., 2025. Safezone*: A graph-based and time-optimal cooperative 3d printing framework. Journal of Computing and Information Science in Engineering 25, 061004.

[27] Stone, R.F., Gammage, M., Wang, J., Cullinan, M., Kovar, D., Sha, Z., 2024b. Multi-nozzle cooperation for micro-cold spray .

[28] Stone, R.F., Zhou, W., Akleman, E., Krishnamurthy, V.R., Sha, Z., 2023. Print as a dance duet: Communication strategies for collision-free arm-arm coordination in cooperative 3d printing, in: International Design Engineering Technical Conferences and Computers and Information in Engineering Conference, American Society of Mechanical Engineers. p. V002T02A081.

[29] Zhang, W., Tong, M., Harrison, N.M., 2020. Scanning strategies effect on temperature, residual stress and deformation by multi-laser beam powder bed fusion manufacturing. Additive Manufacturing 36, 101507.