**Proceedings of the ASME 2025 International Design Engineering Technical Conferences & Computers and Information in Engineering Conference**
**IDETC/CIE 2025**
**August 17–20, 2025, Anaheim, CA**

# IDETC/CIE 2025-169037

# MEDIALPART: MEDIAL AXIS-BASED GEOMETRIC PARTITIONING FOR COOPERATIVE 3D PRINTING

**Ronnie F. P. Stone[1], Matthew Ebert[2], Ergun Akleman [3,4], Vinayak Krishnamurthy [2,4], Zhenghui Sha[1,*]**

[1]Walker Department of Mechanical Engineering, University of Texas at Austin, Austin, TX
[2]J. Mike Walker '66 Department of Mechanical Engineering, Texas A&M University, College Station, TX
[3]Department of Visualization, Texas A&M University, College Station, TX
[4]Department of Computer Science and Engineering (by affiliation), College Station, TX

## ABSTRACT

*In this study, we pose the following question: How can we leverage the geometric information of a part to guide the partition of a print job for cooperative 3D printing (C3DP)? To enable a print job to be cooperatively fabricated by multiple robots, traditional partitioning in C3DP has historically been robot-centric, meaning that the print job is divided according to the robot resources (e.g., the number of robots) and their physical constraints (e.g., arm reachability and kinematics). However, often the user of the C3DP system does not know how to leverage the existing resources, as assigning too many robots can lead to diminishing, or even negative returns in terms of certain performance metrics (e.g., makespan). Thus, we are interested in the inverse problem, where the partition is geometry-centric and inherently suggests what the robot resources should be, given a printing layer. Our hypothesis is that there exists valuable information embedded in the geometric and topological structure of the printing layer that can provide a natural way to partition it, and therefore suggest the number of robots needed for the process, including their location in the workspace. This work elicits a significant paradigm shift from robot-centric to geometry-centric C3DP partitioning. We demonstrate that, when considering the medial axis transform (MAT) of the printing layer boundary, we can take advantage of its corresponding radius function to find the most important subset of the medial axis that induces a natural domain decomposition. We achieve this by iteratively constructing a set of vertices in the branches of the medial axis that faithfully capture the important local geometric features of the printing part. The degree of the selected vertices is then used to infer the most natural number of robots to print that immediate vicinity. Finally, given the suggested number of robots, Voronoi sites are optimally sampled on the boundary of the layer, and the final partition is created from the resulting Voronoi tessellation. The proposed framework is layer-wise, as the cross-section of a part may vary significantly both geometrically and topologically along its height. We show a range of numerical results, demonstrating that our methodological framework offers a robust and intuitive way to decompose the geometric domain of each layer for C3DP. Finally, the proposed medial axis-based methodology is process-agnostic and can be used in different multi-robot coverage problems, including with other additive manufacturing technologies.*

**Keywords: Cooperative 3D Printing, Multi-Robot Additive Manufacturing, Medial Axis Transform.**

## 1. INTRODUCTION

Following the widespread adoption of additive manufacturing (AM) technologies in industrial applications, cooperative 3D printing (C3DP) has attracted substantial research efforts in recent years [1–5]. In C3DP, multiple robots collaborate in a shared workspace to print large-scale parts, offering significant benefits in dimensional scalability and manufacturing speed, two common limitations of conventional AM technologies. Unlike traditional 3D printing, which typically relies on a single robot or agent (e.g., single-nozzle gantry), C3DP introduces a range of new challenges due to its collaborative nature. For example, challenges such as part partitioning, manufacturing scheduling, and multi-robot collision avoidance, require specific methodological frameworks to enable truly safe, scalable, and efficient C3DP systems. The focus of this study is on partitioning (i.e., division of labor), which involves segmenting a given part geometry so that the resulting subregions can be distributed among multiple robots during the C3DP process.

Partitioning strategies in C3DP literature have been historically guided by the desired or simply available number and types of printing robots. In this work, we propose a geometry-centric approach, in which we ask a different question, namely how the geometry of the part can guide us in making informed decisions

*Corresponding author: zsha@austin.utexas.edu
Documentation for `asmeconf.cls`: Version 1.40, May 20, 2025.

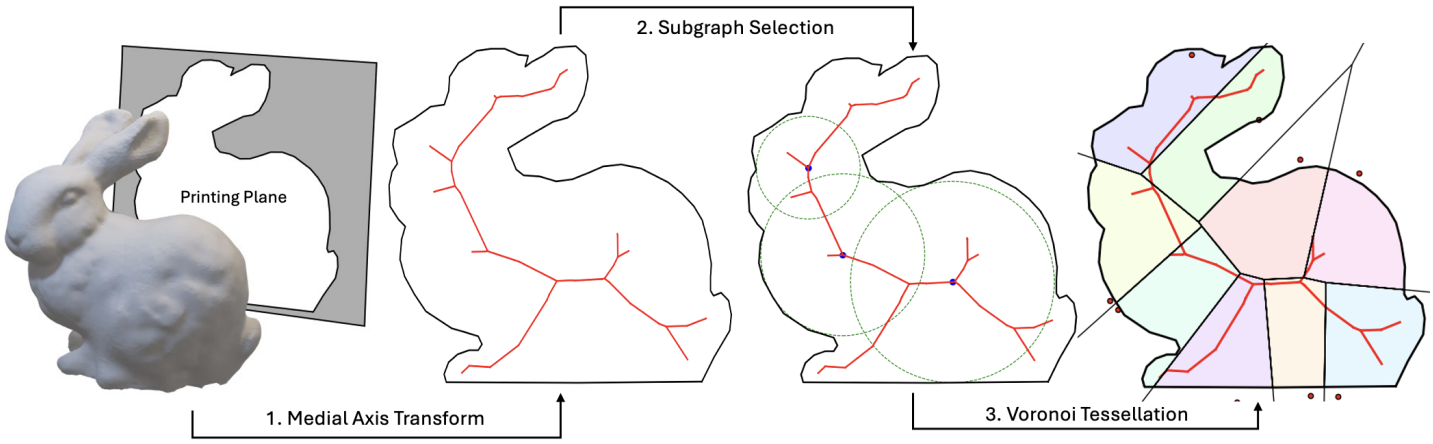


**FIGURE 1: SUMMARY OF THE PROPOSED FRAMEWORK HIGHLIGHTING THE THREE MAIN STEPS: (1) THE PRINTING LAYER ABSTRACTION THROUGH THE MAT (2) THE SELECTION OF THE MOST IMPORTANT VERTICES BASED ON THE RADIUS FUNCTION (3) FINAL VORONOI TESSELLATION ON THE BOUNDARY BASED ON THE SUGGESTED NUMBER OF ROBOTS. NOTE THAT THE PRINTING DIRECTION WAS ARBITRARILY CHOSEN TO HIGHLIGHT AN INTERESTING PLANAR DOMAIN. OUR ALGORITHM IS ORIENTATION-AGNOSTIC.**

on partitioning, robot positioning, and even the number of robots to use in the C3DP process. Such a framework would allow C3DP users to allocate resources more efficiently given arbitrary printing geometries. As a simple example to motivate this, if a design consists of two disjoint parts of similar volume being printed, then a natural partitioning that is cost- and time-optimal may be to have a dual-robot system, where one robot prints one part and the other robot prints the other part. In this scenario, adding a third robot would likely create more problems as collision-avoidance protocols could cause printing delays. Although this simple example is an edge-case scenario in C3DP, it highlights how the geometry of a design may aid us in making decisions on partitioning and, therefore, the required robot resources and placement. Therefore, the goal is to be able to answer this question for any printing layer.

As previously discussed, prior studies in this area focus on robot-centric partitioning schemes, where the division of labor is guided by the number of robots and their respective kinematic constraints, thus simplifying the problem so that the part geometry becomes secondary or is only implicitly considered. However, if we are to leverage part geometry, we need a general mathematical descriptor of shapes. In this work, we show that the medial axis transform (MAT) of the part geometry [6] is a valuable universal descriptor for C3DP due to its ability to capture important local geometric information of the part. The medial axis can be qualitatively understood as the geometric skeleton of a shape, abstracting the geometry by capturing its symmetries, and allowing us to leverage its properties for partitioning applications.

The key idea is to find a subset of the medial axis that best approximates and covers the part, which can then be used to determine the number and position of robots. Our framework finds this subset by ranking the vertices of the medial axis through its radius function, which is computed during the MAT and represents the radius of the maximally inscribed circle (restricted by the boundary) centered at that vertex. The final number of robots $N$ is then determined by assessing the degree (i.e., number of connected edges) of the selected vertices based on the connectivity of the skeleton. Once robot resources are assigned, a final partition is found by optimally sampling $N$ Voronoi sites on the boundary of the part, approximating the important subset of the medial axis (Fig. 1). The sampling of the sites is optimal in that the final division of labor is as even as possible (area-balanced).

As our proposed framework leverages the MAT in the context of planar domains, it is important to emphasize that the approach is layer-wise, meaning that the partitioning process considers one layer at a time. Therefore, the resulting partition for one layer can vary significantly from subsequent layers depending on the part geometry. While this may seem restrictive from a manufacturing perspective, layer-wise C3DP strategies are aligned with recent advances in the literature, which move away from the volumetric (e.g., chunking) approaches. One of the main advantages of layer-wise C3DP is the better control of the manufacturing process itself, which can lead to improved mechanical integrity (e.g., via topological interlocking) and time-optimal makespan. Thus, while our proposed framework requires no additional considerations for geometries that are approximately 2.5D, it may require further consideration from the user in terms of robot resources when printing geometries with drastic variations along the printing direction. For example, our framework may suggest 3 printing robots for the lower layers and then 5 for the upper layers, as the part geometry changes. Hence, similar to other layer-wise approaches in the literature, we leave this implementation detail to the user, who is aware of the limitations of their manufacturing systems.

As such, the key contributions in this paper are as follows:

- We elicit a paradigm shift in C3DP partitioning strategies, from robot- to geometry-centric, by asking the question of how to allocate the printing resources based on the part.
- We propose a geometry-centric framework for C3DP partitioning and robot resource allocation, using the medial axis transform to determine natural domain decomposition at each 2D layer of a print.

- With the suggested number of robots from the partitioning step, we demonstrate how to optimally sample Voronoi sites on the boundary of the part, leading to a Voronoi tessellation that minimizes differences in the amount of work between robots, while preserving the important features of the medial axis, thus preventing reachability issues.

The remainder of the paper is organized as follows. In Section 2, we review related work in the literature, including multi-robot applications beyond C3DP, and highlight the research gap that geometry-centric partitioning bridges. The conceptual framework and the mathematical formulation for the problem are presented in Section 3. Our proposed methodological framework, assuming only a boundary representation of the layer as input, is outlined in Section 4. Then, a set of numerical results for different geometries (e.g., Benchy model) highlighting the robustness and intuitiveness of our framework is provided in Section 5. Finally, we discuss limitations and future research directions in Section 6.

## 2. RELATED WORK

There has been significant interest in decomposing a shape into a simpler representation in order to either simplify geometric computations or extract important features and information. The medial axis has been used to represent shapes using a set of lines for 2D shapes or surfaces for 3D solids [7]. Formally, the medial axis of a shape is the set of interior points that have two or more equidistant and nearest disjoint set of points on the boundary.

Other methods for shape representation are, for example, using the convex hull of a shape, which is the smallest convex shape that completely encompasses the geometry [8]. The convex-hull method works well in approximating nearly convex shapes, but would fail for shapes with large or several concave regions. As an alternative, the alpha shape, building on the convex hull method [9, 10], can create approximations of shapes at various levels of detail. While these methods can describe shapes with different degrees of accuracy and representations, they do not inherently capture critical geometric information such as local symmetries, thicknesses, and curvature. In fact, taking the convex hull as an example, the algorithm was designed to purposefully ignores those traits.

Current work on partitioning for C3DP uses various methods for segmentation, which vary depending on the desired objective (e.g., mechanical integrity). Some approaches focus on using interlocking subvolumes, which are cellular structures that vary across the printing direction [11, 12]. Others focus on partitioning a part into sloped subvolumes, also called chunks, which can be printed separately in a non-layer synchronized fashion [13–15]. For layer-wise C3DP, it is also possible to segment each layer geometry using Voronoi sites as robot positions to minimize the makespan of the printing process by co-considering work distribution and manufacturing schedule [16, 17]. However, these methods are robot-centric in that the partition process and outcomes are constrained by robot resources and kinematics, instead of treating the properties of the part geometry as the driving force for partition. Consequently, they cannot answer the question of optimal allocation of resources without testing for every possible combination of multi-robot setup. While this is theoretically possible, it is not computationally feasible as the C3DP process is, in its most general form, a multi-level optimization problem with a very large design space.

One area of research that is similar to the work done here is mesh refinement for the purpose of finite element analysis (FEA) or computational fluid dynamics (CFD) analysis. Specifically in 2D structures, the medial axis of a shape can be very influential and helpful in creating a mesh, capable of accurately modeling shape behavior [18, 19]. This has been extended to the meshing of general solids, and the benefit is that the medial axis can capture general shape properties [20].

In order to gain a better understanding of geometry-centric partitioning, we also explore how works outside of C3DP have used part geometry for segmentation or navigation in multi-robot applications, which are closely related to our use-case. One example of segmentation for 3D objects is part-based segmentation where a 3D object is decomposed into subvolumes [21]. Other methods also exist such as co-segmentation [22]. In these cases, segmentation may not be possible given robotic constraints, such as if a subvolume is completely enclosed by another subvolumes additional schedule constraints must be considered. An example of using geometric features to make informed decisions in a different robotics application is in path planning, where robots must traverse a cluttered environment [23]. In these cases, the robot must move between points while avoiding obstacles. One common method to achieve this is to use the information from the medial axis, since it will automatically keep robots as far away from obstacles (boundary of a shape) as possible, as well as automatically providing a path for robots to follow [24–26]. It can also be combined with the data collection of environments in real time to create trees for exploration and traversal [27].

There currently exists a research gap in geometry-centric partitioning approaches for C3DP applications. Our review shows that geometry-centric is understudied in the literature and that a medial axis-based approach can be beneficial for C3DP as it has been in other domains, as it provides a simple, yet powerful abstraction of the printing geometry, eliciting a wealth of information such as thickness, curvature, and structural changes, ultimately allowing us to answer fundamental questions about resource allocation in C3DP.

## 3. CONCEPTUAL FRAMEWORK

The fundamental question we explore in this work is whether the geometry of a printing part can provide us with hints on how to partition it for cooperative 3D printing (C3DP). This is a significant paradigm shift from traditional robot-centric partitioning, where the number of robots and their kinematic properties are the driving forces. As such, the first challenge is to select a universal geometric property of domains that can guide us in the partitioning process.

### 3.1. The Medial Axis Transform

Based on our experience and literature analysis, it is evident that the medial axis [6] of a domain is an excellent tool for testing our hypothesis. Before discussing the reasons for its selection, we go over its formal definition and intuition behind the way in

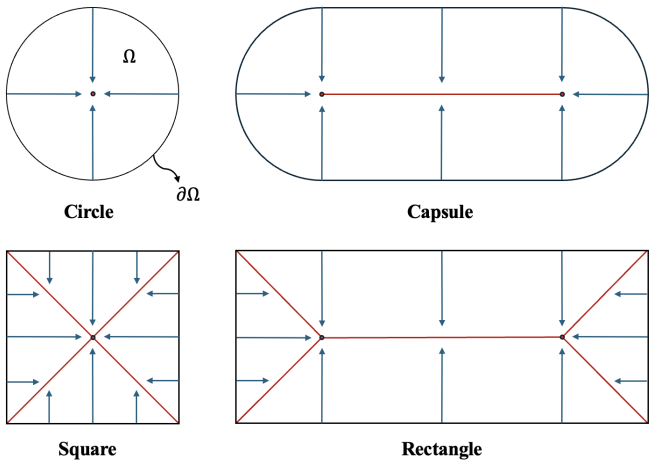


**FIGURE 2: MEDIAL AXIS OF SELECTED PRIMITIVE SHAPES AND GRASSFIRE ANALOGY VISUALIZATION, SHOWING AN INTUITIVE WAY OF THINKING ABOUT THE RADIUS FUNCTION.**

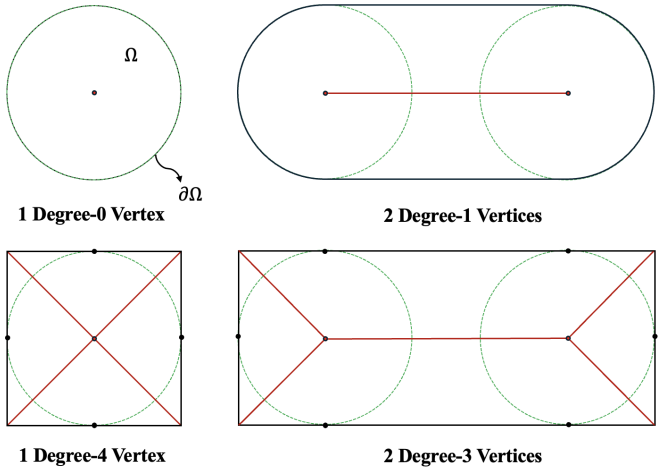


**FIGURE 3: VERTEX CLASSIFICATION DEPENDING ON NUMBER OF DISJOINT CONTACT POINTS BETWEEN INSCRIBED MAXIMAL CIRCLE AND THE BOUNDARY $\partial\Omega$. VERTICES THAT HAVE DEGREE $\geq 3$ ARE CALLED BRANCHING POINTS.**

which it is constructed. Since we are mainly interested in layer-wise printing, the discussion of the medial axis transform (MAT) is limited to planar domains.

Any compact planar domain $\Omega$ has an associated MAT [28]. The MAT of $\Omega$ provides us with its medial axis $\mathcal{M}(\Omega)$, given by a set of points, and a radius function $r$ that maps each point in the medial axis to the minimum distance (in a given metric) to the boundary of the domain $\partial\Omega$. Thus, the medial axis of a domain can be denoted as:

$$\mathcal{M}(\Omega) = \{x \in \Omega \mid \exists y_1, y_2 \in \partial\Omega,\ d(x, y_1) = d(x, y_2)\}, \quad (1)$$

where $y_1 \neq y_2$ and $d(x, y)$ is some metric, such as the Euclidean distance (i.e., L2-norm). Consequently, each point on the medial axis is at the same distance from at least two points on the boundary of the domain.

There exist several equivalent definitions for the medial axis, each highlighting different aspects of how the transform abstracts the domain. An intuitive definition is to think of the medial axis through the lens of the grassfire analogy. From this perspective, the medial axis can be seen as the set of points at which the fire wavefront would meet if the entire boundary of the domain were set on fire simultaneously (see Fig. 2). This analogy is useful in understanding and computing the radius function $r$, which is also provided by the MAT, and formally defined as:

$$r(x) = \min_{y \in \partial\Omega} d(x, y), \quad (2)$$

where, assuming that $d(x, y)$ is the L2-norm, the image of $r(x)$ is the set of radii of maximally inscribed circles centered at $x$. Thus, the radius function provides a measure of the local thickness of the part geometry. Hence, by combining the definitions in Eq. 1 and Eq. 2, we can formally define MAT as:

$$\text{MAT}(\Omega) = \{(x, r(x)) \mid x \in \mathcal{M}(\Omega)\}, \quad (3)$$

Given the formal definition of the MAT in Eq. 3, the question remains: why is it a useful descriptor for C3DP partitioning? The key is that MAT serves as a shape descriptor that encodes local symmetry, thickness, curvature, and complexity, all of which are critical geometric features of any planar domain $\Omega$. Variations in these properties may indicate significant structural changes in the part, naturally suggesting domain decomposition. Following this idea, we leverage information on the branches of MAT to make informed decisions in geometry-centric C3DP partitioning.

### 3.2. Medial Axis Branches

The medial axis can be mathematically represented as a geometric graph $G$ where the degree of each vertex is a positive integer. The branching points (BPs) of the medial axis are the vertices of $V$ that have a degree greater than 2. Geometrically, the degree of the vertices can be related to the radius function $r$, in the sense that the degree of a vertex $v \in V$ is equal to the number of disjoint set of points in $\partial\Omega$ that the maximally inscribed circle centered on $v$ touches (Fig. 3). In the vast majority of cases, the graph $G$ of the medial axis only consists of vertices with degrees lower than or equal to 3. In fact, designing a non-primitive shape that has a vertex of degree higher than 3 in the geometric graph of its medial axis is not trivial. Therefore, in this study, we provide examples and results considering the lower-degree cases (i.e., $\deg(V) \leq 3$). However, this is not a limiting factor and the proposed framework works with branching points of any degree.

As discussed, the branches of the medial axis provide us with a wealth of information on the local geometry of the part and may suggest a natural way to decompose the domain. Physically, using the medial axis as a guide for partitioning is also beneficial, as they can suggest subdomains that are more reachable to robots. This is a direct consequence of how the medial axis captures the local symmetries in the part. However, not all branches of the medial axis are equally important. If we are to use branches as a means to guide us in geometry-centric partitioning, we must be able to identify which are the most relevant for C3DP applications. This

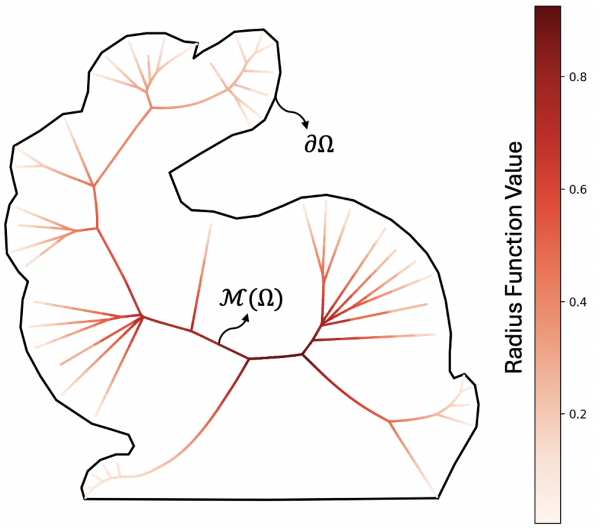


**FIGURE 4: MEDIAL AXIS APPROXIMATION FOR THE STANFORD BUNNY. THE RESULTING SKELETON CONTAINS MULTIPLE BRANCHES.**

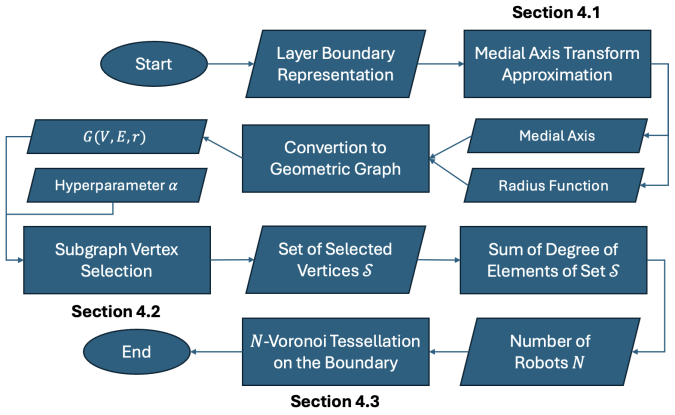


**FIGURE 5: FLOWCHART FOR THE PROPOSED METHODOLOGICAL FRAMEWORK. FROM A LAYER BOUNDARY REPRESENTATION INPUT, WE PROVIDE A NATURAL AND BALANCED FINAL TESSELLATION OF THE DOMAIN.**

is particularly important since generic shapes may have hundreds of branches. Furthermore, given that MAT is generally never computed analytically, numerical approximations may induce the creation of more branches than expected.

To illustrate this point, consider, as an example, a generic planar domain $\Omega$, (e.g., Stanford Bunny cross section [29] shown in Fig. 4). If we reasonably approximate its MAT, we can see that it has many branches with vertices of varying degrees. As such, it is not immediately clear which subgraph of $G$ is most indicative for domain decomposition and, additionally, how we can leverage that information to suggest a reasonable number of robots for partitioning. Thus, our goal is to design a methodological framework that can tackle both of these problems simultaneously. In this paper, we show how the radius function can provide us with hints in answering these questions.

## 4. METHODOLOGY

In this section, we leverage the conceptual framework discussed in the previous section, and provide a complete pipeline for medial axis-based C3DP partitioning (Fig. 5). All the user must provide is the boundary description of the printing layer in the form of any computational design file (e.g., STL). Our methodology provides the best subset (subgraph) of the medial axis to approximate the shape for C3DP partitioning, the suggested number of robots $N$ to be used, their approximate locations, as well as a final balanced (area-wise) partitioning based on $N$.

### 4.1. MAT Approximation and Graph Construction

Given a compact planar domain $\Omega$ in the form of a printing layer, we can approximate its MAT by sampling points along its boundary $\partial\Omega$ and using them as a set of Voronoi sites [28]. From the resulting Voronoi diagram, the MAT is simply the union of the finite ridges (shared edges between Voronoi cells) of the diagram that are completely contained in $\Omega$. In fact, many algorithms leverage the relationship between Voronoi diagrams and MATs to efficiently compute MAT for certain planar domains, such as simple polygons [30]. However, in this work, for simplicity, we evenly distribute the Voronoi sites along the boundary representation of the printing layer. Consequently, the quality of the approximation is determined by the number of Voronoi sites sampled and the complexity of the boundary representation. An example of this process is shown in Fig. 6 for different sampling resolutions.

Following the MAT approximation, we computationally represent it as a graph structure $G_{MAT} = (V, E, r)$, where the edges $E$ are the selected Voronoi ridges, the vertices $V$ are their associated endpoints, and $r\colon V \to \mathbb{R}^+$ is the radius function, assigning each vertex $v \in V$ to the radius of the maximal inscribed ball centered at $v$ and bounded by $\partial\Omega$. The rationale for using a graph representation is that it allows us to more easily manipulate the medial axis, such as the degree of the vertices and their associated radius values. Finally, the next step is to find a representative subgraph of the medial axis by iteratively selecting a subset of points on the medial axis.

### 4.2. Subgraph Selection and Number of Robots

This step aims to find a good subgraph of the medial axis so that we can infer the appropriate number of robots for the partition. We do this through an iterative vertex selection process, in which the vertices of the medial axis are ranked according to a ranking function $f(v)$.

We first define some important terms. Let $\mathcal{O}$ be the set of open vertices in $G$, that is, the vertices that have not yet been selected. Furthermore, let $\mathcal{S}$ be the set of selected vertices. At the beginning of the iterative process, $\mathcal{O} = V$, which is the set of all vertices of $G$ and $\mathcal{S} = \emptyset$, which is the empty set. Thus, the ranking function $f(v)$ at any step can be formally defined as:

$$f(v) = |A(v)| - |A(v) \cap \bigcup_{u \in \mathcal{S}} A(u)|, \tag{4}$$

where $|A(v)|$ is the scalar value of the area of the maximally inscribed circle at vertex $v$, and the second term represents the scalar value of the area given by the intersection of the circle at

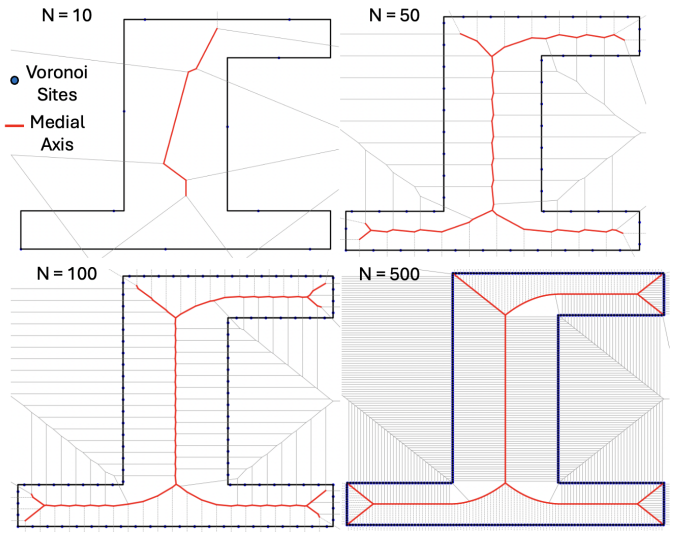


**FIGURE 6: MEDIAL AXIS TRANSFORM APPROXIMATION VIA VORONOI SITE SAMPLING ON THE BOUNDARY. THE FINITE RIDGES CONTAINED IN $\partial\Omega$ MAKE UP THE SHAPE SKELETON.**

vertex $v$ with the union of all circles of the vertices that were already selected $u \in \mathcal{S}$.

Our ranking function is, therefore, defined by these two terms that guide the vertex selection based on two principles (Fig. 7), respectively:

1. Vertices with higher radius function values are preferred as they represent regions of the medial axis with higher local thickness, suggesting that those local geometries are more important for the decomposition.

2. Vertices with inscribed circles that are minimally intersecting those from the set of selected vertices $\mathcal{S}$ are preferred, as it allows for a better coverage of the global planar domain, capturing all important geometric features.

This iterative process continues until a user-defined threshold is met or a maximum number of available robots is reached, which could be limited by the user. Regarding the numerical threshold, we define a hyperparameter $\alpha$ that is the ratio between the values of the ranking function for the selected node in the $k^{th}$ iteration and the first iteration:

$$\alpha = \frac{f(v_k)}{f(v_1)}, \tag{5}$$

where the vertices $v_1$ and $v_k$ are the optimal selections for the first and the current step at $k$, respectively. Thus, if the value of $\alpha$ is smaller than a user-defined threshold, the algorithm stops. Intuitively, this hyperparameter informs that when the selection of another vertex in the medial axis is not significant in the coverage of the printing layer, thus suggesting that the most natural number of robots for partitioning has already been found. In case multiple vertices have the same value for the ranking function, then the one with the highest vertex degree is selected. If the tie still persists, then a random selection between the candidates takes place for that iteration.

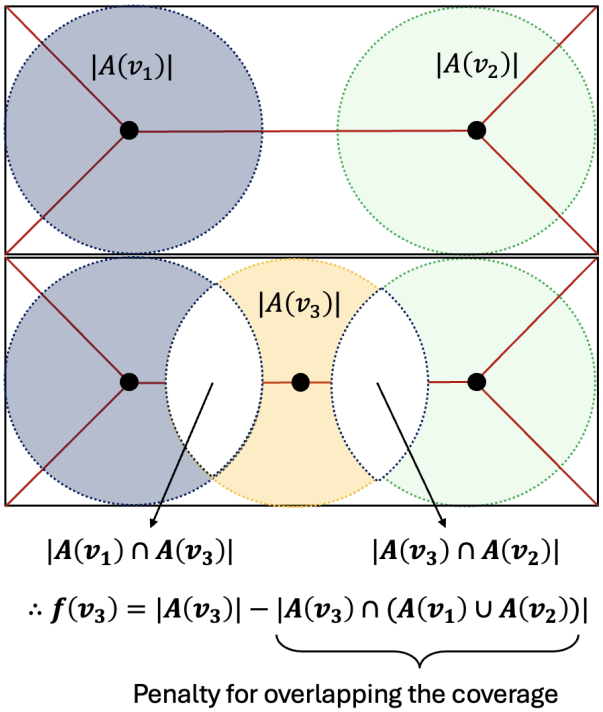


**FIGURE 7: OVERVIEW OF THE TWO GUIDING PRINCIPLES IN THE VERTEX SELECTION. THE FIRST TERM REFERS TO THE AREA OF THE MAXIMALLY INSCRIBED CIRCLE. THE SECOND TERM PROVIDES A PENALTY FOR OVERLAPPING COVERAGE.**

In general, the selection of $\alpha$ is left to the user. Our numerical results have shown that values between 0.1 and 0.3 work well for the decomposition. A high value of $\alpha$ can be used when the user would like to limit the number of printing robots, thus creating a decomposition with fewer partitions. Due to the computational simplicity of the algorithm, different values of $\alpha$ may be easily tested to achieve the desired results. After the process terminates, the selected number of robots is then simply given by the sum of the degrees of the vertices contained in $\mathcal{S}$:

$$N = \sum_{v \in \mathcal{S}} \deg(v), \tag{6}$$

thus capturing the effect of branching points, if any are selected. The rationale is that, at least locally, the degree of the branching point will infer a natural partition for the domain. If the degree is 3, then there is a natural 3-way area decomposition, suggesting that 3 printing robots could be assigned locally. Recall that the vertex degree is defined to be the number of edges of the medial axis connected to that vertex, as in the general graph-theoretical sense.

### 4.3. Voronoi Tessellation on Boundary

Once the number of robots is determined by Equation (6), thus implicitly covering a selected subgraph of the medial axis, we can move forward with the final partition given any defined objective function.

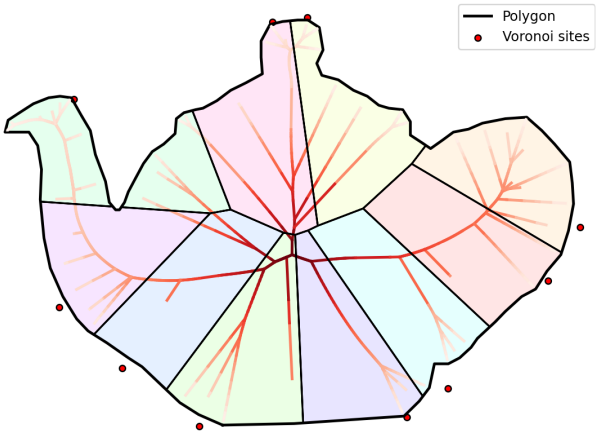


**FIGURE 8: SAMPLE BALANCED PARTITION USING AN ARBITRARY NUMBER OF VORONOI SITES (10) ON THE BOUNDARY. NOTE HOW IT TENDS TO APPROXIMATE THE IMPORTANT BRANCHES OF THE MEDIAL AXIS.**

To optimally position the Voronoi sites along the boundary $\partial\Omega$, we choose the objective of minimizing the area imbalance among the robots. Thus, if we let the areas of the final subregions be denoted $B$, then the objective function can be defined as:

$$\min_{B_1,B_2,\ldots,B_N} \sum_{m=1}^{N} \sum_{n=m+1}^{N} (B_m - B_n)^2 \tag{7}$$

thus minimizing the pairwise difference in area between all subregions using the L2 norm. Naturally, other difference metrics can be used, although in practice we have not observed significant variations in the results. Balancing the partition areas is a standard objective in C3DP applications, as it generally relates to a lower makespan by maximizing parallel printing. Importantly, by sampling Voronoi sites on the boundary, we implicitly approximate the medial axis (Fig. 8). The difference is that, by selecting an appropriate number of robots given the most important subgraph of $G$, the chosen Voronoi sites will probably approximate the chosen subgraph of the medial axis, as originally intended in the conceptual framework. By doing so, it will also minimize any reachability issues for future robot positioning due to the symmetric properties of the medial axis.

The optimization itself is done via a genetic algorithm (GA) approach, as metaheuristic optimizers are suited to such discontinuous and discrete applications. The gene representation is given by a set of indices from the geometric graph $G$. For instance, the gene representation for a 4-robot scenario can be given by:

$$I = [i_1, i_2, i_3, i_4], \tag{8}$$

where $i_j$ is a positive integer and $j$ is strictly less than the size of $\mathcal{S}$. However, the choice of GA is not unique, and other metaheuristic optimizers may be used, such as particle swarm optimization (PSO) or simulated annealing (SA). However, we have shown in prior works that GA behaves well in terms of convergence and sub-optimality factors for similar applications [17].

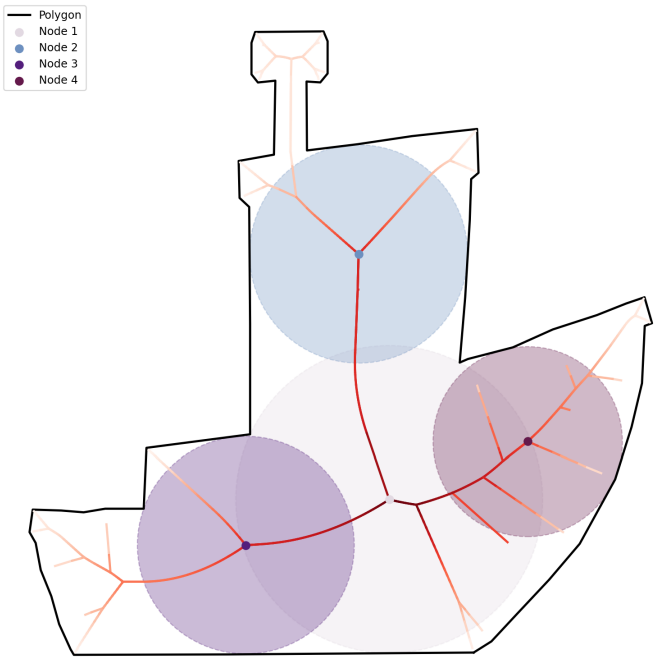


**FIGURE 9: ITERATIVE POINT SELECTION PROCESS FOR THE BENCHY MODEL. THE ALGORITHM TERMINATES AFTER FOUR ITERATIONS DUE TO THE NUMERICAL THRESHOLD STOP CONDITION. THE VALUE OF ALPHA USED IS 0.2. THE SUM OF THE DEGREE OF BRANCHES IS 12, SUGGESTING A NATURAL DECOMPOSITION USING 12 ROBOTS.**

Note that, compared to robot-centric partitioning work, this optimization happens in a design space that is one dimension lower than the robot workspace. That is, instead of looking for Voronoi sites in a 2D space, we are looking along the boundary, which can be considered a 1D artifact. This is a direct consequence of using the medial axis, as there is a one-to-one mapping from the medial axis to the boundary of the printing layer. This reduction in the dimensionality of the design space, while also giving the design space geometric meaning, provides much faster convergence compared to other methods.

## 5. EXPERIMENTS AND RESULTS

For the numerical studies validating our framework, we consider a number of different layer geometries. Using a cross section of the famous Benchy model, we demonstrate our methodology by showing the entire process step by step, highlighting the vertex selection process and how the ranking function behaves in each iteration. Additionally, we show the final tessellation results for cross sections of the Stanford Bunny, Horse, Lizard, and Triangle models, in order to display robustness and generality. Note that the printing directions were arbitrarily chosen to highlight interesting planar domains. Our algorithm does not select a printing orientation, as it is orientation-agnostic, and thus the decision is left to the user.

### 5.1. Benchy Model Case Study

Following the methodology, the first step is to get the MAT approximation of the boundary using the Voronoi site approx-

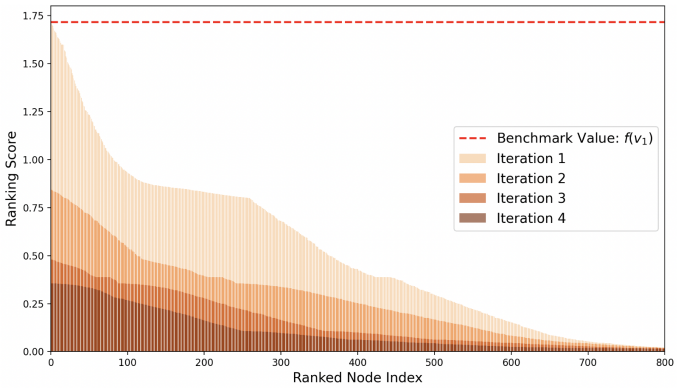


**FIGURE 10: RANKING FUNCTION VALUE FOR ALL CONSIDERED NODES DURING THE ITERATIVE PROCESS. NOTE HOW THE LOCAL PEAKS ARE FLATTENED, SHOWING THAT ALL CRITICAL LOCAL GEOMETRIC INFORMATION IS CAPTURED.**

imation (see Section 4.1). Then, after representing the medial axis as a geometric graph $G$, we run it through the iterative vertex selection process (see Section 4.2) to find a set of vertices that are representative of the most important subgraph of $G$.

Our algorithm leads the Benchy model to converge to 4 degree-3 vertices and thus 12 robots (Eq. 6). Each iteration shows how the algorithm progressively explores the most important subgraph of $G$, capturing the local geometry information to determine a natural domain decomposition (Fig. 9). It is also interesting to see how the ranking function behaves as the iterations progress. We see that, by plotting the vertices and their corresponding ranking function values, the distribution tends to become flatter, showing that the most critical local geometric features have already been captured (Fig. 10).

After selecting the number of robots from the iterative process, we proceed to the Voronoi tessellation using sites on the boundary $\partial\Omega$ of the Benchy model. The final result shows how our algorithm captured the most important branches of the medial axis of the printing part's geometry, thus suggesting a natural number of robots to decompose the domain and approximate the most important subgraph of the medial axis (see Fig. 11).

### 5.2. Generalizability

We show the final tessellations for four distinct geometries, i.e., the cross-section of a Horse, Bunny, Lizard, and Triangle models. We see how the algorithm found a natural way to decompose all of those complex geometries while also offering a physically viable partition for robotic arm-enabled C3DP. The Horse geometry converged to 15 robots total due to its many important branches, while the Bunny and the Triangle geometry converged to 9 robots (Fig. 12). The Lizard model converged to 37 robots due the almost zero variation in radius thickness along the main branch. This result is expected as the medial axis carries no notion of absolute, but only relative scale (see Section 6.2).

The positioning of the robots can be easily inferred from the location of the Voronoi sites on the boundary. However, note that there does not need to be a one-to-one association between the site locations and robot positions. In fact, in previous works, we have

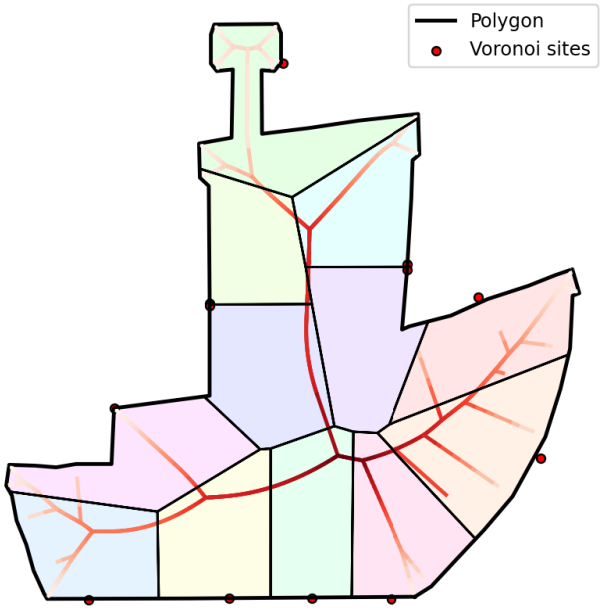


**FIGURE 11: FINAL BENCHY MODEL TESSELLATION USING AN $\alpha$ VALUE OF 0.2. WE SEE HOW THE SUBSEQUENT OPTIMAL PLACEMENT OF THE VORONOI SITES ON THE BOUNDARY PROVIDE A VERY GOOD APPROXIMATION OF THE MOST IMPORTANT SUBGRAPHS OF $G$. THE SUGGESTED NUMBER OF ROBOTS (12) FOR DOMAIN DECOMPOSITION WAS DIRECTLY INFERRED FROM OUR FRAMEWORK.**

provided methods to map the Voronoi sites into viable printer positions (e.g., considering physical constraints) by leveraging the notion of virtual robot positions (see [17]). In this work, we leave this decision to the user, as it may be influenced by specific robot constraints of the C3DP system being used.

It is important to emphasize that not all geometries have a unique way to naturally decompose their domain, and thus, there may be multiple good solutions with different numbers of robots. In fact, this is captured by the choice of the value of $\alpha$. If $\alpha$ is high, the algorithm is more likely to favor domain decompositions for fewer robots, whereas if $\alpha$ is low, a solution with tens of robots may be found instead. The choice of $\alpha$ may also be determined by users, as they are aware of their own manufacturing environment and whether there are enough resources to decompose the domain into multiple sub-regions. Further analysis of the impact of $\alpha$ in the algorithm and how it may impact different geometries (e.g., with convexity measures) may be warranted in future work.

## 6. DISCUSSION

In this section, we discuss topics of interest derived from our analysis of the numerical results. We discuss how robust the framework is against small perturbations on the boundary $\partial\Omega$, which could be due to different computational representations (e.g. STL) of the same printing part. Additionally, we talk about certain classes of geometries in which the algorithm may suggest a large number of robots due to the inherent notion of relative scale embedded in the framework.

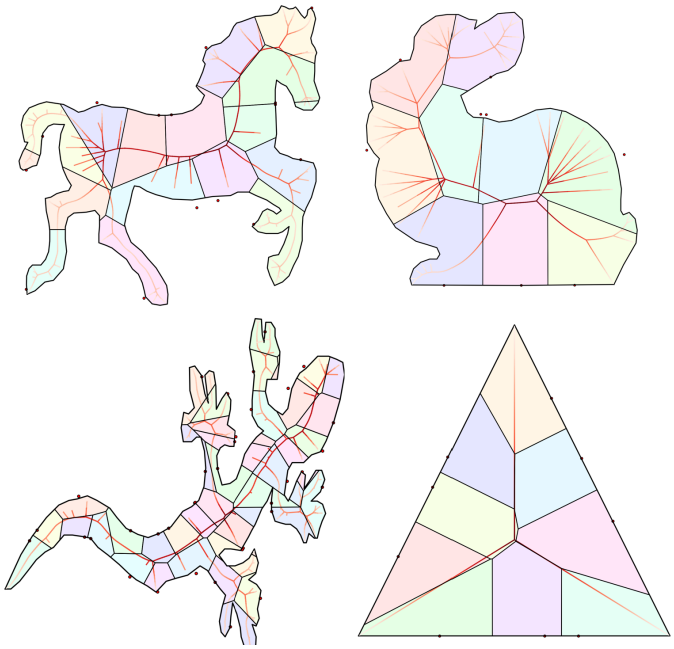

**FIGURE 12: FINAL VORONOI TESSELLATION FOR OTHER COMPLEX GEOMETRIES. THE HORSE, BUNNY, LIZARD, AND TRIANGLE GEOMETRIES USE 15, 9, 37, AND 9 ROBOTS RESPECTIVELY. THIS VARIATION IS EXPECTED AS THE MEDIAL AXIS HAS NO NOTION OF SCALE.**

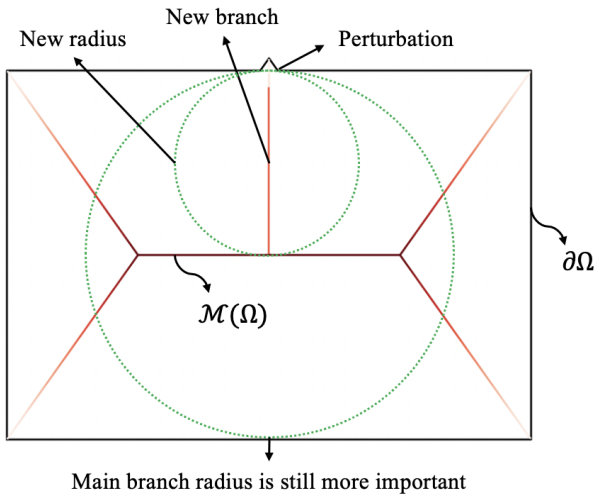


**FIGURE 13: ALTHOUGH THE INTERIOR MAT IS INHERENTLY UNSTABLE DUE TO ITS SENSITIVITY TO PERTURBATIONS IN THE BOUNDARY $\partial\Omega$, OUR FRAMEWORK IS ROBUST TO THESE VARIATIONS AS THE RADIUS FUNCTION WILL NEVER BE HIGHER IN THE NEW BRANCHES WHEN COMPARED TO THE MAIN BRANCH.**

### 6.1. Stability of MAT

It is a well-known fact that MAT is highly sensitive to small perturbations at the boundary of a domain $\partial\Omega$. However, our algorithm should be relatively robust against these perturbations. This is because, although the medial axis may change significantly due to the addition of new branches, the selection of the vertices will be biased towards the branch most representative of the shape itself, which will still be constructed even in the presence of perturbations.

For instance, consider a rectangular geometry with a small perturbation on the boundary. This perturbation will create a new branch, causing a significant change in the resulting medial axis (Fig. 13). However, due to the nature of the ranking function (Eq. 4), it will not modify the vertex selection process because, fundamentally, the radius function along this anomalous branch will not have a higher radius when compared to the main branch. Thus, it will be very unlikely to be considered unless the choice of $\alpha$ is very small and, therefore, impractical for C3DP partitioning applications.

### 6.2. Geometrical Generalizability

Our framework works for any geometry, as it is always possible to extract the interior MAT of a planar domain $\Omega$ regardless of its topology (e.g., non genus-0 layers). However, whether the resulting domain decomposition and, consequently, the suggested number of robots is practical for C3DP, irrespective of the layer geometry, is an important topic of discussion. To answer this, we observe that there is a notion of relative scale embedded in the algorithm, relating the reachability of the printing robots to the maximum value of the radius function. Thus, this scale informs us on the size of the printing robots that should be used, relative to the size of the part. This relationship becomes clear when we analyze a specific class of geometries. This class of geometries can be described by having very small variations in the radius function $r$, that is, the derivative of the radius function along the branches of the medial axis is close to zero. These are mostly slender parts with high aspect ratios.

Consider, for instance, a skinny and long z-beam geometry (Fig. 14). When running this geometry through our proposed methodological framework, the suggested number of robots is large. This can be interpreted in two ways. The first is that the geometry itself does not provide any significant clues on the number of robots that should be used. The second is that our framework inherently applies this notion of relative scale, which relates the size of the part to the size of the robots. Thus, the framework considers implicitly that the robots have reachability close to the maximum of the image of the radius function. In this way, if we put two robots on opposing sides of a medial axis segment, they are able to cover the printing space. This is a natural consequence of a medial axis-based approach, as it is geometry-centric and does not have a notion of absolute scale. One can argue that this provides a good suggestion for the scale of the robots as well, purely based on the geometry of the part. Thus, the proposed framework provides an excellent decomposition, provided that the printing robots can be scaled relative to the thickness of part, which is generally the case in C3DP applications anyway.

Exploring these scaling aspects is something that we would like to address in future research. However, a straightforward modification for such geometries, in case the robots cannot be scaled down, would be to use the convex hull of the original shape, when possible, to avoid issues with sampling points in the

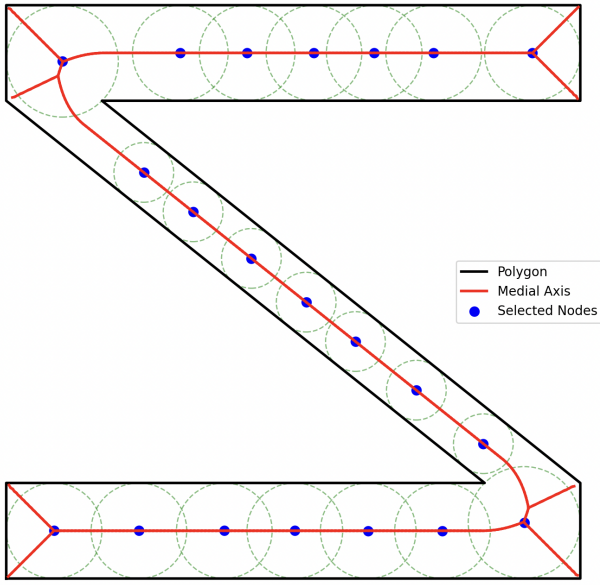


**FIGURE 14: SLENDER GEOMETRIES MAY SUGGEST NO INHERENT BENEFIT FROM LOOKING AT THE GEOMETRY OF THE PRINTING PART. IN THIS Z-BEAM PRINTING LAYER, THE DERIVATIVE OF THE RADIUS FUNCTION IS MOSTLY NEAR ZERO, INDICATING NO VARIATION IN LOCAL THICKNESS.**

iterative algorithm. Alternatively, one could modify the ranking function to use a different coverage metric. It is clear that in these edge cases, the union of intersections is not very indicative of finding critical vertices of the medial axis since most of them have equal values.

## 7. CONCLUSION

To conclude, we presented a medial axis transform-based partitioning approach for C3DP. Our framework represents a novel shift from the robot-centric partitioning paradigm, where the focus is on the number of robots available and their properties, such as mobility and kinematics. The key methodological contribution in this work is the use of the MAT to implicitly find an efficient number of robots to cover the main subgraph of the medial axis. The selected number of robots is then used to optimally sample Voronoi sites on the boundary representation of the printing layer, such that the area imbalance among robots is minimized. In the future, we plan on further exploring the implications of geometry-centric C3DP partitioning, and potentially integrating them with robot constraints.

## ACKNOWLEDGMENTS

This work was partially supported by the Army Research Lab (ARL) via grant number W911NF-24-2-0007. Any opinions, findings, conclusions, or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of ARL.

## REFERENCES

[1] Stone, Ronnie FP, Zhou, Wenchao, Akleman, Ergun, Krishnamurthy, Vinayak Raman and Sha, Zhenghui. "Print as a dance duet: Communication strategies for collision-free arm-arm coordination in cooperative 3d printing." *International Design Engineering Technical Conferences and Computers and Information in Engineering Conference*, Vol. 87295: p. V002T02A081. 2023. American Society of Mechanical Engineers.

[2] Poudel, Laxmi, Marques, Lucas Galvan, Williams, Robert Austin, Hyden, Zachary, Guerra, Pablo, Fowler, Oliver Luke, Sha, Zhenghui and Zhou, Wenchao. "Toward Swarm Manufacturing: Architecting a Cooperative 3D Printing System." *Journal of Manufacturing Science and Engineering* Vol. 144 No. 8 (2022): p. 081004.

[3] Alhijaily, Abdullah, Kilic, Zekai Murat and Bartolo, AN Paulo. "Teams of robots in additive manufacturing: a review." *Virtual and Physical Prototyping* Vol. 18 No. 1 (2023): p. e2162929.

[4] Khosravani, Mohammad Reza and Haghighi, Azadeh. "Large-scale automated additive construction: overview, robotic solutions, sustainability, and future prospect." *Sustainability* Vol. 14 No. 15 (2022): p. 9782.

[5] Rescsanski, Sean, Hebert, Rainer, Haghighi, Azadeh, Tang, Jiong and Imani, Farhad. "Towards intelligent cooperative robotics in additive manufacturing: Past, present, and future." *Robotics and Computer-Integrated Manufacturing* Vol. 93 (2025): p. 102925.

[6] Blum, Harry. "A transformation for extracting new descriptions of shape." *Models for the perception of speech and visual form* (1967): pp. 362–380.

[7] Sherbrooke, Evan C, Patrikalakis, Nicholas M and Brisson, Erik. "An algorithm for the medial axis transform of 3D polyhedral solids." *IEEE transactions on visualization and computer graphics* Vol. 2 No. 1 (1996): pp. 44–61.

[8] Barber, C Bradford, Dobkin, David P and Huhdanpaa, Hannu. "The quickhull algorithm for convex hulls." *ACM Transactions on Mathematical Software (TOMS)* Vol. 22 No. 4 (1996): pp. 469–483.

[9] Akkiraju, Nataraj, Edelsbrunner, Herbert, Facello, Michael, Fu, Ping, Mucke, EP and Varela, Carlos. "Alpha shapes: definition and software." *Proceedings of the 1st international computational geometry software workshop*, Vol. 63. 66. 1995.

[10] Edelsbrunner, Herbert and Mücke, Ernst P. "Three-dimensional alpha shapes." *ACM Transactions On Graphics (TOG)* Vol. 13 No. 1 (1994): pp. 43–72.

[11] Krishnamurthy, Vinayak, Poudel, Laxmi, Ebert, Matthew, Weber, Daniel H., Wu, Rencheng, Zhou, Wenchao, Akleman, Ergun and Sha, Zhenghui. "LayerLock: Layer-Wise Collision-Free Multi-Robot Additive Manufacturing Using Topologically Interlocked Space-Filling Shapes." *Computer-Aided Design* Vol. 152 (2022): p. 103392. DOI https://doi.org/10.1016/j.cad.2022.103392. URL https://www.sciencedirect.com/science/article/pii/S0010448522001300.

[12] Ebert, Matthew, Stone, Ronnie, Koithan, John, Zhou, Wenchao, Pharr, Matt, Estrin, Yuri, Akleman, Ergun, Sha, Zhenghui and Krishnamurthy, Vinayak. "NoodlePrint: Cooperative Multi-Robot Additive Manufacturing with Helically Interlocked Tiles." *Journal of Manufacturing Science and Engineering* (2025): pp. 1–21DOI 10.1115/1.4067617. URL https://asmedigitalcollection.asme.org/manufacturingscience/article-pdf/doi/10.1115/1.4067617/7422467/manu-24-1396.pdf, URL https://doi.org/10.1115/1.4067617.

[13] Poudel, Laxmi, Marques, Lucas Galvan, Williams, Robert Austin, Hyden, Zachary, Guerra, Pablo, Fowler, Oliver Luke, Moquin, Stephen Joe, Sha, Zhenghui and Zhou, Wenchao. "Architecting the cooperative 3d printing system." *International Design Engineering Technical Conferences and Computers and Information in Engineering Conference*, Vol. 83983: p. V009T09A029. 2020. American Society of Mechanical Engineers.

[14] Poudel, Laxmi, Zhou, Wenchao and Sha, Zhenghui. "A generative approach for scheduling multi-robot cooperative three-dimensional printing." *Journal of Computing and Information Science in Engineering* Vol. 20 No. 6 (2020): p. 061011.

[15] McPherson, Jace and Zhou, Wenchao. "A chunk-based slicer for cooperative 3D printing." *Rapid Prototyping Journal* Vol. 24 No. 9 (2018): pp. 1436–1446.

[16] Stone, Ronnie FP, Ebert, Matthew, Zhou, Wenchao, Akleman, Ergun, Krishnamurthy, Vinayak and Sha, Zhenghui. "SafeZone: A Topologically-Aware Voronoi-Based Framework For Fast Collision-Free Cooperative 3D Printing." *International Design Engineering Technical Conferences and Computers and Information in Engineering Conference*. 2024.

[17] Stone, Ronnie FP, Ebert, Matthew, Zhou, Wenchao, Akleman, Ergun, Krishnamurthy, Vinayak and Sha, Zhenghui. "SafeZone*: A Graph-Based and Time-Optimal Cooperative 3D Printing Framework." (2025).

[18] Tam, TKH and Armstrong, Cecil G. "2D finite element mesh generation by medial axis subdivision." *Advances in engineering software and workstations* Vol. 13 No. 5-6 (1991): pp. 313–324.

[19] Ang, Pin Yang and Armstrong, Cecil G. "Adaptive shape-sensitive meshing of the medial axis." *Engineering with computers* Vol. 18 (2002): pp. 253–264.

[20] Quadros, William Roshan. "LayTracks3D: A new approach for meshing general solids using medial axis transform." *Computer-Aided Design* Vol. 72 (2016): pp. 102–117.

[21] Wang, Hao, Lu, Tong, Au, Oscar Kin-Chung and Tai, Chiew-Lan. "Spectral 3D mesh segmentation with a novel single segmentation field." *Graphical models* Vol. 76 No. 5 (2014): pp. 440–456.

[22] Sidi, Oana, Van Kaick, Oliver, Kleiman, Yanir, Zhang, Hao and Cohen-Or, Daniel. "Unsupervised co-segmentation of a set of shapes via descriptor-space spectral clustering." *Proceedings of the 2011 SIGGRAPH Asia Conference*: pp. 1–10. 2011.

[23] Borenstein, Johann and Koren, Yoram. "Real-time obstacle avoidance for fast mobile robots in cluttered environments." *Proceedings., IEEE International Conference on Robotics and Automation*: pp. 572–577. 1990. IEEE.

[24] Xu, Yangsheng, Mattikalli, Raju and Khosla, Pradeep. "Motion planning using medial axis." *IFAC Proceedings Volumes* Vol. 25 No. 28 (1992): pp. 135–140.

[25] Lien, Jyh-Ming, Thomas, Shawna L and Amato, Nancy M. "A general framework for sampling on the medial axis of the free space." *2003 IEEE International Conference on Robotics and Automation (Cat. No. 03CH37422)*, Vol. 3: pp. 4439–4444. 2003. IEEE.

[26] Wilmarth, Steven A, Amato, Nancy M and Stiller, Peter F. "Motion planning for a rigid body using random networks on the medial axis of the free space." *Proceedings of the fifteenth annual symposium on Computational geometry*: pp. 173–180. 1999.

[27] Denny, Jory, Greco, Evan, Thomas, Shawna and Amato, Nancy M. "MARRT: Medial axis biased rapidly-exploring random trees." *2014 IEEE International Conference on Robotics and Automation (ICRA)*: pp. 90–97. 2014. IEEE.

[28] Wolter, Franz-Erich. "Cut locus and medial axis in global shape interrogation and representation." (1993).

[29] Turk, Greg and Levoy, Marc. "Zippered polygon meshes from range images." *Proceedings of the 21st annual conference on Computer graphics and interactive techniques*: pp. 311–318. 1994.

[30] Lee, Der-Tsai. "Medial axis transformation of a planar shape." *IEEE Transactions on pattern analysis and machine intelligence* No. 4 (1982): pp. 363–369.