# Multi-agent Collective Construction using 3D Decomposition

Akshaya Kesarimangalam Srinivasan[1], Shambhavi Singh[2], Geordan Gutow[3], Howie Choset[4] and Bhaskar Vundurthy[5]

*Abstract*— This paper addresses a Multi-Agent Collective Construction (MACC) problem that aims to build a three-dimensional structure comprised of cubic blocks. We use cube-shaped robots that can carry one cubic block at a time, and move forward, reverse, left, and right to an adjacent cell of the same height or climb up and down one cube height. To construct structures taller than one cube, the robots must build supporting stairs made of blocks and remove the stairs once the structure is built. Conventional techniques solve for the entire structure at once and quickly become intractable for larger workspaces and complex structures, especially in a multi-agent setting. To this end, we present a decomposition algorithm that computes valid substructures based on intrinsic structural dependencies. We use Mixed Integer Linear Programming (MILP) to solve for each of these substructures and then aggregate the solutions to construct the entire structure.

Extensive testing on 200 randomly generated structures shows an order of magnitude improvement in the solution computation time compared to an MILP approach without decomposition. Additionally, compared to Reinforcement Learning (RL) based and heuristics-based approaches drawn from the literature, our solution indicates orders of magnitude improvement in the number of pick-up and drop-off actions required to construct a structure. Furthermore, we leverage the independence between substructures to detect which substructures can be built in parallel. With this parallelization technique, we illustrate a further improvement in the number of time steps required to complete building the structure. This work is a step towards applying multi-agent collective construction for real-world structures by significantly reducing solution computation time with a bounded increase in the number of time steps required to build the structure.

## I. INTRODUCTION

There is a growing class of applications in which robots are used to assemble and construct structures [3]. Some robotics technologies for on-site building construction [18] include additive manufacturing [19], automated robotic assembly [20], and bricklaying [21]. It is particularly relevant in settings like open pit mining and extraterrestrial or underwater construction where human presence is difficult or dangerous [5][22][23][24]. Delegating construction to robots in these scenarios can thus be beneficial. Automation can

*This work was not supported by any organization

[1]Akshaya Kesarimangalam Srinivasan is a Masters student in the Robotics Institute at Carnegie Mellon University, USA

[2]Shambhavi Singh is an intern in the Robotics Institute at Carnegie Mellon University, USA, and a student at Birla Institute of Technology and Science, Pilani, India

[3]Geordan Gutow is a Postdoctoral fellow in the Robotics Institute at Carnegie Mellon University, USA

[4]Howie Choset is a Professor with the Robotics Institute at Carnegie Mellon University, USA

[5]Bhaskar Vundurthy is a Postdoctoral fellow in the Robotics Institute at Carnegie Mellon University, USA

also improve construction speed, and efficiency [4]. In these applications, teams of smaller robots are often more effective than a few larger robots as they are cheaper, easier to deploy, and facilitate more parallelization [2][3][25].



(a) Input structure     (b) Structural decomposition

(c) Ordering of substructures     (d) Parallelizable construction

Fig. 1: Visualization of our solution to the MACC problem.

The multi-agent collective construction (MACC) problem aims to construct a given three-dimensional structure using a set of robots in a grid world in the minimum number of time steps [1]. Previously the MACC problem has been solved using optimization [1], heuristics [2][9], [10], or Reinforcement Learning (RL) [6], [11]. All of these approaches have severe limitations. Although [1] finds an optimal solution, it has a very high solution computation time, even for some simple structures. [2] requires less computation time but provides high-cost solutions. While the approach taken in [6] is cheap to evaluate on a novel structure, it requires large amounts of training data and does not reliably produce the desired structure.

We present our work on decomposing an input structure into substructures and using an existing Mixed Integer Linear Programming (MILP) formulation [1] for each substructure to find good-quality solutions in a reasonable amount of time. This is conceptually visualized in Fig. 1.

Our main contributions are as follows:

1) An algorithm to decompose an input structure into

substructures whose construction may be planned independently

2) An algorithm to find an order in which the substructures can be built
3) An approach to leverage dependency between substructures to identify those that can be built simultaneously
4) Extensive numerical results demonstrating the computational improvements over existing methods

The problem formulation is presented in section II. Then, existing approaches to the multi-agent collective construction problem are outlined in section III. Section IV presents the decomposition and bottom-up planning algorithms proposed to solve this problem. The extension of the algorithm to parallel construction is described in section IV-C. Section V provides numerical results demonstrating the approach's effectiveness. Section VI presents some conclusions and possible directions for future work.

## II. PROBLEM FORMULATION

Similar to [1], the MACC problem in this work is set in a 3D grid world with two principal components: a predefined structure comprised of cube-shaped building *blocks* and block-sized *robots*. The robots collaboratively construct the structure by moving these blocks using any of the following permissible actions in a given time step:

1) Move to a free cell in the four compass directions
2) Climb up or down one block height to an adjacent occupied cell
3) Pick up or place a block at an adjacent cell of the same height

The standard rules of gravity apply to all the blocks, and it is assumed that robots always interact with the topmost blocks. In order to access a cell not surrounded by blocks of the same height, the robots construct scaffolding. All scaffolding must be removed after the completion of the entire structure. We further assume an unlimited supply of blocks at the boundary of the grid world and disregard any movement of the robots outside the grid world. In other words, the robots that exit the grid world can enter from any boundary cell in the next time step. The world is initially empty, and the robots start and finish outside the grid world.

This work discusses time-efficient algorithms to obtain a sequence of actions for $N$ robots to collectively construct a predefined structure. We present a comparison of our solutions with state-of-the-art via the following metrics:

- **Computation Time:** Time taken to compute the action sequence for building the structure
- **Makespan:** Total number of time steps to complete building the structure
- **Sum of costs:** Total number of actions for $N$ robots to build the structure

## III. RELATED WORKS

Multi-agent collective construction has been explored in both two [8] and three-dimensional worlds with varying types of agents and construction blocks. For instance, [3] uses a team of quadrotors to build structures made of beams and columns. Harvard's TERMES project, on the other hand, addresses the problem with homogeneous blocks and agents [4]. They show how teams of smaller robots are effective at collectively building structures much larger than themselves. The TERMES project inspired several works that proposed ways to find a sequence of actions for TERMES-like robots to build structures.

[2] solves the problem by performing dynamic programming on a minimum spanning tree that spans the cells of the weighted workspace and restricts the agent's movements to the edges of this tree. This planning method minimizes the number of pick-up and drop-off operations but is restricted to the single agent case. [7] extends this to multiple agents by parallelizing the action sequence. It performs a local search to find a more cost-efficient spanning tree using a recursive algorithm. However, this increases the algorithm's complexity, and the solution is still not optimal.

The distributed multi-agent reinforcement learning algorithm used in [6] extends single-agent advantage actor-critic to enable multiple agents to learn a homogeneous, distributed policy. The learned policy is tested with swarms of various sizes on structures not seen during the nine days of training. Although this approach is complete, some structures were not built accurately. Moreover, the makespan of the test structures was 100 times more than the minimum possible makespan found by optimization approaches.

The work most closely related to the current effort is the optimization approach presented in [1]. The MACC problem is solved using MILP or Constraint Programming (CP). The MILP model treats all robots as one flow through a time-expanded graph. Each decision variable defines the action the robot took and if it was carrying a block at that time step. The dynamics of the world are modeled as constraints for the set of decision variables. The formulation also models each cell in the workspace as pillars that aim to reach their target height. The CP approach uses a simpler network flow structure than the MILP approach. CP models the specifications of the world, such as the vertical dimension, actions, and block-carrying state, as logical and element constraints that better exploit the strength of CP. Both formulations use an objective function that minimizes the sum of costs and externally minimizes makespan. The MILP model computed the globally optimal makespan and the optimal sum of costs for a particular makespan. Both optimization models found feasible solutions to the six test instances with makespans less than 20. However, the MILP formulation required less computation time than the CP formulation. The approach presented in Section IV calls this MILP formulation as a subroutine.

The low makespans of the structures studied in [1] indicate that small structures are simple to construct as they do not need large makespans. However, even for some of the six small structures, the optimization models needed more than five days to compute a solution. As will be further demonstrated in section V, the MILP solution computation time is excessive for complex structures requiring longer

makespans to complete construction. This emphasizes the need to find approaches with a practical solution computation time for structures of varying complexity.

There has been some work, including [15], [14] for solving these large MILP problems. [14] uses a two-level approach to make large MILP problems tractable. It first coarsens binary variables to reduce the number of variables in the MILP problem and forms a semi-coarse model. It then aggregates constraints by partitioning them into groups and adding the violated constraints to the semi-coarse model iteratively till all the constraints in the full model are added. Inspired by 3D model decomposition work [16], this paper reduces the number of variables in the MILP problem by solving for one substructure at a time. Constraints are aggregated at each stage to represent the intermediate structure to be built.

## IV. APPROACH

In this section, we propose a decomposition algorithm to break down predefined structures into simpler substructures. We then determine an order in which it is possible to build the substructures and utilize Mixed Integer Linear Programming (MILP) to compute an optimal sequence of actions for every substructure [1].

### A. Structural Decomposition

In a 3D grid world of dimensions $(X \times Y \times Z)$, we denote predefined structures using $\bar{z}(x,y)$ where $\bar{z}$ indicates the height of the topmost block at every grid location $(x,y)$, $x,y,z,\bar{z} \geq 0$ and $x \leq X, y \leq Y, z \leq \bar{z} \leq Z$. We begin by ensuring that all predefined structures are valid and performing a similar check for substructures.

*Definition 1:* A structure is considered to be **valid** if, for every block of the structure at height $z$ $(z > 1)$, there exists a block at height $(z-1)$.

Definition 1 ensures that any movement in the blocks is achieved only through the robots' permissible actions. A valid structure always has a sequence of permissible actions that constructs the structure. It is worth noting that the robots can utilize blocks at a lower height as scaffolding for higher blocks. Such an action minimizes duplication of efforts from building temporary scaffolding for every block in the structure. As a result, it is preferable to ensure that any useful blocks (for scaffolding) are part of the same substructure as the higher block under consideration. We use the notion of a *shadow region* to identify the relation between neighboring blocks and identify substructures.

*Definition 2:* For a tower of height $z$ located at $(x,y)$, all cells $(x',y',z')$ s.t.

$$|x-x'| + |y-y'| < h$$

and

$$z' \leq h - (|x-x'| + |y-y'|)$$

are part of the **shadow region** of the tower.

Substructures are inherently associated with a specific order in that the validity of a substructure depends on which substructures are already present.

*Definition 3:* Consider a set of substructures represented by $S_1$ to $S_d$ where $'d'$ is the number of substructures. Let $\bigcup_{i=1,2,\cdots,j-1} S_i$ be a valid structure. Then $S_j$ is a **valid substructure** if $\bigcup_{i=1,2,\cdots,j} S_i$ is also a valid structure.

In the context of our problem, we define the basin of attraction of a tower of height $h$ in the structure to be all cells in the shadow region of that tower. Equivalently, all the structure blocks that can help build a tower by acting as direct support or scaffolding are part of its basin of attraction.

The decomposition algorithm iterates through the towers in the input structure in decreasing order of height. At each step, the blocks in the shadow region of the tower that are not already part of another substructure become part of its substructure. This decomposition algorithm is described in Algorithm 1.

---

**Algorithm 1:** Decomposition using basins of attraction

1 $(X,Y,Z) \leftarrow$ grid world dimension;
2 $\bar{z} \leftarrow$ input structure;
3 $towers \leftarrow$ non-zero $\bar{z}$ elements in decreasing order;
4 Initialize the set of all substructures, $sub = \varnothing$;
5 **for** $h$ *in towers* **do**
6     **if** *Topmost block of tower $h$ already in a substructure* **then**
7         continue;
8     **end**
9     Initialize substructure of $h$, $sub_h \leftarrow \varnothing$;
10     Shadow indices, $s_{idx} \leftarrow$ grid cells in shadow region of tower $h$;
11     **for** $s$ *in* $s_{idx}$ **do**
12         **if** $s$ *not in any substructure* **then**
13             add $s$ to $sub_h$;
14         **end**
15     **end**
16     Add $sub_h$ to $sub$;
17 **end**

---

*Theorem 1:* Each substructure in the order they are found by algorithm 1 is a valid substructure.

Proof of Theorem 1 is omitted in this paper for brevity.

### B. Bottom Up Planning

Consider the example shown in Fig. 1. Algorithm 1 identifies five substructures (see Fig. 2), and hence there are 120 different possible orders in which the substructures can be constructed. However, building all substructures will not be possible in some of those orders. For example, consider building the substructures in the order of (2),(4),(5),(1), and (3). This is not possible as (4) requires (1) to be built before it, and (3) cannot be built once (1), (2), (4), and (5) are constructed. Hence, the order of construction of the substructures is not trivial.

To find an order of construction of substructures, i.e., a sequence of substructures, we define the traversability

Fig. 2: Substructures identified by Algorithm 1

property of substructures. As with validity, this is depends on which other substructures are already present.

*Definition 4:* A substructure is **traversable** if, for every block in the substructure, there exists a feasible path from the boundary to at least one neighbor cell at the same height as the block in the presence of previously constructed substructures.

*Definition 5:* A **feasible path** is a sequence of permissible actions, not including picking up blocks from previously constructed substructures.

*Lemma 1:* If there is no neighbor cell at the same height as a block to be placed in a substructure, as a consequence of Algorithm 1, there will always be enough space to build a scaffolding of the required height.

The proof of Lemma 1 is omitted in this paper for brevity.

*Definition 6:* A sequence of substructures is **feasible** if $\forall i \in 1, \cdots, d$, $S_i$ is traversable and $\bigcup_{k=1,2,\cdots,i} S_k$ is a valid structure.

Obtaining a feasible sequence of substructures is treated as an assembly sequencing problem where the substructures are the components, and the input structure is the final product. The key idea is to build substructures in the reverse order in which they can be removed/disassembled from the goal structure while leaving at every step a valid structure. Such an order is obtained via a bottom-up planning algorithm inspired from [12], [13], presented in Algorithm 3.

*Remark 1:* Henceforth, the index of a substructure refers to the order in which the decomposition algorithm found the substructures.

To find the order in which substructures can be removed, we need to check if each block in a substructure is removable. From the problem formulation:

1) To remove a block at $(x, y, z)$, there should be no block at $(x, y, z+1)$
2) To pick up a block at $(x, y, z)$, the robot needs to be at a neighbor cell of $(x, y)$ at the same height $z$
3) If there are no neighbor cells at height $z$, it needs to have enough space to build scaffolding to height $z$

These conditions can be reduced further. Condition 3 is always satisfied by algorithm 1 as stated in Lemma 1.

Thus, the necessary and sufficient conditions to check if a substructure is removable can be reduced to two:

(a) No block in the substructure should have a block from another substructure on top of it
(b) There is a traversable path for a robot from outside the grid to at least one of the neighboring cells for each of the blocks in the substructure, i.e., the substructure should be traversable in the current state of the environment

Condition one implies that the only blocks allowed to be on top of a block $b$ in substructure $S$ are blocks that are part of $S$ too. Thus, the top blocks will be removed while removing $S$, ensuring block $b$ is removable. This yields a notion of dependence between substructures:

*Definition 7:* A substructure $S_i$ is said to be **dependent** on substructure $S_j$, (denoted $S_i \rightarrow S_j$) if $\bigcup_{\substack{k=1,2,\cdots,i \\ k \neq j}} S_k$ is not a valid substructure.

To check condition 2, a traversability matrix is constructed at every stage, considering all the substructures yet to be removed. The $(i, j)^{th}$ element is 1 only if the $(i, j)$ location is reachable from outside the grid. This is determined using dynamic programming starting from the boundary cells (which are always reachable). Every subsequent cell is reachable if a neighboring cell is reachable and the neighbor has a height difference of less than two from the cell (as the robots can only climb or descend one cube height at a time). Once this matrix is obtained, contour polygons of unreachable cells are computed. These contours represent impassable walls in the environment. Blocks are treated as not removable if they are enclosed on all four sides by blocks from another substructure or inside a contour of impassable walls. Note that in certain cases (the presence of a staircase on the wall's interior), scaffolding would, in principle, allow passing these walls. Thus this check is a sufficient but not necessary condition for removability. This algorithm is described in Algorithm 2.

Finally, a feasible sequence of substructures can be generated. Let $O_d$ be the order in which the substructures were found and $O_{new}$ be the new feasible ordering. Algorithm 3 iterates through $O_d$ in reverse order and adds only substructures that can be removed using Algorithm 2 to $O_{new}$. During one pass through $O_d$, if none of the substructures were removable, this implies that a substructure $S_i$ was not traversable due to substructure $S_j$ and $S_i \rightarrow S_j$. In this rare scenario, the two substructures are merged, resulting in a traversable substructure. This is repeated until all substructures are added to $O_{new}$. The final feasible sequence is computed as the reverse of $O_{new}$.

*Remark 2:* None of the 206 different structures considered in section V required merging to obtain a feasible sequence of substructures.

*Lemma 2:* The reverse of the order of removing substructures is a feasible assembly order.

The proof of Lemma 2 is omitted in this paper for brevityx.

*Theorem 2:* Given valid substructures, algorithm 3 always generates a feasible sequence of substructures.

*Lemma 3:* Given a valid input structure, the decomposition and ordering algorithms followed by MILP optimization for each substructure return a legal/feasible action sequence to build the structure if one exists.

*Theorem 3:* Consider a structure decomposed into $d$ substructures. The number of time steps required to construct the structure by constructing substructures sequentially is no more than $d$ times the number of time steps required to construct the structure without decomposition.

**Algorithm 2:** substructure removable check

1   $S \leftarrow$ substructure being checked
2   *traversability_matrix* $\leftarrow$ reachable positions in the (x,y) grid
3   *contours* $\leftarrow$ polygons representing impassable enclosures in the traversability matrix
4   **for** $b \in$ *blocks of S* **do**
5     **if** *b surrounded by blocks from another substructure in all four directions* **then**
6       **return** False;
7     **end**
8     **if** *b inside any contour* $\in$ *contours* **then**
9       **return** False;
10    **end**
11 **end**
12 **return** True;

---

**Algorithm 3:** substructure ordering

1   $O_d \leftarrow$ original order;
2   $O_{new} \leftarrow$ feasible order;
3   **while** *all substructures are not ordered* **do**
4     **for** *sub in reverse order in $O_d$* **do**
5       **if** *sub in $O_{new}$* **then**
6         continue;
7       **end**
8       **if** *removable(sub)* **then**
9         delete *sub* from $O_d$;
10        append *sub* to $O_{new}$;
11       **end**
12     **end**
13     **if** *no sub was removed and all substructures not ordered* **then**
14       Merge last two substructures in $O_d$ that are not in $O_{new}$
15     **end**
16 **end**
17 **return** reverse of $O_{new}$

Proofs of Theorems 2-3 and Lemma 3 are omitted in this paper for brevity.

### C. Parallel Construction

Once the substructures and a feasible order are obtained, any of the conventional approaches ([1], [6] and [2]) can be used to find the sequence of actions to build each substructure. However, if mixed integer linear programming as in [1] is used, some substructures can be built in parallel. This reduces the total time to build the structure.

This is achieved by modifying the bottom-up planner. In sequential construction of substructures: any input is decomposed using Algorithm 1 into substructures ordered using Algorithm 3. For parallel construction, Algorithm 3 is modified such that at every stage, all the substructures that can be removed are determined instead of just the substructure being

considered as per $O_d$ (default order). These substructures can potentially be built in parallel.

Let *SP* be a set of substructures that can be built in parallel. The first structure in *SP* is built as described for sequential construction. For every subsequent substructure $S_i$ in *SP*, the actions taken to build all previous $S_j$ for $j \in 1, \cdots, i-1$ are added as a constraint to the MILP. This ensures that $S_i$'s solution avoids agent-agent collision with the agents building the previous substructure and does not use more agents than permitted. In this approach, in the worst case, if the substructures in *SP* cannot be built in parallel (due to, say, insufficient agents), the ordering algorithm will provide a sequential order.

Parallel execution reduces makespan at the cost of increased constraints in the MILP formulations for each substructure. However, our initial experiments show that this does not significantly affect the solution computation time.

## V. RESULTS

### A. Experimental Setup

We evaluate the effectiveness of our algorithm on a variety of test structures: (a) six test cases used by [1] and [6], (b) one hundred randomly generated structures in a 10x10x4 grid world, and (c) one hundred randomly generated structures in a 7x7x4 grid world.

The MILP approach obtains the action sequence to build each structure or substructure per the experiment. All optimization models are solved using Gurobi 9.0.2, a state-of-the-art solver for MILP on an Intel® Core™ i7-7700K CPU @ 4.20GHz × 8 with 94GB of memory. In every case, the Gurobi model was allowed to run for up to 10,000 seconds for each structure.

### B. Results on an Example Structure

We first demonstrate our method on the example structure from Fig. 1. As shown in Fig. 1 (b), the algorithm gives five substructures. The MILP model gets each substructure as an input and gives a set of action sequences with a maximum of 20 robots for each substructure. Table I presents the metrics of the solution for each substructure. The approach in [1] iterates through increasing makespans until the model becomes feasible and subsequently performs optimization to find a solution. 'Solve Time' denotes the time taken to optimize the final model, and 'Total Solve Time' denotes the time to iterate through all time steps along with the solve time of the model.

TABLE I: Results on the structure used in Fig. 1 illustrating the metrics for each substructure indicated by Sub. No.

| Sub. No. | Makespan | Sum-of-costs | Solve Time | Total Solve Time |
|---|---|---|---|---|
| 1 | 14 | 68 | 37.0s | 123.6s |
| 2 | 14 | 50 | 8.6s | 96.8s |
| 3 | 14 | 42 | 7.4s | 95.5s |
| 4 | 15 | 39 | 6.7s | 111.4s |
| 5 | 10 | 8 | 1.3s | 34.6s |
| Total | 67 | 207 | 61.0s | 461.9s |

## C. Comparison with Other Approaches on Six Test Structures

We next conduct experiments to compare our approach with two existing sub-optimal approaches [2][6]. The experiments use a set of six test structures also used by [1] and [6] for reporting their results. The Reinforcement Learning approach [6] runs a pretrained policy for 100 trials using eight agents for each structure. The authors noted that using more than eight agents detriments the performance.

Table II reports the sum of costs for the tree approach [2] and the average sum of costs for the RL approach (over successful trials). It also reports the sum of costs using the MILP approach with decomposition (our approach). Since these existing methods aim to minimize the number of pick-up and drop-off actions alone, our approach though not optimal achieves costs nearly an order of magnitude smaller than these methods.

TABLE II: Comparison of the sum of costs of solutions for our approach with other non-optimal methods

| Structure | Tree based [2] | RL based [6] | Ours |
|-----------|----------------|--------------|------|
| 1 | 1144 | 3040 | 179 |
| 2 | 836 | 1026 | 128 |
| 3 | 1590 | 3056 | 326 |
| 4 | 2120 | 3252 | 263 |
| 5 | 2180 | 2804 | 381 |
| 6 | 808 | 1276 | 161 |

Finally, we compare our approach with an existing optimal approach [1]. [1] uses the entire structure as input to their MILP formulation. In our experiments, the number of robots is limited to 20 for uniformity between solving with and without decomposition. Table III presents results for the approach from [1] and our decomposition technique with and without parallelism. Decomposition significantly improves the solution computation times over pure MILP while maintaining a similar sum of costs. However, we observe an increase in time steps when constructing the structure via substructures. Introducing parallel construction largely mitigates this increase. On average, parallel construction reduces the number of time steps by 46% compared to pure decomposition with a similar sum of costs and solution computation time. Note that, for structures 4-6, parallel construction significantly increases the solution computation time. However, it is still much faster than MILP without decomposition.

## D. Comparison with Exact Approach on Random Structures

We demonstrate the scope of our approach by testing all three optimization-based approaches on a set of randomly generated structures in two grid world sizes: 10x10x4 and 7x7x4. Structures were generated targeting ranges of occupancy percentage: the number of blocks in the input structure divided by the number of cells in the workspace. Out of the 100 test structures, half had an occupancy percentage between 40% to 60%, one-fourth had less than 40%, and one-fourth had more than 60%. The maximum number of



Fig. 3: The total computation time when using MILP to solve for the entire structure is exponential in the number of optimization variables.

robots permitted is 20 for the 10x10x4 environment and 6 for the 7x7x4 environment.

On average, for a set of test structures in a fixed environment size of 7X7X4, we noted that as the occupancy percentage increased from 30 to 70%, the number of time steps increased from 76 to 122. This happens because the number of variables in the optimization model in the MILP formulation increases significantly, with the number of time steps following a linear trend. However, the solution computation time increases exponentially as the number of variables in the model increases. The exponential increase is illustrated by Fig. 3. This is why decomposition is beneficial for run-time: solving several smaller MILP models (one for each substructure, smaller because substructures can usually be built in fewer time steps than the full structure) is much faster than solving one large model (for the entire structure at once).

For 100 random structures in the 10x10x4 grid world, we again tested our decomposition algorithm with MILP optimization. Table IV shows the results for both environment sizes.

## VI. CONCLUSIONS

In this paper, we presented an algorithm to decompose any input structure into substructures and obtain a feasible order to build the substructures. Experimental results showed that MILP optimization with decomposition has an order of magnitude improvement in the solution computation time compared to MILP without decomposition. However, the former demonstrated a higher makespan when the substructures were built sequentially or with basic parallelization.

Developing more sophisticated algorithms to construct substructures in parallel is a promising future direction. For example, one can determine the action sequence to construct each substructure that can be built in parallel. Then the action sequences can be post-processed to enforce collision and number of agents constraints. Further, the decomposition into substructures can be modified to optimize metrics like the number of scaffolding blocks required or the parallelism provided.

Finally, the similarity between substructures can be leveraged to calculate the action sequence required to build the

## TABLE III: Results for the six test cases

This table presents a comparison of metrics for solutions obtained using our approach with the time-optimal solution presented in [1] on a set of six test structures used by [1] and [6]. Numbers indicated in green show the improvements of our approach compared to the state of the art. Here, **A - MILP-based Exact Approach, B - 3D decomposition with MILP-based Approach and C - Parallel Construction of substructures using 3D decomposition with MILP-based Approach**

| Test Structure | | | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|
| Method | A | B | C | A | B | C | A | B | C |
| Sum of costs | 173 | 176 | 179 | 124 | 128 | 128 | - | 326 | 326 |
| No. of timesteps | 13 | 48 | 17 | 13 | 48 | 14 | - | 106 | 44 |
| Final Computation Time (in sec) | 1030.3 | 16.6 | 11.1 | 61.0 | 11.5 | 10.2 | - | 49.6 | 26.4 |
| Total Computation Time (in sec) | 1115.0 | 241.3 | 259.4 | 139.0 | 235.9 | 198.1 | >10,000 | 377.2 | 318.1 |

| Test Structure | | | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|
| Method | A | B | C | A | B | C | A | B | C |
| Sum of costs | - | 204 | 263 | - | 365 | 381 | 160 | 153 | 161 |
| No. of timesteps | - | 113 | 75 | - | 130 | 90 | 21 | 50 | 40 |
| Final Computation Time (in sec) | - | 36.7 | 610.7 | - | 37.1 | 639.4 | 1215.3 | 15.6 | 265.9 |
| Total Computation Time (in sec) | >10,000 | 31.2 | 758.6 | >10,000 | 27.8 | 688.5 | 1715.2 | 12.2 | 287.9 |

## TABLE IV: Results for tests on random structures

Here, **EA - MILP-based Exact Approach, EAD - 3D decomposition with MILP-based Exact Approach**

| Environment Size | 10x10x4 | | 7x7x4 | |
|---|---|---|---|---|
| Method | A | B | A | B |
| Sum of costs | - | 384 | 83.72 | 97.12 |
| No. of timesteps | - | 84 | 18.01 | 51.17 |
| Final Computation Time | - | 48.10 | 229.41 | 1.48 |
| Total Computation Time | >10,000 | 567.99 | 423.51 | 37.83 |

latest substructure using solutions of previous substructures.

## REFERENCES

[1] Edward Lam, Peter J. Stuckey, Sven Koenig, and T. K. Satish Kumar. 2020. Exact Approaches to the Multi-agent Collective Construction Problem. In Principles and Practice of Constraint Programming: 26th International Conference, CP 2020, Louvain-la-Neuve, Belgium, September 7–11, 2020, Proceedings. Springer-Verlag, Berlin, Heidelberg, 743–758. https://doi.org/10.1007/978-3-030-58475-7_43.

[2] Kumar, T.K.S. & Jung, Sangmook & Koenig, Sven. (2014). A Tree-Based Algorithm for Construction Robots. Proceedings of the International Conference on Automated Planning and Scheduling. 2014. 481-489. 10.1609/icaps.v24i1.13673.

[3] Lindsey, Quentin & Mellinger, Daniel & Kumar, Vijay. (2012). Construction with quadrotor teams. Autonomous Robots. 33. 10.1007/s10514-012-9305-0.

[4] Petersen, Kirstin & Nagpal, Radhika & Werfel, Justin. (2011). TER-MES: An Autonomous Robotic System for Three-Dimensional Collective Construction. 10.15607/RSS.2011.VII.035.

[5] Werfel, Justin & Nagpal, Radhika. (2006). Extended Stigmergy in Collective Construction. IEEE Intelligent Systems. 21. 20-28. 10.1109/MIS.2006.25.

[6] Sartoretti, Guillaume et al. "Distributed Reinforcement Learning for Multi-robot Decentralized Collective Construction." DARS (2018).

[7] Trevor Cai, David Y. Zhang, T.K. Satish Kumar, Sven Koenig, and Nora Ayanian. 2016. Local Search on Trees and a Framework for Automated Construction Using Multiple Identical Robots: (Extended Abstract). In Proceedings of the 2016 International Conference on Autonomous Agents & Multiagent Systems (AAMAS '16). International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, 1301–1302.

[8] "Proceedings 2006 IEEE International Conference on Robotics and Automation [Title page]," Proceedings 2006 IEEE International Conference on Robotics and Automation, 2006. ICRA 2006., 2006, pp. 0_2-0_2, doi: 10.1109/ROBOT.2006.1641149.

[9] Alexander Grushin and James A. Reggia. 2008. Automated design of distributed control rules for the self-assembly of prespecified artificial structures. Robot. Auton. Syst. 56, 4 (April, 2008), 334–359. https://doi.org/10.1016/j.robot.2007.08.006.

[10] Panangadan A, Dyer MG. Construction in a Simulated Environment Using Temporal Goal Sequencing and Reinforcement Learning. Adaptive Behavior. 2009;17(1):81-104. doi:10.1177/1059712308101787

[11] Barros dos Santos, Sergio Ronaldo & Givigi, Sidney & Nascimento Jr, Cairo. (2013). Autonomous construction of structures in a dynamic environment using Reinforcement Learning. SysCon 2013 - 7th Annual IEEE International Systems Conference, Proceedings. 452-459. 10.1109/SysCon.2013.6549922.

[12] L. S. Homem de Mello and A. C. Sanderson, "A correct and complete algorithm for the generation of mechanical assembly sequences," in IEEE Transactions on Robotics and Automation, vol. 7, no. 2, pp. 228-240, April 1991, doi: 10.1109/70.75905.

[13] S. Chakrabarty and J. Wolter, "A structure-oriented approach to assembly sequence planning," in IEEE Transactions on Robotics and Automation, vol. 13, no. 1, pp. 14-29, Feb. 1997, doi: 10.1109/70.554344.

[14] Lin, Fu & Leyffer, Sven & Munson, Todd. (2016). A Two-Level Approach to Large Mixed-Integer Programs with Application to Cogeneration in Energy-Efficient Buildings. Computational Optimization and Applications. 65. 10.1007/s10589-016-9842-0.

[15] M. A. Bragin, P. B. Luh, B. Yan and X. Sun, "A Scalable Solution Methodology for Mixed-Integer Linear Programming Problems

Arising in Automation," in IEEE Transactions on Automation Science and Engineering, vol. 16, no. 2, pp. 531-541, April 2019, doi: 10.1109/TASE.2018.2835298.

[16] Jain, Arjun & Thormählen, Thorsten & Ritschel, Tobias & Seidel, Hans-Peter. (2012). Exploring Shape Variations by 3D-Model Decomposition and Part-based Recombination. Computer Graphics Forum. 31. 631-640. 10.1111/j.1467-8659.2012.03042.x.

[17] Nusse, H. E., Yorke, J. A., & Kostelich, E. J. (1994). Basins of Attraction. In Dynamics: Numerical Explorations: Accompanying Computer Program Dynamics (pp. 269–314). Springer US. https://doi.org/10.1007/978-1-4684-0231-5_7

[18] Gharbia, Marwan & Chang-Richards, Alice & Lu, Yuqian & Zhong, Ray & Li, Heng. (2020). Robotic technologies for on-site building construction: A systematic review. Journal of Building Engineering. 32. 101584. 10.1016/j.jobe.2020.101584.

[19] C. Ye, N. Chen, L. Chen and C. Jiang, "A Variable-Scale Modular 3D Printing Robot of Building Interior Wall," 2018 IEEE International Conference on Mechatronics and Automation (ICMA), 2018, pp. 1818-1822, doi: 10.1109/ICMA.2018.8484433.

[20] Jung, Kyoungmo & Chu, Baeksuk & Hong, Daehie. (2013). Robot-based construction automation: An application to steel beam assembly (Part II). Automation in Construction. 32. 62–79. 10.1016/j.autcon.2012.12.011.

[21] Y. Wu, H. H. Cheng, A. Fingrut, K. Crolla, Y. Yam and D. Lau, "CU-brick cable-driven robot for automated construction of complex brick structures: From simulation to hardware realisation," 2018 IEEE International Conference on Simulation, Modeling, and Programming for Autonomous Robots (SIMPAR), 2018, pp. 166-173, doi: 10.1109/SIMPAR.2018.8376287.

[22] Miri Weiss Cohen & Vitor Nazário Coelho (2021). Open-Pit Mining Operational Planning using Multi Agent Systems. Procedia Computer Science. ISSN 1877-0509, https://doi.org/10.1016/j.procs.2021.08.172.

[23] Khoshnevis, Behrokh. "Automated construction by contour crafting—related robotics and information technologies." Automation in construction 13.1 (2004): 5-19.

[24] Werfel, Justin & Nagpal, Radhika. (2008). Three-Dimensional Construction with Mobile Robots and Modular Blocks. I. J. Robotic Res.. 27. 463-479. 10.1177/0278364907084984.

[25] Silva, Maira Saboia da & Thangavelu, Vivek & Napp, N.. (2018). Autonomous Multi-Material Construction with a Heterogeneous Robot Team.