

Reconfiguration of a 2D Structure Using Spatio-Temporal Planning and Load Transferring

Javier Garcia¹, Michael Yannuzzi¹, Peter Kramer², Christian Rieck², Sándor P. Fekete², and Aaron T. Becker^{1,2}

Abstract—We present progress on the problem of reconfiguring a 2D arrangement of building material by a cooperative group of robots. These robots must avoid collisions, deadlocks, and are subjected to the constraint of maintaining connectivity of the structure. We develop two reconfiguration methods, one based on spatio-temporal planning, and one based on target swapping, to increase building efficiency. The first method can significantly reduce planning times compared to other multi-robot planners. The second method helps to reduce the amount of time robots spend waiting for paths to be cleared, and the overall distance traveled by the robots.

I. INTRODUCTION

A challenge in robotics is to use agents (robots) to change the configuration of a supply of passive building material. A typical task arises from relocating a collection of tiles from a given start configuration into a desired goal configuration in an efficient manner. Reconfiguration time can be decreased by using multiple robots, but this requires careful coordination to avoid collisions, deadlocks, imbalanced task allocation between robots, as well as maintaining important constraints such as connectivity of the structure.

These considerations play an important role when constructing *large-scale* configurations, ranging from plans for kilometer-scale manufacturing structures in space [1], [10], to millimeter-scale smart material [18], and nano-scale assembly with DNA [17]. In such domains, where disconnected components can drift apart, it is often necessary that the structure remains a single component.

In previous work [6], we showed how a sampling-based approach (the RRT*) can be used to enable a single robot to reconfigure a 2D set of connected tiles in complex environments, where multiple obstacles may be present. In this paper, we assume that the reconfiguration sequence is already given, by RRT* or another approach, and focus on multi-robot cooperation to carry it out. This requires methods for coordination and task allocation to improve reconfiguration times without disconnecting the structure.

Fig. 1 shows the hardware platform that motivates the constraints in this theoretical work. Two BILL-E bots, a platform developed by Jenett et al. [7], [9], are standing on a tile structure. These inchworm robots step on the tiles, using

This work was supported by the National Science Foundation under [IIS-1553063, 1849303, 2130793], the DFG project „Space Ants“, grant number FE 407/22-1, and the Alexander von Humboldt Foundation.

¹Electrical Engineering, University of Houston, TX USA

{jgarciagonzalez, mcyannuzzi, atbecker}@uh.edu

²Computer Science, TU Braunschweig, Braunschweig, Germany
{kramer, rieck, fekete}@ibr.cs.tu-bs.de

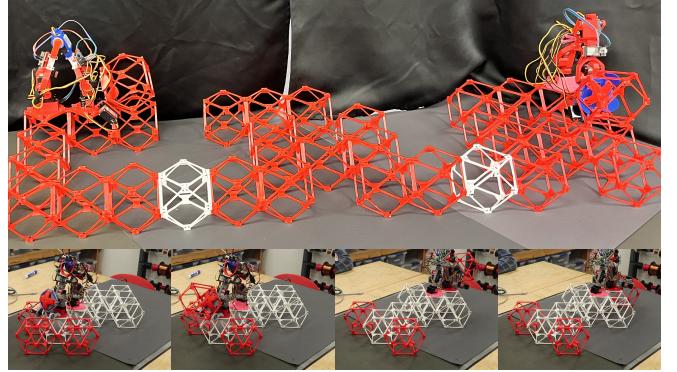


Fig. 1. This paper implements and compares algorithms for automated reconfiguration using multiple robots. Above is a physical representation of the problem, and a sequence of frames showing a BILL-E bot moving a tile. See video overview at <https://youtu.be/tCKMjhkzb8>.

a rotating key on the bottom of their feet to lock on the structure. A gripper located on the front foot can move up and down to pick and place tiles, and the robots can move while carrying the tiles. In order to guarantee that the robots can reach any point in the structure, the latter must remain connected throughout the reconfiguration.

II. RELATED WORK

A. Automated reconfiguration

In [9], the BILL-E bot traverses and alters a tile structure. For a single robot, the assembly sequence is deterministic, and is handled layer by layer in the three-dimensional case. The multi-robot case is handled by treating other robots as obstacles. If a desired path is blocked, a robot waits it to be cleared/constructed before placing their tile. In non-trivial configurations, this approach results in waiting times that increase with the number of robots.

Work on simplifying complex reconfiguration using principles of finite automata is presented in [1], [5] and [12]. Structures can be built, scaled and rotated while respecting the constraint of tile connectivity.

In our case, we exploit the ability to temporarily store tiles at any valid location when reconfiguring structures. One benefit is that this enables creating bridges that act as shortcuts. Deterministically finding these intermediate structures is challenging, so in [6] we used a sampling-based approach. The reconfiguration sequence is updated as lower cost paths are found. While we have only considered one robot (moving only one tile at a time), this paper focuses on performing reconfiguration sequences with $m \in \mathbb{N}$ robots.

B. Multi-robot planning

While increasing the number of robots from 1 to m has potential benefits, it also causes additional difficulties. A formulation as a centralized path planning problem raises the DoF by a factor of m . The increased complexity makes sampling-based approaches harder to compute. Moreover, the practical use of multiple robots relies on *distributed* methods for path planning and motion control.

To reduce the complexity of the problem, de-coupled planning methods have been developed. Otte and Correll [13] segment the joint problem and thus facilitate path computations; segments are only combined if no solution can be found. Wagner and Choset [19] combine segmentation with A* to create the M* algorithm, only considering neighbors when expanding from a given vertex. Additionally, collision information is propagated to find paths without collisions. Furthermore, a modification to the D* Lite algorithm is given by Pent et al. [14], where robots are treated as obstacles and graph costs are updated as the agents move.

C. Multi-agent cooperation

Increased reconfiguration speed can be achieved by considering swarms of cooperating robots, such as described in [2]–[4]. In addition, we can employ the robots’ ability to swap targets as well as tasks. A BILL-E can, e.g., transfer its carried tile to a second BILL-E and proceed to a new target if doing so reduces the overall reconfiguration cost.

We take inspiration from [20], where Wang and Rubenstein use local task swapping to create a shape with a homogeneous swarm by using a hop-count algorithm. Compared to their application, task swapping for the BILL-E bots is more complicated, since they cannot transfer carried tiles directly in their current version. Instead, the tile must be placed where another robot can later pick it up.

III. COMPUTATIONAL COMPLEXITY

We consider the *workspace* to be a rectangular unit grid, where each cell is either free, filled by a tile, or filled by an obstacle. We are interested in reconfiguring a set of tiles from a start to a goal configuration. Both are *connected* components, i.e., for every tile there is a path for the robot on the present configuration to every other tile. These shapes are called *polyominoes*. As neither a robot nor a tile carried by a robot can cross an obstacle, we assume that start and goal configurations are located within the same connected component of free space. With this assumption, there always exists a feasible reconfiguration sequence.

An arrangement is reconfigured by moving a robot to an adjacent position of a tile, picking it up, walking along a path on the remaining configuration, then placing the tile in another location. An ordered series of these operations is called a *reconfiguration sequence*. Distances between workspace positions are defined by the length of the geodesic edge-connected path between them, determined by a *breadth-first search* tree over the configuration. We call the problem of deciding whether the exists a sequence of a specific length that reconfigures a start into a goal configuration by a single

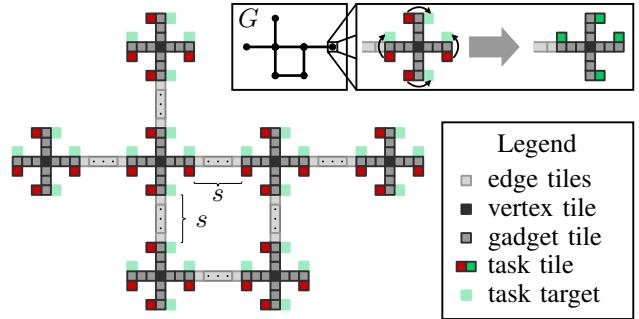


Fig. 2. Symbolic overview of the construction used to show NP-hardness of the BILL-E RECONFIGURATION problem, see proof of Theorem 1.

robot the BILL-E RECONFIGURATION problem. If more than a single robot is considered, i.e., we have m robots available, we refer to this as the **MULTI-ROBOT** or **COOPERATIVE** variant. We make use of the **HAMILTONIAN PATH** problem in grid graphs to prove NP-hardness of deciding the length of optimal reconfiguration sequences. This problem, which asks us to decide whether there exists a path in the input grid graph that visits each vertex exactly once, was shown to be NP-complete by Itai et al. [8].

Before arguing hardness of the reconfiguration problem, we observe that we can certify a correct solution in polynomial time, implying membership in NP. Clearly, it is possible to confirm the validity of an individual move in polynomial time with the number of tiles and robots involved. Validating an arbitrary reconfiguration sequence is therefore possible by individually validating the steps involved.

Observation 1. *We can validate solutions in polynomial time, so BILL-E RECONFIGURATION is in NP.*

Theorem 1. *BILL-E RECONFIGURATION is NP-complete.*

Proof. For an illustration of our construction, we refer to Fig. 2. Given a grid graph $G = (V, E)$, we construct a start and a goal configuration as follows.

We begin by creating an arbitrarily scaled embedding of the vertex positions in V in the grid. To achieve this, we place *vertex tiles* (black) in identical locations in both the start and goal configuration, such that two adjacent vertices $u, v \in V$ will be exactly $s + 7$ units apart in the configurations, as we combine a scale factor $s \in \mathbb{N}_+$ with a requirement of 7×7 units of space around each vertex tile for our construction.

Each of the vertex tiles is then surrounded by a so-called *vertex gadget*, which contains four *task tiles* and *task positions* (red and green) arranged around a 7×7 cross of auxiliary *gadget tiles* (dark gray), centered on the vertex tile itself. A task tile or position in this construction represents a location that is occupied exclusively in either the start or the goal configuration, respectively. We say that a reconfiguration sequence *solves* a vertex gadget exactly if it picks up each task tile and places it at one of the nearby task positions. Therefore, a reconfiguration sequence that transforms the start into the goal configuration must solve each gadget. Finally, we connect every pair of gadgets that

correspond to adjacent vertices in G by a straight path of s many so-called *edge tiles* (light gray), which are identical in the start and goal configurations.

Let t_g now refer to the number of moves required to solve one vertex gadget. By construction, this value is equal for all gadgets. A reconfiguration sequence that solves all vertex gadgets therefore takes at least $t_g \cdot |V|$ moves, not accounting for travel between vertices. As this takes $(s + 7)$ moves per vertex by construction, we conclude that there exists a sequence consisting of exactly $(s + 7) \cdot (|V| - 1) + t_g \cdot |V|$ moves, if there is a Hamiltonian path in G . In this case, no reconfiguration sequence could ever be faster, as building bridges between any two vertex tiles would take more than $(s + 7)$ moves, therefore resulting in a more expensive reconfiguration sequence than on a Hamiltonian path.

If no Hamiltonian path exists, a shorter reconfiguration sequence that travels solely along the existing edges of the underlying graph would visit at least one vertex tile twice, implying a longer sequence. If the robot were to construct at least one additional bridge between vertices, this would immediately imply that it spent at least $(s + 7)$ moves constructing a bridge in addition to the $(s + 7) \cdot (|V| - 1) + t_g \cdot |V|$ moves it would have to spend solving the instance. \square

By creating m copies of the resulting configurations (one for each robot), placing these pairwise sufficiently far apart, and connecting them by a single line of tiles (to create a connected configuration), Corollary 1 is straightforward.

Corollary 1. **MULTI-ROBOT BILL-E RECONFIGURATION** with m robots is NP-complete for every $m \in \mathbb{N}$.

Clearly, every **MULTI-ROBOT BILL-E RECONFIGURATION** instance is an instance of the **COOPERATIVE BILL-E RECONFIGURATION** problem as well, meaning that Corollary 1 explicitly covers cooperative reconfiguration.

IV. METHODS

A. Graph representation

Many problems in path planning can be solved by converting them to graph search problems. The configurations of the system can be discretized into vertices connected by edges or transitions, based on the basic motions of the robots. Then, it becomes a matter of finding a valid path between the start and goal vertices along the provided edges.

A popular deterministic search algorithm is A^* , employed often due to its completeness and ease of implementation [11]. Every iteration, A^* expands a vertex by adding its *neighbors* (vertices that can be reached with one transition) to an open list. The most promising vertex in this list is then expanded and removed from the list. This is repeated until the goal is reached. The promise of vertex v is defined as

$$f(v) = g(v) + \varepsilon \cdot h(v). \quad (1)$$

Here, $g(v)$ is the cost from the start to v , and $h(v)$ is a heuristic (both admissible and consistent) to estimate the cost from v to the goal. Scaling the heuristic by $\varepsilon > 1$ speeds up the pathfinding process, at the cost of optimality [15].

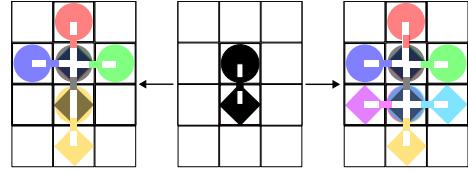


Fig. 3. Starting from a configuration (middle), BILL-E can reach several configurations after a single motion. For clarity, the front foot is drawn as a circle, and configurations are color coded. (Left) The motions in S_5 are • waiting, • stepping forward, • backward, and placing the front foot one tile to the right • or one tile to the left • and moving the back foot one tile forward. (Right) S_7 is expanded to include placing the back foot one tile to the left • or right • and moving the front foot one tile backward.

A significant drawback of such a graph representation, however, is that the number of neighbors each vertex has rapidly increases for systems with more degrees of freedom, making it impractical for systems with large numbers of robots. To deal with this, the M^* variant initially only considers the *limited neighbors* of vertices when expanding them [19]. This approach is equivalent to planning the paths of the robots independently. If a collision results, all vertices leading to it are updated with a collision set, and re-added to the open list. When expanding vertices with collision sets, the paths of the colliding robots are planned jointly. Notably, M^* is most useful in situations when only a subset of robots are in danger of colliding on their optimal paths to their goals. On crowded maps, where the majority of vertices have collision sets, it reverts to regular A^* .

In order to employ graph search methods, we must now precisely define the BILL-E configurations and the feasible transitions for the graph representation. Each vertex v represents a configuration by a tuple of the coordinates of the robot's feet and of the carried tile, if applicable. Neighbors of v_i are all the configurations that the BILL-E can reach after one edge transition starting from the respective vertex. All edge transitions depend on S , the robot's basic motions.

One advantage of using a graph search approach is that changes in hardware can be easily accounted for by changing S . For example, our current BILL-E has a rotation limit of about 180° on its feet. This means that if we want one foot to be able to rotate the robot around, it can only do so in one direction. While our robots are homogeneous, a heterogeneous group with different movement capabilities can be handled by suitably modifying S . Changes in S considerably impact the ability to find a solution.

Two sets of S , one with five motions and the other with seven, for a robot not carrying a tile are shown in Fig. 3. Although a larger S means that the robots are capable of moving into more configurations, it also increases the time spent creating and searching the graph. Hence it can be beneficial to limit the number of basic motions.

B. Temporal A^*

We implement a temporal version of A^* to handle the multi-robot problem with a time-varying structure. We use a *time horizon* H and a *priority queue* $Q = 1, 2, \dots, m$ as the de-coupling methods to reduce the complexity of the problem. The planner determines the next H moves of each

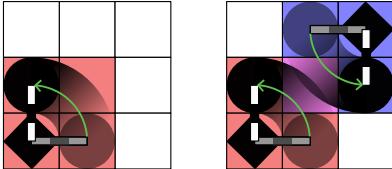


Fig. 4. The collision set of a motion is used to determine validity. (Left) The BILL-E rotates 90° counterclockwise, and the highlighted squares represent the respective collision set. (Right) Example of a collision resulting from moving two BILL-E bots in a certain way.

robot in order according to Q . Similar to [14] and [16], when path planning for one robot, the other robots are treated as obstacles. Higher priority robots, those with IDs appearing first in Q , only consider the positions of lower priority robots at $H = 1$. If $H > 1$, the lower priority robots are assumed to move out of the way so their positions are not considered. Lower priority robots consider the positions of higher priority robots at the H value they are planning for. If the current Q results in a deadlock, the ID of the lower priority robot is shifted up in Q and the paths are planned again. A higher H value can correct deadlock situations but is more computationally expensive. It is also possible for a deadlock to be unavoidable for all permutations of Q and any value of H , making our Temporal A* not complete. This is more likely to happen with smaller S .

Time is embedded into the graph by differentiating vertices that describe the same configuration for a given robot but at different values of time t . This way, *waiting* can be treated as any other member of S . All members of S can be made to take the same amount of time by the central computer. Even if some motions are faster than others, the robots are signaled to wait for all others to be done moving. This synchronous scheme allows us to assign uniform costs to S . Since *waiting* is as costly as moving, the planner is incentivized to try to move robots through different paths to reach their destinations. Having many vertices that lead to the same configurations at different times makes the neighbor sets grow quickly, further limiting the practical values of H .

Collisions are handled at edge transitions. If the transition involves a rotation, the tiles that the robot sweeps over are included in the *collision set* c . If the robot is carrying a tile, c is appropriately expanded. If c for two robots intersect, then there will be a collision, and the transition is labelled as invalid. In this case the lower priority robot will re-plan as explained before. Visuals of different c are provided in Fig. 4.

C. Load transfer

A key aspect of our automated reconfiguration scheme is the ability of the robots to cooperate by exchanging carried tiles. This allows robots to avoid collisions and deadlocks by changing targets, usually to pick-up and drop-off locations closer to their current location, reducing the density of paths.

Reducing the number of tiles traveled is the metric used to determine if a load transfer should take place. The cost notation is as follows: $P_{i,(a,b)}^{(x,y)}$ is the cost for robot i to pick-up a tile at location (x,y) starting from (a,b) . Similarly, $D_{i,(a,b)}^{(x,y)}$ is the cost for the same robot to drop-off a tile. Two

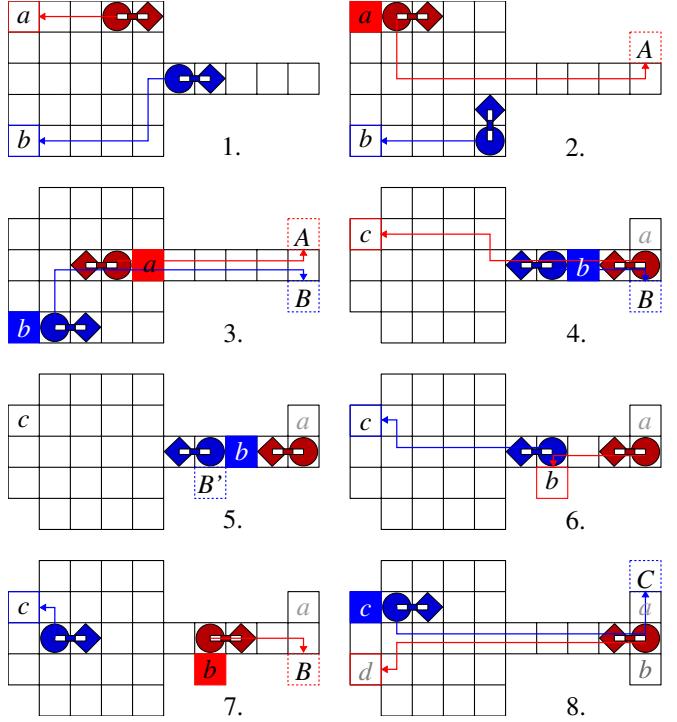


Fig. 5. 1. Each robot moves to pick up its assigned tile. 2. Red picks up tile a and is routed to place it at A . 3. Blue picks up tile b and is routed to place it at B . 4. Red places tile a and is routed to pick up tile c . 5. The condition for load transfer is met at the new B' . 6. Blue places tile b and is routed to tile c that Red was originally going to pick up, while Red is routed to tile b . 7. Red picks up tile b and is routed to drop it off. 8. The building continues as normal.

robots i and j , one carrying a tile to (x_1, y_1) starting from (a_1, b_1) , and the other on its way to pick up a tile at (x_2, y_2) starting from (a_2, b_2) , have an original combined cost of

$$C_{i,j} = D_{i,(a_1,b_1)}^{(x_1,y_1)} + P_{j,(a_2,b_2)}^{(x_2,y_2)}. \quad (2)$$

The load transfer process for the BILL-E bots involves finding secondary locations to drop-off and pick-up their loads since the robots cannot exchange tiles directly in their current version. Following the case above, a third location (x_3, y_3) , with an adjacent tile (a_3, b_3) is considered for the transfer. If the condition

$$D_{i,(a_1,b_1)}^{(x_3,y_3)} + P_{i,(a_3,b_3)}^{(x_2,y_2)} + P_{j,(a_2,b_2)}^{(x_3,y_3)} + D_{j,(a_3,b_3)}^{(x_1,y_1)} < C_{i,j} \quad (3)$$

is met, then robot i moves to place its tile at (x_3, y_3) , where robot j will pick it up later. Both robots then proceed to the other's original targets. In the case that both robots are initially carrying a tile, there is no need to switch tiles because the tiles are identical. Switching drop-off locations may shorten the required travel, i.e., if

$$D_{i,(a_1,b_1)}^{(x_2,y_2)} + D_{j,(a_2,b_2)}^{(x_1,y_1)} < D_{i,(a_1,b_1)}^{(x_1,y_1)} + D_{j,(a_2,b_2)}^{(x_2,y_2)}. \quad (4)$$

Load transfer is particularly useful in situations where a narrow path is present, such as a one-tile width bridge where two robots cannot pass each other. Without transferring, one robot must wait for the other to get out of the way. An example of (3) being met is shown in Fig. 5.

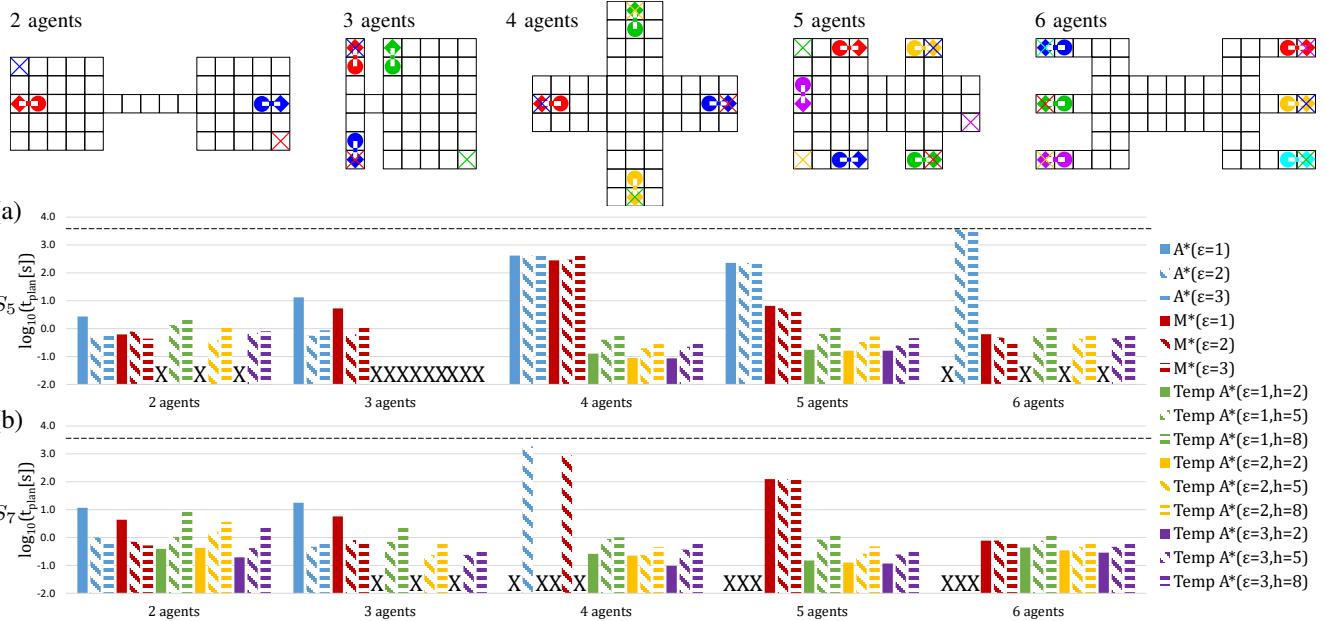


Fig. 6. Planning time t_{plan} for A^* , M^* , and our Temporal A^* on crowded maps with increasing number of robots, shown at the top. The planners must move the robots' front feet (the circles) to the respective colored \times , the back foot can end up anywhere as long as it is a valid position. Results with S_5 in (a) and S_7 in (b). An X indicates the planner was unable to find a solution within the time limit (dashed line). Plots are shown in log scale.

Finding candidate locations where load transfer can take place involves checking valid positions where the tile can be placed, usually along the outer edges of the polyomino. A candidate location closer to the robot dropping off the tile is better, since dropping off a tile takes time. Structure connectivity remains a priority when performing load transfers.

Because this strategy can be computationally expensive with larger configurations, we limit the search to pairs of robots that are within a Manhattan distance d_{tr} of each other. In the case of multiple robots wanting to transfer loads with the same robot, the transfer that lowers cost the most is chosen. The planner used can be our Temporal A^* , M^* , or any other that finds paths to the robots' goals.

The goal selection is revisited every time step to check for load transfers. The planned paths are saved until the goals of the robots change, to avoid unnecessary computation.

V. RESULTS

A. Planner comparison

We first measure the time required for planning by our Temporal A^* implementation, regular A^* , and the M^* variant. For five different maps with increasing number of robots, the three planners are tasked with finding paths to predetermined goal positions. Tests are conducted for movesets S_5 and S_7 , which are illustrated in Fig. 3. Each test is given a maximum of one hour to solve. Results are shown in Fig. 6.

The maps are intentionally crowded to test the planners in challenging scenarios. From the results, it can be seen that larger values of ε (a weighted heuristic [15]) do not substantially decrease the time required by regular A^* or M^* for most maps. Regular A^* takes longer than M^* for all maps, and fails consistently for the tests with S_7 . On the contrary, Temporal A^* consistently finds solutions much faster thanks

to the frequent reduction in neighboring nodes expanded. Larger ε do improve planning times in some cases. Larger H values make Temporal A^* take longer, which is expected as it means more steps are planned and thus considered when moving lower priority robots. However, lower values can prevent the planner from finding a solution, as can be seen in maps 1 and 5 with S_5 , and in map 2 with S_7 .

The performance of the Temporal A^* in terms of time steps required for reconfiguration of the maps, as well as the tiles traveled by all the robots, are shown in Table I for the S_7 case. For the maps with 4 and 6 agents, Temporal A^* can finish builds faster than both regular A^* and M^* . In the other maps it takes slightly longer compared to other planners that found a solution. In terms of total tiles traveled, Temporal A^* consistently makes the agents travel longer paths as it sacrifices optimality to keep the planning decoupled. This can be an important consideration for power requirements, for example. Both of these issues can be partially addressed through the load transferring method. As mentioned before, the size of S is important for Temporal A^* to find solutions. It was not able to solve map 2 with S_5 for any set of parameters, because the limited moveset makes it impossible for the blue and red robots to let the other through, even if they switch priorities. S_7 includes motions where the robots can turn their back feet first, offering more flexibility for the robots to get out of the way.

Unlike regular A^* and M^* , where larger movesets increase planning time by increasing the number of neighbors each vertex has, for Temporal A^* a larger moveset can actually decrease planning times. Its lower planning time performance makes it viable to use first and then switch to another planner like M^* if a solution is not found.

TABLE I: RECONFIGURATION PERFORMANCE (S_7)

	Time steps					Tiles traveled				
	$m = 2$	$m = 3$	$m = 4$	$m = 5$	$m = 6$	$m = 2$	$m = 3$	$m = 4$	$m = 5$	$m = 6$
$A^*(\epsilon = 1)$	23	10	N/A	N/A	N/A	36	25	N/A	N/A	N/A
$A^*(\epsilon = 2)$	23	17	24	N/A	N/A	36	25	46	N/A	N/A
$A^*(\epsilon = 3)$	23	17	N/A	N/A	N/A	36	25	N/A	N/A	N/A
$M^*(\epsilon = 1)$	23	10	N/A	13	28	36	25	N/A	56	104
$M^*(\epsilon = 2)$	23	17	24	13	28	36	25	46	56	106
$M^*(\epsilon = 3)$	23	17	N/A	13	28	36	25	N/A	56	106
Temp $A^*(\epsilon = 1, h = 2)$	30	N/A	14	17	28	59	N/A	50	80	165
Temp $A^*(\epsilon = 1, h = 5)$	24	11	13	14	22	45	28	49	68	130
Temp $A^*(\epsilon = 1, h = 8)$	24	11	13	14	22	45	28	49	68	130
Temp $A^*(\epsilon = 2, h = 2)$	30	N/A	14	17	28	59	N/A	50	80	165
Temp $A^*(\epsilon = 2, h = 5)$	24	11	13	14	22	45	28	49	68	130
Temp $A^*(\epsilon = 2, h = 8)$	24	11	13	14	22	45	28	49	68	130
Temp $A^*(\epsilon = 3, h = 2)$	30	N/A	14	17	28	59	N/A	50	80	165
Temp $A^*(\epsilon = 3, h = 5)$	24	11	14	14	22	45	28	50	68	130
Temp $A^*(\epsilon = 3, h = 8)$	24	11	14	14	22	45	28	52	68	103

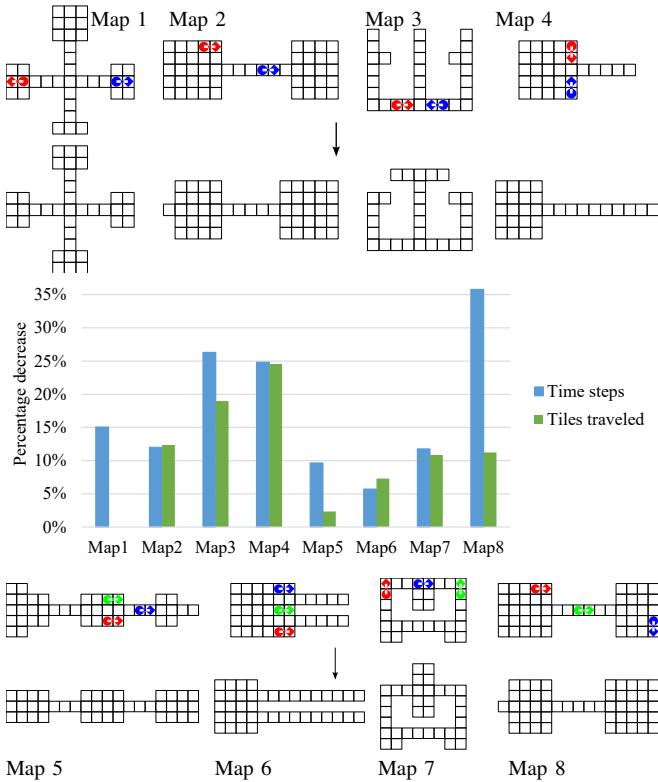


Fig. 7. The load transfer strategy is tested on several maps. The initial configurations are shown on top and the final ones below them (the final positions of the robots do not matter as long as they are valid). The percentage decrease in time steps and tiles traveled, when employing load transferring, is shown in the plot.

B. Load transfer

Load transfer often shortens total tiles traveled by the robots and time steps required for building, as illustrated in Fig. 7. Different maps were created to perform a building sequence with and without load transfer enabled. These maps contain at least one long, single tile segment where the probability for deadlocks is high. Maps 1 to 4 have two robots, and maps 5 to 8 have three. Four tiles are moved in all maps with the exception of maps 3 and 6, where three and ten tiles are moved.

In these tests, our Temporal A* is used as the planner, with $\epsilon = 1$ and $h = 10$, and $d_{tr} = 5$. Load transferring's impact is how much it decreases the required time steps and total tiles traveled, compared to not employing it.

The results for map 1 are interesting in that only the time steps were reduced. The same number of tiles were traveled but the robots did not wait as much for paths to be cleared. For the rest of the maps both metrics were reduced, significantly so for maps 3, 4 and 8.

Although all maps used in these tests allow the building sequence to be completed without cooperation, there are many configurations that require cooperation to be completed, i.e. it is impossible to avoid deadlocks otherwise.

VI. CONCLUSIONS AND FUTURE WORK

We showed that several variants of the BILL-E RECONFIGURATION problem are NP-complete. In particular, the cooperative variant is proven to be hard, even for two robots. We compared three planners, and tested a load transfer strategy to reduce planning time and building time. The three planners are the regular A* algorithm, a multi-robot variant M^* , and a priority-based Temporal A*.

The Temporal A* is faster at planning paths, and can benefit from larger configuration spaces unlike the other two. However, it is not complete and depends on the size of the moveset. In some maps it finished the reconfiguration faster, although it usually results in a larger combined path cost (tiles traveled). While the load transfer can significantly reduce travel costs and time required, it can be computationally expensive so it is limited to pairs of robots within a predefined distance of each other. Overall, we were able to plan paths and perform building sequences more efficiently utilizing both methods discussed.

Future work could focus on combining the temporal A* planning with the randomized sampling methods presented in our prior work, to further resolve deadlock situations in complex and crowded maps. We are also interested in extending these algorithms to 3D structures. Finally, it might be worth considering a transfer to distributed computation of local motion plans, as indicated in Section II-B.

REFERENCES

- [1] A. Abdel-Rahman, A. T. Becker, D. E. Biediger, K. C. Cheung, S. P. Fekete, N. A. Gershenfeld, S. Hugo, B. Jenett, P. Keldenich, E. Niehs, C. Rieck, A. Schmidt, C. Scheffer, and M. Yannuzzi, "Space ants: Constructing and reconfiguring large-scale structures with finite automata," in *Symposium on Computational Geometry (SoCG)*, 2020, pp. 73:1–73:7.
- [2] E. D. Demaine, S. P. Fekete, P. Keldenich, C. Scheffer, and H. Meijer, "Coordinated motion planning: Reconfiguring a swarm of labeled robots with bounded stretch," *SIAM Journal on Computing*, vol. 48, pp. 1727–1762, 2019.
- [3] S. P. Fekete, P. Keldenich, R. Kosfeld, C. Rieck, and C. Scheffer, "Connected coordinated motion planning with bounded stretch," *Autonomous Agents and Multi-Agent Systems*, vol. 37, no. 2, 2023.
- [4] S. P. Fekete, P. Kramer, C. Rieck, C. Scheffer, and A. Schmidt, "Efficiently reconfiguring a connected swarm of labeled robots," in *International Symposium on Algorithms and Computation (ISAAC)*, 2022, pp. 17:1–17:15.
- [5] S. P. Fekete, E. Niehs, C. Scheffer, and A. Schmidt, "Connected reconfiguration of lattice-based cellular structures by finite-memory robots," *Algorithmica*, vol. 84, no. 10, pp. 2954–2986, 2022.
- [6] J. Garcia, M. Yannuzzi, P. Kramer, C. Rieck, and A. T. Becker, "Connected reconfiguration of polyominoes amid obstacles using RRT*," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2022, pp. 6554–6560.
- [7] C. E. Gregg, D. Catanoso, O. I. B. Formoso, I. Kostitsyna, M. E. Ochalek, T. J. Olatunde, I. W. Park, F. M. Sebastianelli, E. M. Taylor, G. T. Trinh, and K. C. Cheung, "Ultralight, strong, and self-reprogrammable mechanical metamaterials," *Science Robotics*, vol. 9, no. 86, 2024.
- [8] A. Itai, C. H. Papadimitriou, and J. L. Szwarcfiter, "Hamilton paths in grid graphs," *SIAM Journal on Computing*, vol. 11, no. 4, pp. 676–686, 1982.
- [9] B. Jenett, A. Abdel-Rahman, K. Cheung, and N. Gershenfeld, "Material–robot system for assembly of discrete cellular structures," *IEEE Robotics and Automation Letters*, vol. 4, no. 4, pp. 4019–4026, 2019.
- [10] B. Jenett, C. Gregg, D. Cellucci, and K. Cheung, "Design of multifunctional hierarchical space structures," in *IEEE Aerospace Conference*, 2017, pp. 1–10.
- [11] S. M. LaValle, *Planning Algorithms*. Cambridge University Press, 2006.
- [12] E. Niehs, A. Schmidt, C. Scheffer, D. E. Biediger, M. Yannuzzi, B. Jenett, A. Abdel-Rahman, K. C. Cheung, A. T. Becker, and S. P. Fekete, "Recognition and reconfiguration of lattice-based cellular structures by simple robots," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2020, pp. 8252–8259.
- [13] M. Otte and N. Correll, "Dynamic teams of robots as ad hoc distributed computers: Reducing the complexity of multi-robot motion planning via subspace selection," *Autonomous Robots*, vol. 42, no. 8, pp. 1691–1713, 2018.
- [14] J.-H. Peng, I.-H. Li, Y.-H. Chien, C.-C. Hsu, and W.-Y. Wang, "Multi-robot path planning based on improved D*Lite algorithm," in *IEEE International Conference on Networking, Sensing and Control (ICNSC)*, 2015, pp. 350–353.
- [15] I. Pohl, "First results on the effect of error in heuristic search," *Machine Intelligence*, vol. 5, pp. 219–236, 1970.
- [16] D. Silver, "Cooperative pathfinding," in *AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment*, vol. 1, no. 1, 2005, pp. 117–122.
- [17] J. Song, Z. Li, P. Wang, T. Meyer, C. Mao, and Y. Ke, "Reconfiguration of DNA molecular arrays driven by information relay," *Science*, vol. 357, no. 6349, 2017.
- [18] P. Thalamy, B. Piranda, and J. Bourgeois, "Engineering efficient and massively parallel 3d self-reconfiguration using sandboxing, scaffolding and coating," *Robotics and Autonomous Systems*, vol. 146, p. 103875, 2021.
- [19] G. Wagner and H. Choset, "M*: A complete multirobot path planning algorithm with performance bounds," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2011, pp. 3260–3267.
- [20] H. Wang and M. Rubenstein, "Shape formation in homogeneous swarms using local task swapping," *IEEE Transactions on Robotics*, vol. 36, no. 3, pp. 597–612, 2020.