



Connected coordinated motion planning with bounded stretch

Sándor P. Fekete¹ · Phillip Keldenich¹ · Ramin Kosfeld¹ · Christian Rieck¹ · Christian Scheffer²

Accepted: 13 September 2023 / Published online: 17 October 2023
© The Author(s) 2023

Abstract

We consider the problem of connected coordinated motion planning for a large collective of simple, identical robots: From a given start grid configuration of robots, we need to reach a desired target configuration via a sequence of parallel, collision-free robot motions, such that the set of robots induces a connected grid graph at all integer times. The objective is to minimize the *makespan* of the motion schedule, i.e., to reach the new configuration in a minimum amount of time. We show that this problem is **NP**-complete, even for deciding whether a makespan of 2 can be achieved, while it is possible to check in polynomial time whether a makespan of 1 can be achieved. On the algorithmic side, we establish simultaneous constant-factor approximation for two fundamental parameters, by achieving *constant stretch* for *constant scale*. Scaled shapes (which arise by increasing all dimensions of a given object by the same multiplicative factor) have been considered in previous seminal work on self-assembly, often with unbounded or logarithmic scale factors; we provide methods for a generalized scale factor, bounded by a constant. Moreover, our algorithm achieves a *constant stretch factor*: If mapping the start configuration to the target configuration requires a maximum Manhattan distance of d , then the total duration of our overall schedule is $\mathcal{O}(d)$, which is optimal up to constant factors.

Keywords Motion planning · Parallel motion · Group coordination · Swarm robotics · Bounded stretch · Scaled shape · Makespan · Connectivity

1 Introduction

Coordinating the motion of a set of objects is a fundamental problem that occurs in a large spectrum of theoretical contexts and practical applications. A typical challenge arises from relocating a large collection of agents from a given start into a desired goal configuration in an efficient manner, while respecting a number of natural constraints, such as avoiding collisions or disrupting the coherence of the arrangement. This problem was also the subject of the 2021 Computational Geometry Challenge, highlighting the high relevance

A preliminary abstract appears in the *32nd International Symposium on Algorithms and Computation (ISAAC)* [20]. This version contains the omitted proofs and further details.




Fig. 1 (Top left) An autonomous, sphere-shaped catom, changing location by rotating around a second catom used as a pivot [30]. (Bottom left) A collective of catoms building a wall [30]. (Right) A configuration of catoms in the process of building a scaffold structure [41]

for the algorithmic community; see [21] for an overview and [9, 27, 44] for successful contributions.

In this paper, we consider a *connected* configuration of objects, e.g., a (potentially large) collective of mobile robots or blocks of building material that can be moved by a large group of robots, which needs to be transformed into a desired target configuration by a sequence of parallel, collision-free motions that keeps the overall arrangement connected at all integer times. Problems of this type occur in many contexts requiring relocation of autonomous agents; the connectivity constraint arises naturally, e.g., for assemblies in space, where disconnected pieces cannot regain connectivity, or for small-scale swarm robots (such as *catoms* in *claytronics* [22]) which need connectivity for local motion, electric power and communication; see Fig. 1.

A crucial algorithmic aspect is *efficiency*: How can we coordinate the robot motions, such that a target configuration is reached in timely or energy-efficient manner? Most previous work has largely focused on sequential schedules, where one robot moves at a time, with objectives such as minimizing the number of moves. In practice, however, robots usually move simultaneously, so we desire a *parallel* motion schedule, with a natural objective of minimizing the time until completion, called *makespan*. How well can we exploit parallelism in a robot swarm to achieve an efficient schedule? As illustrated in Fig. 2, this is where the connectivity constraints make a tremendous difference.

A critical parameter in self-assembly is the robustness of the involved shapes, corresponding to sufficient local connectivity to prevent fragility. This leads to the concept of *scaled shapes*; intuitively, a scale factor of c corresponds to replacing each pixel of a polyomino shape by a quadratic $c \times c$ array of pixels. This has fundamental connections to Kolmogorov and runtime complexity, as shown by Soloveichik and Winfree [36]: “Furthermore, the independence of scale in self-assembly theory appears to play the same crucial role as the independence of running time in the theory of computability... [we] show that the running-time complexity, with respect to Turing machines, is polynomially equivalent to the scale complexity of the same function implemented via self-assembly by a finite set of tile types.” As a consequence, limiting scale has received considerable attention: While Soloveichik and Winfree established unbounded scale in general self-assembly, other work has managed to achieve logarithmic and even (in specific scenarios) constant scale.

As we demonstrate in this paper, achieving optimal makespan for connected reconfiguration is provably hard, even in relatively basic cases. On the positive side, we present methods that are capable of achieving a constant-factor approximation, assuming not

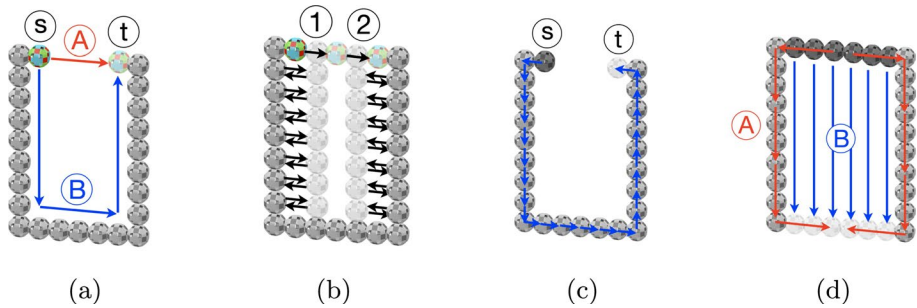


Fig. 2 Reconfiguration with and without connectivity constraints. **a** Relocating the colored particle from s to t , without (red trajectory A) and with connectivity constraint (blue trajectory B). **b** Coordinating many particles to quickly deliver a specific particle to a desired location, while preserving connectivity. Both vertical lines of robots move in parallel towards each other, staying connected to the horizontal line. **c** Reconfiguring an arrangement of *identical* particles in a single, parallel, connected step. In this case, it is possible that the whole arrangement rotates one step, as the robots are indistinguishable; the top-rightmost robot moves onto the goal location and all other robots follow. **d** Reconfiguring an arch-shaped arrangement of identical particles into a U-shaped one, without (motion plan A, shown in red) and with connectivity (motion plan B, shown in blue). Motion plan A splits the configuration, resulting in a makespan of length half the amount of robots in the top horizontal row. In motion plan B everything stays connected, but the horizontal row has to move all the way to the bottom, what results in longer makespan (Color figure online)

more than a generalization of constant scale of start and target configurations. In fact, our method realizes *constant stretch*: If mapping the start configuration to the target configuration requires a maximum Manhattan distance of d , then the total duration of our overall schedule is $\mathcal{O}(d)$. As can be seen from Fig. 2a (where d corresponds to trajectory A), this is less straightforward than in a non-connected setting, even in very basic instances. Instead, this quickly requires coordination of the whole arrangement, as sketched in Fig. 2b.

1.1 Our results

We provide a spectrum of new results for questions arising from efficiently reconfiguring a connected, unlabeled collective of robots from a start configuration C_s into a target configuration C_t , aiming for minimizing the overall makespan and maintaining connectivity in each step. We obtained all of our results in obstacle-free 2D space, and discuss possible extensions as well as open questions in the conclusion. Some of these results are featured in a video [4].

- Deciding whether there is a schedule with a makespan of 1 transforming C_s into C_t can be done in polynomial time, see Theorem 1.
- Deciding whether there is a schedule with a makespan of 2 transforming C_s into C_t is NP-complete, see Theorem 2. This implies NP-hardness of approximating the minimum makespan within a constant of $(\frac{3}{2} - \epsilon)$, for any $\epsilon > 0$, see Corollary 2.
- As our main result, we show that there is a constant c^* such that for any pair of start and target configurations with a (generalized) scale of at least c^* , a schedule with constant stretch can be computed in polynomial time, see Theorem 3 and Corollary 3. This implies that there is a constant-factor approximation for the problem of computing

schedules with minimal makespan restricted to pairs of start and target configurations with a scale of at least c^* , see Corollary 4.

1.2 Related work

Coordinating the motion of many agents plays a central role when dealing with large numbers of moving robots, vehicles, aircraft, or people. How can each agent choose an efficient route that avoids collisions with other agents as they simultaneously move to their destinations? These basic questions arise in many applications, such as ground swarm robotics [32, 33], aerial swarm robotics [8, 43], air traffic control [11], and vehicular traffic networks [19, 34].

Multi-robot coordination dates back to the early days of robotics and computational geometry. The seminal work by Schwartz and Sharir [35] from the 1980 s considers coordinating the motion of disk-shaped objects among obstacles. Their algorithms are polynomial in the complexity of the obstacles, but exponential in the number of disks. Hopcroft et al. [24] and Hopcroft and Wilfong [25] proved PSPACE-completeness of moving multiple robots to a target configuration, showing the significant challenge of coordinating many robots.

There is a vast body of other related work dealing with multi-robot motion planning, both from theory and practice. For a more extensive overview, see [13]. In both discrete and geometric variants of the problem, the objects can be *labeled*, *colored* or *unlabeled*. In the *labeled* case, the objects are all distinguishable and each object has its own, uniquely defined target position. In the *colored* case, the objects are partitioned into k groups and each target position can only be covered by an object with the right color. This was considered by Solovey and Halperin [37], who present and evaluate a practical sampling-based algorithm. In the *unlabeled* case, objects are indistinguishable and target positions can be covered by any object. This was first considered by Kloder and Hutchinson [26], who presented a practical sampling-based algorithm. Turpin et al. [42] give an algorithm for finding a solution in polynomial time, if one exists. This is optimal with respect to the longest distance traveled by any one robot, but only holds for disk-shaped robots under additional restrictive assumptions on the free space. For unit disks and simple polygons, Adler et al. [1] provide a polynomial-time algorithm under the additional assumption that the start and target positions have some minimal distance from each other. Under similar separability assumptions, Solovey et al. [39] provide a polynomial-time algorithm that produces a set of paths that is no longer than $\text{OPT} + 4m$, where m is the number of robots, and OPT denotes the total length of a set of paths of an optimal solution. However, they do not consider the makespan, but only the total path length. On the negative side, Solovey and Halperin [38] prove that the unlabeled multiple-object motion planning problem is PSPACE-hard, even when restricted to unit square objects in a polygonal environment.

For an extensive overview of multi-agent path planning, refer to [40]. Yu and LaValle [45] discuss the relationship of multi-agent path planning and flow problems in *collision-free unit-distance graphs*. These are unit-distance graphs having the additional property that two discs of radius $\sqrt{2}/4$ do not collide when traveling with unit speed on two paths that do not contain the same vertex for the same time step. They consider different problems settings and show, among other results, that if the goal locations can be assigned to arbitrary agents, a solution always exists and the longest path has at most $n + V - 1$ edges, where n and V are the number of agents and vertices of the graph, respectively.

They also mention NP-hardness of the decision problem in case goals are pre-assigned to agents. Charrier et al. [5–7] study reachability and coverage planning problems for connected agents. They show different complexity results for different *topological graphs*, consisting of vertices and two kinds of edges, one for robot motion, the other for communication. They also introduce *sight-movable graphs*, which are undirected topological graphs for which there is a movable path between any pair of vertices if these vertices can communicate with each other. They show that this class admits efficient algorithms for the reachability and coverage problem. Queffelec, Sankur, and Schwarzenrüber [31] study the connected multi-agent path finding problem in partially known environments in which the graph is not known entirely in advance, and show PSPACE-completeness of the problem. Despite the fact that all of this is related to our work, a crucial difference is that we consider the stretch factor as the main performance measure.

There is also a wide range of practical related work. Self-configuration of robots as active agents was studied by Naz et al. [29]. A basic model in which robots are used as building material was introduced by Derakhshandeh et al. [16, 17]. This resembles Claytronics robots like Catoms, see Goldstein and Mowry [22]. In more recent work, Thalamy et al. [41] consider using scaffolding structures for asynchronous reconfiguration.

For an instance of parallel reconfiguration, a lower bound for the time required for *all* robots to reach their destinations is the time it takes to move just *one* robot to its destination in the absence of other robots, i.e., by the maximum distance between a robot's origin and destination. Moving a dense arrangement of robots to their destinations while avoiding collisions may require substantially more time than this lower bound. This motivates the *stretch factor*, which is defined to be the ratio of the time taken by a parallel motion plan divided by the simple lower bound.

In recent work, Demaine et al. [13] provide several fundamental insights into these problems of coordinated motion planning for the scenario with labeled robots without a connectivity constraint. They were able to develop algorithms that (under relatively mild assumptions on the separation between robots) can achieve *constant* stretch factors that are independent of the number of robots. Thus, these algorithms provide an absolute performance guarantee on the makespan of the parallel motion schedule, which implies that the schedule is a constant-factor approximation of the best possible schedule. For densely packed arrangements of robots (without separation assumptions), they prove that a constant stretch factor is no longer possible, and give upper and lower bounds on the worst-case stretch factor. Note that the approaches of [13] cannot be adapted, thus our demand for connectivity requires new algorithmic ideas.

In the methods developed in [3, 13], elementary pieces can achieve arbitrary relative configurations, as shown in Fig. 3, in which colors indicate the final destinations of components.

As discussed in [3, 13], it is straightforward to adapt this to geometric reconfigurations: “Filled” pixels (corresponding to material) get assigned one color (e.g., “black”), while “empty” pixels (corresponding to void) get a second color (e.g., “white”). However, in real-world scenarios (such as in space, or for swarm robots as shown in Figs. 1 and 2), basic components need to stay connected to keep them from drifting apart. How can we develop efficient parallel *connected* reconfiguration schedules, in which all “black” pixels remain connected throughout the process? (Note that we aim for efficient schedules, i.e., *constant stretch*, so approaches like the recent one by Akitaya et al. [2] for transforming any arrangement into the same base configuration are insufficient.)

The concept of scale complexity has received a considerable amount of attention in the context of self-assembly; in all settings, achieving constant scale has required special cases

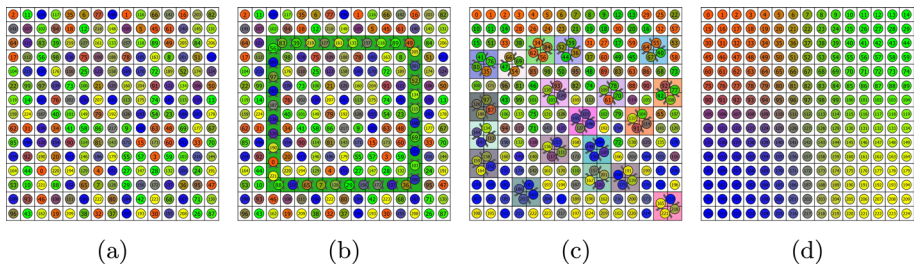


Fig. 3 Parallel reconfiguration, established by [13]: **a** Start configuration. **b** A feasible, parallel reconfiguration move. **c** Parallel reconfiguration moves. **d** Target configuration. See <https://www.ibr.cs.tu-bs.de/users/fekete/Videos/CoordinatedMotionPlanning.mp4> for a video [3]

or operations. As shown by Soloveichik and Winfree [36], the minimal number of distinct tile types necessary to self-assemble a shape, at some scale, can be bounded both above and below in terms of the shape's Kolmogorov complexity, leading to unbounded scale in general. As shown by Demaine et al. [15], allowing additional operations that allow destroying tiles can be exploited to achieve a scale that is only bounded by a logarithmic factor, beating the linear bound without such operations. In a setting of recursive, multi-level *staged* assembly with a logarithmic number of stages (i.e., “hands” for handling subassemblies), Demaine et al. [12] achieved logarithmic scale, and constant scale for more constrained classes of polyomino shapes; this was later improved by Demaine et al. [14] to constant scale for a logarithmic number of stages. More recently, Luchsinger et al. [28] employed repulsive forces between tiles to achieve constant scale in two-handed self-assembly.

2 Preliminaries

We consider *robots* at integer grid positions. A set of n unlabeled robots forms a *configuration* C , corresponding to a vertex-induced subgraph H of the infinite integer grid, with an edge between two grid vertices $v_1, v_2 \in C$ if and only if v_1 and v_2 are on adjacent grid positions, i.e., a distance of 1 apart. A configuration is *connected*, if H is connected. Two configurations C_1 and C_2 *overlap*, if they have at least one position in common. Two robots are *adjacent* if their positions v_1, v_2 are adjacent, and *diagonally adjacent* if their positions are adjacent with a common vertex v such that (v_1, v) and (v, v_2) lie orthogonal.

A robot can move in discrete time steps by changing its location from a grid position v to an adjacent grid position w ; denoted by $v \rightarrow w$. Two moves $v_1 \rightarrow w_1$ and $v_2 \rightarrow w_2$ are called *collision-free* if $v_1 \neq v_2$ and $w_1 \neq w_2$. A *transformation* between two configurations $C_1 = \{v_1, \dots, v_n\}$ and $C_2 = \{w_1, \dots, w_n\}$ is a set of collision-free moves $\{v_i \rightarrow w_i \mid i = 1, \dots, n\}$. Furthermore, robots are allowed to remain in their current position, i.e., we allow the move $v \rightarrow v$. However, even though this is unproblematic for the unlabeled scenario, we do not allow *swaps* between robots, i.e., the two parallel moves $v_1 \rightarrow v_2$ and $v_2 \rightarrow v_1$ cause a collision.

For $M \in \mathbb{N}$, a *schedule* is a sequence $C_1 \rightarrow \dots \rightarrow C_{M+1}$ (also denoted as $C_1 \Rightarrow C_{M+1}$) of transformations, with a *makespan* of M . A *stable schedule* $C_1 \Rightarrow_{\chi} C_{M+1}$ uses only connected configurations. Let C_s, C_t be two connected configurations with equally many robots called *start* and *target configuration*, respectively. A *matching* is a one-to-one mapping between vertices from C_s and C_t . The *diameter* of a matching is the maximal Manhattan

distance between two matched vertices. A *bottleneck matching* is, among all possible matchings, a matching with minimal diameter. The *diameter* d of (C_s, C_t) is the diameter of a bottleneck matching. The *stretch (factor)* of a (stable) schedule is the ratio between its makespan M and the diameter d of (C_s, C_t) .

We consider the CONNECTED COORDINATED MOTION PLANNING PROBLEM that is stated as follows; see Fig. 2 for an illustration. Given a pair (C_s, C_t) of unlabeled connected start and target configurations, and an integer k , we are asked to decide whether there is a stable schedule with a makespan of k that transforms C_s into C_t .

3 Makespan 1 and 2

As a first observation we note that it can be decided in polynomial time whether there is a schedule $C_s \rightarrow C_t$ with a makespan of 1 between a start and a target configuration.

Theorem 1 *For a pair of configurations C_s and C_t , each with n vertices, it can be decided in polynomial time whether there is a schedule with a makespan of 1 transforming C_s into C_t .*

Proof Given two connected configurations C_s and C_t , each with n vertices. We compute the bipartite graph $G_{C_s, C_t} = (V_s \cup V_t, E)$, where V_s and V_t consist of all occupied positions in C_s and C_t . For E , we add an edge if and only if an occupied position in C_t is adjacent (or identical) to an occupied position in C_s .

Consider a perfect matching in G_{C_s, C_t} . By construction, the edges in G_{C_s, C_t} only connect positions from start and target configuration that are at most one unit step apart. As there is a perfect matching, no two robots want to occupy the same position. Furthermore, two paths of length 1 can only cross at a common vertex; this cannot happen, as we consider a perfect matching; therefore, all robots can move along their respective matching edges in parallel without collisions.

If there is no perfect matching in G_{C_s, C_t} , at least one robot would have to move to a position further away. Thus, a makespan of 1 would not be achievable. So, there is a schedule of makespan 1 if and only if G_{C_s, C_t} admits a perfect matching. Because the graph is sparse, this can be checked in $\mathcal{O}(n^{1.5})$ time, using the method of Hopcroft and Karp [23]. \square

Note that, because C_s and C_t have to be connected, a schedule with a makespan of 1 is always stable. Furthermore, it is easy to see that the method described in the proof of Theorem 1 can be applied iteratively to verify whether a suggested schedule with makespan $k \in \mathbb{N}$ is stable, and indeed transforms the given start configuration into the respective goal configuration. Hence, we conclude that the problem is contained in NP.

Corollary 1 *For any pair of configurations, any given (stable) schedule with makespan $k \in \mathbb{N}$ can be verified in polynomial time, i.e., this problem is in NP.*

However, even for a makespan of 2, the same problem becomes provably difficult. We show the following theorem.

Theorem 2 *For a pair of configurations C_s and C_t , each with n vertices, deciding whether there is a stable schedule with a makespan of 2 transforming C_s into C_t is NP-complete.*

The proof is based on a reduction from the NP-hard problem PLANAR MONOTONE 3SAT [10], which asks to decide whether a Boolean 3-CNF formula φ is satisfiable, for which in each clause the literals are either all unnegated or all negated. Note that a general instance of PLANAR MONOTONE 3SAT may contain clauses with three literals as well as clauses with only two literals.

For the following, we refer to Fig. 4. The reduction works as follows: For every instance φ of PLANAR MONOTONE 3SAT, we construct an instance I_φ , consisting of a start configuration C_s and a target configuration C_t . In the figure, we use three different colors to indicate occupied positions in the start configuration (red), in the target configuration (dark cyan), and in both configurations (gray). Therefore, we consider a rectilinear planar embedding of the variable-clause incidence graph G_φ of φ where the variable vertices are placed horizontally in row, and clauses containing unnegated and negated literals are placed above and below, respectively. All variables of φ are represented by a horizontal *variable gadget* (light red). Furthermore, we position two additional *auxiliary gadgets* (light blue) at the top and at the bottom boundary of the instance, which are connected to the variable gadget via bridges at the right boundary. There will be a *separation gadget* (yellow) between each adjacent and nested pair of *clause gadgets* (blue). All clause gadgets are connected via bridges to separation gadgets and possibly to the auxiliary gadgets. Further, there are bridges from a clause gadget to the respectively contained variables.

There is a stable schedule for I_φ transforming the start configuration C_s into the target configuration C_t with a makespan of 2, if and only if φ is satisfiable. In order to transform C_s into C_t , the separation gadgets ensure that in the single intermediate configuration, all clause and auxiliary gadgets are disconnected from each other. Therefore, to satisfy the connectivity constraint, some robots of the variable gadget need to move to provide connectivity between the variable gadget and the clause gadgets. At the same time, our construction ensures that robots representing a variable can either keep connectivity to their unnegated or to their negated literal containing clauses within a makespan of 2, because otherwise the connectivity within the variable gadget would be broken; thus, these movements can be used to determine a valid variable assignment for φ .

3.1 Construction of the gadgets

We start with some observations regarding the target neighborhood of a given robot, i.e., the positions that can be reached by moving a robot twice from a given start position.

Observation 1 For a given position, the set of positions reachable within two moves is given by Fig. 5a.

Because we can swap the roles of the start and the target configurations, an easy consequence is that if there is only one start position in the reachable neighborhood of a target position, then the robot occupying this position in the target configuration is uniquely defined.

In our construction, we place gadgets in such a way that they are “sufficiently far apart”. The reason for this is that robots that belong to a certain gadget only can occupy target positions within the same gadget. In particular, we place our gadgets in such a

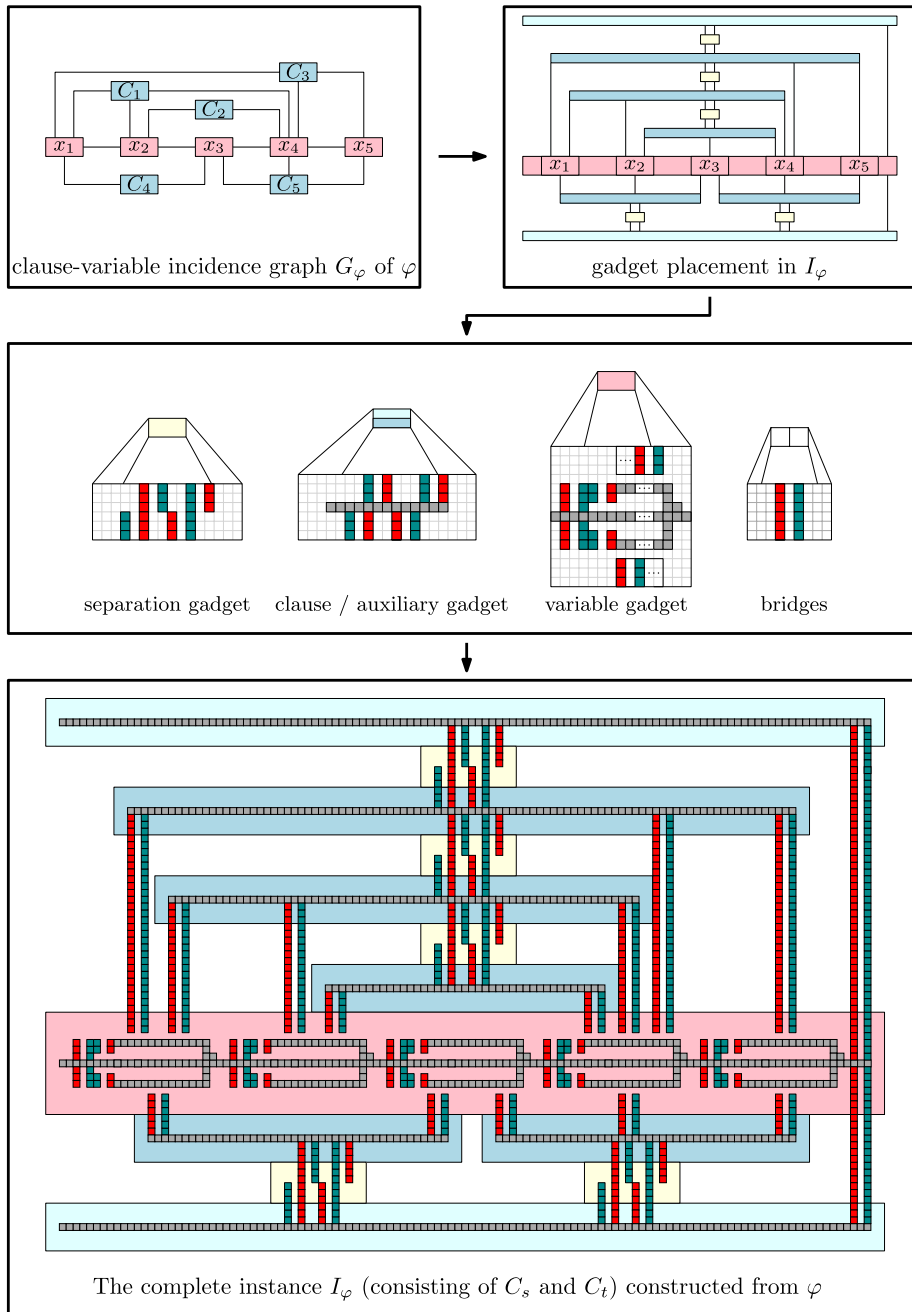



Fig. 4 Symbolic overview of the NP-hardness reduction. The depicted instance is due to the PLANAR MONOTONE 3SAT formula $\varphi = (x_1 \vee x_2 \vee x_4) \wedge (x_2 \vee x_4) \wedge (x_1 \vee x_4 \vee x_5) \wedge (\bar{x}_1 \vee \bar{x}_3) \wedge (\bar{x}_3 \vee \bar{x}_4 \vee \bar{x}_5)$. We use three different colors to indicate occupied positions in the start configuration (red), in the target configuration (dark cyan), and in both configurations (gray) (Color figure online)

Fig. 5 **a** All positions that can be reached by moving the robot initially located on the red square twice. **b** This figure shows a bridge. Note that all white squares have to be free in C_s as well as in C_t (Color figure online)



way that a robot belonging to some gadget would need at least 6 steps to reach another gadget. This allows us to make statements about the movement within gadgets, as well as their interaction with others. We refer to all positions that belong to a gadget as \mathcal{G} , and to the positions between gadgets as \mathcal{B} . In order for the gadgets to interact with each other, we place robots in \mathcal{B} (the bridges) in such a way that their respective target positions are uniquely defined in \mathcal{B} . With this, it is easy to see that gadgets are indeed independent from each other. We call a gadget *solvable*, if there is a schedule that transforms the start configuration into the target configuration with a makespan of 2, and each robot in \mathcal{G} is connected to a robot in \mathcal{B} . Note that we do not necessarily require the intermediate configuration to be connected in order to say that a gadget is solvable; this will not be the case in the separation gadget. Consider any configuration C that is composed of gadgets constructed in this manner, then a necessary condition for the existence of a schedule for C with a makespan of 2 is that every gadget is solvable.

In the following, we introduce different gadgets that are used in the proof showing NP-hardness of the problem. We show that these gadgets are solvable, and we indicate explicit schedules.

Line Gadget A *line gadget* is a rectangular region of size $\ell \times 7$. It consists of $\ell - 4 \geq 3$ horizontally adjacent robots that occupy the same positions in both, the start and the target configuration, as well as a variable number of bridges of length three, as depicted in Fig. 6.


In the NP-hardness proof, line gadgets are used in two ways: As *clause gadgets*, they represent the logic of the satisfiability instance; as *auxiliary gadgets*, they ensure connectivity of the start and the target configuration, respectively.

Lemma 1 *The line gadget is solvable.*

Proof All robots realizing bridges move two steps to the right or to the left to their respective target positions. The line robots do not move at all. Because no robot moves vertically, all configurations are connected. \square

Separation Gadget A *separation gadget* is a rectangular region of size 14×6 , containing the start and the target configuration as shown in Fig. 7a.

Fig. 6 The gray positions indicate the line gadget. The colored positions indicate the start and target configurations of the respective bridges






Fig. 7 **a** shows the separation gadget. **b** visualizes the unique schedule with a makespan of 2

Lemma 2 *The separation gadget is solvable.*

Proof There is a unique schedule with a makespan of 2, as shown in Fig. 7b. This follows from applying Observation 1 to the different target positions. \square

In the intermediate configuration of the unique schedule solving a separation gadget, each pair of bridges is not connected within the gadget, i.e., the intermediate configuration is not connected. For that reason, these gadgets will be used to disconnect the connections between all clause gadgets. To guarantee a fully connected intermediate configuration, a specific movement within the variable gadget is necessary.

Variable Gadget Let n be the number of variables of the PLANAR MONOTONE 3SAT instance. The *variable gadget* is composed of the parts shown in Fig. 8, as follows: There is one *left end* followed by n horizontally aligned *variable arm segments* and one *right end*. In order to capture the respective number of clauses in which a variable is contained, we can independently adjust the width of each arm segment. Bridges are placed in the gray colored hatched parts. We refer to the top and the bottom parts of the arm segments as the *unnegated* and the *negated arm segment*, as they model the unnegated and negated literals, respectively. As an example, consider Fig. 9.




Fig. 8 The variable gadget consists of a number n variable arm segments, and exactly one left and right end each. Bridges to the respective clauses are placed in the hatched areas




Fig. 9 An example construction of the variable gadget with three variables. Note that each arm segment is of different width, due to the number of clauses in which the variable is contained

Lemma 3 *In the variable gadget, it is not possible that both, the unnegated and the negated arm segment of a variable arm segment are simultaneously connected to its respective bridges.*


Proof Without loss of generality, we assume that in the intermediate configuration the robots representing the unnegated arm segment are connected to the respective bridges. We use the labels shown in Fig. 10. We argue that the following properties hold after the first transformation.

- At least one robot of the unnegated arm segment (p_1, \dots, p_ℓ) moved up.
- Robot A_i moved left.
- Robot R_i moved up.

Assume for the sake of contradiction that no robot of the unnegated arm segment moved up. Because the top two bridge robots moved horizontally (by applying Observation 1), every movement of the bridge's bottommost robot would not yield a connection to the arm segment.

The robot A_i has two possible target positions, namely B_i and C_i . Assume that A_i moved down. Then, X_i has the unique target position Y_i , so it moved left—thus, it is isolated in the intermediate configuration. Therefore, p_1 has also moved left. It follows by induction that

Fig. 10 The figure shows the labeling of specific robots of a variable arm segment for the proof of Lemma 3. The positions p_i and n_i denote the unnegated and negated arm segment, respectively



every robot of the unnegated arm segment has moved left, which is a contradiction to the fact that at least one robot moved up. Thus, A_i moved left.

Assume that R_i did not move up. This is a contradiction to A_i having moved left, as otherwise at least A_i would be part of an isolated component in the intermediate configuration.

Similar arguments hold for the case that the robots of the negated arm segment are connected to its respective bridges. Therefore, R_i has to move both up and down in the first transformation. As this is impossible, this concludes the proof. \square

Lemma 4 *A variable gadget is solvable.*

Proof There is a valid schedule with a makespan of 2 given in Fig. 11. The depicted schedule connects the unnegated arm segment with the respective bridges; a similar schedule exists for the negated case. \square

3.2 Completing the NP-hardness reduction

Theorem 2 *For a pair of configurations C_s and C_t , each with n vertices, deciding whether there is a stable schedule with a makespan of 2 transforming C_s into C_t is NP-complete.*

Proof Consider a rectilinear planar embedding of the variable-clause incidence graph G_φ of a given PLANAR MONOTONE 3SAT formula φ . The horizontally aligned variables in G_φ are represented by a variable gadget—each variable is represented by an arm segment. For each clause, we introduce a line gadget with connecting bridges (and extend them if necessary) to the contained variables. We add two or three bridges according to whether the clause contains two or three literals, respectively. Between every nested pair of adjacent clauses, we introduce a separation gadget and connect it to the clauses via bridges. As a last step, we place two additional line gadgets at the top and bottom of the construction, connect them via bridges to the variable gadget, and via separation gadgets to the respective topmost and bottommost line gadgets that represent clauses. For an example, see the bottom part of Fig. 4.

Claim 1 *If the formula φ has a satisfying assignment α , then there is a valid intermediate configuration for I_φ (i.e., the intermediate configuration is connected); therefore, there is a stable schedule transforming C_s into C_t (both given by I_φ) with a makespan of 2.*




Fig. 11 The figure shows a stable schedule with a makespan of 2 as indicated in Lemma 4. Positions with less opacity show the respectively occupied positions in the previous step. **a** shows that start configuration, while **b** indicates the single intermediate configuration, and **c** depicts the respective target configuration

Proof Let α be a satisfying assignment of φ . A valid intermediate configuration for I_φ can be constructed as follows: All separation, line, and variable gadgets are transformed by the schedules given in Lemmas 1, 2 and 4. For the assignment of variable x_i the respective arm segment moves vertically. All bridge robots move to their respective target position due to Observation 1. In this configuration the helping line gadgets are directly connected to the variable gadget, and all clause gadgets are connected to the variable gadget as well. Because the disjoint parts of the separation gadgets are connected to a line gadget on one side, this configuration is connected.

Claim 2 *If there is a connected intermediate configuration for I_φ (such that there is a stable schedule transforming C_s into C_t (both given by I_φ) with a makespan of 2), then φ is satisfiable.*

Proof Because the intermediate configuration is connected, each clause gadget has to be connected to the variable gadget. This is only possible via the bridges at the boundary of the gadget. Due to the construction of the separation gadgets, the clause gadgets have to be *directly* connected to the variable gadget, i.e., a shortest path connecting a robot of each clause gadget with a robot of the variable gadget passes no other gadget of the construction. Thus, at least one of the three bridges of each clause is connected to a respective arm segment of the variable gadget. Because of Lemma 3, for each variable either its unnegated or its negated arm segment can be connected to their respective bridges. Therefore, there is an assignment of $\{1, 0, \perp\}$ for φ , given by the movement of each variable arm segment. We set a variable in φ to 1, if the respective arm segment is connected to the unnegated side, 0 if it is connected to its negated side, and \perp otherwise. For each \perp we arbitrarily choose 0 or 1. Due to the construction, this assignment satisfies φ .

Claims 1 and 2 complete the proof of Theorem 2. □

As a consequence of our construction in the proof of Theorem 2, even approximating the makespan is NP-hard.

Corollary 2 *It is NP-hard to compute for a pair of configurations C_s and C_t , each with n vertices, a stable schedule that transforms C_s into C_t within a constant of $(\frac{3}{2} - \epsilon)$ (for any $\epsilon > 0$) of the minimum makespan.*

To see this, consider an instance that derived from a Boolean formula as constructed in our NP-hardness reduction, and assume that the Boolean formula is satisfiable. Hence, the optimal makespan is 2. Consider a potential approximation algorithm that has a guaranteed approximation factor better than $\frac{3}{2}$. As the makespan is a natural number, and the minimum makespan of a false instance is 3, we would be able to decide with that algorithm whether the Boolean formula is satisfiable, having as a consequence that P=NP.

4 Bounded stretch for arbitrary makespan

Now we describe our algorithm for computing stable schedules with constant stretch, for configurations of constant scale. Again, a scale factor of c corresponds to replacing each pixel of a polyomino shape by a quadratic $c \times c$ array of pixels; this will be defined in the

preliminaries for the algorithm. In the remainder of this section, we describe the different phases of our approach that together show the following.

Theorem 3 *There is a constant c^* such that for any pair of overlapping start and target configurations with a scale of at least c^* , there is a stable schedule of constant stretch.*

For clearer presentation, we do not focus on the specific value of the constant c^* , but only argue its existence.

4.1 Algorithm overview and preliminaries

We provide a high-level overview of our algorithm, as depicted in Fig. 12.

Our algorithm works in five phase, where the first two phases are preparation phases; note that these two phases are not visualized in Fig. 12. In Phase 1, we ensure that the pair (C_s, C_t) overlaps in at least one position. For this, we move C_s towards C_t along a bottleneck matching, such that the respective positions that realize the bottleneck distance, coincide. (This overlap is crucial for successfully constructing the auxiliary structure in Phase 3 of the approach, as we explain later.) Afterwards, in Phase 2, we use another bottleneck matching for mapping the start configuration C_s to the target configuration C_t , minimizing the maximum distance d between a start and a target location. Furthermore, we establish




Fig. 12 Overview of the computed schedule: (Top) Constructing an auxiliary structure, called *scaffold*, in Phase 3 of our algorithm, (Right) the refilling phase (Phase 4), and (Bottom) deconstructing the scaffold in the final Phase 5

the scale in both configurations, set c to be the minimum of both scale values, and compute a suitable tiling whose tile size is $c \cdot d$, and that contain both C_s and C_t .

In Phase 3, we build a scaffolding structure around C_s and C_t , based on the boundaries of cd -tiles of the specific tiling. This structure enables simple connected reconfiguration procedures in the interior of tiles, as described in Phase 4 of our approach, yielding to a reconfiguration approach that consists of a connected configuration in each transformation step. Restricting robot motion to their current and adjacent tiles (due to the tile size) also ensures constant stretch. Note that, as the size of the tiles is related to d , the scaffolding structure is connected.

In Phase 4, we perform the actual reconfiguration of the arrangement. This consists of refilling the tiles of the scaffold structure, achieving the proper number of robots within each tile, based on elementary flow computations. As a subroutine, we transform the robots inside each tile into a canonical “triangle” configuration.

In the final Phase 5, we disassemble the scaffolding structure and move the involved robots to their proper destinations.

We proceed by providing the basic definitions, followed by summarizing the technical key components of our approach.

4.1.1 Preliminaries for the algorithm

A configuration C is c -scaled, if it is the union of $c \times c$ pairwise disjoint squares of vertices. The *scale* of a configuration C is the maximal c such that C is c -scaled; this notion reflects the idea of objects being composed of pixels at a certain resolution. (Note that this is a generalization of the uniform pixel scaling studied in previous literature, which considers a c -grid-based partition instead of an arbitrary union, so it supersedes that definition and leads to a more general set of results; for simplicity, we refrain from using “generalized scale” and stick to “scale” in the rest of the paper.) As we will establish, sufficient scale ensures sufficient locally available building material for connected rearrangement. Let $c, d \in \mathbb{N}$ be the scale and the diameter of the pair (C_s, C_t) , respectively. For $x, y \in \mathbb{N}$, a cd -tile T , or tile T for short, with *anchor vertex* $(x \cdot cd, y \cdot cd) \in V(G)$ is a set of $(cd)^2$ vertices from the grid G with x -coordinates from the range between $x \cdot cd$ and $x \cdot cd + cd - 1$ and y -coordinates from the range between $y \cdot cd$ and $y \cdot cd + cd - 1$. Note that, by definition, two tiles are pairwise disjoint, and all anchor points have the same coordinates mod cd . The *boundary* of T is the set of vertices from T with an x -coordinate equal to $x \cdot cd$ or equal to $x \cdot cd + cd - 1$, or with a y -coordinate equal to $y \cdot cd$ or equal to $y \cdot cd + cd - 1$. The *interior* of T is T without its boundary. The *right*, *top*, *left*, and *bottom sides* of T are the sets of vertices from the boundary of T with maximum x -coordinates, maximum y -coordinates, minimum x -coordinates, and minimum y -coordinates, respectively. The left and right sides of a tile are *vertical sides* and the top and bottom sides are *horizontal sides*. Two tiles T_1, T_2 are *horizontal (vertical) neighbors* if they have two horizontal (vertical) sides $s_1 \subset T_1$ and $s_2 \subset T_2$, such that each vertex from s_1 is adjacent in G to a vertex from s_2 . Two tiles T_1 and T_2 are *diagonal neighbors* if there is another tile T , such that T and T_1 are horizontal neighbors and T and T_2 are vertical neighbors. The *neighborhood* of a tile T is the set of all neighbors of T .

A *start tile* is a tile containing a vertex from the start configuration. A *target tile* is a tile containing a vertex from the target configuration. The cd -tiling \mathcal{T} of (C_s, C_t) is the union of all start tiles including their neighborhoods and all target tiles. The *scaffold* $\Sigma(\mathcal{T})$ is the union of all boundaries of tiles from \mathcal{T} . A cd -tiled configuration, or *tiled configuration* for

short, is a configuration that is a subset of the union of all tiles from \mathcal{T} and a superset of $\Sigma(\mathcal{T})$. The *interior* of a tiled configuration C is the set of all vertices from C not lying on $\Sigma(\mathcal{T})$. The *filling level* of a tile $T \in \mathcal{T}$ is the number of robots in the interior of T . The *filling level* of a tiled configuration C is the mapping of each tile onto its filling level in C .

4.1.2 Technical key components

On a technical level, the five phases can be summarized as follows; cf. Fig. 12.

- (1) *Guaranteeing Overlap* Move the given configurations towards each other along a bottleneck matching to ensure that the (new) pair (C_s, C_t) overlaps in at least one position.
- (2) *Preprocessing* Apply the following three steps: (2.1) Set c to be the minimum of c^* (that is a roughly estimated constant lower bound on the scale for which our approach works) and the minimum scale values of C_s and C_t . (2.2) Compute the diameter d of (C_s, C_t) . (2.3) Compute the tiling \mathcal{T} of (C_s, C_t) .

The algorithmic core of our algorithm consists of the following three phases.

- (3) *Scaffold Construction* Reconfigure the start configuration C_s to a tiled configuration C'_s such that the interior of C'_s is a subset of the start configuration C_s .
- (4) *Refilling Tiles* Reconfigure C'_s to a tiled configuration C'_t , such that the interior of C'_t is a subset of the target configuration C_t .
- (5) *Scaffold Deconstruction* Reconfigure C'_t to C_t .

Note that the scaffold deconstruction is inverse to the scaffold construction. In the following, we give the technical description of our algorithm, and the corresponding correctness analysis. In particular, we first assume that the start and target configurations overlap in at least one position, resulting in an algorithm guaranteeing constant stretch, and adapt this to the case in that an overlap initially does not exist, afterwards.

4.2 Scaffold construction

Lemma 5 *For any configuration C_s of scale c there is a stable schedule of makespan $\mathcal{O}(d)$, transforming C_s into a tiled configuration C'_s , with the interior of C'_s being a subset of C_s .*

Outline of the Construction For the construction, we consider $5 \cdot 5$ different classes, based on x - and y -coordinates modulo $5cd$; see Fig. 13.

We process a single class as follows. For each tile T we consider its *indirect* neighborhood $N[T]$ consisting of all neighbors of T and all neighbors of neighbors of T , i.e., a 5×5 arrangement of tiles centered at T . For constructing the boundary of T , we make use of robots from the interior of a tile in $N[T]$.

Constructing the Boundary of T Works in Two Phases (3.1) Constructing the boundaries of all start tiles. (3.2) Constructing the boundaries of all neighbors of start tiles.

Note that it suffices to construct all boundaries of the start tiles and their neighboring tiles, because each target tile shares a side with a start tile or a side with a tile adjacent to a start tile. Furthermore, the scale condition is only necessary for the construction of the scaffold, i.e., we need to ensure that enough robots are available to build the scaffolding structure. Each additional step of the algorithm works independently of this condition.




Fig. 13 Constructing the scaffold. Tiles with currently constructed boundaries are marked in purple. The zoom into the start configuration C_s shows the indirect neighborhood $N[T]$ of a tile T (middle) for which its boundary is currently constructed. A further zoom (right) into T with an associated robot motion. In each step a robot from the interior of a tile $T' \in N[T]$ is swapped with a free position on the boundary of T based on a path P on a BFS-tree (Color figure online)

For Phase 3.1, we subdivide all start tiles into $5 \cdot 5$ classes of tiles, distinguishing between their anchors' coordinates modulo $5cd$; see the purple squares in Fig. 13. These classes are processed iteratively. In particular, we construct the boundaries of all tiles belonging to a specific class in parallel as follows. Without loss of generality, we assume that the union of C_s with all boundaries of start tiles is connected. Otherwise, C_s lies completely inside a single tile T , and we move all robots simultaneously down until the first robot lies on the boundary of T .

For each tile T of the current class, we consider a 5×5 -neighborhood $N[T]$ of T which consists of T , its neighbors, and all neighbors of neighbors of T , see Fig. 13. The scaffold $\Sigma(N[T])$ is the union of all boundaries of tiles from $N[T]$. The interior of $N[T]$ is $N[T]$ without its scaffold, i.e., $N[T] \setminus \Sigma(N[T])$. The *priority* of a robot r from the interior of $N[T]$ is the shortest-path distance inside the current configuration between r and any vertex from $\Sigma(N[T])$.

Let C be the current configuration, v be a vertex from the boundary of T that does not lie inside C but is adjacent to a robot from C , and r be a robot from the interior of $N[T]$ with highest priority. If C is connected and the robot r lies in the interior of $N[T]$, then there is a path P inside C connecting v and r , see the right of Fig. 13. We push all robots on P along P into the direction of v resulting in the connected configuration $C \cup \{v\} \setminus \{r\}$. This one-step motion is repeated until the interior of $N[T]$ is empty or the boundary of T is a subset of the current configuration.

Correctness of the Approach We need to argue that in Phase 3.1 the construction of the boundary of T stops because the boundary is completed and not because the interior of $N[T]$ is empty.

Consider the vertices from $N(T) := N[T] \setminus T$ to be organized in squares nested inside each other as follows. All vertices $v \in N(T)$ adjacent to the boundary of T belong to layer 1. Let $N_1(T)$ be the vertices from $N(T)$ without the vertices from layer 1. All vertices from $N_i(T)$ that are adjacent to vertices from layer i belong to Layer $i + 1$. Consider a

path $P \subset C_s$ connecting layer c with layer $cd - c$. As C_s is c -scaled, for each $v \in P$, there is a square S of side length c such that $v \in S \subset C_s$. Hence, there is a tile $T' \in N[T]$, such that all these squares cover at least $\frac{(cd-2(c-2))(c-2)}{2} \geq \frac{(c(d-2))(c-2)}{2}$ vertices from C_s , which is lower-bounded by $\frac{c^2d}{4}$ for $c \geq 4$. Constructing the boundary of T takes at most $4cd - 4$ robots to be taken from the interior of $N[T]$. The tile T' is used as part of the interior of at most $5 \cdot 5 = 25$ tiles' neighborhoods $N[T]$. Thus, at most $(4cd - 4)25 < 100cd$ robots are required from the interior of the tile $T' \in N[T]$. Hence, for $\frac{c^2d}{4} \geq 100cd$, i.e., $c \geq 400$ the construction of the boundary of each start tile T is completed in Phase 3.1. Because there is a constant number of tile classes (where each class can be processed in parallel), and the moving distances are bounded by a constant times d , the construction finishes in $\mathcal{O}(d)$. Phase 3.2 is realized analogously.

4.3 Refilling tiles

It remains to modify configurations within and between tiles. To this end, we first establish how to efficiently perform reconfigurations between any two tiled configurations *with the same numbers of robots* in the interior of respective tiles; see Sect. 4.3.1. As a second step, we describe how to relocate robots between tiles such that efficient reconfigurations between any two tiled configurations *with different numbers of robots* in the interior of respective tiles are achieved; see Sect. 4.3.2.

4.3.1 Reconfiguration maintaining the number of robots in tiles

Lemma 6 *Let C'_s, C'_t be two tiled configurations such that for all tiles T , C'_s and C'_t have the same filling levels, i.e., for any tile, the corresponding start and target configurations consist of the same respective numbers of robots. Then there is a stable schedule transforming C'_s into C'_t within a makespan of $\mathcal{O}(d)$.*

In the following we describe reconfigurations that leave all robot movements within the interior of their respective tiles T ; thus, all tiles can be reconfigured in parallel. Therefore, we only have to describe the approach for a start configuration C_s and a target configuration C_t within the interior of a single tile T of a tiled configuration C'_s . The idea is as follows: As the moves are reversible, it is sufficient to prove reachability of a canonical intermediate configuration in $\mathcal{O}(d)$ transformations, which allows us to simplify our argument. In the following, we define the properties of our canonical structure C_Δ and provide the necessary means to reach it from an arbitrary initial state.

Outline of the Reconfiguration First, we compute two stable schedules $C_s \Rightarrow_\chi C_s^m$ and $C_t \Rightarrow_\chi C_t^m$, where C_s^m and C_t^m are *monotone configurations*, where we call a configuration to be monotone, if for every robot r in the interior of T all positions to the left and to the bottom are occupied. These reconfigurations are achieved by a sequence of down and left movements, maintaining connectivity after each move (see Fig. 14, Phase 4.1). Proceeding from these monotone configurations, the robots are arranged into a triangular configuration C_Δ that occupies the lower left positions (defined by a diagonal line with a slope of -1) of the interior of T . This is achieved by swapping pairs of occupied and empty positions within a carefully defined area in several one-step moves along L-shaped paths (see

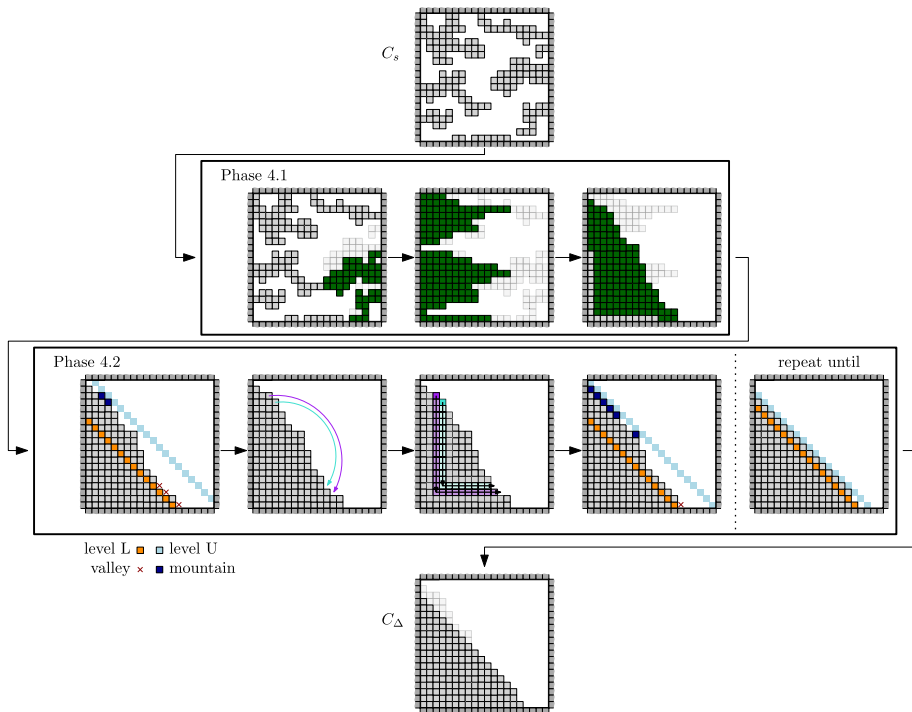


Fig. 14 Reconfiguring the configuration C_s (top) into the canonical triangle configuration C_Δ (bottom). Phase 4.1 achieves a monotonic arrangement; light gray squares indicate previous positions of moved robots (shown in green). Phase 4.2 transforms the monotonic configuration into C_Δ

Fig. 14, Phase 4.2). The property of C_Δ is that it is the same for all initial configurations with equally many robots. Thus, to get the stable schedule $C_s \Rightarrow_\chi C_\Delta \Rightarrow_\chi C_t$, we can simply revert $C_t \Rightarrow_\chi C_\Delta$ and combine the result with $C_s \Rightarrow_\chi C_\Delta$.

Technically, the approach consists of the following four phases, see Fig. 14.

- (4.1) *Monotone Start Configuration* Reconfigure C_s into C_s^m .
- (4.2) *Canonical Triangle* Reconfigure C_s^m into C_Δ .
- (4.3) *Monotone Target Configuration* Reconfigure C_Δ into C_t^m .
- (4.4) *Target Configuration* Reconfigure C_t^m into C_t .

Phase 4.4 corresponds to a reversal of Phase 4.1, and Phase 4.3 to one of Phase 4.2, so we only have to describe the first two phases. We analyze them individually, leading to a proof of Lemma 6. Note that we exclude the corners of a tile, so the robots on the tile's side now form four *non-adjacent* sides. Furthermore, only robots in a tile's interior move.

Constructing the Monotone Start Configuration (Phase 4.1) In the first step, we only consider robots for which the right side of their tile T is the only one to which they are

connected through the interior of T . We iteratively move these robots down until further movement is blocked, i.e., any further down move is not collision-free. In the second and third steps, we move all robots left, followed by moving all robots down, each time until further movement is blocked.

Claim 3 *Phase 4.1 turns a configuration in a tiles' interior into a monotone configuration.*

Proof It is straightforward to see that for, any configuration, Phase 4.1 results in a monotone configuration, so we only argue that it yields a stable schedule.

Without loss of generality, we only consider the down movement of a robot. If the movement of a robot r resulted in a collision, there would be another robot on its target position that cannot move either. Because the robots only move down, it iteratively follows that this movement is blocked because of a robot that is either adjacent to the bottom side or for which the right side is not the only one to which it is connected.

We now argue that a down move of a robot r preserves connectivity to other robots. Assume that there is a robot r' to which r was connected before the move (say, in C_i), but not after it (in C_{i+1}). Consider a path from r to r' in C_i , the last robot r_j on that path which is not connected to r in C_{i+1} , and its adjacent robot r_{j+1} on that path that is still connected to r' ; note that this robot could be r' itself. Because r_j and r_{j+1} are not connected in C_{i+1} , they were horizontally adjacent in C_i . Without loss of generality, this disconnection happened because r_j moved down, whereas r_{j+1} did not. But then all positions below r_{j+1} had to be occupied. Thus, r_j and r_{j+1} are still connected, a contradiction. \square

By simply changing the direction of the robot's movements from down to left, all arguments hold true for the entire phase.

For Phase 4.2, we use the following terminology. The *level* of a position in the tile's interior is the sum of its coordinates. A level is *filled*, if all of its positions are occupied by a robot, and *empty*, if none is occupied by a robot. The highest filled level is denoted by L , the lowest empty level by U . Let \mathcal{M} be the set of all positions on level $U - 1$ occupied by a robot, and \mathcal{V} be the set of all positions on level $L + 1$ that are not occupied by a robot; we call the positions of \mathcal{M} and \mathcal{V} *mountains* and *valleys*, respectively.

Constructing the Canonical Triangle (Phase 4.2) Choose two equally sized subsets $\mathcal{M}' \subseteq \mathcal{M}$ and $\mathcal{V}' \subseteq \mathcal{V}$ and push each robot from \mathcal{M}' to a different position in \mathcal{V}' along an L-shaped path; this can be done simultaneously in one parallel move for all paths. To determine the paths, simply match mountains and valleys, iteratively, in a way that no pair of paths cross each other. This results in reducing U , and raising L , i.e., the two levels move towards each other. We distinguish two cases.

- $U - L > 2$: If $|\mathcal{M}| \geq |\mathcal{V}|$, choose an arbitrary subset $\mathcal{M}' \subset \mathcal{M}$ with $|\mathcal{M}'| = |\mathcal{V}|$. Otherwise, choose an arbitrary subset $\mathcal{V}' \subset \mathcal{V}$ with $|\mathcal{V}'| = |\mathcal{M}|$.
- $U - L = 2$: Note that mountains and valleys are on the same level, and $|\mathcal{V}| \geq |\mathcal{M}|$ hold. Choose $\mathcal{V}' \subset \mathcal{V}$ to be the subset of size $|\mathcal{M}|$ with *smallest* x -coordinates and set $\mathcal{M}' = \mathcal{M}$.

Claim 4 *Phase 4.2 yields a stable schedule for reconfiguring all monotone configurations with the same number of robots into the same triangular configuration.*

Proof We show the following properties.

1. The pushing paths do not cross.
2. After each iteration, $U - L$ decreases by at least 1.
3. The approach terminates in a configuration that only depends on the number of robots.

For the non-crossing property of the paths, we show that the mountains and valleys have pairwise disjoint x - and y -coordinates. If $U - L > 2$, assume for the sake of contradiction that a mountain and a valley have equal x -coordinate. These positions have to be adjacent, which would result in $U - L = 1$, a contradiction. If $U - L = 2$ holds, mountains and valleys are already on the same level, so they have distinct x -coordinates. Similar arguments hold for the y -coordinates. It directly follows from the definition of levels that for all positions $(x_1, y_1), (x_2, y_2) \in \mathcal{M} \cup \mathcal{V}$, $x_1 < x_2$ implies $y_1 > y_2$. Together with the construction of the paths by matching adjacent mountains and valleys, it follows that these paths do not cross.

By construction, $U > L$ always holds. If $\mathcal{M}' = \mathcal{M}$ holds, level $U - 1$ gets empty; therefore, U is decreased. In any case, U cannot be increased, because all positions that get occupied are at level $L + 1$. A similar argument holds if $\mathcal{V}' = \mathcal{V}$ holds, thus, $U - L$ decreases by at least 1 in each iteration.

If $U - L = 2$ holds, there is at most one level that contains occupied and empty positions. All robots on this level will be positioned leftmost; all levels with greater values are empty, and all levels with lower values are full.

Because the paths are non-crossing, the computed schedule is stable. \square

Proof of Lemma 6 Compute stable schedules consisting of Phases 4.1 and 4.2 for the start and the target configuration, C_s and C_t , to reconfigure them into C_Δ . Reverting $C_t \Rightarrow_\chi C_\Delta$ and combining it with $C_s \Rightarrow_\chi C_\Delta$ results in the entire stable schedule transforming C_s into C_t . Because Phase 4.1 as well as Phase 4.2 can be realized in a makespan of $\mathcal{O}(d)$ (because a robot's maximum distance to a side, and the maximum level difference are upper-bounded by $\mathcal{O}(d)$), the makespan of $C_s \Rightarrow_\chi C_t$ is in $\mathcal{O}(d)$. \square

4.3.2 Refilling tiled configurations

Now we describe the final step for reconfiguring a tiled start configuration C'_s into a tiled target configuration C'_t ; because no robots are destroyed or created, this hinges on shifting robots between adjacent tiles, such that the required filling levels are achieved.

Lemma 7 *We can efficiently compute a stable schedule transforming C'_s into C'_t within a makespan of $\mathcal{O}(d)$.*

Outline of the Refilling Phase To compute the schedule of Lemma 7, we transfer robots between tiles, so that each tile T contains the desired number in C'_t . We model this robot transfer by a *supply and demand flow*, see Figs. 15 and 16, followed by partitioning the flow into $\mathcal{O}(1)$ subflows, such that each subflow can be realized within a makespan of $\mathcal{O}(d)$. For realizing a single subflow, we use the approach of Sect. 4.3.1 as a preprocessing step, i.e., to rearrange robots participating in a specific subflow and place them at suitable positions.

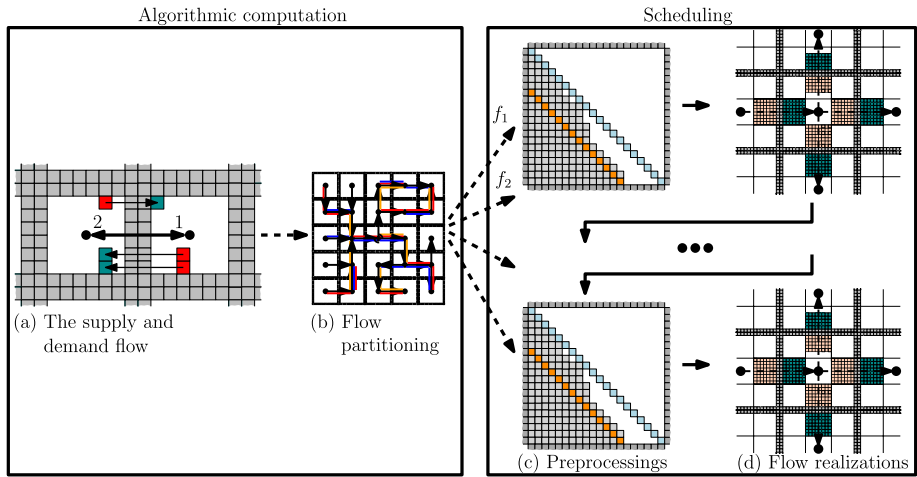



Fig. 15 An overview of the schedule refilling tiles: transforming C'_s into C'_t by realizing a partition of a supply and demand flow that is computed in advance

Fig. 16 A pair of configurations C'_s (red) and C'_t (cyan) and the resulting flow modeling the edges of the bottleneck matching. The left vertex of G has demand 1, the right vertex has supply 1



Modeling Transfer of Robots via a Supply and Demand Flow We model the transfer of robots between tiles as a flow $F : E(G) \rightarrow \mathbb{N}$, using the directed graph $G = (V, E)$ which is dual to the tiling \mathcal{T} . Let B be the bottleneck matching between vertices from the original (non-tiled) C_s and vertices from the final (non-tiled) C_t . In G we have an edge $(u, v) \in E$, if there is at least one matching edge $(r_u, r_v) \in B$, such that r_u lies in the interior of the tile u in configuration C'_s , and r_v lies in the interior of the tile v in configuration C'_t . The flow value $F((u, v))$ of (u, v) is equal to the number of such edges $(r_u, r_v) \in B$. A vertex $v \in V(G)$ has a *demand* of $a > 0$ if the sum of the flow values of outgoing edges from v plus a is equal to the sum of the flows of incoming edges to v . Analogously, v has a *supply* of $a > 0$ if the sum of flow values incoming to v plus a equals the sum of flow values outgoing from v .

Flow Partition and Algorithmic Computation We now define a *flow partition* of F . For $k \in \mathbb{N}$, a k -*subflow* of F is a supply and demand flow $f : E(G) \rightarrow \mathbb{N}$ on G with $f(e) \leq \min\{k, F(e)\}$. A k -*partition* of the flow F is a set $\{f_1, \dots, f_\ell\}$ of k -subflows of F , such that $\sum_{i=1, \dots, \ell} f_i(e) = F(e)$ holds for all edges $e \in E(G)$.

We describe our approach for computing a ϕ -partition of F , with $\phi := \lfloor \frac{(cd-2)^2}{9} \rfloor$; the value ϕ arises from partitioning the interior (made up of $(cd-2)^2$ pixels) of each tile into 9 almost equally sized *subtiles* that are used for realizing a single set of paths as described below, see Fig. 18.

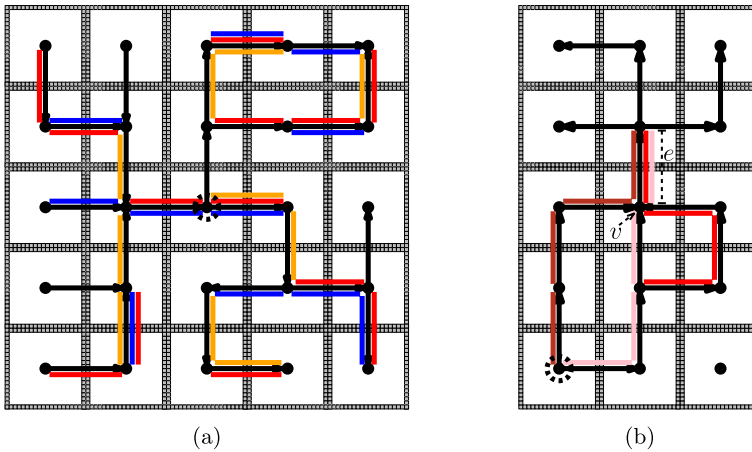


Fig. 17 **a** We model the movements of robots between tiles as paths forming a tree. By greedily assigning these paths to sets (here highlighted by different colors), such that inside each set each edge is contained in no more than 3 paths, we obtain that at most $\Theta(d^2)$ sets are needed. **b** An example of a set containing three paths (dark red, pink, and red; assigned in that order to S_j) having a common edge e caused by a vertex v of e with an incoming degree of 3 (Color figure online)

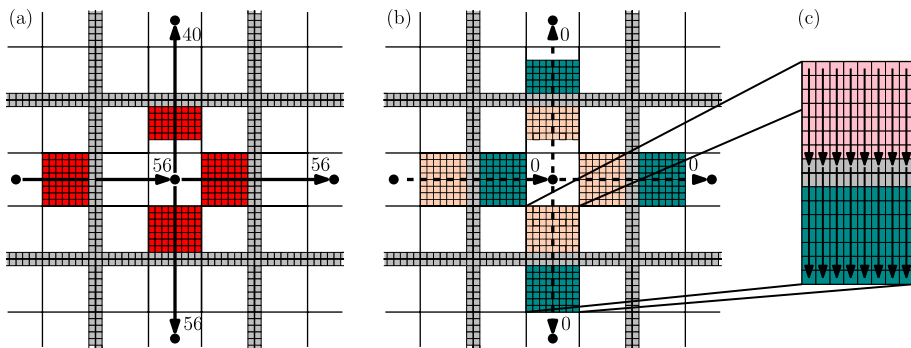


Fig. 18 **a** A portion of a set of paths containing a vertex v : 56 paths are passing v from left to right, while $40 + 56 = 96$ paths are starting in v . **b** The configuration after realizing the set of paths shown in the previous figure. **c** How positions of robots have changed after realization

We compute a 1-partition of G , with each 1-subflow being either a cycle or a path that connects a supply vertex with a demand vertex. Because the robots are unlabeled, we can simplify G by eliminating all cyclic 1-subflows, as they are not necessary to realize this specific transfer of robots; note that this also applies to bidirectional edges. Furthermore, we replace diagonal edges $(v, w) \in E(G)$ by a pair of adjacent edges $(v, u), (u, w) \in E(G)$. After all cyclic subflows are removed, G is a planar, directed forest consisting of 1-subflows that are paths, see Fig. 17a. We process each tree $A \subset G$ that is made up of paths P_1, P_2, \dots separately as follows: We choose an arbitrary vertex of A as its *root* and consider the *link distance* of P_i as the minimal length of the path between a vertex from P_i and the root of A . Let $P_1, P_2, \dots \subseteq G$ be sorted by increasing link distances, which is important for our next argument, regarding that we can partition these paths into constant many subflows each of which is realizable in time linear in d . We greedily assign each path $P_i = P_1, P_2, \dots$

to a set S_j , such that the first edge e_1 of P_i is not part of another path inside S_j . If no such a set S_j exists, we create a new set $S \leftarrow \{P_i\}$. For each tree we use the same sets S_1, S_2, \dots of collected paths, because different trees of G are disjoint. Note, that the construction of the sets S_j allow that an edge is part of at most three paths inside S_j . This is due to the fact that the income degree of the head vertex of a directed edge is at most three in the setting of a grid graph, resulting in at most three outgoing edges.

Finally, we greedily partition $\{S_1, S_2, \dots\}$ into subsets G_1, G_2, \dots called *groups*, made up of $\frac{(cd-2)^2}{9 \cdot 3}$ sets. For each group G_i , we define a subflow f_i by setting $f_i(e)$ as the number of paths from G_i containing the edge e . As for each set S_i and each edge e , there are at most three paths inside S_i containing e , each resulting subflow f_i is a $\left(\frac{(cd-2)^2}{9 \cdot 3} \cdot 3\right) = \phi$ -subflow. Finally, we have to upper-bound the number of resulting subflows, i.e., the number of groups.

Lemma 8 *The constructed ϕ -partition $\{f_1, f_2, \dots\}$ consists of at most 28 subflows.*

Proof The number of robots in the interior of each tile is upper-bounded by $(cd-2)^2$ which is the number of pixels in a tiles' interior. Substituting at most two diagonal edges increases the number of robots leaving or entering a specific tile by at most $2(cd-2)^2$. Hence, the flow value of each edge is upper-bounded by $3(cd-2)^2$ (\star) after substituting diagonal edges.

Suppose our approach constructs 28 groups implying that our approach gets into configuration with $27 \cdot \frac{(cd-2)^2}{9 \cdot 3} =: \lambda$ constructed sets S_1, S_2, \dots and the current path P_i has to be assigned to a new set $S_{\lambda+1}$. This implies that each set S_1, \dots, S_λ contains a path S_i containing the first edge e_1 of P_i . Hence, there are at least $27 \cdot \frac{(cd-2)^2}{9} + 1 > 3(cd-2)^2$ paths containing e_1 , contradicting (\star). \square

Realizing a ϕ -Partition Now we describe how to reconfigure a tiled configuration, such that a ϕ -subflow is removed from G . A ϕ -subflow f_i is *realized* by transforming the current configuration into another configuration, such that for each edge e with $f_i((T, T')) > 0$, the number of robots in the interior of tile T is decreased by $f_i((T, T'))$ and the number of the robots in the interior of tile T' is increased by $f_i((T, T'))$.

Next we realize a specific ϕ -subflow within a makespan of $\mathcal{O}(d)$. In particular, we partition the interior of each tile T into 9 *subtiles* with equal side lengths (up to rounding), see Fig. 18. For each subflow f_i , we place $f_i((T, T'))$ robots inside the middle subtile of T that shares an edge with the boundary of T adjacent to T' . In particular, robots placed in the same subtile are arranged in layers of width $\lfloor \frac{cd-2}{9} \rfloor$ as close as possible to the boundary of the tile, see Fig. 18a. The resulting arrangement of robots inside the subtile of T is a *cluster* and T' the *target tile* of the cluster. By a single application of the approach from Sect. 4.3.1, all clusters of all tiles are arranged simultaneously within a makespan of $\mathcal{O}(d)$. Finally, simultaneously pushing all clusters of all tiles into the direction of their target tiles realizes S_i , see Fig. 18b. Note that not all robots are pushed into the target tile T' but some replace robots on the boundaries between T and T' , see Fig. 18c.

Repeating this approach for each subflow f_i leads to a stable schedule that realizes the entire flow F within a makespan linear in d , i.e., transforms C'_s into C'_t within $\mathcal{O}(d)$ transformation steps, see Fig. 15.

Proof of Lemma 7 The schedule described above (and illustrated in Fig. 18) is correct because of the following. For each tile $T \in V(G)$, the number of the paths starting in T equals the supply of T , while the number of paths ending in T equals the demand of T . As T has either a supply or a demand, but not both, so the tile T is never too full to receive additional robots, and never too empty to deliver robots as required by paths containing T . \square

This concludes the proof of Theorem 3. Finally, we adapt the result to the general case, for which overlap of the start and target configuration is not guaranteed.

Corollary 3 *There is a constant c^* such that for any pair of start and target configurations with a scale of at least c^* , there is a stable schedule of constant stretch.*

Proof In case of a pair (C_s, C_t) , consisting of a start and a target configuration, that does not overlap, our algorithm computes in a first step a minimum bottleneck matching mapping C_s to C_t resulting in a bottleneck distance \bar{d} , and translates C_s into a configuration \bar{C}_s overlapping the target configuration within a makespan of \bar{d} . This results in a bottleneck distance d between \bar{C}_s and C_t which is at most $2\bar{d}$. As Theorem 3 guarantees a makespan linear in d , we obtain a makespan linear in $\bar{d} + d$, i.e., linear in \bar{d} for the overall algorithm. \square

As the diameter of the pair (C_s, C_t) is a lower bound for the makespan of any schedule transforming C_s into C_t , we obtain the following.

Corollary 4 *There is a constant-factor approximation for computing stable schedules with minimal makespan between pairs of start and target configurations with a scale of at least c^* , for some constant c^* .*

All steps in the overall algorithm can be computed in an efficient manner. Critical is the computation of matchings (Phase 2): the underlying geometry can be exploited for computing bottleneck matchings for a set of points in the plane in $\mathcal{O}(n^{1.5} \log n)$ [18]. The required BFS tree (Phase 3) and flow decomposition (Phase 4) can both be computed in straightforward manner; note that Phase 4 does not require a smallest decomposition, and the edge capacities are strongly polynomial in the number n of robots.

5 Conclusion

We have shown that connected coordinated motion planning is challenging even in relatively simple cases, such as unlabeled robots that have to travel a distance of at most 2 units, by establishing NP-completeness. On the other hand, we have shown that (assuming sufficient scale of the swarm), it is possible to compute efficient reconfiguration schedules with constant stretch.

It is straightforward to extend our approach to other scenarios, e.g., to three-dimensional configurations. Other questions appear less clear. Is it possible to achieve constant stretch for arrangements with very small scale factor? We believe that this may hinge on the ability to perform synchronized shifts on long-distance “chains” of robots without delay, which is not a valid assumption for many real-world scenarios. (A well-known example is a line of cars when a traffic light turns green.) As a consequence, the answer may depend on crucial

assumptions on motion control; we avoid this issue in our approach. Can we provide alternative approaches with either weaker scale assumptions or better stretch factors? Can we extend our methods to the labeled case? All these questions are left for future work.

Funding Open Access funding enabled and organized by Projekt DEAL.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

1. Adler, A., de Berg, M., Halperin, D., & Solovey, K. (2015). Efficient multi-robot motion planning for unlabeled discs in simple polygons. *IEEE Transactions on Automation Science and Engineering*, 12(4), 1309–1317. <https://doi.org/10.1109/TASE.2015.2470096>
2. Akitaya, H. A., Arkin, E. M., Damian, M., Demaine, E. D., Dujmovic, V., Flatland, R. Y., Korman, M., Palop, B., Parada, I., van Renssen, A., & Sacristán, V. (2021). Universal reconfiguration of facet-connected modular robots by pivots: The O(1) musketeers. *Algorithmica*, 83(5), 1316–1351. <https://doi.org/10.1007/s00453-020-00784-6>
3. Becker, A. T., Fekete, S. P., Keldenich, P., Konitzny, M., Lin, L., & Scheffer, C. (2018). Coordinated motion planning: The video. In *Symposium on computational geometry (SoCG)*, pp. 74:1–74:6. Video at <https://www.ibr.cs.tu-bs.de/users/fekete/Videos/oordinatedMotionPlanning.mp4>. <https://doi.org/10.4230/LIPIcs.SoCG.2018.74>
4. Bourgeois, J., Fekete, S.P., Kosfeld, R., Kramer, P., Piranda, B., Rieck, C., & Scheffer C. (2022). Space ants: Episode II—Coordinating connected catoms. In *Symposium on computational geometry (SoCG)*, pp. 65:1–65:6. Video at <https://youtu.be/m45jWeCUt9Y>. <https://doi.org/10.4230/LIPIcs.SoCG.2022.65>
5. Charrier, T., Queffelec, A., Sankur, O., & Schwarzenruber, F. (2019). Reachability and coverage planning for connected agents. In *International joint conference on artificial intelligence (IJCAI)*, pp. 144–150. <https://doi.org/10.24963/ijcai.2019/21>
6. Charrier, T., Queffelec, A., Sankur, O., & Schwarzenruber F. (2019). Reachability and coverage planning for connected agents. In *International conference on autonomous agents and multiagent systems (AAMAS)*, pp. 1874–1876. <https://dl.acm.org/doi/10.5555/3306127.3331948>
7. Charrier, T., Queffelec, A., Sankur, O., & Schwarzenruber, F. (2020). Complexity of planning for connected agents. *Autonomous Agents Multi Agent Systems*, 34(2), 44. <https://doi.org/10.1007/s10458-020-09468-5>
8. Chung, S.-J., Paranjape, A. A., Dames, P., Shen, S., & Kumar, V. (2018). A survey on aerial swarm robotics. *IEEE Transactions on Robotics*, 34(4), 837–855. <https://doi.org/10.1109/TRO.2018.2857475>
9. Crombez, L., da Fonseca, G. D., Gerard, Y., Gonzalez-Lorenzo, A., Lafourcade, P., & Libralesso, L. (2021). Shadoks approach to low-makespan coordinated motion planning. In *Symposium on computational geometry (SoCG)*, pp. 63:1–63:9. <https://doi.org/10.4230/LIPIcs.SoCG.2021.63>
10. de Berg, M., & Khosravi, A. (2012). Optimal binary space partitions for segments in the plane. *International Journal on Computational Geometry and Applications*, 22(3), 187–206. <https://doi.org/10.1142/S0218195912500045>
11. Delahaye, D., Puechmorel, S., Tsiotras, P., & Féron, E. (2014). Mathematical models for aircraft trajectory design: A survey. In *Air traffic management and systems*, pp. 205–247. https://doi.org/10.1007/978-4-431-54475-3_12
12. Demaine, E. D., Demaine, M. L., Fekete, S. P., Ishaque, M., Rafalin, E., Schweller, R. T., & Souvaine, D. (2008). Staged self-assembly: Nanomanufacture of arbitrary shapes with O(1) glues. *Natural Computing*, 7(3), 347–370. <https://doi.org/10.1007/s11047-008-9073-0>

13. Demaine, E. D., Fekete, S. P., Keldenich, P., Meijer, H., & Scheffer, C. (2019). Coordinated motion planning: Reconfiguring a swarm of labeled robots with bounded stretch. *SIAM Journal on Computing*, 48(6), 1727–1762. <https://doi.org/10.1137/18M1194341>
14. Demaine, E. D., Fekete, S. P., Scheffer, C., & Schmidt, A. (2017). New geometric algorithms for fully connected staged self-assembly. *Theoretical Computer Science*, 671, 4–18. <https://doi.org/10.1016/j.tcs.2016.11.020>
15. Demaine, E. D., Patitz, M. J., Schweller, R. T., & Summers, S. M. (2011). Self-assembly of arbitrary shapes using RNase enzymes: Meeting the Kolmogorov bound with small scale factor. In *Symposium on theoretical aspects of computer science (STACS)*, pp. 201–212. <https://doi.org/10.4230/LIPIcs.STACS.2011.201>
16. Derakhshandeh, Z., Gmyr, R., Richa, A. W., Scheideler, C., & Strothmann, T. (2015). An algorithmic framework for shape formation problems in self-organizing particle systems. In *International conference on nanoscale computing and communication*, pp. 21:1–21:2. <https://doi.org/10.1145/2800795.2800829>
17. Derakhshandeh, Z., Gmyr, R., Richa, A. W., Scheideler, C., & Strothmann, T. (2016). Universal shape formation for programmable matter. In *Symposium on parallelism in algorithms and architectures (SPAA)*, pp. 289–299. <https://doi.org/10.1145/2935764.2935784>
18. Efrat, A., Itai, A., & Katz, M. J. (2001). Geometry helps in bottleneck matching and related problems. *Algorithmica*, 31, 1–28. <https://doi.org/10.1007/s00453-001-0016-8>
19. Fekete, S. P., Hendriks, B., Tessars, C., Wegener, A., Hellbrück, H., Fischer, S., & Ebers, S. (2011). Methods for improving the flow of traffic. *Organic Computing-A Paradigm Shift for Complex Systems*. https://doi.org/10.1007/978-3-0348-0130-0_29
20. Fekete, S. P., Keldenich, P., Kosfeld, R., Rieck, C., & Scheffer, C. (2021). Connected coordinated motion planning with bounded stretch. In *International symposium on algorithms and computation (ISAAC)*, 9:1–9:16. <https://doi.org/10.4230/LIPIcs.ISAAC.2021.9>
21. Fekete, S. P., Keldenich, P., Krupke, D., & Mitchell, J. S. B. (2022). Computing coordinated motion plans for robot swarms: The CG:SHOP challenge 2021. *ACM Journal of Experimental Algorithmics*, 27, 3.1:1–3.1:12. <https://doi.org/10.1145/3532773>
22. Goldstein, S. C., Mowry, T. C. (2004). Claytronics: A scalable basis for future robots. In *Robosphere 2004*, <http://www.cs.cmu.edu/~claytronics/papers/goldstein-robosphere04.pdf>
23. Hopcroft, J. E., & Karp, R. M. (1973). An $n^{3/2}$ algorithm for maximum matchings in bipartite graphs. *SIAM Journal on Computing*, 2(4), 225–231. <https://doi.org/10.1137/0202019>
24. Hopcroft, J. E., Schwartz, J. T., & Sharir, M. (1984). On the complexity of motion planning for multiple independent objects; PSPACE-hardness of the warehouseman's problem. *International Journal of Robotics Research*, 3(4), 76–88. <https://doi.org/10.1177/027836498400300405>
25. Hopcroft, J. E., & Wilfong, G. T. (1986). Reducing multiple object motion planning to graph searching. *SIAM Journal on Computing*, 15(3), 768–785. <https://doi.org/10.1137/0215055>
26. Kloder, S., & Hutchinson, S. (2006). Path planning for permutation-invariant multi-robot formations. *IEEE Transactions on Robotics and Automation*, 22(4), 650–665. <https://doi.org/10.1109/TRO.2006.878952>
27. Liu, P., Spalding-Jamieson, J., Zhang, B., & Zheng, D. W. (2021). Coordinated motion planning through randomized k-opt. In *Symposium on computational geometry (SoCG)*, pp. 64:1–64:8. <https://doi.org/10.4230/LIPIcs.SoCG.2021.64>
28. Luchsinger, A., Schweller, R. T., & Wylie, T. (2019). Self-assembly of shapes at constant scale using repulsive forces. *Natural Computing*, 18(1), 93–105. <https://doi.org/10.1007/s11047-018-9707-9>
29. Naz, A., Piranda, B., Bourgeois, J., & Goldstein, S. C. (2016). A distributed self-reconfiguration algorithm for cylindrical lattice-based modular robots. In *International Symposium on network computing and applications (NCA)*, pp. 254–263. <https://doi.org/10.1109/NCA.2016.7778628>
30. Pescher, F., Napp, N., Piranda, B., & Bourgeois, J. (2020). GAPCoD: A generic assembly planner by constrained disassembly. In *Autonomous agents and multiagent systems (AAMAS)*, pp. 1028–1036. <https://dl.acm.org/doi/abs/10.5555/3398761.3398881>.
31. Queffelec, A., Sankur, O., & Schwarzenrüber, F. (2023). Complexity of planning for connected agents in a partially known environment. *Theoretical Computer Science*, 941, 202–220. <https://doi.org/10.1016/j.tcs.2022.11.015>
32. Rubenstein, M., Cornejo, A., & Nagpal, R. (2014). Programmable self-assembly in a thousand-robot swarm. *Science*, 345(6198), 795–799. <https://doi.org/10.1126/science.1254295>
33. Sahin, E., & Winfield, A. F. T. (2008). Special issue on swarm robotics. *Swarm Intelligence*, 2(2), 69–72. <https://doi.org/10.1007/s11721-008-0020-6>
34. Schreckenberg, M., & Seltén, R. (2013). *Human behaviour and traffic networks*. New York: Springer. <https://doi.org/10.1007/978-3-662-07809-9>

35. Schwartz, J. T., & Sharir, M. (1983). On the piano movers' problem: III. Coordinating the motion of several independent bodies: the special case of circular bodies moving amidst polygonal barriers. *International Journal of Robotics Research*, 2(3), 46–75. <https://doi.org/10.1177/027836498300200304>
36. Soloveichik, D., & Winfree, E. (2007). Complexity of self-assembled shapes. *SIAM Journal on Computing*, 36(6), 1544–1569. <https://doi.org/10.1137/S0097539704446712>
37. Solovey, K., & Halperin, D. (2014). *k*-color multi-robot motion planning. *International Journal of Robotics Research*, 33(1), 82–97. <https://doi.org/10.1177/0278364913506268>
38. Solovey, K., & Halperin, D. (2016). On the hardness of unlabeled multi-robot motion planning. *International Journal of Robotics Research*, 35(14), 1750–1759. <https://doi.org/10.1177/0278364916672311>
39. Solovey, K., Jingjin, Yu., Zamir, O., & Halperin, D. (2015). Motion planning for unlabeled discs with optimality guarantee. *Robotics: Science and Systems*. <https://doi.org/10.15607/RSS.2015.XI.011>
40. Stern, R., Sturtevant, N. R., Felner, A., Koenig, S., Ma, H., Walker, T. T., Li, J., Atzmon, D., Liron Cohen T. K., Kumar, S., Barták, R., & Boyarski, E. (2019). Multi-agent pathfinding: Definitions, variants, and benchmarks. In *Symposium on combinatorial search (SOCS)*, pp. 151–159. <https://doi.org/10.1609/socs.v10i1.18510>
41. Thalamy, P., Piranda, B., & Bourgeois, J. (2019). Distributed self-reconfiguration using a deterministic autonomous scaffolding structure. In *Autonomous agents and multiagent systems (AAMAS)*, pp. 140–148. <https://dl.acm.org/doi/abs/10.5555/3306127.3331685>
42. Turpin, M., Michael, N., & Kumar, V. (2013). Trajectory planning and assignment in multirobot systems. In *Algorithmic foundations of robotics X*, pp. 175–190. https://doi.org/10.1007/978-3-642-36279-8_11
43. Turpin, M., Mohta, K., Michael, N., & Kumar, V. (2014). Goal assignment and trajectory planning for large teams of interchangeable robots. *Autonomous Robots*, 37(4), 401–415. <https://doi.org/10.1007/s10514-014-9412-1>
44. Yang, H., & Vigneron, A. (2021). A simulated annealing approach to coordinated motion planning. In *Symposium on computational geometry (SoCG)*. pp. 65:1–65:9. <https://doi.org/10.4230/LIPIcs.SoCG.2021.65>
45. Yu, J., & LaValle, S. M. (2012). Multi-agent path planning and network flow. In *Workshop on the algorithmic foundations of robotics (WAFR)*, pp. 157–173. https://doi.org/10.1007/978-3-642-36279-8_10

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Authors and Affiliations

Sándor P. Fekete¹ · Phillip Keldenich¹ · Ramin Kosfeld¹ · Christian Rieck¹ · Christian Scheffer²

✉ Sándor P. Fekete
s.fekete@tu-bs.de

Phillip Keldenich
p.keldenich@tu-bs.de

Ramin Kosfeld
r.kosfeld@tu-bs.de

Christian Rieck
c.rieck@tu-bs.de

Christian Scheffer
christian.scheffer@hs-bochum.de

¹ Department of Computer Science, TU Braunschweig, Braunschweig, Germany

² Department of Electrical Engineering and Computer Science, Bochum University of Applied Sciences, Bochum, Germany