

Multi-Robot Grasp Planning for Sequential Assembly Operations

Mehmet Dogar and Andrew Spielberg and Stuart Baker and Daniela Rus

Abstract—This paper addresses the problem of finding robot configurations to grasp assembly parts during a sequence of collaborative assembly operations. We formulate the search for such configurations as a constraint satisfaction problem (CSP). *Collision constraints* in an operation and *transfer constraints* between operations determine the sets of feasible robot configurations. We show that solving the connected constraint graph with off-the-shelf CSP algorithms can quickly become infeasible even for a few sequential assembly operations. We present an algorithm which, through the assumption of feasible *regrasps*, divides the CSP into independent smaller problems that can be solved exponentially faster. The algorithm then uses local search techniques to improve this solution by removing a gradually increasing number of regrasps from the plan. The algorithm enables the user to stop the planner *anytime* and use the current best plan if the cost of removing regrasps from the plan exceeds the cost of executing those regrasps. We present simulation experiments to compare our algorithm’s performance to a naive algorithm which directly solves the connected constraint graph. We also present a real robot system which uses the output of our planner to grasp and bring parts together in assembly configurations.

I. INTRODUCTION

We are interested in multi-robot systems which can perform sequences of assembly operations to build complex structures. Each assembly operation in the sequence requires multiple robots to grasp multiple parts and bring them together in space in specific relative poses. We present an example in Fig. 1 where a team of robots assemble chair parts by attaching them to each other with a fastener. Once an assembly operation is complete, the semi-assembled structure can be transferred to subsequent assembly operations to be combined with even more parts. We present an example sequence in Fig. 2.

This paper addresses the problem of finding robot base and arm configurations which grasp the assembly parts during a sequence of assembly operations.

The problem imposes a variety of constraints on the robot configurations. Take the assembly operation scene in Fig. 1. We immediately see one type of constraint: the robot bodies must not intersect. In effect, they must “share” the free space. The sequential nature of the task, however, may result in even more constraints. A robot may choose one of two strategies to move a semi-assembled structure from one assembly operation to the next (Fig. 2): The robot can *regrasp*, changing its grasp on the semi-assembled structure, or the robot can *transfer* the semi-assembled structure directly to the next operation, keeping the same grasp.

This work was supported by the Boeing Company.

Computer Science and Artificial Intelligence Lab, Massachusetts Institute of Technology, Cambridge, MA 02139 USA
{mdogar, aespielberg, spbaker, rus}@csail.mit.edu




Fig. 1: Three robots grasping parts at a specific assembly configuration while avoiding collision.

Both strategies have their advantages. If the robot chooses transfer, it avoids extra regrasp operations during execution. Regrasps, on the other hand, make the planning problem easier by decoupling sequence of operations from each other: In Fig. 2, since the robot commits to transfer the structure between assembly operations 1 & 2, it must plan a grasp of the part which works for both operations. The coupling between multiple operations makes it extremely expensive to solve problems with long sequences of assembly operations

Humans use a combination of both strategies during manipulation: we regrasp when we need to, but we are also able to use transfer grasps which work for more than one operation. Given a sequence of assembly operations, how can a team of robots decide when to regrasp and when to transfer? We present a planner with this capability: Our algorithm trades off between regrasps and transfers while generating collision-free robot configurations for each assembly operation.

We formulate multi-robot grasp planning as a constraint satisfaction problem (CSP). In this representation every robotic grasp in every assembly operation becomes a variable. Every variable must be assigned a robot configuration which grasps a particular part or semi-assembled structure. We impose two types of constraints: *collision constraints* between variables of the same assembly operation; and *transfer constraints* between variables in subsequent operations.

Ideally, a plan involves no regrasps and the assembly is transferred between operations smoothly. Trying to find a plan with no regrasps, however, means having transfer constraints between all operations. A complete solution requires

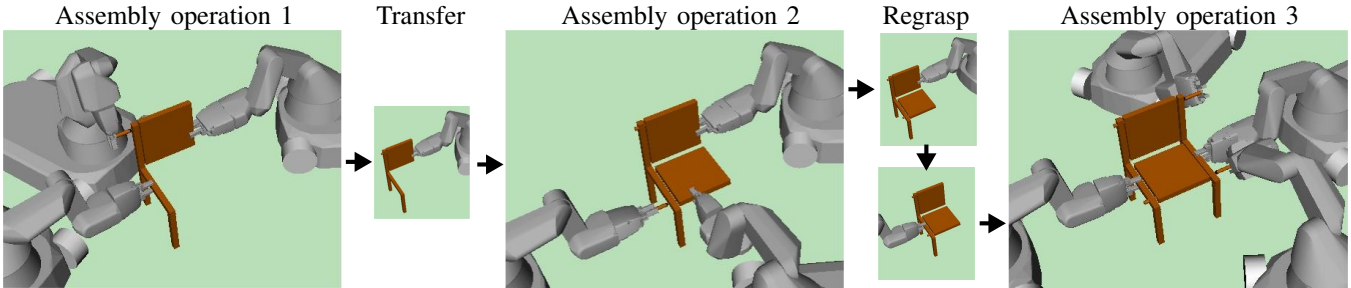


Fig. 2: Sequential assembly operations for of a chair.

solving for all the assembly operations at once. In general, complete CSP solvers display exponential complexity with respect to the number of variables [1]. Solving the multi-robot grasp planning problem then becomes exponentially expensive with increasing number of assembly operations.

Instead, our algorithm starts with a strategy to perform regrasps between all operations. Our key assumption is that, regrasps between any two grasps (possibly through a series of intermediate grasps) are always feasible. This decouples assembly operations from each other. The resulting problem can be solved by solving a small CSP separately for each assembly operation.

After finding this initial solution, our algorithm continues to find solutions with fewer regrasps by imposing sets of transfer constraints. As such solutions are found, the algorithm increases the number of transfer constraints imposed.

Our algorithm is an anytime planner: Given more time, it generates plans with fewer regrasps and more transfers. The algorithm enables the user to stop the planner and use the current best plan if the cost of removing regrasps from the plan exceeds the cost of executing those regrasps.

When imposing a new set of transfer constraints, our algorithm does *not* solve the CSP from scratch: Solutions with fewer (or no) transfer constraints are readily available from previous cycles. We use state-of-the-art local search methods for CSPs, which can be initialized with partial solutions. Local-search methods work only in a locality of the constraint graph and therefore their runtime is not affected by the full size of the CSP [1], leading to fast updates.

A. Related work

Recent work by Lozano-Pérez and Kaelbling [2] also represent sequential manipulation problems as CSPs. These geometric CSPs are formulated by a higher-level task planner. Their focus is on the interface between the task planner and CSP formulation, and they propose methods for constructing the CSPs efficiently. The CSPs are solved by an off-the-shelf solver. We propose an algorithm to solve the CSP itself by using domain-specific assumptions, such as feasible regrasps.

The effectiveness and necessity of regrasping during manipulation have been recognized [3, 4]. We show that assuming feasibility of regrasps we can simplify the CSP solutions of manipulation plans significantly. Structures similar to the

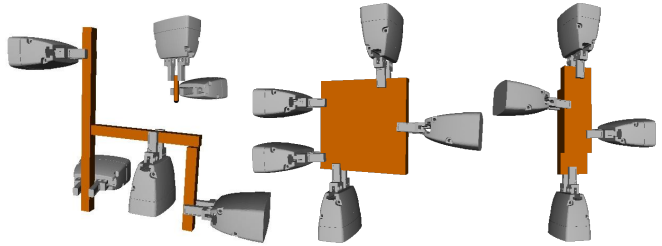


Fig. 3: Example grasps for assembly parts.

grasp-placement space [5] or the *grasp-graph* [6] can be precomputed to satisfy our regrasp feasibility assumption.

Our algorithm takes as input a sequence of relative poses of assembly parts. *Assembly planning* [7, 8] addresses the problem of finding such sequences. In this paper we find *robot configurations* to realize an assembly plan.

Other grasp planners that take into account task constraints [9, 10, 11] and multiple robots [12] exist. Unlike previous work, we focus on planning such grasps in a sequential *and* multi-robot context.

We use complete and local methods to solve CSPs. There is extensive literature in this area but the treatment in Russell and Norvig [1] covers all the methods we use.

II. PROBLEM

An assembly is a collection of simple parts at specific relative poses. A simple part by itself is also a (trivial) assembly. Robots perform an *assembly operation*, $o = (\mathbf{A}_{\text{in}}, a_{\text{out}}, p)$, to produce an output assembly a_{out} from a set of input assemblies \mathbf{A}_{in} . We also assume that a three-dimensional pose in the environment, p , is specified as the location of an operation.

During an assembly operation, input assemblies \mathbf{A}_{in} must be grasped and supported by robots at their respective poses in a_{out} at operation pose p . We assume that a local controller exists to perform the fastening/screwing, once the parts are at the poses specified by the assembly operation.

Note that our definition of an assembly operation also applies to the grasp of a single part a , where $\mathbf{A}_{\text{in}} = \{a\}$ is a singleton, $a_{\text{out}} = a$, and p is the pose of a .

A robot can grasp an assembly by placing its gripper at certain poses on the assembly. We assume we can compute a set of such poses, *grasps*, for each assembly a . We illustrate

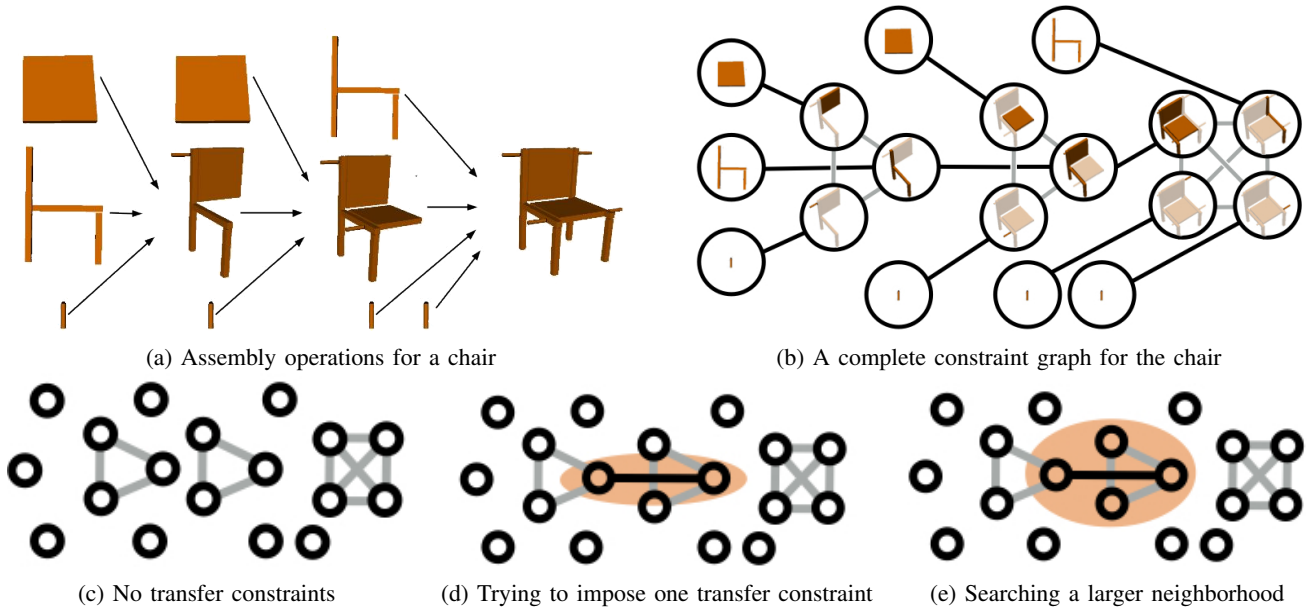


Fig. 4: The chair assembly example.

example grasps for simple parts in Fig. 3. We use \mathbf{Q} to represent the robot configuration space, which includes base pose and arm joint configurations. If a configuration $q \in \mathbf{Q}$ places the robot gripper at a grasping pose for assembly a during operation o , we say that “ q is grasping a during o ”. The robots must avoid collision during assembly operations.

Robots perform a sequence of assembly operations $\mathbf{O} = [o_i]_{i=1}^N$ to gradually build large complex structures: output assemblies of earlier operations are used as inputs in later operations. Robots move the assemblies from one operation to the next.

As an example, we present a sequence of assembly operations to build a chair in Fig. 4a. This example includes eleven operations: three operations in which multiple parts must be assembled, and eight operations where a single part must be grasped at its initial pose. Each arrow indicates an instance where robots move an assembly from one operation to the next.

Given a sequence of assembly operations, we formulate the problem of *multi-robot grasp planning for sequential assembly operations* as finding grasping configurations for all the robots required by the assemblies in all the operations.

A. Moving assemblies between operations

Suppose $o = (\mathbf{A}_{\text{in}}, a_{\text{out}}, p)$ and $o' = (\mathbf{A}'_{\text{in}}, a'_{\text{out}}, p')$ are two operations such that $a_{\text{out}} \in \mathbf{A}'_{\text{in}}$; i.e. the output assembly of o is one of the input assemblies of o' . We call o and o' *sequential operations*. a_{out} must be moved to o' after o is completed. There are two ways this can be done: *transfer* and *regrasp*.

To directly transfer a_{out} , one of the robots grasping an assembly in \mathbf{A}_{in} can keep its grasp and carry a_{out} to o' . There is flexibility; any $a \in \mathbf{A}_{\text{in}}$ may be used for the transfer. For example, after the first assembly operation in Fig. 2, the assembled structure can be transferred either by the grasp on

the back of the chair as in the figure, or alternatively by the grasp on the side of the chair.

The alternative is to regrasp a_{out} after o is completed. Robots can regrasp an assembly in different ways: e.g. by first placing it on the floor in a stable configuration and then grasping it again, or with the help of other robots which can temporarily grasp and support the assembly while it is being regrasped. The important implication for our planning problem is that the new grasp of a_{out} can be different from the grasps of all $a \in \mathbf{A}_{\text{in}}$. An example is the regrasp after the second assembly operation in Fig. 2.

III. CSP FORMULATION

Given a sequence of assembly operations \mathbf{O} , we can formulate multi-robot grasp planning as a CSP.

A CSP is defined by a set of variables \mathbf{X} , a set of possible values $\mathbf{V}(x)$ that each variable x can be assigned with, and a set of constraints specifying consistent assignments of values to variables. A solution to the CSP is an assignment of values to all the variables that is consistent with all the constraints.

Variables: For our problem, we create one variable for the grasp of each input assembly of each assembly operation. We use ${}^o x^a$ to represent the variable corresponding to the grasp of assembly $a \in \mathbf{A}_{\text{in}}$ of operation $o \in \mathbf{O}$.

Values: The set of values for the variable ${}^o x^a$ is the set of robot configurations grasping the assembly:

$$\mathbf{V}({}^o x^a) = \{q \in \mathbf{Q} \mid q \text{ is grasping } a \text{ during } o.\}$$

In general there can be a continuous set of robot configurations grasping a , due to redundancy in the kinematics or due to a continuous representation of grasping gripper poses on a part. We discretize this continuous set by sampling uniformly at a fine resolution.

Constraints: We define two sets of constraints: *collision constraints* and *transfer constraints*. A collision constraint

$c(x, x')$ enforces that two robot configurations assigned to x and x' do not collide. We create a collision constraint $c({}^o x^a, {}^o x^{a'})$ between each pair of variables of the same operation o .

A transfer constraint $t(x, x')$ enforces that robot configurations assigned to x and x' grasp the same part while placing the robot gripper at the same pose on the part.

Given any two sequential assembly operations $o = (\mathbf{A}_{\text{in}}, a_{\text{out}}, p)$, $o' = (\mathbf{A}'_{\text{in}}, a'_{\text{out}}, p')$, and an assembly $a \in \mathbf{A}_{\text{in}}$, we can create a transfer constraint between two variables $t({}^o x^a, {}^{o'} x^{a_{\text{out}}})$. If the CSP with this constraint has a solution, then the assembly a_{out} can be transferred directly from o to o' using the grasp on a . Each different choice of the transfer assembly $a \in \mathbf{A}_{\text{in}}$ corresponds to a different transfer constraint we can impose. Solving for any one of these transfer constraints is sufficient, however.

We can also choose not to create any transfer constraints between o and o' . Our underlying assumption here is that, whatever new grasp is required by a_{out} during o' , it will be feasible to achieve it with a regrasp after o — possibly through a number of intermediate grasps. This is a reasonable assumption in our domain where there is ample space in the environment for robots to change from one feasible grasp to another feasible grasp.

Given a CSP, we can represent the variables and constraints in a constraint graph. In this graph, there is a node for each CSP variable, and an edge between two nodes if a constraint exists between the variables. In Fig. 4b we show a constraint graph for the chair assembly. Each node corresponds to the grasp of a certain part during a certain operation. In the figure, we show the image for the operation inside the node and highlight the image of the part which should be grasped. Light gray edges correspond to collision constraints, and dark edges correspond to transfer constraints. In this graph all operations are connected with transfer constraints: If we can find a solution the robots will not need to perform any regrasps.

A. Solving a CSP

Backtracking search is a widely used and *complete* algorithm for solving CSPs. It searches forward by assigning values to variables such that all assignments obey the constraints. If at any point the algorithm cannot find a value for a variable which obeys the constraints, it backtracks by undoing the most recent assignment. The search continues until an assignment is found for all variables. If there is no solution, backtracking search tries all combinations of value assignments. The worst-case time complexity of backtracking search is exponential in the number of CSP variables. One can use domain-independent heuristics to prune the search space. *Minimum remaining value* and *forward-checking* [1] are two widely used heuristics.

Another approach to solving CSPs is by focusing on a local neighborhood of the constraint graph so that the computation time is not affected by the total size of the graph. These local techniques start with an initial assignment of values to variables, identify the conflict regions in the

constraint graph, and try to resolve the conflicts only in the local neighborhood. One can use different methods in the local neighborhood, e.g. a complete method like the backtracking search or a heuristic-based search like the *min-conflicts* [13] algorithm which greedily minimizes the number of conflicts in the graph.

IV. ALGORITHM

We would like to find solutions which involve a small number of regrasps, since each regrasp in the solution will require extra time to plan and execute.

A naive way to find solutions with minimum number of regrasps would be to create transfer constraints between all operations and try to solve the resulting CSP (e.g. the graph in Fig. 4b) with an algorithm such as backtracking search. If this succeeds we have found a solution with no regrasps. If it fails, we can remove one of the transfer constraints and try to solve the resulting CSP problem again to find a solution with one regrasp. If this fails, we can try removing a different transfer constraint, and if that fails, we can try removing two transfer constraints to find a solution with two regrasps; and so on. We call this the *naive CSP solution*.

The problem with the naive CSP solution is that it tries to solve the most difficult problems first: The CSP graph where operations are connected with transfer constraints make the search space exponentially larger. As we will show in the results, this approach quickly becomes infeasible, requiring hours to solve problems with only a few operations.

Instead, we propose an algorithm (Alg. 1) which works in the opposite direction: it first solves the easiest problem, the constraint graph with no transfer constraints, and then tries to improve the solution by imposing an increasing number of transfer constraints as more time is given.

This approach has two advantages. First, it leads to an *anytime planner* which produces a solution quickly and improves it as more time is given. The planner can be stopped anytime after the initial solution has been achieved and the current solution with the minimum number of regrasps can be used. This, for example, enables the user to stop planning if the planning time spent on imposing new transfer constraints exceeds the time which is required to plan and execute those regrasps. Second, this approach enables the use of local search algorithms to quickly identify easy-to-solve transfer constraints. We would like to solve easy transfer constraints first since we want to minimize the number of regrasps as much as possible before the time allocated to the planner runs out.

A. Generating the “All-Regrasps” Plan

We first assume no transfer constraints between operations. Collision constraints remain, but they only constrain variables within an operation. Hence, the constraint graph is divided into N connected components where N is the number of assembly operations. In Fig. 4c, we show this graph for the chair assembly example.

We solve each of these connected components separately using a complete CSP solver (lines 1-2 in Alg. 1). Any

Algorithm 1 Multi-Robot Grasp Planning for Assembly

Input: $\mathbf{O} = [o_i]_{i=1}^N$ is a sequence of assembly operations.

- 1: **for each** o_i **in** \mathbf{O} **do**
- 2: $sol[o_i] \leftarrow \text{BACKTRACKINGSEARCH}(o_i)$
- 3: $best \leftarrow \{sol[o_i]\}_{i=1}^N$
- 4: **for** $n = 1$ **to** $\text{MaxTransferConstraints}$ **do**
- 5: $best \leftarrow \text{SOLVETRANSFERCONSTRAINTS}(n, best)$
- 6: **procedure** $\text{SOLVETRANSFERCONSTRAINTS}(n, seed)$
- 7: **for** enlarging neighborhood h **do**
- 8: **for each** \mathbf{T} **in** $\text{TransferConstCombinations}(n)$ **do**
- 9: $sol \leftarrow \text{SOLVECSPLOCAL}(\mathbf{T}, h, seed)$
- 10: **if** sol **exists** **then**
- 11: **return** sol

complete CSP solver can be used. We use an implementation of backtracking search with minimum remaining value and forward checking. The combination of solutions for all operations give the overall “all-regrasps” solution (line 3). The “all-regrasps” solution is an inefficient, though valid, plan.

B. Imposing Transfer Constraints

Once the “all-regrasps” solution is found, our algorithm starts imposing a gradually increasing number of transfer constraints (lines 4-5). Fig. 4d shows one example transfer constraint added to the graph. The procedure $\text{SOLVETRANSFERCONSTRAINTS}$ attempts to find a solution to n transfer constraints. If a solution is found, it is recorded as the new best solution, and the algorithm progresses to $n + 1$ transfer constraints. One can stop the algorithm anytime and use the current best solution.

The procedure $\text{SOLVETRANSFERCONSTRAINTS}$ tries to solve for n transfer constraints as quickly as possible. It iterates over all valid n -combinations of transfer constraints (line 8). During this iteration we prioritize combinations which include smaller combinations that we have previously found solutions for. If we cannot find a solution for these prioritized combinations we then try all combinations.

Instead of searching the complete graph and losing time on difficult combinations, our algorithm performs local search (line 9) which succeeds or fails quickly. Local search variables are initialized with values from the current best solution ($seed$). Local search neighborhood size starts small (Fig. 4d) but gets larger (Fig. 4e) if no solution can be found (line 7).

C. Analysis

We analyze several important properties of our algorithm.

1) Completeness:

Proposition 4.1: Algorithm 1 is resolution-complete.

Proof: We use a discrete CSP representation which requires the discretization of the robot configuration space. Assume we are given a resolution with which to discretize. If the algorithm is unable to find a solution with no transfers (as computed in line 3), then the only constraints that the algorithm is unable to satisfy must be those within assembly

operations (i.e. collision constraints). This implies one of the following: either the input problem itself is infeasible, or no solution exists at the given resolution of discretization. At a high enough sampling resolution, the second problem disappears. ■

2) *Optimality:* We define optimality as returning the solution requiring the minimum number of regrasps. We do not necessarily aim for optimality: if the time required to remove more regrasps from the plan is more than the time required to execute those regrasping operations, we would like to stop planning and start execution. For this reason, in our implementation we use the greedy *min-conflicts* algorithm for our local search. In practice we have found it to produce good results, however, min-conflicts does not guarantee optimality and may get stuck in local minima.

Alg. 1, nevertheless, can be turned into an optimal planner if a complete algorithm, e.g. backtracking search, is used to search the local neighborhood.

Proposition 4.2: If a complete local search is used, then Algorithm 1 returns the minimum regrasp solution.

Proof: Our algorithm can terminate immaturely with a suboptimal solution only when it cannot improve the solution via local search for a given number of constraints (line 5 in Alg. 1). However, our algorithm will expand the local neighborhood to include the entire graph before failing (line 7). If the local search is complete, then this becomes a complete graph search, and a complete graph search must always find an improvement if it exists. The algorithm cannot terminate if it has not found an optimal solution, and thus it will always return the optimal solution. ■

3) *Complexity:* The naive CSP solution has an exponential runtime $\mathcal{O}(\exp(n*m))$, where n is the maximum number of robots involved in assembly operations and m is the number of assembly operations ($m * n$ is the total number of CSP variables). By comparison, our algorithm’s initial solution has runtime $\mathcal{O}(m \exp(n))$ — exponential in the number of robots per assembly operation but linear in the number of assembly operations (since each operation can be solved independently). Since n is typically very small in practice, finding initial solutions is generally quick. The complexity associated with improving the initial solution depends on the local search technique used. If a complete method such as backtracking search is used, the complexity of improving the solution will approach the complexity of the naive CSP algorithm as more transfer constraints are imposed. We have, however, found the *min-conflicts* greedy search to be a good trade-off between improvement speed and optimality. As we show in §V min-conflicts improve the solution quickly and reduces the number of regrasps effectively. This is very practical for real-world applications where a small number of regrasps is feasible.

V. EXPERIMENTS AND RESULTS

We implemented and evaluated our algorithm on an example chair assembly operation. We show the sequence of operations in Fig. 4a. The number of robots required by the complex operations are 3, 3, and 4, respectively.

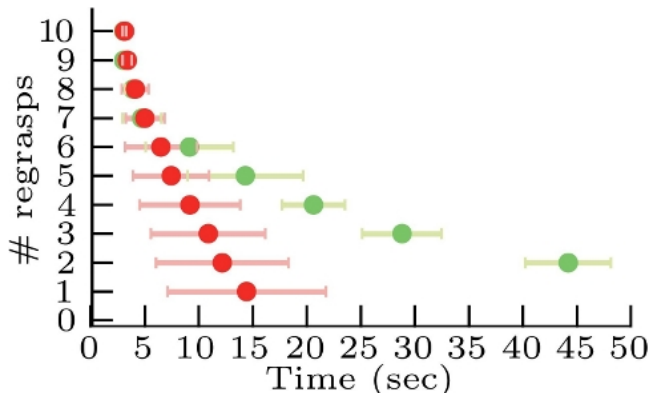


Fig. 5: Time to generate plans with decreasing number of regrasps. In 17 out of 20 runs (red points) our algorithm generated a plan with only one regrasp. In 3 out of 20 runs (light green points) our algorithm generated a plan with two regrasps.

The operations require the semi-assembled structures to be transferred twice and simple parts to be transferred eight times, totaling to ten potential regrasps. We implemented our algorithm and evaluated it in the OpenRAVE environment [14] with four KUKA YouBot robot models¹. We presented the chair parts to the robots in an environment with some obstacles, presented in Fig. 6a. These obstacles make the problem even more constrained making a no-regrasp solution impossible. Particularly one of the chair side parts must be regrasped after its initial grasp.

We ran our algorithm on the chair example 20 times. In 17 of these runs, our algorithm found the optimal solution with one regrasp and in 3 runs it generated a solution with two regrasps. We plot how our algorithm reduces the number of regrasps with time in Fig. 5. We plot the results for the 17 runs with one-regrasp solutions (red points) separately from the 3 runs with two-regrasps solutions (light green points) since they display different and consistent trends. Each data point marks the average time it took our algorithm to produce a plan with the number of regrasps given on the vertical axis. The horizontal bars show the standard deviations. Our algorithm generates the “all-regrasps” solution in about 4 seconds and then improves the solution every few seconds. The difference between the trends that find one-regrasp solutions (red points) and two-regrasp solutions (light green points) stems from the greedy nature of the min-conflicts algorithm. In the 3 runs (15%) which generate two-regrasp results, our algorithm gets stuck at a local minima. Nevertheless, min-conflicts is able to reach the one-regrasp plan in 17 (85%) of the cases.

We present part of an example plan in Fig. 6(b)-(d). The robots are able to transfer the assembly between the complex operations without a regrasp: The left-most robot holding the assembly onto the side of the chair keeps its grasp fixed and transfers the assembly between all three operations.

We also compare the performance of our algorithm with

	Alg. 1	Naive optimal
Chair	18.9 (12.7)	> 3600 (?)
Grid	20.5 (10.1)	> 3600 (?)

TABLE I: Average planning times in seconds. Standard deviations are shown in parantheses. The naive algorithm had to be stopped after running for an hour.

the naive CSP solution mentioned in §IV. This algorithm is optimal, but also naive in that it tries to solve the full CSP at once. In Tab. I we present the planning times of our algorithm (all 20 runs combined) and the naive optimal algorithm for the plans they can generate with the minimum number of regrasps. Our algorithm can generate plans which include one or two regrasps and is at least two orders of magnitude faster. In these experiments we imposed a one-hour time limit, and the naive algorithm exceeded this limit in 5 out of 5 runs.

In Fig. 7 we show another example where robots build a grid structure using square building blocks. Each operation except the first one involves five robots. We ran our planner on this example ten times. It was able to find four transfers in all ten cases. We present the average time it took our algorithm to find these solutions in Tab. I. The naive algorithm, again, exceeded the time limit for these experiments.

The last scene in Fig. 7 shows a possible failure mode for our planner. While the robot configurations are valid, one of the robots become trapped inside the structure. As our planner does not check for reachability to and from these configurations one may expect such problems. A possible solution is to impose a new *reachability constraint* to the robot configurations and use a motion planner to check against this constraint while solving the CSP. Lozano-Pérez and Kaelbling [2] propose methods to perform these kinds of checks in a fast way.

A. Real robot implementation

We are building a real robot team to perform autonomous assembly of complex structures. The algorithm presented in this paper provides our system with the sequences of configurations in which to grasp and assemble parts, enabling fast planning and minimal regrasping operations. In Fig. 8 we present snapshots from the execution of an assembly plan generated by our algorithm. The complete execution can be seen in the video accompanying this paper.

Our system consists of three KUKA Youbot robots, each with an omni-directional base, a 5 degree-of-freedom arm, and a parallel plate gripper. Perception in our system is provided by a motion capture system² which is able to detect and track infra-red reflective markers. We localize our robots and the initial location of assembly parts using such markers. We use an RRT planner [15] to move the robots between configurations generated by our algorithm.

We present the initial placement of parts and robots in Fig. 8a. In Fig. 8b and Fig. 8c two robots use the planned

¹<http://www.youbot-store.com>

²<http://www.vicon.com>

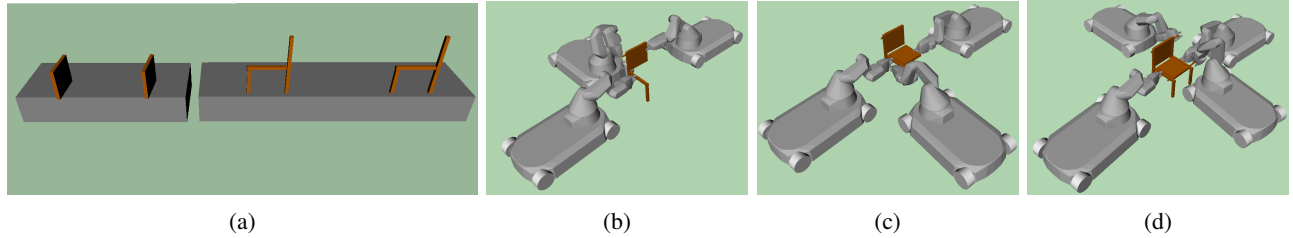


Fig. 6: (a) Initial locations of chair parts. (b)-(d) Solution for the assembly of our chair example.

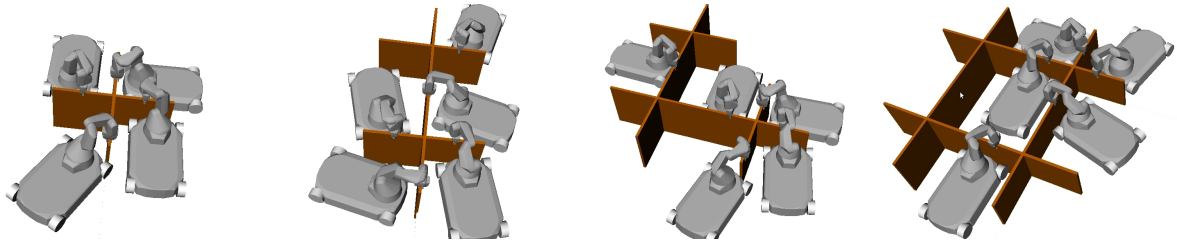


Fig. 7: A plan for a five-robot team building a grid using square blocks.

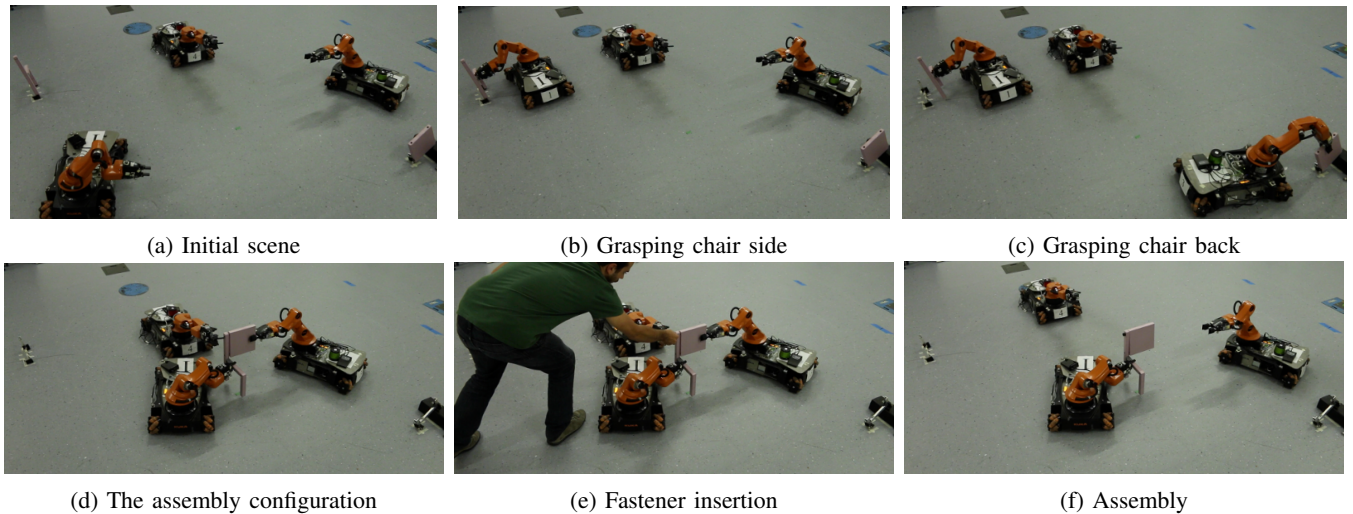


Fig. 8: Multi-robot execution of a chair assembly plan.

configurations to grasp the chair parts. The third robot holds a fastener, and the robots bring these three parts together in Fig. 8d, again using planned configurations. The robots keep the same grasp on the parts through these operations, enabling them to transfer parts between Fig. 8b-Fig. 8d and Fig. 8c-Fig. 8d.

While our robots can successfully use the planner output to bring parts to assembly configurations, they need to perform highly precise manipulation operations to actually insert the fastener and attach the parts to each other. We are currently developing controllers and tools [16] to perform these operations. In this example human intervention is needed (Fig. 8e). Fig. 8f shows the assembled structure, ready to be transferred to the next operation.

VI. FUTURE WORK

While this paper has addressed the high-level planning aspects of coordinated assembly, a system that can go from a design input to complete assemblies with minimal human assistance requires the development of more advanced techniques both in control and reasoning.

Our system is able to plan robot configurations to bring parts into specific relative configurations. However, error accumulated from factors such as loose part grips and localization uncertainty require on-board local controllers for fine manipulation operations.

Our current formulation addresses geometric collision constraints during assembly operations. Other constraints, however, must also be taken account for a robust system capable

of a wide variety of assembly operations. An important class of constraints are the *stability constraints* which require that the grasps on assembly keeps it stable with respect to gravity and other forces that arise during an assembly operation such as screwing.

REFERENCES

- [1] S. J. Russell and P. Norvig, *Artificial Intelligence: A Modern Approach*, 2nd ed., 2003.
- [2] T. Lozano-Pérez and L. P. Kaelbling, "A constraint-based method for solving sequential manipulation planning problems," in *IROS*, 2014.
- [3] T. Lozano-Pérez, J. Jones, E. Mazer, P. O'Donnell, W. Grimson, P. Tournassoud, and A. Lanusse, "Handey: A robot system that recognizes, plans, and manipulates," in *ICRA*, 1987.
- [4] T. Siméon, J.-P. Laumond, J. Cortés, and A. Sahbani, "Manipulation planning with probabilistic roadmaps," *IJRR*, vol. 23, no. 7-8, 2004.
- [5] P. Tournassoud, T. Lozano-Pérez, and E. Mazer, "Regrasping," in *ICRA*, 1987.
- [6] N. Daffe, A. Rodriguez, R. Paolini, B. Tang, S. Srinivasa, M. Erdmann, M. Mason, I. Lundberg, H. Staab, and T. Fuhlbrigge, "Extrinsic dexterity: In-hand manipulation with external forces," in *ICRA*, 2014.
- [7] R. H. Wilson and J.-C. Latombe, "Geometric reasoning about mechanical assembly," *Artificial Intelligence*, vol. 71, no. 2, 1994.
- [8] R. A. Knepper, T. Layton, J. Romanishin, and D. Rus, "Ikeabot: An autonomous multi-robot coordinated furniture assembly system," in *Robotics and Automation (ICRA), 2013 IEEE International Conference on*. IEEE, 2013, pp. 855–862.
- [9] D. Berenson and S. S. Srinivasa, "Grasp synthesis in cluttered environments for dexterous hands," in *Humanoids*, 2008.
- [10] D. Berenson, S. S. Srinivasa, and J. Kuffner, "Task space regions: A framework for pose-constrained manipulation planning," *IJRR*, 2011.
- [11] H. Dang and P. K. Allen, "Semantic grasping: Planning robotic grasps functionally suitable for an object manipulation task," in *IROS*, 2012.
- [12] N. Vahrenkamp, E. Kuhn, T. Asfour, and R. Dillmann, "Planning multi-robot grasping motions," in *Humanoid Robots (Humanoids), 2010 10th IEEE-RAS International Conference on*. IEEE, 2010, pp. 593–600.
- [13] S. Minton, M. D. Johnston, A. B. Philips, and P. Laird, "Minimizing conflicts: a heuristic repair method for constraint satisfaction and scheduling problems," *Artificial Intelligence*, vol. 58, no. 1, pp. 161–205, 1992.
- [14] R. Diankov, "Automated construction of robotic manipulation programs," Ph.D. dissertation, CMU, Robotics Institute, August 2010.
- [15] J. J. Kuffner and S. M. LaValle, "Rrt-connect: An efficient approach to single-query path planning," in *Robotics and Automation, 2000. Proceedings. ICRA'00. IEEE International Conference on*, vol. 2. IEEE, 2000, pp. 995–1001.
- [16] M. Dogar, R. A. Knepper, A. Spielberg, C. Choi, H. I. Christensen, and D. Rus, "Towards coordinated precision assembly with robot teams," *Proceedings of the International Symposium of Experimental Robotics, ISER*, 2014.