

High-Quality Tabletop Rearrangement with Overhand Grasps: Hardness Results and Fast Methods

Shuai D. Han, Nicholas M. Stiffler, Athanasios Krontiris, Kostas E. Bekris, and Jingjin Yu
Computer Science Department

Rutgers, the State University of New Jersey, Piscataway, New Jersey, USA

Email: {shuai.han, nick.stiffler, tdk.krontir, kostas.bekris, jingjin.yu} @ rutgers.edu

Abstract—This paper studies the underlying combinatorial structure of a class of object rearrangement problems, which appear frequently in applications. The problems involve multiple, similar-geometry objects placed on a flat, horizontal surface, where a robot can approach them from above and perform pick-and-place operations to rearrange them. The paper considers both the case where the start and goal object poses overlap, and where they do not. For overlapping poses, the primary objective is to minimize the number of pick-and-place actions and then to minimize the distance traveled by the end-effector. For the non-overlapping case, the objective is solely to minimize the end-effector distance. While such problems do not involve all the complexities of general rearrangement, they remain computationally hard challenges in both cases. This is shown through two-way reductions between well-understood, hard combinatorial challenges and these rearrangement problems. The benefit of the reduction is that there are well studied algorithms for solving these well-established combinatorial challenges. These algorithms can be very efficient in practice despite the hardness results. The paper builds on these reduction results to propose an algorithmic pipeline for dealing with the rearrangement problems. Experimental evaluation shows that the proposed pipeline achieves high-quality paths with regards to the optimization objectives. Furthermore, it exhibits highly desirable scalability as the number of objects increases in both the overlapping and non-overlapping setups.

I. INTRODUCTION

In many industrial and logistics applications, such as those shown in Fig. 1, a robot is tasked to rearrange multiple, similar objects placed on a tabletop into a desired arrangement. In these setups, the robot needs to approach the objects from above and perform a *pick-and-place action* at desired target poses. Such operations are frequently part of product packaging and inspection processes. Efficiency plays a critical role in these domains, as the speed of task completion has a direct impact on financial viability. Beyond industrial robotics, a home assistant robot may need to deal with such problems as part of a room cleaning task. The reception of such a robot by people will be more positive if its solutions are efficient and does not waste time performing redundant actions. Many subtasks affect the efficiency of the overall solution in all of these applications, ranging from perception to the robot’s speed in grasping and transferring objects. But overall efficiency critically depends on the underlying combinatorial aspects of the problem, which relate to the number of pick-and-place actions that the robot performs, the placement of the objects, as well as the sequence of objects transferred.



Fig. 1. Examples of robots deployed in industrial settings tasked to arrange objects in desired configurations through pick and place: (left) ABB’s IRB 360 FlexPicker rearranging pancakes (right) Stäubli’s TP80 Fast Picker robot.

This paper deals with the combinatorial aspects of such object rearrangement tasks. The objective is to understand the underlying structure and obtain high-quality solutions in a computationally efficient manner. The focus is on a subset of general rearrangement problems, which relate to the above mentioned applications. In particular, the setup corresponds to rearranging multiple, similar-geometry, non-stacked objects on a flat, horizontal surface from given initial to target arrangements. The robot can approach the objects from above, pick them up and raise them. At that point, it can move them freely without collisions with other objects.

There are two important variations of this problem. The first requires that the target object poses do not overlap with the initial ones. In this scenario, the number of pick-and-place actions is equal to the number of objects not in their goal pose. Thus, the solution quality is dependent upon the sequence with which the objects are transferred. A good sequence can minimize the distance that the robot’s end-effector travels. The second variant of the problem allows for target poses to overlap with the initial poses, as in Fig. 2. The situation sometimes necessitates the identification of intermediate poses for some objects to complete the task. In such cases, the quality of the solution tends to be dominated by the number of intermediate poses needed to solve the problem, which correlates to the number of the pick-and-place actions the robot must carry out. The primary objective is to find a solution, which uses the minimum number of intermediate poses and among them minimize the distance the robot’s end-effector travels.

Both variations include some assumptions that simplify these instances relative to the general rearrangement problem. The non-overlapping case in particular seems to be quite easy since a random feasible solution can be trivially acquired. Nevertheless, this paper shows that even in this simpler setup, the optimal variant of the problem remains computationally




Fig. 2. An example of an object rearrangement challenge considered in this work where the initial (middle) and final (right) object poses are overlapping and an object needs to be placed at an intermediate location. Images from a V-REP [4] simulation.

hard. This is achieved by reducing the Euclidean-TSP problem [1] to the cost-optimal, non-overlapping tabletop object rearrangement problem. Even in the unlabeled case, where objects can occupy any target pose, the problem is still hard. For overlapping initial and final poses, the paper employs a graphical representation from the literature [2], which leads to the result that finding the minimum number of pick-and-place actions relates to a well-known problem in the algorithmic community, the “Feedback Vertex Set” (FVS) problem [3]. This again indicates the hardness of the challenge.

The benefit of these two-way reductions, beyond the hardness results themselves, is that they suggest algorithmic solutions and provide an expectation on the practical efficiency of the methods. In particular, Euclidean-TSP admits a polynomial-time approximation scheme (PTAS) and good heuristics, which implies very good practical solutions for the non-overlapping case. On the other hand, the FVS problem is APX-hard [3, 5], which indicates that efficient algorithms are harder for the overlapping case. This motivated the consideration of alternative heuristics for solving such challenges that make sense in the context of object rearrangement.

The algorithms proposed here, which arose by mapping the object rearrangement variations to established, well-studied problems, have been evaluated in terms of practical performance. For the non-overlapping case, an alternative solver exists that was developed for a related challenge [6]. The TSP solvers achieve superior performance relative to this alternative when applied to object rearrangement. They achieve sub-second solution times for hundreds of objects. Optimal solutions are shown to be significantly better than the average, random, feasible solution. For the overlapping case, exact and heuristic solvers are considered. The paper shows that practically efficient methods achieve sub-second solution times without a major impact in solution quality for tens of objects.

II. CONTRIBUTION RELATIVE TO PRIOR WORK

Multi-body planning is a related challenge that is itself hard. In the general, continuous case, complete approaches do not scale even though methods exist that try to decrease the effective DOFs [7]. For specific geometric setups, such as unlabeled unit-discs among polygonal obstacles, optimality can be achieved [8], even though the unlabeled case is still hard [9]. Given the hardness of multi-robot planning, decoupled methods, such as priority-based schemes [10] or velocity tuning [11], trade completeness for efficiency. Assembly planning [12, 13, 14] deals with similar problems but few optimality arguments have been made.

Recent progress has been achieved for the discrete variant of the problem, where robots occupy vertices and move along edges of a graph. For this problem, also known as “pebble motion on a graph” [15, 16, 17, 18], feasibility can be answered in linear time and paths can be acquired in polynomial time. The optimal variation is still hard but recent optimal solvers with good practical efficiency have been developed either by extending heuristic search to the multi-robot case [19, 20], or utilizing solvers for other hard problems, such as network-flow [21, 22]. The current work is motivated by this progress and aims to show that for certain useful rearrangement setups it is possible to come up with practically efficient algorithms through an understanding of the problem’s structure.

Navigation among Movable Obstacles (NAMO) is a related computationally hard problem [23, 24, 25, 26], where a robot moves and pushes objects. A probabilistically complete solution exists for this problem [27]. NAMO can be extended to manipulation among movable obstacles (MAMO) [28] and rearrangement planning [29, 30]. Monotone instances for such problems, where each obstacle may be moved at most once, are easier [28]. Recent work has focused on “non-monotone” instances [31, 32, 33, 34, 35, 36]. Rearrangement with overlaps considered in the current paper includes “non-monotone” instances although other aspects of the problem are relaxed. In all these efforts, the focus is on feasibility and no solution quality arguments have been provided. Asymptotic optimality has been achieved for the related “minimum constraint removal” path problem [37, 38], which, however, does not consider negative object interactions.

The *Pickup and Delivery Problem* (PDP) [39, 40] is a well-studied problem in operations research that is similar to tabletop object rearrangement, as long as the object geometries are ignored. The PDP models the pickup and delivery of goods between different parties and can be viewed as a subclass of vehicle routing [41] or dispatching [42]. It is frequently specified over a graph embedded in the 2D plane, where a subset of the vertices are pickup and delivery locations. A PDP in which pickup and delivery sites are not uniquely paired is also known as the NP-hard swap problem [43, 44], for which a 2.5-optimal heuristic is known [44]. Many exact linear programming algorithms and approximations are available [45, 46, 47] when pickup and delivery locations overlap, where pickup must happen some time after delivery. The stacker crane problem (SCP) [48, 6] is a variation of PDP of particular relevance as it maps to the non-overlapping case of labeled object rearrangement. An asymptotically optimal solution for SCP [6] is used as a comparison point in the evaluation section.

This work does not deal with other aspects of rearrangement, such as arm motion [49, 50, 51, 52] or grasp planning [53, 54]. Non-prehensile actions, such as pushing, are also not considered [55, 56]. Similar combinatorial issues to the ones studied here are also studied by integrated task and motion planners, for most of which there are no optimality guarantees [57, 58, 32, 33, 59, 60]. Recent work on asymptotically optimal task planning is at this point prohibitively expensive for practical use [61].

III. PROBLEM STATEMENT

This section formally defines the considered challenges.

A. Tabletop Object Rearrangement with Overhand Grasps

Consider a workspace \mathcal{W} with static obstacles and a set of n movable objects $\mathcal{O} = \{o_1, \dots, o_n\}$. For $o_i \in \mathcal{O}$, \mathcal{C}_i denotes its configuration space. Then, $\mathcal{F}_i \subseteq \mathcal{C}_i$ is the set of collision-free configurations of o_i with respect to the static obstacles in \mathcal{W} . An *arrangement* $R = \{r_1, \dots, r_n\}$ for the objects \mathcal{O} specifies the configurations $r_i \in \mathcal{C}_i$ for each object o_i . A feasible arrangement is one satisfying:

- 1) $\forall r_i \in R, r_i \in \mathcal{F}_i$;
- 2) $\forall r_i, r_j \in R$, if $i \neq j$, then objects o_i and o_j are not in collision when placed at r_i and r_j , respectively.

This work focuses on bounded planar workspaces: $\mathcal{W} \in \mathbb{R}^2$. The setting is frequently referred to as the *tabletop* case, in which the vertical projections of the objects on the tabletop do not intersect. This work assumes that the manipulator is able to employ *overhand* grasps, where an object can be transferred after being lifted above all other objects. In particular, a pick-and-place operation of the manipulator involves four steps:

- a. bringing the end-effector above the object,
- b. grasping and lifting the object,
- c. transfer of the grasped object horizontally to its target (horizontal) location, and
- d. a downward motion prior to releasing the object.

This sequence is defined as a *manipulation action*.

The manipulator is initially at a rest position s_M prior to executing any pick-and-place actions and transitions to a rest position g_M at the conclusion of the rearrangement task. A *rest position* is a safe arm configuration, where there is no collision with objects.

The illustrations that appear throughout the paper assume objects with identical geometry. Nevertheless, the results derived in this paper are not dependent on this assumption, i.e., objects need only be cylindrical in diff. geometry terms.¹

Given the above assumptions, the problem studied in the paper can be summarized as:

Problem 1. Tabletop Object Rearrangement with Overhand grasps (TORO). Given feasible start and goal arrangements R_S, R_G for objects $\mathcal{O} = \{o_1, \dots, o_n\}$ on a tabletop, determine a sequence of collision-free pick-and-place actions with overhand grasps $\mathcal{A} = (a^1, a^2, \dots)$ that transfer \mathcal{O} from R_S to R_G .

A rearrangement problem is said to be *labeled* if objects are unique and not interchangeable. Otherwise, the problem is *unlabeled*. If for two arbitrary arrangements $s \in R_S$ and $g \in R_G$, the objects placed in s and g are not in collision, then the problem is said to have *no overlaps*. Otherwise, the problem is said to have *overlaps*.

This paper primarily focuses on the labeled TORO case and identifies an important subcase:

- TORO with NO overlaps (TORO-NO)

Remark 1. The partition of Problem 1 into the general TORO case and the subcase of TORO-NO is not arbitrary. TORO is structurally richer and harder from a computational perspective. Both version of the problem can be extended to the unlabeled and partially labeled variants. This paper does not treat the labeled and unlabeled variants as separate cases but will briefly discuss differences that arise due to formulation when appropriate.

B. Optimization Criteria

Recall that a manipulation action a^i has four components: an initial move, a grasp, a transport phase, and a release. Since grasping is frequently the source of difficulty in object manipulation tasks, it is assumed in the paper that grasps and releases induce the most cost in manipulation actions. The other source of cost can be attributed to the length of the manipulator's path. This part of the cost is captured through the Euclidean distance traveled by the end effector between grasps and releases. For a manipulation action a^i , the incurred cost is

$$c_{a^i} = c_m d_e^i + c_g + c_m d_l^i + c_r, \quad (1)$$

where c_m , c_g , c_r are costs associated with moving the manipulator, a single grasp, and a single release, respectively. d_e^i and d_l^i are the straight line distances traveled by the end effector in the first (object-free) and third (carrying an object) stages of a manipulation action, respectively.

The total cost associated with solving a TORO instance is then captured by

$$c_T = \sum_{i=1}^{|\mathcal{A}|} c_{a^i} = |\mathcal{A}|(c_g + c_r) + c_m \left(\sum_{i=1}^{|\mathcal{A}|} (d_e^i + d_l^i) + d_f \right), \quad (2)$$

where d_f is the distance between the location of the last release of the end effector and its rest position g_M . Of the two additive terms in (2), note that the first term dominates the second. Because the absolute value of c_g , c_r , and c_m are different for different systems, the assignment of their absolute values is left to practitioners. The focus of this paper is the analysis and minimization of the two additive terms in (2).

C. Object Buffer Locations

The resolution of TORO (Section V) may require the temporary placement of some object(s) at intermediate locations outside those in $R_S \cup R_G$. When this occurs, external buffer locations may be used as temporary locations for object placement. More formally, there exists a set of configurations $B = \{b_1, b_2, \dots\}$, called *buffers*, which are available to the manipulator and do not overlap with object placements in R_S or R_G .

Remark 2. This work, which focuses on the combinatorial aspects of multi-object manipulation and rearrangement, utilizes exclusively buffers that are not on the tabletop. It is understood that the number of external buffers *may* be reduced by attempting to first search for potential buffers within the tabletop. Nevertheless, there are scenarios where the use of external buffers may be necessary.

¹From differential geometry, a cylinder is defined as any ruled surface spanned by a one-parameter family of parallel lines.

IV. TORO WITH NO OVERLAPS (TORO-NO)

When there is no overlap between any pair of start and goal configurations, an object can be transferred directly from its start configuration to its goal configuration. A direct implication is that an optimal sequence of manipulation actions contains exactly $|\mathcal{A}| = |\mathcal{O}| = n$ grasps and the same number of releases. Note that a minimum of n grasps and releases are necessary. This also implies that no buffer is required since using buffers will incur additional grasp and release costs. Therefore, for TORO-NO, (2) becomes

$$c_T = n(c_g + c_r) + c_m \left(\sum_{i=1}^n (d_e^i + d_i^g) + d_f \right), \quad (3)$$

i.e., only the distance traveled by the end effector affects the cost. The problem instance that minimizes (3) is referred to as Cost-optimal TORO-NO. The following theorem provides a hardness result for Cost-optimal TORO-NO.

Theorem IV.1. *Cost-optimal TORO-NO is NP-hard.*

Proof: Reduction from Euclidean-TSP [1]. Let p_0, p_1, \dots, p_n be an arbitrary set of $n + 1$ points in 2D. The set of points induces an Euclidean-TSP. Let d_{ij} denote the Euclidean distance between p_i and p_j for $0 \leq i, j \leq n$. In the formulation given in [1], it is assumed that d_{ij} are integers, which is equivalent to assuming the distances are rational numbers. To reduce the stated TSP problem to a cost-optimal TORO-NO problem, pick some positive $\varepsilon \ll 1/(4n)$. Let p_0 be the rest position of the manipulator in an object rearrangement problem. For each p_i , $1 \leq i \leq n$, split p_i into a pair of start and goal configurations (s_i, g_i) such that (i) $p_i = \frac{s_i + g_i}{2}$, (ii) $s_{i2} = g_{i2}$, and (iii) $s_{i1} + \varepsilon = g_{i1}$. An illustration of the reduction is provided in Fig. 3. The reduced TORO-NO instance is fully defined by p_0 , $R_S = \{s_1, \dots, s_n\}$ and $R_G = \{g_1, \dots, g_n\}$. A cost-optimal (as defined by (3)) solution to this TORO-NO problem induces a (closed) path starting from p_0 , going through each s_i and g_i exactly once, and ending at p_0 . Moreover, each g_i is visited immediately after the corresponding s_i is visited. Based on this path, the manipulator moves to a start location to pick up an object, drop the object at the corresponding goal configuration, and then move to the next object until all objects are rearranged. Denote the loop path as P and let its total length be D .




Fig. 3. Reduction from Euclidean-TSP to cost-optimal TORO-NO

Assume that the Euclidean-TSP has an optimal solution path P_{opt} with a total distance of D_{opt} (an integer). Then P from solving the cost-optimal TORO-NO yields such an optimal path for the TSP. To show this, from P , simply contract the edges $s_i g_i$ for all $1 \leq i \leq n$. This clearly yields a solution to the Euclidean-TSP; let the resulting path be P' with total length

D' . As edges are contracted along P , by the triangle inequality, $D' \leq D$. It remains to show that $D' = D_{opt}$. Suppose this is not the case, then $D' \geq D_{opt} + 1$. However, if this is the case, a solution to the TORO-NO can be constructed by splitting p_i into s_i and g_i along P_{opt} . It is straightforward to establish that the total distance of this TORO-NO path is bounded by $D_{opt} + n\varepsilon < D_{opt} + n * 1/(4n) = D_{opt} + 1/4 < D_{opt} + 1 \leq D' \leq D$. Since this is a contradiction, $D' = D_{opt}$. ■

Remark 3. Note that an NP-hardness proof of a similar problem can be found in [62], as is mentioned in [6]. Nevertheless, the problem is stated for a tree and is non-Euclidean. Furthermore, it is straightforward to show that the decision version of the cost-optimal TORO-NO problem is NP-complete; this non-essential detail is omitted.

Remark 4. Interestingly, TORO-NO also reduces to TSP with very little effort. Because highly efficient TSP solvers are available, the reduction route provides an effective approach for solving TORO-NO. This is not always a feature of NP-hardness reductions. The straightforward algorithm for the computation is outlined in Alg. 1. The inputs to the algorithm are the rest positions of the manipulator and the start and goal configurations of objects. The output is the solution for TORO-NO, represented as a sequence of manipulation actions \mathcal{A} , which has completeness and optimality guarantees.

Algorithm 1: TORONOTSP

Input: Configurations s_M, g_M , Arrangements R_S, R_G .
Output: A sequence of manipulation actions \mathcal{A} .
1 $G_{NO} \leftarrow \text{CONSTRUCTTSPGRAPH}(R_S, R_G, s_M, s_G)$
2 $S_{raw} \leftarrow \text{SOLVETSP}(G_{NO})$
3 $\mathcal{A} \leftarrow \text{RETRIEVEACTIONS}(S_{raw})$
4 **return** \mathcal{A}

At Line 1 of Alg. 1, a graph $G_{NO}(V_{NO}, E_{NO})$ is generated as the input to the TSP problem. The graph is constructed from the TORO-NO instance as follows. A vertex is created for each element of R_S and R_G . Then, a complete bipartite graph is created between these two sets of vertices. A set of vertices $U = \{u_1, \dots, u_{|R_S|}\}$ is then inserted into edges $s_i g_i$ for $1 \leq i \leq |R_S|$. Afterward, s_M (resp., g_M) is added as a vertex and is connected to s_i (resp., g_i) for $1 \leq i \leq |R_S|$. Finally, a vertex u_0 is added and connected to both s_M and g_M . See Fig. 4 for the straightforward example for $|R_S| = 2$.




Fig. 4. An example of G_{NO} for 2 objects. The nodes s_M and g_M denote the initial and final rest position of the manipulator end effector.

Let $w(a, b)$ denote the weight of an edge $(a, b) \in E_{NO}$. For all $1 \leq i, j \leq n, i \neq j$ ($\text{dist}(x, y)$ denotes the Euclidean

distance between x and y in 2D):

$$\begin{aligned} w(s_M, u_0) &= w(g_M, u_0) = 0, w(s_M, s_i) = \text{dist}(s_M, s_i), \\ w(g_M, g_i) &= \text{dist}(g_M, g_i), w(s_i, u_i) = w(u_i, g_i) = 0, \\ w(s_i, g_j) &= \text{dist}(s_i, g_j). \end{aligned}$$

With the construction, a TSP tour through G_{NO} must use $s_M u_0 g_M$ and all $s_i u_i g_i$ for all $1 \leq i \leq |R_S|$. To form a complete tour, exactly $(|R_S| - 1)$ edges of the form $g_i s_j$, where $i \neq j$ must be used. At Line 2, the TSP is solved (Concorde TSP solver [63] is used). This gives a minimum weight solution S_{raw} , which is a cycle containing all $v \in V_{NO}$. The manipulation actions can then be retrieved (Line 3).

An alternative solution to TORO-NO could employ the asymptotically optimal, *SPLICE* algorithm, introduced in [6].

The scenario where objects are unlabeled is a special case of TORO-NO which has significance in real-world applications (e.g., the pancake stacking application). This case is denoted as TORO-UNO (unlabeled, no overlap). Adapting the NP-hardness proof for the TORO-NO problem shows that cost-optimal TORO-UNO is also NP-hard. Similar to the TORO-NO case, the optimal solution only hinges on the distance traveled by the manipulator because no buffer is required and exactly n grasps and releases are needed.

Theorem IV.2. *Cost-optimal TORO-UNO is NP-hard.*

Proof: See Appendix A of [64]. ■

When solving a TORO-UNO instance, Alg. 1 may be used with a few small changes. First, a different underlying graph must be constructed. Denote the new graph as $G_{UNO}(V_{UNO}, E_{UNO})$, where $V_{UNO} = R_S \cup R_G \cup \{s_M, u_0, g_M\}$. For all $1 \leq i, j \leq n$:

$$\begin{aligned} w(s_M, u_0) &= w(g_M, u_0) = 0, w(s_M, s_i) = \text{dist}(s_M, s_i), \\ w(g_M, g_i) &= \text{dist}(g_M, g_i), w(s_i, g_j) = \text{dist}(s_i, g_j). \end{aligned}$$

All other edges are given infinite weight. An example of the updated structure of G_{UNO} for two objects is illustrated in Fig. 5.




Fig. 5. An example of G_{UNO} for 2 objects.

V. TORO WITH OVERLAP (TORO)

Unlike TORO-NO, TORO has a more sophisticated structure and may require buffers to solve. In this section, a *dependency graph* [2] is used to model the structure of TORO, which leads to a classical NP-hard problem known as the *feedback vertex set* problem [3]. The connection then leads to a complete algorithm for optimally solving TORO.

A. The Dependency Graph and NP-Hardness of TORO

Consider a *dependency digraph* $G_{dep}(V_{dep}, A_{dep})$, where $V_{dep} = \mathcal{O}$, and $(o_i, o_j) \in A_{dep}$ iff g_i and s_j overlap. Therefore, o_j must be moved away from s_j before moving o_i




Fig. 6. Illustration of the dependency graph. (a) Two objects are to be moved from s_i to g_i , $i = 1, 2$. Due to the overlap between s_1 and g_2 as well as the overlap between s_2 and g_1 , one of the objects must be temporarily moved aside. (b) The dependency graph capturing the scenario in (a).

to g_i . An example involving two objects is provided in Fig. 6. The definition of dependency graph implies the following two observations.

Observation V.1. *If the out-degree of $o_i \in V_{dep}$ is 0, then o_i can move to g_i without collision.*

Observation V.2. *If G_{dep} is not acyclic, solving TORO requires at least $n + 1$ grasps.*

The dependency graph has obvious similarities to the well known *feedback vertex set* (FVS) problem [3]. A directed FVS problem is defined as follows. Given a strongly connected directed graph $G = (V, A)$, an FVS is a set of vertices whose removal leaves G acyclic. Minimizing the cardinality of this set is NP-hard, even when the maximum in degree or out degree is no more than two [43]. As it turns out, the set of removed vertices in an FVS problem mirrors the set of objects that must be moved to temporary locations (i.e., buffers) for resolving the dependencies between the objects, which corresponds to the additional grasps (and releases) that must be performed in addition to the n required grasps for rearranging n objects. The observation establishes that cost-optimal TORO is also computationally intractable. The following lemma shows this point.

Lemma V.1. *Let the dependency graph of a TORO problem be a single strongly connected graph. Then the minimum number of additional grasps required for solving the TORO problem equals the cardinality of the minimum FVS of the dependency graph.*

Proof: Given the dependency graph, let the additional grasps and releases be n_x and the minimum FVS have a cardinality of n_{fvs} , it remains to show that $n_x = n_{fvs}$. First, if fewer than n_{fvs} objects are removed, which correspond to vertices of the dependency graph, then there remains a directed cycle. By Observation V.2, this part of the problem cannot be solved. This establishes that $n_x \geq n_{fvs}$. On the other hand, once all objects corresponding to vertices in a minimum FVS are moved to buffer locations, the dependency graph becomes acyclic. This allows the remaining objects to be rearranged. This operation can be carried out iteratively with objects whose corresponding vertices have no incoming edges. On a directed acyclic graph (DAG), there is always such a vertex. Moreover, as such a vertex is removed from a DAG, the remaining graph must still be a DAG and therefore must have either no vertex (a trivial DAG) or a vertex without incoming edges. ■

For dependency graphs with multiple strongly connected components, the required number of additional grasps and releases is simply the sum of the required number of such

actions for the individual strongly connected components.

For a fixed TORO problem, let n_{fvs} be the cardinality of the largest (minimal) FVS computed over all strongly connected components of its dependency graph. Then it is easy to see that the maximum number of required buffers is no more than n_{fvs} . The NP-hardness of cost-optimal TORO is established using the reduction from FVS problems to TORO. This is more involved than reducing TORO to FVS because the constructed TORO must correspond to an actual TORO problem.

Theorem V.1. *Cost-optimal TORO is NP-hard.*

Proof: The FVS problem on directed graphs is reduced to cost-optimal TORO. An FVS problem is fully defined by specifying an arbitrary strongly connected directed graph $G = (V, A)$ where each vertex has no more than two incoming and two outgoing edges. A typical vertex neighborhood can be represented as illustrated in Fig. 7(a). Such a neighborhood is converted to a dependency graph neighborhood of object rearrangement as follows. Each of the original vertex $v_i \in V$ becomes an object o_i which has some (s_i, g_i) pair as its start and goal configurations. For each directed arc $v_i v_j$, split it into two arcs and add an additional object o_{ij} . That is, create new arcs $o_i o_{ij}$ and $o_{ij} o_j$ for each original arc $v_i v_j$ (see Fig. 7(b)). This yields a dependency graph that is again strongly connected. Two claims will be proven:

- 1) The constructed dependency graph corresponds to an object rearrangement problem, and
- 2) The minimum number of objects that must be moved away temporarily to solve the problem is the same as the size of the minimum FVS.




Fig. 7. Converting a neighborhood of a graph for an FVS problem to parts of a dependency graph for a TORO problem.

To prove the first claim, assume without loss of generality that the objects have the same footprints on the tabletop. Furthermore, only the neighborhood of o_1 needs to be inspected because it is isolated by the newly added objects. Recall that an incoming edge to o_1 means that the start configuration o_1 blocks the goals of some other objects, in this case o_{21} and o_{31} . This can be readily realized by putting the goal configurations of o_{21} and o_{31} close to each other and have them overlap with the start configuration of o_1 . Note that the goal configurations of o_{21} and o_{31} have no other interactions. Therefore, such an arrangement is always achievable for even simple (e.g., circular or square) footprints. Similarly, for the outgoing edges from o_1 , which mean other objects block o_1 's goal, in this case o_{12} and o_{14} , place the start configurations of o_{12} and o_{14} close to each other and make both overlap with the goal configuration of o_1 . Again, the start configurations of o_{12} and o_{14} have no other interactions.

The second claim directly follows Lemma V.1. Now, given an optimal solution to the reduced TORO problem, it remains to show that the solution can be converted to a solution to the original FVS problem. The solution to the TORO problem provides a set of objects that are moved to temporary locations. This yields a minimum FVS on the dependency graph but not the original graph G . Note that if a newly created object (e.g., o_{ij}) is moved to a temporary place, either object o_i or o_j can be moved since this will achieve no less in disconnecting the dependency graph. Doing this across the dependency graph yields a minimum FVS for G . ■

Remark 5. It is possible to prove that TORO is NP-hard using a similar proof to the TORO-NO case. To make the proof for Theorem IV.1 work here, each p_i can be split into an overlapping pair of start and goal. Such a proof, however, would bury the structure of TORO, which is a more difficult problem. Unlike the Euclidean-TSP problem, which admits $(1 + \epsilon)$ -approximations and good heuristics, FVS problems are APX-hard [3, 5].

B. Algorithmic Solutions for TORO

Feasible algorithm: Once the link between a TORO buffer requirement and FVS is established, an algorithm for solving TORO becomes possible. To do this, an FVS set is found. Then the optimal rearrangement distance is computed for this FVS set. The procedure for doing this is outlined in TOROFVSSINGLE (Alg. 2). At Line 1, the dependency graph G_{dep} is constructed. At Line 3-4, an FVS is obtained for each strongly connected component (SCC) in G_{dep} . Note that if these FVSs are optimal, then the step yields the minimum number of required grasps (and releases) as: $\min |A| = n + |B|$.

The residual work is to find the solution with $n + |B|$ grasps and the shortest travel distance (Line 5). The manipulation actions are then retrieved and returned.

Algorithm 2: TOROFVSSINGLE

Input: Configurations s_M, g_M , Arrangements R_S, R_G

Output: A set of manipulation actions \mathcal{A}

- 1 $G_{dep} \leftarrow \text{CONSTRUCTDEPGRAPH}(R_S, R_G)$
 - 2 $B \leftarrow \emptyset$
 - 3 **for each** SCC **in** G_{dep} **do**
 - 4 $B \leftarrow B \cup \text{SOLVEFVS}(\text{SCC})$
 - 5 $S_{raw} \leftarrow \text{MINDIST}(s_M, g_M, R_S, R_G, G_{dep}, B)$
 - 6 $\mathcal{A} \leftarrow \text{RETRIEVEACTIONS}(S_{raw})$
 - 7 **return** \mathcal{A}
-

The paper explores two exact and three approximate methods as implementations of SOLVEFVS() (Line 4 of Alg. 2). The two exact methods are both based on integer linear programming (ILP) models, similar to those introduced in [65]. They differ in how cycle constraints are encoded: one uses a polynomial number of constraints and the other simply enumerates all possible cycles. Denote these two exact methods as **ILP-Constraint** and **ILP-Enumerate**, respectively. The details of these two exact methods are explained in Appendix B of [64]. With regards to approximate solutions, several heuristic solutions are presented:

- 1) **Maximum Simple Cycle Heuristic (MSCH).** The FVS is obtained by iteratively removing the node that appears on the most number of simple cycles in G_{dep} until no more cycles exist. The simple cycles are enumerated.
- 2) **Maximum Cycle Heuristic (MCH).** This heuristic is similar to MSCH but counts cycles differently. For each vertex $v \in V_{dep}$, it finds a cycle going through v and marks the outgoing edge from v on this cycle. The process is repeated for v until no more cycles can be found. The vertex with the largest cycle count is then removed first.
- 3) **Maximum Degree Heuristics (MDH).** This heuristic constructs an FVS through vertex deletion based on the degree of the vertex until no cycles exist.

Based on FVS, the solution minimizing travel distance can be found by MINDIST() (line 5), which is an LP modeling method inspired by [22] and described in Appendix C of [64].

Complete algorithm: Note that TOROFVSSINGLE() is a complete algorithm for solving TORO but it is not a complete algorithm for solving TORO optimally. With some additional engineering, a complete optimal TORO solver can also be constructed: under the assumption that grasping dominates the traveling costs, simply iterate through all optimal FVS sets and then compute the subsequent minimum distance. After all such solutions are obtained, the optimal among these are chosen. It turns out that doing this enumeration does not provide much gain in solution quality as the optimal distances are very similar to each other.

VI. PERFORMANCE EVALUATION

Experiments are executed on a Intel(R) Core(TM) i7-6900K CPU with 32GB RAM at 2133MHz. Concorde [63] is used for solving the TSP and Gurobi 6.5.1 [66] for ILP models.

A. TORO-NO: Minimizing the Travel Distance

To evaluate the effectiveness of TORONOTSP, random TORO-NO instances are generated in which the number of objects varies. For each choice of number of objects, 100 instances are tried and the average is taken. Although TORONOTSP works on thousands of objects (it takes less than 30 seconds for TORONOTSP to solve instances with 2500 objects), the evaluation is limited to 200 objects². Concerning running time, TORONOTSP is compared with SPLICE [6] which does not compute an exact optimal solution. As shown in Fig. 8, it takes less than a second for TORONOTSP to compute the distance optimal manipulation action set. Fig. 9 illustrates the solution quality of TORONOTSP, SPLICE, and an algorithm that picks a random feasible solution. Notice that the random feasible solution generally has poor quality. SPLICE does well as the number of objects increases, but under-performs compared to TORONOTSP. In conclusion, TORONOTSP provides the best performance on both running time and optimality for practical sized TORO-NO problems.

²State-of-the-art Delta robots have comparable abilities. For example, the Kawasaki YF03 Delta Robot is capable of performing 222 pick-and-place actions per minute (1kg objects).




Fig. 8. Running time comparison of TORONOTSP and SPLICE.




Fig. 9. Optimality of TORONOTSP, SPLICE and a random selection method.

For the unlabeled case (TORO-UNO), the same experiments are carried out. The results appear in Table I. Note that SPLICE no longer applies. The last line of the table is the optimality of random solutions, included for reference purposes. For larger cases, the TSP based method is able to solve for over 500 objects in 30 seconds.

B. TORO: Minimizing the Number of Grasps

To evaluate different FVS minimization methods, dependency graphs are generated by capping the average degree and maximum degree for a fixed object count. To evaluate the running time, the average degree is set to 2 and the maximum degree is set to 4, which creates significant dependencies. The running time comparison is given in Fig. 10 (averaged over 100 runs per data point). Although exact ILP-based methods took more time than heuristics, they can solve optimally for over 30 objects in just a few seconds, which makes them very practical.




Fig. 10. Running time of various methods for optimizing FVS.

When it comes to performance (Fig. 11), ILP-based methods have no competition. Interestingly, the simple cycle based method (MSCH) also works quite well and may be useful

Number of objects	10	50	100	200
Running time (sec)	0.04	0.58	2.43	7.30
Optimality of random solution	1.94	3.72	4.92	6.01

in place of ILP-based methods for larger problems, given that MSCH runs faster.




Fig. 11. Optimality ratio of various methods for optimizing FVS as compared with the optimal ILP-based methods.

The performance is also affected by the average degree for each node, which is directly linked to the complexity of G_{dep} . Fixating on the ILP-Constraint algorithm, experiments with an average degree of 0.5-2.5 are included (2.5 average degree yields rather constrained dependency graphs). As can be observed from Fig. 12, for up to 35 objects, an optimal FVS can be readily computed in a few seconds.




Fig. 12. The running time of ILP-Constraint under varying G_{dep} average degree. Maximum degree is capped at twice the average degree.

Finally, this section emphasizes an observation regarding the number of optimal FVS sets (Fig. 13). By disabling FVSs that are already obtained in subsequent runs, all FVSs for a given problem can be exhaustively enumerated for varying numbers of objects and average degree of G_{dep} . The number of optimal FVSs turns out to be fairly limited.




Fig. 13. The number of optimal FVS solutions in expectation.

C. TORO: Overall Performance

The running time for the entire TOROFVSSINGLE() is provided in Fig. 14. Observe that FVS computation takes almost no time in comparison to the distance minimization step. As expected, higher average degrees in G_{dep} make the computation harder.

Running TOROFVSSINGLE() together with FVS enumeration, a global optimal solution is computed for TORO under




Fig. 14. The total running time for TOROFVSSINGLE().

the assumption that the grasp/release costs dominate. Only solutions with an optimal FVS are considered. The computation time is provided in Fig. 15. The result shows that it gets costly to compute the global optimal solution as the number of objects go beyond 15 for dense setups. It is empirically observed that for the same problem instance and different optimal FVSs, the minimum distance computed by MINDIST() in Alg. 2 has less than 5% variance. This suggests that running TOROFVSSINGLE() just once should yield a solution that is very close to being the global optimum.




Fig. 15. The running time to produce a global optimal solution for TORO.

VII. CONCLUSION

This paper studies the combinatorial structure inherent in tabletop object rearrangement problems. For TORO-NO and TORO-UNO, it is shown that Euclidean-TSP can be reduced to them, establishing their NP-hardness. More importantly, TORO-NO and TORO-UNO can be reduced to TSP with little overhead, thus establishing that they have similar computational complexity and lead to an efficient solution scheme. Similarly, an equivalence was established between dependence breaking of TORO and FVS, which is APX-hard. The equivalence enables subsequent ILP-based methods for effectively and optimally solving TORO instances containing tens of objects with overlapping starts and goals.

Exploring scenarios in which objects are more tightly entangled, i.e., cases with high object density and thus little “buffer” space, remain an open problem. The methods and algorithms in this paper serve as a strong foundation for solving complex untangling and rearrangement tasks on tabletops.

ACKNOWLEDGMENTS

This work is supported by NSF awards IIS-1617744, IIS-1451737 and CCF-1330789, as well as internal support by Rutgers University. Any opinions or findings expressed in this paper do not necessarily reflect the views of the sponsors. The authors would like to thank the anonymous RSS reviewers for their constructive comments.

REFERENCES

- [1] C. H. Papadimitriou, “The Euclidean travelling salesman problem is NP-complete,” *Theoretical Computer Science*, vol. 4, no. 3, pp. 237–244, 1977.
- [2] J. van den Berg, J. Snoeyink, M. Lin, and D. Manocha, “Centralized path planning for multiple robots: Optimal decoupling into sequential plans,” in *Proc. Robotics: Science and Systems*, 2009.
- [3] R. M. Karp, “Reducibility among combinatorial problems,” in *Complexity of computer computations*. Springer, 1972, pp. 85–103.
- [4] M. F. E. Rohmer, S. P. N. Singh, “V-REP: a Versatile and Scalable Robot Simulation Framework,” in *Proc. IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2013.
- [5] I. Dinur and S. Safra, “On the hardness of approximating minimum vertex cover,” *Annals of Mathematics*, pp. 439–485, 2005.
- [6] K. Treleaven, M. Pavone, and E. Frazzoli, “Asymptotically optimal algorithms for one-to-one pickup and delivery problems with applications to transportation systems,” *IEEE Transactions on Automatic Control*, vol. 58, no. 9, pp. 2261–2276, 2013.
- [7] B. Aronov, M. de Berg, A. F. van den Stappen, P. Švestka, and J. Vleugels, “Motion Planning for Multiple Robots,” *Discrete and Computational Geometry*, vol. 22, no. 4, pp. 505–525, 1999.
- [8] K. Solovey, J. Yu, O. Zamir, and D. Halperin, “Motion Planning for Unlabeled Discs with Optimality Guarantees,” in *Proc. Robotics: Science and Systems*, 2015.
- [9] K. Solovey and D. Halperin, “On the hardness of unlabeled multi-robot motion planning,” in *Proc. Robotics: Science and Systems*, 2015.
- [10] J. van den Berg and M. Overmars, “Prioritized Motion Planning for Multiple Robots,” in *Proc. IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2005, pp. 2217–2222.
- [11] S. Leroy, J.-P. Laumond, and T. Siméon, “Multiple Path Coordination for Mobile Robots: A Geometric Algorithm,” in *Proc. International Joint Conferences on Artificial Intelligence*, 1999, pp. 1118–1123.
- [12] R. H. Wilson and J.-C. Latombe, “Geometric Reasoning about Mechanical Assembly,” *Artificial Intelligence*, vol. 71, no. 2, pp. 371–396, 1994.
- [13] D. Halperin, J.-C. Latombe, and R. H. Wilson, “A General Framework for Assembly Planning: the Motion Space Approach,” *Algorithmica*, vol. 26, no. 3-4, pp. 577–601, 2000.
- [14] S. Sundaram, I. Remmler, and N. M. Amato, “Disassembly Sequencing Using a Motion Planning Approach,” in *Proc. IEEE International Conference on Robotics and Automation*, Washington, D.C., May 2001, pp. 1475–1480.
- [15] D. Kornhauser, G. Miller, and P. Spirakis, “Coordinating Pebble Motion on Graphs, the Diameter of Permutation Groups, and Applications,” in *Proc. IEEE Symposium on Foundations of Computer Science*, 1984, pp. 241–250.
- [16] G. Calinescu, A. Dumitrescu, and J. Pach, “Reconfigurations in Graphs and Grids,” *SIAM Journal on Discrete Mathematics*, vol. 22, no. 1, pp. 124–138, 2008.
- [17] V. Auletta, A. Monti, D. Parente, and G. Persiano, “A Linear Time Algorithm for the Feasibility of Pebble Motion on Trees,” *Algorithmica*, vol. 23, pp. 223–245, 1999.
- [18] G. Goralý and R. Hassin, “Multi-Color Pebble Motion on Graphs,” *Algorithmica*, vol. 58, no. 3, pp. 610–636, 2010.
- [19] G. Wagner, M. Kang, and H. Choset, “Probabilistic Path Planning for Multiple Robots with Subdimensional Expansion,” in *Proc. IEEE International Conference on Robotics and Automation*, 2012.
- [20] G. Sharon, R. Stern, A. Felner, and N. R. Sturtevant, “Conflict-based Search for Optimal Multi-agent Pathfinding,” in *Artificial Intelligence*, no. 219, 2015, pp. 40–66.
- [21] J. Yu and S. M. LaValle, “Multi-agent Path Planning and Network Flow,” in *Algorithmic Foundations of Robotics X*. Springer, 2013, pp. 157–173.
- [22] —, “Optimal Multi-Robot Path Planning on Graphs: Complete Algorithms and Effective Heuristics,” *IEEE Transactions on Robotics*, vol. 32, no. 5, pp. 1163–1177, 2016.
- [23] G. Wilfong, “Motion Planning in the Presence of Movable Obstacles,” in *Annals of Mathematics and Artificial Intelligence*, 1991, pp. 131–150.
- [24] P. C. Chen and Y. K. Hwang, “Practical Path Planning Among Movable Obstacles,” in *Proc. IEEE International Conference on Robotics and Automation*, May 1991, pp. 444–449.
- [25] E. Demaine, J. O’Rourke, and M. L. Demaine, “Pushpush and push-1 are NP-hard in 2D,” in *Proc. Canadian Conference on Computational Geometry*, 2000, pp. 211–219.
- [26] D. Nieuwenhuisen, A. F. van der Stappen, and M. H. Overmars, “An Effective Framework for Path Planning amidst Movable Obstacles,” in *Proc. Workshop on the Algorithmic Foundations of Robotics*, 2006.
- [27] J. van den Berg, M. Stilman, J. J. Kuffner, M. Lin, and D. Manocha, “Path Planning Among Movable Obstacles: A Probabilistically Complete Approach,” in *Proc. Workshop on the Algorithmic Foundations of Robotics*, 2008.
- [28] M. Stilman, J. Schamburek, J. J. Kuffner, and T. Asfour, “Manipulation Planning Among Movable Obstacles,” in *Proc. IEEE International Conference on Robotics and Automation*, 2007.
- [29] O. Ben-Shahar and E. Rivlin, “Practical Pushing Planning for Rearrangement Tasks,” *IEEE Transactions*

- on *Robotics and Automation*, vol. 14, no. 4, Aug. 1998.
- [30] J. Ota, “Rearrangement Planning of Multiple Movable Objects,” in *Proc. IEEE International Conference on Robotics and Automation*, 2004.
- [31] G. Havur, G. Ozbilgin, E. Erdem, and V. Patoglu, “Geometric Rearrangement of Multiple Moveable Objects on Cluttered Surfaces: A Hybrid Reasoning Approach,” in *Proc. IEEE International Conference on Robotics and Automation*, 2014.
- [32] S. Srivastava, E. Fang, L. Riano, R. Chitnis, S. Russell, and P. Abbeel, “Combined Task and Motion Planning through an Extensible Planner-Independent Interface Layer,” in *Proc. IEEE International Conference on Robotics and Automation*, 2014.
- [33] C. R. Garrett, T. Lozano-Pérez, and L. P. Kaelbling, “FFR_{ob}: An efficient heuristic for task and motion planning,” in *Proc. Workshop on the Algorithmic Foundations of Robotics*, 2014.
- [34] A. Krontiris, R. Shome, A. Dobson, A. Kimmel, and K. E. Bekris, “Rearranging similar objects with a manipulator using pebble graphs,” in *Proc. IEEE International Conference on Humanoid Robotics*, Madrid, Spain, 2014.
- [35] A. Krontiris and K. E. Bekris, “Dealing with Difficult Instances of Object Rearrangement,” in *Proc. Robotics: Science and Systems*, Rome, Italy, Jul. 2015.
- [36] —, “Efficiently Solving General Rearrangement Tasks: A Fast Extension Primitive for an Incremental Sampling-based Planner,” in *Proc. IEEE International Conference on Robotics and Automation*, Sweden, 2016.
- [37] K. Hauser, “The Minimum Constraint Removal Problem with Three Robotics Applications,” in *International Journal of Robotics Research*, 2014.
- [38] —, “Minimum Constraint Displacement Motion Planning,” in *Proc. Robotics: Science and Systems*, 2013.
- [39] G. Berbeglia, J.-F. Cordeau, I. Gribkovskaia, and G. Laporte, “Static pickup and delivery problems: a classification scheme and survey,” *Top*, vol. 15, no. 1, pp. 1–31, 2007.
- [40] G. Berbeglia, J.-F. Cordeau, and G. Laporte, “Dynamic pickup and delivery problems,” *European Journal of Operational Research*, vol. 202, no. 1, pp. 8–15, 2010.
- [41] G. Laporte, “The vehicle routing problem: An overview of exact and approximate algorithms,” *European Journal of Operational Research*, vol. 59, no. 3, pp. 345–358, 1992.
- [42] N. Christofides and S. Eilon, “An algorithm for the vehicle-dispatching problem,” *Journal of the Operational Research Society*, vol. 20, no. 3, pp. 309–318, 1969.
- [43] M. R. Garey and D. S. Johnson, *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman, 1979.
- [44] S. Anily and R. Hassin, “The swapping problem,” *Networks*, vol. 22, no. 4, pp. 419–433, 1992.
- [45] P. Beullens, D. van Oudheusden, and L. N. van Wassenhove, “Collection and vehicle routing issues in reverse logistics,” in *Reverse Logistics*. Springer, 2004, pp. 95–134.
- [46] A. Hoff and A. Løkketangen, “Creating lasso-solutions for the traveling salesman problem with pickup and delivery by tabu search,” *Central European Journal of Operations Research*, vol. 14, no. 2, pp. 125–140, 2006.
- [47] I. Gribkovskaia, Ø. Halskau, G. Laporte, and M. Vlček, “General solutions to the single vehicle routing problem with pickups and deliveries,” *European Journal of Operational Research*, vol. 180, no. 2, pp. 568–584, 2007.
- [48] G. N. Frederickson, M. S. Hecht, and C. E. Kim, “Approximation algorithms for some routing problems,” in *17th Annual Symposium on Foundations of Computer Science (sfcs 1976)*. IEEE, 1976, pp. 216–227.
- [49] T. Siméon, J.-P. Laumond, J. Cortés, and A. Sahbani, “Manipulation Planning with Probabilistic Roadmaps,” *International Journal of Robotics Research*, no. 23, 2004.
- [50] D. Berenson, S. S. Srinivasa, and J. J. Kuffner, “Task Space Regions: A Framework for Pose-Constrained Manipulation Planning,” *International Journal of Robotics Research*, vol. 30, no. 12, pp. 1435–1460, 2012.
- [51] J. B. Cohen, S. Chitta, and M. Likhachev, “Single- and Dual-arm Motion Planning with Heuristic Search,” in *International Journal of Robotics Research*, 2013.
- [52] M. Zucker, N. Ratliff, A. Dragan, M. Pivtoraiko, M. Klingensmith, C. Dellin, J. A. Bagnell, and S. S. Srinivasa, “CHOMP: Covariant Hamiltonian Optimization for Motion Planning,” *International Journal of Robotics Research*, 2013.
- [53] M. Ciocarlie and P. Allen, “Hand Posture Subspaces for Dexterous Robotic Grasping,” in *International Journal of Robotics Research*, vol. 28, no. 7, 2009.
- [54] J. Bohg, A. Morales, T. Asfour, and D. Kragic, “Data-driven Grasp Synthesis - A Survey,” in *IEEE Transactions on Robotics*, vol. 30, no. 2, Apr. 2014.
- [55] A. Cosgun, T. Hermans, V. Emeli, and M. Stilman, “Push Planning for Object Placement on Cluttered Table Surfaces,” in *Proc. IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2011.
- [56] M. R. Dogar and S. S. Srinivasa, “A Framework for Push-Grasping in Clutter,” in *Proc. Robotics: Science and Systems*, 2011.
- [57] S. Cambon, R. Alami, and F. Grivot, “A Hybrid Approach to Intricate Motion, Manipulation, and Task Planning,” *International Journal of Robotics Research*, no. 28, 2009.
- [58] E. Plaku and G. Hager, “Sampling-based motion and symbolic action planning with geometric and differential constraints,” in *Proc. IEEE International Conference on Robotics and Automation*, 2010.
- [59] M. Gharbi, R. Lallement, and R. Alami, “Combining

- symbolic and geometric planning to synthesize human-aware plans: toward more efficient combined search,” in *Proc. IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2015, pp. 6360–6365.
- [60] N. T. Dantam, Z. K. Kingston, S. Chaudhuri, and L. E. Kavraki, “Incremental Task and Motion Planning: A Constraint-Based Approach,” in *Proc. Robotics: Science and Systems*, 2016.
- [61] W. Vega-Brown and N. Roy, “Asymptotically optimal planning under piecewise-analytic constraints,” in *Proc. Workshop on the Algorithmic Foundations of Robotics*, 2016.
- [62] G. N. Frederickson and D. J. Guan, “Nonpreemptive ensemble motion planning on a tree,” *Journal of Algorithms*, vol. 15, no. 1, pp. 29–60, 1993.
- [63] D. L. Applegate, R. E. Bixby, V. Chvatal, and W. J. Cook, *The Traveling Salesman Problem: A Computational Study*. Princeton, NJ, USA: Princeton University Press, 2007.
- [64] S. D. Han, N. M. Stiffler, A. Krontiris, K. E. Bekris, and J. Yu, “High-Quality Tabletop Rearrangement with Overhand Grasps: Hardness Results and Fast Methods,” in *Proc. Robotics: Science and Systems*, Boston, Massachusetts, U.S.A., Jul. 2017, supplementary material available at <https://arxiv.org/abs/1705.09180>.
- [65] A. Baharev, H. Schichl, and A. Neumaier, “An exact method for the minimum feedback arc set problem,” 2015 (submitted).
- [66] I. Gurobi Optimization, “Gurobi Optimizer Reference Manual,” 2016. [Online]. Available: <http://www.gurobi.com>