

Fast High-Quality Tabletop Rearrangement in Bounded Workspace

Kai Gao¹ Darren Lau² Baichuan Huang¹ Kostas E. Bekris¹ Jingjin Yu¹

Abstract—In this paper, we examine the problem of rearranging many objects on a tabletop in a cluttered setting using overhand grasps. Efficient solutions for the problem, which capture a common task that we solve on a daily basis, are essential in enabling truly intelligent robotic manipulation. In a given instance, objects may need to be placed at temporary positions (“buffers”) to complete the rearrangement, but allocating these buffer locations can be highly challenging in a cluttered environment. To tackle the challenge, a two-step baseline planner is first developed, which generates a primitive plan based on inherent combinatorial constraints induced by start and goal poses of the objects and then selects buffer locations assisted by the primitive plan. We then employ the “lazy” planner in a tree search framework which is further sped up by adapting a novel preprocessing routine. Simulation experiments show our methods can quickly generate high-quality solutions and are more robust in solving large-scale instances than existing state-of-the-art approaches.

source: github.com/arc-1/TRLB

I. INTRODUCTION

We study the problem of rearranging many objects on a tabletop in a cluttered environment using overhand grasps, where the robot, in a pick-n-place operation, may grasp and lift an object, move it around freely, and then place it at a collision free pose. This is known as Tabletop Object Rearrangement with Overhand Grasps (TORO), which is NP-hard to optimally solve [1]. As a task that humans face all the time, solving TORO autonomously and intelligently is essential in enabling smart robots, at home or in factories.

TORO is difficult to optimally solve due to the complex constraints induced by the start and goal poses of the involved objects. The combinatorial constraints are generated by (potential) object collisions. Fig. 1 shows a robot arm rearranging letters to form “ICRA 2022” using a vacuum gripper while having access only to a bounded rectangular area. In this example, when letter “I” is moved to the goal position, it will collide with “R” in its current pose, which results in the *inherent* constraint that “R” must move away before “I” moves into the goal pose. Such inherent constraints require some objects to be temporarily displaced at “buffer” locations, which may induce further *acquired* constraints as there can be many choices of potential buffer locations. Acquired constraints cannot always be avoided, especially when the environment is cluttered. Consider again Fig. 1 for an example. On one hand, if the character “0” is




Fig. 1. A robot arm rearranges word patterns with overhand grasps, minimizing the number of pick-n-place actions. The letters cannot go out of the bounding box and cannot overlap on the tabletop.

transferred to the lower right corner of the workspace, then it needs to be moved before a character “2” is moved to the lower right side. On the other hand, if “0” is moved to the upper left corner of the workspace, then it needs to be transferred before letter “I” can be moved to the goal. While the inherent constraints are fully determined by the start and goal configurations, acquired constraints depend on the choices made by an algorithm, which significantly increases the size of the search space and the cost of computations.

In tackling TORO, our key contribution is the insight that solutions for untangling inherent constraints [1]–[3] can assist in resolving the entire problem, and properly using them to develop a fast high-quality solver for dense TORO problems. We propose *Tabletop Rearrangement with Lazy Buffers* (TRLB), an effective framework for solving TORO in dense/cluttered environments. TRLB first computes a primitive plan only for the inherent constraints and then “lazily” selects acquired constraints without the need for a complete feasibility check. The lazy strategy delays buffer allocation and effectively reduces the number of collision checks. Simulation experiments show that TRLB computes high-quality plans for TORO efficiently and is more robust than existing approaches in challenging and practical problems. Together with the preprocessing routine, TRLB can solve dense 60-object TORO instances in an average of 0.3 seconds, allowing its use in real-time applications. In comparison with state-of-the-art approaches, TRLB computes better solutions two magnitudes faster.

II. RELATED WORK

Rearrangement Approaches: Object rearrangement is a topic of interest within Task and Motion Planning (TAMP). Typical definitions in this domain [2], [4]–[6] involve arranging multiple objects to specific goals. Certain variations such as NAMO (Navigation among Movable Obstacles) [7]–[9], and retrieval problems [10], [11], focus on clearing out a path for a target object while identifying objects that need to be relocated. Rearrangement may be approached either via simple but inaccurate non-prehensile actions, e.g., pushes [4], [12], [13], or more purposeful prehensile actions,

¹Department of Computer Science, Rutgers University, NJ, USA. Email: {kg627, bh417, kb572, jy512}@cs.rutgers.edu.

²Department of Computer Science, Cornell University, NY, USA. Email: {dl755}@cornell.edu.

This work is supported by NSF awards IIS-1845888, CCF-1934924 and IIS-2132972.

such as grasps [6], [14]–[18]. Focusing on the combinatorial challenges, some planners use external space for temporary object storage [1]–[3], [11], while others exploit problem linearity to simplify search [7], [9], [19], [20]. By linking rearrangement to established graph-based problems, efficient algorithms have been obtained for various tasks and objectives [1]–[3]. This paper uses a plan generated given access to external buffer locations as a “primitive plan”, which then guides buffer allocation inside the workspace.

Dependency Graph: We represent the combinatorial constraints of such problems via a dependency graph, which was first applied to general multi-body planning problems [21] and then rearrangement [14], [15]. Choosing different manipulation sequences gives rise to multiple dependency graphs for the same problem instance, which limits scalability. Prior work [1] has applied full dependency graphs to address TORO, showing that the challenge embeds the NP-hard Feedback Vertex Set (FVS) and the Traveling Salesperson (TSP) problems. More recently, some of the authors [2] examined an optimization objective, *running buffers*, which is the size of the external space needed for rearrangement, and also examined an unlabeled setting. Similar structures are also used in other problems, such as packing [22]. Deep neural networks have been also applied to detect the embedded dependency graph of objects in a cluttered space to determine the ordering of object retrieval [10].

Buffer Identification: In rearrangement, collision-free locations are needed for obstacle displacements. Some previous efforts predetermine buffer candidates and place obstacles to accessible candidate locations when necessary [5], [23]. Others decouple the problem into successive subproblems [5], [15]. Intuitively speaking, a valid buffer location needs to avoid other objects at their current poses. Backtracking search further constrains object displacements given paths of future manipulation actions [8], [9]. Nevertheless, these methods are computationally expensive since they deal with inherent and acquired constraints at the same time. To reduce the associated cost, a lazy strategy can be applied [24]–[26], which delays path/configuration collision checking. A similar idea is proposed here. Feasible object locations are obtained with the aid of a “rough schedule” of object manipulation actions, which is computed given only the inherent constraints. By decoupling inherent and acquired constraints, the proposed method computes high-quality solutions efficiently.

III. TABLETOP REARRANGEMENT WITH INTERNAL BUFFERS

A. Problem Statement

Consider a 2D bounded workspace $\mathcal{W} \subset \mathbb{R}^2$ containing a set of n objects $\mathcal{O} = \{o_1, \dots, o_n\}$. Each object is assumed to be an upright *generalized cylinder*. A *feasible arrangement* $\mathcal{A} = \{p_1, \dots, p_n\}$, $p_i = \{x_i, y_i, \theta_i\} \in SE(2)$ is a set of poses for objects in \mathcal{O} , such that (1) each object’s footprint is contained in \mathcal{W} , and (2) no two objects collide.

Consider (overhand) pick-n-place actions to move objects one by one. A pick-n-place action, represented as an ordered



Fig. 2. [Left] A rearrangement task involving three soda cans. [Right] The corresponding dependency graph G .

pair (o, p) , grasps an object o at its current pose and lifts it above all other objects. It then moves it horizontally, and places it at the target pose p within \mathcal{W} . An action is collision-free iff both the current and resulting arrangements are feasible. A plan P from a feasible \mathcal{A}_s to a feasible \mathcal{A}_g is a sequence of collision-free pick-n-place actions transforming \mathcal{A}_s into \mathcal{A}_g . We want to compute feasible plans that minimize the number of pick-n-place actions, which leads to increased system throughput, i.e.,

Problem 1 (TORO w/ Internal Buffers (TORI)). *Given feasible arrangements $\mathcal{A}_s = \{p_1^s, \dots, p_n^s\}$ and $\mathcal{A}_g = \{p_1^g, \dots, p_n^g\}$, find a feasible plan P sequentially moving objects from \mathcal{A}_s to \mathcal{A}_g , which minimizes the number of actions.*

B. Dependency Graph and Internal Buffers

It is not always possible to move an object to its goal pose, which may be occupied by other objects. This leads to *dependencies* between objects, i.e., when the goal pose p_i^g of o_i collides with o_j at its current pose, then o_i depends on o_j and o_j must be moved before moving o_i to its goal.

Dependencies induced by \mathcal{A}_s and \mathcal{A}_g give rise to a dependency graph G [1], [21]. Fig. 2 provides an example graph. When G is acyclic, the instance is called “monotone” and can be solved with at most n actions moving each object once from its start to its goal pose following the topological order of G . Otherwise, some object(s) must be temporarily displaced to break these cycles and complete the task. We refer to these intermediate poses as “*buffers*”, which may be external (i.e., outside \mathcal{W}) or internal (i.e., contained in \mathcal{W}). If the buffers must be internal, the problem is TORI. Otherwise, the problem is TORO with external buffers, denoted as TORE. Solving TORE only requires dealing with inherent constraints defined by G . For instance, to solve the problem in Fig. 2, we can move the Pepsi to an external buffer to break the cycle, move the Coke first and then Fanta to their goal locations and finally bring back the Pepsi into the workspace. With internal buffers, we must find a temporary location for the Pepsi in \mathcal{W} . Due to acquired constraints (as defined in Sec. I) arising from internal buffer selection, TORI, the problem we study in this work, is more challenging than TORE. Intuitively, selecting buffers inside workspace (TORI) is much more difficult and constrained than using buffers outside the workspace (TORE) to store displaced objects.

Nevertheless, we show here that plans can be efficiently derived from the minimum *total buffer* solution to TORE [1] and the minimum *running buffer* solution to TORE [2], which computes the minimum number of *concurrent external buffers* needed to solve a TORE instance. For both objectives, since TORE has been shown to be computationally intractable [1], [2] and is a special case of TORI, TORI is also NP-hard.

IV. ALGORITHMIC SOLUTIONS

This section first describes a rearrangement solver with lazy buffer allocation (Sec. IV-A), where buffer allocation is delayed after getting a “rough” schedule of object movements. To enhance scalability to larger and more cluttered instances, the TRLB framework (Sec. IV-B) recovers from buffer allocation failures. Finally, a preprocessing routine helps with further speedups (Sec. IV-C).

A. Lazy Buffer Allocation

When an object stays at a buffer, it should avoid blocking the upcoming manipulation actions of other objects. Otherwise, either the object in the buffer or the manipulating object has to yield, which increases the number of necessary actions. In other words, we need to carefully choose acquired constraints. If we know the schedule of other objects in advance, a buffer can be selected to minimize unnecessary obstructions. This observation motivates us to “lazily” allocate buffers by ignoring initially acquired constraints and solve the rearrangement problem in two steps: First, compute a *primitive plan*, which is an incomplete schedule ignoring acquired constraints; second, given the incomplete schedule as a reference, generate buffers to optimize the selection of acquired constraints.

1) *Primitive Plan*: To compute a *primitive plan*, we assume enough free space is available so that no acquired constraints will be created. This transforms the problem into a TORE problem, where each object is displaced at most once before it moves to the goal pose. Then, an object $o_i \in \mathcal{O}$ can have three *primitive* actions:

- 1) $(o_i, s \rightarrow g)$: moving from p_i^s to p_i^g ;
- 2) $(o_i, s \rightarrow b)$: moving from p_i^s to a buffer;
- 3) $(o_i, b \rightarrow g)$: moving from a buffer to p_i^g .

A primitive plan is a sequence of primitive actions; computing such a plan is similar to finding a linear vertex ordering [27], [28] of the dependency graph. We use dynamic programming based methods [2] to achieve this, which minimizes the number of total buffers or running buffers.

2) *Buffer Allocation*: Free space inside the workspace \mathcal{W} is scarce in cluttered spaces (e.g., Fig. 5) and acquired constraints must be dealt with through the careful allocation of buffers inside \mathcal{W} . We apply a greedy strategy to find feasible buffers based on a primitive plan (Algo. 1). The general idea is to incrementally add constraints on the buffers until we find feasible buffers for the whole primitive plan or terminate at a step where there are no feasible buffers for the primitive plan. In Algo. 1, $\mathcal{O}_s, \mathcal{O}_g, \mathcal{O}_b$ are the sets of objects currently at start poses, goal poses and buffers respectively.

We start with \mathcal{A}_s where all the objects are at start poses and the buffers are initialized at random poses (line 1). Each action in π indicates an object o_i that is manipulated and the action m performed (line 2). If o_i is moved to a buffer (line 3), then we add it into \mathcal{O}_b (line 4). The current poses of other objects in $\mathcal{O}_s \cup \mathcal{O}_g$ are seen as fixed obstacles for o_i (line 5). If o_i is leaving the buffer (line 6), then other objects in \mathcal{O}_b should avoid the goal pose p_i^g of o_i (line 7).

Algorithm 1: Buffer Allocation

Input : π : a primitive plan; $\mathcal{A}_s = \{p_1^s, \dots, p_n^s\}$: start arrangement; $\mathcal{A}_g = \{p_1^g, \dots, p_n^g\}$: goal arrangement
Output: B : buffers; TerminatingStep: the action step where buffer generation fails, ∞ if Success.

```

1  $\mathcal{O}_s \leftarrow \mathcal{O}; \mathcal{O}_g, \mathcal{O}_b \leftarrow \emptyset; B \leftarrow \text{RandomPoses}(\mathcal{O})$ 
2 for  $(o_i, m) \in \pi$  do
3   if  $m$  is  $s \rightarrow b$  then
4      $\mathcal{O}_b.\text{add}(o_i)$ 
5      $\text{Constraints}[o_i] \leftarrow \text{GetPoses}(\mathcal{O}_s \cup \mathcal{O}_g - \{o_i\})$ 
6   else if  $m$  is  $b \rightarrow g$  then
7     for  $o \in \mathcal{O}_b \setminus \{o_i\}$  do  $\text{Constraints}[o].\text{add}(p_i^g)$ ;
8   else
9     for  $o \in \mathcal{O}_b$  do  $\text{Constraints}[o].\text{add}(p_i^g)$ ;
10     $\text{Success}, B' \leftarrow \text{BufferGeneration}(\mathcal{O}_b, \text{Constraints}, B)$ 
11    if Success then
12       $B \leftarrow B'$ 
13       $\mathcal{O}_s, \mathcal{O}_g, \mathcal{O}_b \leftarrow \text{UpdateState}(\mathcal{O}_s, \mathcal{O}_g, \mathcal{O}_b)$ 
14    else return  $B, \pi.\text{index}(\text{action})$ ;
15 return  $B, \infty$ 
```

If o_i is moving directly from p_i^s to p_i^g (line 8, the “else” corresponds to m being $s \rightarrow g$, e.g., directly go from start to goal), then all buffers for objects in the current \mathcal{O}_b need to avoid p_i^g (line 9). After setting up acquired constraints, we generate new buffers for objects in \mathcal{O}_b to satisfy these constraints by either sampling or solving an optimization problem (line 10). Old buffers in B satisfying new constraints will be directly adopted. If feasible buffers are found (line 11), then buffers and object states will be updated (line 12-13). Otherwise, we return the feasible buffers computed and record the terminating step of the algorithm (line 14). In the case of a failure, we provide a TRLB framework to recover, which is presented in Sec. IV-B.

Fig. 3 illustrates the buffer allocation process via an example. The green, cyan, and transparent discs represent the current poses, goal poses and allocated buffers respectively. When we move o_1 to a buffer B_1 (Fig. 3(b)), it only needs to avoid collision with p_2^s and p_3^s . But as we move o_3 to a buffer, B_1 needs to avoid o_3 ’s buffer B_3 as well. To satisfy the added constraint, B_1 will be reallocated. Since the new buffers B_1 and B_3 (Fig. 3(c)) satisfy the constraints added in the following steps, they need not to be relocated. Note that the buffer originally selected for o_1 but then replaced will not appear in the resulting plan, i.e., o_1 will move directly to the new buffer (Fig. 3(c)-(f)). Algo. 1 works with one strongly connected component of the dependency graph at a time, treating objects in other components as fixed obstacles.

Once the feasible buffers are found, all the primitive actions can be transformed into feasible pick-n-place actions inside the workspace. And therefore, the primitive plan can be transformed into a rearrangement plan moving objects from \mathcal{A}_s to \mathcal{A}_g . The function *BufferGeneration* is implemented by either sampling or solving an optimization problem, both of which are discussed below.

a) *Sampling*: Given the object poses that buffers need to avoid so far, feasible buffers can be generated by sampling poses inside the free space. When objects stay in buffers at

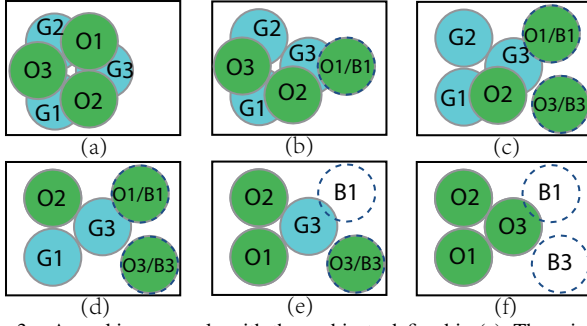


Fig. 3. A working example with three objects defined in (a). The primitive plan is $[(o_1, s \rightarrow b), (o_3, s \rightarrow b), (o_2, s \rightarrow g), (o_1, b \rightarrow g), (o_3, b \rightarrow g)]$. Figures (b)-(f) show the steps of Alg. 1 after each action. The transparent discs with dashed line boundaries (B_i) represent the buffers satisfying constraints up to each step. For each object o_i , the green (O_i) and cyan (G_i) discs represent the current and goal poses respectively.

the same time, we sample buffers one by one and previously sampled buffers will be seen as obstacles for latter ones.

b) Optimization: For cylindrical objects o_i at (x_i, y_i) with radius r_i , and o_j at (x_j, y_j) with radius r_j , they are collision-free when $(x_i - x_j)^2 + (y_i - y_j)^2 \geq (r_i + r_j)^2$ holds. Therefore, the collision constraints obtained in Algo. 1 can be transformed into quadratic inequalities. By further restricting the range of buffer centroids to assure they are in the workspace, the buffer allocation problem can be transformed into a quadratic programming problem with no objective function. For objects with general shapes, collision avoidance cannot be presented by inequalities of object centroids. We can construct the optimization problem with ϕ functions of the objects [29] and solve the problem with gradients.

B. Tabletop Rearrangement with Lazy Buffers (TRLB)

The TRLB framework builds on the insight that a new TORI instance is generated when lazy buffer allocation fails. The new instance has the same goal \mathcal{A}_g as the original one but some progress has been made in solving the TORI task. There are two straightforward implementations of TRLB: forward search and bidirectional search. In the first case, by accepting partial solutions, a rearrangement plan can be computed by developing a search tree T rooted at \mathcal{A}_s . In the search tree T , nodes are feasible arrangements and edges are partial plans containing a sequence of collision-free actions. When buffer allocation fails, we add the resulting arrangement into the tree and resume the rearrangement task from a random node in T . This randomness and the randomness in primitive plan computation and buffer allocation allows TRLB to recover from failures.

In bidirectional search, two search trees rooted at \mathcal{A}_s and \mathcal{A}_g are developed. This more involved procedure is shown in Algo. 2, which computes two search trees that connect \mathcal{A}_s and \mathcal{A}_g . In line 1, the trees are initialized. For each iteration, we first rearrange between a random node \mathcal{A}_{rand} on T_1 to the root node of T_2 (line 3-5). The function *LazyBufferAllocation* refers to the overall algorithm developed in Sec. IV-A. A found path yields a feasible plan for TORI (line 6). Otherwise, we rearrange between the new arrangement \mathcal{A}_{new1} and its nearest neighbor in T_2 (line

Algorithm 2: TRLB with Bidirectional Search

Input : $\mathcal{A}_s, \mathcal{A}_g, max.time$
Output: Search trees: T_1, T_2

```

1  $T_1.root, T_2.root \leftarrow \mathcal{A}_s, \mathcal{A}_g$ 
2 while not exceeding max.time do
3    $\mathcal{A}_{rand} \leftarrow \text{RandomNode}(T_1)$ 
4    $\mathcal{A}_{new1} \leftarrow \text{LazyBufferAllocation}(\mathcal{A}_{rand}, T_2.root)$ 
5    $T_1.add(\mathcal{A}_{new1})$ 
6   if  $\mathcal{A}_{new1}$  is  $T_2.root$  then return  $T_1, T_2$ ;
7    $\mathcal{A}_{near} \leftarrow \text{NearestNode}(\mathcal{A}_{new1}, T_2)$ 
8    $\mathcal{A}_{new2} \leftarrow \text{LazyBufferAllocation}(\mathcal{A}_{near}, \mathcal{A}_{new1})$ 
9    $T_2.add(\mathcal{A}_{new2})$ 
10  if  $\mathcal{A}_{new2}$  is  $\mathcal{A}_{new1}$  then return  $T_1, T_2$ ;
11   $T_1, T_2 \leftarrow T_2, T_1$ 

```

7-9). This paper defines the distance between arrangement nodes as the number of objects at different poses. If a path is found, then we find a feasible rearrangement plan for TORI (line 10). Otherwise, we switch the trees and attempt rearrangement from the opposite side (line 11).

C. Preprocessing

In dense environments, allocating buffers is hard, motivating minimizing the number of running buffers [2], which is generally low even in high density settings if we treat objects as *unlabeled*. Based on this, for each component of the dependency graph that is not a single vertex/cycle, we reduce the running buffer size to at most 1 by first solving an *unlabeled* instance [2]. After preprocessing, we obtain a TORI requiring at most one running buffer. Fig. 4 shows an example of preprocessing. o_1, o_2 and o_3 form a complete graph, where at least two objects need to be placed at buffers simultaneously. We conduct preprocessing of the three-vertex component by moving o_2 to a buffer position, o_1 to p_3^g and o_3 to p_2^g . o_2 will not move to p_1^g since it does not occupy other goal poses. The preprocessing step needs one buffer and the resulting rearrangement problem is monotone.




Fig. 4. A four-object example of preprocessing. The green and cyan discs represent current and goal arrangements respectively. Before preprocessing (left), two buffers need to be allocated synchronously. After preprocessing (right), the problem becomes monotone.

V. EXPERIMENTS

We implemented the algorithms of the TRLB framework in Python. Simulated experiments use environments with different density levels ρ , defined as the proportion of the tabletop surface occupied by objects, i.e., $\rho := (\sum_{o_i \in \mathcal{O}} S_{o_i}) / S_W$, where S_{o_i} is the base area of o_i and S_W is the area of \mathcal{W} .




Fig. 5. [left] 10 cylinders with $\rho = 0.5$, [right] 10 cuboids with $\rho = 0.4$.

Fig. 5 shows a dense 10-cylinder arrangement with $\rho = 0.5$, and a dense 10-cube arrangement with $\rho = 0.4$. The experiments are executed on an Intel® Xeon® CPU at 3.00GHz. Each data point is the average of 30 cases except for unfinished trials, given a time limit of 300 sec. per case.

A. Ablation Study for Cylindrical Objects

We first present experiments with cylindrical objects to compare lazy buffer generation algorithms given different options, including: (1) Primitive plan computation: running buffer minimization (RBM), total buffer minimization (TBM), random order (RO); (2) Buffer allocation methods: optimization (OPT), sampling (SP); (3) High level planners: one-shot (OS), forward search tree (ST), bidirectional search tree (BST); and (4) With or without preprocessing (PP). Here, the one-shot (OS) planner is using primitive plans and buffer allocation (Sec. IV-A) without tree search (Sec. IV-B). In OS, we attempt to compute a feasible rearrangement plan up to $30|\mathcal{O}|$ times before announcing a failure. Notice that at least $|\mathcal{O}|$ actions are required for solving any instance.

A full TRLB algorithm is a combination of components, e.g., RBM-SP-BST stands for using the primitive plans that minimize running buffer size, performing buffer allocation by sampling, maintaining a bidirectional search tree, and doing so without preprocessing.

For evaluation, we first compare the primitive plan computation options, using sampling-based buffer allocation, bidirectional tree search and no preprocessing. TBM and RBM plans are computed using dynamic programming solvers [2]. The results are shown in Fig.6. Even though plans generated by TBM-SP-BST are slightly shorter than RBM-SP-BST, TBM-SP-BST is less scalable as either the density level or the number of objects in the workspace increases. Compared to RBM plans, individual RO plans can be generated almost instantaneously but there is little benefit in computation for the overall algorithm. The results indicate that RBM should be used for primitive plan computation as it results in efficient and high-quality solutions.




Fig. 6. Comparison of primitive planners with 10-40 cylinders and density levels $\rho = 0.3$ (top), 0.5 (bottom) (left: computation time in seconds; middle: success rate; right: number of actions as multiples of $|\mathcal{O}|$).

In Fig. 7, buffer allocation methods are compared using the RBM primitive planner and the OS high-level planner. Optimization-based allocation guarantees completeness and generates high-quality plans but it is computationally expensive. When $\rho = 0.5$, the success rate tends to be low in instances with a small number of objects. That is because for the given density level, the smaller the number of

objects, the larger the object size relative to the environment, and the smaller the configuration space size relative to the environment. Thus, precisely allocating buffer locations with OPT is helpful in these cases.




Fig. 7. Comparison of buffer allocation methods with 10-40 cylinders at density levels $\rho = 0.3$ (top), 0.5 (bottom) (left: computation time in seconds; middle: success rate; right: number of actions as multiples of $|\mathcal{O}|$).

The effectiveness of the high-level planners and preprocessing are shown in Fig. 8, which suggests that ST, BST and preprocessing are all effective in increasing success rate in dense environments. In addition, preprocessing significantly speeds up computation in large scale dense cases at the price of extra actions to execute preprocessing. By simplifying the dependency graph with preprocessing, less time is needed to compute a primitive plan.




Fig. 8. Comparison among high-level frameworks and OS with preprocessing. There are 10-40 cylinders in the workspace at density levels $\rho = 0.3$ (top), 0.5 (bottom) (left: computation time in seconds; middle: success rate; right: number of actions as multiples of $|\mathcal{O}|$).

The robustness of ST and BST are further evaluated with “dense-small” instances where a few objects are packed densely (Fig. 9). The bidirectional search tree has a higher success rate in these cases, especially in 5-object instances.

In addition to the above evaluations, we also tried integrating the preprocessing into the BST framework (RBM-SP-BST-PP), which speeds up the computation: for 60-object instances with $\rho = 0.5$, only 63% of them can be solved by RBM-SP-BST in 300 seconds. All of them, however, can be solved together with the preprocessing and the solution time averaged 0.29 seconds. Similarly to the results in Fig. 8, preprocessing makes the solution plan much longer than necessary (needs around 30% more actions than RBM-SP-BST). Based on the analysis on computation time, success rate, and solution quality, RBM-SP-BST is the best overall combination, and preprocessing significantly speeds up the solver with a reduction of solution quality.




Fig. 9. Comparison between ST and BST frameworks with “dense-small” instances where 5-8 cylinders packed in the environment and density level $\rho = 0.5$. [left] An example with 5 cylinders, [right] success rate of methods.

B. Comparison with Alternatives for Cylindrical Objects

We compare the proposed method RBM-SP-BST with BiRRT(fmRS) [15] and a MCTS planner [16], which, to the best of our knowledge, are state-of-the-art planners for TORI. The MCTS planner is a C++ solver, while the other two methods are implemented in Python. Besides success rate, solution quality, and computation time, we also compare the number of collision checks which are time-consuming in most planning tasks. In Fig. 10, we compare the methods in large scale problems with $\rho = 0.3$. The success rate is 100% for all. Our method, RBM-SP-BST, avoids repeated collision checks due to the use of the dependency graph. BiRRT(fmRS), which only uses dependency graphs locally, spends a lot of time and conducts a lot of collision checks to generate random arrangements. MCTS generates solutions with similar optimality but does so also with a lot of collision checking, which slows down computation. We note that a value of 1 in the right figure (number of actions) is the minimum possible, so both RBM-SP-BST and MCTS compute high-quality solutions, which RBM-SP-BST does slightly better. To sum up, in sparse large scale instances, RBM-SP-BST is two magnitudes faster and conducts much fewer collision checks than the alternatives.




Fig. 10. Comparison of algorithms with 20-100 cylinders at density level $\rho = 0.3$ (left: computation time in seconds; middle: number of collision checks; right: number of actions as multiples of $|\mathcal{O}|$).

Next, we compare the methods in “dense-small” instances (Fig. 11). Here, RBM-SP-BST is the only method that maintains high success rate in these difficult cases.




Fig. 11. Comparison of methods on “dense-small” instances where 5-8 objects are packed in an environment with $\rho = 0.5$ (left: computation time in seconds; middle: number of collisions; right: success rate).

We further compare the performance of RBM-SP-BST and MCTS in lattice rearrangement problems recently studied in the literature [30]. An example with 15 objects is shown in Fig. 12[left]. In the start and goal arrangements, gaps between adjacent objects are set to be 0.01 object radius, and thus buffer allocation is challenging for sampling-based methods. While MCTS tries all actions on each node, RBM-

SP-BST can detect the embedded combinatorial object relationship via the dependency graph and therefore needs less buffer allocations. As shown in Fig. 12[right], RBM-SP-BST has much higher success rate in lattice rearrangement tasks.




Fig. 12. Comparison among methods in lattice instances with 3-35 objects. [left] lattice example; [right] success rate

C. Cuboid Objects

Because the MCTS solver only supports cylindrical objects, we only compare RBM-SP-BST and BiRRT(fmRS) in the cuboids setup (Fig. 5[right]). When $\rho = 0.3$, RBM-SP-BST computes high quality solutions efficiently, while BiRRT(fmRS) can only solve instances with up to 20 cuboids. We mention that, when $\rho = 0.4$, BiRRT(fmRS) cannot solve any instance, but RBM-SP-BST can solve 50-object rearrangement problems in 28.6 secs on average.




Fig. 13. Comparison between methods in cuboid instances with 10-30 cuboids and $\rho = 0.3$ (left: computation time in seconds; middle: success rate; right: number of actions as multiples of $|\mathcal{O}|$).

D. Hardware Demonstration




Fig. 14. Experimental settings in addition to that of Fig 1.

We further demonstrate that the plans computed by TRLB can be readily executed on real robots in a complete vision-planning-control pipeline. Our hardware setup consists of a UR-5e robot arm, an OnRobot VGC 10 vacuum gripper, and an Intel RealSense D435 RGB-D camera. As shown in the accompanying video, TRLB solves all attempted instances (Fig 1 and Fig. 14), which involves concave objects, in an apparently natural and efficient manner.

VI. CONCLUSION AND FUTURE WORK

The TRLB framework proposed in this work employs the dependency graph representation and a lazy buffer allocation approach for efficiently solving the problem of rearranging many tightly packed objects on a tabletop using internal buffers (TORI). Extensive simulation studies show that TRLB computes rearrangement plans of comparable or better quality as the state-of-the-art methods, and does so up to 2 magnitudes faster. Demonstration with real robot experiments shows that the solutions computed by TRLB are not only efficient but also appear natural. As such, the solutions can potentially be deployed as part of home automation and industrial automation solutions in next generation robots.

REFERENCES

- [1] S. D. Han, N. M. Stiffler, A. Krontiris, K. E. Bekris, and J. Yu, "Complexity results and fast methods for optimal tabletop rearrangement with overhand grasps," *The International Journal of Robotics Research*, vol. 37, no. 13-14, pp. 1775–1795, 2018.
- [2] K. Gao, S. W. Feng, and J. Yu, "On minimizing the number of running buffers for tabletop rearrangement," in *Robotics: Sciences and Systems*, 2021.
- [3] S. Bereg and A. Dumitrescu, "The lifting model for reconfiguration," *Discrete & Computational Geometry*, vol. 35, no. 4, pp. 653–669, 2006.
- [4] A. Cosgun, T. Hermans, V. Emeli, and M. Stilman, "Push planning for object placement on cluttered table surfaces," in *2011 IEEE/RSJ international conference on intelligent robots and systems*. IEEE, 2011, pp. 4627–4632.
- [5] R. Wang, K. Gao, D. Nakhimovich, J. Yu, and K. E. Bekris, "Uniform object rearrangement: From complete monotone primitives to efficient non-monotone informed search," in *IEEE International Conference on Robotics and Automation*, 2021.
- [6] R. Wang, Y. Miao, and K. E. Bekris, "Efficient and high-quality prehensile rearrangement in cluttered and confined spaces," *arXiv preprint arXiv:2110.02814*, 2021.
- [7] M. Stilman and J. J. Kuffner, "Navigation among movable obstacles: Real-time reasoning in complex environments," *International Journal of Humanoid Robotics*, vol. 2, no. 04, pp. 479–503, 2005.
- [8] M. Stilman, K. Nishiwaki, S. Kagami, and J. J. Kuffner, "Planning and executing navigation among movable obstacles," *Advanced Robotics*, vol. 21, no. 14, pp. 1617–1634, 2007.
- [9] M. Stilman and J. Kuffner, "Planning among movable obstacles with artificial constraints," *The International Journal of Robotics Research*, vol. 27, no. 11-12, pp. 1295–1307, 2008.
- [10] H. Zhang, Y. Lu, C. Yu, D. Hsu, X. Lan, and N. Zheng, "INVIGORATE: Interactive Visual Grounding and Grasping in Clutter," in *Proceedings of Robotics: Science and Systems*, Virtual, July 2021.
- [11] C. Nam, J. Lee, Y. Cho, J. Lee, D. H. Kim, and C. Kim, "Planning for target retrieval using a robotic manipulator in cluttered and occluded environments," *arXiv preprint arXiv:1907.03956*, 2019.
- [12] J. E. King, V. Ranganeni, and S. S. Srinivasa, "Unobservable monte carlo planning for nonprehensile rearrangement tasks," in *2017 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2017, pp. 4681–4688.
- [13] B. Huang, S. D. Han, A. Boularias, and J. Yu, "Dipn: Deep interaction prediction network with application to clutter removal," in *IEEE International Conference on Robotics and Automation*, 2021.
- [14] A. Krontiris and K. E. Bekris, "Dealing with difficult instances of object rearrangement," in *Robotics: Science and Systems*, vol. 1123, 2015.
- [15] —, "Efficiently solving general rearrangement tasks: A fast extension primitive for an incremental sampling-based planner," in *2016 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2016, pp. 3924–3931.
- [16] Y. Labbé, S. Zagoruyko, I. Kalevatykh, I. Laptev, J. Carpentier, M. Aubry, and J. Sivic, "Monte-carlo tree search for efficient visually guided rearrangement planning," *IEEE Robotics and Automation Letters*, vol. 5, no. 2, pp. 3715–3722, 2020.
- [17] A. Morgan, B. Wen, J. Liang, A. Boularias, A. Dollar, and K. Bekris, "Vision-driven compliant manipulation for reliable, high-precision assembly tasks," *Proceedings of Robotics: Science and Systems (RSS)*, 2021.
- [18] B. Wen, W. Lian, K. Bekris, and S. Schaal, "Catgrasp: Learning category-level task-relevant grasping in clutter from simulation," *arXiv preprint arXiv:2109.09163*, 2021.
- [19] K. Okada, A. Haneda, H. Nakai, M. Inaba, and H. Inoue, "Environment manipulation planner for humanoid robots using task graph that generates action sequence," in *2004 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)(IEEE Cat. No. 04CH37566)*, vol. 2. IEEE, 2004, pp. 1174–1179.
- [20] M. Levihn, J. Scholz, and M. Stilman, "Hierarchical decision theoretic planning for navigation among movable obstacles," in *Algorithmic Foundations of Robotics X*. Springer, 2013, pp. 19–35.
- [21] J. van Den Berg, J. Snoeyink, M. C. Lin, and D. Manocha, "Centralized path planning for multiple robots: Optimal decoupling into sequential plans," in *Robotics: Science and systems*, vol. 2, no. 2.5, 2009, pp. 2–3.
- [22] F. Wang and K. Hauser, "Robot packing with known items and nondeterministic arrival order," *IEEE Transactions on Automation Science and Engineering*, 2020.
- [23] S. H. Cheong, B. Y. Cho, J. Lee, C. Kim, and C. Nam, "Where to relocate?: Object rearrangement inside cluttered and confined environments for robotic manipulation," in *2020 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2020, pp. 7791–7797.
- [24] R. Bohlin and L. E. Kavraki, "Path planning using lazy prm," in *Proceedings 2000 ICRA. Millennium Conference. IEEE International Conference on Robotics and Automation. Symposia Proceedings (Cat. No. 00CH37065)*, vol. 1. IEEE, 2000, pp. 521–528.
- [25] J. Denny, K. Shi, and N. M. Amato, "Lazy toggle prm: A single-query approach to motion planning," in *2013 IEEE International Conference on Robotics and Automation*. IEEE, 2013, pp. 2407–2414.
- [26] K. Hauser, "Lazy collision checking in asymptotically-optimal motion planning," in *2015 IEEE international conference on robotics and automation (ICRA)*. IEEE, 2015, pp. 2951–2957.
- [27] D. Adolphson and T. C. Hu, "Optimal linear ordering," *SIAM Journal on Applied Mathematics*, vol. 25, no. 3, pp. 403–423, 1973.
- [28] Y. Shiloach, "A minimum linear arrangement algorithm for undirected trees," *SIAM Journal on Computing*, vol. 8, no. 1, pp. 15–32, 1979.
- [29] N. Chernov, Y. Stoyan, and T. Romanova, "Mathematical model and efficient algorithms for object packing problem," *Computational Geometry*, vol. 43, no. 5, pp. 535–553, 2010.
- [30] J. Yu, "Rearrangement on lattices with pic-n-swaps: Optimality structures and efficient algorithms," in *Robotics: Sciences and Systems*, 2021.