

On Running Buffer Minimization for Tabletop Rearrangement

Kai Gao Si Wei Feng Jingjin Yu

Abstract—For tabletop object rearrangement problems with overhand grasps, storage space which may be inside or outside the tabletop workspace, or running buffers, can temporarily hold objects which greatly facilitates the resolution of a given rearrangement task. This brings forth the natural question of how many running buffers are required so that certain classes of tabletop rearrangement problems are feasible. In this work, we examine the problem for both the labeled (where each object has a specific goal pose) and the unlabeled (where goal poses of objects are interchangeable) settings. On the structural side, we observe that finding the minimum number of running buffers (MRB) can be carried out on a dependency graph abstracted from a problem instance, and show that computing MRB on dependency graphs is NP-hard. We then prove that under both labeled and unlabeled settings, even for uniform cylindrical objects, the number of required running buffers may grow unbounded as the number of objects to be rearranged increases; we further show that the bound for the unlabeled case is tight. On the algorithmic side, we develop highly effective, exact algorithms for finding MRB for both labeled and unlabeled tabletop rearrangement problems, scalable to over a hundred objects under very high object density. More importantly, our algorithms also compute a sequence witnessing the computed MRB that can be used for solving object rearrangement tasks. Employing these algorithms, empirical evaluations reveal that random labeled and unlabeled instances, which more closely mimics real-world setups, generally have fairly small MRBs.

source: github.com/rutgers-arc-lab/running-buffer
video: youtu.be/hbD-cumF_H4

I. INTRODUCTION

In nearly all aspects of our everyday lives, be it work related, at home, or for play, objects are to be grasped and rearranged, e.g., tidying up a messy desk, cleaning the table after dinner, or solving a jigsaw puzzle. Similarly, many industrial and logistics applications require repetitive rearrangements of many objects, e.g., the sorting and packaging of products on conveyors with robots, and doing so efficiently is of critical importance to boost the competitiveness of the stakeholders. However, even without the challenge of grasping, deciding the sequence of objects for optimizing a rearrangement task is non-trivial. To that end, Han et al. [1] examined the problem of *tabletop object rearrangement with overhand grasps* (TORO), where objects may be picked up, moved around, and then placed at poses that are not in collision with other objects. An object that is picked up but cannot be directly placed at its goal is temporarily stored at a *buffer* location. For example, for the setup given in Fig. 1, using a single manipulator, either the



Fig. 1: A TORO instance where the three soda cans are to be rearranged from the left configuration to the right configuration.

Coke can or the Pepsi can must be moved to a buffer before the task can be completed. They show that computing a pick-n-place sequence that minimizes the use of the total number of buffers is NP-hard and provide fast methods for computing that solution for problems with a couple dozen objects.

In this study, we examine a more practical rearrangement objective which minimizes the number of *running buffers* (RB) in solving a TORO instance. We seek rearrangement plans that minimize the maximum number of objects stored at buffers at any given moment, assuming that each object is moved to a temporary location at most once. We denote this quantity as MRB (minimum running buffer). The objective is important because if the MRB required for solving a TORO instance exceeds the available buffer storage, which is always limited in practice, then the instance is infeasible under the model. Therefore, the structural results and the algorithms that we present may be used not only for computing feasible and high-quality rearrangement plans, but they are also invaluable as a verification tool, e.g., to verify that a certain rearrangement setup will be able to solve most tasks for which it is designed to tackle.

Besides introducing running buffers and the unlabeled dependency graph, this work brings forth several novel technical contributions. First, we show that computing MRB on arbitrary *dependency graphs*, which encode the combinatorial information of TORO instances, is NP-hard. Second, we establish that for an n -object TORO instance, MRB can be lower bounded by $\Omega(\sqrt{n})$ for uniform cylinders, even when all objects are *unlabeled*. This implies that the same is true for the *labeled* setting. Then, we provide a matching algorithmic upper bound $O(\sqrt{n})$ for the unlabeled setting. Last but not least, we develop multiple highly effective and optimal algorithms for computing rearrangement plans with MRB for TORO. In particular, we present a dynamic programming method for the labeled setting, a priority queue-based algorithm for the unlabeled setting, and a much more efficient *depth-first-search dynamic programming* routine that readily scales to instances

The authors are with the Department of Computer Science, Rutgers, the State University of New Jersey, Piscataway, NJ, USA. E-Mails: {kai.gao, siwei.feng, jingjin.yu}@rutgers.edu.

with over a hundred objects for both settings. Furthermore, we provide methods for computing plans with the minimum number of *total buffers* subject to the MRB constraints. These algorithms not only provide the optimal number of buffers but also provide a rearrangement plan that witness the optimal solution.

Related work. As a high utility capability, manipulation of objects in a bounded workspace has been extensively studied, with works devoted to perception/scene understanding [2]–[5], task/rearrangement planning [1], [6]–[15], manipulation [16]–[23], as well as integrated holistic approaches [24]–[28]. As object rearrangement problems often embed within them multi-robot motion planning problems, rearrangement inherits the PSPACE-hard complexity [29]. These problems remain NP-hard even without complex geometric constraints [30]. Considering rearrangement plan quality, e.g. minimizing the number of pick-n-places or the end-effector travel, is also computationally intractable [1].

For rearrangement tasks using mainly prehensile actions, the algorithmic studies of Navigation Among Movable Obstacles [7], [31] result in backtracking search methods that can effectively deal with monotone and other instances with “nice” properties. Via carefully calling monotone solvers, difficult non-monotone cases can be solved as well [11]. Han et al. [1] relates tabletop rearrangement problems to the Traveling Salesperson Problem [32] and the Feedback Vertex Set problem [33], both of which are NP-hard. Nevertheless, integer programming models are shown to quickly compute high quality solutions for practical sized (e.g., 1-2 dozen of objects) problems. Focusing mainly on the unlabeled setting, bounds on the number of pick-n-places are provided under different assumptions on disk objects in [34]. In [14], a complete algorithm is developed that reasons about object retrieval, rearranging other objects as needed, with later work [35] considering plan optimality and sensor occlusion. While objectives in most problems focus on the number of motions, Halperin et al. [36] seeks to minimize the space needed to carry out a rearrangement task in which discs move along straight lines in the workspace.

Non-prehensile rearrangement has also been extensively studied, with singulation as an early focus [37]–[39]. Iterative search was employed in [13] for accomplishing a multitude of rearrangement tasks spanning singulating, separation, and sorting of identically shaped cubes. Song et al. [40] combines Monte Carlo Tree Search with a deep policy network for separating many objects into coherent clusters within a bounded workspace, supporting non-convex objects. More recently, a bi-level planner is proposed [15], engaging both (non-prehensile) pushing and (prehensile) overhand grasping for sorting a large number of objects. Synergies between non-prehensile and prehensile actions have been explored for solving clutter removal tasks [?], [41] and the more challenging object retrieval tasks [?] using a minimum number of pushing and grasping actions.

On the structural side, a central object that we study is the *dependency graph* structure. To our knowledge, dependency

structures for rearrangement problems are first examined in [42] where partial dependencies are used for optimal decoupling of object interactions. Subsequently, the structure was employed for reasoning about and solving challenging rearrangement problems [?], [11], [43]. The full labeled dependency graph, as induced by a rearrangement instance, is first introduced and studied in [1]. This current work introduces the unlabeled dependency graph. We observe that, in the labeled setting, through the dependency graph, the running buffer problems naturally connect to *graph layout* problems [44]–[49], where an optimal linear ordering of graph vertices is sought. Graph layout problems find a vast number of important applications including VLSI design, scheduling [50], and so on. For the unlabeled setting, the dependency graph becomes a planar one for uniform objects with a square or round base. Rearrangement can be tackled through partitioning of the dependency graph using a *vertex separator* [51]–[54]. For a survey on these topics, see [44].

Paper organization. The rest of the paper is organized as follows. We introduce the MRB focused rearrangement problems and discuss the associated dependency graphs in Sec. II. Then, in Sec. III, we establish some basic structural properties of the optimality structure of the problems, and show that minimizing running buffer size on dependency graphs is computationally intractable. We proceed to establish the lower and upper bounds on MRB in Sec. IV and describe our proposed algorithmic solutions in Sec. V. Evaluation follows in Sec. VI. We conclude with Sec. VII.

II. PRELIMINARIES

We describe two practical (labeled and unlabeled) formulations of the tabletop object rearrangement problems using external buffers, and discuss the important *dependency graph* structure for both settings.

A. Labeled Tabletop Rearrangement with External Buffers

Consider a bounded workspace $\mathcal{W} \subset \mathbb{R}^2$ with a set of n objects $\mathcal{O} = \{o_1, \dots, o_n\}$ placed inside it. All objects are assumed to be *generalized cylinders* with the same height. A *feasible arrangement* of these objects is a set of poses $\mathcal{A} = \{x_1, \dots, x_n\}, x_i \in SE(2)$ in which no two objects collide. Let $\mathcal{A}_1 = \{x_1^s, \dots, x_n^s\}$ and $\mathcal{A}_2 = \{x_1^g, \dots, x_n^g\}$ be two feasible arrangements, a tabletop object rearrangement problem [1] seeks a plan using *pick-n-place* operations that move the objects from \mathcal{A}_1 to \mathcal{A}_2 (see Fig. 2(a) for an example with 7 uniform cylinders). In each pick-n-place operation, an object is grasped by a robot arm, lifted above all other objects, transferred to and lowered at a new pose $p \in SE(2)$ where the object will not be in collision with other objects, and then released. A pick-n-place operation can be formally represented as a 3-tuple $a = (i, x', x'')$, denoting that object o_i is moved from pose x' to pose x'' . A full rearrangement plan $P = (a_1, a_2, \dots)$ is then an ordered sequence of pick-n-place operations.

Depending on \mathcal{A}_1 and \mathcal{A}_2 , it may not always be possible to directly transfer an object o_i from x_i^s to x_i^g in a single pick-n-place operation, because x_i^g may be occupied by other objects.




Fig. 2: A 7-object labeled instance with uniform cylinders; we will use this instance as a running example. (a) The green discs (as projections of cylinders) represent the start arrangement \mathcal{A}_1 and the cyan discs represent the goal arrangement \mathcal{A}_2 . (b) The corresponding labeled dependency graph. (c) The corresponding unlabeled dependency graph, which is bipartite and planar.

This creates *dependencies* between objects. If object o_i at pose x_i^g intersects object o_j at pose x_j^s , we say o_i *depends on* o_j . This suggests that object o_j must be moved first before o_i can be placed at its goal pose x_i^g .

It is possible to have circular dependencies, e.g., between objects 3 and 5 in Fig. 2(a). In such cases, some object(s) must be temporarily moved to an intermediate pose to solve the rearrangement problem. Similar to [1], we assume that *external buffers* outside of the workspace are used for assuming intermediate poses, which avoids time-consuming geometric computations if the intermediate poses are to be placed within \mathcal{W} . During the execution of a rearrangement plan, there can be multiple objects that are stored at buffer locations. We call the buffers currently being used as *running buffers* (RB). With the introduction of buffers, there are three types of pick-n-place operations: 1) pick an object at its start pose and place at a buffer, 2) pick an object at its start pose and place at its goal pose, and 3) pick an object from buffer and place at its goal pose. Notice that buffer poses are not important. Naturally, it is desirable to be able to solve a rearrangement problem with the least number of running buffers, yielding the running buffer minimization problem.

Problem 1 (Labeled Running Buffer Minimization (LRBM)). *Given feasible arrangements \mathcal{A}_1 and \mathcal{A}_2 , find a rearrangement plan P that minimizes the maximum number of running buffers used at any given time.*

In an LRBM instance, the set of all dependencies induced by \mathcal{A}_1 and \mathcal{A}_2 can be represented using a directed graph $G_{\mathcal{A}_1, \mathcal{A}_2}^l = (V, A)$, where each $v_i \in V$ corresponds to object o_i and there is an arc $v_i \rightarrow v_j$ for $1 \leq i, j \leq n, i \neq j$ if object o_i depends on object o_j . We call $G_{\mathcal{A}_1, \mathcal{A}_2}^l$ a *labeled dependency graph*. The labeled dependency graph for Fig. 2(a) is given in Fig. 2(b). We can immediately identify multiple circular dependencies in the graph, e.g., between objects 3 and 5, or among objects 7, 2, 6 and 5. It is not difficult to see that the dependency graph abstraction fully captures the information needed to solve a tabletop rearrangement problem with external buffers.

B. Unlabeled Tabletop Rearrangement Problem

In an unlabeled setting, objects are interchangeable. That is, it does not matter which object goes to which goal. For example, in Fig. 2, object 5 can move to the goal for object 6. We call this version the Unlabeled Running Buffer Minimization (URBM) problem, which is intuitively easier. The plan for the unlabeled problem can be represented similarly as the labeled setting; we continue to use labels but do not require matching labels for start and goal poses.

For the unlabeled setting, there clearly remains dependency between start and goal arrangements, but in a different form. We update the *unlabeled* dependency graph for URBM as an undirected *bipartite* graph between the start arrangement and the goal arrangement. That is, $G_{\mathcal{A}_1, \mathcal{A}_2}^u = (V_1 \cup V_2, E)$ where each $v \in V_1$ (resp., $v \in V_2$) corresponds to a start (resp., goal) pose $p \in \mathcal{A}_1$ (resp., $p \in \mathcal{A}_2$). We denote the vertices representing the start and goal poses as *start vertices* and *goal vertices*, respectively. There is an edge between $v_1 \in V_1$ and $v_2 \in V_2$ if the objects at the corresponding poses overlap. The unlabeled dependency graph for Fig. 2(a) is given in Fig. 2(c).

We make a straightforward but important observation of the unlabeled dependency graph when objects are uniform cylinders, which is a key sub-class of TORO problems, e.g., many products can be approximated as uniform cylinders.

Proposition II.1. *For unlabeled tabletop object rearrangement problems where all objects are identical cylinders, the unlabeled dependency graph is a planar bipartite graph with maximum degree 5.*

Proof. The bipartite and planar part come directly from the problem setup. Since we work with uniform cylinders which have uniform disc base, one disc may only touch six non-overlapping discs and non-trivially intersect at most five non-overlapping discs. \square

III. STRUCTURAL ANALYSIS AND NP-HARDNESS

In this section, we highlight some important structural properties of LRBM, including (1) the comparison to minimizing the total number of buffers [1], (2) the solutions of LRBM and the *linear arrangement* [55] or *linear ordering* [56] of its dependency graph, and (3) the hardness of computing MRB for labeled dependency graphs.

A. Running Buffer versus Total Buffer

As mentioned in the introduction, running buffers are related to but different from the total number of buffers required, as studied in [1], to solve a rearrangement problem using external buffers. It was shown that the minimum number of total buffers for solving an LRBM is the same as the size of the minimum *feedback vertex set* (FVS) of the underlying dependency graph. An FVS is a set of vertices the removal of which leaves a graph acyclic. An LRBM with an acyclic dependency graph can be solved without using any buffer. We denote the size of the minimum FVS as MFVS.

As an example, for LRBM, consider a labeled dependency graph that is formed by n copies of 2-cycles. The MFVS is n . On the other hand, the MRB is just 1 for the problem.

That is, only a single external buffer is needed to solve the problem. Therefore, whereas the total number of buffers used has more bearing on global solution optimality, MRB sheds more light on *feasibility*. Knowing the MRB tells us whether a certain number of external buffers will be sufficient for solving a class of rearrangement problems. This is critical for practical applications where the number of external buffers is generally limited to be a small constant.

We give an example where the MRB and MFVS cannot always be optimized simultaneously. For the setup (Fig. 3) where objects have convex footprints, the MFVS, $\{7, 9, 10\}$, has size 3. Using our algorithms, to be detailed later, the MRB is 2 (e.g., with the sequence 10, 8, 4, 5, 3, 6, 7, 1, 2, 9, the interpretation of which is given in Sec. III-B). However, constrained on $\text{MRB} = 2$, the total number of buffers that must be used is at least $4 > 3$. We note that, this is rarely the case; for uniform cylinders, the total number of buffers needed after first minimizing the running buffer is almost always the same as the MFVS size.




Fig. 3: An LRBM instance with uniform thin cuboids (left) and its labeled dependency graph, where the total number of buffers needed is more than the size of the MFVS when the number of running buffers is minimized.

B. Linear Ordering of Graph Vertices and Running Buffer

Given a graph with vertex set V , a *linear ordering* of V is a bijective function $\varphi : \{1, \dots, |V|\} \rightarrow V$. Given a dependency graph $G = (V, A)$ for an LRBM and a linear ordering φ , we may turn it into a plan P by sequentially picking up objects corresponding to vertices $\varphi(1), \varphi(2), \dots$. For each object that is picked up, it is moved to its goal pose if it has no further dependencies; otherwise, it is stored in the buffer. Objects already in the buffer will be moved to their goal pose at the earliest possible opportunity.

For example, given the linear ordering 1, 5, 6, 3, 4, 2, 7 for the dependency graph from Fig. 2(b), first, o_1 can be directly moved to its goal. Then, o_5 is moved to the buffer because it has dependency on o_3 and o_7 (but no longer on o_1). Then, o_6 can be directly moved to the buffer because o_5 is now at a buffer location. Similarly, o_3 can be moved to its goal next. Then, o_4 and o_2 must be moved to buffer, after which o_7 can be moved to its goal directly. Finally, o_2, o_4 , and o_5 can be moved to their respective goals from the buffer. This leads to a maximum running buffer size of 3. This is not optimal; an optimal sequence is 5, 6, 2, 7, 4, 3, 1, with $\text{MRB} = 2$. Both sequences are illustrated in Fig. 4.

From the discussion, we may view the number of running buffers as a function of a dependency graph G and a linear




Fig. 4: Two linear orderings of vertices of the labeled dependency graph from Fig. 2(b). The right one minimizes MRB.

ordering φ , i.e., $\text{RB}(G, \varphi)$ is the number of running buffers needed for rearranging G following the order given by φ . We then have $\text{MRB}(G) = \min_{\varphi} \text{RB}(G, \varphi)$.

C. Intractability of Computing $\text{MRB}(G)$

Since computing MFVS is NP-hard [1], one would expect that computing MRB for a labeled dependency graph, which can be any directed graph, is also hard. We show that this is indeed the case, through examining the interesting relationship between MRB and the *vertex separation problem* (VSP), which is equivalent to path width, gate matrix layout and search number problems as described in Theorem 3.1 in [44], resulting from a series of studies [57]–[59]. Unless $P = NP$, there cannot be an absolute approximation algorithm for any of these problems [49]. First, we describe the vertex separation problem. Intuitively, given an undirected graph $G = (V, E)$, VSP seeks a linear ordering φ of V such that, for a vertex with order i , the number of vertices come no later than i in the ordering, with edges to vertices that come after i , is minimized.

Vertex Separation (VSP)

Instance: Graph $G(V, E)$ and an integer K .

Question: Is there a bijective function $\varphi : \{1, \dots, n\} \rightarrow V$, such that for any integer $1 \leq i \leq n$, $|\{u \in V \mid \exists(u, v) \in E \text{ and } \varphi(u) \leq i < \varphi(v)\}| \leq K$?

As an example, in Fig. 5(a), with the given linear ordering, at the second vertex, both the first and the second vertices have edges crossing the vertical separator, yielding a crossing number of 2. Given a graph G and a linear ordering φ , we define $\text{VS}(G, \varphi) := \max_i |\{u \in V \mid \exists(u, v) \in E \text{ and } \varphi(u) \leq i < \varphi(v)\}|$, VSP seeks φ that minimizes $\text{VS}(G, \varphi)$. Let $\text{MINVS}(G)$, the vertex separation number of graph G , be the minimum K for which a VSP instance has a yes answer, then $\text{MINVS}(G) = \min_{\varphi} \text{VS}(G, \varphi)$. Now, given an undirected graph G and a labeled dependency graph G^ℓ obtained from G by replacing each edge of G with two directed edges in opposite directions, we observe that there are clear similarities between $\text{VS}(G, \varphi)$ and $\text{RB}(G^\ell, \varphi)$, which is characterized in the following lemma.

Lemma III.1. $\text{VS}(G, \varphi) \leq \text{RB}(G^\ell, \varphi) \leq \text{VS}(G, \varphi) + 1$.

Proof sketch. Fixing a linear ordering φ , it is clear that $\text{VS}(G, \varphi) \leq \text{RB}(G^\ell, \varphi)$, since the vertices on the left side of a separator with edges crossing the separator for G corresponds to the objects that must be stored at buffer locations. For example, in Fig. 5(a), past the second vertex from the left,

both the first and the second vertices have edges crossing the vertical “separator”. In the corresponding dependency graph shown in Fig. 5(b), objects corresponding to both vertices must be moved to the external buffer. On the other hand, we have $\text{RB}(G^\ell, \varphi) \leq \text{VS}(G, \varphi) + 1$ because as we move across a vertex in the linear ordering, the corresponding object may need to be moved to a buffer location temporarily. For example, as the third vertex from the left in Fig. 5(a) is passed, the vertex separator drops from 2 to 1, but for dealing with the corresponding dependency graph in Fig. 5(b), the object corresponding to the third vertex from the left must be moved to the buffer before the first and the second objects stored in buffer can be placed at their goals. \square




Fig. 5: (a) An undirected graph and a linear ordering of its vertices. (b) A corresponding labeled dependency graph with the same vertex ordering.

Theorem III.1. *Computing MRB, even with an absolute approximation, for a labeled dependency graph is NP-hard.*

Proof. Given an undirected graph G , we reduce from approximating VSP within a constant to approximating MRB within a constant for a dependency graph G^ℓ from G constructed as stated before, replacing each edge in G as a bidirectional dependency.

Unless $P = NP$, VSP does not have absolute approximation in polynomial time. Henceforth, if $\text{MRB}(G^\ell, \varphi)$ can be approximated within α in polynomial time, which means for graph G , we can find a φ^* in polynomial time such that $\text{RB}(G^\ell, \varphi^*) \leq \text{MRB}(G^\ell) + \alpha$, we then have $\text{VS}(G, \varphi^*) \leq \text{RB}(G^\ell, \varphi^*) \leq \alpha + \text{MRB}(G^\ell) \leq \text{MINVS}(G) + \alpha + 1$, which shows vertex separation can have an absolute approximation, implying $P = NP$. \square

IV. LOWER AND UPPER BOUNDS ON MRB

We proceed to establish bounds on MRB, i.e., what is the lowest possible MRB for LRBM and URBM, and what is the best that we can do to lower MRB? An important outcome is that MRB can grow unbounded with the number of objects, even for URBM when objects are all uniform cylinders. Another very interesting result is that we are able to close the gap between lower and upper bound for URBM for uniform cylinders.

A. Intrinsic MRB Lower Bounds

When there is no restrictions on object footprint, MRB can easily reach the maximum possible $n - 1$ for an n object instance, even in the URBM case. Such an example is given in Fig. 6, where $n = 6$ thin cuboids are aligned horizontally in \mathcal{A}_1 , one above the other. The cuboids are vertically aligned in \mathcal{A}_2 , and every pair of start pose and goal pose induces a collision. Clearly, this yields a bidirectional K_6 labeled

dependency graph in the LRBM case and a $K_{6,6}$ unlabeled dependency graph in the URBM case. For both, $n - 1 = 5$ objects must be moved to buffer before the problem can be resolved. The example clearly can be generalized to an arbitrary number of objects.

Proposition IV.1. *MRB lower bound is $n - 1$ for n objects for both LRBM and URBM, which is the maximum possible, even for uniform convex shaped objects.*




Fig. 6: An instance with 6 cuboids where horizontal and vertical sets represent start and goal poses, respectively.

The lower bound on MRB being $\Omega(n)$ is undesirable, but it is established using objects that are “thin”. Everyday objects are not often like that. An ensuing question of high practical value is then: what happens when the footprint of the objects are “nicer”? Next, we show that, the lower bound drops to $\Omega(\sqrt{n})$ for uniform cylinders, which approximate many real-world objects/products. Further more, we show that this lower bound is tight for URBM (in Section IV-B).

We first establish the $\Omega(\sqrt{n})$ lower bound for URBM. For convenience, assume n is a perfect square, i.e., $n = m^2$ for some integer m . To get to the proof, a grid-like unlabeled dependency graph is used, which we call a *dependency grid*, where \mathcal{A}_1 and \mathcal{A}_2 have fixed grid (rotated by $\pi/4$) patterns, an example of which is given in Fig. 7. We use $\mathcal{D}(w, h)$ to denote a dependency grid with w columns and h rows. Let (x, y) be the coordinate of a vertex $v_{x,y}$ on $\mathcal{D}(w, h)$ with the top left being $(1, 1)$. The parity of $x + y$ determines the partite set of the vertex (recall that unlabeled dependency graph for uniform cylinders is always a planar bipartite graph, by Proposition II.1), which may correspond to a start pose or a goal pose. With this in mind, we simply call vertices of $\mathcal{D}(w, h)$ start and goal vertices; let $v_{1,1}$ be a start vertex.




Fig. 7: A URBM instance (left) and its unlabeled dependency graph (right), a 4×3 dependency grid. Green and cyan indicate start and goal arrangements, respectively.

We use $\mathcal{D}(m, 2m)$ for establishing the lower bound on MRB. We use a *vertex pair* $p_{i,j}$ to refer to two adjacent vertices $v_{i,2j-1}$ and $v_{i,2j}$ in $\mathcal{D}(m, 2m)$. It is clear that a vertex pair contains a start and a goal vertex. We say that a goal vertex

is *filled* if an object is placed at the corresponding goal pose. We say that a start vertex (which belongs to a vertex pair) is *cleared* if the corresponding object at the vertex is picked (either put at a goal or at a buffer) but the corresponding goal in the vertex pair is not filled. At any moment when the robot is not holding an object, the number of objects in the buffer is the same as the number of cleared vertices. For each column i , $1 \leq i \leq m$, let f_i (resp., c_i) be the number of goal (resp., start) vertices in the column that are filled (resp., cleared). Notice that a goal cannot be filled until the object at the corresponding start vertex is removed.

Lemma IV.1. *On a dependency grid $\mathcal{D}(m, 2m)$, for two adjacent columns i and $i+1$, $1 \leq i < m$, if $f_i + f_{i+1} \neq 0$ or $2m$, then $c_i + c_{i+1} \geq 1$. In other words, there is at least one cleared vertex in the two adjacent columns unless $f_i = f_{i+1} = 0$ or $f_i = f_{i+1} = m$.*

Proof. If there is a j , $1 \leq j \leq m$, such that only one of the goal vertices in vertex pairs $p_{i,j}$ and $p_{i+1,j}$ is filled (Fig. 8(a)), then the start vertex in the other vertex pair must be cleared. Therefore, $c_i + c_{i+1} \geq 1$.

On the other hand, if, for each j , $1 \leq j \leq m$, both or neither of the goal vertices in $p_{i,j}$ and $p_{i+1,j}$ is filled, then there is a j , $1 \leq j \leq m-1$, such that both goal vertices in $p_{i,j}$ and $p_{i+1,j}$ are filled but neither of those in $p_{i,j+1}$ and $p_{i+1,j+1}$ is filled (Fig. 8(b)) or the opposite (Fig. 8(c)). Then, for the vertex pairs whose goal vertices are not filled, say $p_{i,j+1}$ and $p_{i+1,j+1}$, one of their start vertices is a neighbor of the filled goal in $p_{i,j}$ and $p_{i+1,j}$. Therefore, at least one of the start vertices in $p_{i,j+1}$ and $p_{i+1,j+1}$ is a cleared vertex. And thus, $c_i + c_{i+1} \geq 1$.




Fig. 8: Some cases discussed in the lemma IV.1. The green and cyan nodes represent the start and goal vertices in the dependency graph. Specifically, the cyan nodes with a dot inside represent the filled vertices and the green nodes with a cross inside represent the cleared vertices. (a) When only one goal vertex in $p_{i,j}$ and $p_{i+1,j}$ is filled up, the start vertex in the other vertex pair is a cleared vertex. (b) When both goal vertices in $p_{i,j}$ and $p_{i+1,j}$ are filled but neither of those in $p_{i,j+1}$ and $p_{i+1,j+1}$ is filled, one of the start vertices $p_{i,j+1}$ and $p_{i+1,j+1}$ is a cleared vertex. (c) The opposite case of (b).

□

Lemma IV.2. *Given a URBM instance with $n = m^2$ objects and whose dependency graph is $\mathcal{D}(m, 2m)$, its MRB is lower bounded by $\Omega(m) = \Omega(\sqrt{n})$*

Proof. We show that there are $\Omega(m)$ cleared vertices when $\lfloor n/3 \rfloor$ goal vertices are filled. Suppose there are q columns in

\mathcal{D} with $1 \leq f_i \leq m-1$. According to the definition of f_i , for each of these q columns, there is at least one goal vertex that is filled and at least one goal vertex that is not.

If $q < \frac{\lfloor n/3 \rfloor}{3(m-1)}$, then there are two columns i and j , such that $f_i = m$ and $f_j = 0$. That is because $\sum_{1 \leq i \leq m} f_i = \lfloor n/3 \rfloor$ and $0 \leq f_i \leq m$ for all $1 \leq i \leq m$. Therefore, for the vertex pairs in each row j , at least one goal vertex is filled but at least one is not. And thus, for each j , there are two adjacent columns $i, i+1$, $1 \leq i < m$, such that there is only one goal vertex in $p_{i,j}$ and $p_{i+1,j}$ is filled and the start vertex in the other vertex pair is cleared (Fig. 8(a)). Therefore, there are at least m cleared vertices in this case.

If $q \geq \frac{\lfloor n/3 \rfloor}{3(m-1)}$, then we partition all the columns in \mathcal{D} into $\lfloor m/2 \rfloor$ disjoint pairs: (1,2), (3,4), ... The q columns belong to at least $\lfloor q/2 \rfloor$ pairs of adjacent columns. Therefore, according to Lemma IV.1, we have $\Theta(m)$ cleared vertices.

In conclusion, there are $\Omega(m)$ cleared vertices when there are $\lfloor n/3 \rfloor$ filled goal. Therefore, the minimum MRB of this instance is $\Omega(m)$. □

Because a URBM always have lower MRB than an LRBM with the same objects and goal placements, the conclusion of Lemma IV.2 directly applies to LRBM. Therefore, we have

Theorem IV.1. *For both URBM and LRBM with n uniform cylinders, MRB is lower bounded by $\Omega(\sqrt{n})$.*

For uniform cylinders, while the lower bound on URBM is tight (as shown in Section IV-B), we do not know whether the lower bound on LRBM is tight; our conjecture is that $\Omega(\sqrt{n})$ is not a tight lower bound for LRBM. Indeed, the $\Omega(\sqrt{n})$ lower bound can be realized when uniform cylinders are simply arranged on a cycle, an illustration of which is given in Fig. 9. For a general construction, for each object o_i , let o_i depend on $o_{(i-1) \bmod n}$ and $o_{(i+\sqrt{n}) \bmod n}$, where n is the number of objects in the instance. From the labeled dependency graph, we can construct the actual LRBM instance where start and goal arrangements both form a cycle. We can show that when $n/2$ objects are at the goal poses, $\Omega(\sqrt{n})$ objects are at the buffer. We omit the proof, which is similar in spirit to that for Lemma IV.2.




Fig. 9: An example of a 9-object LRBM yielding $\Omega(\sqrt{n})$ MRB (left) and the corresponding dependency graph (right).

B. MRB Algorithmic Upper Bounds

We now establish that regardless of how n uniform cylinders are to be rearranged, the corresponding URBM instance admits

a solution using $O(\sqrt{n})$ running buffers. The lower and upper bounds on URBM agree and are therefore tight.

Theorem IV.2. *For URBM with n uniform cylinders, a polynomial time algorithm can compute a plan with $O(\sqrt{n})$ RB, which implies that MRB is bounded by $O(\sqrt{n})$.*

Proof. When objects are all uniform cylinders, the corresponding unlabeled dependency graph $G_{\mathcal{A}_1, \mathcal{A}_2}^u$ is a planar graph with degree no more than 5, by Proposition II.1. We propose an $O(n \log(n))$ -time algorithm SEPPLAN for the setting based on a vertex separator of $G_{\mathcal{A}_1, \mathcal{A}_2}^u$. SEPPLAN can find a rearrangement plan with $O(\sqrt{n})$ running buffers.

In SEPPLAN, given the planar $G_{\mathcal{A}_1, \mathcal{A}_2}^u = (V, E)$, we can partition V into three disjoint subsets A , B and C [51], such that there is no edge connecting vertices in A and B , $|A|, |B| \leq 2|V|/3$, and $|C| \leq 2\sqrt{2|V|}$ (Fig. 10(a)). For the start vertices in C and the neighbors of the goal vertices in C , we remove them from $G_{\mathcal{A}_1, \mathcal{A}_2}^u$. Since there are at most 5 neighbors for each goal vertex, there are at most $10\sqrt{2|V|}$ objects moved to the buffer in this operation. After that, we remove the goal vertices in C which should be isolated now. Let A' , B' be the remaining vertices in A and B . With the removal of C from $G_{\mathcal{A}_1, \mathcal{A}_2}^u$, A' and B' form two independent subgraphs (Fig. 10(b)). We can deal with the subgraphs one after the other by recursively calling SEPPLAN (Fig. 10(c)). Let $\delta(V') := g(V') - s(V')$ where $g(V')$ and $s(V')$ are the number of goal and start vertices in a vertex set V' respectively. Between vertex subsets A and B , we prioritize the one with larger $\delta(\cdot)$ value. With straightforward reasoning, the $10\sqrt{2|V|}$ upper bound of the additional buffer size still holds after we deal with the prioritized subgraph.




Fig. 10: The recursive solver SEPPLAN for URBM. (a) A $O(\sqrt{|V|})$ vertex separator for the planar dependency graph. (b) By removing the start vertices in C and the neighbors of the goal vertices in C , the remaining graph consists of two independent subgraphs and isolated goal poses in C . (c) The problem can be solved by recursively calling SEPPLAN.

Let the maximum RB of the plan found by SEPPLAN be $b(m)$, where $m = 2n$ is the number of vertices in the dependency graph. We have $b(m) \leq b(2m/3) + O(\sqrt{m})$. And therefore, $b(m) = O(\sqrt{m})$. Since the $O(\sqrt{|V|})$ vertex separator can be found in $O(|V|)$ time [51], SEPPLAN can find the rearrangement plan for an arbitrary n -object URBM instance in $O(n \log(n))$ time. Due to limited space, the details of SEPPLAN and this proof are omitted; they will be made available in the extended version of this paper to be put on arXiv. \square

V. FAST ALGORITHMS FOR LRBM AND URBM

In this section, we first describe a dynamic programming-based method for LRBM (Sec. V-A). Then, we propose a priority queue based method in Sec. V-B for URBM. Finally, a significantly faster depth-first modification of DP for computing MRB is provided in Sec. V-C. We mention that, we also developed an integer programming model, denoted TB_{MRB} , for computing the minimum total number of buffers needed subject to the MRB constraint, which is compared with the algorithm that computes the minimum total buffer without the MRB constraint, denoted as TB_{FVS} , from [1]. A brief description of TB_{MRB} is given in Sec. V-D.

A. Dynamic Programming (DP) for LRBM

As mentioned in Sec. III-B, a rearrangement plan in LRBM can be represented as a linear ordering of object labels where objects will be moved out of the start pose based on the order. That is, given an ordering of objects, π , we start with $o_{\pi(1)}$. If $x_{\pi(1)}^g$ is not occupied, then $o_{\pi(1)}$ is directly moved there. Otherwise, it is moved to buffer location. We then continue with the second object in the order, and so on. After we work with each object in the given order, we always check whether objects in buffer can be moved to their goals, and do so if an opportunity presents. We now describe a *dynamic programming* (DP) algorithm for computing an ordering that yields the MRB.

The pseudo-code of the algorithm is given in Algo. 1. The algorithm maintains a search tree T , each node of which represents an arrangement where a set of objects S have left the corresponding start poses. We record the objects currently at the buffer ($T[S].b$) and the minimum running buffer from the start arrangement \mathcal{A}_1 to the current arrangement ($T[S].MRB$). The DP starts with an empty T . We let the root node represent \mathcal{A}_1 (line 1). At this moment, there is no object in the buffer and the MRB is 0 (line 2-3). And then we enumerate all the arrangements with $|S| = 1, 2, \dots$ and finally n (line 4-5). For arbitrary S , the objects at the buffer are the objects in S whose goal poses are still occupied by other objects (line 6), i.e., $\{o \in S \mid \exists o' \in \mathcal{O} \setminus S, (o, o') \in A\}$, where A is the set of arcs in $G_{\mathcal{A}_1, \mathcal{A}_2}^l$. $T[S].MRB$, the minimum running buffer from the root node $T[\emptyset]$ to $T[S]$, depends on the last object o_i added into S and can be computed by enumerating o_i (line 7-20):

$$T[S].MRB = \min_{o_i \in S} \max(T[S \setminus \{o_i\}].MRB, |T[S].b|, |T[S \setminus \{o_i\}].b| + TC(S \setminus \{o_i\}, S)),$$

where the *transition cost* TC is given as

$$TC(S \setminus \{o_i\}, S) = \begin{cases} 1, & o_i \in T[S].b, \\ 0, & \text{otherwise,} \end{cases}$$

with x_i^g currently occupied (line 10), the transition cost is due to objects dependent on o_i cannot be moved out of the buffer before moving o_i to the buffer (line 11). If $T[S].MRB$ is minimized with o_i being the last object in S from the starts,

then $T[S \setminus \{o_i\}]$ is the parent node of $T[S]$ in T (line 14-16). Once $T[\mathcal{O}]$ is added into T , $T[\mathcal{O}].MRB$ is the MRB of the instance (line 17) and the path in T from $T[\emptyset]$ to $T[\mathcal{O}]$ is the corresponding solution to the instance.

Algorithm 1: Dynamic Programming

Input : $G_{\mathcal{A}_1, \mathcal{A}_2}(\mathcal{O}, A)$: labeled dependency graph
Output: MRB: the minimum number of running buffers

```

1  $T.root \leftarrow \emptyset$ 
2  $T[\emptyset].b \leftarrow \emptyset$  % objects currently at the buffer
3  $T[\emptyset].MRB \leftarrow 0$  % current minimum running buffer
4 for  $1 \leq k \leq |\mathcal{O}|$  do
    % enumerate cases where  $k$  objects have left
    % the start poses
5   for  $S \in k\text{-combinations of } \mathcal{O}$  do
6      $T[S].b \leftarrow \{o \in S \mid \exists o' \in \mathcal{O} \setminus S, (o, o') \in A\}$ 
7     % Find the MRB from  $T[\emptyset]$  to  $T[S]$ 
8      $T[S].MRB \leftarrow \infty$ 
9     for  $o_i \in S$  do
10       $parent = S \setminus \{o_i\}$ 
11      if  $o_i \in T[S].b$  then
12         $RB \leftarrow \max(T[parent].MRB, |T[S].b|,$ 
13           $|T[parent].b| + 1)$ 
14      else
15         $RB \leftarrow \max(T[parent].MRB, |T[S].b|)$ 
16      if  $RB < T[S].MRB$  then
17         $T[S].MRB \leftarrow RB$ 
18         $T[S].parent \leftarrow parent$ 
19 return  $T[\mathcal{O}].MRB$ 

```

For the LRBM instance in Fig. 2, Tab. I shows $T[S].MRB$ with different last-object options when $S = \{o_2, o_5, o_6\}$. If the last object o_i is o_5 , then we need to move o_5 into the buffer before moving o_6 out of the buffer. Therefore, even though buffer sizes of the parent node and the current node are both 2, there is a moment when all of the three objects are at the buffer. However, when we choose o_2 or o_6 as the last object to add, the $T[S].MRB$ becomes 2.

TABLE I. $T[S].MRB$ for different last objects ($[p] = [parent]$)

Last object	$T[p].MRB$	$T[p].b$	$T[S].b$	$T[S].MRB$
o_2	1	$\{o_5\}$	$\{o_2, o_5\}$	2
o_5	2	$\{o_2, o_6\}$	$\{o_2, o_5\}$	3
o_6	2	$\{o_2, o_5\}$	$\{o_2, o_5\}$	2

B. A Priority Queue based method for URBM

Similar to LRBM, rearrangement plans in URBM can be represented by a linear ordering of goal vertices in $G_{\mathcal{A}_1, \mathcal{A}_2}^u$. We can compute the ordering that yields MRB by maintaining a search tree like in Algo. 1. Each node $T[V]$ in the tree represents an arrangement where a set of goal vertices V have been filled. The remaining dependencies of $T[V]$ is an induced graph of $G_{\mathcal{A}_1, \mathcal{A}_2}^u$, formed from $V(G_{\mathcal{A}_1, \mathcal{A}_2}^u) \setminus (V \cup N(V))$ where $N(V)$ is the neighbors of V in $G_{\mathcal{A}_1, \mathcal{A}_2}^u$. The running buffer size of $T[V]$ is $|N(V)| - |V|$. Given an induced graph $I(V)$, denote the goal vertices with no more than one neighbor in $I(V)$ as *free goals*. We make two observations. First, given an induced graph $I(V)$, we can always prioritize the free

goals in terms of the order to fill without optimality loss. Second, multiple free goals may be generated as a goal vertex is filled. For example, in the instance shown in Fig. 2(a), when the goal representing c_g^5 is filled, c_g^2 , c_g^3 , and c_g^4 become free goals and can be added to the linear ordering in an arbitrary order. In conclusion, the necessary nodes (nodes without free goals in the induced graph) in the search tree are sparse and enumerating nodes with DP carries much overhead.

As such, instead of exploring the search tree layer by layer like DP, we maintain a sparse tree with a priority queue Q . While each node still represents an arrangement, each edge in the tree represents either an action moving an object to the buffer or multiple actions filling free goal poses. We always pop out and develop the node with the smallest MRB in Q . If a child node of the one that we develop already exist in the tree but is with smaller MRB than previously claimed, we will update the parent of the child node into the node we are developing. The MRB of the node representing \mathcal{A}_2 sets an upper bound of the solution and nodes in Q with larger MRB will be pruned away. The algorithm terminates when Q is empty. We denote this priority queue-based search method PQS.

C. Dynamic Programming with Depth-First Exploration

Both LRBM and URBM can be viewed as solving a series of decision problems, i.e., asking whether we can find a rearrangement plan with $k = 1, 2, \dots$ running buffers. As dynamic programming is applied to solve such decision problems, instead of performing the more standard breadth first exploration of the search tree, we identified that a depth-first exploration is much more effective. We call this variation of dynamic programming, which is a fairly straightforward alteration of a standard DP procedure. Essentially, DFDP fixes a k and checks whether there is a plan requiring no more than k running buffers. As the search tree (see Sec. V-A) is explored, depth-first exploration is used instead of breadth-first. The intuition is that, when there are many rearrangement plans on the search tree that do not use more than k running buffers, depth-first search will quickly find such a solution, whereas a standard DP must grow the full search tree before returning a feasible solution. A similar depth-first exploration heuristic is used in [60].

D. Minimizing Total Buffers Subject to MRB Constraints

Let binary variables $c_{i,j}$ represent $G_{\mathcal{A}_1, \mathcal{A}_2}^\ell$: $c_{i,j} = 1$ if and only if $i = j$ or (i, j) is in the arc set of $G_{\mathcal{A}_1, \mathcal{A}_2}^\ell$. Let $y_{i,j}$ be the binary sequential variables: $y_{i,j} = 1$ if and only if $i = j$ or o_i moves out of the start pose before o_j . MRB can be expressed based on three constraints: First, MRB is at least the size of the running buffer at any moment; Second, an object $o_j \in \mathcal{O}$ is at the goal pose if and only if all the objects $o_k \in \mathcal{O}$ with $c_{j,k} = 1$ have left the start poses. Third, an object $o_j \in \mathcal{O}$ is at the buffer if and only if o_j is neither at the start pose nor at the goal pose. To encode the constraints, we further introduce two sets of binary variables $g_{i,j}$ and $b_{i,j}$. $g_{i,j} = 1$ and $b_{i,j} = 1$ indicate that o_j is at the goal pose and o_j is at

the buffer respectively when the action moving o_i away from the start pose is complete. Therefore, given a running buffer limit k , we have:

$$k \geq \sum_{1 \leq j \leq n} b_{i,j} \quad \forall 1 \leq i \leq n \quad (1)$$

$$\sum_{1 \leq l \leq n} \frac{c_{j,l} y_{i,l}}{n} \leq g_{i,j} \leq \sum_{1 \leq l \leq n} c_{j,l} y_{i,l} \quad \forall 1 \leq i, j \leq n \quad (2)$$

$$\sum_{1 \leq l \leq n} \frac{2 - y_{i,l} - g_{i,l}}{2} \leq b_{i,j} \leq \sum_{1 \leq l \leq n} (2 - y_{i,l} - g_{i,l}) \quad (3)$$

$$\forall 1 \leq i, j \leq n$$

The constraints can be added into the MFVS ILP formulation [1] to find the minimum total buffer size with at most k running buffers. We denote this method as TB_{MRB} when $k = MRB$, and the MFVS method from [1] as TB_{FVS} .

VI. EXPERIMENTAL STUDIES

Our evaluation focuses on uniform cylinders, given their prevalence in practical applications. For simulation studies, instances with different object densities are created, as measured by *density level* $\rho := n\pi r^2 / (h * w)$, where n is the number of objects and r is the base radius. h and w are the height and width of the workspace. In other words, ρ is the proportion of the tabletop surface occupied by objects.

The evaluation is conducted on both random object placements and manually constructed difficult setups (e.g., dependency grids with $MRB = \Omega(\sqrt{n})$). For generating test cases with high ρ value, we invented a physics engine (we used Gazebo) based approach for doing so. Within a rectangular box, we sample placements of cylinders at lower density and then also sample locations for some smaller “filler” objects (see Fig. 11, left). From here, one side of the box is pushed to reach a high density setting (Fig. 11, right), which is very difficult to generate via random sampling. By controlling the ratio of the two types of objects, different density levels can be readily obtained. Fig. 12 shows three random object placements for $\rho = 0.2, 0.4$ and 0.6 .




Fig. 11: Generating dense instances using a physics-engine based simulator through compression of the left scene to the right scene.




Fig. 12: Unlabeled arrangements with $\rho = 0.2, 0.4, 0.6$ respectively.

From two randomly generated object placements with same ρ and n values, a URBM instance can be readily created by superimposing one over the other. LRBM instances can be generated from URBM instances by assigning each object a random label in $[n]$ for both start and goal configurations.

The proposed algorithms are implemented in Python and all experiments are executed on an Intel® Xeon® CPU at 3.00GHz. For solving ILP, Gurobi 9.16.0 [61] is used.

A. Labeled Rearrangement over Random Instances

In Fig. 13, we compare the effectiveness of the DP and DFDP, in terms of computation time and success rate, for different densities. Each data point is the average of 30 test cases minus the unfinished ones, if any, subject to a time limit of 300 seconds per test case. For LRBM, we are able to push to $\rho = 0.4$, which is fairly dense. The results clearly demonstrate that DFDP significantly outperform the baseline DP. Based on the evaluation, both methods can be used to tackle practical sized problems (e.g., tens of objects), with DFDP demonstrating superior efficiency and robustness.




Fig. 13: Performance of DFDP and DP over LRBM. The top row shows the average computation time (s) and the bottom row the success rate, for density levels $\rho = 0.2, 0.3, 0.4$, from left to right. The x -axis denotes the number of objects involved in a test case.

The actual MRB sizes for the same test cases from Fig. 13 are shown in Fig. 14 on the left. We observe that MRB is rarely very large even for fairly large LRBM instances. The size of MRB appears correlated to the size of the largest connected component of the underlying dependency graph, shown in Fig. 14 on the right.




Fig. 14: For LRBM instances with $\rho = 0.2-0.4$ and $n = 20-100$, the left figure shows average MRB size and range. The right figure shows the size of the largest connected component of the dependency graph.

For LRBM with $\rho = 0.3$ and n up to 50, we computed the MFVS sizes using TB_{FVS} (which does not scale to higher ρ and n) and compared that with the MRB sizes, as

shown in Fig. 15 (a). We observe that the MFVS is about twice as large as MRB, suggesting that MRB provides more reliable information for estimating the design parameters of pick-n-place systems. For these instances, we also computed the total number of buffers needed subject to the MRB constraint using TB_{MRB} . Out of about 150 instances, only 1 showed a difference as compared with MFVS (therefore, this information is not shown in the figure). In Fig. 15 (b), we provided computation time comparison between TB_{FVS} and TB_{MRB} , showing that TB_{MRB} is practical, if it is desirable to minimize the total buffers after guaranteeing the minimum number of running buffers.




Fig. 15: (a) Comparison between size of MRB and MFVS. (b) Computation time comparison between TB_{FVS} and TB_{MRB} .

Considering our theoretical findings and the evaluation results, an important conclusion can be drawn here is that MRB is effectively a small constant for random instances, even when the instances are very large. Also, minimizing the total number of buffers used subject to MRB constraint can be done quickly for practical sized problems.

B. Unlabeled Rearrangement over Random Instances

For URBM, we carry out similar performance evaluation as we have done for LRBM. Here, PQS and DFDP are compared. For each combination of ρ and n , 100 random test cases are evaluated. Notably, we can reach $\rho = 0.6$ with relative ease. From Fig. 16, we observe that DFDP is more efficient than PQS, especially for large-scale dense settings. In terms of the MRB size, all instances tested has an average MRB size between 0 and 0.7, which is fairly small (Fig. 17). Interestingly, we witness a decrease of MRB as the number of objects increases, which could be due to the lessening “border effect” of the larger instances. That is, for instances with fewer objects, the bounding square puts more restriction on the placement of the objects inside. For larger instances, such restricting effects become smaller. We mention that the total number of buffers for random URBM cases subject to MRB constraints are generally very small.

C. Manually Constructed Difficult Cases

In the random scenario, the running buffer size is limited. In particular, for LRBM, the dependency graph tends to consist of multiple strongly connected components that can be dealt with independently. We further show the performance of DFDP on the instances with $MRB = \Theta(\sqrt{n})$. We evaluate three kinds of instances: (1) UG: m^2 -object URBM instances whose G_{A_1, A_2}^u are dependency grid $\mathcal{D}(m, 2m)$ (e.g., Fig. 7); (2) LG: m^2 -object LRBM instances whose start and goal arrangements are the same as the instances in (1). (3) LC: m^2 -object




Fig. 16: Performance of DFDP and PQS over URBM. The top row shows the average computation time and the bottom row shows the success rate, for density levels $\rho = 0.4, 0.5, 0.6$, from left to right.




Fig. 17: Average MRB size for URBM instances with $\rho = 0.4 - 0.6$ and $n = 20 - 100$. For $\rho = 0.4$ and 0.5 , the MRB sizes are near zero as the number of objects goes beyond 20.

LRBM instances with objects placed on a cycle (Fig. 9). The computation time and the corresponding MRB are shown in Fig. 18. For LG instances, the labels are randomly assigned. We try 30 test cases and then plot out the average. We observe that the MRB are much larger for these handcrafted instances as compared with random instances with similar density and number of objects.




Fig. 18: For handcrafted cases and different number of objects, the left figure shows the computation time by DFDP and the right figure shows the resulting MRB size.

VII. CONCLUSION AND DISCUSSION

In this work, we investigate the problem of minimizing the number of running buffers (MRB) for solving labeled and unlabeled tabletop rearrangement problems with overhead grasps (TORO), which translates to finding a best linear ordering of vertices of the associated underlying dependency graph. For TORO, MRB is an important quantity to understand as it determines the problem’s feasibility if only external buffers are to be used, which is the case in some real-world applications [1]. Despite the provably high computational complexity that is involved, we provide effective dynamic programming-based algorithms capable of quickly computing MRB for large and dense labeled/unlabeled TORO instances. In addition, we also

provide methods for minimizing the total number of buffers subject to MRB constraints. Whereas we prove that MRB can grow unbounded for both labeled and unlabeled settings for special cases for uniform cylinders, empirical evaluations suggest that real-world random TORO instances are likely to have much smaller MRB values.

We conclude by leaving the readers with some interesting open problems. On the structural side, while LRBM in general is proven to be NP-Hard, the computational intractability of either LRBM with uniform cylinders or URBM in general remains unresolved. As for bounds, the lower and upper bounds of MRB for LRBM for uniform cylinders do not yet agree; can the bound gap be narrowed further?

ACKNOWLEDGMENTS

This work is supported in part by NSF awards IIS-1734419, IIS-1845888, and CCF-1934924. We sincerely thank the anonymous reviewers for their insightful comments and suggestions.

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] A. Saxena, J. Driemeyer, and A. Y. Ng, "Robotic grasping of novel objects using vision," *The International Journal of Robotics Research*, vol. 27, no. 2, pp. 157–173, 2008.
- [3] M. Gualtieri, A. Ten Pas, K. Saenko, and R. Platt, "High precision grasp pose detection in dense clutter," in *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2016, pp. 598–605.
- [4] C. Mitash, K. E. Bekris, and A. Boularias, "A self-supervised learning system for object detection using physics simulation and multi-view pose estimation," in *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2017, pp. 545–551.
- [5] Y. Xiang, T. Schmidt, V. Narayanan, and D. Fox, "Posecnn: A convolutional neural network for 6d object pose estimation in cluttered scenes," in *Robotics: Science and Systems*, 2018.
- [6] O. Ben-Shahar and E. Rivlin, "Practical pushing planning for rearrangement tasks," *IEEE Transactions on Robotics and Automation*, vol. 14, no. 4, pp. 549–565, 1998.
- [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] 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.
- [9] G. Havur, G. Ozbilgin, E. Erdem, and V. Patoglu, "Geometric rearrangement of multiple movable objects on cluttered surfaces: A hybrid reasoning approach," in *2014 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2014, pp. 445–452.
- [10] J. A. Haustein, J. King, S. S. Srinivasa, and T. Asfour, "Kinodynamic randomized rearrangement planning via dynamic transitions between statically stable states," in *2015 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2015, pp. 3075–3082.
- [11] A. Krontiris and K. E. Bekris, "Dealing with difficult instances of object rearrangement," in *Robotics: Science and Systems*, vol. 1123, 2015.
- [12] J. E. King, M. Cagnetti, and S. S. Srinivasa, "Rearrangement planning using object-centric and robot-centric action spaces," in *2016 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2016, pp. 3940–3947.
- [13] E. Huang, Z. Jia, and M. T. Mason, "Large-scale multi-object rearrangement," in *2019 International Conference on Robotics and Automation (ICRA)*. IEEE, 2019, pp. 211–218.
- [14] J. Lee, Y. Cho, C. Nam, J. Park, and C. Kim, "Efficient obstacle rearrangement for object manipulation tasks in cluttered environments," in *2019 International Conference on Robotics and Automation (ICRA)*. IEEE, 2019, pp. 183–189.
- [15] Z. Pan and K. Hauser, "Decision making in joint push-grasp action space for large-scale object sorting," *arXiv preprint arXiv:2010.10064*, 2020.
- [16] R. H. Taylor, M. T. Mason, and K. Y. Goldberg, "Sensor-based manipulation planning as a game with nature," in *Fourth International Symposium on Robotics Research*, 1987, pp. 421–429.
- [17] K. Y. Goldberg, "Orienting polygonal parts without sensors," *Algorithmica*, vol. 10, no. 2, pp. 201–225, 1993.
- [18] K. M. Lynch and M. T. Mason, "Dynamic nonprehensile manipulation: Controllability, planning, and experiments," *The International Journal of Robotics Research*, vol. 18, no. 1, pp. 64–92, 1999.
- [19] M. Dogar and S. Srinivasa, "A framework for push-grasping in clutter," *Robotics: Science and Systems VII*, vol. 1, 2011.
- [20] J. Bohg, A. Morales, T. Asfour, and D. Kragic, "Data-driven grasp synthesis—a survey," *IEEE Transactions on Robotics*, vol. 30, no. 2, pp. 289–309, 2013.
- [21] N. C. Dafe, A. Rodriguez, R. Paolini, B. Tang, S. S. Srinivasa, M. Erdmann, M. T. Mason, I. Lundberg, H. Staab, and T. Fuhlbrügge, "Extrinsic dexterity: In-hand manipulation with external forces," in *2014 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2014, pp. 1578–1585.
- [22] A. Boularias, J. Bagnell, and A. Stentz, "Learning to manipulate unknown objects in clutter by reinforcement," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 29, no. 1, 2015.
- [23] N. Chavan-Dafle and A. Rodriguez, "Prehensile pushing: In-hand manipulation with push-primitives," in *2015 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2015, pp. 6215–6222.
- [24] L. P. Kaelbling and T. Lozano-Pérez, "Hierarchical task and motion planning in the now," in *2011 IEEE International Conference on Robotics and Automation*. IEEE, 2011, pp. 1470–1477.
- [25] S. Levine, C. Finn, T. Darrell, and P. Abbeel, "End-to-end training of deep visuomotor policies," *The Journal of Machine Learning Research*, vol. 17, no. 1, pp. 1334–1373, 2016.
- [26] J. Mahler, J. Liang, S. Niyaz, M. Laskey, R. Doan, X. Liu, J. A. Ojea, and K. Goldberg, "Dex-net 2.0: Deep learning to plan robust grasps with synthetic point clouds and analytic grasp metrics," *arXiv preprint arXiv:1703.09312*, 2017.
- [27] A. Zeng, S. Song, K.-T. Yu, E. Donlon, F. R. Hogan, M. Bauza, D. Ma, O. Taylor, M. Liu, E. Romo, *et al.*, "Robotic pick-and-place of novel objects in clutter with multi-affordance grasping and cross-domain image matching," in *2018 IEEE international conference on robotics and automation (ICRA)*. IEEE, 2018, pp. 3750–3757.
- [28] A. M. Wells, N. T. Dantam, A. Shrivastava, and L. E. Kavraki, "Learning feasibility for task and motion planning in tabletop environments," *IEEE robotics and automation letters*, vol. 4, no. 2, pp. 1255–1262, 2019.
- [29] J. E. Hopcroft, J. T. Schwartz, and M. Sharir, "On the complexity of motion planning for multiple independent objects; pspace-hardness of the 'warehouseman's problem'," *The International Journal of Robotics Research*, vol. 3, no. 4, pp. 76–88, 1984.
- [30] G. Wilfong, "Motion planning in the presence of movable obstacles," *Annals of Mathematics and Artificial Intelligence*, vol. 3, no. 1, pp. 131–150, 1991.
- [31] M. Stilman, J.-U. Schamburek, J. Kuffner, and T. Asfour, "Manipulation planning among movable obstacles," in *Proceedings 2007 IEEE international conference on robotics and automation*. IEEE, 2007, pp. 3327–3332.
- [32] C. H. Papadimitriou, "The euclidean travelling salesman problem is np-complete," *Theoretical computer science*, vol. 4, no. 3, pp. 237–244, 1977.
- [33] R. M. Karp, "Reducibility among combinatorial problems," in *Complexity of computer computations*. Springer, 1972, pp. 85–103.
- [34] S. Bereg and A. Dumitrescu, "The lifting model for reconfiguration," *Discrete & Computational Geometry*, vol. 35, no. 4, pp. 653–669, 2006.
- [35] 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.
- [36] D. Halperin, M. van Kreveld, G. Miglioli-Levy, and M. Sharir, "Space-aware reconfiguration," *arXiv preprint arXiv:2006.04402*, 2020.

- [37] L. Chang, J. R. Smith, and D. Fox, "Interactive singulation of objects from a pile," in *2012 IEEE International Conference on Robotics and Automation*. IEEE, 2012, pp. 3875–3882.
- [38] M. Laskey, J. Lee, C. Chuck, D. Gealy, W. Hsieh, F. T. Pokorny, A. D. Dragan, and K. Goldberg, "Robot grasping in clutter: Using a hierarchy of supervisors for learning from demonstrations," in *2016 IEEE International Conference on Automation Science and Engineering (CASE)*. IEEE, 2016, pp. 827–834.
- [39] A. Eitel, N. Hauff, and W. Burgard, "Learning to singulate objects using a push proposal network," in *Robotics research*. Springer, 2020, pp. 405–419.
- [40] H. Song, J. A. Haustein, W. Yuan, K. Hang, M. Y. Wang, D. Kragic, and J. A. Stork, "Multi-object rearrangement with monte carlo tree search: A case study on planar nonprehensile sorting," *arXiv:1912.07024*, 2019.
- [41] A. Zeng, S. Song, S. Welker, J. Lee, A. Rodriguez, and T. Funkhouser, "Learning synergies between pushing and grasping with self-supervised deep reinforcement learning," in *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2018, pp. 4238–4245.
- [42] 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.
- [43] A. Krontiris and K. E. Bekris, "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.
- [44] J. Díaz, J. Petit, and M. Serna, "A survey of graph layout problems," *ACM Computing Surveys (CSUR)*, vol. 34, no. 3, pp. 313–356, 2002.
- [45] M. R. Garey and D. S. Johnson, "Vertex ordering," in *Computers and Intractability: a Guide to The Theory of NP-Completeness*, 1979, pp. 199–201.
- [46] C. H. Papadimitriou, "The np-completeness of the bandwidth minimization problem," *Computing*, vol. 16, no. 3, pp. 263–270, 1976.
- [47] M. R. Garey, D. S. Johnson, and L. Stockmeyer, "Some simplified np-complete problems," in *Proceedings of the sixth annual ACM symposium on Theory of computing*, 1974, pp. 47–63.
- [48] F. Gavril, "Some np-complete problems on graphs," John Hopkins University, Baltimore, MD, 1977, pp. 91–95.
- [49] H. L. Bodlaender, J. R. Gilbert, H. Hafsteinsson, and T. Kloks, "Approximating treewidth, pathwidth, frontsize, and shortest elimination tree," *Journal of Algorithms*, vol. 18, no. 2, pp. 238–255, 1995.
- [50] T.-h. Shin, H. Oh, and S. Ha, "Minimizing buffer requirements for throughput constrained parallel execution of synchronous dataflow graph," in *16th Asia and South Pacific Design Automation Conference (ASP-DAC 2011)*. IEEE, 2011, pp. 165–170.
- [51] R. J. Lipton and R. E. Tarjan, "A separator theorem for planar graphs," *SIAM Journal on Applied Mathematics*, vol. 36, no. 2, pp. 177–189, 1979.
- [52] J. R. Gilbert, J. P. Hutchinson, and R. E. Tarjan, "A separator theorem for graphs of bounded genus," *Journal of Algorithms*, vol. 5, no. 3, pp. 391–407, 1984.
- [53] N. Alon, P. Seymour, and R. Thomas, "A separator theorem for graphs with an excluded minor and its applications," in *Proceedings of the twenty-second annual ACM symposium on Theory of computing*, 1990, pp. 293–299.
- [54] U. Elsner, *Graph partitioning-a survey*. Techn. Univ., 1997.
- [55] Y. Shiloach, "A minimum linear arrangement algorithm for undirected trees," *SIAM Journal on Computing*, vol. 8, no. 1, pp. 15–32, 1979.
- [56] D. Adolphson and T. C. Hu, "Optimal linear ordering," *SIAM Journal on Applied Mathematics*, vol. 25, no. 3, pp. 403–423, 1973.
- [57] L. M. Kirousis and C. H. Papadimitriou, "Searching and pebbling," *Theoretical Computer Science*, vol. 47, pp. 205–218, 1986.
- [58] N. G. Kinnarsley, "The vertex separation number of a graph equals its path-width," *Information Processing Letters*, vol. 42, no. 6, pp. 345–350, 1992.
- [59] M. R. Fellows and M. A. Langston, "On search decision and the efficiency of polynomial-time algorithms," in *Proceedings of the twenty-first annual ACM symposium on Theory of computing*, 1989, pp. 501–512.
- [60] 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," *arXiv preprint arXiv:2101.12241*, 2021.
- [61] L. Gurobi Optimization, "Gurobi optimizer reference manual," 2021. [Online]. Available: <http://www.gurobi.com>