

On the Utility of Buffers in Pick-n-Swap Based Lattice Rearrangement

Kai Gao

Jingjin Yu

Abstract—We investigate the utility of employing multiple buffers in solving a class of rearrangement problems with pick-n-swap manipulation primitives. In this problem, objects stored randomly in a lattice are to be sorted using a robot arm with $k \geq 1$ swap spaces or buffers, capable of holding up to k objects on its end-effector simultaneously. On the structural side, we show that the addition of each new buffer brings diminishing returns in saving the end-effector travel distance while holding the total number of pick-n-swap operations at a minimum. This is due to an interesting recursive cycle structure in random m -permutation, where the largest cycle covers over 60% of objects. On the algorithmic side, we propose fast algorithms for 1D and 2D lattice rearrangement problems that can effectively use multiple buffers to boost solution optimality. Numerical experiments demonstrate the efficiency and scalability of our methods, as well as confirm the diminishing return structure as more buffers are employed.

Introduction video: <https://youtu.be/KtBxoARGaVQ>

I. INTRODUCTION

Efficient multi-object rearrangement is an essential skill for robots to master. The quality of a rearrangement solution largely depends on how often the end-effector picks up objects and, to a lesser extent, on the end-effector travel. Whereas most existing robot end-effectors can only pick and place a single object at a time, end-effectors capable of holding multiple objects have begun to pop up (see, e.g., Fig. 1, top row). For these new end-effector types, new algorithms must be developed to exploit their capabilities fully. Toward this goal, the *pick-n-swap* primitive, which allows an end-effector to hold one object and make object swaps temporarily, is systematically studied in [1] for objects stored in regular, lattice-like storage (see, e.g., Fig. 1, bottom row). Exploring unique properties of the pick-n-swap problem on lattices, the study developed efficient algorithms for computing optimal or near-optimal solutions for minimizing the number of pick-n-swaps and end-effector travel.

In this work, we study the pick-n-swap problem further. As it is shown [1] that holding one object and performing object swaps can significantly enhance the efficiency in solving rearrangement tasks (as compared with pick-n-place without swapping capability), a natural follow-up question is whether an additional gain is possible by making the end-effector even more capable. Specifically, we want to understand the benefit, if any, of allowing the end-effector to hold multiple objects at a time. We call each unit of capacity the end-effector has in temporarily holding an object as a *buffer*, i.e.,

K. Gao and J. Yu are with the Department of Computer Science, Rutgers, the State University of New Jersey, Piscataway, NJ, USA. Emails: {kai.gao, jingjin.yu}@rutgers.edu. This work is partly supported by NSF awards IIS-1845888 and IIS-2132972, and an Amazon Research Award.




Fig. 1. [top left] A robot from Universal Robots equipped with dual Robotiq grippers. [top right] A dual OnRobot vacuum gripper setup. [bottom] Illustrations of 1D and 2D lattice rearrangement problems.

an end-effector with k buffers can temporarily hold k objects. Intuitively, having more buffers will allow more flexibility in rearrangement and thus reduce task execution time.

On the structural side, as it turns out, for rearranging randomly distributed objects stored in lattice-like storage, when the pick-n-swap time dominates the total cost, each additional buffer after the first provides increasingly diminished returns in improving the overall efficiency. Due to the observation, we can conclude that there is little to gain after having 2-3 buffers, regardless of the total number of objects. On the algorithmic side, building on a *cycle-following* algorithm proposed in [1] for a single buffer, we have developed effective and fast algorithms using dynamic programming (DP) for handling an arbitrary number of buffers for 1D lattices (i.e., objects form a single row), which is subsequently extended to 2D settings. We have also developed search-based algorithms for handling the general objective, allowing arbitrary cost trade-offs between the pick-n-swaps and end-effector travel.

More concretely, the contributions of this work are:

- We establish a tight lower bound of the expected length of the largest cycle in a random permutation. The result is supported by a simple yet elegant proof and has the potential to be of broader interest beyond the immediate scope of this work.
- For the multi-buffer pick-n-swap-based rearrangement problem on lattices, for sequentially minimizing the number of pick-n-swaps and end-effector travel, applying the aforementioned structural result, we establish that the utility of having k buffers exponentially decays with k . Quantitatively, having 3 buffers is sufficient to reap over 90% of the possible efficiency gain, implying that a more

complex end-effector design may be unnecessary.

- We propose an effective dynamic programming (DP) method for solving multi-buffer pick-n-swap problems. Our method computes high-quality solutions for 1D and 2D lattices. Extensive numerical evaluations not only confirm our method's effectiveness but also confirm that adding more buffers provides a diminished return in efficiency gain. Our DP methods may be readily extended to higher dimensions.
- We further develop an algorithm based on Monte Carlo tree search that supports arbitrary trade-offs between pick-n-swap costs and end-effector travel costs, providing a general solution to the multi-buffer pick-n-swap problem.

Organization. The rest of the paper is organized as follows. In Sec. II, we introduce previous works in various related problem domains. In Sec. III, we provide a formal definition of the problem we study and review a crucial cycle structure in single buffer lattice rearrangement. Then, in Sec. IV, we establish some structural properties on parallelizing rearrangement tasks with multiple buffers. We describe our proposed algorithmic solutions in Sec. V. The evaluation follows in Sec. VI. We conclude with Sec. VII.

II. RELATED WORKS

Autonomous robotic object rearrangement apply to both household settings [2]–[5] and logistics [6]–[10]. While rearrangement in shelf-like scenarios [2], [9]–[12] focus more on robot-object collision avoidance, the ones in the tabletop [3], [5], [8], [13]–[17] or conveyor scenario [18], [19] tend to use overhand grasps and emphasize manipulation efficiency. Similar to the tabletop setting, a recent work [1] studies rearrangement in lattices while allowing a pick-n-swap operation instead of the more standard pick-n-place operation. Rearrangement with non-prehensile actions (e.g. pushes) [8], [13], [16], [17], [20] does not require accurate contact between the robot and objects but the resulting arrangements tend to be less predictable. In contrast, rearrangement with prehensile actions (e.g. pick-n-place) [2], [3], [7], [9]–[11], [13], [14], [21]–[23] requires accurate grasping poses but more predictable manipulation enables long-horizon planning, leading to increased efficiency.

Based on different scenarios and applications, the objective of rearrangement varies, with an initial focus on seeking feasible solutions in difficult and general rearrangement problems [21], [22]. Gradually, the focus shifts to minimizing the number of manipulations to increase the system throughput [3], [7], [9], [10], [16]. Additionally, some works [14] take the traveling distance of the gripper into consideration, minimizing the total execution time. In this paper, we investigate lattice rearrangement with pick-n-swaps. Taking both the number of operations and the traveling distance of the end-effector into account, we seek to reduce the execution time of the rearrangement task.

In many research domains including robotic manipulation and operations research, capacity constraints frequently arise that add significant complexity. In rearrangement problems, a recent work [24] studies the *running buffer size*, which

is the size of the needed free space for temporary object displacement in the rearrangement task. In the field of printed circuit board assembly, manipulation with multi-head grippers [25]–[27] enables the robot arm to carry multiple objects simultaneously and saves traveling time of the arm. Most of these works tend to assume that the picking locations or the placing positions are located close to each other. In our setting, locations of picking and placing are more entangled, making the problem challenging to solve.

III. PRELIMINARIES

A. Problem Formulation

Consider a d -dimensional lattice with size $m_1 \times m_2 \times \dots \times m_d$, with each cell storing exactly one object. The storage yields an arrangement \mathcal{A} of the labeled objects. A robot arm is tasked to move objects from a start arrangement \mathcal{A}_s to a desired goal arrangement \mathcal{A}_g with pick-n-swap operations. Each pick-n-swap can be represented as a 3-tuple (ℓ, i, j) , where the robot moves to the picking (lattice) position ℓ , swap a held object i with the object j stored in the cell. It is possible that $i = \varepsilon$ or $j = \varepsilon$, i.e., empty. In this paper, we work with a k -buffer arm that can carry at most k objects when traveling over the lattice. A rearrangement plan $P = \{p_0, p_1, p_2, \dots, p_N, p_{N+1}\}$ is a sequence of pick-n-swaps moving objects from \mathcal{A}_s to \mathcal{A}_g with $p_0 = p_{N+1} = (r, \varepsilon, \varepsilon)$, where r is the rest position of the robot arm. Correspondingly, the robot arm starts from a rest position r in the lattice, executes pick-n-swaps one by one in P , and moves back to r after finishing the task. We assume the rest position in the lattice is $(1, \dots, 1)$.

Similar to [1], we evaluate the quality of the rearrangement plan based on the execution time. Assume that the time cost of one pick-n-swap is c_p and that spent by the end-effector to travel a unit distance (the distance between adjacent lattice cells) is c_t , then the cost function can be represented as:

$$J_T(P) = c_p N + c_t \sum_{i=0}^N \text{dist}(p_i.\ell, p_{i+1}.\ell) \quad (1)$$

where $\text{dist}(p_i.\ell, p_{i+1}.\ell)$ measures the Euclidean distance between two successive pick-n-swap positions in the lattice. In practical scenarios, c_p tends to be much larger than c_t as object picking and placing involves prehensile manipulation. Therefore, while our paper does examine the general objective, the discussion focuses primarily on solutions that first minimize the number of pick-n-swaps.

Based on the definitions above, the problem investigated in this paper can be formulated as follows.

Problem III.1 (Multi-Buffer Lattice Rearrangement (MBLR)). *Given the current arrangement \mathcal{A}_s , and the goal arrangement \mathcal{A}_g , compute a rearrangement plan P for a k -buffer robot arm minimizing $J_T(P)$.*

Specifically, we denote MBLR in d -dimensional lattice as d -MBLR. It is assumed that \mathcal{A}_s is a random arrangement of the objects and \mathcal{A}_g is an ordered arrangement.

The proof of Theorem IV.1 is rather complex. In this work, using rudimentary analysis, we establish a lower bound fairly close to the Golomb-Dickman Constant.

Theorem IV.2. *The expected size of the largest cycle in a random permutation of m letters is lower bounded by $0.607m$ asymptotically.*

Proof sketch. Let P_m^i , $1 \leq i \leq m$, denote the probability that the largest cycle of a random permutation on m letters has a size i . The key idea behind the technical proof is to compute P_m^i for $\lfloor m/3 \rfloor < i \leq m$ and tally all the entries, which gives the result. We refer the readers to [29] for the complete and somewhat involved proof. \square

We note that Thm. IV.1 applies recursively since if we remove the largest cycle from a random permutation, the rest of the permutation is still random. This suggests that, for reasonably large m , the largest $k = 3$ cycles occupies over $(1 - (1 - 0.624)^3) > 94.6\%$ of all elements.

Alternatively, this can also be observed by estimating the number of cycles in a random permutation of m letters, which equals $H_m - 1$ [1], in which $H_m \approx \ln m$ is the m -th harmonic number. For two different numbers m_1 and m_2 , $m_1 \geq m_2$, the number of cycles differs by $H_{m_1} - H_{m_2} \approx \ln \frac{m_1}{m_2}$. If we set $H_{m_1} - H_{m_2} = 3$, then $m_2 \approx 0.05m_1$. In other words, three cycles will occupy about 95% of elements, agreeing with the estimate using Thm. IV.1.

C. Task Parallelism with Multiple Buffers

The analysis so far shows the solution structure of MBLR when the number of pick-n-swaps is minimized. In this subsection, we will discuss task parallelization based on *cycle groups*, each of which is a group of cycles with range overlaps in 1-MBLR [1]. In high-dimensional instances, all cycles belong to a single cycle group. When paralleling the rearrangement of objects i and j , there are three different cases: (1) i and j are in the same cycle; (2) i and j are in different cycle groups; (3) i and j are in the same cycle group but in different cycles.

When rearranging objects i and j , which are in the same permutation cycle, according to Prop. IV.1, if the rearrangement of i and j is parallelized, then the number of pick-n-swaps in the plan exceeds the minimum.

When objects i and j are in different cycle groups, task parallelization does not reduce the total distance since the range of cycles in different cycle groups does not overlap. Additionally, for 1-MBLR, we prove the independence of cycle groups in the rearrangement plan computation.

Proposition IV.2. *In 1-MBLR, rearranging objects from different cycle groups can be considered separately.*

See [29] for the detailed technical proof.

Finally, for objects in the same cycle group but in different cycles, using the running example from Fig. 3, we show parallel rearrangement using multiple buffers reduces the traveling cost without increasing the number of pick-n-swaps. In the example, the cycles (574) and (86) can be

handled in parallel. If the robot arm has two buffers, it can place object 7 on its way to cell 8, and place object 6 on its way back to cell 4. In this way, the traveling distance is reduced by 2 (total: $14 \rightarrow 12$). The new plan using two buffers is illustrated in Fig. 4.



Fig. 4. Using a second buffer (indicated by the orange arrows), parallelism can be realized by arranging (574) and (86) in the same cycle group.

V. METHODS

In this section, we propose fast algorithms for MBLR. Based on properties established in Sec. IV, we introduce dynamic programming (DP) based algorithms for computing MBLR solutions with the minimum number of pick-n-swaps that seek to minimize end-effector travel, for one- and two-dimensional lattices. Our methods also apply to higher dimensions. Additionally, for MBLR with the general objective function (1), we propose an algorithm based on Monte Carlo Tree Search (MCTS) for minimizing it.

A. DP based Algorithm for 1-MBLR Minimizing the number of Pick-n-Swaps (DP-MBLR)

Algo. 1 shows the pipeline of the DP algorithm for 1-MBLR. Since rearrangement plans for different cycle groups can be computed independently (Prop. IV.2), without loss of optimality, we first decompose the instance into cycle groups to simplify the computation (Line 1). In Line 3, rearrangement tasks of objects in the same cycle will be assigned to the same buffer in order to minimize the number of pick-n-swaps (Prop. IV.1). Based on the assignment, we compute the optimal rearrangement plan minimizing the number of pick-n-swaps and traveling distance for each buffer using the exact algorithm for single-buffer [1] (Line 5). After that, we merge the task sequences to obtain a high-quality rearrangement plan for the cycle group. Finally, we concatenate plans for cycle groups with the cycle-switching algorithm [1] to minimize the additional cost when traveling among cycle groups.

Algorithm 1: DP-MBLR for 1-MBLR

Input : \mathcal{A}_s : Start arrangement in a 1-dimensional lattice arrangement.
 k : the number of buffers

Output: P : The rearrangement plan for the instance

- 1 $G_1, G_2, \dots, G_n \leftarrow \text{CycleGroupDecomposition}(\mathcal{A}_s)$
 - 2 **for** $1 \leq i \leq n$ **do**
 - 3 $C_1, C_2, \dots, C_k \leftarrow \text{CycleAssignment}(G)$
 - 4 **for** $1 \leq j \leq k$ **do**
 - 5 $s_j \leftarrow \text{SingleBufferRearrangement}(C_j)$
 - 6 $P_i \leftarrow \text{TaskSequencesMerge}(\{s_1, \dots, s_k\})$
 - 7 $P \leftarrow \text{CycleGroupSwitching}(\{P_1, \dots, P_n\})$
-

1) *Task Assignment to Buffers*: In this step, we assign permutation cycles to k buffers maximizing the minimum number of pick-n-swaps assigned to each buffer. This problem is equivalent to Bottleneck Multiple Subset Sum Problem (B-MSSP) [30]. For example, consider an MBLR instance with two buffers for four permutation cycles $c_1 - c_4$, which contain 10, 5, 5, 2 objects respectively. In this case, the needed pick-n-swaps for each cycle is 11, 6, 6, 3. The optimal task assignment to buffers is $(c_1, c_4), (c_2, c_3)$, with the minimum number of pick-n-swaps to be 12. Given k buffers and c cycles, when $c \leq k$, we assign one cycle to each of the first c buffers. Otherwise, we solve this problem with an integer linear programming (ILP) model [30].

2) *Merging Task Sequences of Buffers*: For each cycle group, we propose a DP-based method for minimizing the total travel cost given the task sequences. Let T be a k -tuple (t_1, t_2, \dots, t_k) representing a subproblem where only the first t_i tasks are executed for buffer i . Let f be the index of the buffer that the last executed task belongs to. For each combination of T and f ($T[f] \geq 1$), $DP[T, f].cost$ counts the traveling distance from the moment when the arm leaves the rest position to the moment when the arm finishes the $T[f]^{th}$ task of buffer f . It is computed with the following equation.

$$DP[T, f].cost = \min_{\substack{T'=(t_1, \dots, t_{f-1}, \dots, t_k) \\ T'[f'] \geq 1}} (DP[T', f'].cost + dist((T', f'), (T, f)))$$

where $(T', f')((T, f))$ represents the position of the $T'[f']^{th}$ task of buffer f' (the $T[f]^{th}$ task of buffer f). Let T_c be the k -tuple representing the original problem. The minimum travel cost is derived based on the equation below:

$$\min_{T_c[f] \geq 1} DP[T_c, f].cost + dist((T_c, f), r)$$

where r is the rest position. In a k -buffer scenario, the computation time of the dynamic programming approach is $O(m^k)$, which is polynomial in the number of objects m .

For the running example from Fig. 3, using 2 buffers, DP-MBLR first decomposes objects into cycles groups $G_1 = \{(231)\}$ and $G_2 = \{(574), (86)\}$. For G_1 , the robot arm moves objects to goal positions in the ordering of 2, 3, and 1 with buffer 1. For G_2 , DP-MBLR assigns one cycle to each buffer: $C_1 = \{(574)\}$, $C_2 = \{(86)\}$. Note that according to Prop. IV.2, even though buffer 1 has claimed the tasks in G_1 , it does not affect the total cost no matter which cycle in G_2 is assigned to buffer 1. After DP-MBLR computes a single buffer plan in G_2 , it merges the rearrangement plans inside G_2 and between cycle groups as shown in Fig. 5.




Fig. 5. DP-MBLR on the running example (Fig. 3) with two buffers. This yields the plan shown in Fig. 4, reducing traveling distance from 16 to 14.

B. Extension of DP-MBLR to High Dimensions

DP-MBLR can be easily extended to d -MBLR ($d > 1$) with a three-step pipeline. First, assign cycles to buffers with the ILP solver in Sec. V-A.1. Second, for each buffer, we compute a plan with an asymptotically distance-optimal algorithm in the single buffer scenario [1]. Finally, we merge task sequences of buffers with dynamic programming as Sec. V-A.2.

C. Monte Carlo Tree Search for General Objective

To decide on the next pick-n-swap in optimizing the general form of $J_T(P)$ in Eq. (1), our MCTS planner maintains a search tree where each node represents a rearrangement state containing three parts of information: object arrangement, objects in buffers, and the robot position. Each edge represents a pick-n-swap. For each node s , let $c(s)$ be the cost from the root node to s with the same c_p and c_t in $J_T(P)$.

For each state, we reduce the action set with two rules without loss of optimality:

- Only execute pick-n-swaps at cells c where the target object $A_g[c]$ is not placed inside.
- When executing a pick-n-swap at a cell i , if i is in buffers, then place i .

Additionally, for 1-MBLR, we notice that it is inefficient to bypass cell i without placing object i if it is in buffers. Based on the observation, the next pick-n-swap position is limited to the range between the closest goal positions on the left and right (if any) of objects in the buffers. Take the instance in Fig. 3 as an example, if the arm is located at 7 and the objects in buffers are 1 and 4, then the range of the next pick-n-swap location is $[4, 8]$.

In the selection stage of a node s , the MCTS planner chooses an action a based on the following formula in the same spirit as the upper confidence bound (UCB):

$$\operatorname{argmin}_a \left(\frac{w(f(s, a))}{n(f(s, a))} - C \sqrt{\frac{\log(n(s))}{n(f(s, a))}} \right) \quad (3)$$

where $f(s, a)$ is the child node of s after action a , $w(s)$ and $n(s)$ are the total cost and times of visits at s .

VI. EVALUATION

In this section, we show experiments on the lattice rearrangement problem in two different scenarios, MBLR with minimum pick-n-swaps and MBLR with general cost functions, with a focus on the former. The experiments are executed on an Intel® Xeon® CPU at 3.00GHz. Each data point is the average of 30 test cases except for unfinished trials, if any, given a time limit of 600 seconds for each test case. Given a random m -permutation, the workspace dimension of 1-MBLR and 2-MBLR in the experiments are $1 \times m$ and $\lceil \sqrt{m} \rceil \times \lceil \sqrt{m} \rceil$ respectively (see Fig. 1).

A. MBLR with Minimum Number of Pick-n-Swaps

For 1-MBLR with the minimum number of pick-n-swaps, we first compare DP-MBLR with a brute force exact algorithm OPT-MBLR where we enumerate all possible cycle

assignment and cycle-switching options to get the minimum distance cost. Fig. 6[left] shows the computation time. When $m = 20, 25$, OPT-MBLR fails in some test cases. In these cases, the computation time is recorded as 600 seconds. Fig. 6[right] shows the distance cost of computed solutions as proportions of the shortest distance in the single buffer setting. We observe that 2-3 buffers provide significant savings in end-effector travel. The results suggest that compared with the exact algorithm for multiple buffers, DP-MBLR is much more scalable with little loss of optimality.




Fig. 6. Comparison between DP-MBLR and OPT-MBLR. [left] Computation time comparison. [right] Distance cost of computed solutions as proportions of the shortest distance in the single buffer setting.

Fig. 7 shows the performance of DP-MBLR on 1-MBLR with 2-5 buffers. In terms of computation time, the dominant component of DP-MBLR is TaskSequenceMerge, which is polynomial in m in the worst case. As a result, DP-MBLR is highly scalable. Regarding optimality, compared with the single buffer setting, lattice rearrangement with two buffers saves travel distance by 25%–30%. For 3 buffers, the saving is about 30%–35%. Adding more buffers beyond three, however, provides negligible additional benefits as predicted by our theoretical analysis. That is, as the largest three cycles cover more than 93% of objects and each cycle requires at most one unique buffer, additional buffers can barely make any difference.




Fig. 7. Performance of DP-MBLR on 1-MBLR instances with different numbers of buffers and objects. [left] Computation time (secs). [right] Distance cost as a proportion of that in optimal solutions with one buffer.

Fig. 8 shows the performance of DP-MBLR in 2-MBLR. Compared with single buffer rearrangement, 2-MBLR solutions with two buffers save distance cost by 15%–20% and those with three or more buffers save distance cost by 20%–25%. While there is slightly less efficiency gain using multiple buffers in 2-MBLR, the diminishing return from additional buffers is consistent with the results in 1-MBLR. The results show the efficiency of DP-MBLR in 2-MBLR and the effectiveness of robot arms with 2 or 3 buffers in two-dimensional lattices.




Fig. 8. Performance of DP-MBLR in 2-MBLR instances with different number of objects. [left] Computation time (secs). [right] Distance cost as a proportion of that in optimal solutions with one buffer.

B. MBLR with General Objective Functions

To evaluate the performance of MCTS-MBLR, we solve MBLR with two different cost functions where the execution time of pick-n-swaps is not dominant in the rearrangement process. When $c_p : c_t = 1 : 1$, both the number of pick-n-swaps and the traveling distance affect the cost in a computed solution. When $c_p : c_t = 1 : 10^5$, the traveling cost is dominant. Fig. 9 shows the cost of solutions computed by MCTS-MBLR as a proportion of that in the optimal single-buffer solutions. As the number of objects increases, rearrangement with multiple buffers has the potential to save more cost on traveling. However, in the meanwhile, the computation time for each step of the decision decreases. Therefore, when $m = 20$, less saving on traveling is seen by using multiple buffers.




Fig. 9. Cost of solutions computed by MCTS-MBLR as a proportion of that in optimal single-buffer solutions. [Left] $c_p : c_t = 1 : 1$, [Right] $c_p : c_t = 1 : 10^5$.

VII. CONCLUSION

In this work, we study the problem of employing multiple buffers to perform pick-n-swap-based rearrangement on lattices where the robot arm can hold multiple objects simultaneously when traveling. In analyzing the effectiveness of additional swap spaces, we prove multiple structural results on parallelizing rearrangement tasks with multiple buffers. Specifically, using the Golomb-Dickman constant, we establish an exponentially diminishing return in each additional buffer when the number of pick-n-swaps is minimized. We also provide a rudimentary proof of how a tight lower bound of the Golomb-Dickman constant can be obtained. Using the obtained properties, we propose an efficient algorithm minimizing the number of pick-n-swaps based on dynamic programming. For general cost functions, we also propose a Monte Carlo Tree Search-based method. Experiments show the effectiveness of using multiple buffers compared with using one buffer in lattice rearrangement problems and further verify the properties we obtain.

REFERENCES

- [1] J. Yu, "Rearrangement on lattices with pic-n-swaps: Optimality structures and efficient algorithms," in *Robotics: Sciences and Systems*, 2021.
- [2] K. Wada, S. James, and A. J. Davison, "Reorientbot: Learning object reorientation for specific-posed placement," *arXiv preprint arXiv:2202.11092*, 2022.
- [3] K. Gao, D. Lau, B. Huang, K. E. Bekris, and J. Yu, "Fast high-quality tabletop rearrangement in bounded workspace," *arXiv preprint arXiv:2110.12325*, 2021.
- [4] A. Zeng, P. Florence, J. Tompson, S. Welker, J. Chien, M. Attarian, T. Armstrong, I. Krasin, D. Duong, V. Sindhwani, *et al.*, "Transporter networks: Rearranging the visual world for robotic manipulation," in *Conference on Robot Learning*. PMLR, 2021, pp. 726–747.
- [5] M. Danielczuk, A. Mousavian, C. Eppner, and D. Fox, "Object rearrangement using learned implicit collision functions," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2021, pp. 6010–6017.
- [6] M. Szegedy and J. Yu, "On rearrangement of items stored in stacks," in *International Workshop on the Algorithmic Foundations of Robotics*. Springer, 2020, pp. 518–533.
- [7] 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 *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2021, pp. 6621–6627.
- [8] S. Han, B. Huang, C. Song, S. W. Feng, M. Xu, A. Boularias, and J. Yu, "Toward fully automated metal recycling using computer vision and non-prehensile manipulation," in *Proceedings of the 17th IEEE International Conference on Automation Science and Engineering*, 2021.
- [9] 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.
- [10] R. Wang, K. Gao, J. Yu, and K. Bekris, "Lazy rearrangement planning in confined spaces," *arXiv preprint arXiv:2203.10379*, 2022.
- [11] 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.
- [12] E. R. Vieira, D. Nakhimovich, K. Gao, R. Wang, J. Yu, and K. E. Bekris, "Persistent homology for effective non-prehensile manipulation," *arXiv preprint arXiv:2202.02937*, 2022.
- [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] S. D. Han, N. M. Stiffler, A. Kroutiris, 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.
- [15] R. Shome, K. Solovey, J. Yu, K. Bekris, and D. Halperin, "Fast, high-quality dual-arm rearrangement in synchronous, monotone tabletop setups," in *International Workshop on the Algorithmic Foundations of Robotics*. Springer, 2018, pp. 778–795.
- [16] B. Huang, S. D. Han, J. Yu, and A. Boularias, "Visual foresight trees for object retrieval from clutter with nonprehensile rearrangement," *IEEE Robotics and Automation Letters*, vol. 7, no. 1, pp. 231–238, 2021.
- [17] 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," in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2020, pp. 9433–9440.
- [18] B. Carlisle, K. Goldberg, A. Rao, and J. Wiegley, "A pivoting gripper for feeding industrial parts," in *Proceedings of the 1994 IEEE International Conference on Robotics and Automation*. IEEE, 1994, pp. 1650–1655.
- [19] S. D. Han, S. W. Feng, and J. Yu, "Toward fast and optimal robotic pick-and-place on a moving conveyor," *IEEE Robotics and Automation Letters*, vol. 5, no. 2, pp. 446–453, 2019.
- [20] B. Huang, S. D. Han, A. Boularias, and J. Yu, "Dipn: Deep interaction prediction network with application to clutter removal," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2021, pp. 4694–4701.
- [21] A. Kroutiris and K. E. Bekris, "Dealing with difficult instances of object rearrangement," in *Robotics: Science and Systems*, vol. 1123, 2015.
- [22] —, "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.
- [23] 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.
- [24] K. Gao, S. W. Feng, and J. Yu, "On minimizing the number of running buffers for tabletop rearrangement," in *Robotics: Sciences and Systems*, 2021.
- [25] M. Grunow, H.-O. Günther, M. Schleusener, and I. O. Yilmaz, "Operations planning for collect-and-place machines in pcb assembly," *Computers & Industrial Engineering*, vol. 47, no. 4, pp. 409–429, 2004.
- [26] M. Moghaddam and S. Y. Nof, "Parallelism of pick-and-place operations by multi-gripper robotic arms," *Robotics and Computer-Integrated Manufacturing*, vol. 42, pp. 135–146, 2016.
- [27] J. Luo and J. Liu, "An milp model and clustering heuristics for led assembly optimisation on high-speed hybrid pick-and-place machines," *International Journal of Production Research*, vol. 52, no. 4, pp. 1016–1031, 2014.
- [28] S. W. Golomb, "Random permutations," *Bulletin of the American Mathematical Society*, vol. 70, p. 747, 1964.
- [29] K. Gao and J. Yu, "On the utility of buffers in pick-n-swap based lattice rearrangement," 2022. [Online]. Available: <https://arxiv.org/pdf/2209.05390.pdf>
- [30] A. Caprara, H. Kellerer, and U. Pferschy, "The multiple subset sum problem," *SIAM Journal on Optimization*, vol. 11, no. 2, pp. 308–319, 2000.