

# Towards computing low-makespan solutions for multi-arm multi-task planning problems

Valentin N. Hartmann<sup>1,2</sup>, Marc Toussaint<sup>2</sup>

<sup>1</sup>Visus, University of Stuttgart, Germany, valentin.hartmann@ipvs.uni-stuttgart.de

<sup>2</sup>Learning and Intelligent Systems Lab, TU Berlin, Germany

## Abstract

We propose an approach to find low-makespan solutions to multi-robot multi-task planning problems in environments where robots block each other from completing tasks simultaneously.

We introduce a formulation of the problem that allows for an approach based on greedy descent with random restarts for generation of the task assignment and task sequence. We then use a multi-agent path planner to evaluate the makespan of a given assignment and sequence. The planner decomposes the problem into multiple simple subproblems that only contain a single robot and a single task, and can thus be solved quickly to produce a solution for a fixed task sequence. The solutions to the subproblems are then combined to form a valid solution to the original problem.

We showcase the approach on robotic stippling and robotic bin picking with up to 4 robot arms. The makespan of the solutions found by our algorithm are up to 30% lower compared to a greedy approach.

## 1 Introduction

In many common applications, a given set of tasks need to be completed. Examples include, e.g., unloading a dishwasher, laser welding, bin picking, inspection, or unloading a delivery truck. Such problems could often benefit from parallelization of tasks by using multiple robots. This requires deciding which robot works on which tasks, and in which order the tasks should be completed to minimize the total needed time. Both these things are nontrivial problems.

For a single robot, similar problems were previously solved in laser welding (Saha et al. (2006); Kovács (2016)) by formulating them as traveling salesman problem with neighborhoods. Problems involving multiple mobile robots have been tackled by treating the problem as search on graphs (Yu and LaValle (2016); Ma and Koenig (2016)) or with conflict based search (Hönig et al. (2018)).

A common assumption in planning for multi-agent systems in task and motion planning is synchronicity of tasks, i.e., each agent starts and finishes its task at the same times (Pan et al. (2021); Toumi et al. (2022); Zhang et al. (2022)). Planning with this synchronicity assumption does not work well if tasks have different lengths and leads to idle times for the robots if they block each other from starting a task.

Chen et al. (2022) tackle finding low-makespan solutions for assembly scenarios, i.e. scenarios where the order of tasks is highly constrained using mixed integer linear programs to compute the task assignment. However, the resulting computation times are in the thousands of hours, and the algorithm has only been demonstrated for up to three robots.

To tackle multi robot task assignment for offshore missions, temporal planning is used in Carreno et al. (2020), but the robots tend not to interfere with each other much, whereas in our setting, the space is much more congested.

Finally, we refer to Stern (2019) for a survey on multi robot motion planning. We particularly want to point out priority based search (Ma et al. (2019)), which is most similar to our approach. However, we do not introduce the priority on the path-level, but on the tasks for each robot.

To compute a good solution to the combined multi-agent task assignment and path planning problem, we decompose the task assignment and the multi-agent planning problems. Particularly, we serialize the multi-robot task sequence to obtain a sequence that defines precedence constraints and assigns a robot to each task. For this sequence, we then plan a path to determine its makespan. This enables us to proceed with a standard greedy search with random restarts to find an improved sequence.

We evaluate this approach on robotic stippling and robotic bin-picking with multiple robots and compare it to a greedy approach and an approach using only a single arm.

## 2 Problem Formulation & Notation

We consider settings with  $R$  robots, indexed by  $r \in \{1, \dots, R\}$ , with configuration spaces  $\mathcal{C}^r$ , and  $M$  objects, indexed by  $m \in \{1, \dots, M\}$  with poses in  $SE(3)$ . We use  $\mathcal{X} = \mathcal{C}^1 \times \dots \times \mathcal{C}^R \times SE(3)^M$  to denote the configuration space and  $x : [0, t_{\text{final}}] \rightarrow \mathcal{X}$  to denote a path through this configuration space. We use  $s^r = (s_1^r, \dots, s_{N_r}^r)$  to denote the sequence of tasks that robot  $r$  needs to fulfil, and  $s = s^1 \times \dots \times s^R \in \mathbb{S}$  for the combined sequence containing every robot. The set  $\mathbb{S}$  consists of all valid sequences that complete all  $N$  tasks. The tasks in  $s^r$  must be completed by robot  $r$  in the order in which they are contained in  $s^r$ , i.e.,  $T_i^r < T_{i+1}^r$ , where  $T_i^r$  is the time at which robot  $r$  finishes task  $s_i^r$ . Similar to  $s$ , we use  $T$  to denote all finishing times  $T = T^1 \times \dots \times T^R$ , with




Figure 1: Example of a serialization  $\hat{s}$  of a sequence  $s$ .

$T^r = (T_1^r, \dots, T_{N_r}^r)$ . We want to solve

$$\min_{s, x, T} \max_r T_{N_r}^r \quad (1a)$$

$$\text{s.t. } x(0) = x_0 \quad (1b)$$

$$\forall r g(x^r(t), s_i^r) \leq 0 \quad t \in [T_{i-1}^r, T_i^r] \quad \forall i \in (1, \dots, N_r) \quad (1c)$$

$$\forall r h(x^r(T_i^r), s_i^r) = 0 \quad \forall i \in (1, \dots, N_r) \quad (1d)$$

$$x(t) \in \mathcal{C}_{\text{free}} \quad \forall t \in [0, \max_r T_{N_r}^r] \quad (1e)$$

$$s \in \mathbb{S}, \quad (1f)$$

i.e., want to find a sequence of task-assignments and the corresponding paths that minimize the makespan.

In (1), (1c) are the constraints on the path that are active in the current task (e.g., robot  $r$  is moving object  $m$ ), (1d) constrains the path of robot  $r$  such that task  $s_i^r$  is completed at time  $T_i^r$  via the constraint function  $h$  (e.g., a pose-constraint on object  $m$ ), (1e) ensures that the found path is collision free, and (1f) ensures that the sequence of tasks is a valid sequence that fulfills all tasks (and possibly constrains the order of tasks). This is a Logic-Geometric program (Toussaint et al. (2018)).

This problem could be formulated as a multi-agent traveling salesman problem, and could theoretically be solved using mixed integer programming. However, the path planning problem for a multi-robot system itself is hard, and a purely optimization based approach is thus unsuitable, as many of the local optima are infeasible due to violation of (1e)

In the following, we assume that:

- Each task can be handled by one robot, and
- There exists a ‘resting’ pose for each robot in which no other robot is blocked from completing its tasks, where the robot can return to after competing its task.

The first assumption is made in this work for simplicity, but can be removed by enabling planning for multiple robots for a task (as done in Hartmann et al. (2022)). The second assumption enables completeness of the planning subroutine.

### 3 Method

To make the problem tractable, we decompose the problem into ‘generating a robot-task sequence’ and ‘evaluating the sequence’. It is then possible to use search methods, e.g., greedy descent to improve an initial sequence.

While  $s$  determines the sequence of tasks per robot  $r$ , it does not determine their relative ordering across robots. Instead of searching over  $s$ , we search over a *serialized* sequence  $\hat{s}$  (see Fig. 1 for an illustration). If  $s$  is partially ordered, i.e., the tasks have dependencies between each other, such a serialization can be obtained via topological search.

Given  $s$ , such a serialized assignment is defined as  $\hat{s} = (s_{i(j)}^{r_{k(j)}}, r_{k(j)})_{j=0}^N$ , where  $k(j)$  maps the index  $j$  of a

---

#### Algorithm 1: Greedy descent with restarts for multi robot task planning

---

```

1  $x_{\min} \leftarrow \emptyset$ ,  $t_{\min} \leftarrow \infty$ 
2 for  $j \leftarrow 0$  to  $\text{max\_outer\_iter}$  do
3    $\hat{s} \leftarrow \text{InitializeSequence}()$ 
4    $x_{\text{inner, min}}, t_{\text{inner, min}} \leftarrow \text{PlanGivenSequence}(\hat{s})$ 
5   for  $i \leftarrow 0$  to  $\text{max\_inner\_iter}$  do
6      $\hat{s}_{\text{cand}} \leftarrow \text{GenerateNeighbor}(\hat{s})$ 
7      $x, t \leftarrow \text{PlanGivenSequence}(\hat{s}_{\text{cand}})$ 
8     if  $t < t_{\text{inner, min}}$ 
9        $x_{\text{inner, min}}, t_{\text{inner, min}} \leftarrow x, t$ 
10     $\hat{s} \leftarrow \hat{s}_{\text{cand}}$ 
11    if  $t < t_{\min}$ 
12       $x_{\min}, t_{\min} \leftarrow x, t$ 
13 return  $x_{\min}, t_{\min}$ 

```

---



---

#### Algorithm 2: GenerateNeighbor( $\hat{s}$ )

---

```

1  $r \leftarrow \text{rnd}(0, 1)$ 
2 if  $r < \frac{1}{3}$  return  $\text{SwapRobot}(\hat{s})$ 
3 else if  $\frac{1}{3} \leq r < \frac{2}{3}$  return  $\text{SwapRandomElements}(\hat{s})$ 
4 else return  $\text{ReverseSubtour}(\hat{s})$ 

```

---

task in sequence  $\hat{s}$  to the index of the robot that the task is assigned to, and  $i(j)$  maps the index  $j$  of the sequence  $\hat{s}$  to the index of the sequence  $s^r$ . The ordering in the serialized sequence imposes that task  $s_{i(j)}^{r_{k(j)}}$  must be finished before task  $s_{i(j+1)}^{r_{k(j+1)}}$ , i.e.,  $T_{i(j)}^{r_{k(j)}} < T_{i(j+1)}^{r_{k(j+1)}}$ . For any sequence  $s$ , many possible serializations  $\hat{s}$  exist. As we search over  $\hat{s}$ , we can cover all possible precedence constraints.

We then compute a valid path for the multi-agent system for a serialized sequence using an approach similar to Hartmann et al. (2022), i.e., by planning the tasks in  $\hat{s}$  one-by-one, and treating the previously computed paths as fixed.

#### 3.1 Time embedded multirobot path planning

To compute a plan for a given sequence, we build on the work from Hartmann et al. (2022): In the path planning subroutine, we decompose the full problem, and solve sub-problems with a single robot and a single task only. That is, we sequentially plan the robot/task-pairings in  $\hat{s}$ , while treating previously planned paths as fixed. The precedence constraints encoded in  $\hat{s}$  are respected in this planning step, i.e., we enforce that the tasks are fulfilled in the order specified in  $\hat{s}$  by introducing a lower bound on the finishing time of the plan, i.e.,  $T_i^r < T_{i+1}^r$ .

After completion of a task, an ‘escape’ path is planned to the pose that does not block any other robot. Being able to move to this non-blocking pose guarantees that there is always a valid path for subsequent tasks and robots, since there is a time in the future where all robots that are not involved in the current task are at their non-blocking pose. This escape path is discarded when planning the paths for the next task for the robot. Figure 2 illustrates how the paths are planned using time-embedded path planning.

We use ST-RRT\* (Grothe et al. (2022)) which allows the specification of a maximum velocity for time-embedded

**Algorithm 3:** PlanGivenSequence( $\hat{s}$ )

---

```

1 sol  $\leftarrow$  {}
2 for  $i \leftarrow 0$  to len( $\hat{s}$ ) do
3    $(x_i, t_i) \leftarrow$  PlanTaskForRobot( $\hat{s}_i$ , sol)
4   sol  $\leftarrow (x_i, t_i)$ 
5  $(x, T) \leftarrow$  GetPathsAndMakespan(sol)
6 return  $(x, T)$ 

```

---




Figure 2: Illustration of a single planning step in the sequential planning for planning task  $j$ : a pick and place task consisting of a ‘pick’ and a ‘place’ action.

path planning, and run it until a solution was found, and continue optimizing the path for a fixed number of iterations. This results in a good path, but typically not in the optimal one due to the restriction of the runtime. We post-process the resulting path via shortcutting (Hauser and Ng-Thow-Hing (2010) adapted to the time-embedded case), and smooth it with an optimizer to minimize accelerations. This results in paths that can be directly executed on a real robot.

### 3.2 Algorithms & Practical details

We give pseudocode for the proposed approach in Algorithms 1 to 3. We implement a greedy descent with restarts, i.e., the greedy descent in the inner loop is repeated `max_outer_iter` times (Alg. 1, Line 2). The serialized assignment is randomly initialized (Alg. 1, Line 3) before planning for the given sequence (Alg. 1, Line 4). In the inner loop, a neighbouring sequence is generated (Alg. 1, Line 6), and a plan is generated. If the makespan is lower than the current best makespan that was found in the inner loop so far, the best makespan of the inner loop, the corresponding path, and the current best serialized sequence are updated (Alg. 1, Line 8). If the makespan is lower than the best makespan achieved so far, the best overall makespan and corresponding path are updated. Planning the path for a given sequence  $\hat{s}$  proceeds iteratively for each task (Alg. 1, Line 11) until all tasks are planned. Previous solutions are respected in the planning process by storing the previous solutions (Alg. 1, Line 4), and considering them as constraints for subsequent tasks.

We make the following implementation choices, and small modification to the core idea presented before to speed up the search and planning process:

**Generating the initial sequence** We generate the initial sequence  $\hat{s}$  by going through the robots in round-robin, and assigning a random feasible task to the current robot. This results in all the agents alternating between each other.

**Altering a sequence** Greedy descent requires a function that generates a ‘neighbouring’ sequence from a given se-




Figure 3: **Left and middle:** Positions of the tip of the pen for the stippling scenarios. **Right:** Bin picking scenario with the goal positions of the objects in low opacity.

quence<sup>1</sup>. We do this with equal probability by either (i) Swapping elements of the sequence (Alg. 2, Line 3), (ii) Reversing subtours (Alg. 2, Line 4), or (iii) Swapping robots (Alg. 2, Line 2). These operations are inspired by typical approaches used in greedy searches on the traveling salesman problem (Abdoun, Abouchabaka, and Tajani (2012)).

**Early stopping of motion planning** A lower bound on the makespan of the remaining sequence can be computed by using the euclidean distance between poses and optimistically assume movement at maximum velocity to obtain a lower bound on the finishing time of a task. This lower bound can then be used to avoid the costly computation of paths that can not improve the current solution.

**Caching of subsequences** If we already computed a valid plan for a sequence  $\hat{s}_1$ , and are now planning for a sequence  $\hat{s}_2$ , and the first  $m$  elements of two sequences  $\hat{s}_1$  and  $\hat{s}_2$  are the same, we can reuse the first  $m$  plans that were computed for this subsequence, and only need to compute the plans for the remaining tasks.

## 4 Experiments

We demonstrate our approach on the example of robotic stippling (where a task corresponds to ‘go to pose’)<sup>2</sup>, and multi-arm bin picking (where one task is ‘pick obj, place obj’, i.e., two actions). The stippling and the bin-picking setting are shown in Fig. 3.

We analyze several versions of the stippling scenarios: for the grid, we compare a two and a four robot version; for the LIS logo, we consider a ‘small’ version, where dots can not be made with two arms at the same time, and that can not be well parallelized, and a ‘large’ version, where some dots can be made simultaneously.

We compare our algorithm to a baseline of a single robot, and an initialization that alternates the robot-tasks in a greedy manner, i.e., always chooses the task next that has the minimum estimated distance to the current robot pose.

<sup>1</sup>If precedence constraints are present in the original sequence  $s$  (e.g. when stacking boxes, the lowest box must be placed first), these constraints could be enforced by, e.g., rejection sampling.

<sup>2</sup>‘Robotic stippling’ uses robot arms equipped with pens to make dots on paper.

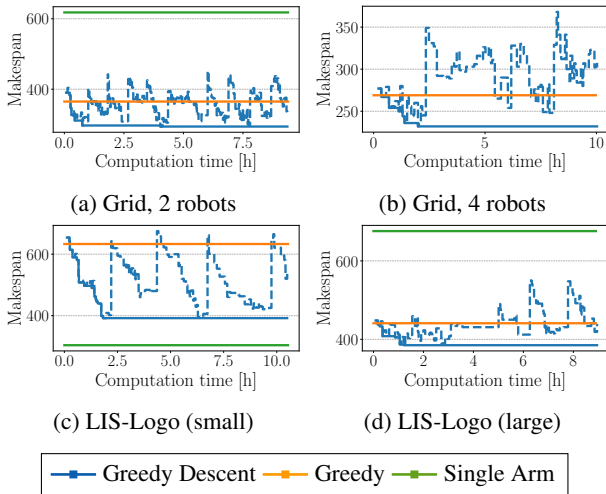


Figure 4: Makespan of current sequence (dashed line), and minimum makespan so far (solid line) at a computation time on the stippling scenarios. Both greedy, and single arm are not anytime, so do not improve the initial solution. For the four arm grid setting, there is no comparisons to the single arm baseline, as not all dots can be made by all robots.

Videos of the execution of some generated paths in simulation and on real robots can be found at <https://vhartmann.com/seq-opt>.

#### 4.1 Results

Figure 4 shows the evolution of the makespan over the computation time for the stippling scenarios. For the four-arm scenario, the single-arm baseline is not available, as not every dot can be made by every arm, and the baseline-sequence is thus not feasible.

The average planning time for an individual sequence was roughly 180s for the grid with two robots (11s per task), 543s for the grid with four robots (33s per task), 288s for the small logo (9.6s per task) and 185s for the large logo (6s per task).

Figure 5 shows a sequence of snapshots from two robots executing an optimized sequence from the bin-picking scenario, and Fig. 6 shows the evolution of the makespan. Here, the average planning time for an individual sequence was roughly 110s (consisting of 6 tasks, each containing 2 actions), i.e. 18s per task, respectively 9s per action.

In these experiments, it can be seen that the improvement of using multiple arms is more pronounced in the case where tasks can be completed in parallel (i.e. in the grid, the large LIS logo, and the bin-picking scenarios) compared to the case where the arms completely block each other (small LIS logo), where the single-arm version is quicker than the one using two arms. In this case, our algorithm finds a solution that minimizes the ‘switches’ between robots as well, but does not completely reach it due to the limited number of iterations of the inner loop of the search.

Comparing the results of the grid with two and four arms, it can be seen that the planning time increases considerably, as the space becomes more occupied by the robots, and the

planning queries become more complex (and thus slower). The algorithm is still able to handle this, and finds a good solution, even though many fewer iterations are done in the same computation time.

## 5 Discussion & Limitations

Both multi-agent task assignment and planning are hard problems. We showed the algorithm on multiple scenarios with up to four robots. The setting with four robots is extremely crowded, and planning times are thus high, which leads to less iterations of the search algorithm. The approach we show here works well in practice for our problems, even though it is not time-optimal.

To achieve time-optimal plans, one would have to take into account dynamics, and consider an approach that enables joint planning of multiple agents. Settings where many robots are trying to occupy the same space (as the four arm stippling-scenario) could benefit most from such a joint approach.

We noted that the lower bound that we compute tends to be optimistic, and does not help much in reducing computation time. It might be worth the effort to compute a more accurate lower bound for the time it takes to complete a task.

We currently only consider single poses to fulfill a task, and do not allow ‘neighborhoods’, i.e., the set of poses that fulfill the constraints associated with a task. We expect this to be interesting for future work.

Further, we want to investigate more informed ways of generating candidate sequences that take into account the geometry of the scene, and the tasks that should be done by each robot.

Finally, we restricted our problem setting here to tasks that can be solved by a single robot. In the future, we want to investigate more complex task and motion planning settings, e.g. tasks that require cooperation, or non-prehensile manipulation.

## 6 Conclusion

We presented an approach that enables the optimization of the task sequence and assignment in multi-agent systems. The approach can lead to improvements of the makespan of up to 50% over a single robot doing all tasks, and up to a 30% improvement over a greedy alternating assignment.

We believe that the work can benefit from more research on how to speed up the multi-agent path planning, and from exploring other search approaches that account for the computationally expensive motion planning.

## Acknowledgements

The authors would like to thank Wolfgang Hönig for helpful comments on the draft of this work. We would also like to thank the anonymous reviewers for their helpful comments.

The research has been supported by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany’s Excellence Strategy – EXC 2120/1 – 390831618, and the DAAD.




Figure 5: Snapshots of two robots working through a task sequence from the bin picking scenario.




Figure 6: Makespan of current sequence (dashed line), and minimum makespan so far (solid line) at a computation time on the bin-picking scenario.

## References

- Abdoun, O.; Abouchabaka, J.; and Tajani, C. 2012. Analyzing the performance of mutation operators to solve the travelling salesman problem. *arXiv preprint arXiv:1203.3099*.
- Carreno, Y.; Pairet, È.; Petillot, Y.; and Petrick, R. P. 2020. Task allocation strategy for heterogeneous robot teams in offshore missions. In *Proc. of the 19th Int. Conf. on Autonomous Agents and MultiAgent Systems*, 222–230.
- Chen, J.; Li, J.; Huang, Y.; Garrett, C.; Sun, D.; Fan, C.; Hofmann, A.; Mueller, C.; Koenig, S.; and Williams, B. C. 2022. Cooperative task and motion planning for multi-arm assembly systems. *arXiv preprint arXiv:2203.02475*.
- Grothe, F.; Hartmann, V. N.; Orthey, A.; and Toussaint, M. 2022. ST-RRT\*: Asymptotically-optimal bidirectional motion planning through space-time. In *Proc. of the IEEE Int. Conf. on Robotics and Automation (ICRA)*.
- Hartmann, V. N.; Orthey, A.; Driess, D.; Oguz, O. S.; and Toussaint, M. 2022. Long-horizon multi-robot rearrangement planning for construction assembly. *IEEE Transactions on Robotics*.
- Hauser, K., and Ng-Thow-Hing, V. 2010. Fast smoothing of manipulator trajectories using optimal bounded-acceleration shortcuts. In *Proc. of the IEEE Int. Conf. on Robotics and Automation (ICRA)*, 2493–2498. IEEE.
- Hönig, W.; Kiesel, S.; Tinka, A.; Durham, J.; and Ayanian, N. 2018. Conflict-based search with optimal task assignment. In *Proc. of the Int. Joint Conference on Autonomous Agents and Multiagent Systems*.
- Kovács, A. 2016. Integrated task sequencing and path planning for robotic remote laser welding. *International journal of production research* 54(4):1210–1224.
- Ma, H., and Koenig, S. 2016. Optimal target assignment and path finding for teams of agents. In *Proc. Int. Conf. on Autonomous Agents & Multiagent Systems*, AAMAS ’16, 1144–1152.
- Ma, H.; Harabor, D.; Stuckey, P. J.; Li, J.; and Koenig, S. 2019. Searching with consistent prioritization for multi-agent path finding. AAAI Press.
- Pan, T.; Wells, A. M.; Shome, R.; and Kavraki, L. E. 2021. A general task and motion planning framework for multiple manipulators. In *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 3168–3174. IEEE.
- Saha, M.; Roughgarden, T.; Latombe, J.-C.; and Sánchez-Ante, G. 2006. Planning tours of robotic arms among partitioned goals. *The International Journal of Robotics Research* 25(3):207–223.
- Stern, R. 2019. Multi-agent path finding—an overview. *Artificial Intelligence: 5th RAAI Summer School, Dolgoprudny, Russia, July 4–7, 2019, Tutorial Lectures* 96–115.
- Toumi, I.; Orthey, A.; von Rohr, A.; and Vien, N. A. 2022. Multi-arm bin-picking in real-time: A combined task and motion planning approach. *arXiv preprint arXiv:2211.11089*.
- Toussaint, M.; Allen, K. R.; Smith, K. A.; and Tenenbaum, J. B. 2018. Differentiable physics and stable modes for tool-use and manipulation planning. In *Proc. of Robotics: Science and Systems (R:SS)*.
- Yu, J., and LaValle, S. M. 2016. Optimal multirobot path planning on graphs: Complete algorithms and effective heuristics. *IEEE Transactions on Robotics* 32(5):1163–1177.
- Zhang, H.; Chan, S.-H.; Zhong, J.; Li, J.; Koenig, S.; and Nikolaidis, S. 2022. A mip-based approach for multi-robot geometric task-and-motion planning. In *Conf. on Automation Science and Engineering*, 2102–2109. IEEE.