



HAL
open science

Shadoks Approach to Low-Makespan Coordinated Motion Planning (CG Challenge)

Loïc Crombez, Guilherme D. da Fonseca, Yan Gérard, Aldo Gonzalez-Lorenzo, Pascal Lafourcade, Luc Libralesso

► To cite this version:

Loïc Crombez, Guilherme D. da Fonseca, Yan Gérard, Aldo Gonzalez-Lorenzo, Pascal Lafourcade, et al.. Shadoks Approach to Low-Makespan Coordinated Motion Planning (CG Challenge). 37th International Symposium on Computational Geometry, SoCG 2021, June 7-11, 2021, Buffalo, NY, USA (Virtual Conference), Jun 2021, Buffalo, United States. 10.4230/LIPIcs.SoCG.2021.63 . hal-03485486

HAL Id: hal-03485486

<https://hal.science/hal-03485486v1>

Submitted on 28 Oct 2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

Shadoks Approach to Low-Makespan Coordinated Motion Planning

Loïc Crombez   




Université Clermont-Auvergne and LIMOS, France

Guilherme D. da Fonseca   

Aix Marseille Université and LIS, France

Yan Gerard   

Université Clermont-Auvergne and LIMOS, France

Aldo Gonzalez-Lorenzo   

Aix Marseille Université and LIS, France

Pascal Lafourcade   

Université Clermont-Auvergne and LIMOS, France

Luc Libralesso   

Université Clermont-Auvergne and LIMOS, France

Abstract

This paper describes the heuristics used by the **Shadoks**¹ team for the CG:SHOP 2021 challenge. This year's problem is to coordinate the motion of multiple robots in order to reach their targets without collisions and minimizing the makespan. Using the heuristics outlined in this paper, our team won first place with the best solution to 202 out of 203 instances and optimal solutions to at least 105 of them.

2012 ACM Subject Classification Theory of computation → Computational geometry; Computing methodologies → Motion path planning

Keywords and phrases heuristics, motion planning, digital geometry, shortest path

Category CG challenge

Supplementary Material vimeo.com/515040997 and github.com/gfonsecabr/shadoks-robots

Funding *Loïc Crombez*: This work has been sponsored by the French government research program “Investissements d’Avenir” through the IDEX-ISITE initiative 16-IDEX-0001 (CAP 20-25).

Guilherme D. da Fonseca: This work is supported by the French ANR PRC grant ADDS (ANR-19-CE48-0005).

Yan Gerard: This work is supported by the French ANR PRC grants ADDS (ANR-19-CE48-0005) and ACTIVmap (ANR-19-CE19-0005).

Pascal Lafourcade: This work is supported by the French ANR PRC grant MobiS5 (ANR-18-CE39-0019), DECRYPT (ANR-18-CE39-0007), and SEVERITAS (ANR-20-CE39-0005).

Luc Libralesso: This work is supported by the French ANR PRC grant DECRYPT (ANR-18-CE39-0007).

Acknowledgements We would like to thank H el ene Toussaint, Rapha el Amato, Boris Lonjon, and William Guyot-L enat from LIMOS, as well as the Qarma and TALEP teams and Manuel Bertrand from LIS, who continue to make the computational resources of the LIMOS and LIS clusters available to our research. We would also like to thank the challenge organizers and other competitors for their time, feedback, and making this whole event possible.

¹ The team name comes from the animated television series Les Shadoks https://en.wikipedia.org/wiki/Les_Shadoks.

1 Introduction

We explain some heuristics used by the **Shadoks** team to solve the CG:SHOP 2021 challenge that considers a coordinated motion planning problem in the two-dimensional grid \mathbb{Z}^2 . The goal is to move a set of n labeled unit squares called *robots* between given start and target grid cells without collisions.

More formally, the input consists of a set of *obstacles* O and a set of n *robots* $R = \{r_1, \dots, r_n\}$. Each obstacle is a lattice point and each robot r_i is a pair of lattice points s_i, d_i respectively called *start* and *target*. A path P_i of length m is a sequence of $m + 1$ lattice points $p_i(t)$ for a time $t = 0, \dots, m$. A solution of makespan m is an assignment of *paths* P_i of length m to each robot r_i satisfying the following constraints.

- (i) $p_i(0) = s_i$ and $p_i(m) = d_i$ for $1 \leq i \leq n$,
- (ii) $\|p_i(t) - p_i(t-1)\| \leq 1$ for $1 \leq i \leq n$ and $1 \leq t \leq m$,
- (iii) $p_i(t) \notin O$ for $1 \leq i \leq n$ and $0 \leq t \leq m$,
- (iv) (*collision constraint*) $p_i(t) \neq p_j(t)$ for $i \neq j$ and $0 \leq t \leq m$, and
- (v) (*overlap constraint*) if $p_i(t) = p_j(t-1)$ then $p_i(t) - p_i(t-1) = p_j(t) - p_j(t-1)$.

The objective of the problem is to minimize the makespan² m . A trivial lower bound is obtained by ignoring constraints (iv) and (v) and finding shortest paths. Since the problem is symmetric with respect to the time, we may always exchange start-target positions and reverse the paths, keeping the best solution found. For more details about the problem, see the overview [9] and the related paper [8].

The challenge CG:SHOP 2021 provided 203 instances containing between 10 and 9000 robots, out of which 202 of our solutions were the best ones among all the 17 teams who participated. To our surprise, we succeed in finding 105 solutions that match the trivial lower bound. Our strategy consists of two steps, presented in Section 2 and 3: finding a feasible solution and reducing its makespan.

Literature review The *multi-agent path finding* problem (MAPF) has been well studied over the last 20 years. This problem occurs in many industrial applications that involve agents that have to reach destinations without colliding with each other [16]. For instance, in automated warehouses [15], autonomous vehicles, and robotics [2]. Well-studied approaches include search-based solvers (for instance: HCA* [18]) that make plans agent by agent according to a predefined agent order. When an agent is scheduled, it “reserves” times and locations and the algorithm schedules the next agent. They can also plan all n robots at the same time, thus having for each time-step up to 5^n possibilities (for instance: Enhanced Partial Expansion A* [13]). We may also cite rule-based solvers that identify scenarios and apply rules to move agents, for instance the push-and-rotate algorithm [7]. Some algorithms model the MAPF problem using network flows, integer linear programs [24], or SAT instances [21, 22].

One of the most popular methods is CBS (*conflict-based search*) [17]. This method starts by solving a relaxation of the original problem in which agent collisions are ignored. This relaxation is relatively easy to solve as it consists of running a shortest-path algorithm for each agent. If the resulting plan contains a time t and coordinate c where two agents r_1, r_2 collide, then the algorithm forbids either r_1 or r_2 from being at the coordinate c at time t . This results in a search tree that is explored until it is depleted (thus finding the optimal

² The challenge also considered the objective of minimizing the sum of the distances, but we did not optimize our solutions for this version.

solution for the problem). This method has been reported to achieve excellent results and multiple improvements have been made. We may cite for instance, addition of heuristics [10], merge-and-restart and conflict prioritization [3], some sub-optimal variants [1], and some techniques such as a branch-and-cut-and-price algorithm [14]. For more information about the multi-agent path finding problem, we refer the reader to surveys about the MAPF variants and instances [20] and about the MAPF solvers [11].

At the beginning of the challenge, we tried some of the aforementioned approaches (notably CBS) to solve the challenge instances. To our surprise, CBS did not perform well. Indeed, the challenge instances are much denser than the ones in the literature. These instances are usually sparse in the number of agents (there are from 2 to 120 agents placed in grids with over 100,000 cells in these instances, where instances of the challenge contain hundreds or thousands of agents placed in grids never larger than 100×100). This structural difference has a dramatic effect on the performance of CBS and in our experiments, CBS fails to find solutions for most of the challenge instances (with the exception of some very small ones).

The remainder of the paper is organized as follows. In Section 2 we consider the problem of obtaining feasible solutions of moderate makespan. These solutions are optimized later on, with the techniques described in Section 3. Details on implementing the algorithms are described in Section 4. Section 5 describes the results we obtained for some challenge instances and presents a comparison with the strategies used by other teams.

2 Initial Solutions

Feasibility is guaranteed for the challenge instances since the number of obstacles is finite and every start and target are located in the unbounded region of space. In this section, we show how to obtain feasible solutions with a moderate makespan. We divide the heuristics in two categories. In Section 2.1, we compute the solution one step at a time, considering multiple robots simultaneously. In Section 2.2, we compute the solutions one robot at a time.

The heuristics of the first category are not guaranteed to find a solution, but when they do they often find solutions of lower makespan than those of the second category. The algorithms of the second category are guaranteed to find a solution, but the resulting makespan may potentially be high.

2.1 Step by Step Computation

The problem of finding a solution for coordinated motion planning in a given number of steps can be modeled as an Integer Linear Problem (ILP) or equivalently as a SAT problem (see [21, 22] and references therein). While applying such an approach is intractable even for small instances, it can be adapted to find an initial solution. The general idea of the *Greedy* solver is to plan only a small number k of steps for the robots such that the overall distance to the targets decreases as much as possible and repeat until reaching the targets.

Our ILP model considers a Boolean variable $x_{r,P}$ for each robot r and for each possible path P of length k starting at the position of the robot. Constraints of having one and only one path per robot and avoiding obstacles and collisions between robots are easily expressed as linear inequalities. The objective function we maximize is the sum of all the variables with weight

$$\text{weight}(r, P) = \left(\delta_r(p(0)) - \delta_r(p(k)) \right) \cdot \left((\delta_r(p(0)))^2 + 1 \right),$$

where $p(0)$ and $p(k)$ are the first and the last positions of path P and $\delta_r(p)$ is the obstacle-avoiding distance from a point p to the target of robot r . The first factor encourages the solution to push the robots towards their targets, since it is better to get closer to target. The second factor prioritizes moving robots that are farther from their targets and we add one so that robots that are already at their target position are encouraged to remain there.

In practice, we set $k = 3$ and we only perform the first step of the planned moves so that the robots can *anticipate* the moves of the other robots. Using the CPLEX [6] library to solve these problems, we can handle instances with up to roughly 200 robots. Note that this Greedy algorithm is not guaranteed to find a solution, and it fails to solve instances with *corridors* such as `universe_bg_000`.

2.2 Robot by Robot Computation

The algorithms in this section compute the solution one robot at a time using an A* search. The search happens in 3-dimensional space where each robot state has integer coordinates of the form (x, y, t) for position coordinates x, y and time t . There are 5 possible *movements*, all of which increase t by one unit. One movement keeps the position x, y unchanged, while the other 4 movements increment or decrement one of the two coordinates. A movement is feasible if it does not violate any of the problem constraints, considering the current path of the other robots.

We refer to the *bounding box* as an integer axis-aligned rectangular region containing all the start, target positions and obstacles inside its strict interior (not on the boundary). Given a set of obstacles and a bounding box, the *depth* of a position p is the minimum obstacle-avoiding distance from p to a position outside the bounding box. One may note that, being in its strict interior, the start and target positions have depth at least one. Indeed the subsequent algorithms require a way for a robot to go around these positions. A bounding box guarantees that there exists at least one such way: its border, i.e. the positions of depth 0. The minimum depth of start and target positions of a bounding box might become an issue for performance; 1 to 3, depending on the algorithm, seem to be good compromises in our experience.

All algorithms in this section are based on a *storage network* N . A storage network is a set N of positions outside a predetermined bounding box such that for every position p in N , there exists a path that avoids all other positions of N and goes from p to some point in the bounding box. Each robot r_i is assigned to a distinct element of N , called the *storage* of r_i .

Initially, we set the path of each robot to be stationary at the start position. We sort the robots by *increasing start depth* and for each robot in order, we use A* search to find the shortest path from start to storage, replacing the previous stationary path. The order by which the robots are sorted guarantees that such a path exists.

After finding paths from start to storage for every robot, we proceed to the next phase of the algorithm. We now sort the robots by *decreasing target depth*. Again, the order of the robots guarantees that a path from storage to target exists. However, we do not compute such a path. Instead, we compute a path from start to target directly, whose existence is guaranteed by the existence of a path from start to storage and another one from storage to target. The following paragraphs describe the design of four different storage networks.

Cross. In the *Cross* strategy, we define the storage network N as the set of columns of even x coordinate lying directly above or below the bounding box and the set of rows of even y coordinate lying directly to the left or right of the bounding box, hence the name Cross. Then, we compute a maximal cardinality matching between the robots and N . We




■ **Figure 1** Cross storage network for the `small_free_016` instance colored based on start and target locations, respectively.

tried both minimum-weight matching and greedy matchings, minimizing a weight function that considers the distance from start to storage as well as the distance from storage to target. In the greedy matching version, robots are assigned a storage ordered by decreasing start-to-target distance. The result is represented in Figure 1.

Cootie Catcher. The previous strategy works very well for small or sparse instances. However, the different directions of the flow of robots from start to storage make the solutions inefficient for large dense instances. The *Cootie Catcher* strategy computes the storage using only the start location, in order to better exploit parallel movement of the robots. The storage network shape consisting of four diamonds is presented in Figure 2. For instances without obstacles, the strategy is guaranteed to find a path from start to storage using at most $w/2 + \mathcal{O}(1)$ steps, where w is the largest bounding box side. Surprisingly, this strategy also works well for many instances with obstacles.


Dichotomy. The weakness of the previous method is that robots may be assigned storage in a location that is opposite to the direction from start to target. Furthermore, the parallel



■ **Figure 2** Cootie Catcher storage network for the `small_free_016` instance colored based on start and target locations, respectively.

movement of the robots make it unlikely that a robot will be able to take any significant shortcuts before it reaches the storage. In order to exploit parallel movements while taking the target location into consideration, we developed the *Dichotomy* strategy. The strategy only works for instances without obstacles.

We translate the coordinate system so that the origin is the center of the bounding box. The robots are partitioned into two sets called *left side* and *right side* according to the sign of the target location's x -coordinate. Left-side robots are assigned storage with positive x -coordinate while right-side robots are assigned storage with negative x -coordinate, as represented in Figure 3.



■ **Figure 3** Dichotomy storage network for the `small_free_016` instance colored based on start and target locations, respectively.

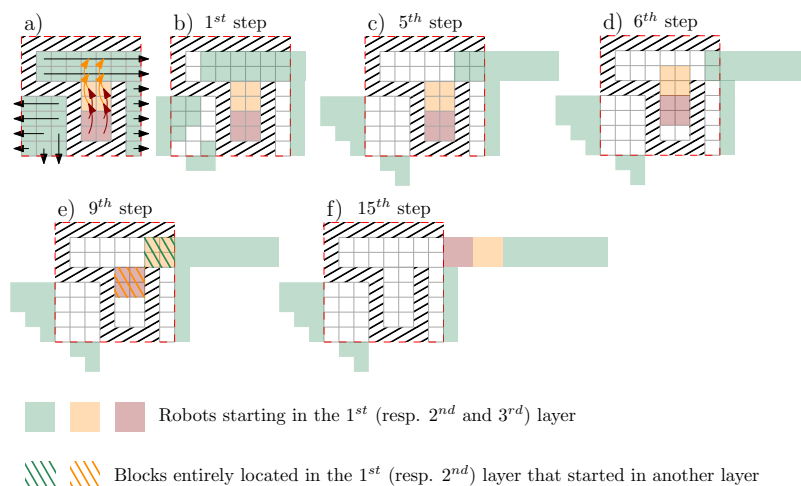
The algorithm performs the following steps, described only for the robots with non-negative y -coordinate for simplicity, as the other half is analogous.

1. Each robot goes up from start position (x, y) to position $(x, 2y)$.
2. If the robot target is on the right side, the robot moves up one more row. At this point, the even y rows contain left-side robots and the odd y rows contain the right-side robots.
3. If a right-side (resp. left-side) robot is still inside the bounding box, then it moves to the right (left) as far as needed to leave enough space for the other robots to its left on the same row to move out of the bounding box. Otherwise, a right-side (left-side) robot moves right (left) in order to leave enough space for the robots on the same row to move to a position of positive (negative) x -coordinate.

Going from start to storage takes at most $3w/2 + \mathcal{O}(1)$ movements for a $w \times w$ bounding box. Instead of sorting the robots by decreasing target depth as usual, we sort the robots by absolute value of the target x -coordinate and then determine the paths from start to target using A* as usual.

Escape. This strategy focuses on instances with obstacles, especially on dense instances where the obstacles create bottlenecks. The goal of the *Escape* strategy is to clear the bounding box as quickly as possible. To do so, we move the robots by blocks as large as possible making efficient use of parallel movements. The Escape strategy defines layers inside the bounding box. Robots located in the first layer will move in a straight line to reach the outside of the bounding box. No intersection is allowed between the path taken by robots located in the first layer. Then, the second layer is defined. It consists in blocks adjacent to the first layer, ideally as large as possible, that will move in a straight line into the first layer.

All robots located in the same block will move towards the first block in parallel motion. Either all robots from a same block are set in motion at once, or none are. Then, in the same way, a third layer is defined, consisting of blocks that will move into the first or second layer. Layers are added until covering the location of every robot as represented in Figure 4. We used a naive algorithm to define the layers, and partially redefined them by hand for the most complicated instances and the unsatisfying results. Outside the bounding box, robots emerging from three columns are stored on only two of these as can be seen in Figure 5.




■ **Figure 4** Escape strategy. a) In green (resp. orange and red), the first (resp. second and third) layer. The orange arrows show where the second layer will move, the red arrows show the same for the third layer. b-c) The robots in the first layer are moving, but the robots of the second and third layer are still stuck. d) The robots from the second layer are now free to move towards the first layer, this also allows the robots from the third layer to move. e) The robots that used to be located in the third layer, are now located in the second layer and are waiting for their path to be clear of robots. f) Final disposition outside the bounding box.

3 Improving Solutions

In this section we discuss the two heuristics that we used to reduce the makespan of a given feasible solution. The first heuristic makes local changes to the solution, which remains feasible throughout the process, and possibly reduces the makespan. The second heuristic destroys the feasibility of the solution and either finds another solution of reduced makespan, or no feasible solution at all. Throughout, let m be the makespan of the input solution.

Feasible Optimizer. The idea of the *Feasible Optimizer* is the following. We iteratively remove the path of a robot r from the solution, and then use the A* algorithm to find a new (hopefully different) path for r . The A* algorithm may be tuned in several ways to produce different paths, and we do so in such a way that the makespan of the solution never increases and also that a robot is only allowed to move at time m if it already did so in the original path. This way, not only the makespan but also the number of robots moving at time m never increase. Next, we list some examples on how to modify the A* search.

- Find the path from start to target that reaches the target as quickly as possible but breaks ties using the sum of random weights given to each grid cell the robot passes through.



■ **Figure 5** Escape storage network for the `medium_007` and `buffalo_003` instances, respectively.

- Reversing the direction of time and then finding a path from target to start that reaches the start as quickly as possible. In the original time direction, that means that the robot will remain at the start for as long as possible.
- In the reversed case, force the robot to stay at target for a certain number of steps.

Conflict Optimizer. The previous optimization strategy may take very long to reduce the makespan. Next, we describe a more aggressive approach that leaves the feasible solution space and works far better than we expected. The algorithm uses a modified A* search that allows for a robot to go over another robot's path, which we call a *conflict*. We start by creating a *queue* with all the robots that move at makespan time m . While the queue is not empty, we repeat the following procedure for a robot r popped from the front of the queue. Let $q(r)$ be the number of times a robot r has been popped out of the queue. We define the *weight* of a robot r as $1 + (q(r))^2$.

1. Erase r 's path.
2. Find a path for r from start to target that arrives no later than time $m - 1$ and minimizes the weighted sum of conflicting robots.
3. Add all conflicting robots to the queue.

For sparse instances, the Conflict Optimizer can even be used to compute solutions from scratch by choosing an initial makespan and putting all the robots in the queue.

4 Algorithm Engineering

In this section, we describe different techniques used to efficiently implement and apply the previously described heuristics. All heuristics have been executed multiple times extensively using randomization whenever possible. Furthermore, different rotations have been applied, as well as reversing the start and the target of the instances.

Multiple executions were used to produce over ten thousand solution files total. All solution files were saved with a timestamp on the file name. That allowed us to find initial solutions that would optimize better. Even though we only optimized for makespan, this large volume of solutions allowed us to obtain solutions with a sufficiently low sum of the

distances to obtain the third place in that category. Developing tools to efficiently organize and view all these solutions were an important part of the team strategy.

The heuristics have been coded in C++, most of which are available publicly on github. The tools have been coded in python. We executed the code on several Linux machines, both personal computers and high performance computing clusters at the LIS and LIMOS labs.

The A* search is in the heart of most of our heuristics. Hence, a lot of work has been done to improve its performance. Sometimes we used deterministic A*, breaking ties by the coordinates of the position, but more often we used a randomized A* algorithm where ties are broken by the sum of the weight of the positions in the path, which are assigned randomly. The A* algorithm needs a distance function as a lower bound and a collision detection, which are described next.

Distance queries The A* algorithm is guided by a lower bound to the distance to target. In the case without obstacles the lower bound we used is simply the L_1 distance, which can be calculated in $\mathcal{O}(1)$ time. While this lower bound is still valid for instances with a set O of obstacles, it is inefficient because it does not take the obstacles into account. Instead, we used the obstacle-avoiding L_1 distance.


Calculating the obstacle-avoiding L_1 distance from scratch is a slow process. Since this computation happens many times during the execution of our heuristics, it is essential to be able to compute it quickly. To this purpose, we need to use a data structure to compute the distance query(p) from a query point p to a target (in our case, given at preprocessing time). Existing data structures for the problem [4, 5] seem hard to implement. Instead, we designed a simple data structure that takes $\mathcal{O}(\log w)$ query time for obstacles inside a square of side w . The storage requirement may potentially be close to w^2 , but in our case it's significantly less, generally close to $\mathcal{O}(|O|)$.

9	8	7	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20
8	7	6	5	∞	7	8	9	10	11	12	13	14	15	16	17	18	19
7	6	5	4	∞	6	7	8	9	10	11	12	13	14	15	16	17	18
6	5	4	3	∞	5	6	7	8	9	10	11	12	13	14	15	16	17
5	4	3	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16
4	3	2	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
3	2	1	0	1	2	3	4	5	∞	∞	∞	∞	∞	∞	14	15	16
4	3	2	1	2	3	4	5	∞	∞	∞	19	18	17	16	15	16	17
5	4	∞	∞	∞	4	5	6	∞	∞	∞	18	19	18	17	16	17	18
6	5	∞	∞	∞	5	6	7	∞	∞	∞	17	18	∞	∞	∞	∞	19
7	6	7	8	7	6	7	8	∞	∞	∞	16	17	∞	∞	∞	∞	20
8	7	8	9	8	7	8	9	∞	∞	∞	15	16	∞	∞	∞	∞	21
9	8	9	10	9	8	9	10	11	12	13	14	15	∞	∞	∞	∞	22
10	9	10	11	10	9	10	11	12	13	14	15	16	17	18	19	20	21

■ **Figure 6** Distance with obstacles. Only yellow positions are stored.

Given two consecutive points (x, y) , $(x + 1, y)$ on the same line we have $\text{query}(x, y) - \text{query}(x + 1, y) \in \{-1, 0, 1\}$. Furthermore, if x is to the left (resp., right) of the bounding box, then $\text{query}(x, y) - \text{query}(x + 1, y) = 1$ (resp. $= -1$). Hence, for each line y in the bounding box we only store the points (x, y) such that $\text{query}(x, y) \neq (\text{query}(x - 1, y) + \text{query}(x + 1, y))/2$, as shown in Figure 6. All the remaining queries for line y can be calculated by interpolating or extrapolating these stored values, which can be located in $\mathcal{O}(\log w)$ time using binary search on a sorted vector.

Queries for a point (x, y) above or below the bounding box are answered by using the closest line of the bounding box and the fact that $\text{query}(x, y) = \text{query}(x, y') + |y - y'|$.



■ **Figure 7** Improved makespan over computation time using the Conflict Optimizer.

Collision detection Fast collision detection is a key point to the performance of the A* algorithm. We used an internal storage Hopscotch hash table implemented by Thibaut Goetghebuer-Planchon and distributed under the MIT license [12]. Given a position and time, we stored the robot in that position (or the list of robots in that position for the Conflict Optimizer). Hence, collision detection reduces to a small number of hash table lookups.

Avoiding Conflict Optimizer stalls The Conflict Optimizer is arguably the most significant contribution of this work. However, it may stall at sub-optimal solutions. To reduce this problem, we may use the following approaches. (i) Reverse start and target as well as the paths in the solution. (ii) Use the Feasible Optimizer to shuffle the solution. (iii) Use randomized paths in the A* search. (iv) Insert the robots that are inserted simultaneously in the queue using a random order.

5 Results

Tables 1 and 2 show the makespan obtained using different heuristics on some selected challenge instances and the makespan lower bound. The Feasible Optimizer column corresponds to the best optimization it obtained starting from different solutions. The Conflict Optimizer column corresponds to the optimization of the solution obtained by the Feasible Optimizer. Figure 7 shows the improvement obtained by the Conflict Optimizer over a little more than one hour of execution. Near the end of the challenge, some solutions kept improving very slowly with the Conflict Optimizer. Instances with a few thousand robots like `sun_007`, `clouds_008`, and `large_free_007` were consistently giving 1 unit of makespan improvement for every 10 to 20 hours of computation through several weeks.

Comparison with other teams The two other teams UNIST [23] and gitastrophe [19] on the podium of the CG:SHOP 2021 challenge used a two-phase strategy similar to ours: first compute an initial solution and then optimize it.

In general, the initial solution is computed through a storage network of the robots outside the bounding box of obstacles, start, and target positions. The existence of a solution is guaranteed by using the depth of the start and target positions. We noticed that gitastrophe used the same trick we did: first compute a partial solution going from start

instance	n	w	initial solution				optimizer		lower bound
			Gree.	Cross	Coot.	Dich.	Feas.	Conf.	
small_free_002	40	10	17	22	27	22	17	15	15
small_free_003	70	10	20	31	27	26	20	16	14
small_free_010	200	20	34	46	54	45	33	32	32
small_free_015	280	20	.	60	68	65	51	40	32
small_free_016	320	20	63	68	77	68	60	47	36
medium_free_007	630	30	148	89	103	95	81	60	52
medium_free_009	800	40	93	97	124	109	81	71	71
medium_free_012	1000	50	.	114	125	127	96	94	94
microbes_004	1250	50	.	132	159	135	125	91	91
buffalo_free_003	1440	60	.	149	165	158	125	87	78
london_night_005	1875	50	.	179	190	173	157	124	92
universe_bg_005	2000	50	.	194	198	177	173	141	82
galaxy_c2_008	3000	100	.	198	258	234	168	163	163
large_free_004	3938	75	.	274	276	256	240	204	127
large_free_005	5000	100	.	260	316	293	252	184	184
large_free_007	6000	100	.	297	343	325	295	236	189
sun_009	7500	100	.	424	395	361	354	345	187
large_free_009	9000	100	.	514	440	391	378	374	182

■ **Table 1** Makespan of different heuristics for selected instances without obstacles.

instance	n	w	initial solution				optimizer		lower bound
			Gree.	Cross	Coot.	Esca.	Feas.	Conf.	
small_005	63	10	27	28	32	37	25	20	18
sun_000	143	20	32	39	46	61	29	27	27
small_011	183	20	56	60	70	67	48	40	37
small_016	276	20	.	67	72	79	57	43	36
medium_005	407	30	.	119	110	106	94	74	58
london_night_002	825	50	.	149	162	165	142	94	84
microbes_002	958	50	.	111	135	173	97	89	89
clouds_001	912	50	.	117	138	159	94	83	83
medium_014	1165	40	.	180	161	180	161	151	73
algae_004	1113	50	.	139	160	191	121	84	79
buffalo_004	1404	60	.	136	164	195	120	104	104
large_003	1906	100	.	172	224	250	154	154	154
large_004	2034	100	.	431	391	381	.	381	185
large_005	3223	75	.	398	310	317	.	299	141
universe_bg_007	3820	100	.	224	289	323	202	184	184
large_007	4706	100	.	753	497	491	497	471	215
microbes_008	5643	100	.	329	359	425	322	279	188
algae_009	7311	100	.	500	439	441	.	421	176
large_009	8595	100	.	398	387	566	.	352	176

■ **Table 2** Makespan of different heuristics for selected instances with obstacles.

to storage and then compute new paths going from start to target since the existence of such paths is guaranteed. `gitastrophe` also used a minimum weight matching to assign each robot to a storage position. The main difference between the teams during this phase is in the choice of the storage network (points within pairwise L_∞ distance at least 2 for `gitastrophe` and points having even coordinates for UNIST). Among the different storage networks that we used (Cross, Cootie Catcher, Dichotomy, Escape), we noticed that none of the four is always the best one.

For the optimization of an initial solution, there are again some similarities, but differences are more significant. The standard strategy is to randomly remove the paths of a sample of robots before recomputing them with the hope of an improvement. `gitastrophe` experimented different ways to choose the samples: according to their makespan, relative distance, or conflicts with a given robot. UNIST used a sample of only one robot as we did in our *Feasible Optimizer* but UNIST adds also an original simulated annealing optimization step. Neither UNIST nor `gitastrophe` have used an optimization algorithm where the sample of robots to recompute evolves dynamically as in our *Conflict Optimizer*. The reason could be that this strategy destroys the solution feasibility without any guarantee of recovery. However, the Conflict Optimizer is probably the main ingredient that gave us a dramatic advantage over the other teams.

References

- 1 Max Barer, Guni Sharon, Roni Stern, and Ariel Felner. Suboptimal variants of the conflict-based search algorithm for the multi-agent pathfinding problem. In *Seventh Annual Symposium on Combinatorial Search*, 2014.
- 2 Roman Barták, Jiří Švancara, Věra Škopková, David Nohejl, and Ivan Krasičenko. Multi-agent path finding on real robots. *AI Communications*, 32(3):175–189, 2019. doi:10.3233/AIC-190621.
- 3 Eli Boyarski, Ariel Felner, Roni Stern, Guni Sharon, Oded Betzalel, David Tolpin, and Eyal Shimony. ICBS: The improved conflict-based search algorithm for multi-agent pathfinding. In *Eighth Annual Symposium on Combinatorial Search*, 2015.
- 4 Danny Z. Chen, Rajasekhar Inkulu, and Haitao Wang. Two-point L_1 shortest path queries in the plane. *Journal of Computational Geometry*, 7(1):473–519, 2016. doi:10.20382/jocg.v7i1a20.
- 5 Danny Z Chen and Haitao Wang. L_1 shortest path queries among polygonal obstacles in the plane. In *30th International Symposium on Theoretical Aspects of Computer Science, STACS 2013*, 2013. doi:10.4230/LIPIcs.STACS.2013.293.
- 6 IBM ILOG Cplex. V12. 1: User’s manual for CPLEX. *International Business Machines Corporation*, 46(53):157, 2009.
- 7 Boris De Wilde, Adriaan W. Ter Mors, and Cees Witteveen. Push and rotate: a complete multi-agent pathfinding algorithm. *Journal of Artificial Intelligence Research*, 51:443–492, 2014. doi:10.1613/jair.4447.
- 8 Erik D Demaine, Sándor P Fekete, Phillip Keldenich, Henk Meijer, and Christian Scheffer. Coordinated motion planning: Reconfiguring a swarm of labeled robots with bounded stretch. *SIAM Journal on Computing*, 48(6):1727–1762, 2019. doi:10.1137/18M1194341.
- 9 Sándor Fekete, Phillip Keldenich, Dominik Krupke, and Joseph S. B. Mitchell. Cg:shop 2021 coordinated motion planning. In *37th International Symposium on Computational Geometry, SoCG 2021*, volume 189 of *LIPIcs*, 2021.
- 10 Ariel Felner, Jiaoyang Li, Eli Boyarski, Hang Ma, Liron Cohen, TK Satish Kumar, and Sven Koenig. Adding heuristics to conflict-based search for multi-agent path finding. In *Proceedings of the International Conference on Automated Planning and Scheduling*, volume 28, 2018.
- 11 Ariel Felner, Roni Stern, Solomon Eyal Shimony, Eli Boyarski, Meir Goldenberg, Guni Sharon, Nathan Sturtevant, Glenn Wagner, and Pavel Surynek. Search-based optimal solvers for the

- multi-agent pathfinding problem: Summary and challenges. In *Tenth Annual Symposium on Combinatorial Search*, 2017.
- 12 Thibaut Goetghebuer-Planchon. A C++ implementation of a fast hash map and hash set using hopscotch hashing. <https://github.com/Tessil/hopscotch-map>, 2020.
 - 13 Meir Goldenberg, Ariel Felner, Roni Stern, Guni Sharon, Nathan Sturtevant, Robert C. Holte, and Jonathan Schaeffer. Enhanced partial expansion A*. *Journal of Artificial Intelligence Research*, 50:141–187, 2014.
 - 14 Edward Lam, Pierre Le Bodic, Daniel Damir Harabor, and Peter J. Stuckey. Branch-and-cut-and-price for multi-agent pathfinding. In *International Joint Conferences on Artificial Intelligence*, pages 1289–1296, 2019. doi:10.24963/ijcai.2019/179.
 - 15 Jiaoyang Li, Andrew Tinka, Scott Kiesel, Joseph W Durham, TK Kumar, and Sven Koenig. Lifelong multi-agent path finding in large-scale warehouses. *arXiv preprint arXiv:2005.07371*, 2020.
 - 16 Hang Ma, Sven Koenig, Nora Ayanian, Liron Cohen, Wolfgang Hönig, TK Kumar, Tansel Uras, Hong Xu, Craig Tovey, and Guni Sharon. Overview: Generalizations of multi-agent path finding to real-world scenarios. *arXiv preprint arXiv:1702.05515*, 2017.
 - 17 Guni Sharon, Roni Stern, Ariel Felner, and Nathan R Sturtevant. Conflict-based search for optimal multi-agent pathfinding. *Artificial Intelligence*, 219:40–66, 2015. doi:10.1016/j.artint.2014.11.006.
 - 18 David Silver. Cooperative pathfinding. *AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment*, 1:117–122, 2005.
 - 19 Jack Spalding-Jamieson, Paul Liu, Brandon Zhang, and Da Wei Zheng. Coordinated motion through good starts and better improvements. In *37th International Symposium on Computational Geometry, SoCG 2021*, volume 189 of *LIPICs*, 2021.
 - 20 Roni Stern, Nathan Sturtevant, Ariel Felner, Sven Koenig, Hang Ma, Thayne Walker, Jiaoyang Li, Dor Atzmon, Liron Cohen, TK Kumar, et al. Multi-agent pathfinding: Definitions, variants, and benchmarks. *arXiv preprint arXiv:1906.08291*, 2019.
 - 21 Pavel Surynek. Unifying search-based and compilation-based approaches to multi-agent path finding through satisfiability modulo theories. In *International Joint Conferences on Artificial Intelligence*, pages 1177–1183, 2019.
 - 22 Pavel Surynek, Ariel Felner, Roni Stern, and Eli Boyarski. Efficient SAT approach to multi-agent path finding under the sum of costs objective. In *22nd European Conference on Artificial Intelligence, ECAI 2016*, volume 285 of *Frontiers in Artificial Intelligence and Applications*, pages 810–818, 2016. doi:10.3233/978-1-61499-672-9-810.
 - 23 Hyeyun Yang and Antoine Vigneron. A simulated annealing approach to coordinated motion planning. In *37th International Symposium on Computational Geometry, SoCG 2021*, volume 189 of *LIPICs*, 2021.
 - 24 Jingjin Yu and Steven M LaValle. Planning optimal paths for multiple robots on graphs. In *2013 IEEE International Conference on Robotics and Automation*, pages 3612–3617. IEEE, 2013.