

Optimal and Bounded Suboptimal Any-Angle Multi-agent Pathfinding

Konstantin Yakovlev¹, Anton Andreychuk², Roni Stern³

Abstract—Multi-agent pathfinding (MAPF) is the problem of finding a set of conflict-free paths for a set of agents. Typically, the agents’ moves are limited to a pre-defined graph of possible locations and allowed transitions between them, e.g. a 4-neighborhood grid. We explore how to solve MAPF problems when each agent can move between any pair of possible locations as long as traversing the line segment connecting them does not lead to the collision with the obstacles. This is known as any-angle pathfinding. We present the first optimal any-angle multi-agent pathfinding algorithm. Our planner is based on the Continuous Conflict-based Search (CCBS) algorithm and an optimal any-angle variant of the Safe Interval Path Planning (TO-AA-SIPP). The straightforward combination of those, however, scales poorly since any-angle path finding induces search trees with a very large branching factor. To mitigate this, we adapt two techniques from classical MAPF to the any-angle setting, namely Disjoint Splitting and Multi-Constraints. Experimental results on different combinations of these techniques show they enable solving over 30% more problems than the vanilla combination of CCBS and TO-AA-SIPP. In addition, we present a bounded-suboptimal variant of our algorithm, that enables trading runtime for solution cost in a controlled manner.

I. INTRODUCTION

Multi-agent pathfinding (MAPF) is the problem of finding a set of conflict-free paths for a set of agents. Its applications include automated warehouses [1], traffic control [2], digital entertainment [3], etc. Different optimization variants of MAPF have proven to be NP-Hard [4], yet efficient optimal MAPF solvers have been proposed such as Conflict Based Search (CBS) [5], SAT-MDD [6], BCP [7] and others.

Most prior work focused on the classical version of the problem which makes many simplifying assumptions [8]. The most common of them are that *i*) the agents’ moves are limited to a given graph of possible locations and the transitions between them, and *ii*) all moves take one time unit. Indeed, these limitations are not well-suited for a variety of robotic applications and, thus, some works have begun to explore how to solve MAPF without them. In particular, the Continuous Conflict-based Search (CCBS) algorithm guarantees completeness and optimality while avoiding the need to discretize time and can handle agent actions with different durations [9]. Still, CCBS as well as many other previously proposed optimal MAPF solvers rely on the first assumption, i.e., that the agents move over a pre-defined graph of locations and allowed transitions between them.

¹Konstantin Yakovlev is with Federal Research Center for Computer Science and Control RAS and with AIRI yakovlev@airi.net

²Anton Andreychuk is with AIRI andreychuk@airi.net

³Roni Stern is with Ben-Gurion University of the Negev sternron@bgu.ac.il




Fig. 1. Two optimal solutions of the same MAPF instance: the one composed of the cardinal moves only (left) and the one with any-angle moves (right). The cost of the latter is 22% lower.

We focus on relaxing this assumption, allowing each agent to move between any pair of locations – see Fig. 1. This type of pathfinding is often called *any-angle pathfinding* [10]. Algorithms such as Anya [11] and TO-AA-SIPP [12] have been proposed for optimal single-agent any-angle path finding, and a suboptimal any-angle MAPF solver was also proposed [13]. **We propose AA-CCBS, the first any-angle MAPF algorithm that is guaranteed to return cost optimal solutions.**

AA-CCBS integrates CCBS with TO-AA-SIPP. While it is guaranteed to return optimal solutions, AA-CCBS scales poorly since any-angle path finding induces search trees with a very large branching factor. To mitigate this, we propose several enhancements to AA-CCBS based on techniques from classical MAPF, namely disjoint splitting (DS) [14] and multi-constraints (MC) [15]. In particular, we suggest several novel variants of how MC can be applied for our any-angle setup, without compromising the theoretical guarantees.

We conduct a thorough empirical evaluation of different variants of AA-CCBS across standard benchmarks. The results show that the our enhancements AA-CCBS work very well, solving significantly more problems under a given time limit than vanilla AA-CCBS. We also show how AA-CCBS can be easily generalized to be a bounded-suboptimal algorithm, allowing a controlled way to trade off runtime for solution quality. Our experimental results demonstrate that allowing even a small amount of suboptimality allows solving many more problems.

II. RELATED WORKS

Lifting certain Classical MAPF [8] assumptions to make the resultant solvers more suitable to real-world applications has been studied recently. In [13] a prioritized any-angle MAPF planner was introduced and in [16] this planner was enhanced to support robots of different sizes and moving speeds. Other techniques for planning with kinematic constraints, including the post-processing of the plans produced by the classical MAPF solvers, were also proposed [17], [18], [19]. Yan and Li [20] proposed an involved three-stage solver based on Priority-Based Search [21], that handles not only kinematic constraints but accelerating and decelerating actions. All these solvers, unlike ours, are not optimal.

Optimal MAPF solvers that go beyond Classical MAPF setting also exist. A vast majority of them are adaptations of the seminal CBS algorithm [5]. In [22] individual roadmaps for the robots are constructed and then used by CBS (for collision detection each path in each roadmap is split into the uniform time resolution segments). A dedicated technique is suggested to decide when to re-construct these roadmaps. A combination of CBS with trajectory optimization methods was suggested to construct collision-free paths for a team of heterogeneous agents (quadrotors) [23], [24]. Kinodynamic variants of CBS were considered in [25], [26], [27]. In this work, we also leverage CBS, specifically its variant adapted to non-discretized timeline [9], to conduct multi-agent any-angle path finding with optimality guarantees.

III. PROBLEM STATEMENT

Consider a set of vertices V residing in the metric space, e.g., a grid where the vertices are the set of centers of the grid cells. There are n designated start and goal vertices, and n agents initially located at the start vertices. The agents are modeled as disks of a certain radius. Two types of actions for each agent are possible: wait at the current vertex (for an arbitrary amount of time) and move between the vertices. A move is valid only if the given line-of-sight function, $los : V \times V \rightarrow \{true, false\}$, returns *true*. This function checks whether an agent can follow a straight-line segment connecting the vertices without colliding with an obstacle, e.g., whether the set of grid cells swept by the disk-shaped agent contains only free cells. The cost of each action is its duration. The duration of a wait action may be arbitrary. The duration of a move action equals the distance between the vertices comprising this move (we assume that the agents start/stop instantaneously and move with the unit speed).

An individual plan for an agent is a sequence of timed actions $\pi = ((a_0, t_0), (a_1, t_1), \dots, (a_n, t_n))$, where a_i is a move action, t_i is the starting time of this action, and each action must start where the previous one ends. The wait actions are encoded implicitly in the plan, i.e. if $t_i + dur(a_i) < t_{i+1}$ then the agent is waiting at the endpoint of a_i for a time of $t_{i+1} - t_i - dur(a_i)$. Here $dur(a_i)$ denotes the duration of action a_i . The plans of distinct agents are said to be in conflict if there exists a time moment when the distance between the agents, executing these plans, is less than sum of their radii. The any-angle multi-agent pathfinding problem

(AA-MAPF) asks to find a set of n plans transferring the agents from their start vertices to the goal ones, such that each pair of plans is collision-free. The cost of the individual plan, $c(\pi)$, is the time when the agent reaches its goal. We wish to solve the problem optimally w.r.t. sum-of-cost objective, which is the sum of durations of the agents' plans, $SOC = \sum_{i=1}^n c(\pi_i)$.

IV. BACKGROUND

Our planner is based on a combination of CCBS, [9] and TO-AA-SIPP [12]. Next, we briefly introduce these methods.

A. CCBS

CCBS is a variant of CBS [5] adapted to non-discretized timeline. It reasons over the continuous time-intervals instead of distinct isolated time steps. CCBS works by finding plans for each agent separately, detecting *conflicts* between these plans, and resolving them by replanning for the individual agents subject to specific *constraints*. A conflict between the two agents in CCBS is defined by (a_i, a_j, t_i, t_j) , representing that a collision occurs when executing the actions a_i and a_j at time moments t_i and t_j , respectively. CCBS resolves a conflict by imposing a constraint on an agent and replan for that agent. The constraint imposed by CCBS is a tuple $(a, [t, t'])$, stating that the agent is prohibited from taking action a in the time range $[t, t']$. The latter is called the *unsafe interval*. Several approaches were considered to compute t' , the first time moment an agent can safely start executing the action, including closed-loop formulas [28].

To guarantee optimality, CCBS runs two search processes: a high-level search to choose which agent to constraint and a low-level search to find a path for that agent that satisfies its constraints. CCBS uses a best-first search for the high-level search and the Safe-Interval Path Planning (SIPP) algorithm [29] for the low-level search. SIPP is a specific adaptation of A* for space-time domains. It was originally proposed for single-agent pathfinding among dynamic obstacles, but can be easily adapted to handle CCBS constraints, as the latter can be viewed as dynamic obstacles that temporarily appear in the workspace.

B. TO-AA-SIPP

TO-AA-SIPP is a variant of SIPP that aims at finding an optimal any-angle path for an agent navigating in an environment with dynamic obstacles. As in regular SIPP, its search node is identified by a tuple $(v, [t_l, t_u])$, where v is a graph vertex and $[t_l, t_u]$ is a safe time interval that dictates when the agent may safely reside at v .

The principal difference between TO-AA-SIPP and SIPP is that the former does not iteratively build a search tree by expanding search nodes. Instead, it generates all search nodes beforehand (this is possible due to the usage of time intervals as the nodes' identifiers) and iteratively tries to identify the correct parent relationships between the nodes to obtain the optimal solution. The authors refer to this technique as "inverted expansions". TO-AA-SIPP was proven to be complete and optimal.




Fig. 2. An example of vanilla AA-CCBS multi-constraint.

V. AA-CCBS

To solve AA-MAPF, we propose to use TO-AA-SIPP as the low-level planner for CCBS. We call this algorithm AA-CCBS. Unfortunately, AA-CCBS struggles to solve even simple instances. E.g., the small AA-MAPF instance with 3 agents depicted in Fig. 1 requires expanding 6, 660 high-level nodes and consequently a similar number of low-level calls to TO-AA-SIPP. To enable AA-CCBS to scale more gracefully, we propose two vital enhancements: Multi-Constrains (MC) and Disjoint Splitting (DS).

A. Multi-Constraints in Any-Angle MAPF

CCBS resolves a conflict by adding a single constraint to one of the agents and replanning. However, adding multiple constraints (multi-constraint) when resolving a single conflict can reduce the number of high-level search iterations [30], [15]. In particular, Walker et al. [15] proposed the Time-Annotated Biclique (TAB) method for adding multiple constraints in domains with non-uniform actions.

For two conflicting actions, (a_i, t_i) and (a_j, t_j) , TAB iterates over all actions that have the same source vertices as either a_i or a_j and identifies the subsets of them, A_i and A_j , that are *mutually conflicting*, i.e. each pair of actions from $A_i \times A_j$ lead to a conflict (if they start at t_i, t_j respectively). The multi-constraint (MC) added to the first agent comprises the constraints $(a_i, [t, t'])$ where $[t, t']$ is the largest time interval that is *fully included* into the unsafe intervals induced by a_i and *all* $a_j \in A_j$. The MC added to the second agent is defined similarly. TAB is applicable in AA-CCBS. We call this variant MC1.

Example. Consider two actions sets $\{a_{i_1}, a_{i_2}\}$, $\{a_{j_1}, a_{j_2}\}$ and time moments t_i, t_j . Let $[t_i, t'] = \text{unsafe}(a_{i_1}, a_{j_1}) \cap \text{unsafe}(a_{i_1}, a_{j_2})$ and $[t_i, t''] = \text{unsafe}(a_{i_2}, a_{j_1}) \cap \text{unsafe}(a_{i_2}, a_{j_2})$, where *unsafe* denotes the unsafe interval of a_{i_k} starting at t_i w.r.t. a_{j_l} starting at t_j . Then, the MC associated with the first action set is $\{(a_{i_1}, [t_i, t']); (a_{i_2}, [t_i, t''])\}$. It dictates that the agent is not allowed to perform either a_{i_1} or a_{i_2} in the respective time intervals. Fig. 2 shows an illustrative example.

The problem with MC1 is that in AA-MAPF the number of mutually-conflicting actions may be large. Thus, after intersecting their unsafe intervals the final unsafe interval for a certain action is likely to get trimmed. For example, consider action a_{11} of Fig. 2, which has an unsafe interval




Fig. 3. Actions comprising different versions of multi-constraints. Please note that the time intervals of the actions in MC2 and MC3 are notably larger compared to MC1. Thus MC2 and MC3 are expected to exhibit more pruning power.

of $[0; 0.522]$ in the MC, despite its unsafe interval w.r.t a_{21} , ending at 1.106, is twice longer. A shorter unsafe interval means that the power of constraint diminishes. To mitigate this issue we suggest two modifications.

First, we suggest to consider only a limited set of actions to form A_i (A_j similarly). The source of each action should coincide with the source of the initially conflicting action a_i . The destination must be a vertex that is swept by the agent when executing a_i , i.e. the agent's body intersects this vertex (grid cell). Thus the resulting actions form a “stripe” along a_i – see Fig. 3 (the cells that form a stripe are highlighted).

The second modification is that we filter out the actions from the MC that lead to trimming the original unsafe interval. In particular, if $\text{unsafe}(a_i, a'_j) \not\subseteq \text{unsafe}(a_i, a_j)$ then a'_j is excluded from A_j . Here a_i and a_j are the original conflicting actions, and $\{a'_j\}$ is the set of actions that are in conflict with all $a \in A_i$. The combination of those two enhancements is dubbed MC2.

The detailed pseudo-code of constructing MC2 is shown in Alg. 1. After initialization (Lines 1-3), we identify the swept cells (Line 4) and form the candidate sets of actions $\hat{A}_i, \hat{A}_j$ (Line 5). Then, in Line 6, we merge them into $\hat{A}$ in such a way that the actions from $\hat{A}_i$ and $\hat{A}_j$ alternate (this merging procedure is denoted as *zip* in code). Next, we iterate over $\hat{A}$ and filter out actions that lead to trimming the unsafe interval of a_i (Lines 10-12). After pruning the candidate sets we iterate over the actions once again to compute their unsafe intervals (Lines 15-21). The resulting sets of action-interval pairs form MC2.

Next, we suggest one more way of composing the action sets of MC, when we include into the constrained set actions that start at the *different* graph vertices. This is possible as by definition the set of actions comprising A_i and A_j must be mutually conflicting but the definition does not specify what source and target vertices of these actions should be. In specifics, we enlarge the action sets of MC2 with the actions that have the same target vertex (as the original conflicting actions) and the source vertices that lie on the same stripe as before – see Fig. 3 (right). The rationale is that these actions are likely to lead to collisions that will happen between the same agents and nearly in the same place. Thus it is natural

Algorithm 1: Forming MC2 for AA-CCBS

Input: Conflicting actions (and times) a_i, t_i, a_j, t_j
Output: A pair of multi-constraints $\{A_i, A_j\}$

- 1 $[t_i, t'] \leftarrow$ unsafe interval of a_i w.r.t. (a_j, t_j) ;
- 2 $[t_j, t''] \leftarrow$ unsafe interval of a_j w.r.t. (a_i, t_i) ;
- 3 $A_i \leftarrow \emptyset$; $A_j \leftarrow \emptyset$;
- 4 $C_{\{i,j\}} \leftarrow$ vertices swept by $a_{\{i,j\}}$;
- 5 $\hat{A}_{\{i,j\}} \leftarrow$ valid actions that start at the source of $a_{\{i,j\}}$ and end in $C_{\{i,j\}}$;
- 6 $\hat{A} \leftarrow \text{zip}(\hat{A}_i, \hat{A}_j)$;
- 7 **for each** $\hat{a} \in \hat{A}$ **do**
- 8 **if** $\text{source}(\hat{a}) = \text{source}(a_i)$ **then**
- 9 **if** $\text{IsConflict}(\hat{a}, A_j) = \text{true}$ **then**
- 10 $[t_j, \hat{t}] \leftarrow$ unsafe interval of a_j w.r.t. $(\hat{a}, t_i)$;
- 11 **if** $t'' > \hat{t}$ **then**
- 12 remove $\hat{a}$ from $\hat{A}_i$;
- 13 **else**
- 14 handling a_j similarly (Lines 9-12);
- 15 **for each** $a \in \hat{A}_i$ **do**
- 16 $t_{\min} \leftarrow \infty$;
- 17 **for each** $a' \in \hat{A}_j$ **do**
- 18 $[t_i, t_{\text{cur}}] \leftarrow$ unsafe interval of a w.r.t. (a', t_j) ;
- 19 $t_{\min} \leftarrow \min\{t_{\min}, t_{\text{cur}}\}$;
- 20 $A_i \leftarrow A_i \cup \{(a, [t_i, t_{\min}])\}$;
- 21 Process $\hat{A}_j$ similarly (Lines 15-20);
- 22 **return** $\{A_i, A_j\}$

to include them in the same multi-constraint. The pseudo-code for constructing MC3 largely repeats the one for MC2 (with certain additions) and is omitted for the sake of space.

B. Disjoint Splitting

Disjoint splitting (DS) is a powerful technique for reducing search effort in CBS [14] that was shown to be effective in CCBS as well [31]. In DS for CCBS, a conflict (a_i, a_j, t_i, t_j) is resolved in two ways: one CCBS child gets a regular (negative) constraint on agent i , while the other one - a *positive* constraint to agent i and a negative constraint to agent j . The positive constraint comes in the conventional CCBS form of $(a, [t, t'])$ but dictates the agent must perform action a in a time moment belonging to $[t, t']$. Thus action a becomes a *landmark* in the sense that the low-level search has to, first, find a plan to the source of a , then execute a , and then plan to the goal (or to the next landmark).

In this work, we suggest two options for implementing DS for AA-CCBS: using DS as-is and using DS with MC. The first one is straightforward, while the second requires more attention.

To integrate DS with MCs that start at the same vertex but end in different ones, like MC1 or MC2, we need to impose a combined positive constraint that will dictate an agent to perform *any* of the actions comprising the constrained set.

Thus, we need to reason over the multiple distinct start locations for the consecutive search, which is not trivial. Integrating DS with MC3 is even more involved as here we need to handle landmarks that are composed of the actions that start at different locations.

To simplify the integration of DS and MC in AA-CCBS we suggest to avoid creating landmarks out of multi-constraints, while still operating with conventional negative multi-constraints. In particular, consider a conflicting pair of timed actions, (a_i, t_i, a_j, t_j) . As in the original DS, we impose a singular negative constraint on the first agent in the left CCBS child and the corresponding singular positive constraint in the right child. Additionally, the second agent gets a negative multi-constraint in the right child, which is computed with respect to the action (a_i, t_i) only. In this way, there is no need to handle positive MCs, while the regular (singular) positive constraint is enforced with the negative MC (in one of the descendants in CCBS search tree).

Theorem 1: AA-CCBS with any of the enhancements — MC1, MC2, MC3, and DS — is sound, solution-complete, and optimal, i.e., it is guaranteed to return a valid solution if such exists, and the solution it returns is cost-optimal.

Proof outline: Soundness, solution-completeness, and optimality for vanilla AA-CCBS is trivial, following the proofs of CCBS [9] and TO-AA-SIPP [12]. To prove that the proposed improvements (MC1, MC2, MC3, and DS) preserve these properties, it is sufficient to show that AA-CCBS with them always resolves conflicts with a *sound pair of constraints* [32]. A pair of constraints is sound if in any valid solution, at least one of these constraints is satisfied. That is, any solution in which both constraints are violated is not valid. In MC1, MC2, and MC3, by construction, for every pair of multi-constraints A_i and A_j they return it holds that any pair of constraints $(a_i^*, T_i^*) \in A_i$ and $(a_j^*, T_j^*) \in A_j$ form a sound pair of constraints. Thus, any solution that violates a constraint in A_i and a constraint in A_j cannot be valid. A similar argument holds for DS due to the negative constraints it imposes.

VI. EXPERIMENTAL RESULTS

We implemented AA-CCBS in C++¹ and conducted the experimental evaluation on five different maps from the well-known in the community MovingAI benchmark [8]: empty-16-16, random-32-32-20, maze-32-32-4, den321d and warehouse-10-20-10-2-2. For each map, the benchmark provides 25 scenarios. Each scenario is a set of start-goal locations. We take two agents, run the algorithm, and then increment the number of agents until the algorithm is not able to find a solution within a time-limit of 300 seconds. In the latter case, we proceed to the next scenario. We evaluate six versions of AA-CCBS: AA-CCBS (vanilla), AA-CCBS+MC1, AA-CCBS+MC2, AA-CCBS+MC3, AA-CCBS+DS, AA-CCBS+DS+MC3.

Fig. 4 shows the results. The left pane shows the number of solved instances (with the breakdowns for each map).

¹<https://github.com/PathPlanning/Continuous-CBS/tree/AA-CCBS>




Fig. 4. The results of evaluation of all the versions of AA-CCBS. a) shows the amount of totally solved instances b) show the time required to find a solution c) demonstrates the comparison of the three best-performing methods in terms of HL expansions.

The results indicate that MC1 and MC2 do not provide significant improvement over vanilla AA-CCBS. MC3, DS, and DS+MC3, however, notably outperform it, especially the ones that use DS. The central pane shows the number of tasks solved within a certain time cap. The X-axis shows the number of tasks and Y-axis shows the time needed to solve this amount of tasks. One can note that MC3 outperforms all other variants of AA-CCBS for easy tasks, i.e. for the ones that need less than 0.1 seconds to be solved. Meanwhile DS variants of AA-CCBS clearly outperform MC3 when the time, needed to solve an instance, is greater than 1 second. The right pane of Fig. 4 depicts the number of the high-level iterations of AA-CCBS with MC3, DS, DS+MC3 for each MAPF instance that was solved by all these methods. As seen on the plot, this number is lower for MC3 for nearly half of the instances. This also confirms that MC3 solves some (i.e. easy) instances faster than DS and MC2+DS, while it is outperformed on the other (hard) ones.

Overall the obtained results confirm that both MC and DS are valuable techniques to enhance AA-CCBS performance. MCs (specifically, MC3) are well-suited for the easy instances while the DS – for the harder ones. Generally, we would recommend to use AA-CCBS+DS+MC3 to solve any-angle MAPF optimally.

VII. BOUNDED-SUBOPTIMAL AA-CCBS

Given a real number $w > 1$ a bounded suboptimal (BS) solution is the one whose cost does not exceed the cost of the optimal one by more than a factor of w . Finding a BS solution is often much easier and faster than finding an optimal one. Creating a BS variant of AA-CCBS with any of our enhancements is straightforward: replacing the best-first search in the AA-CCBS high-level with a Focal search [33].² I.e. at each high-level iteration we, first, form a subset of nodes whose cost does not exceed the cost of the most promising node (the one which regular AA-CCBS would have picked) by a factor of w and, second, choose from this subset the node with the smallest number of conflicts breaking ties in favor of nodes containing more constraints. Thus we force the search to more greedily resolve the conflicts, while still maintaining the bound on the suboptimality of the resultant solution.

²Creating a BS low-level search is not trivial in our case and goes beyond the scope of this paper.




Fig. 5. The results of the evaluation of AA-CCBS and its modifications with different suboptimality factors.

We implemented such a BS version of AA-CCBS with the three best-performing enhancements, DS, MC3, DS+MC3, and evaluated them on the same problem instances as before. Fig. 5 shows the results. In the upper part the box-and-whiskers plot of the resultant solution cost depending on the suboptimality bound are shown. The table at the bottom shows the number of the solved instances across all maps.

Indeed, MC3 notably increases the performance of AA-CCBS when searching for bounded suboptimal solutions (i.e. when $w > 1$), while DS variants are not even able to outperform vanilla AA-CCBS. We hypothesize that this is due to the positive constraints imposed by DS do not lead to immediate eliminating the conflict that causes these constraints. When searching for optimal solutions this is compensated by reducing the branching factor. However the search for suboptimal solutions is more greedy by design and branching factor becomes much less important. Thus, DS loses much of its power, while retaining the drawback of not immediately eliminating the conflict. At the same time, MC3 imposes constraints on a larger number of actions, and as the search greedily goes deeper it actually eliminates the conflict. Moreover, due to the requirement to satisfy all the imposed positive constraints, the actual costs of the solutions found by DS versions are on average more than twice higher

than the ones of MC3. The actual overhead in solution costs of MC3 version is not higher than 2% on average even when the suboptimality bound is $w = 1.25$. The obtained results provide further support that MC3 is a useful technique: not only does it speed the search for optimal solutions on easy instances, but it also significantly increases the performance of bounded-suboptimal AA-CCBS.

VIII. CONCLUSION AND FUTURE WORK

In this work, we presented AA-CCBS, the first optimal any-angle MAPF algorithm. We showed how to incorporate existing CCBS enhancements, namely Disjoint Splitting and Multi-Constraints, into AA-CCBS to allow it to scale better. Specifically, we introduced three ways to implement Multi-Constraints, providing stronger pruning of the high-level search tree. AA-CCBS can also be easily adapted to return bounded-suboptimal solutions. Experimental results confirm that the suggested enhancements have a huge impact on the algorithm's performance. Future work can explore more sophisticated procedures of forming multi-constraints, identifying which sets of actions should be considered in each step, as well as adapting incremental search techniques in the any-angle low-level search.

REFERENCES

- [1] P. R. Wurman, R. D'Andrea, and M. Mountz, "Coordinating hundreds of cooperative, autonomous vehicles in warehouses," *AI magazine*, vol. 29, no. 1, pp. 9–9, 2008.
- [2] A. Parks-Young and G. Sharon, "Intersection management protocol for mixed autonomous and human-operated vehicles," *IEEE Transactions on Intelligent Transportation Systems*, vol. 23, no. 10, pp. 18315–18325, 2022.
- [3] H. Ma, J. Yang, L. Cohen, T. K. S. Kumar, and S. Koenig, "Feasibility study: Moving non-homogeneous teams in congested video game environments," in *AIIDE*, 2017.
- [4] J. Yu and S. M. LaValle, "Optimal multirobot path planning on graphs: Complete algorithms and effective heuristics," *IEEE Transactions on Robotics*, vol. 32, no. 5, pp. 1163–1177, 2016.
- [5] G. Sharon, R. Stern, A. Felner, and N. R. Sturtevant, "Conflict-based search for optimal multi-agent pathfinding," *Artificial Intelligence*, 2015.
- [6] P. Surynek, A. Felner, R. Stern, and E. Boyarski, "Efficient sat approach to multi-agent path finding under the sum of costs objective," in *ECAI*, pp. 810–818, 2016.
- [7] E. Lam, P. Le Bodic, D. Harabor, and P. J. Stuckey, "Branch-and-cut-and-price for multi-agent path finding," *Computers & Operations Research*, vol. 144, 2022.
- [8] R. Stern, N. R. Sturtevant, A. Felner, S. Koenig, H. Ma, T. T. Walker, J. Li, D. Atzmon, L. Cohen, T. S. Kumar, *et al.*, "Multi-agent pathfinding: Definitions, variants, and benchmarks," in *Proceedings of the 12th Annual Symposium on Combinatorial Search (SoCS 2019)*, pp. 151–158, 2019.
- [9] A. Andreychuk, K. Yakovlev, P. Surynek, D. Atzmon, and R. Stern, "Multi-agent pathfinding with continuous time," *Artificial Intelligence*, 2022.
- [10] A. Nash, K. Daniel, S. Koenig, and A. Felner, "Theta*: Any-angle path planning on grids," in *Proceedings of The 22nd AAAI Conference on Artificial Intelligence (AAAI 2007)*, pp. 1177–1183, 2007.
- [11] D. Harabor and A. Grastien, "An optimal any-angle pathfinding algorithm," in *International Conference on Automated Planning and Scheduling*, pp. 308–311, 2013.
- [12] K. Yakovlev and A. Andreychuk, "Towards time-optimal any-angle path planning with dynamic obstacles," in *Proceedings of the 31st International Conference on Automated Planning and Scheduling (ICAPS 2021)*, pp. 405–414, 2021.
- [13] K. Yakovlev and A. Andreychuk, "Any-angle pathfinding for multiple agents based on sipp algorithm," in *International Conference on Automated Planning and Scheduling*, pp. 586–594, 2017.
- [14] J. Li, D. Harabor, P. J. Stuckey, A. Felner, H. Ma, and S. Koenig, "Disjoint splitting for multi-agent path finding with conflict-based search," in *International Conference on Automated Planning and Scheduling*, pp. 279–283, 2019.
- [15] T. T. Walker, N. R. Sturtevant, and A. Felner, "Generalized and sub-optimal bipartite constraints for conflict-based search," in *Proceedings of the 34th AAAI Conference on Artificial Intelligence (AAAI 2020)*, pp. 7277–7284, 2020.
- [16] K. Yakovlev, A. Andreychuk, and V. Vorobyev, "Prioritized multi-agent path finding for differential drive robots," in *Proceedings of the 2019 European Conference on Mobile Robots (ECMR 2019)*, pp. 1–6, IEEE, 2019.
- [17] H. Ma, W. Hönig, T. K. S. Kumar, N. Ayanian, and S. Koenig, "Lifelong path planning with kinematic constraints for multi-agent pickup and delivery," in *Proceedings of the 33rd AAAI Conference on Artificial Intelligence (AAAI 2019)*, pp. 7651–7658, 2019.
- [18] W. Hönig, T. S. Kumar, L. Cohen, H. Ma, H. Xu, N. Ayanian, and S. Koenig, "Multi-agent path finding with kinematic constraints," in *Proceedings of The 26th International Conference on Automated Planning and Scheduling (ICAPS 2016)*, pp. 477–485, 2016.
- [19] Z. A. Ali and K. Yakovlev, "Prioritized sipp for multi-agent path finding with kinematic constraints," in *Proceeding of the 6th International Conference on Interactive Collaborative Robotics (ICR 2021)*, pp. 1–13, Springer, 2021.
- [20] J. Yan and J. Li, "Multi-agent motion planning with bézier curve optimization under kinodynamic constraints," *IEEE Robotics and Automation Letters*, 2024.
- [21] H. Ma, D. Harabor, P. J. Stuckey, J. Li, and S. Koenig, "Searching with consistent prioritization for multi-agent path finding," in *Proceedings of the 33rd AAAI Conference on Artificial Intelligence (AAAI 2019)*, pp. 7643–7650, 2019.
- [22] I. Solis, J. Motes, R. Sandström, and N. M. Amato, "Representation-optimal multi-robot motion planning using conflict-based search," *IEEE Robotics and Automation Letters*, vol. 6, no. 3, pp. 4608–4615, 2021.
- [23] M. DeBord, W. Hönig, and N. Ayanian, "Trajectory planning for heterogeneous robot teams," in *Proceedings of the 2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2018)*, pp. 7924–7931, IEEE, 2018.
- [24] W. Hönig, J. A. Preiss, T. S. Kumar, G. S. Sukhatme, and N. Ayanian, "Trajectory planning for quadrotor swarms," *IEEE Transactions on Robotics*, vol. 34, no. 4, pp. 856–869, 2018.
- [25] J. Kottlinger, S. Almagor, and M. Lahijanian, "Conflict-based search for multi-robot motion planning with kinodynamic constraints," in *Proceedings of the 2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2022)*, pp. 13494–13499, IEEE, 2022.
- [26] A. Moldagaliyeva, J. Ortiz-Haro, M. Toussaint, and W. Hönig, "db-cbs: Discontinuity-bounded conflict-based search for multi-robot kinodynamic motion planning," *International Conference on Robotics and Automation (ICRA)*, 2024.
- [27] L. Wen, Y. Liu, and H. Li, "Cl-mapf: Multi-agent path finding for car-like robots with kinematic and spatiotemporal constraints," *Robotics and Autonomous Systems*, vol. 150, p. 103997, 2022.
- [28] T. T. Walker and N. R. Sturtevant, "Collision detection for agents in multi-agent pathfinding," 2019.
- [29] M. Phillips and M. Likhachev, "SIPP: Safe interval path planning for dynamic environments," in *Proceedings of The 2011 IEEE International Conference on Robotics and Automation (ICRA 2011)*, pp. 5628–5635, 2011.
- [30] J. Li, P. Surynek, A. Felner, H. Ma, and S. Koenig, "Multi-agent path finding for large agents," in *Proceedings of the 33rd AAAI Conference on Artificial Intelligence (AAAI 2019)*, pp. 7627–7634, 2019.
- [31] A. Andreychuk, K. Yakovlev, E. Boyarski, and R. Stern, "Improving continuous-time conflict based search," in *Proceedings of the 35th AAAI Conference on Artificial Intelligence (AAAI 2021)*, pp. 11220–11227, 2021.
- [32] D. Atzmon, R. Stern, A. Felner, G. Wagner, R. Barták, and N.-F. Zhou, "Robust multi-agent path finding and executing," *Journal of Artificial Intelligence Research*, vol. 67, pp. 549–579, 2020.
- [33] J. Pearl and J. H. Kim, "Studies in semi-admissible heuristics," *IEEE transactions on pattern analysis and machine intelligence*, no. 4, pp. 392–399, 1982.