

Conflict-Based Lazy Search for Fast Multimanipulator Planning

Dongliang Zheng¹, Member, IEEE, Zhipeng Wang², Member, IEEE, Siqi Wang³,
Yuxi Lu⁴, Member, IEEE, Bin He, Senior Member, IEEE, Hesheng Wang⁵, Senior Member, IEEE,
and Panagiotis Tsiotras⁶, Fellow, IEEE

Abstract—Employing multiple manipulators can boost efficiency and accomplish tasks that a single manipulator cannot do. However, real-time planning for multiple manipulators in a cluttered workspace still poses significant challenges for planning algorithms. This article proposes a new planning algorithm called Conflict-Based Lazy Search (CBLS) for multimanipulator planning. CBLS is built on Conflict-Based Search (CBS), an efficient multiagent pathfinding (MAPF) algorithm that has shown an order of magnitude speedup over previous approaches [1], [2]. CBS addresses MAPF by solving many single-agent pathfinding (SAPF) problems. Thus, its planning time directly depends on the efficiency of the SAPF algorithm adopted. Our CBLS algorithm enhances CBS with precomputation and lazy search. First, a lazily evaluated graph with controlled sparsity is precomputed for a single manipulator. Second, we propose the Lazy Edged-based A* (LEA*) for efficient SAPF. Since edge evaluation is the computational bottleneck of manipulator planning, LEA* uses lazy search and an edge queue to reduce the number of edge evaluations. We show that LEA* is optimally vertex efficient and has improved edge efficiency compared to A*. We apply the proposed CBLS to multimanipulator planning problems and show its superior performance by comparing it with CBS and a sampling-based algorithm, namely, RRT-Connect.

Index Terms—Conflict-based search (CBS), lazy search, manipulator planning, multimanipulator, path planning.

I. INTRODUCTION

FINDING optimal collision-free paths for multiagent systems is an NP-hard problem, especially for multimanipulator systems operating in cluttered environments. Sampling-based motion planning methods such as PRM [3], RRT* [4], RRT# [5], and BIT* [6] make the problem tractable by approximating the search space using graphs or trees. They can also be viewed as graph-based methods. Graph search is used in PRM to find resolution-optimal paths, and local graph exploitation is used in RRT* to find asymptotically optimal paths. PRM builds a roadmap by drawing samples to cover the search space approximately and connecting neighboring samples using edges. RRT* builds a rapidly exploring tree by connecting the tree to incrementally added samples.

Sampling-based motion planning methods are well-suited for planning in high-dimensional spaces. They find approximate solutions quickly by using graphs with a small number of samples. Earlier studies for multiagent path planning combine the search space of each agent to form a “joint agent” and solve a single-agent planning problem [7], [8]. The dimensionality of the search space is the sum of the dimensionality of each agent. However, the complexity of planning grows exponentially with the dimensionality of the search space. These methods usually find solutions that are far from optimal, and the convergence rate to the optimal solution is slow.

Conflict-Based Search (CBS) [1], [9] is a dedicated search-based algorithm for multiagent pathfinding (MAPF). Instead of treating multiple agents as one “joint agent,” CBS iterates between: 1) independent single-agent pathfinding (SAPF); 2) conflict checking among all agent paths; and 3) conflict resolution. CBS uses a two-level scheme. At the high-level, CBS checks for any conflict among single-agent paths and resolves conflicts by adding constraints. At the low-level, single-agent paths consistent with the constraints are computed.

CBS requires an existing graph for graph search and usually assumes the graph is a grid world. For multimanipulator planning using CBS, the first step is to construct a graph. Implicit graphs can be defined using motion primitives (MPs) [10], [11]. MPs specify the allowable actions of the manipulator. The manipulator moves from one vertex to its neighboring vertex

Received 26 April 2025; revised 6 September 2025 and 28 December 2025; accepted 13 June 2026. Recommended by Technical Editor J. D. Sanchez Torres and Senior Editor M. Indri. This work was supported in part by the Science and Technology Commission of Shanghai Municipality (STCSM) under Grant 25ZR1404010, in part by the Shanghai Pujiang Program under Grant 25PJA140, in part by the National Natural Science Foundation of China under Grant 62088101, in part by the STCSM under Grant 2021SHZDZX0100, and in part by the Fundamental Research Funds for the Central Universities. (Corresponding author: Zhipeng Wang.)

Dongliang Zheng, Zhipeng Wang, Siqi Wang, Yuxi Lu, and Bin He are with the Shanghai Research Institute for Intelligent Autonomous Systems, Shanghai Institute of Intelligent Science and Technology, State Key Laboratory of Autonomous Intelligent Unmanned Systems, and Frontiers Science Center for Intelligent Autonomous Systems of Ministry of Education, Tongji University, Shanghai 201203, China (e-mail: dzheng46@tongji.edu.cn; wangzhipeng@tongji.edu.cn; yuxilu@tongji.edu.cn; 2310881@tongji.edu.cn; binhe@tongji.edu.cn).

Hesheng Wang is with the Department of Automation, Shanghai Jiao Tong University, Shanghai 200240, China (e-mail: wanghesheng@sjtu.edu.cn).

Panagiotis Tsiotras is with the School of Aerospace Engineering and Institute for Robotics and Intelligent Machines, Georgia Institute of Technology, Atlanta, GA 30332 USA (e-mail: tsiotras@gatech.edu).

Color versions of one or more figures in this article are available at <https://doi.org/10.1109/TMECH.2026.3704386>.

Digital Object Identifier 10.1109/TMECH.2026.3704386

by applying such a motion primitive. Using a graph search algorithm, [7] finds a sequence of MPs that connect the start and goal configuration of the manipulator. Other works, such as PRM, use an explicit graph, where graph vertices are random samples drawn in the configuration space, and neighboring vertices within a radius are connected using edges. In this article, we construct an explicit graph for manipulator planning, and introduce a parameter to control the sparsity of the graph. The proposed graph construction method generates increasingly dense graphs for various requirements. It can be applied to different numbers of manipulators and environments with different obstacle configurations and thus is environment-independent and robot-centric.

Graph search is performed for SAPF after graph construction. Under the CBS framework, the MAPF planning time is proportional to the SAPF planning time [1]. Numerous algorithms have been developed for SAPF. The A* [12] algorithm is a popular algorithm that is vertex optimal, that is, any other algorithm that finds the same shortest path will expand at least as many vertices when using the same heuristic. The weighted-A* algorithm finds a bounded suboptimal solution with an inflated heuristic cost [10].

One practical issue with robot path planning is that the edge costs may be unknown when starting the graph search. For example, in unknown environments, it is not known whether the edges are in collision with obstacles or not. The procedure of computing the cost of an edge is called *edge evaluation*. Edge evaluations are performed online as part of the graph search algorithm to compute the edge cost (e.g., perform collision checking). Even if the environment is known before starting the graph search, evaluating all edges before starting the graph search is unnecessary, as only part of the edges will be visited when searching for the solution. Thus, performing edge evaluation online saves time compared to evaluating all edges before graph search. In many robotic motion planning problems, edge evaluation (collision checking) is the major computational bottleneck [13]. Lazy search algorithms that aim to reduce the number of edge evaluations have been developed [14], [15]. Lazy search algorithms adopt a heuristic edge cost. Such a cost provides a lower bound of the true edge cost but is easier to compute. This heuristic cost is used to guide the graph search so that edge evaluation is done only when necessary. LazySP is proven to be optimal in the sense that minimizes edge evaluations [16].

Lazy search algorithms reduce edge evaluations at the expense of extra graph operations (vertex expansion, calculating the current best plan), which introduce extra computational overhead compared to A*. This article proposes a new efficient graph search algorithm, namely, Lazy Edged-based A* (LEA*), for SAPF. LEA* is vertex optimal and has improved edge efficiency compared to A*. It introduces less computational overhead compared to other lazy search algorithms.

The contributions of this article are summarized as follows.

- 1) We propose Conflict-Based Lazy Search (CBLs), a new MAPF algorithm focusing on multimanipulator path planning. CBLs enhances CBS with a precomputed, lazily evaluated graph and a new SAPF algorithm using lazy search.

- 2) We propose the LEA* algorithm for solving SAPF problems. LEA* uses lazy search and an edge queue to reduce the number of edge evaluations. We show the completeness, optimality, and optimal vertex efficiency of LEA*. We also show the improved edge efficiency over the A* algorithm through numerical comparisons.
- 3) A lazily-evaluated graph for efficient online multimanipulator planning. A joint space graph with controlled sparsity is precomputed for manipulator planning.
- 4) Detailed simulation and experiment studies of the proposed algorithm are conducted to compare LEA* with previous SAPF algorithms. We also compare CBLs with CBS and RRT-connect [17] demonstrating CBLs's superior performance.

II. RELATED WORKS

A. Lazy Search for SAPF

Graphs offer a powerful abstraction tool for robot path and motion planning. When combined with sampling-based methods, algorithms such as PRM* [4], RRT* [4], [18], and BIT* [6] can solve planning problems in high-dimensional spaces. They construct an explicit graph of robot configurations and find the shortest path by exploiting this graph. Implicit graphs may be defined using state lattices or motion primitives [19], [20], [21]. Motion primitives are precomputed and beneficial in dealing with kinematic, nonholonomic, and differential constraints. Planning problems for ground vehicles and micro aerial vehicles are studied in [20] and [21], respectively.

The A* algorithm is a popular search algorithm that is vertex optimal [12] but may lead to excessive edge evaluations. Recent studies show that the number of edge evaluations can be reduced by employing a lazy approach [14], [15], [22], [23]. For instance, LWA* [14] uses a one-step lookahead to postpone edge evaluation and uses duplicated vertices in the vertex queue. First, a valid vertex with an estimated cost is popped from the vertex queue. Then, after edge evaluation, it is inserted into the queue again with the true cost. LazySP [15] uses an infinite-step lookahead and is shown to be edge optimal (evaluating the minimum number of edges). LRA* [16] interpolates between LWA* and LazySP and uses a constant lookahead in the interval $[1, \infty]$. As the lookahead steps increase, the number of edge evaluations required to find the shortest path decreases, while the additional graph operations increase. These lazy search algorithms achieve better edge efficiency at the cost of considerable overhead compared to A*.

In [24], the generalized lazy search algorithm (GLS) that unifies LWA*, LazySP, and LRA* was proposed. GLS introduces an EVENT function and a SELECTOR function. It admits various realizations through different choices of EVENT and SELECTOR functions. With specific EVENT and SELECTOR instances, GLS reduces exactly to the LazySP or LRA* algorithm. GLS may employ informed EVENT and SELECTOR functions to improve the planning efficiency of LazySP/LRA*. However, this requires additional prior information (e.g., priors on edge validity learned from experience), which may be unavailable or may deviate from the true edge collision probabilities. A

lazy incremental algorithm dealing with dynamically changing graphs was introduced in [25] by combining the idea of lazy search and lifelong planning [26].

B. MAPF and Conflict-Based Search

To solve MAPF problems, earlier works search the augmented space of all agents. The single and dual-arm planning problem was studied in [7]. It uses the ARA* algorithm to search an implicit graph defined by motion primitives. By combining individual agents as a single ‘joint agent’, the dimensionality of planning space becomes too large. These methods may find solutions that are far from the optimal one, and the convergence rate to the optimal solution is slow.

CBS [1] uses a two-level planning paradigm. It uses the fact that SAPF is much simpler than MAPF. The low-level of CBS solves many SAPF problems and uses high-level planning to resolve conflicts between agents. Several enhancements of CBS have been proposed where admissible heuristics are developed to guide the high-level search [27], [28]. Disjoint splitting for conflict resolution was proposed in [29] to reduce duplication of search efforts. The CBS-Budget (CBSB) algorithm was proposed in [30]. A Class-Ordered A* algorithm, which finds the shortest path with a minimal number of conflicts that is upper bounded in terms of length, is used for the low-level planning of CBSB. Continuous-time Conflict-Based Search was proposed in [2] to remove time discretization, where Safe Interval Path Planning was used for single-agent pathfinding in CBS.

All methods mentioned above consider a grid world. Multi-robot kinodynamic motion planning was studied in [31] based on CBS and discontinuity-bounded A*. The proposed method is used for multiple vehicle planning. Multimaneulator planning was studied in [9]. Incomplete constraints that allow fast search space pruning are introduced to prune the search space of CBS. The resulting generalized ECBS algorithm obtains bounded suboptimality. Note that the method proposed in this article and [9] are complementary. The authors in [9] used incomplete constraints as a heuristic to reduce the number of SAPF problems needed by CBS, and used weighted-A* for SAPF [10], while our method mainly focuses on more efficient SAPF methods. Our article uses complete constraints (vertex and edge constraints); incomplete constraints could be integrated following [9].

C. Manipulator Planning

Manipulator motion planning for human-manipulator interaction and collaboration was studied in [32], where human motion prediction was utilized for proactive manipulator planning. In [33], the multiarm collaborative pick-and-place task was studied. The closed-chain constraint is considered using a probabilistic roadmap algorithm. The task considers a relatively structured environment, reducing the complexity of obstacle avoidance. Prioritized planning was used in [34], where a dual arm is divided into a main arm and a slave arm. Based on the RRT algorithm, the main arm planning is performed first, then the main arm is treated as a dynamic obstacle with a known trajectory in the slave arm planning. An asymptotically optimal multirobot motion planning algorithm was proposed in [8],

where an informed search is performed on the tensor product of roadmaps. Multimaneulator planning using CBS is studied in [35], where online-generated experience is exploited to reduce single-manipulator planning time at the low level of CBS.

III. PROBLEM FORMULATIONS

Given n manipulators working in a shared environment, we label them as $a_1, \dots, a_n$. Let the joint angles of manipulator i at time k be $\dot{j}_{i,k}$, where $\dot{j}_{i,k}$ is a vector with dimension equal to the number of joints of manipulator i . The set of joint angles for all n manipulators at time k is given by $J_k = \{\dot{j}_{1,k}, \dot{j}_{2,k}, \dots, \dot{j}_{n,k}\}$. The set J_k is a *valid joint angle set* if there is no manipulator self-collision, manipulator-obstacle collision, or manipulator-manipulator collision at time k . A joint angle path for all n manipulators $P = (J_1, \dots, J_K)$ is a *valid joint angle path* if the joint angle set J_k is valid for all time $k = 1, \dots, K$ along the path.

The multimaneulator planning problem is defined as finding a valid path P from a start joint angle set J_s to a goal set J_g . The cost of P is the sum of all individual agent path lengths. We adopt conflict-based search (CBS) [1] to solve the multimaneulator planning problem. CBS is a two-level algorithm. The low-level solves a series of SAPF problems and finds optimal paths for the individual agents. For each agent a_i , SAPF involves searching a precomputed graph to find a valid shortest length path $p_i = \{\dot{j}_{i,1}, \dots, \dot{j}_{i,K}\}$ from the start $\dot{j}_{i,1}$ to goal $\dot{j}_{i,K}$. Each path is free from manipulator self-collisions and manipulator-obstacle collisions, but overlooks manipulator-manipulator collisions. The high-level planner checks for conflicts among all individual paths. Any manipulator-manipulator collision results in a conflict. If no conflict is found, a valid path has been found. Otherwise, the conflict is split into new constraints. Low-level SAPF problems with the new constraints are then solved.

We use lazy search to solve the low-level SAPF. Collision checking is one of the computational bottlenecks of manipulator planning. To reduce the total number of collision checks, lazy search methods use an edge cost heuristic before computing the true edge cost. A graph G is given by a vertex set V and an edge set E . We define the edge cost heuristic $\hat{w}(e)$ for an edge e as the length of the edge. The true edge cost $w(e)$ is given by

$$w(e) = \begin{cases} \hat{w}(e), & \text{if } e \text{ is not in collision,} \\ \infty, & \text{if } e \text{ is in collision.} \end{cases} \quad (1)$$

In the following sections, we first summarize our method for constructing a graph for manipulator planning. Then, the lazy search algorithm along with its properties is introduced. Finally, we combine the two components with CBS to form the CBLS algorithm.

IV. LAZILY EVALUATED SPARSE GRAPH

CBLS uses a precomputed graph for online pathfinding. Without loss of generality, we consider the path planning problem for several identical manipulators. To deal with different manipulators, we need to compute one graph for each manipulator model.

Algorithm 1: SparseSampling.

Input: parameters N , max_count , m_dist
Output: vertex set V

```

1  $V \leftarrow \emptyset$ ;
2  $\text{count} \leftarrow 0$ ;
3 while  $V.\text{size} < N$  and  $\text{count} < \text{max\_count}$  do
4    $v \leftarrow \text{Sampling}$ ;
5   if  $\text{selfCollision}(v)$  then
6      $\text{continue}$ ;
7    $d \leftarrow \text{mindist}(V, v)$ ;
8   if  $d < m\_dist$  then
9      $\text{count} = \text{count} + 1$ ;
10     $\text{continue}$ ;
11    $\text{count} \leftarrow 0$ ;
12    $V \leftarrow V \cup \{v\}$ 
13 return  $V$ ;
```

The vertex sampling method is given in Algorithm 1. In Line 4, we draw a sample from the joint angle space of one manipulator using uniform random sampling. This step is repeated until a self-collision-free sample v is found (Lines 5–6). We introduce the parameter m_dist as the lower-bound distance between sampled vertices. The mindist function returns the minimum distance d between v and all other vertices in V (Line 7). Vertex v is rejected if d is less than m_dist (Lines 8–10). Otherwise, v is added to the vertex set V . We use count to keep track of the number of consecutive rejected vertices. In Line 3, the algorithm terminates if a maximum sample number N is reached or count reaches max_count .

The parameter m_dist controls the sparsity of the vertex set, where all vertices are at least m_dist away from each other. Starting with a large m_dist , we use a small number of samples to cover the joint angle space of the manipulator. Using the most recent vertex set V and gradually reducing m_dist , we obtain an incrementally denser graph using Algorithm 1, where the vertices are still approximately equally spaced. Algorithm 1 is only for vertex sampling. To construct a graph, we connect each vertex to its neighbors following the PRM* algorithm.

We check self-collisions during graph construction, while manipulator–obstacle and manipulator–manipulator collisions are handled via online lazy search. Because specific obstacles are not considered during graph construction, the resulting graph is environment-independent and can be reused across environments with different obstacle configurations. We exploit offline computation and online lazy search to reduce online computation. Other online algorithms, such as RRT* and motion primitive-based algorithms, also need to perform online self-collision checking. Previous MAPF formulations often assume a grid world that is fixed to a global frame, and all agents move within this grid world. In contrast, in our manipulator planning setting, the graph is agent-centric. Each manipulator has its own graph. Two manipulators will occupy different workspaces even if they have the same joint angles.

Algorithm 2: LEA*.

Input: start v_s , goal v_g for one arm
Output: planned path for one arm

```

1  $Q_E \leftarrow \{(v_s, v_i) | \forall v_i \in \text{Succ}(v_s)\}$ ;
2 while  $Q_E \neq \emptyset$  do
3    $(v_0, v_1) \leftarrow Q_E.\text{Pop}$ ;
4   if  $g(v_g) \leq f((v_0, v_1))$  then
5      $\text{return RetrievePath}$ 
6    $\text{collisionFree} \leftarrow \text{CollisionCheck}((v_0, v_1))$ ;
7   if  $\text{collisionFree}$  then
8      $g_{\text{new}} \leftarrow g(v_0) + w((v_0, v_1))$ ;
9     if  $g_{\text{new}} < g(v_1)$  then
10       $g(v_1) \leftarrow g_{\text{new}}$ ;
11       $v_1.\text{parent} \leftarrow v_0$ ;
12       $Q_E \leftarrow Q_E \cup \{(v_1, v_i) | \forall v_i \in \text{Succ}(v_1)\}$ ;
```

V. LEA* ALGORITHM FOR SAPF

In this section, we provide a detailed description of the proposed LEA* algorithm. LEA* is an efficient algorithm for SAPF, which, in turn, helps reduce the planning time for CBS. LEA* uses lazy search and edge queue for tree expansion. We design LEA* to have minimal changes compared to the original A* algorithm, thus having restrained computational overhead, by searching the lazily evaluated graph (Section IV).

Similar to A*, we define the *cost-to-come* and the heuristic *cost-to-go* of the vertex v as $g(v)$ and $h(v)$, which are the cost from v_s to v given the current search tree and the estimated cost from v to v_g , respectively. Here, $h(\cdot)$ is an admissible and consistent heuristic. The total estimated cost of the path passing through v is given by

$$f(v) = g(v) + h(v). \quad (2)$$

An edge is given by $e = (v_0, v_1)$, where v_0 is the source vertex and v_1 is the target vertex. We define the total estimated cost of the path passing through edge e as

$$f(e) = g(v_0) + \hat{w}(e) + h(v_1). \quad (3)$$

The LEA* algorithm is given in Algorithm 2. In Line 1, the edges from the starting vertex v_s to its successors v_i are added to the edge queue Q_E . The edges in Q_E are ordered by their f -values given by (3). In Line 3, the best edge $e = (v_0, v_1)$ with the smallest f -value is popped from Q_E . The f -value is used in the termination condition in Line 4. If the cost-to-come of goal vertex $g(v_g)$ is less than or equal to $f((v_0, v_1))$, we have found the shortest path. We can retrieve the shortest path from v_s to v_g using the parent information of the vertices. Collision checking is performed in Line 6. If the edge is collision-free, i.e., no manipulator self-collision or manipulator-obstacle collision, the new cost-to-come of v_1 , g_{new} , is computed in Line 8. If g_{new} is better than the current cost-to-come $g(v_1)$, $g(v_1)$ and the parent of v_1 are updated (Line 9–11). Finally, in Line 12, the edges from v_1 to its successors v_i are added to the edge queue Q_E .

A simple illustration of the LEA* and its difference with the A* algorithm is given in Fig. 2. Fig. 2(a)–(c) are the steps of




Fig. 1. Manipulator planning setup. Given a valid start-goal configuration pair of the manipulators, the CBLs algorithm generates a collision-free joint angle path. A calibrated external camera is used to track the locations of the obstacles.

LEA* and Fig. 2(d)–(f) are the steps of A*. At the beginning of LEA*, four edges are added to the edge queue as shown in Fig. 2(a). Note that these edges are only lazily evaluated using the estimated cost $\hat{w}$, and $\hat{w}$ is used to compute their f -values according to (3). The best edge e_4 is selected for edge evaluation. In Fig. 2(b), e_4 is evaluated using collision checking, and the next edges are added to the queue. Repeating this process, the best edge e_7 in the current queue is selected for evaluation. For A*, all outgoing edges are evaluated when expanding a vertex. In this example, by using lazy evaluation and an edge queue, LEA* only evaluated two edges to find the solution, while A* evaluated eight edges.

A. Algorithm Analysis

In this subsection, we analyze the completeness, optimality, and vertex efficiency of LEA*. We first show that LEA* is complete.

Proposition 1: If at least one solution exists for the single-agent graph search problem, LEA* will return a solution. Otherwise, it will return that no solution exists.

Proof: We first show that the algorithm terminates with $Q_E = \emptyset$ and $g(v_g) = \infty$ when no solution exists. All edges have positive edge costs. The new edge (v_1, v_i) is added to Q_E in Line 12, Algorithm 2, only when we find a better path to v_1 (Line 9). Since the cost-to-come to every vertex is decreasing and lower bounded, Line 12 will only run a finite number of iterations. Thus, $Q_E = \emptyset$ after executing Line 3 a finite number of iterations. Therefore, LEA* will terminate with $Q_E = \emptyset$ and $g(v_g) = \infty$ if no solution exists.

Next, we consider the case when the problem has a solution. LEA* terminates when $g(v_g) \leq \min_{e \in Q_E} f(e)$ (Line 4) or $Q_E = \emptyset$ (Line 2). If the algorithm terminated with $g(v_g) \leq \min_{e \in Q_E} f(e)$, we have $g(v_g) < \infty$, which implies a solution has been found.

Let $(e_0, e_1, \dots, e_{n-1})$ (equivalently, $(v_s, v_1, v_2, \dots, v_g)$) be a solution path. Initially, $e_0 \in Q_E$. After evaluating e_0 , e_1 is added to Q_E . Similarly, after evaluating e_1 , edge e_2 is added to Q_E . If the algorithm terminates with $Q_E = \emptyset$, it must have evaluated

$e_0, e_1, \dots, e_{n-1}$. Thus, the algorithm has found this solution, and the returned solution is at least as good as $(e_0, e_1, \dots, e_{n-1})$. $\square$

Next, we prove the optimality properties of LEA*.

Proposition 2: If at least one solution exists for the graph search problem, LEA* finds the minimum cost solution.

Proof: Let $\tau^* = (v_s, v_1^*, v_2^*, \dots, v_{n-1}^*, v_g)$ (equivalently $(e_0^*, e_1^*, \dots, e_{n-2}^*, e_{n-1}^*)$) be an optimal solution path with path cost c^* . Let $\tau = (v_s, v_1, v_2, \dots, v_{n-1}, v_g)$ (equivalently $(e_0, e_1, \dots, e_{n-2}, e_{n-1})$) be the solution path returned by LEA*. The cost of τ is c . To prove optimality, we need to show that c is equal to c^* .

To reach a contradiction, let us assume $c^* < c$. Note that $\max\{f(e_0^*), f(e_1^*), \dots, f(e_{n-2}^*), f(e_{n-1}^*)\} \leq c^*$. Also, $g(v_g) = c \leq \min_{e \in Q_E} f(e)$ holds when LEA* terminates. At the beginning of LEA*, $e_0^* \in Q_E$. Therefore, e_0^* must have been evaluated before $c \leq \min_{e \in Q_E} f(e)$ is true. After evaluating e_0^* , the vertex v_1^* is added to the expansion tree, and edge e_1^* is added to Q_E . By repeating this analysis, $e_1^*, \dots, e_{n-2}^*, e_{n-1}^*$ must all have been evaluated before $c \leq \min_{e \in Q_E} f(e)$ is true. After evaluating $e_1^*, \dots, e_{n-1}^*$, we have found a path to v_g and $f(v_g) \leq c^*$. Since the $f(v_g)$ is nonincreasing, the algorithm will never return a path with cost $c > c^*$. Therefore $c = c^*$. $\square$

The next result illustrates the optimal vertex efficiency property of LEA*.

Proposition 3: LEA* has the same vertex efficiency as A*. Furthermore, the evaluated edge set of LEA* is a subset of the evaluated edge set of A*.

Proof: For LEA*, we call evaluating any outgoing edge of vertex v as “expanding v .” A* expands v by evaluating all outgoing edges of v . Vertex v is expanded in LEA* if a subset of its outgoing edges is evaluated.

From the previous section, we know that LEA* finds an optimal path. Let $\tau^* = (e_0^*, e_1^*, \dots, e_{n-2}^*, e_{n-1}^*)$ be the path found by LEA*, and c^* be the cost of τ^* . Let $T = (V, E)$ and $T^e = (V^e, E^e)$ be the expansion tree grown by A* and LEA*, respectively. Let W^v and W^e be the set of vertices expanded by A* and LEA*, respectively. By showing $W^e \subseteq W^v$, we prove the optimally efficient search of LEA*. Note that if $W^e \subseteq W^v$, we have $V^e \subseteq V$.

Assume $W^e \not\subseteq W^v$. Then, there exists v , such that $v \notin W^v$, $v \in W^e$, $v \in V^e$ and $v \in V$. To show this, we start with $v_0 \in W^e$, and $v_0 \notin W^v$. Then, $v_0 \in V^e$ (vertex must be in the tree for it to be expanded). If $v_0 \notin V$, we find v_1 , which is the parent of v in T^e . Then, we have $v_1 \in V^e$, $v_1 \in W^e$, and $v_1 \notin W^v$. If $v_1 \notin V$, we repeat this process by finding its parent vertex in V^e , and one of these parents v_i must satisfy $v_i \in V$ since V and V^e share the same root vertex.

From $v \notin W^v$, we have $c^* \leq f(v) = g(v) + h(v)$. From $v \in W^e$, we have that there exists $v_i \in \text{Succ}(v)$, such that $f((v, v_i)) = g(v) + \hat{w}((v, v_i)) + h(v_i) \leq c^*$. Note that $f(v) \leq f((v, v_i))$. Then, we have

$$c^* \leq f(v) \leq f((v, v_i)) \leq c^* \quad (4)$$

which only holds when $f(v) = f((v, v_i)) = c^*$. Note that

$$\max\{f(e_0^*), f(e_1^*), \dots, f(e_{n-2}^*), f(e_{n-1}^*)\} \leq c^*. \quad (5)$$




Fig. 2. Illustration of the steps of LEA* ((a)–(c)) and A* ((d)–(f)). The set below each figure is the current edge queue or vertex queue. Dashed red lines are edges that are lazily evaluated. Solid red lines are evaluated edges. In (a), four edges are added to the edge queue, and e_4 is selected for evaluation. In (b), e_4 is evaluated, the next edges are added to the edge queue, and e_7 is selected. By using lazy evaluation and an edge queue, LEA* only evaluated two edges while A* evaluated eight edges.

Using (5), in order for (v, v_i) to be evaluated by LEA*, there exist $e_i^* = (v_0, v_1) \in \tau^*$, such that $f(e_i^*) = c^*$ and $g(v) < g(v_0)$. Otherwise, all $e_i^* \in \tau^*$ have a higher priority than (v, v_i) and (v, v_i) will not be evaluated. Using $f(v) = c^*$ and $g(v) < g(v_0) \leq g(v_g)$, v has a higher priority than v_g . Therefore, v must be expanded by A* before v_g is selected from the vertex queue. Therefore, $v \in W^v$, contradicting $v \notin W^v$. Thus, $W^e \subseteq W^v$ and LEA* is optimally vertex efficient.

Since A* evaluates all outgoing edges of $v \in W^v$ and LEA* evaluates a subset of outgoing edges of $v \in W^e$, LEA* evaluates fewer edges than A*.

LEA* is an efficient SAPF algorithm. To integrate it into CBLS, a time dimension is added, similar to time-expanded A* [37], and a constraint set is added. Denote the constraint set as $\mathcal{C}$, which will be clarified in Section VI. We add two lines after Line 5 to check if $(v_0, v_1) \in \mathcal{C}$ or not. If $(v_0, v_1) \in \mathcal{C}$, this edge has a conflict with other manipulators and is skipped. The algorithm moves to its next iteration.

VI. CBLS ALGORITHM FOR MAPF

CBLS uses a precomputed graph and makes use of LEA* for MAPF. It solves constrained SAPF problems for each manipulator, checks for any conflict between planned paths, and resolves the conflict by adding new constraints for SAPF problems. It repeats this process until a solution is found or no solution can be found.

In the original CBS algorithm, vertex conflicts and edge conflicts are defined for grid world problems. For manipulator planning, we treat a vertex conflict as a special case of an edge conflict. Therefore, we only consider edge conflicts. After solving the SAPF problems for each manipulator, we obtain the paths of agent i and k which are given by $p_i = \{j_{i,1}, j_{i,2}, \dots, j_{i,L}\}$ and $p_k = \{j_{k,1}, j_{k,2}, \dots, j_{k,M}\}$, respectively. Note that joint angle

Algorithm 3: CBLS.

Input: start and goal queries for n arms

Output: planned paths for n arms

```

1 for  $i = 1, \dots, n$  do
2    $\mathcal{C}_i \leftarrow \emptyset$ ;
3    $(s_i, \text{cost}_i) = \text{LEA}^*(\text{arm}_i, \mathcal{C}_i)$ ;
4    $\text{Root}.\mathcal{C}_i \leftarrow \mathcal{C}_i$ ;  $\text{Root}.s_i \leftarrow s_i$ ;
5    $\text{Root}.\text{cost}_i = \text{cost}_i$ ;
6  $Q \leftarrow \{\text{Root}\}$ ;
7 while  $Q \neq \emptyset$  do
8    $P \leftarrow Q.\text{pop}$ ;
9    $\text{Conflict} \leftarrow \text{CheckConflict}(P)$ ;
10  if  $\text{Conflict} = \emptyset$  then
11    return  $P$ 
12  for  $\text{arm}_i \in \text{Conflict}$  do
13     $\mathcal{C}'_i \leftarrow P.\mathcal{C}_i \cup \text{Conflict}.\text{arm}_i$ ;
14     $(s'_i, \text{cost}'_i) = \text{LEA}^*(\text{arm}_i, \mathcal{C}'_i)$ ;
15     $P_{\text{new}} \leftarrow P$ ;
16     $P_{\text{new}}.\mathcal{C}_i \leftarrow \mathcal{C}'_i$ ;  $P_{\text{new}}.s_i \leftarrow s'_i$ ;
17     $P_{\text{new}}.\text{cost}_i \leftarrow \text{cost}'_i$ ;
18     $Q.\text{push}(P_{\text{new}})$ ;

```

j is a vertex v of the graph since we sample vertices in the joint angle space. We assume that each segment of the path (edge) takes one unit of time to travel. This assumption is reasonable as graph vertices are approximately equally spaced. Then p_i takes L units of time, and p_k takes M units of time. If $L > M$, we can always append an appropriate number of $j_{k,M}$ to p_k such that p_i and p_k have the same length. This is also true for the case $L < M$.




Fig. 3. Planning results for a 7DOF manipulator.

When checking for conflicts between p_i and p_k , we sequentially check for any collisions between edges in p_i and p_k . If a collision happens between $(j_{i,\ell}, j_{i,\ell+1})$ and $(j_{k,\ell}, j_{k,\ell+1})$, a conflict is formed given by $((j_{i,\ell}, j_{i,\ell+1}), (j_{k,\ell}, j_{k,\ell+1}), \ell)$. This conflict states that agent i and k cannot move along the edge $(j_{i,\ell}, j_{i,\ell+1})$ and $(j_{k,\ell}, j_{k,\ell+1})$ simultaneously at time ℓ . To resolve this conflict, a new constraint $((j_{i,\ell}, j_{i,\ell+1}), \ell)$ is added to the constraint set C_i and a new constraint $((j_{k,\ell}, j_{k,\ell+1}), \ell)$ is added to C_k . In the next iteration of CBLS, SAPF problems with the updated constraint set are solved for agents i and k .

The CBLS algorithm is given by Algorithm 3. In Lines 1–5, individual SAPF problems for each agent are solved using LEA*. Initially, the constraint set c_i for each agent i is empty. s_i and $cost_i$ are the planned path and path cost for agent i , respectively. The nodes in the conflict tree save the planning solutions for all agents. The initial solution is saved in the *Root* node. The queue Q is initialized with the *Root* node (Line 6). The nodes in Q are prioritized by the total path cost of all agents. In Line 8, the node with the lowest cost is popped from Q . Manipulator-manipulator collision checking for paths in P is performed in Line 9. If no conflict is found among the paths, a solution is returned (Lines 10–11). Otherwise, the conflict is resolved in Lines 12–18. For each arm_i affected by the conflict, the constraint set is updated (Line 13), and a SAPF problem is solved for arm_i using LEA* (Line 14). We update P with the new solution of arm_i (Lines 16–17), and a new node P_{new} is added to the queue Q (Line 18).

VII. SIMULATION AND EXPERIMENTS

In this section, we first provide simulation results for LEA* and CBLS. For single-manipulator planning, we compare three SAPF algorithms: LEA, A*, and LazySP. For multimaniplator planning, we consider scenarios with up to five manipulators operating in close proximity and compare the proposed CBLS algorithm with CBS and RRT-connect. CBS uses A* for low-level SAPF planning and RRT-connect plans in the composite state space. Then, real-world experiments based on the CBLS algorithm are conducted with two UR5 manipulators.

A. LEA* Simulation Results

We consider tabletop planning problems for a 7-DOF KUKA manipulator. The PyBullet simulator [38] is adopted for the simulations. The planning environment is shown in Fig. 4. Cubic obstacles are randomly placed on the tabletop. We generate different environments by varying the number, the location, and the size of obstacles in each environment. The environment is




Fig. 4. Single manipulator tabletop example.

divided into sparse, medium, and cluttered environments where the number of obstacles are 4, 8, and 12, respectively. A back wall and a ceiling are used to limit the operational space of the manipulator. The width, depth, and height of each obstacle are sampled uniformly and randomly from their respective intervals. The locations of the obstacles on the table are also sampled. We sampled 30 environments in total, and graphs with different sizes were precomputed. We consider small graphs, medium graphs, and large graphs with $N = 1000$, $N = 5000$, and $N = 10000$, respectively. We use the Manhattan distance to sum the path lengths across individual degrees of freedom and across different manipulators. We set $max_count = 2000$, $m_dist = 3, 2.7, 2.5$, respectively, and the connecting radius $r = 5, 4.5, 4.2$, respectively. By randomly sampling 50 start-goal queries for each environment and graph combination, the total number of planning problems is 4500.

The planning results for A*, LazySP, and LEA* are given in Table I. The three algorithms were tested using the same 4500 problems. The statistics vary considerably across different planning problems, and Table I reports the average results. A*, LazySP, and LEA* use a cost-to-go heuristic $h(\cdot)$ to guide the graph search. While searching with an admissible heuristic guarantees path optimality, using an inflated heuristic has been shown to reduce planning time at the cost of optimality [6], [7]. Bounded suboptimal solutions can be obtained by inflating $h(\cdot)$ with a factor ϵ [39]. Results with inflation factor $\epsilon = 2$ are also reported in Table I.

For A* and LEA*, the planning time is approximately proportional to the number of evaluated edges during the search. Using lazy edge evaluation, LEA* reduces edge evaluation significantly compared to A*. LazySP assumes all edges are collision-free and uses A* to find the best path. Then, the best path is checked for collision. If a collision happens, the graph is updated, and a new planning problem is solved using A*. LazySP is shown to evaluate the minimum number of edges

TABLE I
PLANNING RESULTS FOR A 7-DOF MANIPULATOR

	Planning time (s)			Edge evaluations			Path length		
	N=1,000	N=5,000	N=10,000	N=1,000	N=5,000	N=10,000	N=1,000	N=5,000	N=10,000
A* ($\epsilon=1$)	0.6773	1.7215	3.5747	331.29	917.36	1929.2	23.9935	22.1431	21.5779
LazySP ($\epsilon=1$)	0.1060	0.3527	0.6567	15.77	15.11	17.69	23.9935	22.1431	21.5779
LEA* ($\epsilon=1$)	0.0763	0.1123	0.1755	51.45	82.55	129.48	23.9935	22.1431	21.5779
A* ($\epsilon=2$)	0.2633	0.3800	0.4879	147.34	235.69	309.07	24.4210	22.4675	21.7728
LazySP ($\epsilon=2$)	0.0431	0.0433	0.0472	15.12	13.20	14.61	24.4210	22.4675	21.7728
LEA* ($\epsilon=2$)	0.0296	0.0234	0.0238	19.79	16.40	17.41	24.4210	22.4675	21.7728




Fig. 5. Results for different parameter N . The results are obtained by testing LEA* ($\epsilon = 2$) on the same set of problems by varying N .

(edge optimal) at the cost of repeatedly calling the A* algorithm in the inner loop. As shown in Table I, even if LazySP evaluates the minimum number of edges, LEA* is up to three times faster than LazySP for $N = 5000$ and $N = 10000$, as LazySP requires more graph operations such as priority queue updating and vertex expansions. LEA* uses the least amount of time to find the same solution compared to the other two algorithms. It performs 7-DOF manipulator planning at 30–50 Hz.

In Table I, the planning time and edge evaluations are further reduced with an inflation factor $\epsilon = 2$. More results for different inflation factors are shown in Fig. 3. With only a slight increase in path length, the planning time decreases considerably, and the edge evaluations of LEA* approach those of LazySP. A study regarding different graph sizes is provided in Fig. 5. LEA* with $\epsilon = 2$ is applied to the same set of planning problems with varying graph sizes N . As the graph size increases, the success rate improves and the path length decreases. Note that the planning time also decreases on average. This improvement is due to the inflation factor, which helps focus the search. Also, a denser graph reduces detours between the start and goal. As shown in Table I, the planning time may increase with larger N when $\epsilon = 1$.

B. CBLS Simulation Results

Planning problems for two-manipulator, three-manipulator, four-manipulator, and five-manipulator systems were studied.




Fig. 6. Multimanipulator planning.

The planning environments are shown in Fig. 6. Fig. 6(c) represents a shelf environment. In the other three environments, the obstacles are randomly sampled following the single manipulator case. We construct a graph with $N = 5000$, 10000 , 20000 , and 40000 vertices, respectively. We set $\text{max_count} = 2000$, $\text{m_dist} = 2.7, 2.5, 2.2, 2$, respectively, and the connecting radius $r = 4.5, 4.2, 4.0, 3.5$, respectively. For each environment, we sampled 20 start-goal queries.

CBLS, CBS, and RRT-connect were used to solve the same set of planning problems. Both CBLS and CBS used the same precomputed graphs (Section IV). The only difference is that CBLS uses LEA* for SAPF while CBS uses A* for SAPF. An additional time dimension is added to A* and LEA*. Each vertex can move to its neighbors or stay at its current position at the next time step.

The RRT-connect algorithm is a representative algorithm for solving high-dimensional planning problems. It is probabilistically complete but does not guarantee properties such as asymptotic optimality. For many planning problems, it is one of the fastest algorithms for finding a feasible solution; however, the resulting solutions are typically of low quality [36]. By combining the joint angle spaces of all manipulators, we use RRT-connect to solve planning problems with state spaces of up to 35 dimensions. Since RRT-connect is a randomized algorithm, it produces different solutions for the same problem across runs. We solve each problem ten times using RRT-connect and use the average results for comparison. CBLS and CBS are deterministic algorithms. For each planning problem, CBLS/CBS returns the same solution, with nearly identical planning times across runs.

TABLE II
PLANNING RESULTS FOR THREE MANIPULATORS

	Planning time (s)				Path length				Edge evaluations			
	N=5,000	10,000	20,000	40,000	N=5,000	10,000	20,000	40,000	N=5,000	10,000	20,000	40,000
CBS ($\epsilon=1$)	0.8770	1.3198	1.9067	2.1713	60.9560	51.6858	48.4792	46.1714	3817.1	7063.5	10786	13214
CBLS ($\epsilon=1$)	0.2586	0.3313	0.4719	0.6278	60.9560	51.6858	48.4792	46.1714	1266.2	1647.5	2162.3	2888.5
CBS ($\epsilon=2$)	0.2692	0.3967	0.4201	0.4863	62.2471	53.8270	50.9448	48.5605	1292.3	2071.1	2228.8	2660.8
CBLS ($\epsilon=2$)	0.0985	0.1107	0.1321	0.1982	62.2471	53.8270	50.9448	46.5605	383.65	395.97	342.09	432.5
RRT-connect	0.7631				83.0737				109.4276			

TABLE III
PLANNING RESULTS FOR FIVE MANIPULATORS

	Planning time (s)				Path length				Edge evaluations			
	N=5,000	10,000	20,000	40,000	N=5,000	10,000	20,000	40,000	N=5,000	10,000	20,000	40,000
CBS ($\epsilon=1$)	0.3895	0.2646	0.4970	0.4327	85.7357	76.3592	71.5091	70.048	2096.6	1411.9	2709.3	2778.6
CBLS ($\epsilon=1$)	0.2417	0.2113	0.2971	0.3918	85.7357	76.3592	71.5091	70.048	1125.7	823.52	1103.5	1536.1
CBS ($\epsilon=2$)	0.1025	0.1206	0.1486	0.2647	88.9099	81.1278	76.0302	74.7359	334.35	441.09	601.39	1254.2
CBLS ($\epsilon=2$)	0.1231	0.1278	0.1467	0.2208	88.9099	81.1278	76.0302	74.7359	171.92	143.99	141.74	193.29
RRT-connect	2.5072				174.0136				280.6029			




Fig. 7. Multimaneulator planning Success Rate (SR). SR decreases as the number of manipulators increases. SR increases as the graph size parameter N increases.

We set the maximum planning time to 15 seconds. The planning query fails if the algorithm returns no solution or if the planning time exceeds 15 seconds. The success rate of multimaneulator planning is given in Fig. 7. For CBLS, the planning success rate increases as the graph size N increases. The success rate decreases with the number of manipulators. RRT-Connect achieves high success rates for two- and three-manipulator planning; however, its performance drops rapidly for four and five manipulators, as the planning time often exceeds the 15-s limit.

Multimaneulator planning results are given in Tables II and III. The planning results of two manipulators and four manipulators can be found at <https://youtu.be/qC--BcjcdwA>. The planning time of CBLS outperforms CBS and RRT-connect for all cases, on average. Compared to RRT-connect, which directly composes the DOF of all manipulators, our CBLS combines precomputation, fast SAPF, and CBS. CBLS and CBS also find paths with shorter lengths compared to RRT-connect. In CBLS and CBS, edge evaluation requires checking the collision of an edge for a single manipulator, while in RRT-connect, edge evaluation requires checking the collision of all manipulators. From the results of Tables II and III, it is evident that using an inflated heuristic is beneficial. It greatly decreases the planning time, returning bounded suboptimal solutions that only increase the (resolution) optimal path length by a small factor.




Fig. 8. Planning with dynamic obstacles.

Dynamic obstacles can be handled by replanning. At every time step, the predicted trajectories of the moving obstacles within a future time horizon are used for planning. Beyond this horizon, the dynamic obstacles are not considered. Thus, the planned multimaneulator path is guaranteed to be collision-free within the upcoming horizon. As the manipulators move along the planned path, the obstacle trajectories for the next horizon become available. Replanning is triggered if any potential collision is detected. This procedure is repeated until the manipulators reach their goals. Path planning for three manipulators, while avoiding dynamic obstacles, is demonstrated in Fig. 8. The planning time is correlated with the duration of the predicted obstacle trajectories. In this example, the planning frequency is above 5 Hz, and the time horizon can be set to 0.2 s. With a faster planner, less information about the moving obstacles is required.

C. CBLS Experimental Results

We performed extensive real-world experiments using the CBLS algorithm for two-manipulator planning. The environment setup is shown in Figs. 1 and 9. Two UR5 manipulators were used, and obstacles (shown as wooden bricks) were randomly placed. Similar to the KUKA manipulator example, a simulation environment was built to replicate the experiment environment. We use an external camera and QR codes to identify and localize obstacles. After calibrating the camera and the manipulator system, we compute the obstacle locations in real time in the manipulator's coordinate frame. This approach




Fig. 9. Two planning examples. CBLS plans collision-free paths for the UR5 manipulators to move from their starting configurations to the goals (shown in transparent mode).

provides a relatively accurate model of the environment. Future work includes conducting more complex experiments, adopting more sophisticated perception methods, and planning under uncertainty. Our experiments show that CBLS achieves real-time planning. Manipulator self-collision, manipulator-obstacle collision, and manipulator-manipulator collision were avoided successfully during the experiments. Snapshots of the two manipulators switching goals are shown in Fig. 9. Animations of the experiments and simulations can be found at <https://youtu.be/qC--BcjdWA>.


VIII. CONCLUSION

This article introduces the CBLS algorithm for fast planning in high-dimensional spaces. CBLS is based on CBS, a start-of-the-art backbone for many MAPF algorithms and is tailored for manipulator planning. It introduces two enhancements to CBS for solving SAPF problems more efficiently. First, a lazily evaluated graph is used with controlled sparsity, which is precomputed for online planning. This graph is environment-independent and can be applied to any number of manipulators. Second, Lazy Edged-based A* (LEA*) is utilized as the low-level planner of CBS for efficient SAPF. It is shown that LEA* is complete, finds the optimal solution, is optimally vertex efficient, and reduces the number of collision checks using lazy search. Various simulations are conducted to demonstrate the theory. The advantages of LEA* and CBLS are illustrated by comparing them with their respective competing algorithms. Practical applications of CBLS on a two-manipulator system are also conducted to validate the proposed methods.

REFERENCES

- [1] G. Sharon, R. Stern, A. Felner, and N. R. Sturtevant, "Conflict-based search for optimal multi-agent pathfinding," *Artif. Intell.*, vol. 219, pp. 40–66, 2015.
- [2] A. Andreychuk, K. Yakovlev, P. Surynek, D. Atzmon, and R. Stern, "Multi-agent pathfinding with continuous time," *Artif. Intell.*, vol. 305, pp. 1–32, 2022.
- [3] L. E. Kavraki, P. Svestka, J. C. Latombe, and M. H. Overmars, "Probabilistic roadmaps for path planning in high-dimensional configuration spaces," *IEEE Trans. Robot. Automat.*, vol. 12, no. 4, pp. 566–580, Aug. 1996.
- [4] S. Karaman and E. Frazzoli, "Sampling-based algorithms for optimal motion planning," *Int. J. Robot. Res.*, vol. 30, pp. 846–894, 2011.
- [5] O. Arslan and P. Tsotras, "Use of relaxation methods in sampling-based algorithms for optimal motion planning," in *Proc. IEEE Int. Conf. Robot. Automat.*, Karlsruhe, Germany, May 2013, pp. 2421–2428.
- [6] J. D. Gammell, T. D. Barfoot, and S. S. Srinivasa, "Batch Informed Trees (BIT*): Informed asymptotically optimal anytime search," *Int. J. Robot. Res.*, vol. 39, no. 5, pp. 543–567, 2020.
- [7] B. Cohen, S. Chitta, and M. Likhachev, "Single-and dual-arm motion planning with heuristic search," *Int. J. Robot. Res.*, vol. 33, no. 2, pp. 305–320, 2014.
- [8] R. Shome, K. Solovey, A. Dobson, D. Halperin, and K. E. Bekris, "dRRT*: Scalable and informed asymptotically-optimal multi-robot motion planning," *Auton. Robots*, vol. 44, no. 3, pp. 443–467, 2020.
- [9] Y. Shaoul, R. Veerapaneni, M. Likhachev, and J. Li, "Unconstraining multi-robot manipulation: Enabling arbitrary constraints in ECBS with bounded sub-optimality," in *Proc. Int. Symp. Combinatorial Search*, Alberta, Canada, 2024, pp. 109–117.
- [10] B. J. Cohen, G. Subramania, S. Chitta, and M. Likhachev, "Planning for manipulation with adaptive motion primitives," in *Proc. IEEE Int. Conf. Robot. Automat.*, Shanghai, China, May 2011, pp. 5478–5485.
- [11] D. M. Saxena, M. S. Saleem, and M. Likhachev, "Manipulation planning among movable obstacles using physics-based adaptive motion primitives," in *Proc. IEEE Int. Conf. Robot. Automat.*, Xi'an, China, 2021, pp. 6570–6576.
- [12] P. E. Hart, N. J. Nilsson, and B. Raphael, "A formal basis for the heuristic determination of minimum cost paths," *IEEE Trans. Syst. Sci. Cybern.*, vol. 4, no. 2, pp. 100–107, Jul. 1968.
- [13] S. M. LaValle, *Planning Algorithms*, Cambridge, U.K.: Cambridge University Press, 2006.
- [14] B. Cohen, M. Phillips, and M. Likhachev, "Planning single-arm manipulations with n-arm robots," in *Proc. Robot.: Sci. Syst.*, Berkeley, CA, USA, Jul. 2014, pp. 1–9.
- [15] C. Dellin and S. Srinivasa, "A unifying formalism for shortest path problems with expensive edge evaluations via lazy best-first search over paths with edge selectors," in *Proc. Int. Conf. Automated Plan. Scheduling*, London, U.K., Jun. 2016, vol. 26, pp. 459–467.
- [16] A. Mandalika, O. Salzman, and S. Srinivasa, "Lazy receding horizon A* for efficient path planning in graphs with expensive-to-evaluate edges," in *Proc. Int. Conf. Automated Plan. Scheduling*, Delft, The Netherlands, Jun. 2018, vol. 28, pp. 476–484.
- [17] J. J. Kuffner and S. M. LaValle, "RRT-connect: An efficient approach to single-query path planning," in *Proc. IEEE Int. Conf. Robot. Automat.*, San Francisco, CA, USA, 2000, pp. 995–1001.
- [18] L. Jiang, S. Liu, Y. Cui, and H. Jiang, "Path planning for robotic manipulator in complex multi-obstacle environment based on improved_RRT," *IEEE/ASME Trans. Mechatron.*, vol. 27, no. 6, pp. 4774–4785, Dec. 2022.
- [19] M. Likhachev and D. Ferguson, "Planning long dynamically feasible maneuvers for autonomous vehicles," *Int. J. Robot. Res.*, vol. 28, no. 8, pp. 933–945, 2009.
- [20] M. Pivtoraiko, R. A. Knepper, and A. Kelly, "Differentially constrained mobile robot motion planning in state lattices," *J. Field Robot.*, vol. 26, no. 3, pp. 308–333, 2009.
- [21] S. Liu, N. Atanasov, K. Mohta, and V. Kumar, "Search-based motion planning for quadrotors using linear quadratic minimum time control," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, Vancouver, Canada, Sep. 2017, pp. 2872–2879.
- [22] R. Bohlin and L. E. Kavraki, "Path planning using lazy PRM," in *Proc. IEEE Int. Conf. Robot. Automat.*, San Francisco, CA, USA, Apr. 2000, pp. 521–528.
- [23] K. Hauser, "Lazy collision checking in asymptotically-optimal motion planning," in *Proc. IEEE Int. Conf. Robot. Automat.*, Seattle, WA, USA, May 2015, pp. 2951–2957.
- [24] A. Mandalika, S. Choudhury, O. Salzman, and S. Srinivasa, "Generalized lazy search for robot motion planning: Interleaving search and edge evaluation via event-based toggles," in *Proc. Int. Conf. Automated Plan. Scheduling*, Berkeley, CA, USA, Jul. 2019, vol. 29, pp. 745–753.
- [25] J. Lim, S. Srinivasa, and P. Tsotras, "Lazy lifelong planning for efficient replanning in graphs with expensive edge evaluation," in *IEEE/RSJ Int. Conf. Intell. Robots Syst.*, Kyoto, Japan, Oct. 2022, pp. 8778–8783.
- [26] S. Koenig, M. Likhachev, and D. Furcy, "Lifelong planning A*," *Artif. Intell.*, vol. 155, no. 1, pp. 93–146, 2004.
- [27] E. Boyarski et al., "ICBS: Improved conflict-based search algorithm for multi-agent pathfinding," in *Proc. Int. Joint Conf. Artif. Intell.*, 2015, pp. 740–746.
- [28] J. Li, A. Felner, E. Boyarski, H. Ma, and S. Koenig, "Improved heuristics for multi-agent pathfinding with conflict-based search," in *Proc. Int. Joint Conf. Artif. Intell.*, 2019, pp. 442–449.

- [29] J. Li, D. Harabor, P. J. Stuckey, A. Felner, H. Ma, and S. Koenig, "Dis-joint splitting for multi-agent pathfinding with conflict-based search," in *Proc. Int. Conf. Automated Plan. Scheduling*, Berkeley, CA, USA, 2019, pp. 279–283.
- [30] J. Lim and P. Tsotras, "CBS-Budget (CBSB): A complete and bounded suboptimal search for multi-agent path finding," *Artif. Intell.*, vol. 346, 2025, Art. no. 104349.
- [31] A. Moldagalieva, J. Ortiz-Haro, M. Toussaint, and W. Hönig, "db-CBS: Discontinuity-bounded conflict-based search for multi-robot kinodynamic motion planning," in *Proc. IEEE Int. Conf. Robot. Automat.*, Yokohama, Japan, 2024, pp. 14569–14575.
- [32] W. Liu, K. Eltouny, S. Tian, X. Liang, and M. Zheng, "Integrating uncertainty-aware human motion prediction into graph-based manipulator motion planning," *IEEE/ASME Trans. Mechatron.*, vol. 29, no. 4, pp. 3128–3136, Aug. 2024.
- [33] K. Jang, J. Baek, S. Park, and J. Park, "Motion planning for closed-chain constraints based on probabilistic roadmap with improved connectivity," *IEEE/ASME Trans. Mechatron.*, vol. 27, no. 4, pp. 2035–2043, Aug. 2022.
- [34] W. Shi, K. Wang, C. Zhao, and M. Tian, "Obstacle avoidance path planning for the dual-arm robot based on an improved RRT algorithm," *Appl. Sci.*, vol. 12, no. 8, pp. 1–17, 2022.
- [35] Y. Shaoul, I. Mishani, M. Likhachev, and J. Li, "Accelerating search-based planning for multi-robot manipulation by leveraging online-generated experiences," in *Proc. Int. Conf. Automated Plan. Scheduling*, Alberta, Canada, 2024, pp. 523–531.
- [36] A. Orthey, C. Chamzas, and L. E. Kavraki, "Sampling-based motion planning: A comparative review," *Annu. Rev. Control, Robot., Auton. Syst.*, vol. 7, pp. 1–25, 2023.
- [37] D. Silver, "Cooperative pathfinding," in *Proc. AAAI Conf. Artif. Intell. Interactive Digit. Entertainment*, 2005, pp. 117–122.
- [38] E. Coumans and Y. Bai, "Pybullet, A Python module for physics simulation for games, robotics and machine learning," 2016–2024. [Online]. Available: <http://pybullet.org>
- [39] J. Pearl, *Heuristics: Intelligent Search Strategies for Computer Problem Solving*, Reading, MA, USA: Addison-Wesley Longman Publishing, 1984.



Dongliang Zheng (Member, IEEE) received the B.Eng. degree in automation from Northeastern University, Shenyang, China, in 2015, the M.S. degree in control engineering from Shanghai Jiao Tong University, Shanghai, China, in 2018, and the Ph.D. degree in aerospace engineering from the Georgia Institute of Technology, Atlanta, GA, USA, in 2024.

He is currently an Associate Professor with the Shanghai Research Institute for Intelligent Autonomous Systems, Tongji University, Shanghai, China. His research interests include robot motion planning and autonomy.



Zhipeng Wang (Member, IEEE) received the M.S. degree in mechanical manufacturing and automation from Zhejiang University, Hangzhou, China, in 2011, and the Ph.D. degree in control science and engineering from Tongji University, Shanghai, China, in 2015.

He is currently an Associate Professor with the Department of Control Science and Engineering, Tongji University. His research interests include robot motion planning, mechatronics, and dynamics.



Siqi Wang received the B.E. degree in control science and engineering from Tongji University, Shanghai, China, in 2023. He is currently working toward the Ph.D. degree with the Shanghai Research Institute for Intelligent Autonomous Systems, Tongji University.

His research interests include the motion planning and safe collaborative operation of dual-arm robots.



Yuxi Lu (Member, IEEE) received the B.S. and M.S. degrees in engineering from the Tokyo University of Science, Tokyo, Japan, in 2019 and 2021, respectively, and the Ph.D. degree in medical engineering from the Graduate School of Science and Engineering, Chiba University, Chiba, Japan, in 2024.

He is currently an Assistant Professor with the Shanghai Research Institute for Intelligent Autonomous Systems, Tongji University, Shanghai, China. His current research interests include medical robotics and soft robotic systems.



Bin He (Senior Member, IEEE) received the Ph.D. degree in mechanical and electronic control engineering from Zhejiang University, Hangzhou, China, in 2001.

From 2001 and 2003, he held Postdoctoral research appointments with the State Key Lab of Fluid Power Transmission and Control, Zhejiang University. He is currently a Professor with the College of Electronics and Information Engineering, Tongji University, Shanghai, China. His current research interests include intelligent robot control, biomimetic microrobots, and wireless networks.



Hesheng Wang (Senior Member, IEEE) received the B.Eng. degree in electrical engineering from the Harbin Institute of Technology, Harbin, China, in 2002, and the M.Phil. and Ph.D. degrees in automation and computer-aided engineering from The Chinese University of Hong Kong, Hong Kong, in 2004 and 2007, respectively.

He is currently a Distinguished Professor with the Department of Automation, Shanghai Jiao Tong University, Shanghai, China. His research

interests include visual servoing, intelligent robotics, computer vision, and autonomous driving.

Dr. Wang is an Associate Editor for *Robotic Intelligence and Automation* and the *International Journal of Humanoid Robotics*, a Senior Editor for IEEE/ASME TRANSACTIONS ON MECHATRONICS, and an Editor-in-Chief for *Robot Learning*. He was an Associate Editor for IEEE TRANSACTIONS ON ROBOTICS from 2015 to 2019, and for IEEE TRANSACTIONS ON AUTOMATION SCIENCE AND ENGINEERING from 2021 to 2023. He was an Editor for the Conference Editorial Board of the IEEE Robotics and Automation Society from 2022 to 2024, the General Chair of IEEE ROBIO 2022 and IEEE RCAR 2016, and Program Chair of the IEEE ROBIO 2014 and IEEE/ASME AIM 2019. He will be the General Chair of IEEE/RSJ IROS 2025.



Panagiotis Tsotras (Fellow, IEEE) received the M.S. degree in aerospace engineering from Virginia Tech, VA, USA, in 1987, the Eng. Dipl. degree in mechanical engineering from NTUA, Athens, Greece, in 1986, and the M.S. degree in mathematics, and the Ph.D. degree in aerospace engineering from Purdue University, IN, USA, in 1992 and 1993, respectively.

He is the David and Andrew Lewis Chair Professor with the Daniel Guggenheim School of Aerospace Engineering at the Georgia Institute

of Technology, Atlanta, GA, USA. He has held visiting research appointments with MIT, JPL, INRIA Rocquencourt, and Mines ParisTech. His research interests include optimal control of nonlinear systems and ground, aerial, and space vehicle autonomy.

Dr. Tsotras is the recipient of the NSF CAREER award, the Outstanding Aerospace Engineer award from Purdue, and the Technical Excellence Award in Aerospace Control from IEEE. He has served on the Editorial Boards of the Transactions on Automatic Control, the *IEEE Control Systems Magazine*, the *AIAA Journal of Guidance, Control and Dynamics*, the *Dynamic Games and Applications*, and *Dynamics and Control*. He is a Fellow of AIAA and AAS.