

db-CBS: Discontinuity-Bounded Conflict-Based Search for Multi-Robot Kinodynamic Motion Planning

Akmaral Moldagalieva¹, Joaquim Ortiz-Haro^{1,2}, Marc Toussaint¹, and Wolfgang Hönig¹

Abstract—This paper presents a multi-robot kinodynamic motion planner that enables a team of robots with different dynamics, actuation limits, and shapes to reach their goals in challenging environments. We solve this problem by combining Conflict-Based Search (CBS), a multi-agent path finding method, and discontinuity-bounded A*, a single-robot kinodynamic motion planner. Our method, db-CBS, operates in three levels. Initially, we compute trajectories for individual robots using a graph search that allows bounded discontinuities between precomputed motion primitives. The second level identifies inter-robot collisions and resolves them by imposing constraints on the first level. The third and final level uses the resulting solution with discontinuities as an initial guess for a joint space trajectory optimization. The procedure is repeated with a reduced discontinuity bound. Our approach is anytime, probabilistically complete, asymptotically optimal, and finds near-optimal solutions quickly. Experimental results with robot dynamics such as unicycle, double integrator, and car with trailer in different settings show that our method is capable of solving challenging tasks with a higher success rate and lower cost than the existing state-of-the-art.

I. INTRODUCTION

Multi-robot systems have a wide-range of real-world applications including delivery, collaborative transportation, and search-and-rescue. One of the essential requirements to enhance the autonomy of a team of robots in these settings is being able to reach the goal quickly while avoiding collisions with obstacles and other robots. Moreover, the planned motions are required to respect the robots' dynamics which impose constraints on their velocity or acceleration. Considering the complexity of Multi-Robot Motion Planning (MRMP), the majority of existing solutions either make simplified assumptions like ignoring the robot dynamics or actuation limits, produce highly suboptimal solutions, or do not scale well with the number of robots.

In this paper, we present an efficient, probabilistically-complete and asymptotically optimal motion planner for a heterogeneous team of robots which takes into account constraints imposed by robot dynamics, see Fig. 1 for an example. Our method leverages the Multi-Agent Path Finding (MAPF) optimal solver Conflict-Based Search (CBS), the single-robot kinodynamic motion planner discontinuity-bounded A* (db-A*), and nonlinear trajectory optimization.




Fig. 1. Db-CBS solving multi-robot kinodynamic motion planning. Left: problem instance with a team of 8 heterogeneous robots: unicycle 1st order (rectangle), double integrator (circle), car with trailer (rectangle with box in the end). Right: Real-world experiment with four flying robots that have to pass through a narrow window to swap their positions.

We are motivated by the success of solving multi-agent path finding problems efficiently with bounded suboptimality guarantees using the family of Conflict-Based Search (CBS) algorithms. Our algorithm, called Discontinuity-Bounded Conflict-Based Search (db-CBS), performs a three-level search. On the low-level the single-robot motion is planned separately for each robot relying on pre-computed motion primitives obeying the robot's dynamics. The computed trajectories may result in inter-robot collisions, which are resolved one-by-one in the second-level search. Different from CBS, we allow the trajectories to be dynamically infeasible up to a given discontinuity bound during these two searches. Then, we use the output trajectories as an initial guess for a joint space trajectory optimization on the third level. This is repeated with a lower discontinuity bound in case of failure or if a lower-cost solution is desired. In this way, we combine the advantages of informed discrete search and optimization to find a near-optimal solution quickly.

The main contribution of this work is a new kinodynamic motion planner for multi-robot systems. We compare our method with other kinodynamic motion planning methods with the identical objective of computing time-optimal trajectories. Finally, we demonstrate that our planned motions can be safely executed on a team of flying robots.

II. RELATED WORK

Multi-Agent Path Finding (MAPF) assumes a discrete state space represented as a graph. A robot can move from one vertex along an edge to an adjacent neighbor in one step; robots cannot occupy the same vertex or traverse the same edge at the same timestep. For collision checking, vertices are often placed in a grid, but any-angle extensions exist [1]. Finding an optimal MAPF solution is NP-hard [2].

¹ Technical University of Berlin, Berlin, Germany
moldagalieva@tu-berlin.de

² Machines in Motion Laboratory, New York University, USA
Code: <https://github.com/IMRCLab/db-CBS>

We thank Nan Cai for help with flight experiments.

The research was partially funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) - 448549715.

Current algorithms plan paths for each robot individually and then resolve conflicts by fixing velocities [3], assigning priorities to robots [4] or adding path constraints as in CBS [5]. Although MAPF has practical applications in real-world problems, it ignores robot dynamics and can result in kinodynamically infeasible solutions. The planned paths may be followed using a controller and plan-execution framework [6]. However, this is suboptimal, may result in collisions, and does not generalize to all robot types.

Kinodynamic motion planning remains challenging, even for single robots. Search-based methods with motion primitives have been adapted to a variety of robotic systems, including high-dimensional systems [7]. Those motion primitives are on a state lattice, thus they are able to connect states properly and can be designed manually [8]. Afterwards, any variant of discrete path planning can be used without modifications. Sampling-based planners expand a tree of collision-free and dynamically feasible motions [9–11]. Even though solutions have probabilistic completeness guarantees, they are suboptimal and require post-processing to smooth the trajectory. Optimization-based planners rely on an initial guess and optimize locally, for example by using sequential convex programming (SCP) [12]. Hybrid methods can combine search and sampling [13], search and optimization [14] or sampling and optimization [15]. In this work, we extend db-A* [16] - a hybrid method connecting ideas from search, sampling, and optimization.

For multi-robot kinodynamic motion planning, single-robot planners applied to the joint space can be used, but do not scale well beyond a few robots. Extensions of CBS can tackle the Multi-Robot Motion Planning (MRMP) problems. For example, discretized workspace grid cells can be connected via predefined motion primitives [17]. However, computing primitives remains challenging for higher-order dynamics. Another option is to use model predictive control [18] or a sampling-based method [19] as the low-level planner. The latter, called K-CBS, uses any sampling-based planner and merges the search space of two robots if the number of conflicts between the robots exceeds a threshold. In the worst case, all robots may be merged but probabilistic completeness is achieved. In Safe Multi-Agent Motion Planning with Nonlinear Dynamics and Bounded Disturbances (S2M2) [20], Mixed-Integer Linear Programs (MILP) and control theory are combined by computing piecewise linear paths that a controller should be able to track within a safe region. These safe regions can be large, which makes the method incomplete.

Db-CBS, also leverages CBS, but we use a different single-robot planner and trajectory optimization, resulting in an efficient algorithm with completeness guarantees.

III. PROBLEM DEFINITION

We consider N robots and denote the state of the i^{th} robot with $\mathbf{x}^{(i)} \in \mathcal{X}^{(i)} \subset \mathbb{R}^{d_x^{(i)}}$, which is actuated by controlling actions $\mathbf{u}^{(i)} \in \mathcal{U}^{(i)} \subset \mathbb{R}^{d_u^{(i)}}$. The workspace the robots operate in is given as $\mathcal{W} \subseteq \mathbb{R}^{d_w}$ ($d_w \in \{2, 3\}$). The collision-free space is $\mathcal{W}_{\text{free}} \subseteq \mathcal{W}$.

We assume that each robot $i \in \{1, \dots, N\}$ has dynamics:

$$\dot{\mathbf{x}}^{(i)} = \mathbf{f}^{(i)}(\mathbf{x}^{(i)}, \mathbf{u}^{(i)}). \quad (1)$$

The Jacobian of $\mathbf{f}^{(i)}$ with respect to $\mathbf{x}^{(i)}$ and $\mathbf{u}^{(i)}$ are assumed to be available in order to use gradient-based optimization.

With zero-order hold discretization, (1) can be framed as

$$\mathbf{x}_{k+1}^{(i)} \approx \text{step}(\mathbf{x}_k^{(i)}, \mathbf{u}_k^{(i)}) \equiv \mathbf{x}_k^{(i)} + \mathbf{f}^{(i)}(\mathbf{x}_k^{(i)}, \mathbf{u}_k^{(i)})\Delta t, \quad (2)$$

where Δt is sufficiently small to ensure that the Euler integration holds.

We use $\mathbf{X}^{(i)} = \langle \mathbf{x}_0^{(i)}, \mathbf{x}_1^{(i)}, \dots, \mathbf{x}_{K^{(i)}}^{(i)} \rangle$ as a sequence of states of the i^{th} robot sampled at times $0, \Delta t, \dots, K^{(i)}\Delta t$ and $\mathbf{U}^{(i)} = \langle \mathbf{u}_0^{(i)}, \mathbf{u}_1^{(i)}, \dots, \mathbf{u}_{K^{(i)}-1}^{(i)} \rangle$ as a sequence of actions applied to the i^{th} robot for times $[0, \Delta t), [\Delta t, 2\Delta t), \dots, [(K^{(i)} - 1)\Delta t, K^{(i)}\Delta t)$. Our goal is to move the team of N robots from their start states $\mathbf{x}_s^{(i)} \in \mathcal{X}^{(i)}$ to their goal states $\mathbf{x}_f^{(i)} \in \mathcal{X}^{(i)}$ as fast as possible without collisions. This can be formulated as the following optimization problem:

$$\begin{aligned} & \min_{\{\mathbf{X}^{(i)}\}, \{\mathbf{U}^{(i)}\}, \{K^{(i)}\}} \sum_{i=1}^N K^{(i)} & (3) \\ \text{s.t.} & \begin{cases} \mathbf{x}_{k+1}^{(i)} = \text{step}(\mathbf{x}_k^{(i)}, \mathbf{u}_k^{(i)}) & \forall i \forall k, \\ \mathbf{u}_k^{(i)} \in \mathcal{U}^{(i)}, \mathbf{x}_k^{(i)} \in \mathcal{X}^{(i)} & \forall i \forall k, \\ \mathcal{B}^{(i)}(\mathbf{x}_k^{(i)}) \in \mathcal{W}_{\text{free}} & \forall i \forall k, \\ \mathcal{B}^{(i)}(\mathbf{x}_k^{(i)}) \cap \mathcal{B}^{(j)}(\mathbf{x}_k^{(j)}) = \emptyset & \forall i \neq j \forall k, \\ \mathbf{x}_0^{(i)} = \mathbf{x}_s^{(i)}, \mathbf{x}_{K^{(i)}}^{(i)} = \mathbf{x}_f^{(i)} & \forall i, \end{cases} \end{aligned}$$

where $\mathcal{B}^{(i)} : \mathcal{X}^{(i)} \rightarrow 2^{\mathcal{W}}$ is a function that maps the configuration of the i^{th} robot to a collision shape. The objective is to minimize the sum of the arrival time of all robots.

IV. APPROACH

A. Background

We first describe db-A* and CBS in more detail, since our work generalizes them into a multi-robot case operating in continuous space and time.

db-A* is a generalization of A* for kinodynamic motion planning of a single robot that searches a graph of implicitly defined *motion primitives*, i.e., precomputed motions respecting the robot's dynamics [16].

A single motion primitive can be defined as a tuple $\langle \mathbf{X}, \mathbf{U}, K \rangle$, consisting of state sequences $\mathbf{X} = \langle \mathbf{x}_0, \dots, \mathbf{x}_K \rangle$ and control sequences $\mathbf{U} = \langle \mathbf{u}_0, \dots, \mathbf{u}_{K-1} \rangle$, which obey the dynamics $\mathbf{x}_{k+1} = \text{step}(\mathbf{x}_k, \mathbf{u}_k)$. These motion primitives are used as graph edges to connect states, representing graph nodes, with user-configurable discontinuity δ .

As in A*, db-A* explores nodes based on $f(\mathbf{x}) = g(\mathbf{x}) + h(\mathbf{x})$, where $g(\mathbf{x})$ is the cost-to-come. The node with the lowest f -value is expanded by applying valid motion primitives. The output of db-A* is a δ -discontinuity-bounded solution as defined below.

Definition 1: Sequences $\mathbf{X} = \langle \mathbf{x}_0, \dots, \mathbf{x}_K \rangle$, $\mathbf{U} = \langle \mathbf{u}_0, \dots, \mathbf{u}_{K-1} \rangle$ are δ -discontinuity-bounded solutions iff the following conditions hold:

$$\begin{aligned} d(\mathbf{x}_{k+1}, \text{step}(\mathbf{x}_k, \mathbf{u}_k)) &\leq \delta \forall k, \\ \mathbf{u}_k &\in \mathcal{U}, \quad \mathbf{x}_k \in \mathcal{X}, \quad \mathcal{B}(\mathbf{x}_k) \in \mathcal{W}_{\text{free}} \forall k, \\ d(\mathbf{x}_0, \mathbf{x}_s) &\leq \delta, \quad d(\mathbf{x}_K, \mathbf{x}_f) \leq \delta, \end{aligned} \quad (4)$$

where d is a metric $d : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$, which measures the distance between two states.

To repair discontinuities in the trajectory, db-A* combines discrete search with gradient-based trajectory optimization in an iterative manner, decreasing δ with each iteration.

CBS is an optimal MAPF solver which works in two levels. The high-level search finds conflicts between single-robot paths, while at the low-level individual paths are updated to fulfill time-dependent constraints using a single-robot planner. Each high-level node contains a set of constraints and, for each robot, a path which obeys the given constraints. The cost of the high-level node is equal to the sum of all single-robot path costs. When a high-level node P is expanded, it is considered as the goal node, if its individual robot paths have no conflicts. If the node P has no conflicts, then its set of paths is returned as the solution. Otherwise, two P_1 and P_2 child nodes are added. Both child nodes inherit constraints of P and have an additional constraint to resolve the last conflict. For example, if the collision happened at vertex s at time t between robots i and j , then the node P_1 has an additional constraint $\langle i, s, t \rangle$, which prohibits the robot i to occupy vertex s at time t . Likewise, the constraint $\langle j, s, t \rangle$ is added to P_2 . Afterwards, given the set of constraints for each node, the low-level planner is executed for the affected robot, only.

B. db-CBS

Our planning approach, db-CBS, adapts CBS to the kinodynamic case. The structure of db-CBS consists of the following steps: i) a single-robot motion with discontinuous jumps is computed for each robot using db-A*, ii) collisions between individual motions are resolved one-by-one, iii) discontinuous motions are repaired into smooth and feasible trajectories with optimization. These steps are repeated until a solution is found or the desired solution quality is achieved. Although the two-level framework of CBS remains unchanged in our approach, the definition of conflicts and constraints need to be changed for the continuous domain. We define a conflict as $C = \langle i, j, \mathbf{x}_k^{(i)}, \mathbf{x}_k^{(j)}, k \rangle$ for a collision between robot i with state $\mathbf{x}_k^{(i)}$ and robot j with state $\mathbf{x}_k^{(j)}$ identified at time k . The resulting constraint for robot i is $\langle i, \mathbf{x}_k^{(i)}, k \rangle$, which prevents it to be within a distance of δ to state $\mathbf{x}_k^{(i)}$ at time k . Similarly, the constraint for robot j is $\langle j, \mathbf{x}_k^{(j)}, k \rangle$. The notion of a discontinuity δ defines the constraint as an actual volume (around a point), which is crucial for the efficiency and completeness guarantees.

The high-level search of db-CBS is shown in Alg. 1. Its major changes compared to CBS are highlighted. Db-CBS finds a solution in an iterative manner and decreases the

value of the discontinuity allowed in the final path with each iteration (Line 5). `AddPrimitives` adds more motion primitives from a pre-computed set of motions and from the optimization output. In each iteration, the root node of the constraint tree is initialized by running the low-level planner for each individual robot separately with the given set of motion primitives, current δ and no constraints (Line 8). Each single-robot motion consists of a concatenation of motion primitives that avoid robot-obstacle collisions as well as the specified constraints. If single-robot motions are found successfully, they are validated for inter-robot collision (Line 13) by checking sequentially at each timestep if a collision between any two robots occurred. The checking terminates when the earliest conflict is detected. This conflict is resolved by creating constraints and computing trajectories with the low-level planner for each constrained robot (Line 20-Line 24). Db-CBS uses an extension of db-A* that can directly consider dynamic constraints efficiently for the low-level search. If the motions have no conflicts, they are used as an initial guess for the trajectory optimization (Line 15). The trajectory optimization is performed in the joint configuration space with:

$$\begin{aligned} \min_{\{\mathbf{x}^{(i)}, \{\mathbf{U}^{(i)}\}, \Delta t} \sum_{i=1}^N \sum_{k=0}^{K-1} J(\mathbf{x}_k^{(i)}, \mathbf{u}_k^{(i)}) \quad (5) \\ \text{s.t.} \quad \begin{cases} \begin{bmatrix} \mathbf{x}_{k+1}^{(1)} \\ \mathbf{x}_{k+1}^{(2)} \\ \vdots \\ \mathbf{x}_{k+1}^{(N)} \end{bmatrix} = \begin{bmatrix} \mathbf{x}_k^{(1)} \\ \mathbf{x}_k^{(2)} \\ \vdots \\ \mathbf{x}_k^{(N)} \end{bmatrix} + \begin{bmatrix} \mathbf{f}^{(1)}(\mathbf{x}_k^{(1)}, \mathbf{u}_k^{(1)}) \\ \mathbf{f}^{(2)}(\mathbf{x}_k^{(2)}, \mathbf{u}_k^{(2)}) \\ \vdots \\ \mathbf{f}^{(N)}(\mathbf{x}_k^{(N)}, \mathbf{u}_k^{(N)}) \end{bmatrix} \Delta t \quad \forall k, \\ \begin{bmatrix} \mathbf{u}_k^{(1)}, \mathbf{u}_k^{(2)}, \dots, \mathbf{u}_k^{(N)} \end{bmatrix}^\top \in \bar{\mathcal{U}} \quad \forall k, \\ \begin{bmatrix} \mathbf{x}_k^{(1)}, \mathbf{x}_k^{(2)}, \dots, \mathbf{x}_k^{(N)} \end{bmatrix}^\top \in \bar{\mathcal{X}} \quad \forall k, \\ \mathbf{g} \left(\begin{bmatrix} \mathbf{x}_k^{(1)}, \mathbf{x}_k^{(2)}, \dots, \mathbf{x}_k^{(N)} \end{bmatrix}^\top \right) \geq \mathbf{0} \quad \forall k, \\ \mathbf{x}_0^{(i)} = \mathbf{x}_s, \quad \mathbf{x}_{K^{(i)}}^{(i)} = \mathbf{x}_f \quad \forall i. \end{cases} \end{aligned}$$

Here, the cost is $J(\mathbf{x}_k^{(i)}, \mathbf{u}_k^{(i)}) = \Delta t + \beta \|\mathbf{u}_k^{(i)}\|^2$ with a regularization β , and $\bar{\mathcal{U}} = \times_i \mathcal{U}^{(i)}$, $\bar{\mathcal{X}} = \times_i \mathcal{X}^{(i)}$ are the joint action space and joint state space, respectively. The optimization variables are state and control actions of all robots, and the length of the time interval Δt . The signed-distance function $\mathbf{g}(\cdot)$ performs robot-robot and robot-obstacle collision avoidance for each state. Since our ultimate objective is to minimize the sum of arrival times, we add the goal constraints at different timesteps $K^{(i)}$ for each robot based on the discontinuity-bounded initial guess. The trajectory optimization problem (5) can be solved with nonlinear, constrained optimization. In our work we use differential dynamic programming (DDP), adding the goal and collision constraints with a squared penalty method.

C. db-A* For Dynamic Obstacles

In CBS and db-CBS, constraints arise directly from conflicts and represent states to be avoided during the low-level search, which requires planning with dynamic obstacles. In

Algorithm 1: db-CBS: Discontinuity-Bounded Conflict-Based Search (High-Level Search)

```

/* Main changes compared to CBS are highlighted. */
Input:  $\{\mathbf{x}_s^{(i)}\}, \{\mathbf{x}_f^{(i)}\}, \mathcal{W}_{\text{free}}, N$ 
Result:  $\{\mathbf{X}^{(i)}\}, \{\mathbf{U}^{(i)}\}$ 
1  $\mathcal{M}^{(r)} = \emptyset \ \forall r \in \{1, \dots, N\}$   $\triangleright$  Initial set of motion primitives
2 for  $n = 1, 2, \dots$  do
3   for  $r \in N$  do
4      $\mathcal{M}^{(r)} \leftarrow \mathcal{M}^{(r)} \cup \text{AddPrimitives}(r)$   $\triangleright$  Add motion
       primitives for each robot dynamics
5      $\delta_n \leftarrow \text{ComputeDelta}()$   $\triangleright$  Update discontinuity bound
6      $\mathcal{S} \leftarrow \text{Node}(\text{solution} : \emptyset, \text{constraints} : \emptyset)$   $\triangleright$  Root node
7     foreach robot  $i \in \mathcal{R}$  do
8        $\mathcal{S}.\text{solution}[i] \leftarrow \text{db-A}^*(\mathbf{x}_s^{(i)}, \mathbf{x}_f^{(i)}, \mathcal{W}_{\text{free}}, \mathcal{M}^{(i)}, \delta, \text{None})$ 
        $\triangleright$  Single-robot planner with no constraints
9      $\mathcal{S}.\text{cost} \leftarrow \text{GetSolutionCost}(\mathcal{S}.\text{solution})$   $\triangleright$  Update cost
10     $\mathcal{O} \leftarrow \{\mathcal{S}\}$   $\triangleright$  Initialize open priority queue
11    while  $|\mathcal{O}| > 0$  do
12       $P \leftarrow \text{PriorityQueuePop}(\mathcal{O})$   $\triangleright$  Lowest solution cost
13       $C \leftarrow \text{GetEarliestConflict}(P.\text{solution})$   $\triangleright$  Check for
       collisions between individual motions
14      if  $C = \emptyset$  then
15         $\{\mathbf{X}^{(i)}\}, \{\mathbf{U}^{(i)}\} \leftarrow \text{Optimization}(P.\text{solution})$ 
16        if  $\{\mathbf{X}^{(i)}\}, \{\mathbf{U}^{(i)}\}$  successfully computed then
17          Report  $(\{\mathbf{X}^{(i)}\}, \{\mathbf{U}^{(i)}\})$   $\triangleright$  New solution found
18        else
19           $\langle i, j, \mathbf{x}^{(i)}, \mathbf{x}^{(j)}, k \rangle \leftarrow C$   $\triangleright$  Extract conflict
20          foreach  $c \in \{i, j\}$  do
21             $P' \leftarrow \text{Node}(\text{solution} : P.\text{solution}, \text{constraints} :$ 
               $P.\text{constraints} \cup \{c, \mathbf{x}^{(c)}, k\})$ 
22             $P'.\text{solution}[c] \leftarrow$ 
               $\text{db-A}^*(\mathbf{x}_s^{(c)}, \mathbf{x}_f^{(c)}, \mathcal{W}_{\text{free}}, \mathcal{M}^{(c)}, \delta, P'.\text{constraints})$ 
23             $P'.\text{cost} = \text{GetSolutionCost}(P'.\text{solution})$ 
24            PriorityQueueInsert  $(\mathcal{O}, P')$ 

```

CBS, the most common approach is to use A* in a space-time search space or SIPP [21]. Inspired by SIPP, we extend db-A* in order to solve single-robot path planning which is consistent with constraints. Our extended db-A* is given in Alg. 2. The notation $\mathbf{x} \oplus m$ indicates that the motion m is applied to state \mathbf{x} .

Recall that a constraint $\langle i, \mathbf{x}_c, k \rangle$ enforces $d(\mathbf{x}_c, \mathbf{x}_k^{(i)}) > \delta$. We handle all constraints during the expansion of neighboring states, where we only include motions that are at least δ away from any constrained state (Line 13 - Line 16). In addition, avoiding dynamic obstacles might require to reach a state with a slower motion. Thus, we keep a list of safe arrival time, parent node, and motion to reach this state from the parent (Line 3). When a potentially better path is found, we keep previous solution motions rather than removing them (Line 26). Conceptually, this is similar to SIPP, except that we do not store safe arrival intervals but potential arrival time points, since not all robots are able to wait. Empirically, storing the arrival times compared to always creating new nodes is significantly faster in our experiments.

D. Properties and Extensions

CBS is a complete and optimal algorithm for MAPF, i.e., if a solution exists it will find it and the first reported solution has the lowest possible cost (sum of costs of all agents).

Algorithm 2: db-A* with Dynamic Obstacles

```

/* Main changes compared to db-A* are highlighted. */
Input:  $\mathbf{x}_s, \mathbf{x}_f, \mathcal{W}_{\text{free}}, \mathcal{M}, \delta, \mathcal{C}$ 
Result:  $\mathbf{X}, \mathbf{U}, K$  or Infeasible
1  $\mathcal{T}_m \leftarrow \text{NearestNeighborInit}(\mathcal{M})$   $\triangleright$  Use start states of
   motions (excl. position)
2  $\mathcal{T}_n \leftarrow \text{NearestNeighborInit}(\{\mathbf{x}_s\})$   $\triangleright$  capture explored
   vertices (incl. position)
3  $\mathcal{O} \leftarrow \{\text{Node}(\mathbf{x} : \mathbf{x}_s, g : 0, h : h(\mathbf{x}_s), A : \{(g : 0, p : \text{None}, a :$ 
    $\text{None})\})\}$   $\triangleright$  Initialize open priority queue
4 while  $|\mathcal{O}| > 0$  do
5    $n \leftarrow \text{PriorityQueuePop}(\mathcal{O})$   $\triangleright$  Lowest  $f$ -value
6   if  $d(n.\mathbf{x}, \mathbf{x}_f) \leq \delta$  then
7     return  $\mathbf{X}, \mathbf{U}, K$   $\triangleright$  Traceback solution
    $\triangleright$  Find applicable motion primitives with discontinuity up to  $\alpha\delta$ 
8    $\mathcal{M}' \leftarrow \text{NearestNeighborQuery}(\mathcal{T}_m, n.\mathbf{x}, \alpha\delta)$ 
9   foreach  $m \in \mathcal{M}'$  do
10    if  $n.\mathbf{x} \oplus m \notin \mathcal{W}_{\text{free}}$  then
11      continue  $\triangleright$  entire motion is not collision-free
12     $g_t \leftarrow n.g + \text{cost}(m)$   $\triangleright$  tentative  $g$ -score for this action
13    foreach  $c \in \mathcal{C}$  do
14       $\mathbf{x}' = m[\lfloor \frac{c.k - n.g}{\Delta t} \rfloor]$   $\triangleright$  Motion primitive state for checking
15      if  $d(\mathbf{x}', c.\mathbf{x}) \leq \delta$  then
16        continue loop Line 9  $\triangleright$   $m$  does not obey all constraints
17     $\triangleright$  find already explored nodes within  $(1 - \alpha)\delta$ 
18     $\mathcal{N}' \leftarrow \text{NearestNeighborQuery}(\mathcal{T}_n, n.\mathbf{x} \oplus m, (1 - \alpha)\delta)$ 
19    if  $\mathcal{N}' = \emptyset$  then
20      PriorityQueueInsert  $(\mathcal{O}, \text{Node}(\mathbf{x} : n.\mathbf{x} \oplus m, g :$ 
         $g_t, h : h(n.\mathbf{x} \oplus m), A : \{(g : g_t, p : n, a : m)\})$ 
21      NearestNeighborAdd  $(\mathcal{T}_n, n.\mathbf{x} \oplus m)$ 
22    else
23      foreach  $n' \in \mathcal{N}'$  do
24        if  $g_t < n'.g$   $\triangleright$  Motion is better than a known motion then
25           $n'.g = g_t$   $\triangleright$  Update cost
26           $n'.A = n'.A \cup \{(g : g_t, p : n, a : m)\}$   $\triangleright$  Add parent
27          PriorityQueueUpdate  $(\mathcal{O}, n')$ 
28 return Infeasible

```

The proof [5] uses two arguments: i) for completeness, it is shown that the search will visit all states that can contain the solution, i.e., no potential solution paths are pruned, and ii) for optimality, it is shown that CBS visits the solutions in increasing order of costs (f -value) and can therefore not have missed a lower-cost solution once it terminates.

The continuous time and continuous space renders the full enumeration proof argument infeasible for MRMP problems. Instead, we consider probabilistic completeness (PC; the probability of finding a solution if one exists is 1 in the limit over search effort) and asymptotic optimality (AO; the cost difference between the reported solution and an optimal solution approaches zero in the limit over search effort). We assume that there exists a finite $\delta > 0$ such that the trajectory with discontinuity is optimized successfully. We now show that db-CBS is both probabilistically complete and asymptotically optimal and therefore an anytime planner, similar to sampling-based methods such as SST*.

Theorem 1: The db-CBS motion planner in Alg. 1 is asymptotically optimal, i.e.

$$\lim_{n \rightarrow \infty} P(\{c_n - c^* > \epsilon\}) = 0, \ \forall \epsilon > 0, \quad (6)$$

where c_n is the cost in iteration n and c^* is the optimal cost.

Proof Sketch: Each iteration in Line 2 operates on a discrete search problem over the graph implicitly defined by the finite set of motion primitives. We find an optimal solution, if one exists, within this discretization by [5, Theorem 1]. Subsequent iterations add more motion primitives and reduce δ as in [16]; therefore, a larger discrete search graph is produced, for which the optimality proof of CBS still holds. In the limit ($|\mathcal{M}| \rightarrow \infty, \delta \rightarrow 0$), db-CBS will report the optimal solution. This argument is the same as the detailed version in [22, Theorem 1 and 2], for combining PRM and CBS. A formal version of this proof, including bounding the actual probability, is given in [16, Theorem 2]. ■

Theorem 2: The db-CBS motion planner in Alg. 1 is probabilistically complete.

Proof: This directly follows from Theorem 1, since asymptotic optimality implies probabilistic completeness. ■

This result compares to the baseline algorithms as follows. SST* is probabilistically complete and asymptotically near-optimal, where the near-optimality stems from a fixed hyperparameter similar to the time-varying δ in db-CBS. K-CBS is probabilistically complete and could achieve asymptotic optimality by using an incremental cost bound similar to AO-x [23]. S2M2 is neither probabilistically complete nor asymptotically optimal, because it uses pessimistic approximations and priority-based search. We show concrete examples in which S2M2 fails to find a solution in Sec. V.

We note that db-CBS can naturally plan for heterogeneous teams of robots, where individual team members may have unique dynamics or collision shapes. This does not require any algorithmic changes, as the low-level planner can use different motion primitives and constraints are defined for each robot themselves, and conflicts are detected using arbitrary collision shapes. On the optimization side, we rely on nonlinear optimization and stacked states which naturally extends to varied state and action dimensions.

V. RESULTS

We compare our method with other multi-robot kinodynamic motion planners on the same problem instances. We consider robots with three to five-dimensional states: unicycle, 2nd order unicycle, double integrator, and car with trailer, see [16] for dynamics and bounds. For the double integrator, we use $v \in [-0.5, 0.5]$ m/s, $a \in [-2.0, 2.0]$ m/s².

For testing scenarios we focus on cases where the workspace has small dimensions, thus it is challenging for robots to find collision-free motions in a time-optimal way.

In each environment, we compare our method with K-CBS [19], S2M2 [20], and joint space SST* [11]. We analyze success rate (p), computational time until the first solution is found (t), and cost of the first solution (J). An instance is not solved successfully if an incorrect solution is returned, or no solution is found after 5 min. To compare relative improvements over baselines, we use a notion of regret, i.e., $J^{s2m2} = (J^{s2m2} - J^{dbcs}) / J^{s2m2}$.

We use OMPL [24] and FCL [25] for K-CBS, SST* and db-CBS. For fair comparison, we set the low-level planner of K-CBS to SST*, and modify the collision checking in order




Fig. 2. Top row: S2M2, K-CBS, db-CBS with *swap* (left to right). Middle row: SST*, S2M2, db-CBS with *alcove*. Last row shows instances where only db-CBS is able to return a solution: 4 cars with trailer with *swap*, 4 unicycles with *swap*, and 8 unicycles with *random* start and goal states.

to support robots of any shape. For S2M2, we use the Euler integration to propagate the robot state and bound control actions. For K-CBS and S2M2 we use the respective publicly available implementations from the authors. For K-CBS, we disable merging. When comparing with S2M2, we use a spherical first order unicycle model and quadrupled angular velocity bounds $\omega \in [-2, 2]$ rad/s, since the provided code was unable to solve most of our very dense planning problems otherwise. Moreover, there is a trade-off in setting a parameter for bloated obstacles; a small gain leads to collisions with the environment, while a large gain results in infeasible solutions. We carefully tuned this and other hyperparameters to maximize the success rate of all algorithms (see our code for exact values). For **db-CBS** Algs. 1 and 2 are implemented in C++, and the benchmarking script is written in Python. For optimization, we extend Dynoplan [26], which uses the DDP solver Crocoddyl [27]. We use a workstation (AMD Ryzen Threadripper PRO 5975WX @ 3.6 GHz, 64 GB RAM, Ubuntu 22.04). Code, benchmark, and video are available on github, see first page.

A. Canonical Examples

We compare the algorithms on canonical examples often seen in the literature, see rows 1–3 in Table I and Fig. 2.

1) *Swap*: the environment has no obstacles and two robots have to swap their positions, i.e. each robot’s start position is the goal position of the other robot. In this environment all baselines succeed with SST* and S2M2 being the fastest to find a solution. However, their solution costs are higher compared to db-CBS, even though less than of K-CBS.

2) *Alcove*: the environment requires one robot to move into an alcove temporarily in order to let the second robot pass. All baselines and db-CBS are able to solve this problem, except K-CBS.

TABLE I
CANONICAL EXAMPLES, SCALABILITY WITH VARYING ROBOT NUMBERS, AND HETEROGENEOUS SYSTEMS.
MEDIAN VALUES OVER 10 TRIALS PER ROW. BOLD ENTRIES ARE THE BEST FOR THE ROW, —NO SOLUTION FOUND, * NOT TESTED.

#	Instance	SST*				S2M2				K-CBS				db-CBS			
		p	$t[s]$	$J[s]$	$r[\%]$	p	$t[s]$	$J[s]$	$r[\%]$	p	$t[s]$	$J[s]$	$r[\%]$	p	$t[s]$	$J[s]$	$r[\%]$
1	swap	0.3	0.1	29.8	49	1.0	0.1	16.0	5	1.0	21.5	30.4	50	1.0	3.6	15.2	0
2	alcove	0.9	4.2	45.5	50	1.0	0.5	29.2	22	0.0	—	—	—	1.0	44.0	22.8	0
3	at goal	0.5	0.5	33.8	47	0.0	—	—	—	0.0	—	—	—	1.0	58.9	17.6	0
4	rand (N=2)	0.1	0.3	18.7	58	0.7	0.1	27.5	8	1.0	2.0	65.0	63	1.0	6.0	25.6	0
5	rand (N=4)	0.0	—	—	—	0.3	0.4	52.0	13	0.6	61.9	144.4	59	1.0	28.7	55.1	0
6	rand (N=8)	0.0	—	—	—	0.0	—	—	—	0.0	—	—	—	0.1	256.6	146.9	0
7	rand hetero (N=2)	0.1	157.2	33.6	29	*	*	*	*	0.5	17.2	50.9	62	1.0	2.2	14.2	0
8	rand hetero (N=4)	0.0	—	—	—	*	*	*	*	0.2	145.6	89.9	65	0.8	7.5	37.1	0
9	rand hetero (N=8)	0.0	—	—	—	*	*	*	*	0.0	—	—	—	0.3	271.8	86.4	0

TABLE II
RUNTIME OF SST* (*), K-CBS (†), AND DB-CBS(‡) IN SECONDS ON
THE SWAP PROBLEM (MEDIAN OVER 10 TRIALS).

N	unicycle 1 st			double int.			car with trailer			unicycle 2 nd		
	*	†	‡	*	†	‡	*	†	‡	*	†	‡
1	0.2	1.0	1.2	0.0	1.0	0.2	0.3	1.0	2.5	1.5	14.6	7.0
2	2.4	2.0	2.1	—	2.0	0.4	—	10.1	5.3	—	39.3	11.3
3	—	34.6	4.1	—	—	0.5	—	—	11.9	—	—	20.9
4	—	214.1	8.0	—	—	1.1	—	—	86.0	—	—	23.5

3) *At Goal*: similar to *Alcove*, but one robot is already at its goal state and needs to move away temporarily to let the second robot pass. Just SST* and db-CBS are able to find a solution, but with high difference between their solution cost. The failure of S2M2 is due to collision violation.

Both *alcove* and *at goal* should be solvable by K-CBS for finite merge bounds (a feature that was not readily available in the provided code). However, in this case the algorithm essentially becomes joint space SST* and would produce solutions with a high cost as seen in Table I.

B. Scalability

1) *Robot Number*: we generate obstacles, start and goal states for up to eight robots, see rows 4–6 in Table I. S2M2 and K-CBS are able to solve this problem with up to four robots, however inconsistently. In addition, both of them are failing to return a solution for eight robot cases, while db-CBS scales up successfully.

2) *Heterogeneous Systems*: we generate obstacles, start and goal states for a team of up to eight heterogeneous robots, see rows 7–9 in Table I. S2M2 is not tested, since it does not support all robot dynamics considered here. While SST* solves instances with two robots with significantly lower cost compared to K-CBS, it struggles to scale to more robots. K-CBS is able to handle up to four heterogeneous robots, but the solution cost is very high. The success rate and solution cost are the best for db-CBS across all settings.

Note that in these random instances the shown median of J is not directly comparable unless the success rate $p = 1$. In the table the positive regret r for all baselines shows that db-CBS consistently computes lower-cost solutions.

3) *State Dimension*: in Table II we report the solution quality of *swap* problem instances with different dynamics

and varying robot numbers (1–4), while keeping the start and goal states unchanged. None of the methods, except db-CBS, succeed to solve the problem with more than two robots across different dynamics. However, K-CBS is able to solve up to four unicycles (1st order), while the success rate of SST* is very low if the robot number is greater than one.

Note that all problem instances consider 5×5 environments, which is significantly smaller compared to scenarios S2M2 and K-CBS were previously validated on. We verified that db-CBS shows anytime behavior empirically; similar to [16] the first solution already has a low cost and subsequent improvements are small and not shown for brevity.

C. Real-Robot Demo

The real-world experiments are conducted inside a $3.5 \times 3.5 \times 2.75$ m³ room. We use Bitcraze Crazyflie 2.1 drones and control them using CrazySwarm2 [28]. We consider two scenarios with four robots modeled with 2D double integrator dynamics (and thus do not include S2M2 in the discussion). We first test the *swap* example. K-CBS and SST* can not solve this problem similar to results observed in Sec. V-B.3, while db-CBS takes 12.1 s to return a trajectory with 11.4 s cost. The second scenario has a wall with a small window and two robots on each side of the wall. The robots have to pass through the small window, see Fig. 1. Only db-CBS finds a solution within 80 s with 21.1 s cost.

VI. CONCLUSION

In this paper, we present db-CBS, an efficient, probabilistically complete and asymptotically optimal motion planner for a heterogeneous team of robots that considers robot dynamics and control bounds. Db-CBS solves the multi-robot kinodynamic motion planning problem by finding collision-free trajectories with a bounded discontinuity and optimizes them in joint configuration space. Empirically, we show that db-CBS finds solutions with a significantly higher success rate and better solution cost compared to the existing state-of-the-art. Finally, we validate our planner on two real-world challenging problem instances on flying robots.

In the future, we will solve problems with more robots and over longer time horizons, for example by combining db-A* with a more advanced version of CBS, such as Enhanced CBS (ECBS), and by repairing discontinuities locally.

REFERENCES

- [1] K. S. Yakovlev and A. Andreychuk, "Any-angle pathfinding for multiple agents based on SIPP algorithm," in *International Conference on Automated Planning and Scheduling (ICAPS)*, 2017, p. 586.
- [2] J. Yu and S. M. LaValle, "Structure and intractability of optimal multi-robot path planning on graphs," in *AAAI Conference on Artificial Intelligence (AAAI)*, 2013, pp. 1443–1449.
- [3] R. Cui, B. Gao, and J. Guo, "Pareto-optimal coordination of multiple robots with safety guarantees," *Autonomous Robots*, vol. 32, no. 3, pp. 189–205, 2012.
- [4] M. Cáp, P. Novák, M. Selecký, J. Faigl, and J. Vokffnek, "Asynchronous decentralized prioritized planning for coordination in multi-robot system," in *International Conference on Intelligent Robots and Systems (IROS)*, 2013, pp. 3822–3829.
- [5] G. Sharon, R. Stern, A. Felner, and N. R. Sturtevant, "Conflict-based search for optimal multi-agent path finding," in *AAAI Conference on Artificial Intelligence (AAAI)*, 2012.
- [6] W. Hönig, T. K. S. Kumar, L. Cohen, H. Ma, H. Xu, N. Ayanian, and S. Koenig, "Multi-agent path finding with kinematic constraints," in *International Conference on Automated Planning and Scheduling (ICAPS)*, 2016, pp. 477–485.
- [7] M. Dharmadhikari, T. Dang, L. Solanka, J. Loje, H. Nguyen, N. Khedekar, and K. Alexis, "Motion primitives-based path planning for fast and agile exploration using aerial robots," in *International Conference on Robotics and Automation (ICRA)*, 2020, pp. 179–185.
- [8] B. J. Cohen, S. Chitta, and M. Likhachev, "Search-based planning for manipulation with motion primitives," in *International Conference on Robotics and Automation (ICRA)*, 2010, pp. 2902–2908.
- [9] L. Kavraki, P. Svestka, J.-C. Latombe, and M. Overmars, "Probabilistic roadmaps for path planning in high-dimensional configuration spaces," *IEEE Transactions on Robotics (T-RO)*, vol. 12, no. 4, pp. 566–580, 1996.
- [10] S. M. LaValle and J. J. K. Jr., "Randomized kinodynamic planning," *The International Journal of Robotics Research (IJRR)*, vol. 20, no. 5, pp. 378–400, 2001.
- [11] Y. Li, Z. Littlefield, and K. E. Bekris, "Asymptotically optimal sampling-based kinodynamic planning," *The International Journal of Robotics Research (IJRR)*, vol. 35, no. 5, pp. 528–564, 2016.
- [12] J. Schulman, Y. Duan, J. Ho, A. X. Lee, I. Awwal, H. Bradlow, J. Pan, S. Patil, K. Goldberg, and P. Abbeel, "Motion planning with sequential convex optimization and convex collision checking," *The International Journal of Robotics Research (IJRR)*, vol. 33, no. 9, pp. 1251–1270, 2014.
- [13] B. Sakcak, L. Bascetta, G. Ferretti, and M. Prandini, "Sampling-based optimal kinodynamic planning with motion primitives," *Autonomous Robots*, vol. 43, no. 7, pp. 1715–1732, 2019.
- [14] R. Natarajan, H. Choset, and M. Likhachev, "Interleaving graph search and trajectory optimization for aggressive quadrotor flight," *IEEE Robotics and Automation Letters (RA-L)*, vol. 6, no. 3, pp. 5357–5364, 2021.
- [15] S. Choudhury, J. D. Gammell, T. D. Barfoot, S. S. Srinivasa, and S. A. Scherer, "Regionally accelerated batch informed trees (RA-BIT*): A framework to integrate local information into optimal path planning," in *International Conference on Robotics and Automation (ICRA)*, 2016, pp. 4207–4214.
- [16] W. Hönig, J. O. de Haro, and M. Toussaint, "db-A*: Discontinuity-bounded search for kinodynamic mobile robot motion planning," in *International Conference on Intelligent Robots and Systems (IROS)*, 2022, pp. 13 540–13 547.
- [17] L. Cohen, T. Uras, T. K. S. Kumar, and S. Koenig, "Optimal and bounded-suboptimal multi-agent motion planning," in *International Symposium on Combinatorial Search (SoCS)*, 2021.
- [18] A. Tajbakhsh, L. T. Biegler, and A. M. Johnson, "Conflict-based model predictive control for scalable multi-robot motion planning," *CoRR*, vol. abs/2303.01619, 2023. arXiv: 2303.01619.
- [19] J. Kottinger, S. Almagor, and M. Lahijanian, "Conflict-based search for multi-robot motion planning with kinodynamic constraints," in *International Conference on Intelligent Robots and Systems (IROS)*, 2022, pp. 13 494–13 499.
- [20] J. Chen, J. Li, C. Fan, and B. C. Williams, "Scalable and safe multi-agent motion planning with nonlinear dynamics and bounded disturbances," in *AAAI Conference on Artificial Intelligence (AAAI)*, 2021, pp. 11 237–11 245.
- [21] S. Hu, D. D. Harabor, G. Gange, P. J. Stuckey, and N. R. Sturtevant, "Multi-agent path finding with temporal jump point search," in *International Conference on Automated Planning and Scheduling (ICAPS)*, 2022, pp. 169–173.
- [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 (RA-L)*, vol. 6, no. 3, pp. 4608–4615, 2021.
- [23] K. Hauser and Y. Zhou, "Asymptotically optimal planning by feasible kinodynamic planning in a state-cost space," *IEEE Transactions on Robotics (T-RO)*, vol. 32, no. 6, pp. 1431–1443, 2016.
- [24] I. A. Şucan, M. Moll, and L. E. Kavraki, "The open motion planning library," *IEEE Robotics & Automation Magazine*, vol. 19, no. 4, pp. 72–82, 2012.
- [25] J. Pan, S. Chitta, and D. Manocha, "FCL: A general purpose library for collision and proximity queries," in *International Conference on Robotics and Automation (ICRA)*, 2012, pp. 3859–3866.
- [26] J. O.-H. et. al., *Dynoplan*, <https://github.com/quimortiz/dynoplan>.
- [27] C. Mastalli, R. Budhiraja, W. Merkt, G. Saurel, B. Hammoud, M. Naveau, J. Carpentier, L. Righetti, S. Vijayakumar, and N. Mansard, "Crocodyl: An efficient and versatile framework for multi-contact optimal control," in *International Conference on Robotics and Automation (ICRA)*, 2020, pp. 2536–2542.
- [28] J. A. Preiss, W. Hönig, G. S. Sukhatme, and N. Ayanian, "Crazyswarm: A large nano-quadcopter swarm," in *International Conference on Robotics and Automation (ICRA)*, 2017, pp. 3299–3304.