

# Optimizing Space Utilization for More Effective Multi-Robot Path Planning

Shuai D. Han<sup>1</sup> Jingjin Yu<sup>1</sup>

**Abstract**—We perform a systematic exploration of the principle of Space Utilization Optimization (SUO) as a heuristic for planning better individual paths in a decoupled multi-robot path planner, with applications to both one-shot and life-long multi-robot path planning problems. We show that the heuristic set, SU-I, preserves single path optimality and significantly reduces congestion that naturally happens when many paths are planned without coordination. Integration of SU-I into complete planners brings dramatic reductions in computation time due to the significantly reduced number of conflicts and leads to sizable solution optimality gains in diverse evaluation scenarios over medium and large maps, for both one-shot and life-long problem settings.

## I. INTRODUCTION

Recent years have witnessed a dramatic acceleration in the deployment of multi-robot systems for general logistic tasks [1], especially in the domain of shipping and warehousing [2]. Fast-paced expansion is predicted across the board, with the warehouse domain alone expecting a 14% year-over-year growth in the next five years [3]. This in turn demands the push for enhancing the scalability of multi-agent and multi-robot systems, which in the end boils down to achieving the maximum possible output attainable. Holding other variables constant, maximizing system throughput is most readily achieved by increasing robot density and plan optimality, which calls for faster and better computational methods for Multi-Robot Path Planning (MRPP) and Life-long Multi-Robot Path Planning (LMPP) problems.

Toward the development of more efficient and higher performance multi-robot systems catering to the rapidly growing need of automation, in this work, following the decoupled planning paradigm [4], we perform a systematic study of an intuitive principle for the design of better performing heuristics for MRPP. The decoupled setting generally involves two planning phases, where the first phase plans individual robot paths ignoring other robots and the second one resolves robot-robot conflicts within some spatio-temporal window. Traditionally, this phase is executed by running single robot path plannings ignoring other robots. The *Space Utilization Optimization* (SUO) principle tackles the first planning phase, seeking to make robots use the free space “evenly”.

Based on vertex, edge, and temporal usage information, SU-I, as our implementation of the SUO principle, builds a global heuristic that tracks how the free space is being used among all participating robots. We then exploit applying SU-I as both an *estimated cost-to-go*, for reducing congestion, and as part of the *cost-to-come*, for reducing conflicts along the

<sup>1</sup>S. D. Han is with Amazon Robotics. The work is done while he was a PhD candidate at Rutgers. J. Yu is with the Department of Computer Science, Rutgers University. Emails: {shuai.han, jingjin.yu}@rutgers.edu. This work is supported by NSF award IIS-1845888 and an Amazon Research Award.

entire robot path. Theoretically, we prove that SU-I does not compromise individual path optimality while simultaneously achieves its design goal of providing better space utilization. In practice, SU-I leads to significant improvement in initial path quality, resulting in over 40% reduction of path conflicts.

The introduction of SU-I brings notable improvement to a multitude of MRPP and LMPP benchmarks. SU-I, which may be applied as an orthogonal heuristic to many MRPP algorithms, leads to sizable gains in both computation time and solution optimality when combined with efficient methods for the second phase of a decouple planner for path scheduling, like DDM [5] and ECBS [6]. For example, using SU-I with the database-driven collision resolution from [5] leads to 15%+ reduction in computation time and improves solution optimality by roughly 25%. For LMPP, SU-I, with additional planning horizon management, could reduce the computation time of a state-of-the-art methods [7] by more than 65%, while keeping the same level of optimality.

**Related Work.** Multi-Robot Path Planning (MRPP), or equivalently, Multi-Agent Path Finding (MAPF) [8], has been actively studied for decades from many angles including computational complexity and effective algorithm design [4], [9], [10]. Until recently, studies on MRPP mainly focus on *one-shot* or *static* problems, where  $n$  robots are to reach  $n$  specific goals. Many algorithms for computing high-quality or optimal solutions have been proposed. Decoupled solutions [4] dominate the algorithmic attack, with methods using techniques including independence detection [11], sub-dimensional expansion [12], conflict-based search [13], [14], among others. Methods have also been proposed through the reduction to other problems including satisfiability (SAT) [15], Answer Set Programming (ASP) [16], and multi-commodity flow [17]. There also exists prioritized methods [18]–[21], splitting heuristics [22], and divide-and-conquer approaches [23], [24] which achieve good scalability but at the cost of either completeness or optimality. In [25], an any time algorithm is proposed to quickly find a feasible solution, which is subsequently improved. A learning-assisted approach [26] has been developed to automatically select the algorithm for solving MRPP challenges.

With the rise of multi-robot applications in the logistics domain [2], *dynamic* or *life-long* MRPP variants, or LMPP, have attracted attention in the past few years. Recent work has focused on dynamic warehouse setups, pursuing both better planning algorithms [27] and robust execution schedules [28]. Prioritized planning method with a flexible priority sequence has also been developed [29].

The general idea behind the SUO principle, i.e., better usage of the shared free space, has been explored under both single and multi-robot settings. For a single robot

exploring an obstacle-laden domain, a path ensemble can increase the chance of succeeding in finding a longer horizon plan [30], [31]. Path diversity is one relevant factor in MRPP resolution [32]. A reservation table is used for collision avoidance [33]. Coordination of moving directions has been studied [34]–[36]. Survivability is examined under a probabilistic framework for multi-robot systems [37]. A heuristic based on path conflicts expedites the solution process of an MRPP algorithm [6].

Despite the fact that the SUO principle is intuitive, to our knowledge, our exploitation of the principle, building on our initial work [5], makes novel contributions to the field. In contrast to [5], this study (i) thoroughly exploits the SUO principle with the introduction of the SU-I heuristics with a number of variations; (ii) proves SU-I’s collision avoidance properties: it finds the shortest paths while minimizes certain collision-based metrics; (iii) integrates SUO to MRPP and LMPP algorithms and empirically shows that using SU-I benefits both computation time and solution optimality.

## II. PRELIMINARIES

### A. Problem Statement

Multi-Robot Path Planning (MRPP) seeks collision-free paths that efficiently route robots. Consider an undirected graph  $\mathcal{G}(V, E)$  and  $n$  robots with start configuration  $S = \{s_1, \dots, s_n\} \subseteq V$  and goal configuration  $G = \{g_1, \dots, g_n\} \subseteq V$ . Each robot has start and goal vertices  $s_i, g_i$ . We define a *path* for robot  $i$  as a map  $P_i : \mathbb{N} \rightarrow V$  where  $\mathbb{N}$  is the set of non-negative integers. A feasible  $P_i$  must be a sequence of vertices that connects  $s_i$  and  $g_i$ : 1)  $P_i(0) = s_i$ ; 2)  $\exists T_i \in \mathbb{N}$ , s.t.  $\forall t \geq T_i, P_i(t) = g_i$ ; 3)  $\forall t > 0, P_i(t) = P_i(t-1)$  or  $(P_i(t), P_i(t-1)) \in E$ .

Here, we first define the single robot version of the problem.

**Problem 1. Single-Robot Path Planning (SPP).** *Given graph  $\mathcal{G}$ , start vertex  $s$ , goal vertex  $g$ ; find a feasible path  $P$ .*

Following the feasibility definition of  $P_i$ , we denote  $T_i$  as the *length* of  $P_i$ . We call  $P_i$  the *shortest path* for robot  $i$  if and only if it minimizes  $T_i$ . Given a set of paths  $\{P_1, \dots, P_n\}$  for all robots, we call them *collision-free* if and only if there are no simultaneous occupancy of the same vertex or edge. That is,  $\forall i, j : 1 \leq i < j \leq n, P_i, P_j$  must satisfy: 1)  $\forall t \geq 0, P_i(t) \neq P_j(t)$ ; 2)  $\forall t > 0, (P_i(t-1), P_i(t)) \neq (P_j(t), P_j(t-1))$ .

We say paths are *independent* if they are feasible single robot paths but not necessarily collision-free.

The traditional one-shot MRPP problem is defined as

**Problem 2. Multi-Robot Path Planning (MRPP).** *Given graph  $\mathcal{G}$ , start configuration  $S$ , goal configuration  $G$ ; find a collision-free path set  $\{P_1, \dots, P_n\}$ .*

An *optimal* solution for MRPP may minimize the *makespan*  $\max_{1 \leq i \leq n} |T_i|$  or *sum-of-cost*  $\sum_{1 \leq i \leq n} T_i$ .

Apart from MRPP, we also study the *life-long* variation LMPP where each robot has a *list of goal vertices*. We denote the goal configuration as  $\mathbf{G} = \{g_1, \dots, g_n\}$  where  $g_i = (g_i^1, g_i^2, \dots)$ . Here,  $g_i^k$  is the  $k$ -th goal in robot  $i$ ’s goal list  $\mathbf{g}_i$ . Note that for an actual LMPP instance,  $\mathbf{g}_i$  is often a list

that is constantly updated. In LMPP, the second condition for a feasible path  $P_i$  becomes: 2)  $\exists T_i^1 < T_i^2 < \dots \in \mathbb{N}$ , s.t.  $P_i(T_i^k) = g_i^k$ .

With all other conditions inherited from MRPP, we have

**Problem 3. Life-long Multi-Robot Path Planning (LMPP).** *Given graph  $\mathcal{G}$ , start configuration  $S$ , goal configuration (list)  $\mathbf{G}$ ; find a collision-free path set  $\{P_1, \dots, P_n\}$ .*

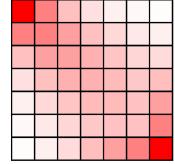
LMPP algorithms often optimizes *throughput*, i.e., the average number of goal reaches in a unit time step. Given a large  $T \in \mathbb{N}$ , the throughput can be expressed by  $(\sum_{1 \leq i \leq n} \arg \max_k (T_i^k | T_i^k \leq T)) / T$ .

For practicality and simplifying explanation, we assume  $\mathcal{G}$  is a 4-connected grid. All algorithms and heuristics proposed in this paper apply to arbitrary graphs with unit edge cost.

### B. Importance of SPP in Solving MRPP and LMPP

Solving SPP is a stepping stone toward MRPP and LMPP. Many MRPP solvers incorporate SPP planners as sub-routines, e.g., decoupled planners generally use a two-phase approach to first plan independent paths and then resolve collisions. Such methodologies are popular for two reasons. First, unlike multi-robot planning which is hard to optimize [38], SPP has been thoroughly studied and can be solved efficiently and optimally using A\* search with simple heuristics. Second, in a practical setting with relatively low robot density, usually only small sub-groups of robots have local interactions at any given time, which can be quickly resolved.

However, such a two-phase approach tends to overuse parts of the free space while leaving other parts underutilized. To demonstrate the effect, we plan shortest paths from the top left corner to the bottom right corner on a  $7 \times 7$  grid with randomized node exploration and randomized tie-breaking. The shades of cells in the figure visualizes the probability a cell is to be used. The paths concentrate along the diagonal connecting the two corners, leaving the top right and bottom left corners underutilized. This effect gets more pronounced when obstacles exist and the available path choices are limited.



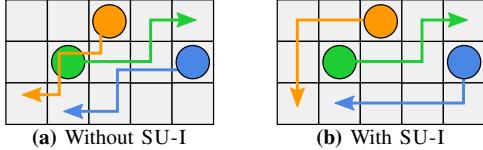
LMPP, more general and practical than MRPP, is often solved iteratively using existing MRPP algorithms. Thus, unbalanced graph utilization issue also negatively impacts LMPP solvers.

## III. SUO: PRINCIPLES AND IMPLEMENTATION

### A. Optimizing Space Utilization: Principles

From the discussion in Section II-B, robots’ individual paths should be spread “evenly” across the free space to reduce congestion. We call this idea the *Space Utilization Optimization (SUO) principle*. In this work, we develop a *first SUO implementation*, SU-I (SU-First), which serves as a global heuristic to help generate independent paths for MRPP/LMPP algorithms to better utilize graph resources. The example in Fig. 1 demonstrates that SU-I could potentially enhance both computational efficiency and optimality.

In SU-I, each path uses vertex, edge, and temporal information from other paths to avoid more congested areas



**Fig. 1:** Two sets of individual paths in a  $5 \times 3$  grid. The robots' start configuration is shown as colored disks. The planned individual paths are drawn as colored lines. (a) Randomly generated paths may result in congestion. (b) When using SU-I, which “spreads out” the robots' paths, the number of conflicts is reduced.

or to reduce the sum of conflicts along the path. SU-I improves upon its predecessor [5] which only considered vertex information. In Appendix [39], we use several examples to highlight the essence of vertex, edge, and temporal information when applying SUO.

### B. Path Planning with SU-I: High Level Procedure

We now describe the implementation of SU-I, including theoretical guarantees and evaluation results to show that SU-I with the proposed path planning procedure indeed optimizes space utilization. Since LMPP algorithms are often based on MRPP solvers, we only describe SU-I in the context of MRPP; SU-I is directly applicable to LMPP.

---

#### Algorithm 1: Generate independent paths with SU-I

---

```

Input: Graph  $\mathcal{G}$ ,  $n$  robots with starts  $S$  and goals  $G$ 
1  $\pi \leftarrow \text{GETORDERBYDISTANCETOGOAL}(\mathcal{G}, S, G)$ 
2  $S \leftarrow \pi(S), G \leftarrow \pi(G)$ 
3 for  $1 \leq i \leq n$  do  $P_i \leftarrow \text{None}$ 
4 for number of planning iterations  $r$  do
5   for  $1 \leq i \leq n$  do
6      $\mathcal{T} \leftarrow \text{BUILDSUO}(P_1, \dots, P_{i-1}, P_{i+1}, \dots, P_n)$ 
7      $P_i \leftarrow \text{FINDPATH}(\mathcal{G}, s_i, g_i, \mathcal{T})$ 
8 return  $\pi^{-1}(P_1, \dots, P_n)$ 

```

---

Alg. 1 shows the high level procedure that uses SU-I to generate individual paths. In lines 1-2, robots are sorted in *descending* distance-to-goal order. Then, paths for the robots are planned sequentially (lines 5-7) while using SU-I to avoid previously planned paths. The rationale for the descending ordering is that planning longer paths first helps quickly collect graph utilization information. As indicated in line 4, the path planning procedure is repeated for  $r$  iterations as more iterations capture the space utilization more accurately, improving the collision avoidance. Both robot ordering and number of iterations are tested in Section V.

### C. SU-I Heuristic Construction and Computation

At line 6, before each time we call SPP planner, BUILDSUO first builds an SU-I lookup table  $\mathcal{T}$  to memorize the space usage of existing paths. Denoting  $\mathcal{T}(v, t)$  (resp.  $\mathcal{T}(v_1, v_2, t)$ ) as the expected usage of vertex  $v$  (resp. directed edge  $v_1, v_2$ ) at time step  $t$ ,  $\mathcal{T}$  aggregates vertex, edge, and temporal information over all BUILDSUO input paths:

$$\begin{aligned} \mathcal{T}(v, t) &= \sum_P [(v) \preceq P_i], \\ \mathcal{T}(v_1, v_2, t) &= \sum_P [(v_1, v_2) \preceq P_i]. \end{aligned}$$

Here,  $[\cdot]$  on the right hand side of the equations is an *indicator* variable: it is 1 if the expression inside is true, otherwise it takes 0. We use  $\preceq$  to denote a sub-sequence relationship.

SU-I is mainly used to generate initial paths for MRPP algorithms. Since the initial paths can be modified (e.g. delayed or diverted for collision avoidance) in a full MRPP algorithm, we add integer parameters  $\alpha_L \geq 0$  and  $\alpha_H \geq 0$  to reason about graph usage in adjacent time steps and handle robot movement uncertainty:

$$\mathcal{T}(v, t) = \sum_P [(v) \preceq P(t - \alpha_H : t + \alpha_L)],$$

$$\mathcal{T}(v_1, v_2, t) = \sum_P [(v_1, v_2) \preceq P(t - \alpha_H : t + \alpha_L)].$$

For each path, the occupancy of graph utility at time  $t$  can now expand its influence to time steps in  $(t - \alpha_L, t + \alpha_H)$ . The values of  $\alpha_L, \alpha_H$  are empirically determined based on the MRPP algorithm itself; in Section V, we observe that such an temporal reasoning feature delivers better collision avoidance results when the movement of robots is uncertain.

We use  $\mathcal{T}$  to calculate SU-I heuristic and then use it during SPP path planning (line 7). Given  $(v_1, v_2, t)$  which means a robot moves from vertex  $v_1$  to vertex  $v_2$  at time  $t$ , with  $n$  as number of robots, the SU-I heuristic value is

$$H_{\text{SU-I}}(v_1, v_2, t) = \beta_v \frac{\mathcal{T}(v_2, t)}{n} + \beta_e \frac{\mathcal{T}(v_2, v_1, t)}{n}.$$

Here, the first term indicates the amount of vertex conflicts the robot encounters for the inferred state transition; the second term is associated with head-to-head edge collisions. Parameters  $\beta_v, \beta_e > 0$  balances between vertex and edge information;  $\beta_v + \beta_e = 1$ . When each vertex/edge is visited for at most once by each single robot path, we have

**Lemma III.1.** *When  $\beta_v + \beta_e = 1$ ,  $0 \leq H_{\text{SU-I}}(\cdot) < 1$ .*

The condition  $H_{\text{SU-I}}(\cdot) < 1$  is essential for SU-I to behave as a tie-breaker, which facilitates a good balance between single path optimality and congestion avoidance.

**Remark 1.** A special case for constructing  $\mathcal{T}$  is to ignore temporal information, i.e., instead of using  $(v, t)$  and  $(v_1, v_2, t)$  as lookup table keys, we use  $(v)$  and  $(v_1, v_2)$ . Thus,  $\mathcal{T}(v)$  (resp.,  $\mathcal{T}(v_1, v_2)$ ) simply records the total number of times the vertex  $v$  (resp., the edge  $(v_1, v_2)$ ) is used by the existing paths. This leads to smaller lookup tables but potentially worse collision avoidance as a result.

**Remark 2.** Note that for the actual implementation, in Alg. 1, line 6,  $\mathcal{T}$  is not re-constructed but updated based on the previous iteration which makes the computational complexity for the construction step  $O(|P|(\alpha_L + \alpha_H))$ .

Line 7 uses standard A\* to find a path from  $s_i$  to  $g_i$ . In the next two subsections, we discuss two ways to integrate SU-I into A\*. For simplicity, we now assume SU-I only uses vertex information (i.e.,  $\beta_v = 1, \beta_e = 0$ ) without temporal information (see Remark 1), unless otherwise specified.

### D. SU-I as Part of Estimated Cost-To-Go

We use  $H_{\text{short}}(v, g_i)$  to denote the shortest path distance between  $v$  and  $g_i$  in  $\mathcal{G}$ .  $H_{\text{short}}(v, g_i)$  is graph-dependent and can be calculated before path planning. For a grid without obstacles,  $H_{\text{short}}(v, g_i)$  is the Manhattan distance heuristic. The heuristic we use in A\* search is

$$\mathcal{H}(v) = H_{\text{short}}(v, g_i) + H_{\text{SU-I}}(v).$$

**Lemma III.2.** *A path planned using A\* search with  $\mathcal{H}$  as heuristic is a shortest path from  $s_i$  to  $g_i$ .*

*Proof.* See Appendix [39].  $\square$

$\mathcal{H}$  not only ensures a shortest path is found; the path also minimizes the maximum single step conflict. Given path  $P_i$ , its maximum single step conflict can be represented as

$$C_{\text{single}}(P_i) = \max_{0 < t < |P_i|} \mathcal{T}[P_i(t)].$$

**Lemma III.3.** A path  $P_i$  planned using A\* search with  $\mathcal{H}$  heuristic is a shortest path minimizing  $C_{\text{single}}(P_i)$ .

*Proof.* See Appendix [39].  $\square$

Now, given paths  $P_1, \dots, P_n$  returned from Alg. 1,  $\text{Im}(f)$  as the image of function  $f$ , we denote the maximum conflict on a single vertex as

$$C_{\text{single}} = \max_{v \in \mathcal{G}} \sum_{1 \leq i \leq n} [v \in \text{Im}(P_i)],$$

we reach the following property.

**Lemma III.4.**  $C_{\text{single}}$  cannot increase after the first SU-I planning iteration.

*Proof.* See Appendix [39].  $\square$

Lemma III.4 directly leads to the following theorem.

**Theorem III.1.**  $C_{\text{single}}$  will converge as the number of planning iterations increases.

#### E. SU-I as Part of Cost-To-Come

With the default transition cost as 1 for all states during A\* search, when aggregating SU-I into cost-to-come, we define the new transition cost leading to vertex  $v$  as

$$C(v) = 1 + H_{\text{SU-I}}(v) / (\max_{1 \leq i \leq n} H_{\text{short}}(s_i, g_i) + 1).$$

The collision avoidance property of using SU-I with cost-to-come is different from that of using SU-I with cost-to-go. We hereby define the number of vertex collisions on  $P_i$  as

$$C_{\text{path}}(P_i) = \sum_{1 \leq j \leq n, j \neq i} |\text{Im}(P_i) \cap \text{Im}(P_j)|.$$

**Lemma III.5.** A path  $P_i$  planned using A\* search with  $C$  as transition cost and an admissible heuristic is the shortest path which minimizes  $C_{\text{path}}(P_i)$ .

*Proof.* See Appendix [39].  $\square$

Given resulting paths as  $P_1, \dots, P_n$  when using SU-I as cost-to-come, we denote the total number of collisions as

$$C_{\text{path}} = \sum_{1 \leq i \leq n} C_{\text{path}}(P_i),$$

we find the following property.

**Lemma III.6.**  $C_{\text{path}}$  cannot increase after the first SU-I iteration.

*Proof.* See Appendix [39].  $\square$

Lemma III.6 directly leads to the following theorem.

**Theorem III.2.**  $C_{\text{path}}$  will converge as the number of planning iterations increases.

**Remark 3.** Regardless of whether SU-I is used as part of the cost-to-come or the cost-to-go, even though the properties mentioned above are only proved for SU-I without temporal information, similar properties exist for SU-I with temporal information when using state-time A\*. Instead of just considering vertex conflicts, all lemmas in this section remain true when edge conflicts are considered (i.e.  $\beta_e > 0$ ), except for Lemma III.4 and Lemma III.6. In evaluation, we empirically observe that running multiple planning iterations considering edge conflicts is still beneficial.

## IV. SPACE UTILIZATION OPTIMIZATION APPLICATION

The paths generated by SU-I can be directly used as input to some MRPP/LMPP algorithms. The more balanced graph utilization and reduced conflicts facilitate collision resolution, improving both computation time and solution optimality.

SU-I can also be combined with time-based divide-and-conquer to provide better intermediate goals. With a baseline structure adapted from [7], we first propose a *horizon cut* technique to reduce unnecessary node explorations and then use SU-I to further enhance its performance.

### A. Baseline Bounded-Horizon Search for LMPP

It is well known that solving an entire LMPP instance in one-shot is impractical, not only because the long lists of goals makes the problem computationally demanding, but also because the goal lists could be dynamically updated in real world scenarios, which invalidates the current solution and brings the need for online re-planning.

Given the above factors, LMPP is usually solved by a bounded-horizon approach. The basic idea is to plan the paths for the  $h$  time steps, execute the paths, and then iteratively re-plan and execute. Here,  $h$  is called the *planning horizon*. The pseudocode of such a baseline horizon-based structure [7] is provided in Alg. 2; readers may ignore lines 7-14 for now as we will discuss later. In the beginning, the goal list  $\mathbf{g}_i$  for each robot  $i$  is shortened until there is at most one unreachable goal for horizon  $h$  (lines 1-6); we use  $(-1)$  to index the last element in a sequence, and  $+$  to indicate sequence concatenation. Then, in line 15, the current state (denoted as  $S$ ) and the modified goal list are sent to a windowed MRPP solver. The behavior of the windowed solver is to output an  $h$ -step collision-free path set, where each robot  $i$  aims to traverse the shortened goal list  $\hat{\mathbf{g}}_i$  in order. The selection of the windowed MRPP solver is flexible. In this work, we use Bounded-Horizon Enhanced Collision Based Search [7] with suboptimality factor  $w = 1.5$  and treat it as a black box.

---

#### Algorithm 2: BOUNDEDHORIZONSEARCH

---

```

Input: Graph  $\mathcal{G}$ , current state  $S$ , goal lists  $\mathbf{G}$ , horizon  $h$ .
Output: A  $h$ -step solution
1  $\forall 1 \leq i \leq n, \hat{\mathbf{g}}_i \leftarrow (s_i), d_i \leftarrow 0$ 
2 for  $1 \leq i \leq n$  do
3   for  $g \in \mathbf{g}_i$  do
4      $d_i \leftarrow d_i + \text{SHORTESTDISTANCE}(\hat{\mathbf{g}}_i(-1), g)$ 
5      $\hat{\mathbf{g}}_i \leftarrow \hat{\mathbf{g}}_i + (g)$ 
6     if  $d_i \geq h$  then break
7 if using horizon cut then
8    $\forall 1 \leq i \leq n, P_i \leftarrow ()$ 
9   for  $1 \leq i \leq n$  do
10    if using SU-I then
11       $H_{\text{SU-I}} \leftarrow \text{BUILD SUO}(P_1, \dots, P_n)$ 
12       $P_i \leftarrow \text{FINDPATH}(\mathcal{G}, \hat{\mathbf{g}}_i(-2), \hat{\mathbf{g}}_i(-1), H_{\text{SU-I}})$ 
13    else  $P_i \leftarrow \text{FINDPATH}(\mathcal{G}, \hat{\mathbf{g}}_i(-2), \hat{\mathbf{g}}_i(-1))$ 
14     $\hat{\mathbf{g}}_i(-1) \leftarrow P_i(-(|P_i| - d_i + h - 1))$ 
15 return  $\text{WINDOWED SOLVER}(S, \hat{\mathbf{g}}_1(1 :), \dots, \hat{\mathbf{g}}_n(1 :), h)$ 

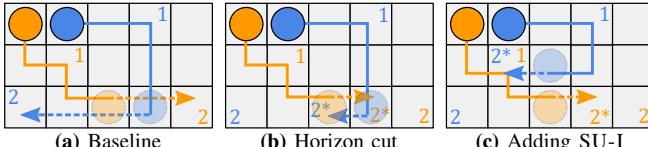
```

---

### B. Horizon Cut and SU-I Integration

Due to using SPP algorithm as subroutine, the baseline bounded-horizon method wastes computation power due to unnecessary reasoning about paths outside horizon  $h$ :

although these paths are not collision-checked, they are planned by the windowed solver and are discarded afterwards. To overcome this weakness, we propose *horizon cut*, which reduces the number of search nodes generated outside of  $h$  while still ensures that the robots move toward their future goals. Horizon cut further truncates the goal list by changing the last goal to a vertex that is in-between the second last goal (i.e.,  $\hat{g}_i(-2)$ ) and the last goal (see Alg. 2, lines 13-14). In our implementation, we find a shortest path between  $\hat{g}_i(-2)$  and  $\hat{g}_i(-1)$  and select the vertex at  $t = h + 1$ . The dashed paths in Fig. 2a and Fig. 2b demonstrate that we can avoid planning redundant paths when using horizon cut.

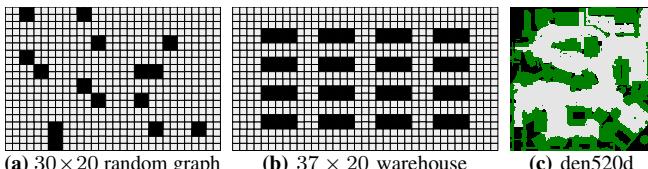


**Fig. 2:** An example comparing bounded-horizon baseline, horizon cut, and horizon cut with SU-I. The robots each has 2 goals to reach, marked as colored numbers. The planned configuration after the current horizon ( $h = 4$ ) is visualized as transparent disks. (a) The paths planned by the baseline bounded-horizon method. The dashed part is planned but not executed. (b) Using horizon cut, we avoid planning unnecessary steps by setting the last goals as  $2^*$ . (c) With SU-I, we have a better selection of  $2^*$  so that the conflicts in the next planning horizon may be avoided beforehand.

We then integrate SU-I to help select better target vertices (see Alg. 2, lines 11-12, which is a similar procedure as Alg. 1). By selecting target vertices on paths which better utilize graph resources, we can avoid conflicts in the future planning iterations. We visualize the effect in Fig. 2c, where we anticipate less conflicts between the two robots in the next planning and execution iteration when using SU-I.

## V. EVALUATION

We performed comprehensive evaluation of SU-I and associated algorithms on randomly generated graphs, warehouse-style graphs, and large DAO maps (Fig. 3 shows a subset of graphs). All experiments are performed on an Intel® Core™ i7-6900K CPU. Data points are averaged over 30 to 100 runs on randomly generated problem instances.



**Fig. 3:** Example of graphs used for evaluation. The bright cells visualize vertices. The black and green cells visualize obstacles.

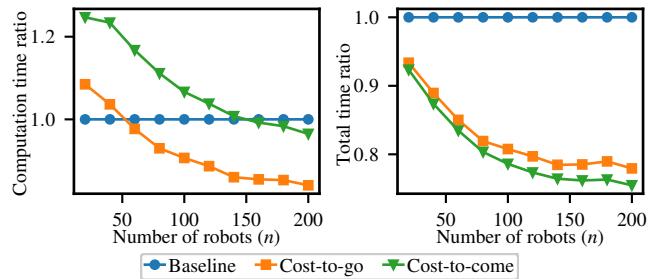
Our evaluation focuses on testing SU-I with complete MRPP algorithms. We also conducted analytical experiments, showing that the graph utilization conflicts indeed converge to some minimal value with SU-I, especially when considering both vertex and edge information with a descending ordering of robots. For details, see Appendix [39].

### A. SU-I in a Full MRPP Algorithm

Combining SU-I with an existing collision resolution method makes MRPP solver more efficient and optimal. For

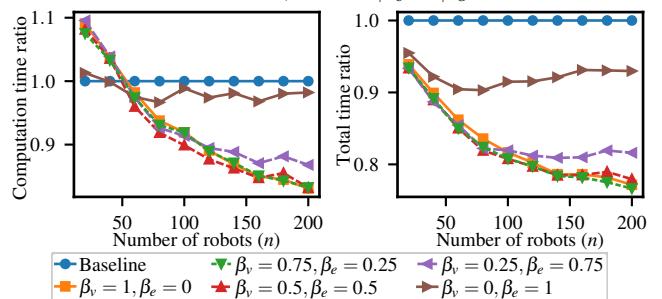
this evaluation, we use the database-driven collision avoidance routine from [5]. For test cases are in  $30 \times 20$  grids with 10% obstacles (see an random example in Fig. 3a), we report computation time and solution optimality (based on sum-of-cost) with varied number of robots. For both metrics, lower is better. All data points are normalized in terms of the baseline algorithm's performance where SU-I is not used.

Starting with  $\beta_v = \beta_e = 0.5$ ,  $r = 1$ , and no temporal information, we first compare using SU-I as a part of estimated cost-to-go and cost-to-come. Shown in Fig. 4, using SU-I as cost-to-go significantly reduces computation time (by 15%+) when robots have interactions. As a comparison, using SU-I as cost-to-come only makes the algorithm slightly more efficient when the number of robots is large. Both SU-I variations significantly improve the optimality, by up to 25%. The efficiency difference was due to SU-I as cost-to-come minimizes collisions along single paths and thus has a larger search space (state-time). Hence forth, we use SU-I as part of estimated cost-to-go by default.



**Fig. 4:** Comparison between using SU-I as a part of estimated cost-to-come and cost-to-go in a full MRPP algorithm.

We then try different  $\beta_v$ ,  $\beta_e$  to demonstrate that using vertex and edge information together is beneficial. Shown in Fig. 5, the performance improvement of using vertex information alone over the baseline is already significant. While using edge information alone is weaker than vertex, combining the two elements pushes both computation time and solution cost even lower. From now on, we set  $\beta_v = \beta_e = 0.5$ .



**Fig. 5:** SU-I MRPP solver evaluation with different  $\beta_v$ ,  $\beta_e$  values.

In Fig. 6, we added temporal information with different  $\alpha_L$  and  $\alpha_H$ . Due to the modifications to the initial paths in the second planning phase, using “soft” temporal information is beneficial: the green line ( $\alpha_L = \alpha_H = 0$ ) under-performs SU-I without temporal information, while the brown line ( $\alpha_L = 2$ ,  $\alpha_H = 15$ ) improves solution quality. Computation time wise, using temporal information adds overhead since it takes a longer time to construct SU-I. So whether to use temporal information is a choice to be made by practitioners.

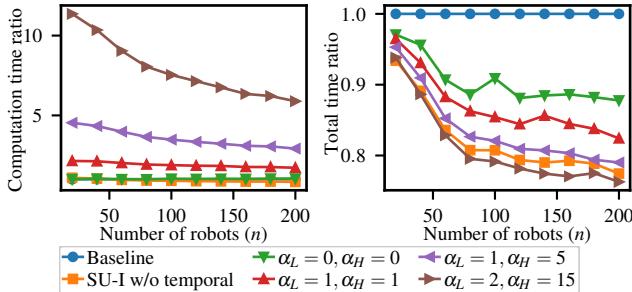


Fig. 6: Comparison between different  $\alpha_L$ ,  $\alpha_H$  values.

Fig. 7 shows that using multiple planning iterations generates better solutions when  $r \leq 4$ .

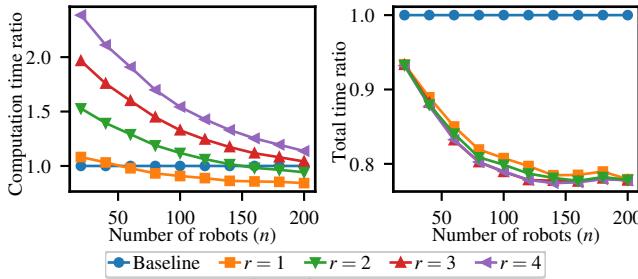


Fig. 7: Comparison between different planning iteration  $r$ .

### B. LMPP Bounded-Horizon Search with SU-I

We evaluate LMPP algorithms discussed in Sec. IV-A in a warehouse-style environment (see Fig. 3b). Each robot starts from a random vertex and is given a random list of goal vertices. The goal lists are continuously extended to make sure each robot always has future goals. The horizon-based planners with  $h = 5$  and ECBS [6] suboptimality factor  $w = 1.5$  are called iteratively until a total of 10000 goals are reached. The evaluated methods are the bounded-horizon search [7] (the baseline), with horizon cut, and with SU-I. We report the total computation time in Fig. 8, and the system throughput in Table I. The throughput is calculated as the average number of goals reached in a single time step; the higher, the better. The data shows that using horizon cut can effectively decrease the computation time by more than 50%. At the same time, using SU-I makes the computation time even lower (by about 65%), while keeps a same level of throughput as the baseline.

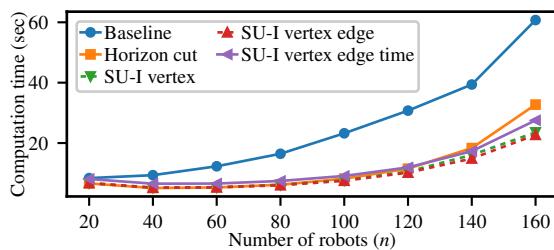


Fig. 8: Computation time of bounded-horizon methods.

TABLE I: Throughput of bounded-horizon methods

$n$	Baseline	Horizon cut	SU-I w/o temporal	SU-I w/ temporal
40	1.981	1.945	1.968	1.974
80	3.737	3.647	3.710	3.719
120	5.421	5.273	5.370	5.387
160	6.873	6.714	6.888	6.896

### C. Bounded-Horizon Search with SU-I on MRPP

As in the last evaluation, we directly apply bounded-horizon search to MRPP by considering MRPP as a special case for LMPP where all goal lists have length 1. The tested graph is DAO den520d (see Fig. 3c), a public MRPP benchmark [40]. The map size is  $257 \times 256$  with 28178 vertices. We set  $h = 50$ . Fig. 9 shows that SU-I remains effective in large environments with a sparse robot setup. Note that when not using horizon cut and SU-I, the baseline bounded-horizon search is significantly slower due to its unnecessary exploration. When we divide the solution quality metric over under-estimated lower-bounds (see Table II), we find that our method is able to generate solutions very close to the optimal.

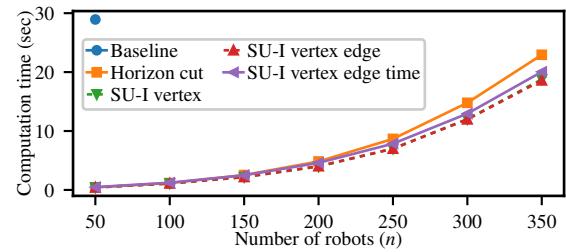


Fig. 9: Bounded-horizon search's computation time on a den520d.

TABLE II: Optimality ratio of MRPP solutions in den520d

$n$	50	100	150	200	250	300	350
Makespan	1.0002	1.0011	1.0020	1.0031	1.0040	1.0052	1.0064
Sum-of-time	1.0019	1.0023	1.0034	1.0051	1.0057	1.0076	1.0090

Experiments in Section V-A and Section V-C were repeated with different randomly generate graphs and more DAO maps (e.g., brc200d, lak201d). The comparison between methods show similar trends as demonstrated in this section.

## VI. CONCLUSION

With SU-I, we performed an in-depth exploration optimizing the space utilization for planning better individual robot paths in the first phase of a modern decoupled MRPP and LMPP pipeline. In addition to proving that SU-I's desirable properties as a heuristic, we validated the effectiveness of SU-I in reducing computation load and improving optimality for one-shot and life-long multi-robot path planning.

Together with [5], this research opens up a new direction in multi-robot path planning. In a sense, SUO, and its implementation, SU-I, are taking the decoupled multi-robot path planning paradigm a step further by reducing possible robot-robot interactions, making the process more like planning single robot paths with loose interactions. In future research, it would be interesting to exploit the SUO principle further to observe how far we can further minimize robot-robot interaction to boost the performance of the system. One immediate direction is to add weights to SU-I so that non-optimal single robot paths will be generated and gauge the trade-off between optimality loss at the first phase and the gain (in computation time and optimality) in the second phase of a decouple multi-robot path planner. Apart from handling uncertainty in time domain, we also plan to add mechanisms to treat space uncertainty.

**Acknowledgement.** We sincerely thank Joseph W. Durham for many insightful discussions of the work.

## REFERENCES

[1] A. Dekhne, G. Hastings, J. Murnane, and F. Neuhaus, “Automation in logistics: Big opportunity, bigger uncertainty,” *McKinsey Q*, pp. 1–12, 2019.

[2] P. R. Wurman, R. D’Andrea, and M. Mountz, “Coordinating hundreds of cooperative, autonomous vehicles in warehouses,” *AI Magazine*, vol. 29, no. 1, pp. 9–19, 2008.

[3] “Warehouse automation market with post-pandemic (covid-19) impact by technology, by industry, by geography - forecast to 2026,” <https://www.researchandmarkets.com/r/s6basv>, accessed: 2021-02-25.

[4] M. Erdmann and T. Lozano-Perez, “On multiple moving objects,” *Algorithmica*, vol. 2, no. 1, pp. 477–521, 1987.

[5] S. D. Han and J. Yu, “Ddm: Fast near-optimal multi-robot path planning using diversified-path and optimal sub-problem solution database heuristics,” *IEEE Robotics and Automation Letters*, vol. 5, no. 2, pp. 1350–1357, 2020.

[6] M. Barer, G. Sharon, R. Stern, and A. Felner, “Suboptimal variants of the conflict-based search algorithm for the multi-agent pathfinding problem,” in *Seventh Annual Symposium on Combinatorial Search*, 2014.

[7] J. Li, A. Tinka, S. Kiesel, J. W. Durham, T. S. Kumar, and S. Koenig, “Lifelong multi-agent path finding in large-scale warehouses,” in *Proceedings of the 19th International Conference on Autonomous Agents and MultiAgent Systems*, 2020, pp. 1898–1900.

[8] R. Stern, N. R. Sturtevant, A. Felner, S. Koenig, H. Ma, T. T. Walker, J. Li, D. Atzman, L. Cohen, T. S. Kumar, *et al.*, “Multi-agent pathfinding: Definitions, variants, and benchmarks,” in *Twelfth Annual Symposium on Combinatorial Search*, 2019, pp. 151–159.

[9] O. Goldreich, “Finding the shortest move-sequence in the graph-generalized 15-puzzle is NP-hard,” 1984, laboratory for Computer Science, Massachusetts Institute of Technology, Unpublished manuscript.

[10] Y. Guo and L. E. Parker, “A distributed and optimal motion planning approach for multiple mobile robots,” in *Proceedings IEEE International Conference on Robotics and Automation*, vol. 3, 2002, pp. 2612–2619.

[11] T. Standley and R. Korf, “Complete algorithms for cooperative pathfinding problems,” in *Proceedings International Joint Conference on Artificial Intelligence*, 2011, pp. 668–673.

[12] G. Wagner and H. Choset, “Subdimensional expansion for multirobot path planning,” *Artificial Intelligence*, vol. 219, pp. 1–24, 2015.

[13] E. Boyarski, A. Felner, R. Stern, G. Sharon, D. Tolpin, O. Betzalel, and E. Shimony, “Icbs: Improved conflict-based search algorithm for multi-agent pathfinding,” in *Twenty-Fourth International Joint Conference on Artificial Intelligence*, 2015.

[14] L. Cohen, T. Uras, T. Kumar, H. Xu, N. Ayanian, and S. Koenig, “Improved bounded-suboptimal multi-agent path finding solvers,” in *International Joint Conference on Artificial Intelligence*, 2016.

[15] P. Surynek, “Towards optimal cooperative path planning in hard setups through satisfiability solving,” in *Pacific Rim International Conference on Artificial Intelligence*. Springer, 2012, pp. 564–576.

[16] E. Erdem, D. G. Kisa, U. Oztok, and P. Schüller, “A general formal framework for pathfinding problems with multiple agents,” in *Twenty-Seventh AAAI Conference on Artificial Intelligence*, 2013.

[17] J. Yu and S. M. LaValle, “Optimal multi-robot path planning on graphs: Complete algorithms and effective heuristics,” *IEEE Transactions on Robotics*, vol. 32, no. 5, pp. 1163–1177, 2016.

[18] J. P. Van Den Berg and M. H. Overmars, “Prioritized motion planning for multiple robots,” in *2005 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2005, pp. 430–435.

[19] M. Bennewitz, W. Burgard, and S. Thrun, “Finding and optimizing solvable priority schemes for decoupled path planning techniques for teams of mobile robots,” *Robotics and autonomous systems*, vol. 41, no. 2, pp. 89–99, 2002.

[20] M. Saha and P. Isto, “Multi-robot motion planning by incremental coordination,” in *2006 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2006, pp. 5960–5963.

[21] J. van Den Berg, J. Snoeyink, M. C. Lin, and D. Manocha, “Centralized path planning for multiple robots: Optimal decoupling into sequential plans,” in *Robotics: Science and systems*, vol. 2, no. 2.5, 2009, pp. 2–3.

[22] T. Guo, S. D. Han, and J. Yu, “Spatial and temporal splitting heuristics for multi-robot motion planning,” in *IEEE International Conference on Robotics and Automation*, 2021.

[23] J. Yu, “Constant factor time optimal multi-robot routing on high-dimensional grids,” *Robotics: Science and Systems*, 2018.

[24] M. Szegedy and J. Yu, “Rubik tables and object rearrangement,” *The International Journal of Robotics Research*, 2022, to appear.

[25] K. Vedder and J. Biswas, “X\* anytime multiagent planning with bounded search,” in *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems*, 2019, pp. 2247–2249.

[26] D. Sigurdson, V. Bulitko, S. Koenig, C. Hernandez, and W. Yeoh, “Automatic algorithm selection in multi-agent pathfinding,” *arXiv preprint arXiv:1906.03992*, 2019.

[27] H. Ma, J. Li, T. Kumar, and S. Koenig, “Lifelong multi-agent path finding for online pickup and delivery tasks,” in *Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems*, 2017, pp. 837–845.

[28] W. Höning, S. Kiesel, A. Tinka, J. W. Durham, and N. Ayanian, “Persistent and robust execution of map schedules in warehouses,” *IEEE Robotics and Automation Letters*, vol. 4, no. 2, pp. 1125–1131, 2019.

[29] K. Okumura, M. Machida, X. Défago, and Y. Tamura, “Priority inheritance with backtracking for iterative multi-agent path finding,” in *Proceedings International Joint Conference on Artificial Intelligence*, 2019, pp. 535–542.

[30] M. S. Branicky, R. A. Knepper, and J. J. Kuffner, “Path and trajectory diversity: Theory and algorithms,” in *2008 IEEE International Conference on Robotics and Automation*. IEEE, 2008, pp. 1359–1364.

[31] L. H. Erickson and S. M. LaValle, “Survivability: Measuring and ensuring path diversity,” in *2009 IEEE International Conference on Robotics and Automation*. IEEE, 2009, pp. 2068–2073.

[32] R. A. Knepper and M. T. Mason, “Path diversity is only part of the problem,” in *2009 IEEE International Conference on Robotics and Automation*. IEEE, 2009, pp. 3224–3229.

[33] D. Silver, “Cooperative pathfinding,” in *Proceedings of the aaai conference on artificial intelligence and interactive digital entertainment*, vol. 1, no. 1, 2005, pp. 117–122.

[34] K.-H. C. Wang, A. Botea, *et al.*, “Fast and memory-efficient multi-agent pathfinding,” in *ICAPS*, 2008, pp. 380–387.

[35] M. R. Jansen and N. R. Sturtevant, “Direction maps for cooperative pathfinding,” in *AIIDE*, 2008.

[36] L. Cohen, T. Uras, and S. Koenig, “Feasibility study: Using highways for bounded-suboptimal multi-agent path finding,” in *International Symposium on Combinatorial Search*, vol. 6, no. 1, 2015.

[37] Y.-H. Lyu, Y. Chen, and D. Balkcom, “ $k$ -survivability: Diversity and survival of expendable robots,” *IEEE Robotics and Automation Letters*, vol. 1, no. 2, pp. 1164–1171, 2016.

[38] J. Yu, “Intractability of optimal multi-robot path planning on planar graphs,” *IEEE Robotics and Automation Letters*, vol. 1, no. 1, pp. 33–40, 2016.

[39] S. D. Han and J. Yu, “Optimizing space utilization for more effective multi-robot path planning,” <https://arxiv.org/abs/2109.04677>, 2021.

[40] N. R. Sturtevant, “Benchmarks for grid-based pathfinding,” *IEEE Transactions on Computational Intelligence and AI in Games*, vol. 4, no. 2, pp. 144–148, 2012.