

# MS\*: A New Exact Algorithm for Multi-agent Simultaneous Multi-goal Sequencing and Path Finding

Zhongqiang Ren<sup>1</sup>, Sivakumar Rathinam<sup>2</sup> and Howie Choset<sup>1</sup>

**Abstract**—In multi-agent applications such as surveillance and logistics, fleets of mobile agents are often expected to coordinate and safely visit a large number of goal locations as efficiently as possible. The multi-agent planning problem in these applications involves allocating and sequencing goals for each agent while simultaneously producing conflict-free paths for the agents. In this article, we introduce a new algorithm called MS\* which computes an optimal solution for this multi-agent problem by fusing and advancing state of the art solvers for multi-agent path finding (MAPF) and multiple travelling salesman problem (mTSP). MS\* leverages our prior subdimensional expansion approach for MAPF and embeds the mTSP solvers to optimally allocate and sequence goals for agents. Numerical results show that our new algorithm can solve the multi-agent problem with 20 agents and 50 goals in a minute of CPU time on a standard laptop.

## I. INTRODUCTION

Multi-agent applications such as information acquisition in surveillance or package delivery in logistic services require agents to safely visit a large number of goal locations as efficiently as possible. A fundamental problem in these applications, referred to as the Multi-agent simultaneous multi-goal sequencing and path finding (MSMP), aims to allocate and sequence goals for each agent and plan conflict-free paths for the agents while minimizing the sum of the travel costs. MSMP generalizes the single-agent path finding problem in [3] to multiple agents. MSMP also extends the multi-agent path finding (MAPF) problem [10], [20] found in the robotics community as well as the well-known multiple traveling salesman problem (mTSP) [12] addressed in the optimization community. Specifically, MAPF generally assumes the number of goals is equal to the number of agents and are mainly about finding conflict-free paths for the agents. On the other hand, the mTSP is a combinatorial problem that only addresses the allocation and sequencing of goals for agents without considering the conflicts between the paths of agents. The MSMP considers both path finding and goal sequencing concurrently (Fig. 1).

While several variants of the MAPF and mTSP have been addressed in the literature in isolation, we are not aware of any exact algorithm that finds an optimal solution for the MSMP problem. In this work, we take our first step to address this gap and propose a new algorithm named as




Fig. 1: An illustration of the challenges in MSMP. The stars, triangles and circles represent the starts, goals and destinations respectively. Each goal must be visited in the middle of some agent’s path while the planned paths of agents must also be conflict-free.

MS\*<sup>1</sup> and a variant named MS\*-c. We show that MS\*-c can solve the MSMP problem to optimality. In MS\* and MS\*-c, we create a new state space and handle conflicts via subdimensional expansion [25] while simultaneously allocating and sequencing goals for agents via state of the art mTSP solvers. Specifically, the subdimensional expansion method dynamically modifies the dimension of the new search space based on agent-agent conflicts and defers planning in the joint space until necessary. Concurrently, the complexity in goal allocation and sequencing is addressed by embedding the mTSP solvers in the form of (1) heuristics that underestimate the cost-to-go from any state, and (2) individual optimal policies that constructs the low dimensional search space for subdimensional expansion. To further reduce the computational complexity, we also define new dominance rules in the proposed state space to prune any partial solution that can not to lead to an optimal solution. Numerically, we perform simulations on instances from [20] with at most 20 agents and 50 goals to verify the performance of the proposed approach.

The rest of the article is organized as follows. Sec. II provides a literature review related to MSMP. We formally define the MSMP problem in Sec. III. Next, we describe the MS\* algorithm in Sec. IV, and discuss its completeness and optimality properties in Sec. V. Numerical results are presented in Sec. VI with conclusions in Sec. VII.

## II. PRIOR WORK

**Multi-agent path finding**, as its name suggests, computes an ensemble of collision-free paths that navigate a set of agents from their respective starts to destinations. Finding optimal paths for MAPF is NP-hard [26] as the problem complexity

<sup>1</sup> Zhongqiang Ren and Howie Choset are with the Robotics Institute and the Department of Mechanical Engineering at Carnegie Mellon University, 5000 Forbes Ave., Pittsburgh, PA 15213, USA. Emails: {zhongqir, choset}@andrew.cmu.edu

<sup>2</sup> Sivakumar Rathinam is with the Department of Mechanical Engineering, Texas A&M University, College Station, TX 77843-3123. Email: srathinam@tamu.edu

<sup>1</sup> Multi Steiner\* (MS\*): M refers to multi-agents and S refers to a Steiner traveling salesman path or a tour for each agent.

increases exponentially with respect to the number of agents. MAPF methods tend to fall on a spectrum from decentralized to centralized, trading off completeness and optimality for scalability. Conflict-based search [19] and subdimensional expansion [25] are widely known approaches for MAPF that lie in the middle of the spectrum. They are complete and guaranteed to provide optimal solutions; however, they typically handle scenarios where the number of goals or tasks is equal to the number of agents.

Subdimensional expansion [25], our prior method, bypasses the curse of dimensionality by dynamically modifying (while attempting to bound) the dimension of the search space based on agent-agent conflicts. This approach applies to many algorithms, and therefore inherits completeness and optimality, if the underlying algorithm already has these features, while scaling well on average in less crowded environments. Applying subdimensional expansion to A\* results in M\*. M\* begins by computing an individual optimal policy for every agent and lets them follow the policies towards their destinations, ignoring any other agents. When a conflict is found, the subset of agents in conflict are coupled together and a corresponding new search space with increased dimensions is formed where a new search is conducted in order to resolve the conflict locally.

**Multi-goal sequencing** is related to the traveling salesman problem (TSP), which aims to find an optimal tour in a given graph is one of the most well known NP-Hard problems in the literature [1]. Several methods have been developed [1] for the TSP ranging from exact techniques (branch and bound, branch and price) to heuristics and approximation algorithms which trade off optimality for computational efficiency. There are numerous generalizations of the TSP which include multiple agents with motion, fuel, capacity and other resource constraints [4], [22]. Multiple TSP (mTSP) is much harder than solving a single TSP since goal locations must be partitioned and allocated to each agent in addition to finding an optimal sequence of the assigned locations for each agent to visit. Our prior work in this area has developed several state of the art approximation algorithms for homogeneous [11], [16] and heterogeneous [4], [15], [2] agents using Lagrangian relaxation and primal-dual techniques. All of these methods ignore conflicts between the agent's paths. This issue is explicitly handled in this work.

**Integrated goal allocation and path finding** extends MAPF to consider multiple goals [9], [23], [10], [7], [8]. However, most of the prior work either consider the *allocation* aspect of the problem: assigning exactly one goal to each agent from a set of goals, or the *sequencing* aspect of the problem: a set of goals are pre-assigned to each agent and the visiting order is to be determined. In this work, the multi-goal sequencing in MSMP involves both allocation and sequencing.

### III. PROBLEM DESCRIPTION

Let index set  $I_N = \{1, 2, \dots, N\}$  denote a set of  $N$  agents. All agents move in a workspace represented as a finite graph  $G = (V, E)$ , where the vertex set  $V$  represents the possible locations for agents and the edge set  $E = V \times V$

denotes the set of all the possible actions that can move an agent between any two vertices in  $V$ . An edge between  $u, v \in V$  is denoted as  $(u, v) \in E$  and the cost of an edge  $e \in E$  is a non-negative real number  $\text{cost}(e) \in \mathbb{R}^+$ .

In this article, we use superscript  $i \in I_N$  over a variable to represent the specific agent to which the variable belongs (e.g.  $v^i \in V$  means a vertex corresponding to agent  $i$ ). Let  $v_o^i \in V$  denote the initial vertex of agent  $i$ . There are  $N$  destination vertices in  $G$  denoted by the set  $V_d \subseteq V$ . Each agent must terminate its path at any one of these destination vertices. Note here that each destination needs to be assigned to exactly one agent and vice-versa. In addition, let  $V_g = \{t_1, t_2, \dots, t_M\} \subseteq V$  denote the set of  $M$  goal vertices that must be visited by at least one of the agents in the middle of its path. For the rest of the article, we will use destinations and goals to represent vertices in  $V_d$  and  $V_g$  respectively.

Let  $\pi^i(v_1^i, v_\ell^i)$  denote a path for agent  $i$  that connects vertices  $v_1^i$  and  $v_\ell^i$  via a sequence of vertices  $(v_1^i, v_2^i, \dots, v_\ell^i)$  in the graph  $G$ . Let  $g^i(\pi^i(v_1^i, v_\ell^i))$  denote the cost associated with the path. This path cost is the sum of the costs of all the edges present in the path, i.e.,  $g^i(\pi^i(v_1^i, v_\ell^i)) = \sum_{j=1,2,\dots,\ell-1} \text{cost}(v_j^i, v_{j+1}^i)$ .

All agents share a global clock. Each action, either wait or move, requires one unit of time for any agent. Any two agents  $i, j \in I_N$  are claimed to be in conflict if one of the following two cases happens. The first case is a ‘‘vertex conflict’’ where two agents occupy the same vertex at the same time. The second case is an ‘‘edge conflict’’ where two agents travel through the same edge from opposite directions between times  $t$  and  $t + 1$  for some  $t$ .

The multi-agent simultaneous multi-goal sequencing and path finding (MSMP) problem aims to find a set of conflict-free paths for the agents such that (1) each goal is visited at least once by some agent, (2) the path for each agent starts at its initial vertices and terminates at one of the destinations, and (3) the sum of the cost of the paths is a minimum.

### IV. MS\* ALGORITHM

#### A. Preliminaries

Let  $\mathcal{G} = (\mathcal{V}, \mathcal{E}) = \underbrace{G \times G \times \dots \times G}_{N \text{ times}}$  denote the joint graph

which is the cartesian product of  $N$  copies of  $G$ , where each vertex  $v \in \mathcal{V}$  represents a joint vertex and  $e \in \mathcal{E}$  represents a joint edge that connects a pair of joint vertices. The joint vertex corresponding to the initial vertices of all agents is  $v_o = (v_o^1, v_o^2, \dots, v_o^N)$ . In addition, let  $\pi(u, v), u, v \in \mathcal{V}$  represent a joint path, which is a tuple of  $N$  individual paths of same length, i.e.,  $\pi(u, v) = (\pi^1(u^1, v^1), \dots, \pi^N(u^N, v^N))$ .

Let  $a \in \{0, 1\}^M$  denote a binary vector of length  $M$  that indicates the visiting status of goals, i.e., for  $m \in \{1, 2, \dots, M\}$ , the  $m$ -th component of  $a$ , denoted as  $a(m)$ , is equal to one if goal  $t_m \in V_g$  is visited by some agent, and  $a(m) = 0$  otherwise. To compare two binary vectors, a *dominance* relationship is established as follows.

*Definition 1 (Binary Dominance):* For any two binary vectors  $a_k$  and  $a_l$ ,  $a_k$  dominates  $a_l$  ( $a_k \succeq_b a_l$ ), if both the following conditions hold:

- $a_k(m) \geq a_l(m), \forall m \in \{1, 2, \dots, M\}$ .
- $a_k(m) > a_l(m), \exists m \in \{1, 2, \dots, M\}$ .

The state space in this work is the cartesian product of the joint graph and the space of binary vectors, *i.e.*,  $\mathcal{G} \times \{0, 1\}^M$ . A state  $s = (v, a)$  is a tuple which includes a joint vertex  $v \in \mathcal{V}$  occupied by the agents and a binary vector,  $a \in \{0, 1\}^M$ . For the rest of the work, we use  $v(s_k)$  and  $a(s_k)$  to denote the joint vertex and binary vector contained in state  $s_k$  respectively. The initial state of all agents is  $s_o = (v_o, 0^M)$ . The search algorithm conducts A\*-like [5] search and every state  $s_k$  represents a partial solution from  $s_o$  to  $s_k$ , which contains a joint path from  $v_o$  to  $v(s_k)$ . The algorithm terminates when one of the states in the final set  $S_f := \{(v_f, 1^M)\}$  is expanded, where  $v_f$  is a joint vertex that represents a permutation of the destinations in  $V_d$ . Additionally, for any state  $s_k$  during the search process, as in A\*, let  $g(s_k)$  denotes the least cost of any joint path found thus far from  $s_o$  to  $s_k$ . To compare two states, we define a dominance rule between them as follows.

*Definition 2 (State Dominance):* For any two states  $s_k$  and  $s_l$  with  $v(s_k) = v(s_l)$ ,  $s_k$  dominates  $s_l$  if either of the following two conditions holds:

- $a(s_k) \succeq_b a(s_l), g(s_k) \leq g(s_l)$ , or
- $a(s_k) = a(s_l), g(s_k) < g(s_l)$ .

If  $s_k$  does not dominate  $s_l$ ,  $s_l$  is then non-dominated by  $s_k$ . Let  $\alpha(v) = \{s_k | v(s_k) = v\}$  denote a set of states generated by the algorithm at joint vertex  $v$ . Initially,  $\alpha(v_o)$  contains  $s_o$  only, and  $\alpha(v) = \emptyset, \forall v \in \mathcal{V}, v \neq v_o$ . When a new state  $s_l$  is generated by the algorithm during the search,  $s_l$  is inserted into  $\alpha(v(s_l))$  only when  $s_l$  is non-dominated by any states in  $\alpha(v(s_l))$ .

A heuristic function is also introduced to guide the search, which maps a state to a non-negative real number (*i.e.*  $h(s) \geq 0$ ) that underestimates the cost-to-go before finding a solution. In addition, let  $f(s) := g(s) + h(s)$  represent the  $f$ -value of state  $s$ , and at any stage of the algorithm, let OPEN denote the open list which contains candidate states to be expanded. As commonly done, we sort the candidate states in OPEN based on their  $f$ -values.

In addition, let collision set  $I_C(s) \subseteq I_N$  represent a set of agents in conflict at state  $s$ . To detect conflicts, collision function  $\Psi : \mathcal{V} \times \mathcal{V} \rightarrow 2^{I_N}$  is introduced to check if there is any vertex or edge conflicts given two neighboring joint vertices. Collision function  $\Psi$  returns either an empty set if no conflict is detected, or the set of agents that are in conflict.

### B. Key features of the Algorithm

MS\* (Algorithm 1) conducts a M\*-like [25] search in the new state space to find an optimal joint path for the agents. However, there are **three crucial aspects** where MS\* differs from M\* that makes it applicable to the MSMP. **First**, to construct individual policies and the heuristic value from a state to any (final) state in  $S_f$ , we ignore the conflicts between the agent's paths and solve the resulting mTSP. The mTSP in this paper is solved by transforming it to an equivalent TSP and using the LKH algorithm [6] (a state of the art TSP solver). The solution to the transformed TSP

can then be readily converted into a path for each agent as explained later in section IV-D. **Second**, when a conflict is found, back-propagating only the subset of agents involved in the conflict may not be sufficient to guarantee an optimal solution for MSMP. Therefore, we have two variants of the MS\* algorithm, one which back-propagates in a similar way as the standard M\* and another called MS\*-c which back-propagates the entire set of agents. While the numerical results show that MS\* finds optimal solutions for almost all of the instances, our proof of optimality currently only applies to MS\*-c. **Third**, we use the state dominance rules to prune any state from OPEN, which can not be part of an optimal solution.

---

### Algorithm 1 Pseudocode for MS\* (MS\*-c)

---

```

1: Initialize OPEN with  $s_o = (v_o, 0^M)$ 
2: while OPEN not empty do
3:    $s_k = (v_k, a_k) \leftarrow \text{OPEN.pop}()$ 
4:   if  $s_k \in S_f$  then
5:      $\pi \leftarrow \text{Reconstruct}(s_k)$ 
6:     return  $\pi$ 
7:    $S^{ngh} \leftarrow \text{GetNeighbor}(s_k)$ 
8:   for all  $s_l = (v_l, a_l) \in S^{ngh}$  do
9:     add  $s_k$  to back_set( $s_l$ )
10:    if  $\Psi(v_k, v_l) \neq \emptyset$ 
11:       $I_C(s_l) \leftarrow I_C(s_l) \cup \Psi(v_k, v_l)$ 
12:      (Notes: for MS*-c,  $I_C(s_l) \leftarrow I_N$ )
13:      BackProp( $s_k, I_C(s_l)$ )
14:      continue
15:     $f(s_l) \leftarrow g(s_l) + h(s_l)$ 
16:    if  $s_l$  is non-dominated by any states in  $\alpha(v_l)$  then
17:      add  $s_l$  to OPEN
18:      parent( $s_l$ )  $\leftarrow s_k$ 
19:    else (Notes:  $s_l$  is dominated by states in  $\alpha(v_l)$ )
20:       $A \leftarrow \{s'_l | s'_l \in \alpha(v_l), s'_l \succeq s_l\}$ 
21:      for all  $s'_l \in A$  do
22:        BackProp( $s_k, I_C(s'_l)$ )
23:        add  $s_k$  to back_set( $s'_l$ )
24: return Failure (no solution)

```

---

MS\* starts by inserting the initial state  $s_o$  into OPEN. In every round of the search, a state  $s_k$  with a minimum  $f$ -value is popped from OPEN. If  $s_k \in S_f$ , the algorithm outputs a solution by simply tracing the parents of the states iteratively from  $s_k$  to  $s_o$  and terminates. Otherwise,  $s_k$  is further expanded.

To expand  $s_k$ , all neighboring state in the limited neighbor set (denoted as  $S^{ngh}$ ), of  $s_k$  are explored (see Sec. IV-C). For each  $s_l \in S^{ngh}$ , the collision checking function  $\Psi(v(s_k), v(s_l))$  is invoked and the resulting collision set  $I_C(s_l)$  is back-propagated, which updates the collision set of the parent states recursively. As noted before, back-propagating  $I_C(s)$  may not be enough to guarantee that the algorithm returns an optimal solution. Thus, in MS\*-c, the index set  $I_N$ , which contains all agents, is back-propagated instead of  $I_C(s_l)$ , when there is conflict detected at state  $s_l$ . Any neighboring state  $s_l$  that leads to a conflict is discarded after the collision sets are back-propagated (line 13). If  $s_l$  is not in conflict,  $s_l$  is then check for dominance against states in  $\alpha(v(s_l))$ . If  $s_l$  is not dominated, it is added to OPEN for

future expansion. Otherwise,  $s_l$  is discarded and for all states  $s_l'$  that dominates  $s_l$ , *i.e.*,  $\{s_l' | s_l' \in \alpha(v_l), s_l' \succeq s_l\}$ ,  $I_C(s_l')$  is backpropagated to  $s_k$ . By doing so, MS\* (and MS\*-c) keeps updating the collision set of ancestor states of  $s_l$  after  $s_l$  is pruned by dominance [18]. Finally, MS\* either finds a solution or terminates when OPEN becomes empty, which indicates there is no solution for the problem. More details of the algorithm are presented in the following subsections.

### C. State Expansion

To expand a state, we leverage the concept of “limited neighbors”, which was originally introduced in our prior work [25]. The limited neighbors  $S^{ngh}$  of a state  $s_k$  is a set of neighboring states  $s_l$  that can be reached from  $s_k$  and is determined by  $I_C(s_k)$ . Denote  $v_k = v(s_k)$  and for each agent  $i$ , if  $i$  is not part of the collision set ( $i \notin I_C(s)$ ), agent  $i$  is only allowed to follow the path  $\phi^i(v_k^i, v_d^i)$  which moves it from  $v_k^i$  to its assigned destination  $v_d^i$  (see Sec. IV-D). If agent  $i$  is in the collision set ( $i \in I_C(s_k)$ ), it is allowed to explore all possible neighbors of  $v_k^i$  in  $G$ . The binary vector  $a(s_l)$  is updated correspondingly if any un-visited goals get visited in  $v_l^i$  by some agent  $i \in I_N$ , where  $v_l = v(s_l)$ .

---

#### Algorithm 2 Pseudocode for BackProp

---

```

1: INPUT:  $s, I$ 
2: if  $I \not\subseteq I_C(s)$  then
3:    $I_C(s) \leftarrow I \cup I_C(s)$ 
4:   if  $s \notin \text{OPEN}$  then
5:     add  $s$  to OPEN
6:   for all  $s_k \in \text{back\_set}(s)$  do
7:     BackProp( $s_k, I_C(s)$ )

```

---

$S^{ngh}$  of a state  $s$  varies once  $I_C(s)$  changes, which dynamically modifies the sub-space embedded in the overall state space that can be reached from  $s$ . Note that  $I_C(s)$  is updated recursively when Algorithm 2 is called. To back-propagate collision sets, a data structure “back\_set” is defined at every state. Intuitively, the back\_set at state  $s$  contains all parent states from which  $s$  is reached. When  $I_C(s)$  is enlarged,  $I_C(s)$  is back-propagated to every state in back\_set( $s$ ). When back-propagating a collision set  $I_C(s)$  to a state  $s_k \in \text{back\_set}(s)$ , if  $I_C(s_k)$  is not a super set of  $I_C(s)$ , then  $I_C(s_k)$  is updated by taking the union of  $I_C(s)$  (line 3). In addition,  $s_k$  is re-opened, which makes MS\* expand  $s_k$  again with a possibly larger limited neighbor set.

The aforementioned state expansion strategy is “aggressive” in a sense that, at a state  $s$ , only agents  $i \in I_C(s)$  are allowed to explore all possible moves and agents  $i \notin I_C(s)$  are restricted to their individual optimal policies. However, in MSMP, all agents are “coupled” in the space of binary vectors  $\{0, 1\}^M$  as a goal visited by one agent does not need to be visited by any other agents. Therefore, when agent  $i \in I_C(s)$  changes its next move at state  $s$ , the optimal policies of other agents  $j \notin I_C(s)$  may change. This requires agents  $j \notin I_C(s)$  to consider all possible moves as well in order to guarantee the optimality. Based on this observation, we propose a “conservative” (-c) expansion strategy and name

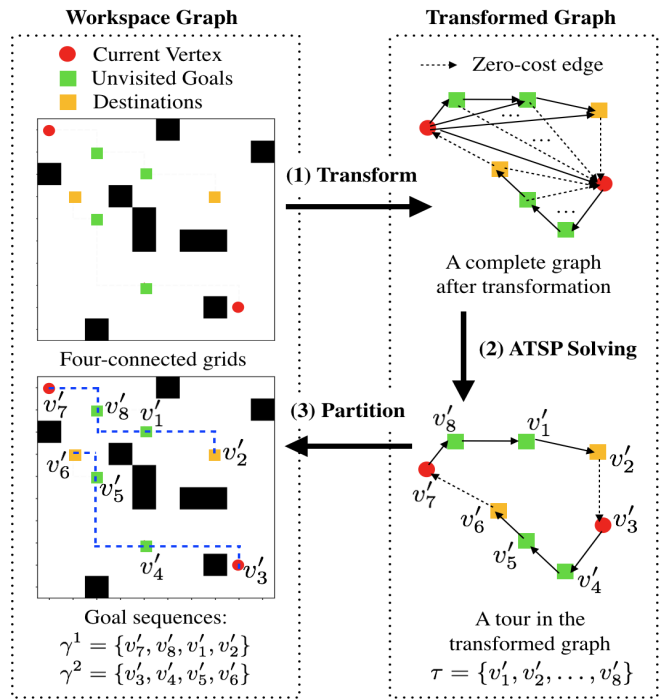


Fig. 2: The workflow of multi-goal sequencing.

the corresponding MS\* algorithm with this strategy as MS\*-c. In MS\*-c, when a collision is detected between agents  $I_C(s) \subseteq I_N$  at state  $s$ , instead of back-propagating  $I_C(s)$  as MS\* does, MS\*-c back-propagates  $I_N$ , the index set that includes all agents. By doing so, all possible combinations of individual agent moves are considered in the state expansion. Notice here that MS\*-c still differs from a naive A\* search in the state space because for states where collision sets are not back-propagated, agents are still restricted to follow their optimal policies.

### D. Multi-goal Sequencing, Policies and Heuristics

**Multi-goal Sequencing.** Ignoring the possible conflicts, the (multi-)goal sequencing procedure (Fig. 2) takes a state  $s$  as input and outputs a set of *goal sequences* for all agents. Each goal sequence is an ordered list of vertices in  $G'$ , which starts from current vertex of an agent, as specified in  $v(s)$ , visits a subset of un-visited goals assigned to the agent in order, and ends at some destination vertex in  $V_d$ . This procedure has three steps:

- The goal sequencing problem in  $G$  is transformed [13] to a single asymmetrical traveling salesman problem (ATSP) in a transformed graph  $G'$ .
- Solve the ATSP in  $G'$  and get an optimal tour  $\tau_*$  in  $G'$ .
- Partition the optimal tour  $\tau_*$  in  $G'$  into  $N$  goal sequences for the agents in  $G$ .

The transformation step leverages our prior work [13]. We use symbols with prime (e.g.  $G'$ ) to denote the variables after transformation. Given a state  $s$ , let  $V_g(s)$  denote the subset of goals that are not visited, *i.e.*,  $V_g(s) = \{t_m : a(m) = 0, \forall m = 1, 2, \dots, M, a = a(s)\}$ . Let  $G' = (V', E')$  represent a transformed graph, where  $V'$  is defined as the

union of (i) the current vertices of the agents, (ii) the unvisited goals in  $V_g(s)$ , and (iii) the destinations  $V_d$ , *i.e.*,  $V' = \{v^i | v = v(s), \forall i \in I_N\} \cup V_g(s) \cup V_d$ . Here  $G'$  is a complete (*i.e.* fully-connected) directed graph and the travel cost,  $cost'(v'_a, v'_b)$ , for any edge joining  $v'_a \in V'$  and  $v'_b \in V'$  is defined in Table I, where  $g(\pi_*)$  denotes the cost of a min-cost path in  $G$  from  $v'_a$  to  $v'_b$ , which can be computed by any single-agent min-cost path algorithms such as A\* [5].

$cost'(v'_a, v'_b)$	$v'_b \in v$	$v'_b \in V_g(s)$	$v'_b \in V_d$
$v'_a \in v$	$\infty$	$g(\pi_*)$	$g(\pi_*)$
$v'_a \in V_g(s)$	$\infty$	$g(\pi_*)$	$g(\pi_*)$
$v'_a \in V_d$	0	$\infty$	$\infty$

TABLE I: Edge cost definition in the transformed graph  $G'$ .

The LKH algorithm [24]<sup>2</sup>, which is one of the best heuristic-based solver currently available for TSP, is used to find a tour in  $G'$ . The LKH, in addition to providing a near-optimal solution, also provides a tight lower bound which can be used as an underestimate for the cost to go from state  $s$ . It has been shown in [13] that the near-optimal tour consists of exactly  $N$  zero cost edges. By construction, each of these zero cost edges connects a destination in  $V_d$  to a current vertex of an agent. This allows us to remove all the zero cost edges from the tour and obtain a *goal sequence*  $\gamma^i := (u_1^i, u_2^i, \dots, u_{\ell_i}^i)$ ,  $i \in I_N$  for each agent  $i = 1, \dots, N$ , where  $u_1^i = v^i$  is the current vertex of agent  $i$ ,  $u_{\ell_i}^i$  is the assigned destination of agent  $i$  and all other intermediate vertices  $u_2^i, \dots, u_{\ell_i-1}^i$  are un-visited goals assigned to agent  $i$  in order.

**Policy Construction and Heuristic Value.** The goal sequence  $\gamma^i$  can be further extended to a path  $\pi^i(v^i, u_{\ell_i}^i)$  for agent  $i$ , where any two adjacent vertices in  $\gamma^i$  is replaced with a min-cost path between them in  $G$ . The path  $\pi^i(v^i, u_{\ell_i}^i)$  for agent  $i$  provides an individual policy  $\phi^i$ , which maps a vertex to the next vertex along the path and thus direct agent  $i$  from  $v^i$  to  $u_{\ell_i}^i$  if no conflict detected along the path. Consequently, with goal sequencing at state  $s$ , a joint path  $\pi$  is found for all agents, which moves the agents from state  $s$  to a final state in  $S_f$ . Since any possible conflicts are ignored, the sum of the cost of individual paths in  $\pi$  provides an underestimate of the cost-to-go at state  $s$  and is used as  $h(s)$ .

**Remark.** The procedure used here is one method to for goal sequencing (or the mTSP). Several other exact algorithms, such as branch-and-cut [21] can be used in place of the transformation/LKH method. An exact algorithm provides an optimal tour for goal sequencing, which in turn allows MS\* to provide an optimal solution for MSMP. Bounded sub-optimal tour leads to inflated heuristic values for states in MS\*. As MS\* select states based on  $f$ -values as A\* does, inflated heuristic values lead to bounded sub-optimal solutions for the MSMP [14]. In the ensuing analysis, to outline the main ideas, we assume that the goal sequencing procedure returns an optimal solution.

## V. ANALYSIS

We provide a sketch proof of the properties of MS\*-c assuming the (multi-)goal sequencing returns an optimal tour.

<sup>2</sup><http://akira.ruc.dk/~keld/research/LKH/>

More discussion can be found in [17].

**Lemma 1:** Given a state  $s$ , let  $C(\tau_*)$  denote the cost of the tour  $\tau_*$  computed by the goal sequencing procedure and let  $\pi_m(v(s), v_f)$  denote a joint path in  $G$  with the minimum cost that starts from  $s$ , visits all un-visited goals and ends at destinations. Then  $C(\tau_*) = g(\pi_m(v(s), v_f))$ .

**Lemma 2:** If there is no feasible solution, MS\*-c reports failure in finite time.

**Theorem 1:** MS\*-c either computes the optimal solution for MSMP or reports failure in finite time.

**Proof:** With Lemma 1, if no conflict, MS\* restricts agents to move along the joint path computed by the goal sequencing procedure and the resulting joint path is guaranteed to be optimal. If there are conflicts at state  $s$  along the joint path, with Alg. 2, all ancestor states of  $s$  are updated with  $I_C = I_N$  and re-inserted into OPEN for re-expansion with all possible moves considered. Among these moves, MS\*-c uses an underestimating heuristic and expands states as in A\*, and thus computes an optimal solution. With Lemma 2, if no solution, MS\*-c reports failure in finite time. ■

## VI. NUMERICAL RESULTS

### A. Test Settings

We implemented both MS\* and MS\*-c in Python and tested on a computer with a CPU of Intel Core i7 and 16GB RAM. We selected grids from [20] and make them four-connected. Each action of an agent, either wait or move, incurs a unit cost. For each grid, we generated a set of  $M$  goals with  $M = 10, 20, \dots, 50$  at random location without overlapping with agents' starts or destinations. We tested with number of agents  $N = 5, 10, 20$ . For each test instance, the time bound is *one minute*. We tested both MS\* and MS\*-c, and they return the same solution cost for any instance when both of them succeed. In addition, the success rates of MS\* is higher than MS\*-c within 5% for all the instances. We thus report only the results of MS\* below to make the article concise.

### B. Success Rate

We first measured the success rates of MS\* with varying  $M$  and  $N$ . As shown in Fig. 3, as  $M$  increases, MS\* achieves relatively stable success rates except the rightmost game grid, where the initial goal sequencing procedure times out with  $M = 50$ . When  $N$  increases, the chance that agents collide increases when they follow the policies computed by the goal sequencing procedure, which lowers the success rates. The types of the grid also affects the success rates.

### C. Multi-goal Sequencing Time

To get a better sense of the computational burden of (multi-)goal sequencing in MS\*, we measured the run time of each goal sequencing call and took the average over all the succeeded test instances. From Fig. 3, as one can expect, as  $N$  or  $M$  increases, each call of goal sequencing takes more time. When  $M = 50$  and  $N = 20$ , there are 20 agents + 50 goals + 20 destinations = 90 vertices in the transformed graph  $G'$ . In random, room and maze grid, it takes roughly 30 seconds to compute the transformed graph

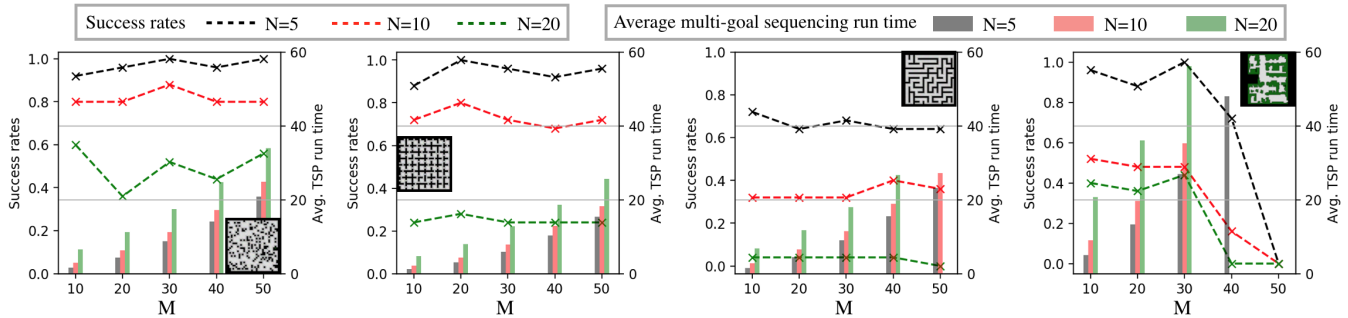


Fig. 3: The left vertical axis of each plot represents the success rates of MS\* as a function of both the number of agents ( $N$ ) and number of goals ( $M$ ). The right vertical axis of each plot (bar plots) represents the average run time (in seconds) of goal sequencing per call as  $M$  varies. The grids used from the left to the right are of type (size): random (32x32), room (32x32), maze (32x32), game map (65x81)

and call LKH to compute a set of goal sequences, which is more than half of the time limit and is very burdensome.

## VII. CONCLUSION

New algorithms, MS\* and its variant MS\*-c were presented for a multi-agent simultaneous multi-goal sequencing and path finding (MSMP) problem. The completeness and optimality properties of the algorithms were showed. Numerical results verifies the performance of both algorithms. There are several directions for future work. One can focus on variants of MSMP and MS\* can be suitably modified to handle the variants. One can also focus on fast approximation algorithms for goal sequencing which can improve the performance of the MS\*.

## REFERENCES

- [1] D. L. Applegate, R. E. Bixby, V. Chvatal, and W. J. Cook. *The Traveling Salesman Problem: A Computational Study (Princeton Series in Applied Mathematics)*. Princeton University Press, Princeton, NJ, USA, 2007.
- [2] Jungyun Bae and Sivakumar Rathinam. A primal-dual approximation algorithm for a two depot heterogeneous traveling salesman problem. *Optimization Letters*, pages 1–17, 2015.
- [3] K. Chour, S. Rathinam, and R. Ravi. S\*: A heuristic information-based approximation framework for multi-goal path finding. In *International Conference on Automated Planning and Scheduling*, 2021.
- [4] R. Doshi, S. Yadlapalli, S. Rathinam, and S. Darbha. Approximation algorithms and heuristics for a 2-depot, heterogeneous hamiltonian path problem. *International Journal of Robust and Nonlinear Control*, 21(12):1434–1451, 2011.
- [5] P. E. Hart, N. J. Nilsson, and B. Raphael. A formal basis for the heuristic determination of minimum cost paths. *IEEE Transactions on Systems Science and Cybernetics*, 4(2):100–107, 1968.
- [6] K. Helsgaun. An effective implementation of the Lin-Kernighan traveling salesman heuristic. *European Journal of Operational Research*, 126(1):106–130, 2000.
- [7] Christian Henkel, Jannik Abbenseth, and Marc Toussaint. An optimal algorithm to solve the combined task allocation and path finding problem. *arXiv preprint arXiv:1907.10360*, 2019.
- [8] Wolfgang Hönig, Scott Kiesel, Andrew Tinka, Joseph Durham, and Nora Ayanian. Conflict-based search with optimal task assignment. In *Proceedings of the International Joint Conference on Autonomous Agents and Multiagent Systems*, 2018.
- [9] Minghua Liu, Hang Ma, Jiaoyang Li, and Sven Koenig. Task and path planning for multi-agent pickup and delivery. In *AAMAS*, pages 1152–1160, 2019.
- [10] Hang Ma and Sven Koenig. Optimal target assignment and path finding for teams of agents. In *Proceedings of the 2016 International Conference on Autonomous Agents & Multiagent Systems*, pages 1144–1152, 2016.
- [11] W. Malik, S. Rathinam, and S. Darbha. An approximation algorithm for a symmetric generalized multiple depot, multiple travelling salesman problem. *Operations Research Letters*, 35:747–753, 2007.
- [12] P. Oberlin, S. Rathinam, and S. Darbha. Today’s traveling salesman problem. *IEEE Robotics and Automation Magazine*, 17(4):70–77, December 2010.
- [13] Paul Oberlin, Sivakumar Rathinam, and Swaroop Darbha. A transformation for a multiple depot, multiple traveling salesman problem. In *2009 American Control Conference*, pages 2636–2641. IEEE, 2009.
- [14] J. Pearl and J. H. Kim. Studies in semi-admissible heuristics. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, PAMI-4(4):392–399, 1982.
- [15] S. Rathinam, R. Ravi, J. Bae, and K. Sundar. Primal-Dual 2-Approximation Algorithm for the Monotonic Multiple Depot Heterogeneous Traveling Salesman Problem. In Susanne Albers, editor, *17th Scandinavian Symposium and Workshops on Algorithm Theory (SWAT 2020)*, volume 162 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 33:1–33:13, Dagstuhl, Germany, 2020. Schloss Dagstuhl–Leibniz-Zentrum für Informatik.
- [16] S. Rathinam, R. Sengupta, and S. Darbha. A resource allocation algorithm for multivehicle systems with nonholonomic constraints. *IEEE Transactions Automation Science and Engineering*, 4:98–104, 2007.
- [17] Zhongqiang Ren, Sivakumar Rathinam, and Howie Choset. Ms\*: A new exact algorithm for multi-agent simultaneous multi-goal sequencing and path finding. *arXiv preprint arXiv:2103.09979*, 2021.
- [18] Zhongqiang Ren, Sivakumar Rathinam, and Howie Choset. Subdimensional expansion for multi-objective multi-agent path finding. *arXiv preprint arXiv:2102.01353*, 2021.
- [19] Guni Sharon, Roni Stern, Ariel Felner, and Nathan R Sturtevant. Conflict-based search for optimal multi-agent pathfinding. *Artificial Intelligence*, 219:40–66, 2015.
- [20] Roni Stern, Nathan Sturtevant, Ariel Felner, Sven Koenig, Hang Ma, Thayne Walker, Jiaoyang Li, Dor Atzmon, Liron Cohen, TK Kumar, et al. Multi-agent pathfinding: Definitions, variants, and benchmarks. *arXiv preprint arXiv:1906.08291*, 2019.
- [21] Kaarthik Sundar and Sivakumar Rathinam. An exact algorithm for a heterogeneous, multiple depot, multiple traveling salesman problem. In *Unmanned Aircraft Systems (ICUAS), 2015 International Conference on*, pages 366–371. IEEE, 2015.
- [22] Kaarthik Sundar and Sivakumar Rathinam. Generalized multiple depot traveling salesmen problem-polyhedral study and exact algorithm. *Computers & Operations Research*, 70:39–55, 2016.
- [23] Pavel Surynek. Multi-goal multi-agent path finding via decoupled and integrated goal vertex ordering. *arXiv preprint arXiv:2009.05161*, 2020.
- [24] Renato Tinós, Keld Helsgaun, and Darrell Whitley. Efficient recombination in the lin-kernighan-helsgaun traveling salesman heuristic. In *International Conference on Parallel Problem Solving from Nature*, pages 95–107. Springer, 2018.
- [25] Glenn Wagner and Howie Choset. Subdimensional expansion for multirobot path planning. *Artificial Intelligence*, 219:1–24, 2015.
- [26] Jingjin Yu and Steven M LaValle. Structure and intractability of optimal multi-robot path planning on graphs. In *Twenty-Seventh AAAI Conference on Artificial Intelligence*, 2013.