

Multi-agent Path Planning and Network Flow

Jingjin Yu and Steven M. LaValle

Abstract. This paper connects multi-agent path planning on graphs (roadmaps) to network flow problems, showing that the former can be reduced to the latter, therefore enabling the application of combinatorial network flow algorithms, as well as general linear program techniques, to multi-agent path planning problems on graphs. Exploiting this connection, we show that when the goals are permutation invariant, the problem always has a feasible solution path set with a longest finish time of no more than $n + V - 1$ steps, in which n is the number of agents and V is the number of vertices of the underlying graph. We then give a complete algorithm that finds such a solution in $O(nVE)$ time, with E being the number of edges of the graph. Taking a further step, we study time and distance optimality of the feasible solutions, show that they have a pairwise Pareto optimal structure, and again provide efficient algorithms for optimizing two of these practical objectives.

1 Introduction

Consider the problem illustrated in Fig. 1, which inspired the authors to pursue this research. As an exercise (26-1 in [7]), the *escape problem* is to determine, given $m \leq n^2$ evaders placed on m different points of an $n \times n$ grid, whether there are m vertex disjoint paths from these m locations to m different points on the boundary of the grid. Intended as a demonstration of applications of *maximum flow* algorithms (Ch. 26 of [7]), it undoubtedly mimics multi-agent path planning problems on

Jingjin Yu

Department of Electrical and Computer Engineering,
University of Illinois, Urbana-Champaign
e-mail: jyu18@uiuc.edu

Steven M. LaValle

Department of Computer Science,
University of Illinois, Urbana-Champaign
e-mail: lavalle@uiuc.edu



Fig. 1 Examples of the escape problem on a 6×6 grid. The black discs are the initial evader locations. The goal is to plan disjoint paths for the evaders to reach different vertices on the boundary of the grid. a) An instance with solution given as the bold edges. b) An instance without a solution.

graphs. Intrigued by the elegant network flow based solution to the escape problem, we wonder: How tightly are these two classes of problems intertwined and how we may take advantage of the relationship?

In this paper, we explore and exploit the connection between multi-agent path planning on *collision-free unit-distance graphs* (or CUGs, see Section 2 for the definition) and network flow. We begin by showing that multi-agent path planning on CUGs is closely related to a class of problems called *dynamic network flow* or *network flow over time*. We then focus on the permutation invariant multi-agent path planning problem on CUGs (by permutation invariant, we mean that goals are not pre-assigned to agents. Instead, we only require that each goal is reached by a unique agent), establishing that such problems always have solutions. To solve the problem algorithmically, an adapted maximum flow algorithm is provided which plans collision free paths for all agents with worst time complexity $O(nVE)$, in which n is the number of agents, V is the number of vertices of the CUG and E is the number of edges of the CUG. Moreover, we guarantee that the last agent takes time no more than $n + V - 1$ to reach its goal, assuming that agents travel at unit speed. Next, we construct efficient algorithms for obtaining temporally and spatially optimal solutions. For example, our algorithm for shortest overall time has running time $O(nVE \log V)$. We also show that these temporal and spatial objectives cannot be optimized simultaneously (i.e., they have a *Pareto optimal* structure).

As a universal subroutine in multi-agent systems, collision-free path planning for multiple agents finds applications in tasks spanning assembly [18, 31], evacuation [4, 37], formation control [2, 36, 39, 41, 43], localization [14], object transportation [29, 38], search and rescue [20], and so on. Given its importance, path planning for multi-agent systems has remained as a subject of intense study for many decades. Due to the vast size of the available literature, we only mention a most related subset of the research in this field and refer the readers to [5, 24, 26] and the references therein for a more comprehensive review of the subject.

When all agents are treated as a single agent with a high dimensional configuration space, the problem can be solved using cylindrical algebraic decomposition [6] or Canny's roadmap algorithm [3], in theory. Such *coupled* approaches suffer from

the curse of dimensionality; even when sampling based methods [22, 25] are used, instances involving only a small number of agents can be computationally challenging. This difficulty prompts the study of methods that seek to explore local features whenever possible to avoid working with too many agents at a time. Among these, *decoupled* planning is the most popular, which generally performs coordination of robot motion after deciding a path for each robot [17, 21, 32, 35, 40, 47]. In contrast, priority based methods force an order on agents to significantly reduce the search space [10, 45]. Some more recent works using decoupling heuristics include applying optimal decoupling techniques to exploit problem instances with low degrees of coupling [46], using *push-and-swap* primitives to avoid unnecessary exploration of search space [28], and heuristics aimed at performance guarantees (completeness is lost) [48].

Our algorithmic efforts in this paper focus on the *permutation invariant* multi-agent path planning problem on CUGs. Such formulations, in both discrete and continuous forms, are extensively studied as formation control problems [2, 36, 39, 41, 43], among others. On research that appears mostly related to this aspect of our paper, a discrete grid abstraction model for formation control was studied in [30]. To plan the paths, a three-step process was used in [30]: 1) Target assignment, 2) Path allocation, 3) Trajectory scheduling. Although it was shown that the process always terminates, no characterization of solution complexity was offered. In contrast, we provide very efficient algorithms that solve a strictly more general class of problems with optimality assurance. On the continuous side, a novel *formation space* approach was employed to represent the entire formation of robot teams with a single polynomial of which the roots correspond to the unassigned configurations for the robots in the formation [23].

We delay the literature review on network flow, from which we devise our time expansion construction for multi-agent path planning, to Section 3. The basic idea of applying time expansion to robotics problem is far from new [10, 34]. To the best of our knowledge, however, the research presented here is an original attempt at proposing a general time expansion technique, connecting it to network flow, and making full use of the benefits that come with this approach. We also note that our exact and complete algorithms all come with low constants in their respective worst case time complexity because they are derived from well studied fully combinatorial algorithms¹. Our simulation result, which we omit due to the length limit, confirms this assertion.

There are three main contributions. First, we formally establish the link between multi-agent path planning on graphs and network flow, showing how multi-agent path planning can be reduced to network flow problems, thereby enabling the potential application of powerful tools from combinatorial optimization to path planning for multiple agents in a principled way. Second, for the planning problem in which agents do not have pre-specified goals, we give fast and complete algorithms for finding collision free path sets that deliver every agent to a different goal. Third, we

¹ A *fully combinatorial algorithm* is an algorithm that only adds, subtracts, and compares values; no multiplication and division operations are allowed (i.e., ordered *group* operations versus ordered *field* operations) [16].

study time and distance optimality of the feasible solutions to the aforementioned problem, show that they have a pairwise Pareto optimal structure, and again provide efficient algorithms for optimizing two of these practical objectives.

The rest of the paper is organized as follows. In Section 2, we define two multi-agent path planning problems on CUGs. Section 3 starts with a quick review of network flow problems and then proceeds to show the reduction from multi-agent path planning on CUGs to network flow. Concentrating our efforts on the permutation invariant multi-agent path planning problem, Section 4 begins with a key construction that allows us to tightly bound the time steps required for a time-expanded network to have a feasible solution, which in turn enables efficient algorithms. Section 5 takes a further step and studies solution optimality on three natural objectives, showing the objectives have a Pareto optimal structure. We conclude in Section 6. Proof sketches of key theorems are provided; full proofs are omitted due to length².

2 Multi-agent Path Planning Problems on Collision-Free Unit-Distance Graphs

Let $G = (V, E)$ be a connected, undirected, simple graph (i.e., no multi-edges), in which $V = \{v_i\}$ is its vertex set and $E = \{(v_i, v_j)\}$ is its edge set. Let $A = \{a_1, \dots, a_n\}$ be a set of agents with initial and goal locations on G given by the injective maps $x_I : A \rightarrow V$ and $x_G : A \rightarrow V$, respectively. Note that A is essentially an index set; $x_I(A)$ and $x_G(A)$ are the set of initial and goal locations, respectively. We require that $x_I(A)$ and $x_G(A)$ be disjoint. For convenience, we let $n = |A|$ and use V, E to denote the cardinality of the sets V, E , respectively, since the meaning is usually clear from the context. Let σ be a bijection³ that acts on x_G , a *feasible path* for a single agent a_i is a map $p_i : \mathbb{Z}^+ \rightarrow V$ with the following properties⁴: 1. $p_i(0) = x_I(a_i)$. 2. For each i , there exists a smallest $k_{\min} \in \mathbb{Z}^+$ such that $p_i(k_{\min}) = (\sigma \circ x_G)(a_i)$ for some fixed σ . That is, the end point of the path p_i is some goal vertex. 3. For any $k \geq k_{\min}$, $p_i(k) \equiv (\sigma \circ x_G)(a_i)$. 4. For any $0 \leq k < k_{\min}$, $(p_i(k), p_i(k+1)) \in E$ or $p_i(k) = p_i(k+1)$. Intuitively, think of the domain of the paths as discrete time steps. We say that two paths p_i, p_j are in *collision* if there exists $k \in \mathbb{Z}^+$ such that $p_i(k) = p_j(k)$ (meet) or $(p_i(k), p_i(k+1)) = (p_j(k+1), p_j(k))$ (head-on). If $p(k) = p(k+1)$, the agent stays at vertex $p(k)$ during the time interval $[k, k+1]$.

As mentioned, in this paper, we work with a specific type of graph called the collision-free unit-distance graph (CUG): A CUG is a connected, undirected graph G satisfying the following: 1. Every edge is of unit length; 2. Given any two distinct edges (u_1, v_1) and (u_2, v_2) of G with $u_1 \neq u_2, v_1 \neq v_2$, two disc shapes (or spherical for 3D or more) agents of radius less than $\sqrt{2}/4$ traveling at unit speed through these edges (starting simultaneously at u_1, u_2 , respectively) will never collide. A

² A comprehensive extended version of this paper, including all proofs and extensions, is permanently available at <http://arxiv.org/abs/1204.5717>.

³ σ is introduced to unify the problem formulations; its use will become clear shortly. For now, the reader may think of it simply as the identity map.

⁴ $\mathbb{Z}^+ := \mathbb{N} \cup \{0\}$.

radius of $\sqrt{2}/4$ is the largest possible for two adjacent agents to travel along an “L” shaped path. One can easily verify that any graph with unit edge length and no acute angles between adjacent edges is a CUG. Hence, a connected 2D grid with holes is a CUG. Since subgraphs of 2D grids are easy to draw and visualize, we generally use subgraphs of 2D grids when we create examples in this paper. With the above setup, the *multi-agent path planning on CUGs* problem is defined as follows.

Problem 1. [Multi-agent Path Planning on CUGs] Given a 4-tuple (G, A, x_I, x_G) in which G is a CUG, find a set of paths $P = \{p_1, \dots, p_n\}$ such that p_i 's are feasible paths for respective agents a_i 's with σ being the identity map and no two paths p_i, p_j are in collision.

We require the graph to be a CUG so that it is suitable for multi-agent path planning. We formalize the rationale in the following observation.

Observation 1. Let p_i, p_j be two paths that are not in collision (as a partial solution to Problem 1). Then two disc shaped agents⁵ of radius less than $\sqrt{2}/4$, starting at the same time and moving along these respective paths with unit speed, will never collide.

Observation 1 shows that a solution to Problem 1 provides a path set for disc agents with radius $\sqrt{2}/4$ in A to reach their respective goals without a collision. It is easy to see that not all instances of this problem are solvable. Moreover, the decision version of Problem 1 (i.e., is there a solution that takes the agents to goals within K time steps) is NP-complete⁶. If we remove the assumption that all agents must reach their respective goals and allow permutation invariant paths (i.e., as long as each goal gets occupied by a unique agent in the end), Problem 1 becomes the *permutation invariant multi-agent path planning on CUGs* problem.

Problem 2. [Permutation Invariant Multi-agent Path Planning on CUGs] Given a 4-tuple (G, A, x_I, x_G) in which G is a CUG, find a set of paths $P = \{p_1, \dots, p_n\}$ such that p_i 's are feasible paths for respective agents a_i 's for an arbitrary (but fixed) permutation σ and no two paths p_i, p_j are in collision.

Problem 2 models the problem in which multiple identical or indistinguishable agents need to be deployed for serving requests at different locations (for example, formation control). This problem always has a solution: We simply plan and execute one path at a time and use more “remote” goal vertices earlier to avoid possible blocking of later paths; a formal result on the existence of such a choice of paths is given in Section 4.

⁵ Or spherical agents with radius less than $\sqrt{2}/4$, for dimensions higher than 2.

⁶ The lengthy proof is out of scope for the current paper. For curious readers, the NP-hardness proof is similar to that from [42]; the NP membership proof, which is non trivial, leads to efficient complete algorithms for solving Problem 1.

3 Multi-agent Path Planning on CUGs and Network Flow

3.1 Network Flow

In this subsection we give a brief review of network flow problems and algorithms pertinent to our problems. For surveys on network flow, see [1, 13]. We start with the classic static network flow problems.

Static Network Flow. A *network* $\mathcal{N} = (G, u, c, S)$ consists of a directed graph $G = (V, E)$ with $u, c : E \rightarrow \mathbb{Z}^+$ as the maps defining the capacities and costs on edges, respectively, and $S \subset V$ as the set of sources and sinks. We let $S = S^+ \cup S^-$, with S^+ denoting the set of sources and S^- denoting the set of sink vertices. Following conventions, for a vertex $v \in V$, let $\delta^+(v)$ (resp. $\delta^-(v)$) denote the set of edges of G going to (resp. leaving) v . A feasible (static) S^+, S^- -flow on this network \mathcal{N} is a map $f : E \rightarrow \mathbb{Z}^+$ that satisfies edge capacity constraints,

$$\forall e \in E, \quad f(e) \leq u(e), \quad (1)$$

the flow conservation constraints at non terminal vertices,

$$\forall v \in V \setminus S, \quad \sum_{e \in \delta^+(v)} f(e) - \sum_{e \in \delta^-(v)} f(e) = 0, \quad (2)$$

and the flow conservation constraints at terminal vertices,

$$\sum_{v \in S^+} \left(\sum_{e \in \delta^-(v)} f(e) - \sum_{e \in \delta^+(v)} f(e) \right) = \sum_{v \in S^-} \left(\sum_{e \in \delta^+(v)} f(e) - \sum_{e \in \delta^-(v)} f(e) \right). \quad (3)$$

The quantity on either side of (3) is called the *value* of the flow.

The classic (single-commodity) *maximum flow* problem asks the question: Given a network \mathcal{N} , what is the maximum value of flow that can be pushed through the network (i.e., seeking to maximize the value of the flow)? The *minimum cost maximum flow* problem further requires the flow to have minimum total cost among all maximum flows. That is, we want to find the flow among all maximum flows such that the quantity

$$\sum_{e \in E} c(e) \cdot f(e) \quad (4)$$

is minimized. Given integer inputs, integer maximum flow always exists, and many polynomial time algorithms exist for finding such a solution [9, 15]. The minimum cost maximum flow problem is equivalent to the *minimum cost circulation* problem, which is also solvable in polynomial time [44].

When additional structure is put on S , additional questions arise. If we limit the supply (resp. demand) of the source (resp. sink) vertices, we obtain a type of the flow problem called the *transshipment* problem. To formalize this, let $d : V \rightarrow \mathbb{Z}$ be the supplies on the vertices of G . Given a vertex $v \in V$, a positive $d(v)$ suggests that

the vertex has positive supply ($v \in S^+$) and a negative one suggests that the vertex has positive demand ($v \in S^-$). For all other vertices v , $d(v) = 0$. The basic version of the transshipment problem asks for a feasible flow through the network that also respects the supply/demand requirements

$$\begin{aligned} \forall v \in S^+, \quad & \sum_{e \in \delta^-(v)} f(e) - \sum_{e \in \delta^+(v)} f(e) = d(v), \\ \forall v \in S^-, \quad & \sum_{e \in \delta^+(v)} f(e) - \sum_{e \in \delta^-(v)} f(e) = d(v). \end{aligned} \tag{5}$$

The transshipment problem becomes the *evacuation* problem when $|S^-| = 1$ and the demand of the single sink vertex is equal to the total supply of the source vertices. The transshipment problem and the evacuation problem, as special cases of the maximum flow problem, can be solved with maximum flow algorithms mentioned above. If we instead require that vertices of S^+, S^- are paired up as $(s_1, s'_1), \dots, (s_k, s'_k)$ and that commodity of type i can be injected only into s_i and taken out at s'_i , we get the *multi-commodity flow* problem. Optimality questions as these from the single-commodity case can be asked here as well. Unlike in the single commodity case, finding integer maximum flow for multi-commodity problems is NP-hard in general and MAX SNP-hard (NP-hard to approximate below a certain multiple of optimal flow value) even for some simple restrictions [8].

Dynamic Network Flow. If we consider that flowing commodities through edges takes some time to complete, the problem becomes a *dynamic network flow* problem, which sometimes is also called *network flow over time*. There are two common variations of the dynamic network flow model: *Discrete time* and *continuous time*. In a *discrete time* model, flows enter and exit from vertices at integer time steps $t = 0, 1, \dots, T$. For a given edge $e = (u, v) \in E$, we may view the cost $c(e)$ as the time that is required to pass an amount of flow (not exceeding the capacity) from the tail u to the head v of the edge e . Therefore, we may interpret a (static) flow network \mathcal{N} as a dynamic one without any change of notations. In the closely related *continuous time* model, which we do not use in this paper, a *flow rate* is assigned to each edge, designating how fast a unit of flow can pass through the edge. The constraints imposed in the static network flow model generally apply to dynamic network flow models, except that dynamic network flow further requires that at any time, the flow passing through any edge cannot exceed the edge capacity.

Given a dynamic flow network, a question similar to the single-commodity maximum flow problem is the following: Starting at $t = 0$, what is the maximum units of flow the can reach the sinks on or before time $t = T$? It turns out that this problem can be solved using static flow algorithms such as Edmonds-Karp [9] over a *time-expanded network*. For example, given the dynamic flow network in Fig. 2(a), its time-expanded network with $T = 4$ is given in Fig. 2(b). To compute a flow over the time expanded network, we first add a *super source* and connect it using




Fig. 2 a) The (static) flow network with source s^+ and sink s^- . The numbers on the edges are the costs/time delay for passing through these edges. We may assume that the capacities are all unit capacities. b) The time-expanded network with 5 copies of the original vertices ($T = 4$). All edges have unit capacity. There is a forward edge between two vertices u and v at time steps t and t' , respectively (e.g. x at $t = 0$ and y at $t' = 1$), if one of the following is true: 1. $e = (u, v)$ is an edge of the static network with $c(e) = t' - t$ (the black edges, which retain the costs as $c(e)$'s); 2. u, v are the same vertex of the static network and $t' - t = 1$ (the green edges, which have unit costs). The green edges are also called *holdover* edges since traveling through a green edge is the same as the agent not actually moving.

outgoing edges to all copies of source vertices at $t = 0$, and add a *super sink* and connect all copies of sink vertices for all t to it using outgoing edges (the super source, super sink, and additional edges are not shown in Fig. 2(b)). With this construction, a static flow on the time-expanded network corresponds to a dynamic flow on the dynamic flow network.

Lemma 2. *For a sufficiently large T , a flow for a dynamic flow network \mathcal{N} is feasible if and only if the corresponding static flow on the time-expanded network of \mathcal{N} is feasible.*

A proof of Lemma 2 can be found in [12]. Note that determining a minimally sufficient T required by Lemma 2, which directly affects the running time of the resulting algorithm, is non-trivial. The standard maximum flow algorithms have time complexity depending polynomially on T and are therefore *pseudopolynomial* in general. For a special class of problems, the *quickest transshipment problem*, of which the goal is finding the quickest feasible flow for a transshipment problem over a dynamic network, a strongly polynomial time algorithm⁷ exists [19]. However, the algorithm requires calling subroutines (for example, submodular function optimization routines) that are not *combinatorial* algorithms and also has with large constant terms when it comes to asymptotic time complexity.

⁷ An algorithm is a strongly polynomial algorithm if: 1. The number of operations in the arithmetic model of computation is bounded by a polynomial in the number of integers in the input instance, and 2. The space used by the algorithm is bounded by a polynomial in the size of the input [16].

3.2 Equivalence between Multi-agent Path Planning on CUGs and Maximum Network Flow

In this subsection, we establish a reduction from the problems of our interest to multi-commodity network flow. For illustration purposes, we use the simple graph G in Fig. 3(a), with initial locations $\{s_i^+\}, i = 1, 2$ and goal locations $\{s_i^-\}, i = 1, 2$. An instance of Problem 1 is given by $(G, \{a_1, a_2\}, x_I : a_i \mapsto s_i^+, x_G : a_i \mapsto s_i^-)$. To apply maximum flow algorithms, we construct from G a time-expanded directed graph G' , part of which is shown in Fig. 3(c). We construct Fig. 3(c) as follows.

Since we cannot create an infinite time-expanded network, we need to specify the required number of time steps. For now assume that this number is some sufficiently large T (that is, if a flow with value F is achievable with an arbitrarily long time expansion, then F is also achievable with only T time steps). After fixing a T , we create $2T + 1$ copies of vertices from G , with indices $0, 1, 1', \dots$, as shown in Fig. 3(c). For each vertex $v \in G$, we denote these copies $v(0) = v(0)', v(1), v(1)', v(2), \dots, v(T)'$. For each edge $(u, v) \in G$ and time steps $t, t + 1, 0 \leq t < T$, we then add the gadget shown in Fig. 3(b) between $u(t)', v(t)'$ and $u(t + 1), v(t + 1)$ (arrows from the gadget are omitted from Fig. 3(c) since they are too small to draw). This gadget ensures that two agents cannot travel in opposite directions on an edge in the same time step. For the gadget, we assign unit capacity to all edges, unit cost to the horizontal middle edge, and zero cost to the other four edges. To finish the construction of Fig. 3(c), for each vertex $v \in G$, we add one edge between every two successive copies (i.e., we add the edges $(v(0), v(1)), (v(1), v(1)'), \dots, (v(T), v(T)')$). These correspond to the green and blue edges in Fig. 3(c). For all green edges, we assign them unit capacity and cost; for all blue edges, we assign them unit capacity and zero cost.




Fig. 3 a) A simple CUG G . b) A gadget for splitting an undirected edge through time steps. c) Part of the time-expanded network ($T = 2$).

The graph Fig. 3(c) is the main piece of G' , which is mostly done with the exception of the set S . We may simply let $S^+ = \{u(0) : u \in \{s_i^+\}\}$ and $S^- = \{v(T)' : v \in \{s_i^-\}\}$. That is, S^+ contains the first copies of the initial locations and S^- the last copies of the goal locations. The network $\mathcal{N}' = (G', u, c, S^+ \cup S^-)$ is now complete; we have reduced Problem 1 to an integer maximum multi-commodity flow problem on \mathcal{N}' with each agent from A as a single type of commodity.

Theorem 3. *Given an instance of Problem 1 with input parameters (G, A, x_I, x_G) , there is a bijection between its solutions (with maximum number of time steps up to T) and the integer maximum multi-commodity flow solutions of flow value n on the time-expanded network \mathcal{N}' constructed from (G, A, x_I, x_G) with T time steps.*

PROOF SKETCH. Fixing a T , one can establish a bijection between solutions on the original network and solutions on the time-expanded network with T time steps. \square

Since integer maximum multi-commodity flow is NP-hard, the above construction does not directly offer an efficient solution to Problem 1. Moving to Problem 2, allowing an arbitrary permutation σ to act on x_G means that we may treat all agents as a single type of commodity. Theorem 3 then implies that Problem 2 is equivalent to the quickest transshipment problem, which is solvable in polynomial time using subroutines for optimizing submodular functions. In the next section, we show that we can do better by bounding the required time steps for finding a feasible solution to Problem 2 and then apply more standard combinatorial algorithms for network flow to solve it.

4 Efficient Combinatorial Algorithms for Permutation Invariant Multi-agent Path Planning on CUGs

If we choose to apply combinatorial network flow algorithms over the time-expanded network to find solutions to Problem 2, the first priority is to determine the required number of time steps necessary to find a solution; otherwise we cannot declare that the algorithm is complete. We now provide a tight bound on T . Let (G, A, x_I, x_G) be an instance of Problem 2. We first prove some intermediate results on path sets over G . To distinguish these paths from the solution path set, denote them as $Q = \{q_1, \dots, q_n\}$. For convenience, $head(q_i)$, $tail(q_i)$, and $len(q_i)$ denote the start vertex, end vertex, and length of q_i , respectively. With a slight abuse of notation, $V(\cdot)$, $E(\cdot)$ denote the vertex set and edge set of the input parameter, which can be either a path, q_i , or a set of paths, such as Q . To start off, we want a path set Q with the following properties:

Property 1. For all $1 \leq i \leq n$, $head(q_i) \in x_I(A)$ and $tail(q_i) \in x_G(A)$. For any two paths q_i, q_j , $head(q_i) \neq head(q_j)$ and $tail(q_i) \neq tail(q_j)$.

Property 2. Each path q_i is a shortest path between $head(q_i)$ and $tail(q_i)$ on G .

Property 3. The total length of the path set Q is minimal.

Property 4. If we orient the edges of every path $q_i \in Q$ from $head(q_i)$ to $tail(q_i)$, no two paths share a common edge oriented in different directions.

Lemma 4. *There exists a set of paths $Q = \{q_1, \dots, q_n\}$ that satisfies Properties 1-4.*

PROOF SKETCH. Properties 1 and 2 are merely restrictions to have the initial and goal vertices paired up using shortest paths. Since we work with a discrete problem, a path set satisfying Property 3 always exists. Property 4 is implied by Property 3:

If on the contrary that two edges from two paths are oriented differently, switching destinations on those paths will reduce the total path length by two. \square

The technique from the proof of Lemma 4 can be generalized to show that oriented paths cannot form a directed cycle.

Proposition 5. *A path set Q that satisfies Properties 1-3 induces a directed acyclic graph (DAG) structure on $E(Q)$.*

A *standalone goal vertex* is a vertex $v \in x_G(A)$ such that there is a single path $q \in Q$ containing v .

Corollary 6. *A path set Q that satisfies Properties 1-3 has a standalone goal vertex.*

PROOF SKETCH. The absence of standalone vertices implies the existence of a directed cycle, contradicting Proposition 5. \square

The existence of a standalone goal vertex allows the construction of a path set which decomposes into paths that can be sequentially scheduled without colliding into each other. We characterize such a path set as one with an additional property.

Lemma 7. *There exists a path set Q satisfying Properties 1-4 and the following additional property:*

Property 5. Let $Q_i := \{q_i, \dots, q_n\}$. For any $1 \leq i \leq n$, restricting to Q_i , among all possible paths connecting an initial location (of Q_i) to a standalone goal location (of Q_i) using oriented edges from $E(Q_i)$, q_i is one shortest such.

PROOF SKETCH. Using an arbitrary shortest path set Q_0 , one can iteratively construct a new path set Q using edges of $E(Q_0)$, starting with a “furthest” standalone vertex, that satisfies Property 5. \square

If we schedule agents using a path set Q satisfying properties 1-5, there can never be cases where two agents block each other, as a direct consequence of Lemma 4. There is still the possibility that one agent blocks another. The following theorem shows that such blocking can be minimized.

Theorem 8. *Given an instance of Problem 2 with input parameters (G, A, x_I, x_G) and let ℓ be the largest pairwise distance between a member of $x_I(A)$ and a member of $x_G(A)$,*

$$\ell = \max_{\forall u \in x_I(A), v \in x_G(A)} \text{dist}(u, v). \quad (6)$$

A time-expanded network \mathcal{N}' with $T = n + \ell - 1$ is necessary and sufficient for a feasible solution to Problem 2 to exist.

PROOF SKETCH. Starting with a path set Q satisfying Properties 1-5 and letting path q_i start at time step $t = i - 1$, one can show that no collision can occur. \square

Since ℓ cannot be larger than V , the number of vertices of G , the following corollary is immediate.

Corollary 9. *For every instance of Problem 2, a feasible solution exists.*

In particular, the construction in the proof of Lemma 4 yields a complete (may not be efficient) algorithm for Problem 2. In addition to confirming that any maximum flow algorithm over the time expanded network \mathcal{N}' with $T = n + \ell - 1$ is also complete algorithm, Theorem 8 enables us to show that such algorithms are efficient.

Theorem 10. *Problem 2 is solvable using a combinatorial algorithm in strongly polynomial time.*

Using the Ford-Fulkerson algorithm [11], the time complexity is $O(nVE)$. In practice, even better running times are possible. If G is a planar graph, we have $E \sim O(V)$ and $\ell \sim O(V^{\frac{1}{2}})$. The time complexity then becomes $O(\text{MF}(n, V(n + V^{\frac{1}{2}} - 2), V(n + V^{\frac{1}{2}} - 2))) \sim O(\text{MF}(n, V(n + V^{\frac{1}{2}}), V(n + V^{\frac{1}{2}})))$. Since in our case $n < V$, Ford-Fulkerson gives us the running time $O(nV(n + V^{\frac{1}{2}})) = O(n^2V + nV^{\frac{3}{2}})$.

5 Optimal Solutions

In this section, we present optimal solutions for the permutation invariant multi-agent path planning problem. After introducing several temporal and spatial objectives of practical importance, we apply techniques from network flow to obtain optimal solutions for two of these objectives. Since these objectives are different from the basic version of Problem 2, the time bound T may be different. Lastly, we show that these objectives possess a Pareto optimal structure and they cannot be optimized simultaneously.

5.1 Optimizing over the Feasible Solutions

Having found *feasible* solutions to Problem 2, we turn the focus to the *optimality* of these solutions for practical purposes. As mentioned in Section 2, we intend to use the formulation as a model for scenarios such as multi-robot servicing. For many applications, time optimality is a top priority. Optimizing over the feasible solutions to Problem 2 (that is, we require that all goals are reached), there are two natural criteria for measuring time optimality:

Objective 1. Minimizing the average time it takes for all agents to reach their goals.

Objective 2. Minimizing the time it takes for the last agent to reach its goal.

In terms of agents (robots or people) serving requests, Objective 1 seeks to minimize the average time before a request gets served. The sufficient condition on the time bound T for this objective is given below.

Theorem 11. *There exists an optimal solution for Objective 1 in a time-expanded network with $T = (n - 1)(n - 2)/2 + V$.*

The second objective, minimizing the time that its last goal is reached, provides a lower bound on the time that is required to reach all goals. Solutions optimizing this objective are useful in providing worst servicing time estimate or guarantee. Solutions to the quickest transshipment problem [19] yield optimal solutions to this objective. However, we can avoid using submodular function optimization routines if we have a polynomial bound on T , which is provided in the following corollary of Theorem 8.

Corollary 12. *There exists an optimal solution for Objective 2 in a time-expanded network with $T = n + \ell - 1$.*

To see that Corollary 12 is true, note that $T = n + \ell - 1$ is sufficient for finding a feasible solution, which must have completion time as large as that of a solution to Objective 2. With the bound on T , running $\log T$ rounds (via binary search) of maximum flow over time-expanded network with different time horizon then gives us an optimal solution to Objective 2. The running time is then bounded by $O(\text{MF}(n, V^2, VE) \log V)$, which is strongly polynomial. In particular, with Ford-Fulkerson, the running time becomes $O(nVE \log V)$. After time optimality, another very useful solution property to optimize is the total distance traveled by the agents, i.e., spatial optimality:

Objective 3. Minimizing the total distance traveled by the agents on G .

Because we work with a CUG, if an agent a_i actually moves along path p_i between time steps t and $t + 1$, $p_i(t)$ must be different from $p_i(t + 1)$. These correspond to the black edges in the time-expanded network (see Fig. 3(b)). Thus, to optimize this objective, we can find the shortest total distance traveled by all agents via setting the cost of the holdover edge (green edges in Fig. 3(b)) to zero and then running minimum cost maximum flow algorithm over the time-expanded network. The method is again strongly polynomial, with complexity $O(V^2E \log V)$, due to the following corollary.

Corollary 13. *There exists an optimal solution for Objective 3 in a time-expanded network with $T = n + \ell - 1$.*

5.2 Pareto Optimality between the Objectives

From the discussion in the previous subsection, we observe that each of the three objectives is of practical importance. At this point, one might be tempted to seek solutions that optimize multiples of these objectives simultaneously. We show that




Fig. 4 a) Overlaying grid on a workspace with obstacles. b) Adapting a roadmap to obtain a graph that can be used with our multi-agent path planning algorithms.

this is not possible for each pair of these objectives. In the following theorem, we say that two objectives are *compatible* if and only if they can be optimized simultaneously. Otherwise, we say the objectives are *incompatible*.

Theorem 14. *Over the feasible solutions to Problem 2, Objectives 1-3 are pairwise incompatible.*

6 Conclusion, Future Work, and Open Problems

In this paper, we established the close link between two classes of problems: Multi-agent path planning on CUGs and network flow. Focusing on the permutation invariant versions of the multi-agent path planning problem, we proved a tight bound on the number of time steps necessary and sufficient for a feasible path set to exist in the time-expanded network, enabling efficient algorithmic solutions to these problems. We then explored optimality issues, demonstrating that the time-expansion bound generally carry over to yield strongly polynomial algorithms for optimizing two of these practical objective functions. Interestingly, each pair of these objectives cannot be optimized simultaneously.

Given our study, an immediate question or criticism is the applicability of the results to problems beyond CUGs. After all, real agents, whether robots or people, do not always live on a discrete graph. To answer this question, we have research under way that explores the idea of overlaying the CUGs on the actual workspace. That is, we may first create a roadmap over the workspace that captures the connectivity and then discretize the roadmap over which the statement of Observation 1 continues to hold (as long as the edges are close to unit length the angle between two edges is obtuse, similar version of Observation 1 can be stated) [33]. A basic solution (Fig. 4(a)) may be to put a grid on the roadmap and delete vertices inside or close to obstacles. To overcome the issue of the inherited Manhattan metric of grids, we may adapt the grid to align with the geodesics of the environment. For example, for a two dimensional workspace with polygonal obstacles, we can arrange the grid edges to follow edges of the visibility graph [27] of the environment when possible.

When clearance is tight, we may start with a maximum clearance roadmap [33] and add the vertices carefully (see Fig. 4(b)). Note that because the workspace is often two dimensional, these preparations can be computed relatively efficiently.

Many interesting open problems remain. Although finding a distance optimal solution to Problem 1 using a time-expanded network is impractical due to its intrinsic hardness, the network flow approach might still produce efficient methods that yield basic feasible solutions since the time-expanded network has a forward only structure. In addition, approximation algorithms on integer multi-commodity flow could lead to better heuristics for optimal solution search. Along this line, we only touched the most essential results in the field of network flow, which are but the tips of an iceberg. It would not be surprising that results from the vast amount of network flow literature could be readily carried over to tackle path planning problems, either along the structure we proposed in this paper or in some other forms. As an example, for Problem 2, since Objectives 1-3 are all of practical concerns but incompatible, it is desirable to seek solutions that provide performance guarantees on each of these objectives. Network flow methods, closely relate to linear programming, appear to be promising tools for such parametric optimization tasks.

Acknowledgments. The authors thank Max Katsev for double checking the proofs and the anonymous reviewers for their constructive suggestions. This work was supported in part by NSF grants 0904501 (IIS Robotics) and 1035345 (Cyberphysical Systems), DARPA STOMP grant HR0011-05-1-0008, and MURI/ONR grant N00014-09-1-1052.

References

1. Aronson, J.E.: A survey on dynamic network flows. *Annals of Operations Research* 20(1), 1–66 (1989)
2. Balch, T., Arkin, R.C.: Behavior-based formation control for multirobot teams. *IEEE Transaction on Robotics and Automation* 14(6), 926–939 (1998)
3. Canny, J.F.: *The Complexity of Robot Motion Planning*. MIT Press, Cambridge (1988)
4. Chalmet, L.G., Francis, R.L., Saunders, P.B.: Network models for building evacuation. *Management Science* 28(1), 86–105 (1982)
5. Choset, H., Lynch, K.M., Hutchinson, S., Kantor, G., Burgard, W., Kavraki, L.E., Thrun, S.: *Principles of Robot Motion: Theory, Algorithms, and Implementations*. MIT Press, Cambridge (2005)
6. Collins, G.E.: Quantifier Elimination for Real Closed Fields by Cylindrical Algebraic Decomposition. In: Brakhage, H. (ed.) *GI-Fachtagung 1975*. LNCS, vol. 33, pp. 134–183. Springer, Heidelberg (1975)
7. Cormen, T.H., Leiserson, C.E., Rivest, R.L., Stein, C.: *Introduction to Algorithms*, 2nd edn. MIT Press, Cambridge (2001)
8. Costa, M.-C., Létocart, L., Roupin, F.: Minimal Multicut and Maximal Integer Multiflow: A Survey. *European Journal of Operational Research* 162, 55–69 (2005)
9. Edmonds, J., Karp, R.M.: Theoretical improvements in algorithmic efficiency for network flow problems. *J. ACM* 19(2), 248–264 (1972)
10. Erdmann, M.A., Lozano-Pérez, T.: On multiple moving objects. In: *Proceedings IEEE International Conference on Robotics & Automation*, pp. 1419–1424 (1986)

11. Ford, L.R., Fulkerson, D.R.: Maximal flow through a network. Research Memorandum RM-1400, The RAND Corporation (November 1954)
12. Ford, L.R., Fulkerson, D.R.: Constructing maximal dynamic flows from static flows. *Operations Research* 6, 419–433 (1958)
13. Ford, L.R., Fulkerson, D.R.: *Flows in Networks*. Princeton University Press, New Jersey (1962)
14. Fox, D., Burgard, W., Kruppa, H., Thrun, S.: A probabilistic approach to collaborative multi-robot localization. *Autom. Robots* 8(3), 325–344 (2000)
15. Goldberg, A.V., Tarjan, R.E.: A new approach to the maximum flow problem. In: *STOC 1986: Proceedings of the Eighteenth Annual ACM Symposium on Theory of Computing*, pp. 136–146. ACM, New York (1986)
16. Grötschel, M., Schrijver, A., Lovász, L.: *Complexity, Oracles, and Numerical Computation*. Springer (1988)
17. Guo, Y., Parker, L.E.: A distributed and optimal motion planning approach for multiple mobile robots. In: *Proceedings IEEE International Conference on Robotics and Automation*, pp. 2612–2619 (2002)
18. Halperin, D., Latombe, J.-C., Wilson, R.: A general framework for assembly planning: The motion space approach. *Algorithmica* 26(3-4), 577–601 (2000)
19. Hoppe, B., Tardos, É.: The quickest transshipment problem. *Mathematics of Operations Research* 25(1), 36–62 (2000)
20. Jennings, J.S., Whelan, G., Evans, W.F.: Cooperative search and rescue with a team of mobile robots. In: *Proceedings IEEE International Conference on Robotics & Automation* (1997)
21. Kant, K., Zucker, S.: Towards efficient trajectory planning: The path velocity decomposition. *International Journal of Robotics Research* 5(3), 72–89 (1986)
22. Kavraki, L.E., Svestka, P., Latombe, J.-C., Overmars, M.H.: Probabilistic roadmaps for path planning in high-dimensional configuration spaces. *IEEE Transactions on Robotics & Automation* 12(4), 566–580 (1996)
23. Kloder, S., Hutchinson, S.: Path planning for permutation-invariant multirobot formations. *IEEE Transactions on Robotics* 22(4), 650–665 (2006)
24. Latombe, J.-C.: *Robot Motion Planning*. Kluwer, Boston (1991)
25. LaValle, S.M.: Rapidly-exploring random trees: A new tool for path planning. TR 98-11, Computer Science Dept., Iowa State University (October 1998)
26. LaValle, S.M.: *Planning Algorithms*. Cambridge University Press, Cambridge (2006), <http://planning.cs.uiuc.edu/>
27. Lozano-Pérez, T., Wesley, M.A.: An algorithm for planning collision-free paths among polyhedral obstacles. *Communications of the ACM* 22(10), 560–570 (1979)
28. Luna, R., Bekris, K.E.: Push and swap: Fast cooperative path-finding with completeness guarantees. In: *Twenty-Second International Joint Conference on Artificial Intelligence*, pp. 294–300 (2011)
29. Mataric, M.J., Nilsson, M., Simsarian, K.T.: Cooperative multi-robot box pushing. In: *Proceedings IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 556–561 (1995)
30. Miklic, D., Bogdan, S., Fierro, R., Nestic, S.: A discrete grid abstraction for formation control in the presence of obstacles. In: *Proceedings IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 3750–3755 (2009)
31. Njaji, B.: *Theory of Automatic Robot Assembly and Programming*. Chapman and Hall (1992)
32. O’Donnell, P.A., Lozano-Pérez, T.: Deadlock-free and collision-free coordination of two robot manipulators. In: *Proceedings IEEE International Conference on Robotics & Automation*, pp. 484–489 (1989)

33. O'Dúnlaing, C., Yap, C.K.: A retraction method for planning the motion of a disc. *Journal of Algorithms* 6, 104–111 (1982)
34. Peasgood, M., Clark, C., McPhee, J.: A complete and scalable strategy for coordinating multiple robots within roadmaps. *IEEE Transactions on Robotics* 24(2), 283–292 (2008)
35. Peng, J., Akella, S.: Coordinating Multiple Robots with Kinodynamic Constraints along Specified Paths. In: Boissonat, J.-D., Burdick, J., Goldberg, K., Hutchinson, S. (eds.) *Algorithmic Foundations of Robotics V. STAR*, vol. 7, pp. 221–237. Springer, Heidelberg (2004)
36. Poduri, S., Sukhatme, G.S.: Constrained coverage for mobile sensor networks. In: *Proceedings IEEE International Conference on Robotics & Automation* (2004)
37. Rodriguez, S., Amato, N.M.: Behavior-based evacuation planning. In: *Proceedings IEEE International Conference on Robotics and Automation*, pp. 350–355 (2010)
38. Rus, D., Donald, B., Jennings, J.: Moving furniture with teams of autonomous robots. In: *Proceedings IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 235–242 (1995)
39. Shucker, B., Murphey, T., Bennett, J.K.: Switching rules for decentralized control with simple control laws. In: *American Control Conference* (July 2007)
40. Siméon, T., Leroy, S., Laumond, J.-P.: Path coordination for multiple mobile robots: A resolution complete algorithm. *IEEE Transactions on Robotics & Automation* 18(1) (February 2002)
41. Smith, B., Egerstedt, M., Howard, A.: Automatic generation of persistent formations for multi-agent networks under range constraints. *ACM/Springer Mobile Networks and Applications Journal* 14(3), 322–335 (2009)
42. Surynek, P.: An optimization variant of multi-robot path planning is intractable. In: *The Twenty-Fourth AAAI Conference on Artificial Intelligence (AAAI 2010)*, pp. 1261–1263 (2010)
43. Tanner, H., Pappas, G., Kumar, V.: Leader-to-formation stability. *IEEE Transactions on Robotics and Automation* 20(3), 443–455 (2004)
44. Tardos, É.: A strongly polynomial minimum cost circulation algorithm. *Combinatorica* 5(3), 247–255 (1985)
45. van den Berg, J., Overmars, M.: Prioritized motion planning for multiple robots. In: *Proceedings IEEE/RSJ International Conference on Intelligent Robots and Systems* (2005)
46. van den Berg, J., Snoeyink, J., Lin, M., Manocha, D.: Centralized path planning for multiple robots: Optimal decoupling into sequential plans. In: *Proceedings Robotics: Science and Systems* (2009)
47. Švestka, P., Overmars, M.H.: Coordinated path planning for multiple robots. *Robotics and Autonomous Systems* 23, 125–152 (1998)
48. Wang, K.-H.C., Botea, A.: Tractable multi-agent path planning on grid maps. In: *Proceedings of the 21st International Joint Conference on Artificial Intelligence, IJCAI 2009*, pp. 1870–1875. Morgan Kaufmann Publishers Inc., San Francisco (2009)