

A^* for Bounding Shortest Paths in the Graphs of Convex Sets

Kaarthik Sundar ¹, Sivakumar Rathinam ²

¹Los Alamos Research Laboratory, New Mexico, 88220, USA

²Texas A&M University, College Station, TX 77843, USA

kaarthik@lanl.gov, srathinam@tamu.edu

Abstract

We present a novel algorithm that fuses the existing convex-programming based approach with heuristic information to find optimality guarantees for the Shortest Path Problem in the Graph of Convex Sets (SPP-GCS). Our method, inspired by A^* , initiates a best-first-like procedure from a designated subset of vertices and iteratively expands it until further growth is neither possible nor beneficial. Traditionally, obtaining solutions with bounds for an optimization problem involves solving a relaxation, modifying the relaxed solution to a feasible one, and then comparing the two solutions to establish bounds. However, for SPP-GCS, we demonstrate that reversing this process can be more advantageous, especially with Euclidean travel costs. We present numerical results to highlight the advantages of our algorithm over the existing approach in terms of the sizes of the convex programs solved and computation time.

Introduction

The Shortest Path Problem (SPP) is one of the most important and fundamental problems in discrete optimization (Lawler 1979; Korte and Vygen 2018). Given a graph, the SPP aims to find a path between two vertices in the graph such that the sum of the cost of the edges in the path is minimized. In this paper, we concern ourselves with a generalization of the SPP, recently introduced in (Marcucci et al. 2024b), where each vertex is associated with a convex set and the cost of the edge joining any two vertices depends on the choice of the points selected from each of the respective convex sets. In this generalization, referred to as the *Shortest Path Problem in the Graph of Convex Sets* (SPP-GCS), the objective is to find a path and choose a point from each convex set associated with the vertices in the path such that the sum of the cost of the edges in the path is minimized (Fig. 1).

SPP-GCS reduces to the standard SPP if the point to be selected from each convex set is given. Also, SPP-GCS reduces to a relatively easy-to-solve convex optimization problem if the path is given. Selecting an optimal point in each set, as well as determining the path, makes the SPP-GCS NP-Hard in the general case (Marcucci et al. 2024b).

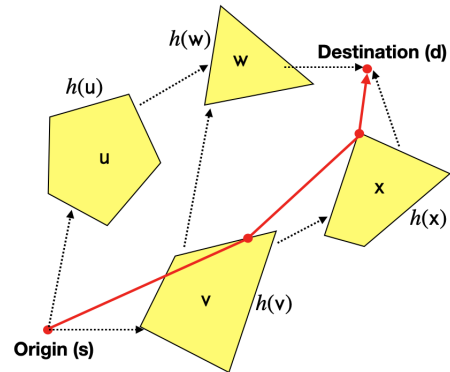


Figure 1: Illustration of the SPP-GCS. The graph has six vertices, with the origin and the destination corresponding to singleton sets. The dotted lines represent the edges in the graph. Each vertex t in the graph is associated with a heuristic value $h(t)$. A feasible path is shown with solid (red) line segments.

We are interested in SPP-GCS because it naturally models SPPs in the geometric domain and planning problems with neighborhoods, with a variety of applications in motion planning (Marcucci et al. 2023, 2024a; Natarajan et al. 2024), sensor coverage (Forney, Burdick, and Rimon 2024; Dumitrescu and Mitchell 2003), and hybrid control (Marcucci et al. 2024b). For example, the SPP in the presence of obstacles in 3D can be posed as an SPP-GCS by decomposing the free space into convex sets and then planning on the corresponding graph of these sets. The sensor coverage problem in 2D (Forney, Burdick, and Rimon 2024), where the robot needs to find a suitable sequence of neighborhoods to traverse while choosing a point from each neighborhood, can also be formulated as a variant of SPP-GCS.

The importance of formulating a planning problem as an SPP-GCS stems primarily from the following observations: finding optimality guarantees for SPPs in a geometric domain or for planning problems with neighborhoods has been computationally challenging, even in special cases. For example, computing the optimum or a tight lower bound for an SPP in 3D with convex or axis-aligned obstacles remains difficult (Mitchell 2017). While theoretical bounds

using polynomial-time approximation schemes are available (Mitchell 2017), we are not aware of any implementations of such schemes with numerical results. Recently, significant progress (Marcucci et al. 2024b) has been made towards addressing this issue where the planning problem is posed as a SPP-GCS and bounds are obtained by solving a tight convex relaxation of the SPP-GCS. Feasible solutions can also be found by either rounding the relaxed solutions or using a sampling-based heuristic (as employed in this paper). Optimality guarantees can then be obtained by comparing the convex relaxation cost (a lower bound) with the length (an upper bound) of a feasible solution. Since the crux of the procedure in (Marcucci et al. 2024b) lies in solving a convex relaxation of the SPP-GCS, this paper is motivated by the need to explore faster methods for solving this relaxation, with the ultimate goal of generating optimality guarantees for planning problems in geometric domains.

One way to develop fast algorithms for an optimization problem is to incorporate heuristic information into the search process. The celebrated A^* algorithm (Hart, Nilsson, and Raphael 1968) accomplishes this for the SPP. The objective of this paper is to propose a new algorithm that combines the existing convex-programming based approach in (Marcucci et al. 2024b) with heuristic information to find lower bounds and optimality guarantees for the SPP-GCS.

We note that there is a parallel development (Chia et al. 2024) that aims to incorporate heuristic information into a best-first search process to find optimal solutions for the SPP-GCS. In this process, sub-optimal solutions are pruned based on a domination criterion, which is non-trivial to check for general convex sets and travel costs. However, when the sets are polytopical and the travel costs are linear, algorithms (Chia et al. 2024) are available to conservatively prune sub-optimal solutions. This search process and its variants seeks to converge to the optimum from above, which is complementary to the lower-bounding approach we pursue in this paper.

We refer to our algorithm as A^* for Graphs of Convex Sets (A^* -GCS). Just like in A^* , we assume that each vertex (or each convex set) has an associated heuristic cost to go from the vertex to the destination (Fig. 1). A^* -GCS follows the spirit of A^* and initiates a best-first-like search procedure from a designated set S (containing the origin) and iteratively builds on it until further growth is neither possible nor beneficial. A key step in A^* -GCS lies in the way we grow S in each iteration. To choose vertices for addition to S , we employ a relaxed solution to a generalization of the SPP-GCS that encompasses the vertices in S and their neighbors. In the special case when the convex sets reduce to singletons and the heuristic information is consistent, A^* -GCS reduces to A^* . On the other hand, if heuristic information is not available, A^* -GCS reduces to iteratively solving a convex relaxation of SPP-GCS until the termination conditions are satisfied.

Unlike A^* , which typically starts its search from a closed set S containing only the origin, A^* -GCS can start its search from any set S as long as it induces a cut — meaning S contains the origin but not the destination. In this paper, we consider two initial choices for S to start the search: one where

S only contains the origin, and another where S contains vertices informed by applying A^* to a representative point (e.g. the centroid) from each convex set. One of the key insights we infer from the numerical results with Euclidean travel costs is that the second choice where A^* -GCS starts with the set S informed by A^* provides the best trade-off between achieving good quality bounds and smaller computation times. In general, A^* -GCS works on relatively smaller convex programs because its vertices are informed by the A^* algorithm, which only explores a subset of vertices in the graph. This leads to reduced computation times compared to solving a convex relaxation on the entire graph.

Typically, to obtain optimality guarantees for an optimization problem, we solve a relaxation of the problem, modify the relaxed solution to obtain a feasible solution, and then compare the two solutions to establish bounds. This is the approach followed in (Marcucci et al. 2024b). *In this paper, we demonstrate that reversing this process can be more effective in obtaining bounds for the SPP-GCS, especially with Euclidean travel costs. Specifically, we first implement A^* using a representative point from each convex set to obtain a feasible solution. We then solve a convex program on subsets of vertices informed by A^* to find a relaxed solution and subsequently compare the solutions.* In hindsight, this approach seems natural; since A^* quickly finds good feasible solutions for the SPP-GCS with a heuristic, it is reasonable to expect that good lower bounds can also be found in the neighborhood of the vertices visited by A^* , based on the strong formulation in (Marcucci et al. 2024b).

After presenting our algorithm with theoretical guarantees, we provide extensive computational results on the performance of A^* -GCS for instances derived from mazes, axis-aligned bars and 3D villages to highlight the advantages of our algorithm over the existing approach in terms of the sizes of the convex programs solved and computation time.

Problem Statement

Let V denote a set of vertices, and E represent a set of directed edges joining vertices in V . Let each vertex $v \in V$ be associated with a non-empty, compact convex set $\mathcal{X}_v \subset \mathbb{R}^n$. Given vertices $u, v \in V$, the cost of traveling the edge $e = (u, v)$ from u to v depends on the choice of the points in the sets \mathcal{X}_u and \mathcal{X}_v . Specifically, given points $x_u \in \mathcal{X}_u$ and $x_v \in \mathcal{X}_v$, let $cost(x_u, x_v)$ denote the travel cost of edge $e = (u, v)$. In this paper, we set $cost(x_u, x_v)$ to be equal to the Euclidean distance between the points x_u and x_v , i.e., $cost(x_u, x_v) := \|x_u - x_v\|_2$.

Any path in the graph (V, E) is a sequence of vertices (v_1, v_2, \dots, v_k) for some positive integer k such that $v_i \in V$, $i = 1, \dots, k$ and $(v_i, v_{i+1}) \in E$, $i = 1, \dots, k - 1$. Given a path $\mathcal{P} := (v_1, v_2, \dots, v_k)$ and points $x_{v_i} \in \mathcal{X}_{v_i}$, $i = 1, \dots, k$, the cost of traveling \mathcal{P} is defined as $\sum_{i=1}^{k-1} cost(x_{v_i}, x_{v_{i+1}})$. Given an origin $s \in V$ and destination $d \in V$, the objective of the SPP-GCS is to find a path $\mathcal{P} := (v_1, v_2, \dots, v_k)$ and points $x_{v_i} \in \mathcal{X}_{v_i}$, $i = 1, \dots, k$, such that $v_1 = s$, $v_k = d$, and the cost of traveling \mathcal{P} is minimized. The optimal cost for SPP-GCS is denoted as $C_{opt}(s, d)$.




Figure 2: Setup in the SPP*-GCS showing the cut-set S , and subsets N_S and S' . N_S contains all the vertices in $V \setminus S$ that are adjacent to S , and S' is any nonempty subset of N_S . A feasible path for SPP*-GCS is shown with solid (red) line segments.

Generalization of SPP-GCS and its Relaxation

We first introduce some notations specific for A^* -GCS. For each vertex $v \in V$, let $h(v)$ denote an underestimate to the optimal cost for the SPP-GCS from v to d , *i.e.*, $h(v) \leq C_{opt}(v, d)$. We refer to such underestimates as admissible (just like in A^*). Note that for the destination d , the underestimate $h(d) = 0$. We also refer to $h(\cdot)$ as a heuristic function for A^* -GCS. Given any $S \subset V$, let N_S denote the set of all the vertices in $V \setminus S$ that are adjacent to at least one vertex in S , *i.e.* $N_S := \{v : u \in S, v \notin S, (u, v) \in E\}$. N_S is also referred to as the neighborhood of S .

Now, consider a closely related problem to SPP-GCS referred to as SPP*-GCS defined on non-empty subsets $S \subset V$ and $S' \subseteq N_S$ such that $s \in S$ and $d \notin S$. We will refer to any set $S \subset V$ that satisfies $s \in S$ and $d \notin S$ as a *cut-set* in this paper (refer to Fig. 2). We will restrict our attention to the vertices in $\bar{V} := S \cup S'$ and the edges in $\bar{E} := \{(u, v) : u \in S, v \in \bar{V}, (u, v) \in E\}$. The objective of SPP*-GCS is to find a terminal vertex v_l , a path $\mathcal{P}_g := (v_1, v_2, \dots, v_l)$ in the graph (\bar{V}, \bar{E}) and points $x_{v_i} \in \mathcal{X}_{v_i}$, $i = 1, \dots, l$ such that $v_1 = s$, $v_l \in S'$, and the sum of the cost of traveling \mathcal{P}_g and the heuristic cost, $h(v_l)$, is minimized. The optimal cost for SPP*-GCS is denoted as $C_{opt}^*(S, S')$. The following Lemma shows the relationship between SPP*-GCS and SPP-GCS.

Lemma 1. *SPP*-GCS is a generalization of SPP-GCS.*

Proof. Without loss of generality, we assume there is at least one feasible path from s to d in the graph (V, E) . Therefore, if $S := V \setminus \{d\}$, then $N_S := \{d\}$. Hence, $C_{opt}^*(V \setminus \{d\}, \{d\}) = C_{opt}(s, d)$ as $h(d) = 0$. \square

Non-Linear Program for SPP*-GCS

To formulate SPP*-GCS, we use two sets of binary variables to specify the path and a set of continuous variables to choose the points from the convex sets corresponding to the vertices in the path. The first set of binary variables, denoted

by y_{uv} for $(u, v) \in \bar{E}$, determines whether the edge (u, v) is selected in a solution to the SPP*-GCS. Here, $y_{uv} = 1$ implies that the edge (u, v) is chosen, while $y_{uv} = 0$ implies the opposite. The second set of variables, denoted by α_v for $v \in S'$, determines whether the vertex v is selected as a terminal or not. The continuous variable x_u for $u \in \bar{V}$ specifies the point chosen in the convex set corresponding to vertex u . For any vertex $u \in \bar{V}$, let $\delta^+(u)$ be the subset of all the edges in \bar{E} leaving u , and let $\delta^-(u)$ be the subset of all the edges in \bar{E} coming into u . The non-linear program for SPP*-GCS is as follows:

$$C_{opt}^*(S, S') := \min \sum_{(u,v) \in \bar{E}} \text{cost}(x_u, x_v) y_{uv} + \sum_{v \in S'} \alpha_v h(v) \quad (1)$$

$$\sum_{v \in S'} \alpha_v = 1 \quad (2)$$

$$\sum_{v \in \bar{V}} y_{sv} = 1 \quad (3)$$

$$\sum_{(u,v) \in \bar{E}} y_{uv} = \alpha_v \text{ for } v \in S' \quad (4)$$

$$\sum_{(u,v) \in \delta^-(v)} y_{uv} = \sum_{(v,u) \in \delta^+(v)} y_{vu} \text{ for } v \in S \setminus \{s\} \quad (5)$$

$$x_u \in \mathcal{X}_u \text{ for all } u \in \bar{V} \quad (6)$$

$$y_{uv} \in \{0, 1\} \text{ for all } (u, v) \in \bar{E} \quad (7)$$

$$\alpha_v \in \{0, 1\} \text{ for all } v \in S' \quad (8)$$

Constraints (2) state that exactly one vertex in S' must be chosen as a terminal. Constraints (3)-(5) state the standard flow constraints for a path from s to a terminal in S' . Constraint (6) states that for any vertex $u \in \bar{V}$, its corresponding point must belong to \mathcal{X}_u . The main challenge in solving the above formulation arises from the multiplication of the travel costs and the binary variables in (1). To address this challenge, we follow the approach in (Marcucci et al. 2024b) and reformulate SPP*-GCS as a Mixed Integer Convex Program (MICP) in the next subsection.

MICP for SPP*-GCS

In this paper, as we are dealing with travel costs derived from Euclidean distances, the objective in (1) can be re-written as:

$$C_{opt}^*(S, S') = \min \sum_{(u,v) \in \bar{E}} \|x_u y_{uv} - x_v y_{uv}\|_2 + \sum_{v \in S'} \alpha_v h(v).$$

In (Marcucci et al. 2024b), the bi-linear terms $x_u y_{uv}$ and $x_v y_{uv}$ in the objective above are replaced with new variables, and the constraints in (5),(6) are transformed to form a MICP. Before we present the MICP corresponding to SPP*-GCS, we first need to define the *perspective* of a set. The perspective of a compact, convex set $\mathcal{X} \subset \mathbb{R}^n$ is defined as $\bar{\mathcal{X}} := \{(x, \lambda) : \lambda \geq 0, x \in \lambda \mathcal{X}\}$. Now, let $x_u y_{uv} = z_{uv}$ and $x_v y_{uv} = z'_{uv}$ for all $(u, v) \in \bar{E}$. By following the same procedure outlined in (Marcucci et al. 2024b), we derive a Mixed-Integer Convex Program (MICP) for the SPP*-GCS as follows:

$$\min \sum_{(u,v) \in \bar{E}} \|z_{uv} - z'_{uv}\|_2 + \sum_{v \in S'} \alpha_v h(v) \quad (9)$$

$$\sum_{v \in S'} \alpha_v = 1 \quad (10)$$

$$\sum_{v \in \bar{V}} y_{sv} = 1 \quad (11)$$

$$\sum_{(u,v) \in \bar{E}} y_{uv} = \alpha_v \text{ for } v \in S' \quad (12)$$

$$\sum_{(u,v) \in \delta^-(v)} (z'_{uv}, y_{uv}) = \sum_{(v,u) \in \delta^+(v)} (z_{vu}, y_{vu}) \text{ for } v \in S \setminus \{s\} \quad (13)$$

$$(z_{uv}, y_{uv}) \in \bar{\mathcal{X}}_u \text{ for all } u \in \bar{V}, (u,v) \in \bar{E} \quad (14)$$

$$(z'_{uv}, y_{uv}) \in \bar{\mathcal{X}}_v \text{ for all } v \in \bar{V}, (u,v) \in \bar{E} \quad (15)$$

$$y_{uv} = \{0, 1\} \forall (u,v) \in \bar{E}, \alpha_v = \{0, 1\} \forall v \in S' \quad (16)$$

Lemma 2. *The MICP formulation in (9)-(16) for the SPP*-GCS has an optimal value equal to $C_{opt}^*(S, S')$.*

Proof. This result follows the same proof as Theorem 5.7 in (Marcucci et al. 2024b). \square

We now arrive at the relaxation used in our algorithm:

Convex Relaxation for SPP*-GCS

$$R_{opt}^*(S, S') = \min \sum_{(u,v) \in \bar{E}} \|z_{uv} - z'_{uv}\|_2 + \sum_{v \in S'} \alpha_v h(v)$$

subject to the constraints in (10)-(15) and,

$$\begin{aligned} 0 \leq y_{uv} \leq 1 \text{ for all } (u,v) \in \bar{E}, \\ 0 \leq \alpha_v \leq 1 \text{ for all } v \in S'. \end{aligned} \quad (17)$$

A*-GCS

A*-GCS (Algorithm 1) starts with the input set $S := S_{init}$ and iteratively grows S until further growth is not possible or is not beneficial. This set S during any iteration of A*-GCS will always be a cut-set. A*-GCS keeps track of the growth of S in two phases (*Phase 1* and *Phase 2*). The phases are determined based on the presence of the destination in the neighborhood of S .

In *Phase 1*, d is not a member of the neighborhood of S (line 15 of Algorithm 1). Therefore, in this phase, we can find lower bounds by solving the relaxation (line 17 of Algorithm 1), but we cannot convert a relaxed solution to a feasible solution for SPP-GCS. Additionally, we add a vertex $v \in S'$ to S if any edge (u,v) leaving S has $y_{uv} > 0$ (line 19 of Algorithm 1, Algorithm 3).

In *Phase 2*, d is part of the neighborhood of S , and therefore, it is possible to use the relaxation to find a lower bound and a feasible solution for SPP-GCS (lines 21-22, 32-33 of Algorithm 1). During this phase, we expand S until one of the following two termination conditions is met: 1) all the

vertices in $V \setminus \{d\}$ have already been added to S , or, effectively in this condition, S cannot grow any further (line 23 of Algorithm 1), 2) the relaxation cost of traveling through any vertex in $S' := N_S \setminus \{d\}$ becomes greater than or equal to the relaxation cost of traveling directly to d , or, effectively in this condition, growing S further may not lead to better lower bounds (line 28 of Algorithm 1). Also, the subroutine for expanding S in *Phase 2* (line 31 of Algorithm 1) mirrors that of *Phase 1*.

The following are the *key features* of A*-GCS:

- The initial set S_{init} can be any subset of V as long as it is a cut-set. Specifically, S_{init} can either consist solely of the origin (similar to how we initiate A* for the SPP) or be generated by an algorithm. In this paper, we also consider S_{init} generated by A* as follows: Suppose \bar{S}_{A^*} is the closed set (which also includes d , the last vertex added to \bar{S}_{A^*}) found by A* when implemented on the centroids of all the convex sets associated with the vertices in V ; let $S_{A^*} := \bar{S}_{A^*} \setminus d$ and then assign $S_{init} := S_{A^*}$.
- A*-GCS can also be preemptively stopped at the end of any iteration. A*-GCS will always produce a lower bound at the end of any iteration in any phase and may produce a feasible solution during an iteration if S is processed in *Phase 2*.
- The algorithm for updating the set (Algorithm 3) based on the fractional values of the relaxed solutions is user-modifiable. Therefore, different variants of A*-GCS can be developed tailored to specific applications.

Remark 1. A*-GCS initiated with $S := \{s\}$ mimics A* in the special case when each of the convex sets is a singleton. Refer to the supplementary material in (Sundar and Rathinam 2024) for more details.

Remark 2. While A*-GCS can be used to generate both bounds and feasible solutions, in this paper, we use A*-GCS to primarily find bounds. Also, we generate feasible solutions using the following two-step heuristic which performs reliably well for SPP-GCS: 1) Given any convex set, choose its centroid as its representative point, and run A* on (V, E) with the chosen points to find a path \mathcal{P}_{A^*} . 2) Let $S := V \setminus \{d\}$ and assign $y_{uv} = 1$ for each edge $(u,v) \in \mathcal{P}_{A^*}$ in the formulation (1)-(8). Solve the resulting convex program to obtain the point corresponding to each vertex visited by \mathcal{P}_{A^*} .

Theoretical Results

We will first demonstrate the termination of A*-GCS in a finite number of iterations and subsequently establish the validity of the bounds it generates. Without loss of generality, we assume there is at least one path from s to d in the input graph $G = (V, E)$. This assumption helps avoid certain trivial corner cases that may arise when updating the subsets (lines 19, 31 of Algorithm 1) or the bounds (lines 18, 26 of Algorithm 1).

Lemma 3. A*-GCS will terminate after completing at most $|V| - 1$ iterations of both *Phase 1* and *Phase 2*.

Proof. If the algorithm enters *Phase 1*, each iteration of the phase must add at least one vertex from N_S to S since there

Algorithm 1: A^* -GCS

```
1 Inputs:
2  $G = (V, E)$  // Input graph
3  $s, d \in V$  // Origin and destination
4  $\mathcal{X}_u \forall u \in V$  // Input convex sets
5  $h(v) \forall v \in V$  // admissible lower bounds
6  $S_{\text{init}} \subset V$  // Input cut-set
7 Output:
8  $C_{lb}$  // Lower bound for SPP-GCS
9  $Sol_f$  // Feasible solution for SPP-GCS
10 Initialization:
11  $S \leftarrow S_{\text{init}}$ 
12  $C_{lb} \leftarrow 0$ 
13 Main Loop:
14 Phase 1
15 while  $d \notin N_S$  do
16    $S' \leftarrow N_S$ 
17    $R_{opt}^*(S, S') \leftarrow \text{SolveRelaxation}(S, S')$ 
18    $C_{lb} \leftarrow \max(C_{lb}, R_{opt}^*(S, S'))$ 
19    $S \leftarrow \text{UpdateSubset}(S, S', \mathcal{F}_{nd}^*)$ 
20 Phase 2
21  $R_{opt}^*(S, \{d\}), \mathcal{F}_d^* \leftarrow \text{SolveRelaxation}(S, \{d\})$ 
22  $Sol_f \leftarrow \text{UpdateFeasibleSol}(Sol_f, \mathcal{F}_d^*)$ 
23 while  $\{d\} \neq N_S$  do
24    $S' \leftarrow N_S \setminus \{d\}$ 
25    $R_{opt}^*(S, S'), \mathcal{F}_{nd}^* \leftarrow \text{SolveRelaxation}(S, S')$ 
26    $C_{lb} \leftarrow \max(C_{lb}, \min(R_{opt}^*(S, S'), R_{opt}^*(S, \{d\})))$ 
27    $Sol_f \leftarrow \text{UpdateFeasibleSol}(Sol_f, \mathcal{F}_{nd}^*)$ 
28   if  $R_{opt}^*(S, S') \geq R_{opt}^*(S, \{d\})$  then
29     break
30   else
31      $S \leftarrow \text{UpdateSubset}(S, S', \mathcal{F}_{nd}^*)$ 
32      $R_{opt}^*(S, \{d\}), \mathcal{F}_d^* \leftarrow \text{SolveRelaxation}(S, \{d\})$ 
33      $Sol_f \leftarrow \text{UpdateFeasibleSol}(Sol_f, \mathcal{F}_d^*)$ 
34 return  $C_{lb}, Sol_f$ 
```

Algorithm 2: SolveRelaxation(S, S')

```
1 Solve the relaxation in (17) given  $S$  and  $S'$ 
2  $R_{opt}^*(S, S') \leftarrow$  Optimal relaxation cost
3  $\mathcal{F}^* \leftarrow$  Optimal solution to the relaxation
4 return  $R_{opt}^*(S, S'), \mathcal{F}^*$ 
```

is at least one edge (u, v) leaving S such that $y_{uv} > 0$. If the algorithm enters *Phase 2* and the termination condition in 28 is not met, $R_{opt}^*(S, S')$ is some finite value, and therefore, there is at least one edge (u, v) leaving S and entering S' such that $y_{uv} > 0$. In both the phases, at most $|V| - 1$ vertices can be added to S before one of the termination conditions (lines 23, 28 of Algorithm 1) become valid. Hence proved. \square

Lemma 4. Consider any subset $S \subset V$ such that $s \in S$ and $d \notin S$. Let the path in an optimal solution to the SPP-GCS be denoted as $\mathcal{P}^* := (v_1^* = s, v_2^*, \dots, v_k^* = d)$ and let the optimal points in the corresponding convex sets be denoted as $x_{v_i^*}, i = 1, \dots, k$. For some $p \in \{2, \dots, k\}$, let v_p^* be such that $v_i^* \in S$ for $i = 1, \dots, p - 1$ and $v_p^* \in N_S$. Let S'

Algorithm 3: UpdateSubset($S, S', \mathcal{F}_{nd}^*$)

```
// For any  $(u, v) \in \bar{E}$ , let the optimal
// value of  $y_{uv}$  in  $\mathcal{F}_{nd}^*$  be  $y_{uv}^*$ .
1  $O_S \leftarrow \{v : u \in S, v \in S', y_{uv}^* > 0\}$ 
2  $S \leftarrow S \cup O_S$ 
3 return  $S$ 
```

Algorithm 4: UpdateFeasible(Sol_f, \mathcal{F}_d^*)

```
1 Use randomized rounding (Marcucci et al. 2024b) or other
// heuristics to convert  $\mathcal{F}_d^*$  into a feasible solution,  $Sol_f^*$ .
2 if  $\text{Cost}(Sol_f^*) < \text{Cost}(Sol_f)$  then  $Sol_f \leftarrow Sol_f^*$ ;
3 return  $Sol_f$ 
```

be any subset of N_S that contains v_p^* . Then, the optimal cost of the SPP-GCS is at least equal to the optimal relaxation cost of SPP*-GCS defined on S and S' , i.e., $C_{opt}(s, d) \geq R_{opt}^*(S, S')$.

Proof. Now, the cost of the given optimal solution to SPP-GCS is:

$$\begin{aligned} C_{opt}(s, d) &= \sum_{i=1}^{k-1} \text{cost}(x_{v_i^*}, x_{v_{i+1}^*}) \\ &= \sum_{i=1}^{p-1} \text{cost}(x_{v_i^*}, x_{v_{i+1}^*}) + \sum_{i=p}^{k-1} \text{cost}(x_{v_i^*}, x_{v_{i+1}^*}) \\ &\geq \sum_{i=1}^{p-1} \text{cost}(x_{v_i^*}, x_{v_{i+1}^*}) + h(v_p^*). \quad (\because h(v_p^*) \text{ is admissible}) \end{aligned} \tag{18}$$

Now, note that the path $(v_1^* = s, v_2^*, \dots, v_p^*)$ and the points in the corresponding convex sets $x_{v_i^*}, i = 1, \dots, p$ is a feasible solution to SPP*-GCS defined on S and S' . Therefore, the equation above (18) further reduces to $C_{opt}(s, d) \geq C_{opt}^*(S, S') \geq R_{opt}^*(S, S')$. Hence proved. \square

Theorem 1. Consider any subset $S \subset V$ such that $s \in S$ and $d \notin S$. Then, $C_{opt}(s, d) \geq R_{opt}^*(S, N_S)$.

Proof. This theorem follows by applying Lemma 4 with $S' := N(S)$. Hence proved. \square

Theorem 2. Consider any subset $S \subset V$ such that $s \in S$ and $d \in N_S$. Then, $C_{opt}(s, d) \geq \min(R_{opt}^*(S, N_S \setminus \{d\}), R_{opt}^*(S, \{d\}))$.

Proof. Consider the edge (v_{p-1}^*, v_p^*) in the optimal path in Lemma 4 that leaves S for the first time. Either $v_p^* = d$ or $v_p^* \in N_S \setminus \{d\}$. If $v_p^* = d$, then applying Lemma 4 for $S' = \{d\}$ leads to $C_{opt}(s, d) \geq R_{opt}^*(S, \{d\})$. On the other hand, if $v_p^* \in N_S \setminus \{d\}$, then applying Lemma 4 for $S' = N_S \setminus \{d\}$ leads to $C_{opt}(s, d) \geq R_{opt}^*(S, N_S \setminus \{d\})$. Therefore, $C_{opt}(s, d) \geq \min(R_{opt}^*(S, \{d\}), R_{opt}^*(S, N_S \setminus \{d\}))$. Hence, proved. \square

Numerical Results

Set up: We initially test our algorithm on graphs generated from mazes and axis-aligned bars in 2D (Fig. 3). Due to space constraints, the procedures for generating these graphs are given in the supplementary material in (Sundar and Rathinam 2024). The convex sets in the GCS generated from any maze and axis-aligned instances are line segments (Fig. 3a) and unit square regions (Fig. 3b) respectively. The size of the GCS generated for each of the maps is provided in Table 1. The destination for each instance is always chosen to be the centroid of the top-rightmost square in the instance.

Map No.	Type	$ V $	$ E $
1	Maze	121	425
2	Maze	415	1020
3	Maze	1615	3649
4	Maze	6417	14239
5	Bars	769	6525
6	Bars	1790	15912

Table 1: 2D instances generated from mazes and bars.

We will compare the performance of the lower bounds generated by A^* -GCS with the lower bound generated by the baseline algorithm that solves the convex relaxation of SPP-GCS on the entire graph. The upper bound¹ for both algorithms is the length of the feasible solution generated by the two-step heuristic. Note that the size of the cut-set during any iteration of A^* -GCS determines the sizes of the convex relaxations we will solve. Therefore, we will keep track of the sizes of the cut-sets in addition to the run times of the algorithms. Also, the baseline is equivalent to implementing (17) with $S := V \setminus \{d\}$; therefore, the size of the cut-set corresponding to the baseline is $|V| - 1$. All the algorithms were implemented using the Julia programming language (Bezanson et al. 2017) and run on an Intel Haswell 2.6 GHz, 32 GB, 20-core machine. Gurobi (Gurobi Optimization, LLC 2023) was used to solve all the convex relaxations.

Heuristics for generating underestimates: We designed two heuristics to generate underestimates for A^* -GCS. Given a set, the first heuristic (h_1) computes the shortest Euclidean distance between any point in the set and the destination, ignoring all other vertices/edges in the graph. h_1 is very fast and requires negligible computation time; it is by default present in A^* -GCS and is already part of all the run times reported in the results section. While h_1 finds consistent underestimates, its bounds may not be strong.

The second heuristic (h_2), referred to as the *expand and freeze* heuristic, computes underestimates of the optimum through an iterative process. It starts from the destination and proceeds in the reverse direction, adding a subset of vertices U in each iteration to the set containing the destination while determining the underestimate for each vertex in U .

¹We did not use the rounding heuristic in (Marcucci et al. 2024b) to generate upper bounds because its performance, especially for maps 5, 6 and the larger 2D/3D maps, was significantly worse compared to the bounds produced by the two-step heuristic.

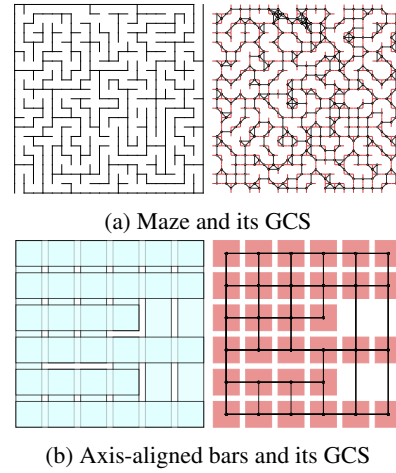


Figure 3: 2D maps and their GCS. In the GCS corresponding to the Maze, the red segments show the convex sets and the black lines show the edges. Similarly, in the GCS corresponding to the axis-aligned bars, the shaded squares show the convex sets and the black lines show the edges.

Specifically, in each iteration, the subset U and its corresponding underestimates are obtained by solving a convex relaxation of SPP*-GCS. Any vertex that is incident to any fraction of an edge in the relaxed solution is included in U . To limit the size of the relaxations, as soon as the size of the set containing the destination reaches a limit (say n_{max}), we shrink all the vertices in the set into a new destination vertex and repeat the process until all the (reachable) vertices are visited. In this paper, we set n_{max} as 100. Generally, for our maps, h_2 produced tighter estimates than h_1 , but h_2 can be inconsistent. h_2 can run either in the order of seconds or minutes, depending on the map. We specify its computation times for all the maps in each of the results sections. We also study the impact of using a convex combination of the two heuristics, *i.e.*, for any vertex $v \in V$, we define the combined heuristic estimate as $h(v) := (1 - w)h_1(v) + wh_2(v)$ where w is the weighting factor.

Comparisons based on number of iterations and the sizes of the cut-sets for A^* -GCS variants: We compare the performance of A^* -GCS for two choices of S_{init} : one with $S_{init} := \{s\}$ and another with $S_{init} := S_{A^*}$ (refer to the key features of A^* -GCS on page 4). We fix the origin for each map in Table 1 at the centroid of the left-bottommost unit square and implement the variants of A^* -GCS on the corresponding graphs for heuristic weight $w = 1$. Both choices of S_{init} produced the same bounds for all the six maps. Therefore, for both the variants, we compare the total number of iterations for both phases of A^* -GCS, the size of the cut-set ($|S|$) upon termination of A^* -GCS, and the total run time (in secs.) as shown in Table 2. While $|S|$ for both variants of A^* -GCS are relatively closer, A^* -GCS initialized with $S_{init} := \{s\}$ required significantly more iterations and run time to terminate compared to the A^* -GCS initialized with $S_{init} := S_{A^*}$. This trend remained the same for any other choice of origin in these maps. Therefore, for the remaining

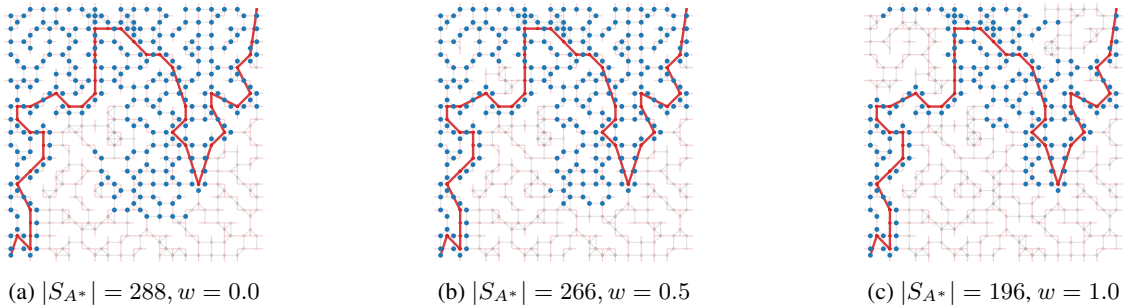


Figure 4: The feasible path (thick red lines) and the cut-set vertices (shown in color blue) produced by the two-step A^* based heuristic for different weights. There are 415 vertices in this GCS corresponding to map no. 2. A^* -GCS in this case terminated in 1 iteration for all the weights and produced the same bounds. By first finding the cut-set (S_{A^*}) and then applying A^* -GCS with $S_{init} := S_{A^*}$, we reduced the size of the relaxation with $|V| - 1 = 414$ vertices into a relaxation with $|S_{A^*}|$ vertices.

results in this paper, we will only focus on the variant of A^* -GCS initialized with $S_{init} := S_{A^*}$. Note that in this variant, the algorithm will never enter Phase 1 since the destination d is already present in the neighborhood of S_{A^*} . Henceforth, the number of iterations of A^* -GCS will simply refer to the number of iterations in Phase 2 of A^* -GCS.

Map No.	$S_{init} := \{s\}$			$S_{init} := S_{A^*}$		
	n_{iter}	$ S $	time	n_{iter}	$ S $	time
1	27	36	0.23	3	44	0.07
2	117	149	3.4	1	196	0.1
3	383	422	43.98	2	518	0.69
4	1726	2699	1077.99	2	4003	4.48
5	110	585	126.3	1	654	2.8
6	171	1657	962.94	1	1693	24.06

Table 2: No. of iterations (n_{iter}), cut-set size ($|S|$) upon termination and run time (in secs.) for the A^* -GCS variants.

Map No.	1	2	3	4	5	6
Time	14.4	21.5	48.4	151.7	43.3	94.4

Table 3: Heuristic h_2 run time in secs.

Impact of the quality of underestimates: For each map in Table 1, we generated 100 instances, with each instance choosing a different origin for the path. The run times for heuristic h_2 are shown in Table 3. Note that this heuristic is run only once and is used as an input for all the 100 instances for any heuristic weight $w > 0$. Given an instance I , its optimality gap is defined as $100 \times \frac{C_f - C_{lb}}{C_{lb}}$, where C_f is the length of the feasible solution obtained by the two-step algorithm in Remark 2, and C_{lb} is the bound returned by A^* -GCS. In Table 4, we present the results (average size of cut-sets, average run times in seconds, and the average optimality gap) at the end of the first iteration of A^* -GCS (referred to as A^* -GCS₁) as well as upon the termination of A^* -GCS (referred to as A^* -GCS_∞). Even with a heuris-




Figure 5: 3D village map.

tic weight $w = 0$ (which corresponds to the first heuristic h_1 that provides relatively weak bounds), A^* -GCS is able to reduce the size of the cut-sets (which, in turn, *reduces the sizes of convex relaxations to solve*) by more than 45%, on average (except in map no. 6). In fact, by selecting the cut-set informed by A^* with $w = 0$ and completing just one iteration of A^* -GCS, we achieve optimality gaps in a shorter time that closely match those provided by the baseline algorithm (except for map no. 6). Refer to Fig. 4 for an illustration. For map. no. 6 the underestimates provided by h_1 was much weaker compared to h_2 , and hence A^* -GCS performed significantly better when the heuristic weight was increased to 1.

For single query problems, the results for the 2D maps show that the default heuristic (h_1) in A^* -GCS is sufficient to provide benefits in terms of reduction in problem sizes and computation time. On the other hand, for multi-query problems, using h_2 can further significantly reduce the computation time while providing similar optimality gaps. For example, in map no. 6, A^* -GCS₁ is at least six times faster than the baseline algorithm for $w = 1$. A^* -GCS performed the best in map no. 4 for all heuristic weights; specifically, A^* -GCS₁ found optimal solutions with average run times considerably lower compared to the baseline for this map. As the heuristic weight increased, the size of the cut-sets as well as the run times decreased on average.

A^* -GCS on large 3D village maps: Next, A^* -GCS was tested on two large 3D village maps (Fig. 5) created using the meshcat library in Python. We followed a procedure

Map No.	Algorithm	$w = 0$			$w = 0.25$			$w = 0.5$			$w = 0.75$			$w = 1.0$		
		$ S $	time	gap	$ S $	time	gap	$ S $	time	gap	$ S $	time	gap	$ S $	time	gap
1	Baseline	120.0	0.1	0.1	120.0	0.1	0.1	120.0	0.1	0.1	120.0	0.1	0.1	120.0	0.1	0.1
	A^*-GCS_∞	37.3	0.1	0.1	35.4	0.1	0.1	32.8	0.0	0.1	31.1	0.0	0.1	30.4	0.0	0.1
	A^*-GCS_1	36.1	0.0	3.4	34.1	0.0	3.4	31.6	0.0	3.4	29.8	0.0	3.4	28.9	0.0	3.4
2	Baseline	414.0	0.4	0.3	414.0	0.4	0.3	414.0	0.4	0.3	414.0	0.3	0.3	414.0	0.3	0.3
	A^*-GCS_∞	225.9	0.2	0.3	201.2	0.2	0.3	174.5	0.1	0.3	151.3	0.1	0.3	123.9	0.1	0.3
	A^*-GCS_1	225.9	0.2	0.5	201.2	0.1	0.4	174.5	0.1	0.3	151.2	0.1	0.4	123.9	0.1	0.3
3	Baseline	1614.0	1.0	0.0	1614.0	1.0	0.0	1614.0	1.0	0.0	1614.0	1.0	0.0	1614.0	1.0	0.0
	A^*-GCS_∞	683.6	1.0	0.0	610.5	0.8	0.0	531.5	0.7	0.0	467.1	0.6	0.0	402.9	0.5	0.0
	A^*-GCS_1	683.6	0.5	0.3	610.5	0.4	0.2	531.5	0.3	0.1	467.1	0.3	0.1	402.9	0.2	0.1
4	Baseline	6416.0	5.9	0.0	6416.0	5.9	0.0	6416.0	5.9	0.0	6416.0	5.8	0.0	6416.0	5.9	0.0
	A^*-GCS_∞	2858.2	3.4	0.0	2575.9	3.0	0.0	2309.3	2.7	0.0	2054.3	2.3	0.0	1786.8	1.9	0.0
	A^*-GCS_1	2858.2	1.7	0.0	2575.9	1.4	0.0	2309.3	1.3	0.0	2054.3	1.1	0.0	1786.8	1.0	0.0
5	Baseline	768.0	3.8	2.6	768.0	3.8	2.6	768.0	3.9	2.6	768.0	3.9	2.6	768.0	3.8	2.6
	A^*-GCS_∞	343.9	2.8	2.6	314.6	2.5	2.6	291.7	2.1	2.5	265.3	1.8	2.6	241.2	1.3	2.6
	A^*-GCS_1	340.9	1.4	3.3	312.8	1.3	3.4	290.3	1.2	3.0	264.6	1.2	3.0	240.5	1.0	2.8
6	Baseline	1789.0	20.5	2.4	1789.0	20.6	2.4	1789.0	20.5	2.4	1789.0	20.6	2.4	1789.0	20.6	2.4
	A^*-GCS_∞	911.2	24.0	2.4	809.3	19.4	2.4	713.3	13.3	2.4	616.8	10.2	2.4	515.1	6.4	2.4
	A^*-GCS_1	882.9	9.4	4.5	788.4	7.1	3.9	694.4	6.0	3.7	601.2	3.6	4.2	508.7	2.8	2.9

Table 4: Average size of the cut-set ($|S|$), average run times (in secs) and average optimality gap (in %) for different levels of heuristic weights for maps in Table 1.




Figure 6: Histogram showing the empirical distribution of the fraction of vertices solved by A^*-GCS_∞ for the 200 instances of the 3D maps. Here, $|S|$ denotes the cut-set size and $|V|$ represents the total number of vertices in the graph.

similar to the one described for maps in Table 1 to develop the GCS for these village maps. The convex sets in these maps, resulting from the intersection of adjacent cubes, are axis-aligned squares (shared sides of adjacent cubes) computed using the algorithm in (Zomorodian and Edelsbrunner 2000). The GCS corresponding to the first 3D map has 6010 vertices and 166,992 edges, while the GCS for the second 3D map has 11,258 vertices and 321,603 edges. In both

maps, the destination was chosen at one of the corners of the 3D region. The computation time for heuristic h_2 for these two maps was 53.30 secs and 89.10 secs respectively.

We generated 100 instances for each 3D map, varying the origins of the paths. Across all 200 instances, A^*-GCS_∞ produced approximately the same optimality gaps as the baseline but, on average, explored only 20.38% of all vertices in the graphs. The histogram showing the fraction of vertices solved by the convex program using A^*-GCS_∞ relative to the total number of vertices for all instances is presented in Fig 6. Specifically, for 63% of the first 3D map instances, A^*-GCS_∞ ran faster than the baseline, with an average savings of 334.27 seconds (around a 61% reduction). Similarly, for 63% of the second 3D map instances, A^*-GCS_∞ ran faster than the baseline, with an average savings of 904.49 seconds (around a 65% reduction). For the remaining instances, even though A^*-GCS_∞ required more time, A^*-GCS_1 produced optimality gaps comparable to the baseline with smaller computation times (refer to the supplementary material in (Sundar and Rathinam 2024)).

Conclusions

In this paper, we show how to effectively combine the existing convex-programming based approach with heuristic information to obtain near-optimal solutions for SPP-GCS. While our results clearly illustrate the benefits of A^*-GCS for travel costs based on Euclidean distances, we do not claim that the proposed approach is superior to directly solving the relaxation of SPP-GCS for every instance. Nevertheless, the proposed approach is the first of its kind for solving convex relaxations of SPP-GCS and opens a new avenue for research into related problems.

Acknowledgments

This material is based upon work partially supported by the National Science Foundation under Grant No. 2120219. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

References

- Bezanson, J.; Edelman, A.; Karpinski, S.; and Shah, V. B. 2017. Julia: A fresh approach to numerical computing. *SIAM Review*, 59(1): 65–98.
- Chia, S. Y. C.; Jiang, R. H.; Graesdal, B. P.; Kaelbling, L. P.; and Tedrake, R. 2024. GCS*: Forward Heuristic Search on Implicit Graphs of Convex Sets. arXiv:2407.08848.
- Dumitrescu, A.; and Mitchell, J. S. 2003. Approximation algorithms for TSP with neighborhoods in the plane. *Journal of Algorithms*, 48(1): 135–159.
- Fourney, E.; Burdick, J. W.; and Rimon, E. D. 2024. Mobile Robot Sensory Coverage in 2-D Environments: An Optimization Approach with Efficiency Bounds. arXiv:2405.15100.
- Gurobi Optimization, LLC. 2023. Gurobi Optimizer Reference Manual.
- Hart, P. E.; Nilsson, N. J.; and Raphael, B. 1968. A formal basis for the heuristic determination of minimum cost paths. *IEEE Transactions on Systems Science and Cybernetics*, SSC-4(2): 100–107.
- Korte, B.; and Vygen, J. 2018. *Shortest Paths*, 159–175. Berlin, Heidelberg: Springer Berlin Heidelberg. ISBN 978-3-662-56039-6.
- Lawler, E. L. 1979. Shortest Path and Network Flow Algorithms. In Hammer, P.; Johnson, E.; and Korte, B., eds., *Discrete Optimization I*, volume 4 of *Annals of Discrete Mathematics*, 251–263. Elsevier.
- Marcucci, T.; Nobel, P.; Tedrake, R.; and Boyd, S. 2024a. Fast Path Planning Through Large Collections of Safe Boxes. *IEEE Transactions on Robotics*, 40: 3795–3811.
- Marcucci, T.; Petersen, M.; von Wrangel, D.; and Tedrake, R. 2023. Motion planning around obstacles with convex optimization. *Science Robotics*, 8(84): eadf7843.
- Marcucci, T.; Umenberger, J.; Parrilo, P.; and Tedrake, R. 2024b. Shortest Paths in Graphs of Convex Sets. *SIAM Journal on Optimization*, 34(1): 507–532.
- Mitchell, J. 2017. Shortest Paths and Networks. In *In Toth, C.D., O’Rourke, J., and Goodman, J.E. (Eds.), Handbook of Discrete and Computational Geometry (3rd ed.)*. Chapman and Hall/CRC.
- Natarajan, R.; Liu, C.; Choset, H.; and Likhachev, M. 2024. Implicit Graph Search for Planning on Graphs of Convex Sets. *Robotics: Science and Systems (RSS)*.
- Sundar, K.; and Rathinam, S. 2024. A* for Graphs of Convex Sets. arXiv:2407.17413.
- Zomorodian, A.; and Edelsbrunner, H. 2000. Fast software for box intersections. In *Proceedings of the sixteenth annual symposium on Computational geometry*, 129–138.