

SHORTEST PATHS IN GRAPHS OF CONVEX SETS*

TOBIA MARCUCCI[†], JACK UMENBERGER[†], PABLO A. PARRILO[†], AND RUSS
TEDRAKE[†]

Abstract. Given a graph, the shortest-path problem requires finding a sequence of edges with minimum cumulative length that connects a source vertex to a target vertex. We consider a variant of this classical problem in which the position of each vertex in the graph is a continuous decision variable constrained in a convex set, and the length of an edge is a convex function of the position of its endpoints. Problems of this form arise naturally in many areas, from motion planning of autonomous vehicles to optimal control of hybrid systems. The price for such a wide applicability is the complexity of this problem, which is easily seen to be NP-hard. Our main contribution is a strong and lightweight mixed-integer convex formulation based on perspective operators, that makes it possible to efficiently find globally optimal paths in large graphs and in high-dimensional spaces.

Key words. Shortest-path problem, graph problems with neighborhoods, mixed-integer convex programming, perspective formulation, optimal control.

MSC codes. 52B05, 90C11, 90C25, 90C35, 90C57, 93C55, 93C83.

1. Introduction. The Shortest-Path Problem (SPP) is one of the most important and ubiquitous problems in combinatorial optimization. In its single-source single-target version, this problem asks for a path of minimum length connecting two prescribed vertices of a graph, where the length of a path is defined as the sum of the lengths of its edges. Typically, the edge lengths are fixed scalars, given as problem data, and the assumptions made on their values have a dramatic impact on the problem complexity [43, Chapters 6 to 8]. In this paper we introduce the SPP in Graph of Convex Sets (GCS), a variant of the SPP in which the edge lengths are convex functions of continuous variables representing the position of the vertices (see Figure 1). More precisely, a GCS is a directed graph in which each vertex is paired with a convex set. The spatial position of a vertex is a continuous variable, constrained to lie in the corresponding convex set. The length of an edge is a given convex function of the position of the vertices that this edge connects. When looking for a path of minimum length in a GCS, we then have the extra degree of freedom of optimizing the position of the vertices visited by the path. According to the literature, this problem could also be classified as an SPP *with neighborhoods*; we call it SPP in GCS to highlight the crucial role that convexity plays in the developments of this paper.

Many problems of practical interest can be formulated as SPPs in GCS: for some of those the convex sets and the edge-length functions are naturally suggested by the application, for others the construction of the GCS requires more thinking. As an example of the former class of problems, scheduling the flight of a drone with limited batteries is immediately cast as an SPP in GCS like the one in Figure 1. The start region is on the left, the goal region is on the right, and the remaining regions can be used for recharging. Pairs of regions that are close enough for the drone to fly between are connected by an edge. The objective is to minimize the overall length of the flight. Optimal control of discrete-time hybrid dynamical systems [4] is a main application that we target in this paper, and is an example of a problem whose formulation as an

* Submitted to the editors DATE.

Funding: National Science Foundation, award no. EFMA-1830901. Department of the Navy, Office of Naval Research, award no. N00014-18-1-2210 and award no. N00014-17-1-2699.

[†] Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, Cambridge, MA (tobiam@mit.edu, umnbrgr@mit.edu, parrilo@mit.edu, russt@mit.edu).




FIG. 1. Example of an SPP in GCS. The source set is on the left and the target set is on the right. The graph edges are arrows, and the shortest path is shown in dashed green. The dotted red lines connect the optimal positions of the vertices along the shortest path.

SPP in GCS is nontrivial. In this case we let the convex sets live in the joint state and control space of the dynamical system. Each discrete time step corresponds to an edge transition in the GCS, and the edge lengths quantify, e.g., the energy consumed to move between states (a length that is infinite if the motion is not compatible with the system dynamics). This is explained in detail in Section 8.

1.1. Contributions. The following are the main contributions of this article.

Problem statement (Section 2). The SPP in GCS represents an unexplored class of problems at the interface of combinatorial and convex optimization. It lends itself to a simple problem statement and, at the same time, it is a versatile framework that includes as special cases many problems of practical relevance.

Mixed-integer convex formulation (Section 5). The SPP in GCS is easily seen to be NP-hard (Section 3). Our main contribution is the formulation of this problem as a strong and lightweight Mixed-Integer Convex Program (MICP). This program extends in a natural way the classical network-flow formulation of the SPP, and it allows us to efficiently find shortest paths in large graphs (hundreds of vertices) and high-dimensional spaces (tens of dimensions). In addition, the design principles of this MICP can be applied to improve existing mixed-integer formulations of other graph problems with neighborhoods, which are limited to small graphs and sets in two or three dimensions (see Appendix A).

Set-based convex relaxation of bilinear constraints (Section 7). The main building block of our MICP is a tight and compact convex relaxation for a class of bilinear constraints that emerge naturally in our problem. This relaxation is *set based*, in the sense that it does not rely on the explicit constraints that define the sets in our GCS, but it works directly with their abstract set representations. This makes our MICP usable even when these sets are black boxes accessible only through a separation oracle. This relaxation is similar in spirit to the Lovász-Schrijver one [31], and is based on perspective operators (a popular tool in mixed-integer optimization [9, 47, 18, 22]).

Control applications (Section 8). Computation times are the main limitation to a widespread application of mixed-integer optimization in control of hybrid systems [38, 46, 35]. Our shortest-path formulation of these problems is substantially different from the state of the art [37, 34], as we do not use binary variables to encode the discrete mode in which the system is at each time step but, instead, we use them to select the transitions between modes. This different parameterization yields slightly larger but much stronger MICPs that, in our computational experiments, are orders of magnitude faster to solve.

1.2. Related graph problems. In this subsection we overview a few variants of classical graph problems that are closely related to our problem formulation.

Graph problems with neighborhoods. Graph problems where the vertices are allowed to move within corresponding sets are often called problems with neighborhoods. The SPP with neighborhoods has been analyzed in [13] under stringent assumptions that ensure polynomial-time solvability: the sets are disjoint rectilinear polygons in the plane, and the edge lengths penalize the \mathcal{L}_1 distance between the vertices. The applications we target with this paper, however, do not verify any of these hypotheses. A special case of the SPP with neighborhoods is the touring-polygon problem, which asks for the shortest path between two points that visits a set of polygons in a given order [14]. Our problem differs from this in that our sets are convex and the order in which we visit them is not predefined. Other problems akin to the touring polygon, but substantially different from the SPP in GCS, are the safari, the zookeeper, and the watchman route; see [29, Part IV] and the references therein.

The Traveling-Salesman Problem (TSP) and the Minimum-Spanning-Tree Problem (MSTP) are the two combinatorial problems that have been studied most extensively in their variants with neighborhoods [3, 50]. Exact algorithms for these generally rely on expensive mixed-integer nonconvex optimization [19, 5, 7], and do not scale beyond two or three dimensions. Although the techniques we propose in this paper are particularly well suited to the structure of the SPP, they can be used without modifications to formulate other graph problems with neighborhoods as very tractable MICPs (see Appendix A).

Graph problems with clusters. Generalized Steiner problems [15] (otherwise known as generalized network-design problems [17, 40]) can be thought as the discrete counterpart of the graph problems with neighborhoods: the vertex set is partitioned into clusters and the problem constraints are expressed in terms of these clusters, rather than the original vertices. A clustered version of the SPP has been presented in [30]: each vertex in the graph is assigned a nonnegative weight, and the total vertex weight incurred by the shortest path within each cluster must not exceed a given value. The problem we analyze in this paper can be approximated as an SPP with clusters in a natural way. In low-dimensional spaces, this approximation can be computationally efficient and sufficiently accurate for practical applications. However, this strategy is infeasible in high dimensions, where covering a volume of space with a cluster requires an exponential number of points.

Euclidean shortest paths. Another related variant of the SPP is the Euclidean SPP [29], where we look for a continuous path that connects two points and avoids given polygonal obstacles. In two dimensions, this problem can be reduced to a discrete search and is solvable in polynomial time [32]. In three dimensions or more the problem is NP-hard [8, Theorem 2.3.2], and common algorithms rely on a grid discretization of the space [27]. More recently, a moment-based technique that handles semialgebraic obstacles has been proposed in [16].

2. Problem statement. We start with a formal statement of the SPP in GCS. Let $G := (\mathcal{V}, \mathcal{E})$ be a directed graph with vertex set \mathcal{V} and edge set \mathcal{E} . For each vertex $v \in \mathcal{V}$, we have a nonempty compact convex set $\mathcal{X}_v \subset \mathbb{R}^n$ and a point \mathbf{x}_v contained in it.¹ The length of an edge $e = (u, v) \in \mathcal{E}$ is determined by the location of the points

¹The results presented in this paper are easily extended to the case in which the sets \mathcal{X}_v do not have common dimension n .

\mathbf{x}_u and \mathbf{x}_v via the expression $\ell_e(\mathbf{x}_u, \mathbf{x}_v)$. The *edge length* function ℓ_e takes values in $\mathbb{R}_{\geq 0} \cup \{\infty\}$ and is assumed to be proper, closed, and convex. Note that, despite its name, we do not assume ℓ_e to be a valid metric, and properties like symmetry or the triangle inequality are not required to hold. Given a source vertex s and a target vertex $t \neq s$, an s - t path p is a sequence of distinct vertices (v_0, \dots, v_K) such that $v_0 = s$, $v_K = t$, and $(v_k, v_{k+1}) \in \mathcal{E}$ for all $k = 0, \dots, K-1$. We denote with $\mathcal{E}_p := \{(v_0, v_1), \dots, (v_{K-1}, v_K)\}$ the set of edges traversed by this path, and with \mathcal{P} the set of all s - t paths in the graph G . The SPP in GCS is then stated as

$$\begin{aligned}
 (2.1a) \quad & \text{minimize} && \sum_{e=(u,v) \in \mathcal{E}_p} \ell_e(\mathbf{x}_u, \mathbf{x}_v) \\
 (2.1b) \quad & \text{subject to} && p \in \mathcal{P}, \\
 (2.1c) \quad & && \mathbf{x}_v \in \mathcal{X}_v, \quad \forall v \in p.
 \end{aligned}$$

The decision variables are the discrete path p and the continuous values \mathbf{x}_v . The cost (2.1a) minimizes the total path length. Constraint (2.1c) is enforced only for the vertices visited by the path, since the positions of the other vertices are irrelevant.

The edge length used in Figure 1 is the Euclidean distance:

$$(2.2) \quad \ell_e(\mathbf{x}_u, \mathbf{x}_v) := \|\mathbf{x}_v - \mathbf{x}_u\|_2.$$

With this choice the polygonal line connecting the points \mathbf{x}_v along a shortest path is as straight as possible, perfectly straight if $(s, t) \in \mathcal{E}$. Conversely, if the edge length is the Euclidean distance squared,

$$(2.3) \quad \ell_e(\mathbf{x}_u, \mathbf{x}_v) := \|\mathbf{x}_v - \mathbf{x}_u\|_2^2,$$

straight trajectories may be suboptimal if they require long steps $\mathbf{x}_v - \mathbf{x}_u$. Note also that by letting ℓ_e take infinite value outside a convex set \mathcal{X}_e we are effectively enforcing the edge constraint $(\mathbf{x}_u, \mathbf{x}_v) \in \mathcal{X}_e$. This will be used in Section 8 to formulate optimal-control problems as SPPs in GCS: there the edge constraints will couple the vertex positions according to the system dynamics.

3. Complexity analysis. If we fix the vertex positions \mathbf{x}_v , problem (2.1) simplifies to the classical SPP with scalar nonnegative edge lengths, which is easily solvable using, e.g., Linear Programming (LP). Similarly, if we fix the path p , problem (2.1) simplifies to a convex program that can be efficiently solved for most convex sets \mathcal{X}_v and edge lengths ℓ_e . In this section we show that the simultaneous optimization of the vertex positions and the path makes the SPP in GCS an NP-hard problem.

Recall that an s - t path $p := (v_0, \dots, v_K)$ is said to be Hamiltonian if it visits every vertex in the graph (i.e., if $K = |\mathcal{V}| - 1$), and a graph is Hamiltonian if it contains such a path. The Hamiltonian-Path Problem (HPP) asks if a given graph is Hamiltonian. As an example, the graph in Figure 1 is not Hamiltonian.

THEOREM 3.1. *The SPP in GCS (2.1) is NP-hard.*

Proof. We show that the HPP is polynomial-time reducible to the SPP in GCS. The thesis will then follow since the HPP is NP-complete [26]. We construct an SPP in GCS that shares the same graph G as the given HPP. We let the source $\mathcal{X}_s := \{0\}$ and target $\mathcal{X}_t := \{1\}$ sets be singletons on the real line, while we define $\mathcal{X}_v := [0, 1]$ for all $v \neq s, t$. The length of each edge is the Euclidean distance squared (2.3).

Given these choices, the optimal positioning of the vertices for a fixed path p is given by $\mathbf{x}_{v_k} = k/K$ for $k = 0, \dots, K$. The length of this path is $K(1/K)^2 = 1/K$. We conclude that an optimal path is one for which K is maximized, and is Hamiltonian if and only if G is Hamiltonian. This reduction operates in polynomial time. \square

This simple reduction shows that, even if the convex sets \mathcal{X}_v are one-dimensional intervals, the SPP in GCS can be a hard problem. Nonetheless, one might wonder if additional assumptions on the problem data could turn the SPP in GCS into a problem that is solvable in polynomial time.

- What if the graph G is acyclic? In case of an acyclic graph the HPP is solvable in linear time [2, Section 4.4], and our hardness proof is not valid.
- What if the sets \mathcal{X}_v are disjoint? In fact, some graph problems with neighborhoods can be solved more efficiently in case of disjoint neighborhoods [13, 7].
- What if the edge lengths ℓ_e are positively homogeneous? An edge length like the Euclidean distance (2.2) could not be used in our reduction since it would not force the optimal path p to visit as many vertices as possible.

It turns out that all these questions have a negative answer. This is summarized in the following theorem, whose proof is omitted since it is a long and relatively straightforward adaptation of the complexity analysis of the Euclidean SPP from [8].

THEOREM 3.2. *Assume that the graph G is acyclic, the sets \mathcal{X}_v are disjoint, and the edge lengths ℓ_e are positively homogeneous. The SPP in GCS (2.1) is NP-hard.*

4. Convex-analysis background. This section introduces two basic concepts in convex analysis: perspective operators (homogenization) and valid inequalities (duality). These are the main tools that we will use in the design and the analysis of our MICP. Our goal here is to set the notation and collect some important definitions and properties; for a comprehensive introduction to these topics see [42, Parts II and III] or [25, Chapters III and IV].

4.1. Perspective operators. There is a natural construction that maps a set in n dimensions to a cone in $n + 1$ dimensions. This is sometimes called *homogenization*, or the *cone over* the set. Here we call it *perspective*, for coherence with the name commonly used for the same operation applied to functions [25, Section IV.2.2].

DEFINITION 4.1. *We define the perspective of a closed convex set $\mathcal{X} \subseteq \mathbb{R}^n$ as*

$$\tilde{\mathcal{X}} := \text{cl}\{(\mathbf{x}, \lambda) : \lambda \geq 0, \mathbf{x} \in \lambda\mathcal{X}\},$$

where cl denotes the closure of the set.

Remark 4.2. The closure operation in Definition 4.1 is unnecessary for bounded sets \mathcal{X} . While, when the set \mathcal{X} is unbounded, it ensures that the perspective $\tilde{\mathcal{X}}$ contains all its limit points with $\lambda = 0$ [42, Theorem 8.2].

Importantly, the perspective operation preserves convexity, and the set $\tilde{\mathcal{X}}$ is a closed convex cone. The next example shows that the perspective of a set represented in conic form can be computed very easily.

Example 4.3. Let $\mathcal{X} := \{\mathbf{x} : \mathbf{A}\mathbf{x} + \mathbf{b} \in \mathcal{K}\}$, for some matrix \mathbf{A} , vector \mathbf{b} , and closed convex cone \mathcal{K} . We have $\tilde{\mathcal{X}} = \{(\mathbf{x}, \lambda) : \lambda \geq 0, \mathbf{A}\mathbf{x} + \mathbf{b}\lambda \in \mathcal{K}\}$.

This example has great practical relevance since, informally, it tells us that if a conic-optimization solver can handle the set \mathcal{X} then it can also handle its perspective $\tilde{\mathcal{X}}$. For instance, we see that the perspective of a polyhedral, ellipsoidal, and spectrahedral set can be represented through a set of linear, second-order-cone, and

semidefinite constraints, respectively. More in general, if the set \mathcal{X} is bounded, the formal equivalence of optimizing over \mathcal{X} and its perspective $\tilde{\mathcal{X}}$ can be established using the ellipsoid method [21, Chapter 4], since the separation problems for these two sets are easily seen to be equivalent.

The next definition uses the construction from [42, Page 39] to describe what the perspective operation does to a convex function.

DEFINITION 4.4. *We define the perspective of a closed convex function $f : \mathbb{R}^n \rightarrow \mathbb{R} \cup \{\infty\}$ as the unique function \tilde{f} whose epigraph is the perspective of the epigraph of f , i.e.,*

$$\tilde{f}(\mathbf{x}, \lambda) := \inf\{\sigma : (\mathbf{x}, \sigma, \lambda) \in \widetilde{\text{epi } f}\},$$

where $\text{epi } f := \{(\mathbf{x}, \sigma) : f(\mathbf{x}) \leq \sigma\}$.²

Since its epigraph is closed and convex, the perspective function \tilde{f} is closed and jointly convex in \mathbf{x} and λ .

Remark 4.5. For $\lambda > 0$, noticing that $\lambda \text{epi } f = \{(\mathbf{x}, \sigma) : \lambda f(\mathbf{x}/\lambda) \leq \sigma\}$, we have that the perspective function is $\tilde{f}(\mathbf{x}, \lambda) = \lambda f(\mathbf{x}/\lambda)$. For $\lambda < 0$, we immediately see that $\tilde{f}(\mathbf{x}, \lambda) = \infty$. The behavior for $\lambda = 0$ is more complicated [42, Corollary 8.5.2], but for the scope of this paper it suffices to note that if f is proper then, by the closedness of \tilde{f} , we must have $\tilde{f}(\mathbf{0}, 0) = 0$.

Although Definition 4.4 might seem unsuitable for numerical optimization, the perspectives of most common functions f can be minimized using standard solvers. In fact, given a conic representation of the epigraph of f , we can compute the epigraph of \tilde{f} as in Example 4.3, and minimize \tilde{f} using a slack variable.

The next two examples draw further useful parallels between the perspective operation applied to sets and to functions.

Example 4.6. Let \mathcal{X} be a nonempty closed convex set and g be a finite convex function. Define $f(\mathbf{x}) := g(\mathbf{x})$ if $\mathbf{x} \in \mathcal{X}$ and $f(\mathbf{x}) := \infty$ otherwise. We have $\tilde{f}(\mathbf{x}, \lambda) = \tilde{g}(\mathbf{x}, \lambda)$ if $(\mathbf{x}, \lambda) \in \tilde{\mathcal{X}}$ and $\tilde{f}(\mathbf{x}, \lambda) = \infty$ otherwise.

Example 4.7. For a set $\mathcal{X} := \{\mathbf{x} : f_i(\mathbf{x}) \leq 0 \text{ for all } i \in \mathcal{I}\}$, where the functions f_i are closed and convex, we have $\tilde{\mathcal{X}} = \{(\mathbf{x}, \lambda) : \tilde{f}_i(\mathbf{x}, \lambda) \leq 0 \text{ for all } i \in \mathcal{I}\}$. Equivalently, using Remark 4.5, we have $\tilde{\mathcal{X}} = \text{cl}\{(\mathbf{x}, \lambda) : \lambda > 0, \lambda f_i(\mathbf{x}/\lambda) \leq 0 \text{ for all } i \in \mathcal{I}\}$.

4.2. Valid inequalities. A second cone that is naturally associated with a convex set is the cone of its valid inequalities. This will play an important role in the analysis of our MICP in Section 7. We report here a formal definition and a useful property.

DEFINITION 4.8. *We define the cone of valid inequalities of a set $\mathcal{X} \subseteq \mathbb{R}^n$ as*

$$\mathcal{X}^\circ := \{(\mathbf{a}, b) : \mathbf{a}^\top \mathbf{x} + b \geq 0 \text{ for all } \mathbf{x} \in \mathcal{X}\}.$$

The cone \mathcal{X}° is easily seen to be closed and convex, even when \mathcal{X} is neither closed nor convex. Note also that the cone of valid inequalities is closely related to the *polar set*, but the latter lives in n dimensions.

The next lemma relates the two operations defined in this section. Recall that the *dual cone* of a closed convex cone \mathcal{K} is the set $\mathcal{K}^* := \{\mathbf{a} : \mathbf{a}^\top \mathbf{x} \geq 0 \text{ for all } \mathbf{x} \in \mathcal{K}\}$.

LEMMA 4.9. *Let \mathcal{X} be a closed convex set. The closed convex cones $\tilde{\mathcal{X}}$ and \mathcal{X}° are dual to each other.*

²More precisely the sets $\text{epi } \tilde{f}$ and $\widetilde{\text{epi } f}$ are isomorphic: $\text{epi } \tilde{f} := \{(\mathbf{x}, \lambda, \sigma) : (\mathbf{x}, \sigma, \lambda) \in \widetilde{\text{epi } f}\}$.

Proof. The perspective cone $\tilde{\mathcal{X}}$ can be equivalently defined as the closure of the cone generated by $\mathcal{X} \times \{1\}$. By applying [42, Corollary 11.7.2] to the latter set, we obtain $\tilde{\mathcal{X}} = \{(\mathbf{x}, \lambda) : \mathbf{a}^\top \mathbf{x} + b\lambda \geq 0 \text{ for all } (\mathbf{a}, b) \in \mathcal{X}^\circ\}$. This shows that $\tilde{\mathcal{X}} = (\mathcal{X}^\circ)^*$. The other direction follows from the bipolar theorem [42, Theorem 14.1]. \square

5. Mixed-integer convex formulation. We now present the main contribution of this paper: the formulation of the SPP in GCS (2.1) as a strong and lightweight MICP. This program is designed in two steps. First, in Section 5.2, we extend the network-flow formulation of the classical SPP (recalled in Section 5.1) to our setting. This yields an optimization problem with bilinear equality constraints. Second, in Section 5.3, we construct a convex relaxation tailored to these bilinear constraints and we formulate our MICP. The relaxation technique used in this section will be described at a higher level of generality and thoroughly analyzed in Section 7.

5.1. Network-flow formulation of the SPP. The starting point for the design of our MICP is the network-flow formulation of the SPP with scalar nonnegative edge lengths (see, e.g., [2, Section 4.1]):

$$\begin{aligned}
 (5.1a) \quad & \text{minimize} && \sum_{e \in \mathcal{E}} l_e y_e \\
 (5.1b) \quad & \text{subject to} && \sum_{e \in \mathcal{E}_s^{\text{out}}} y_e = 1, \quad \sum_{e \in \mathcal{E}_t^{\text{in}}} y_e = 1, \\
 (5.1c) \quad & && \sum_{e \in \mathcal{E}_v^{\text{in}}} y_e = \sum_{e \in \mathcal{E}_v^{\text{out}}} y_e, \quad \sum_{e \in \mathcal{E}_v^{\text{out}}} y_e \leq 1, \quad \forall v \in \mathcal{V} - \{s, t\}, \\
 (5.1d) \quad & && y_e \geq 0, \quad \forall e \in \mathcal{E}.
 \end{aligned}$$

In this LP the decision variables y_e parameterize a path p , with $y_e = 1$ if the edge e is traversed by p and $y_e = 0$ otherwise. The scalar $l_e \geq 0$ represents the length of the edge e . The sets $\mathcal{E}_v^{\text{in}} := \{(u, v) \in \mathcal{E}\}$ and $\mathcal{E}_v^{\text{out}} := \{(v, u) \in \mathcal{E}\}$ collect the edges incoming and outgoing vertex v . Without loss of generality, we assume $|\mathcal{E}_s^{\text{in}}| = |\mathcal{E}_t^{\text{out}}| = 0$, i.e., the source and the target have no incoming and outgoing edges, respectively. Interpreting the value of y_e as the *flow* carried by the edge e , constraint (5.1b) asks that one unit of flow is injected in the source and ejected from the target. For all the other vertices, constraint (5.1c) enforces the *flow conservation* and a *degree constraint*. The latter enforces a limit of one to the total flow traversing the vertex.

Remark 5.1. Note that we do not explicitly require the flows y_e to be binary, but we only enforce their nonnegativity in (5.1d). This is because all the basic feasible solutions of the LP (5.1) can be shown to have binary value (see, e.g., [2, Section 11.12]), and the constraints $y_e \in \{0, 1\}$ would not affect the optimal value of this program.

Remark 5.2. Since we assumed the edge lengths l_e to be nonnegative, the degree constraint in (5.1c) is actually redundant for the LP (5.1), as well as for problem (5.2) below. However, as we will see in Section 5.4, this constraint is not redundant for our final MICP. Therefore we include it in our formulation from the start.

5.2. Biconvex formulation. As an intermediate step towards our MICP, we formulate the SPP in GCS as a *biconvex* optimization problem. Specifically, a nonlinear program whose nonconvexity comes only from products between the vertex locations and the flow variables parameterizing a path. Note that this is consistent with the observation from Section 3 that the SPP in GCS simplifies to a convex program if we fix either the vertex locations or the path through the graph.

A natural attempt to extend the LP (5.1) to the SPP in GCS is to proceed as done for other graph problems with neighborhoods [19, 5, 7]: include the vertex locations \mathbf{x}_v among our decision variables, enforce the constraint $\mathbf{x}_v \in \mathcal{X}_v$ for all $v \in \mathcal{V}$, and substitute the addends in the cost (5.1a) with $\ell_e(\mathbf{x}_u, \mathbf{x}_v)y_e$. However, one immediate issue with this approach is that the latter product is undefined if $\ell_e(\mathbf{x}_u, \mathbf{x}_v) = \infty$ and $y_e = 0$, while we would like the cost contribution of the edge e to always be zero if $y_e = 0$. Perspective functions give us a convenient and rigorous way to “turn on and off” the length of an edge using the corresponding flow variable.

Let us introduce two auxiliary variables $\mathbf{z}_e := y_e \mathbf{x}_u$ and $\mathbf{z}'_e := y_e \mathbf{x}_v$ per edge $e = (u, v)$, and consider the perspective function $\tilde{\ell}_e(\mathbf{z}_e, \mathbf{z}'_e, y_e)$.³ When the flow y_e is positive, this function coincides with the product above:

$$\tilde{\ell}_e(\mathbf{z}_e, \mathbf{z}'_e, y_e) = \ell_e(\mathbf{z}_e/y_e, \mathbf{z}'_e/y_e)y_e = \ell_e(y_e \mathbf{x}_u/y_e, y_e \mathbf{x}_v/y_e)y_e = \ell_e(\mathbf{x}_u, \mathbf{x}_v)y_e,$$

where the first equality comes from Remark 4.5. When the flow y_e is zero, the function $\tilde{\ell}_e$ is well defined and correctly evaluates to zero, even when $\ell_e(\mathbf{x}_u, \mathbf{x}_v) = \infty$. In fact, $y_e = 0$ implies $\mathbf{z}_e = \mathbf{z}'_e = \mathbf{0}$, and $\tilde{\ell}_e(\mathbf{0}, \mathbf{0}, 0) = 0$ as discussed in Remark 4.5.

Overall, we then have the following biconvex formulation of the SPP in GCS:

(5.2a)	minimize	$\sum_{e \in \mathcal{E}} \tilde{\ell}_e(\mathbf{z}_e, \mathbf{z}'_e, y_e)$	
(5.2b)	subject to	constraints of problem (5.1),	
(5.2c)		$\mathbf{x}_v \in \mathcal{X}_v,$	$\forall v \in \mathcal{V},$
(5.2d)		$\mathbf{z}_e = y_e \mathbf{x}_u, \mathbf{z}'_e = y_e \mathbf{x}_v,$	$\forall e = (u, v) \in \mathcal{E}.$

The decision variables are the flows y_e , the vertex positions \mathbf{x}_v , and the auxiliary variables \mathbf{z}_e and \mathbf{z}'_e . The role of the latter is to match the vertices \mathbf{x}_u and \mathbf{x}_v when $y_e = 1$, and collapse to zero when $y_e = 0$. This behavior is driven by the bilinear equality constraints (5.2d), which are the only nonconvexity in our formulation and whose convexification is the focus of the next subsection. Before that, let us formally verify that, as mentioned in Remark 5.1 for the LP (5.1), forcing the flows y_e to be binary does not affect the optimal value of the biconvex program (5.2).

PROPOSITION 5.3. *For any local minimum $L \in \mathbb{R}_{\geq 0}$ of problem (5.2), there exists a feasible point of (5.2) with cost equal to L and such that $y_e \in \{0, 1\}$ for all $e \in \mathcal{E}$.*

Proof. Given a local minimizer of (5.2) with cost L , we fix the vertex positions \mathbf{x}_v . This reduces problem (5.2) to an LP of the form (5.1). The optimal value of this LP must be L , otherwise we would have found a descent direction and our solution of (5.2) would not be locally optimal. Furthermore, because of Remark 5.1, we can assume that the optimal flows of this LP are binary. Paired with the previously fixed variables \mathbf{x}_v , these binary flows yield a feasible solution of (5.2) with cost L . \square

5.3. Convex relaxation of the bilinear constraints. The biconvex program (5.2) is our first formulation of the SPP in GCS that can be tackled numerically. However, the bilinear constraints (5.2d) make this optimization problem challenging to solve, even just locally. In this subsection we show how to reformulate problem (5.2) as a lightweight and strong MICP, that can be reliably solved to global optimality using branch-and-bound algorithms.

³We are slightly abusing notation here: since in Definition 4.4 we defined the perspective of functions with a single argument, to be precise, we should write $\tilde{\ell}_e((\mathbf{z}_e, \mathbf{z}'_e), y_e)$.

The next lemma allows us to construct a tight envelope around the constraints of the biconvex program (5.2) through a small number of perspective cones. In its statement we let $\mathcal{E}_v := \mathcal{E}_v^{\text{in}} \cup \mathcal{E}_v^{\text{out}}$ denote the set of edges incident with vertex $v \in \mathcal{V}$. Recall also that a valid constraint for an optimization problem is a constraint that is verified by all the feasible points.

LEMMA 5.4. *For some vertex $v \in \mathcal{V}$, assume that the linear inequality*

$$(5.3) \quad \sum_{e \in \mathcal{E}_v} c_e y_e + d \geq 0$$

is valid for problem (5.2). Partitioning the summation over \mathcal{E}_v in incoming and outgoing edges, we have that the following convex constraint is also valid for (5.2):

$$(5.4) \quad \left(\sum_{e \in \mathcal{E}_v^{\text{in}}} c_e \mathbf{z}'_e + \sum_{e \in \mathcal{E}_v^{\text{out}}} c_e \mathbf{z}_e + d \mathbf{x}_v, \sum_{e \in \mathcal{E}_v} c_e y_e + d \right) \in \tilde{\mathcal{X}}_v.$$

Proof. Constraint (5.4) requires two conditions to hold. One is (5.3), which is assumed. The other is verified by multiplying both sides of $\mathbf{x}_v \in \mathcal{X}_v$ from (5.2c) by the left-hand side of (5.3), and then using the bilinear constraints (5.2d). \square

Remark 5.5. If the valid constraint (5.3) holds with equality, Lemma 5.4 simply amounts to multiplying this equality by \mathbf{x}_v , and it gives us a valid linear equality of the form $\sum_{e \in \mathcal{E}_v^{\text{in}}} c_e \mathbf{z}'_e + \sum_{e \in \mathcal{E}_v^{\text{out}}} c_e \mathbf{z}_e + d \mathbf{x}_v = \mathbf{0}$.

Remark 5.6. Generating new valid constraints by multiplying existing ones is a standard procedure at the core of many relaxation techniques [36, 44, 31, 28, 39]. Lemma 5.4 will be analyzed at a higher level of generality in Section 7, where its similarities with existing methods will be clearly drawn.

Lemma 5.4 lifts any valid linear constraint on the flows incident with vertex v into a convex constraint that envelops the feasible set of problem (5.2). Our MICP is obtained by applying this lemma to each flow constraint in the LP (5.1), and by replacing the constraints of the biconvex program (5.2) with the envelope resulting from this process. Let us first state our MICP and then prove its equivalence to the SPP in GCS (Theorem 5.7 below):

$$\begin{aligned}
 (5.5a) \quad & \text{minimize} && \sum_{e \in \mathcal{E}} \tilde{\ell}_e(\mathbf{z}_e, \mathbf{z}'_e, y_e) \\
 (5.5b) \quad & \text{subject to} && \sum_{e \in \mathcal{E}_s^{\text{out}}} y_e = 1, \quad \sum_{e \in \mathcal{E}_t^{\text{in}}} y_e = 1, \\
 (5.5c) \quad & && \sum_{e \in \mathcal{E}_v^{\text{out}}} y_e \leq 1, \quad \forall v \in \mathcal{V} - \{s, t\}, \\
 (5.5d) \quad & && \sum_{e \in \mathcal{E}_v^{\text{in}}} (\mathbf{z}'_e, y_e) = \sum_{e \in \mathcal{E}_v^{\text{out}}} (\mathbf{z}_e, y_e), \quad \forall v \in \mathcal{V} - \{s, t\}, \\
 (5.5e) \quad & && (\mathbf{z}_e, y_e) \in \tilde{\mathcal{X}}_u, \quad (\mathbf{z}'_e, y_e) \in \tilde{\mathcal{X}}_v, \quad \forall e = (u, v) \in \mathcal{E}, \\
 (5.5f) \quad & && y_e \in \{0, 1\}, \quad \forall e \in \mathcal{E}.
 \end{aligned}$$

Constraint (5.5d) is obtained as in Remark 5.5 from the flow conservation in (5.1c), and (5.5e) is the result of applying Lemma 5.4 to the nonnegativity constraint (5.1d).

Note that the application of the same technique to the equalities (5.5b) and to the degree constraint in (5.1c) would give us

$$(5.6a) \quad \mathbf{x}_s = \sum_{e \in \mathcal{E}_s^{\text{out}}} \mathbf{z}_e, \quad \mathbf{x}_t = \sum_{e \in \mathcal{E}_t^{\text{in}}} \mathbf{z}'_e,$$

$$(5.6b) \quad \mathbf{x}_v - \sum_{e \in \mathcal{E}_v^{\text{out}}} \mathbf{z}_e \in \left(1 - \sum_{e \in \mathcal{E}_v^{\text{out}}} y_e \right) \mathcal{X}_v, \quad \forall v \in \mathcal{V} - \{s, t\}.$$

However these constraints would be redundant since the vertex positions \mathbf{x}_v for $v \in \mathcal{V}$ do not appear in the rest of problem (5.5). Note also that the combination of (5.6) and (5.5e) implies the constraints $\mathbf{x}_v \in \mathcal{X}_v$ for all $v \in \mathcal{V}$, which would also be redundant for our MICP. The convex relaxation of (5.5) is obtained simply by dropping the integrality constraint (5.5f) (the nonnegativity of the flows y_e is imposed by the cost and also by (5.5e)). Observe that, unlike the biconvex program (5.2), the optimal value of the MICP can decrease if the flows are allowed to be fractional.

THEOREM 5.7. *The MICP (5.5) has optimal value equal to the SPP in GCS (2.1). An optimal path p for problem (2.1) is recovered from the solution of (5.5) through the relation $\mathcal{E}_p := \{e \in \mathcal{E} : y_e = 1\}$. An optimal positioning of the vertices is reconstructed for the source and the target as in (5.6a), and for all the other vertices by letting \mathbf{x}_v be any point such that (5.6b) holds.*

In Section 7.3 we will see that this theorem follows from a simple geometric property of Lemma 5.4. Here we give a direct proof that explicitly illustrates the logic behind our formulation.

Proof. The only flows that can satisfy the constraints in (5.5) are such that the set \mathcal{E}_p defined above describes the vertex-disjoint union of an s - t path and cycles. However, the presence of cycles can be excluded since the edge lengths ℓ_e are non-negative, and traversing a cycle that is disjoint from the main path cannot decrease the cost. Therefore, at optimality, \mathcal{E}_p identifies a path p . For all the edges $e \notin \mathcal{E}_p$, constraint (5.5e) simplifies to $\mathbf{z}_e = \mathbf{z}'_e = \mathbf{0}$, and the corresponding cost addends give $\tilde{\ell}_e(\mathbf{0}, \mathbf{0}, 0) = 0$. For the vertices $v \notin p$, constraint (5.5d) is trivially satisfied and (5.6b) reads $\mathbf{x}_v \in \mathcal{X}_v$. For an edge $e = (u, v)$ along the path p , constraint (5.5e) becomes $\mathbf{z}_e \in \mathcal{X}_u$ and $\mathbf{z}'_e \in \mathcal{X}_v$, and the cost addend is $\tilde{\ell}_e(\mathbf{z}_e, \mathbf{z}'_e, 1) = \ell_e(\mathbf{z}_e, \mathbf{z}'_e)$. Denoting with $f = (v, w) \in \mathcal{E}_p$ the edge after e in the path, the flow conservation (5.5d) reads $\mathbf{z}'_e = \mathbf{z}_f$. Finally, the conditions in (5.6) give us $\mathbf{x}_u = \mathbf{z}_e$ and $\mathbf{x}_v = \mathbf{z}'_e$ for all edges $e = (u, v) \in \mathcal{E}_p$. \square

Remark 5.8. If the sets \mathcal{X}_v are singletons, the SPP in GCS simplifies to the SPP with nonnegative edge lengths. In this case our MICP reduces to the LP (5.1) and its convex relaxation is exact, as discussed in Remark 5.1.

Remark 5.9. For the edge lengths ℓ_e and convex sets \mathcal{X}_v that typically appear in practice, the MICP (5.5) can be solved to global optimality with standard solvers (see the discussion in Section 4.1). However, problem (5.5) can be tackled numerically even when the sets in our GCS are not defined by explicit constraints (e.g., convex inequalities). For example, each set \mathcal{X}_v may be very complex and accessible only through an oracle that, given a point \mathbf{x}_v , either certifies that $\mathbf{x}_v \in \mathcal{X}_v$ or returns a separating hyperplane. In fact, such an oracle is easily adapted to checking membership to the perspective cones in (5.5e), and this black-box access to the problem constraints is sufficient for efficient optimization algorithms like the ellipsoid method [21].

5.4. Degree constraints. In Remark 5.2 we anticipated that, although redundant for the LP (5.1), the degree constraints (5.5c) play an important role in our MICP. This is illustrated in the next example, which shows how the optimal flows from (5.5) can induce cycles if the degree constraints are not enforced.

Example 5.10. Consider a graph with vertices $\mathcal{V} := \{s, 1, 2, t\}$ and edges $\mathcal{E} := \{(s, 1), (1, 2), (2, 1), (1, t)\}$. Define the sets $\mathcal{X}_s := \{-1\}$, $\mathcal{X}_1 := [-1, 1]$, $\mathcal{X}_2 := \{0\}$, and $\mathcal{X}_t := \{1\}$. Let the length of each edge be the Euclidean distance squared (2.3). The optimal value of this SPP in GCS is 2 and the optimal path is $p = (s, 1, t)$. However, if we do not enforce the degree constraints (5.5c), our MICP has optimal value equal to 1 and its optimal flows are $y_e = 1$ for all $e \in \mathcal{E}$, i.e., they induce the cycle $(1, 2, 1)$.

In case of an acyclic graph G , the issues just described do not arise and the degree constraints are redundant for our MICP and its convex relaxation. Nonetheless, we still include them in our formulation since they are computationally light and their explicit presence can trigger the use of specialized generalized-upper-bound branching rules in the solver [10, Section 9.2].

6. Alternative formulations. Multiple alternative MICP formulations of the SPP in GCS can be designed and finding the most effective one is a tradeoff between the size of the program and the tightness of its convex relaxation. The MICP (5.5) is very compact: it has only $O(|\mathcal{E}|)$ binary variables, $O(n|\mathcal{E}|)$ continuous variables, and $O(n(|\mathcal{V}| + |\mathcal{E}|))$ constraints (assuming that the cones in (5.5e) are described by $O(n)$ constraints). In addition, in the experiments in Section 9, we will see that the relaxation of our MICP is typically very tight (although a carefully designed instance in Section 9.4 shows that our relaxation can, in principle, be arbitrarily loose).

A simple alternative way to reformulate the biconvex problem (5.2) as an MICP is to enforce the integrality constraints $y_e \in \{0, 1\}$ for all $e \in \mathcal{E}$, and relax each bilinear constraint (5.2d) independently using a McCormick envelope [36]. With our notation, this amounts to replacing (5.2d) with

$$(6.1) \quad (z_e, y_e) \in \tilde{\mathcal{B}}_u, (z'_e, y_e) \in \tilde{\mathcal{B}}_v, (x_u - z_e, 1 - y_e) \in \tilde{\mathcal{B}}_u, (x_v - z'_e, 1 - y_e) \in \tilde{\mathcal{B}}_v,$$

where, for each $v \in \mathcal{V}$, we let \mathcal{B}_v be an axis-aligned box that contains \mathcal{X}_v . Especially if the convex sets \mathcal{X}_v are defined by many constraints, this MICP is more compact than ours. However, as we will see in Section 9, this formulation has loose convex relaxation and its solution times are generally much larger than with our approach.

At the other end of the spectrum, a variety of stronger but potentially more expensive formulations could be devised. For example, we have found that subtour-elimination constraints like [48, Section 2.2] can tighten the relaxation of our MICP for some classes of problems. Alternatively, we could formulate our MICP by grouping the constraints in (5.2) vertex by vertex, and by computing the convex hull of each group (see Section 7.2 below). We could also use more expensive semidefinite relaxations [31, 28, 39] of the bilinear constraints (5.2d). In our computational experience, the MICP (5.5) represents the best compromise between a lightweight and a strong formulation, and its solution times are lower than any other formulation we have tested.

7. Analysis of the mixed-integer formulation. In this section we describe and analyze at a more abstract level the method used in Section 5.3 to formulate the SPP in GSC as an MICP. We show that Lemma 5.4 can be used to design convex relaxations of a large class of bilinear constraints, and we connect this result to existing

relaxation techniques for nonconvex optimization. Finally, we give a simpler geometric proof of the validity of our MICP (already shown in Theorem 5.7).

7.1. Set-based relaxation of bilinear constraints. Our first step in this analysis is to show that Lemma 5.4 is, in fact, a general-purpose relaxation technique for nonconvex sets of the form

$$(7.1) \quad \mathcal{S} := \{(\mathbf{x}, \mathbf{y}, \mathbf{Z}) : \mathbf{x} \in \mathcal{X}, \mathbf{y} \in \mathcal{Y}, \mathbf{Z} = \mathbf{x}\mathbf{y}^\top\},$$

where $\mathcal{X} \subseteq \mathbb{R}^n$ and $\mathcal{Y} \subseteq \mathbb{R}^m$ are closed convex sets. In particular, here \mathcal{X} takes the place of a generic set \mathcal{X}_v in our GCS, while \mathcal{Y} plays the role of the linear constraints on the flow variables incident with vertex v (see Remark 7.2 below for more details).

A natural approach to construct a convex envelope around the set \mathcal{S} is to multiply all the valid inequalities $\mathbf{a}^\top \mathbf{x} + b \geq 0$ for the set \mathcal{X} by all the valid inequalities $\mathbf{c}^\top \mathbf{y} + d \geq 0$ for the set \mathcal{Y} , and then use the bilinear equality $\mathbf{Z} = \mathbf{x}\mathbf{y}^\top$ to linearize these products. This gives us an infinite family of valid linear inequalities for \mathcal{S} , which form our convex relaxation:

$$(7.2) \quad \mathcal{S} \subseteq \mathcal{S}' := \{(\mathbf{x}, \mathbf{y}, \mathbf{Z}) : \mathbf{a}^\top \mathbf{Z} \mathbf{c} + d \mathbf{a}^\top \mathbf{x} + b \mathbf{c}^\top \mathbf{y} + bd \geq 0 \\ \text{for all } (\mathbf{a}, b) \in \mathcal{X}^\circ \text{ and } (\mathbf{c}, d) \in \mathcal{Y}^\circ\}.$$

Note that the conditions $\mathbf{x} \in \mathcal{X}$ and $\mathbf{y} \in \mathcal{Y}$ are implied by the inequalities in (7.2) that correspond to $(\mathbf{0}, 1) \in \mathcal{Y}^\circ$ and $(\mathbf{0}, 1) \in \mathcal{X}^\circ$, respectively.

The relaxation (7.2) is not obviously implementable on a computer, since it involves an infinite number of constraints. However, if one of the two sets is a polytope (i.e., a bounded polyhedron) then the convex set \mathcal{S}' can be efficiently described by a finite number of perspective-cone constraints.

PROPOSITION 7.1. *Let \mathcal{Y} be a polytope with halfspace representation $\{\mathbf{y} : \mathbf{c}_i^\top \mathbf{y} + d_i \geq 0 \text{ for all } i \in \mathcal{I}\}$. We have*

$$(7.3) \quad \mathcal{S}' = \{(\mathbf{x}, \mathbf{y}, \mathbf{Z}) : (\mathbf{Z} \mathbf{c}_i + d_i \mathbf{x}, \mathbf{c}_i^\top \mathbf{y} + d_i) \in \tilde{\mathcal{X}} \text{ for all } i \in \mathcal{I}\}.$$

Proof. To recover (7.2) from (7.3) we first use Lemma 4.9 to rewrite the membership to $\tilde{\mathcal{X}}$ as $\mathbf{a}^\top (\mathbf{Z} \mathbf{c}_i + d_i \mathbf{x}) + b (\mathbf{c}_i^\top \mathbf{y} + d_i) \geq 0$ for all $(\mathbf{a}, b) \in \mathcal{X}^\circ$. Then we notice that listing only the valid inequalities (\mathbf{c}_i, d_i) for $i \in \mathcal{I}$ is equivalent to listing all the valid inequalities $(\mathbf{c}, d) \in \mathcal{Y}^\circ$. In fact, since \mathcal{Y} is bounded, any vector $(\mathbf{c}, d) \in \mathcal{Y}^\circ$ can be expressed as $\sum_{i \in \mathcal{I}} \alpha_i (\mathbf{c}_i, d_i)$ for some nonnegative coefficients α_i . Using these coefficients, the inequality $\mathbf{a}^\top \mathbf{Z} \mathbf{c} + d \mathbf{a}^\top \mathbf{x} + b \mathbf{c}^\top \mathbf{y} + bd \geq 0$ is seen to be implied by the inequalities generated by (\mathbf{c}_i, d_i) for $i \in \mathcal{I}$. \square

We then have two descriptions of the relaxation \mathcal{S}' : the symmetric one (7.2) that clearly exposes the logic behind the technique, and the asymmetric one (7.3) that is computationally efficient, provided that one of the two sets is polytopic and has a small number of facets. The asymmetric relaxation generalizes Lemma 5.4, with $\mathbf{c}_i^\top \mathbf{y} + d_i \geq 0$ taking the place of the flow inequality (5.3). The asymmetric relaxation is also set based, in the sense that it does not rely on the explicit constraints defining \mathcal{X} , but it works directly with its abstract set representation. Besides making the analysis very concise, this has also the practical advantages discussed in Remark 5.9.

Remark 7.2. The constraints of the biconvex program (5.2) can be restated in terms of the set \mathcal{S} as follows. First, we collect in the vector $\mathbf{y}_v := (y_e)_{e \in \mathcal{E}_v}$ the flows incident with vertex v . Second, we let \mathcal{Y}_v be the polytope defined by the linear

constraints acting on \mathbf{y}_v . Constraint (5.1b) and the flow nonnegativity (5.1d) make \mathcal{Y}_s and \mathcal{Y}_t unit simplices (recall that $|\mathcal{E}_s^{\text{in}}| = |\mathcal{E}_t^{\text{out}}| = 0$). For $v \neq s, t$, the polytope \mathcal{Y}_v is defined by the flow nonnegativity (5.1d) together with the conservation and degree constraints in (5.1c). Third, we stack in the columns of the matrix \mathbf{Z}_v the auxiliary variables \mathbf{z}'_e for $e \in \mathcal{E}_v^{\text{in}}$ and \mathbf{z}_e for $e \in \mathcal{E}_v^{\text{out}}$, so that the bilinear constraints (5.2d) take the form $\mathbf{Z}_v = \mathbf{x}_v \mathbf{y}_v^\top$. By defining the sets \mathcal{S}_v as in (7.1), the constraints of problem (5.2) become $(\mathbf{x}_v, \mathbf{y}_v, \mathbf{Z}_v) \in \mathcal{S}_v$ for all $v \in \mathcal{V}$. Our relaxation of the SPP in GCS is then obtained by replacing the constraint sets \mathcal{S}_v with \mathcal{S}'_v defined as in (7.3).

7.2. Tightness of the relaxation \mathcal{S}' . Ideally, we would like our relaxation to be as tight as possible, and the set \mathcal{S}' to coincide with the convex hull of \mathcal{S} . This equality holds, for example, when \mathcal{X} and \mathcal{Y} are intervals on the real line, in which case \mathcal{S}' simplifies to the McCormick envelope [36]. However, the inclusion $\text{conv } \mathcal{S} \subset \mathcal{S}'$ can be strict in general. In fact, for polytopic sets \mathcal{X} and \mathcal{Y} , our approach of multiplying valid inequalities simplifies to the first level of the Reformulation-Linearization Technique (RLT) [44], which does not yield the convex hull of \mathcal{S} if, e.g., $\mathcal{X} := \mathcal{Y} := [0, 1]^2$.

The convex hull of \mathcal{S} can be efficiently described when \mathcal{Y} is a polytope with a small number of extreme points $\{\hat{\mathbf{y}}_j\}_{j \in \mathcal{J}}$. Specifically, by using disjunctive-programming techniques [9], it can be verified that

$$(7.4) \quad \text{conv } \mathcal{S} = \left\{ \sum_{j \in \mathcal{J}} (\mathbf{x}_j, \lambda_j \hat{\mathbf{y}}_j, \mathbf{x}_j \hat{\mathbf{y}}_j^\top) : \sum_{j \in \mathcal{J}} \lambda_j = 1, (\mathbf{x}_j, \lambda_j) \in \tilde{\mathcal{X}} \text{ for all } j \in \mathcal{J} \right\}.$$

Note that this (lifted) description is convex and also set based. While our relaxation \mathcal{S}' has size proportional to the number $|\mathcal{I}|$ of facets of \mathcal{Y} , this description of the convex hull has size proportional to the number $|\mathcal{J}|$ of extreme points of \mathcal{Y} . For the SPP in GCS, the polytopes \mathcal{Y}_v have $O(|\mathcal{E}_v^{\text{in}}| + |\mathcal{E}_v^{\text{out}}|)$ facets and only $O(|\mathcal{E}_v^{\text{in}}| |\mathcal{E}_v^{\text{out}}|)$ extreme points, and this difference can be relatively small if the graph is sparse. However, in our experience the MICPs obtained with our method provide a better tradeoff between strength and size, and are typically much faster to solve.

Remark 7.3. That our relaxation \mathcal{S}' is not always the convex hull of \mathcal{S} should be fully expected. In fact, for $\mathcal{X} := [0, 1]^n$ and $\mathcal{Y} := [0, 1]^m$, the bilinear program

$$(7.5) \quad \text{minimize } \mathbf{p}^\top \mathbf{x} + \mathbf{q}^\top \mathbf{y} + \mathbf{x}^\top \mathbf{R} \mathbf{y} \quad \text{subject to } \mathbf{x} \in \mathcal{X}, \mathbf{y} \in \mathcal{Y},$$

is NP-hard [41], and equivalent to minimizing a linear function over \mathcal{S} . The equality $\mathcal{S}' = \text{conv } \mathcal{S}$ would then allow us to solve an NP-hard problem in polynomial time.

7.3. Geometric proof of Theorem 5.7. In the proof of Theorem 5.7 we have shown the correctness of the MICP (5.5) by analyzing all the feasible values that the variables in this program can take. We now present a simple property of the relaxation \mathcal{S}' that will lead to a geometric and more concise proof of Theorem 5.7. This result will also generalize a known property of RLT.

LEMMA 7.4. *Let \mathcal{Y} be a polytope and $\hat{\mathbf{y}}$ one of its extreme points. We have $(\mathbf{x}, \hat{\mathbf{y}}, \mathbf{Z}) \in \mathcal{S}$ if and only if $(\mathbf{x}, \hat{\mathbf{y}}, \mathbf{Z}) \in \mathcal{S}'$.*

Proof. One direction follows from $\mathcal{S} \subseteq \mathcal{S}'$. For the other direction we show that if $(\mathbf{x}, \hat{\mathbf{y}}, \mathbf{Z}) \in \mathcal{S}'$ then $\mathbf{Z} = \mathbf{x} \hat{\mathbf{y}}^\top$. Since $\hat{\mathbf{y}}$ is an extreme point of \mathcal{Y} , there are m linearly independent inequalities that are active at $\hat{\mathbf{y}}$. Let $\mathbf{C} \in \mathbb{R}^{m \times m}$ and $\mathbf{d} \in \mathbb{R}^m$ collect the coefficients (\mathbf{c}_i, d_i) of these inequalities, so that $\mathbf{C} \hat{\mathbf{y}} + \mathbf{d} = \mathbf{0}$. For the same inequalities, the constraints in (7.3) give us $\mathbf{Z} \mathbf{c}_i + d_i \mathbf{x} = \mathbf{0}$ or, equivalently, $\mathbf{Z} \mathbf{C}^\top + \mathbf{x} \mathbf{d}^\top = \mathbf{0}$. We then have $\mathbf{Z} \mathbf{C}^\top = \mathbf{x} \hat{\mathbf{y}}^\top \mathbf{C}^\top$ and, since \mathbf{C} is invertible, $\mathbf{Z} = \mathbf{x} \hat{\mathbf{y}}^\top$. \square

Alternative proof of Theorem 5.7. As in the first proof of Theorem 5.7, note that the optimal solution of the MICP (5.5) is such that the edges traversed by a unit of flow identify a path p . Note also that, for all $v \in \mathcal{V}$, the flow vectors \mathbf{y}_v corresponding to a path p are extreme points of the polytopes \mathcal{Y}_v . Then the validity of the MICP (5.5) follows since our relaxation is exact in these points by Lemma 7.4. \square

Remark 7.5. Consider the bilinear program (7.5) with polytopes \mathcal{X} and \mathcal{Y} , and the additional constraint $\mathbf{y} \in \{0, 1\}^m$. Assuming $\mathcal{Y} \subseteq [0, 1]^m$, the first-level RLT is known to yield a valid mixed-integer linear formulation of this program [1, Theorem 1]. Lemma 7.4 extends this result to generic closed convex sets \mathcal{X} . In fact, $\mathcal{Y} \subseteq [0, 1]^m$ ensures that any vector $\mathbf{y} \in \mathcal{Y} \cap \{0, 1\}^m$ is an extreme point of \mathcal{Y} , and the relaxation \mathcal{S}' is exact in correspondence of these points.

7.4. Related relaxation techniques. The basic idea of generating new valid constraints by multiplying existing ones is classical, and has many incarnations: from the simple McCormick envelope [36] to semidefinite hierarchies for polynomial optimization [28, 39], passing through RLT [44]. Among this family of techniques, the Lovász-Schrijver hierarchy [31] is the closest to ours, since it is set based and includes constraints of the form (7.3); see [31, Theorem 1.6 and Conditions (iii) to (iii'')]. However, this hierarchy focuses on binary optimization and symmetric quadratic maps, and its naive application to the bilinear set \mathcal{S} would produce multiple redundant variables and constraints. Our approach leverages the bilinear structure of the set \mathcal{S} , that emerges naturally in the SPP in GCS, to construct a relaxation \mathcal{S}' that is smaller and as tight as the first level of the Lovász-Schrijver hierarchy, without semidefinite constraints. (As discussed in Section 6, our practical experience is that higher levels of the hierarchy and semidefinite constraints lead to MICPs that, although stronger, are significantly slower to solve.)

If both the sets \mathcal{X} and \mathcal{Y} are polytopes, the convex hull of \mathcal{S} in (7.4) is also a polytope, and its extreme points are $\text{ext } \mathcal{S} = \{(\mathbf{x}, \mathbf{y}, \mathbf{x}\mathbf{y}^\top) : \mathbf{x} \in \text{ext } \mathcal{X}, \mathbf{y} \in \text{ext } \mathcal{Y}\}$. In general, this yields an exponential-size description of $\text{conv } \mathcal{S}$. Nevertheless, if the sets \mathcal{X} and \mathcal{Y} have further special structure then specialized techniques can be applied to efficiently generate additional valid inequalities for $\text{conv } \mathcal{S}$; see, e.g., the techniques developed for network-interdiction problems [11], pooling problems [23], bipartite bilinear programs [12], and bipartite boolean quadratic programs [45].

The recent work [51] shows how perspective functions can be used to allow the multiplication of nonlinear convex constraints in the RLT algorithm. However, the relaxation in that work is not set based, and requires an explicit analysis of all the possible products of basic cone inequalities.

8. Control applications. A main application of the framework presented in this paper is optimal control of discrete-time dynamical systems. In this section we show how two simple control problems can be cast as SPPs in GCS. These examples illustrate some basic modeling techniques that can also be applied to control problems involving more complex discrete decision making.

8.1. Minimum-time control. Consider the linear dynamical system $\mathbf{s}_{\tau+1} = \mathbf{A}\mathbf{s}_\tau + \mathbf{B}\mathbf{a}_\tau$, where $\mathbf{s}_\tau \in \mathbb{R}^g$ and $\mathbf{a}_\tau \in \mathbb{R}^r$ are the system state and control action at time step τ . Given an initial state \mathbf{s}_0 , we look for a sequence of controls that drives the system state to the origin in the minimum number T of time steps. At each time τ , the state and control pair $(\mathbf{s}_\tau, \mathbf{a}_\tau)$ is constrained in a compact convex set \mathcal{D} .

To formulate this problem as an SPP in GCS we proceed as in Figure 2a. The vertices \mathcal{V} in our graph are ordered in a sequence. The source s is the first vertex




FIG. 2. Graphs for the formulation of the optimal-control problems in Section 8 as SPPs in GCS.

and the target t is the last. The number of vertices is equal to $\bar{T} + 1$, where \bar{T} is a given upper bound on the optimal time horizon T . Each vertex that is not the target has two outgoing edges: one that connects it to the next vertex in the sequence and one that goes to the target. For each $v \in \mathcal{V}$, the continuous variable \mathbf{x}_v represents a state and control pair $(\mathbf{s}_v, \mathbf{a}_v)$. These variables are constrained by the following sets: $\mathcal{X}_s := \mathcal{D} \cap (\{\mathbf{s}_0\} \times \mathbb{R}^r)$ for the source, $\mathcal{X}_t := \{(\mathbf{0}, \mathbf{0})\}$ for the target (the value of \mathbf{a}_t is actually irrelevant), and $\mathcal{X}_v := \mathcal{D}$ for all the other vertices. To minimize the number of edges in the optimal path (i.e., the time steps to reach the origin), the length of each edge (u, v) is 1 if $\mathbf{s}_v = \mathbf{A}\mathbf{s}_u + \mathbf{B}\mathbf{a}_u$ and infinite otherwise. (See Example 4.6 for the perspective of such a function.)

The solution of the MICP (5.5) gives us a path $p := (v_0, \dots, v_K)$. The optimal time horizon is $T := K$, and the corresponding control sequence is $\mathbf{a}_\tau := \mathbf{a}_{v_\tau}$ for $\tau = 0, \dots, T-1$. The state trajectory is retrieved similarly, and is such that $\mathbf{s}_T := \mathbf{s}_t = \mathbf{0}$.

8.2. Control of hybrid systems. PieceWise-Affine (PWA) systems are a popular framework for modeling hybrid dynamics. Loosely speaking, almost any dynamical system whose nonlinearity is exclusively due to discrete logics can be written in PWA form [24]. Among the many applications of PWA systems, we have automotive [6], power electronics [20], and robotics [33]. Given a finite collection $\{\mathcal{D}_\nu\}_{\nu \in \mathcal{N}}$ of compact convex subsets of the state and control space, a PWA system has dynamics $\mathbf{s}_{\tau+1} = \mathbf{A}_{\nu_\tau} \mathbf{s}_\tau + \mathbf{B}_{\nu_\tau} \mathbf{a}_\tau + \mathbf{c}_{\nu_\tau}$ if $(\mathbf{s}_\tau, \mathbf{a}_\tau) \in \mathcal{D}_{\nu_\tau}$. The index $\nu_\tau \in \mathcal{N}$ represents the system discrete *mode* at time τ , which is itself a decision variable. We consider the problem of driving a PWA system from a given initial state \mathbf{s}_0 to the origin, in a fixed number T of time steps. The objective is to minimize the sum of the stage costs $\gamma(\mathbf{s}_\tau, \mathbf{a}_\tau)$ for $\tau = 0, \dots, T-1$. The function γ is convex and finite.

We model this problem through the GCS in Figure 2b. The source s is the leftmost vertex and the target t is the rightmost. In between, we have T layers with $|\mathcal{N}|$ vertices each. The source is connected via an edge to each vertex in the first layer, and all the vertices in the last layer are connected to the target. Each pair of consecutive layers is fully connected. Also in this case the continuous variables \mathbf{x}_v represent state and control pairs $(\mathbf{s}_v, \mathbf{a}_v)$. The source is paired with the set $\mathcal{X}_s := \{(\mathbf{s}_0, \mathbf{0})\}$, the target with $\mathcal{X}_t := \{(\mathbf{0}, \mathbf{0})\}$, and the ν th vertex v of each layer with $\mathcal{X}_v := \mathcal{D}_\nu$. To enforce the initial conditions, the edges (s, v) outgoing from the source have zero length if $\mathbf{s}_v = \mathbf{s}_s$, and infinite length otherwise. (Note that here the values of both \mathbf{a}_s and \mathbf{a}_t are irrelevant.) The length of any other edge (u, v) , where u is the ν th vertex in its layer, is $\gamma(\mathbf{s}_u, \mathbf{a}_u)$ if $\mathbf{s}_v = \mathbf{A}_\nu \mathbf{s}_u + \mathbf{B}_\nu \mathbf{a}_u + \mathbf{c}_\nu$ and infinite otherwise.

A shortest path $p := (v_0, \dots, v_K)$ has now $T + 2$ vertices. The optimal control at

time $\tau = 0, \dots, T - 1$ is $\mathbf{a}_\tau := \mathbf{a}_{v_{\tau+1}}$. The state trajectory is defined similarly.

Remark 8.1. Frequently in optimal control we need to enforce convex terminal constraints of the form $\mathbf{s}_T \in \mathcal{D}_T$, as well as convex terminal penalties $\gamma_T(\mathbf{s}_T)$. These are easily incorporated in our construction through a suitable modification of the set \mathcal{X}_t and the lengths of the edges incoming to the target vertex.

Remark 8.2. The size of the GCS we just constructed is linear in the time horizon T and quadratic in the number $|\mathcal{N}|$ of discrete modes. Conversely, common formulations for these problems have size linear in both T and $|\mathcal{N}|$ [34]. We will see in Section 9.3 that the greater strength of our MICPs can be well worth this price.

9. Numerical results. This section collects multiple numerical experiments. We start in Section 9.1 with a simple two-dimensional problem. Section 9.2 presents a statistical analysis of the performance of our MICP on large-scale instances of the SPP in GCS. In Section 9.3 we compare our approach with state-of-the-art mixed-integer formulations for control. Finally, in Section 9.4 we use a carefully designed problem to show how symmetries in the GCS can loosen the relaxation of our MICP.

The code necessary to reproduce these results is available at <https://github.com/TobiaMarcucci/shortest-paths-in-graphs-of-convex-sets>. All the experiments are run using the commercial solver MOSEK 10.0 with default options on a laptop computer with processor 2.4 GHz 8-Core Intel Core i9 and memory 64 GB 2667 MHz DDR4. A mature implementation of the techniques presented in this paper is also provided by the open-source software Drake [49].

9.1. Two-dimensional example. We consider the two-dimensional problem in Figure 3a. We have a graph G with $|\mathcal{V}| = 9$ vertices, $|\mathcal{E}| = 22$ edges, and multiple cycles. The source $\mathcal{X}_s := \{\boldsymbol{\theta}_s\}$ and target $\mathcal{X}_t := \{\boldsymbol{\theta}_t\}$ sets are single points, while the remaining regions are full dimensional. The geometry of the sets \mathcal{X}_v and the edge set \mathcal{E} can be deduced from Figure 3a. As edge lengths we consider the Euclidean distance (2.2) and the Euclidean distance squared (2.3), whose corresponding shortest paths are shown in Figure 3a in orange and blue. As expected, the first path is almost straight, while the lengths of the segments in the second are better balanced.

In Figure 3b we compare the optimal values of the SPP in GCS, the relaxation of our MICP (5.5), and the relaxation of the McCormick formulation (6.1). Both relaxations are Second-Order-Cone Program (SOCPs), and for the McCormick one the bounding boxes \mathcal{B}_v are chosen as small as the corresponding sets \mathcal{X}_v allow. We run this comparison for different values of a parameter $\sigma > 0$ that controls the volume of the sets \mathcal{X}_v . The value $\sigma = 1$ corresponds to the GCS in Figure 3a. While for $\sigma \neq 1$ each set \mathcal{X}_v is shrunk or enlarged via a uniform scaling, with scale factor σ , relative to a fixed Chebyshev center of the set.

When the edge length is the Euclidean distance (2.2), the top panel in Figure 3b shows that our relaxation is exact for all values of σ . This was expected for σ close to zero, since by Remark 5.8 our relaxation is exact when the sets are singletons. Similarly, the problem is trivial for very large σ , when the regions are so big that, no matter the discrete path we take, we can always reach the target via a straight line. However, that our relaxation is exact for all the intermediate values of σ is not an obvious result. The McCormick relaxation is also exact for small σ , but gives a trivial lower bound of zero when the sets are large.

With the Euclidean length squared (2.3), both relaxations are still guaranteed to be tight as σ goes to zero. This is confirmed by the bottom panel of Figure 3b. When σ is very large, we have seen in Section 3 that our problem is equivalent to the HPP, and




FIG. 3. Two-dimensional SPP in GCS from Section 9.1. The tightness of the convex relaxation of our MICP (5.5) is analyzed for two edge lengths (the Euclidean distance (2.2) and the Euclidean distance squared (2.3)) and different sizes of the sets \mathcal{X}_v (parameterized by the scalar σ). As a baseline, we also report the optimal value of the relaxation of the McCormick formulation (6.1).

the argument from Theorem 3.1 shows that its optimal value is $\|\theta_t - \theta_s\|_2^2 / K = 11.6$, where $K = 7$ is the number of edges in the longest s - t path in the graph in Figure 3a. A close inspection of the bottom of Figure 3b reveals that, for large σ , our relaxation yields the lower bound $\|\theta_t - \theta_s\|_2^2 / (|\mathcal{V}| - 1) = 10.1$, which corresponds to the simple inequality $K \leq |\mathcal{V}| - 1$. (Using a duality argument, it can be verified that our relaxation always recovers this bound.) Conversely, the lower bound provided by the McCormick relaxation is again equal to zero.

9.2. Large-scale random instances. We present a statistical analysis of the performance of our formulation. We generate a variety of random large-scale SPPs in GCS, and we analyze the relaxation tightness and the solution times of the MICP (5.5) as functions of various problem parameters. We stress that generating random graphs representative of the “typical” SPP in GCS we might encounter in practice is a difficult operation. Inevitably, the instances we describe below are not completely representative, and our algorithm might perform worse or better on other classes of random graphs. Our goal here is to show that our MICP is not limited to small-scale problems.

We construct an SPP in GCS as follows. We set $\mathcal{X}_s := \{\mathbf{0}\}$ and $\mathcal{X}_t := \{\mathbf{1}\}$. The rest of the sets \mathcal{X}_v are axis-aligned cubes with volume Λ and center drawn uniformly at random in $[0, 1]^n$. Given a number $|\mathcal{E}|$ of edges, we construct the edge set in two steps. First we generate multiple s - t paths such that every vertex $v \neq s, t$ is traversed exactly by one path. These are determined via a random partition of the set $\mathcal{V} - \{s, t\}$: the number of sets in the partition (number of paths) is drawn uniformly from the interval $[1, |\mathcal{V}| - 2]$, and also the number of vertices in each set (length of each path) is a uniform random variable. Then we extend the edge set by drawing edges uniformly at random from the set $\{(u, v) \in \mathcal{V}^2 : v \neq s, u \neq t, u \neq v\}$ until a desired cardinality $|\mathcal{E}|$ is reached. As edge lengths we consider the Euclidean distance (2.2) and the Euclidean distance squared (2.3), which both make our formulation (5.5) a mixed-integer SOCP.




FIG. 4. Projection onto two dimensions of a random instance of the SPP in GCS from Section 9.2. The problem parameters have nominal value.




FIG. 5. Relaxation gap versus MICP solution time for the 500 random instances described in Section 9.2. Two edge lengths are analyzed: the Euclidean distance (2.2) and the Euclidean distance squared (2.3). For each edge length, 100 nominal instances are generated with the nominal problem parameters, and four other batches of 100 instances each are obtained by increasing a different subset of the parameters. Our relaxation is almost always exact with the Euclidean length. While, with the Euclidean length squared, it is more sensitive to the dimension n of the space and the density of the graph G . (Note the different horizontal scales of the two plots.)

For each edge length, we first solve 100 random instances with the following nominal parameters: volume $\Lambda = 0.01$, $n = 4$ dimensions, $|\mathcal{V}| = 50$ vertices, and $|\mathcal{E}| = 100$ edges. Then we solve four other batches of 100 problems where, in each batch, a different subset of these parameters is increased by a factor of 5. Specifically, these additional batches test our formulation in case of large sets \mathcal{X}_v (Λ from 0.01 to 0.05), high dimensions (n from 4 to 20), dense graphs ($|\mathcal{E}|$ from 100 to 500), and large graphs ($|\mathcal{V}|$ and $|\mathcal{E}|$ from 50 and 100 to 250 and 500). To give an idea of what these problems look like, the projection onto two dimensions of a GCS generated using the nominal parameters is shown in Figure 4.

Figure 5 shows the relaxation gap (cost gap between the MICP and its relaxation, normalized by the MICP cost) versus the MICP solution time for all the instances described above. As observed in the previous example, the Euclidean edge length (2.2) results in easier programs: our relaxation is tight in almost all the instances and the solution times are relatively low. The squared edge length (2.3) leads to more challenging problems, even though the maximum relaxation gap and runtime are only 2.1% and 0.66s in the nominal case. When the volume of the cubes \mathcal{X}_v is increased to $\Lambda = 0.05$ these values increase to 9.1% and 1.12s, and the performance of our MICP is minimally affected. Note that this is not in contrast with the previous

example, where we analyzed the regime of extremely large sets \mathcal{X}_v . Note also that the volume of the sets does not affect the MICP size. The growth of the space dimension to $n = 20$ increases the size of our programs, and also loosens the relaxation. The largest relaxation gap is 28.9%, and our MICP takes 72s to be solved in the worst case. Similarly, when the number $|\mathcal{E}|$ of edges is increased to 500 the maximum relaxation gap and runtime become 32.9% and 174s. This is due to the combination of the quadratic edge length and the large number of cycles that we have in a graph with high density of edges $|\mathcal{E}|/|\mathcal{V}|$. To show this, in the last batch of problems we keep $|\mathcal{E}| = 500$ and we increase the number of vertices to $|\mathcal{V}| = 250$. This increases the MICP size further but makes the graph sparser, reducing the maximum relaxation gap and runtime to 5.3% and 5.4s.

Also for the problems in this analysis our formulation outperforms the McCormick one in (6.1). With the nominal parameters, the McCormick median (maximum) runtime is 12.9 (4.3) times larger than ours for the Euclidean length (2.2), and 10.3 (2.7) times larger for the Euclidean length squared (2.3). This performance difference grows larger for the other batches of problems, where the McCormick formulation reaches our time limit of one hour very often. The slowness of the McCormick approach is due to its loose relaxation: even with the nominal parameters, we have a median (maximum) relaxation gap of 29% (52%) for (2.2), and 34% (58%) for (2.3).

9.3. Optimal control. We apply the method from Section 8.2 to solve the optimal-control problem shown in Figure 6a. We have a mechanical system with position $\mathbf{q} \in \mathbb{R}^2$, velocity $\mathbf{v} \in \mathbb{R}^2$, and force $\mathbf{a} \in \mathbb{R}^2$. The system has the dynamics of a double integrator: $\mathbf{q}_{\tau+1} = \mathbf{q}_\tau + \mathbf{v}_\tau$ and $\mathbf{v}_{\tau+1} = \mathbf{v}_\tau + \eta \mathbf{a}_\tau$, where η is a scalar parameter that regulates the system controllability. The system state at time τ is $\mathbf{s}_\tau := (\mathbf{q}_\tau, \mathbf{v}_\tau)$. The initial position is $\mathbf{q}_0 := (0.5, -3.5)$ (green plus at the bottom left of Figure 6a), the initial velocity is $\mathbf{v}_0 := \mathbf{0}$. At each time step $\tau = 1, \dots, T - 1$, the position \mathbf{q}_τ must belong to one of the seven regions in Figure 6a, while the velocity and the controls are limited by the constraints $\|\mathbf{v}_\tau\|_\infty \leq 1$ and $\|\mathbf{a}_\tau\|_\infty \leq 1$. The goal is to reach the point $\mathbf{q}_T := (6.5, 3.5)$ (green cross at the top right of Figure 6a) with zero velocity \mathbf{v}_T in $T := 30$ time steps. The cost function is the sum of the stage costs $\gamma(\mathbf{s}_\tau, \mathbf{a}_\tau) := \|\mathbf{v}_\tau\|_2^2/5 + \|\mathbf{a}_\tau\|_2^2$.

We let the parameter η vary between the seven regions. The five regions in the range $-5 \leq q_2 \leq 5$ (light blue in Figure 6a) have $\eta = 1$. While in the other two regions (red in Figure 6a) we make the system more expensive to control by setting $\eta = 0.1$. Since the parameter η varies with the state, the system dynamics is PWA and the control problem falls into the class considered in Section 8.2. The GCS beneath this problem (depicted in Figure 2b) has $|\mathcal{V}| = 212$ vertices and $|\mathcal{E}| = 1435$ edges, and the convex sets \mathcal{X}_v live in \mathbb{R}^6 . Also in this case problem (5.5) is a mixed-integer SOCP.

Figure 6a shows the optimal trajectory $(\mathbf{q}_0, \dots, \mathbf{q}_T)$ (white circles) and the optimal controls $(\mathbf{a}_0, \dots, \mathbf{a}_{T-1})$ (blue arrows). Geometrically, the red regions would be shortcuts to the goal, but the low controllability in these areas makes it too expensive not to fall out of the feasible set. The optimal strategy is then to follow a winding trajectory and incur a cost of 9.37.

As a baseline, we first solve the problem using the state-of-the-art perspective formulation from [37, Section 6] (see also [34, Section 5.2.2]). At each time step τ , this expresses the system state \mathbf{s}_τ as a convex combination of one auxiliary variable \mathbf{s}_τ^ν per region $\nu = 1, \dots, 7$. The control \mathbf{a}_τ is decomposed similarly. When the coefficients b_τ^ν of this combination are required to be binary, the solver is forced to make a hard selection of the region in which the system must be at each time step.




FIG. 6. Control problem from Section 9.3 of driving a dynamical system from start (green plus) to goal (green cross). The light-blue and red regions have high and low controllability, respectively. The optimal positions \mathbf{q}_τ are white circles, the optimal controls \mathbf{a}_τ are blue arrows. The triangles are the auxiliary variables \mathbf{q}_τ^ν whose convex combination yields \mathbf{q}_τ . The opacity of the triangles equals the optimal value of the variables b_τ^ν that serve as weights in this convex combination.

When the coefficients b_τ^ν can be fractional, the system evolves according to a convex combination of the dynamics in each region. Figure 6b illustrates the solution of the convex relaxation of this formulation (which, thanks to a perspective reformulation of the stage cost, is also an SOCP). It reports the position \mathbf{q}_τ , the barely visible controls \mathbf{a}_τ , and the auxiliary copies \mathbf{q}_τ^ν of the position vector. The latter have triangular markers and opacity equal to the value of the indicator b_τ^ν . As it can be seen, this relaxation is insensitive to the arrangement of the regions, and its optimal trajectory heads straight to the goal. Also the indicator variables b_τ^ν are uninformative, and take nonzero value in the regions with low controllability (visible triangles in the red regions). The optimal value of this relaxation is 0.67, which is only 7% of the MICP value (93% relaxation gap). The MICP solution time is 1011s \approx 17min.

The convex relaxation of our formulation is much tighter: its optimal value is 7.46, which is 80% of the MICP value (20% relaxation gap). This has a dramatic effect on computation times that are now reduced to 7.1s. To make a plot comparable to Figure 6b we leverage the structure of our GCS in Figure 2b. The equivalent of the indicator variable b_τ^ν is the total flow traversing the ν th vertex in the τ th layer of the graph. Similarly, the position of the same vertex plays the role of the auxiliary variables ($\mathbf{s}_\tau^\nu, \mathbf{a}_\tau^\nu$), which can then be combined using the coefficients b_τ^ν to get candidate values for the state \mathbf{s}_τ and the control \mathbf{a}_τ . Figure 6c illustrates these values, and shows that the trajectory reconstructed from our relaxation resembles the MICP solution in Figure 6a much more closely. All the markers in the regions with low controllability are now invisible, indicating that our relaxation correctly identifies these as regions of high cost. The visible points \mathbf{q}_τ^ν are clustered along the optimal trajectory of the MICP, suggesting that our relaxation contains detailed information about the optimal path to reach the goal.




FIG. 7. Instance of the SPP in GCS from Section 9.4 that shows how symmetries in the GCS can deteriorate the convex relaxation of our MICP. For the relaxation, the cost contribution of edge e is obtained by multiplying the flow y_e by the distance between \bar{z}_e and \bar{z}'_e . Since only the mean of $\bar{z}'_{(1,3)}$ and $\bar{z}'_{(2,3)}$ is required to match $\bar{z}_{(3,t)}$, the cost is minimized by moving these two points closer to $\bar{z}_{(1,3)}$ and $\bar{z}_{(2,3)}$, respectively.

9.4. Symmetries in the GCS. We conclude by showing how symmetries in the GCS can deteriorate the convex relaxation of our MICP and, in principle, make it arbitrarily loose. We illustrate this through the following carefully designed problem.

We consider the SPP in GCS depicted in Figure 7a. We have an acyclic graph with $|\mathcal{V}| = 5$ vertices and $|\mathcal{E}| = 5$ edges. All the sets \mathcal{X}_v are singletons $\{\theta_v\}$, except for \mathcal{X}_3 which is a full-dimensional rectangle. As an edge length, we use the Euclidean distance (2.2). Solving this problem, we obtain the optimal path $p = (s, 1, 3, t)$ with length 7.4 (the symmetric solution $p = (s, 2, 3, t)$ would also be optimal). The corresponding vertex positions are connected by an orange line in Figure 7a.

Figure 7b illustrates the solution of the relaxation of the MICP (5.5). For each edge e , we connect the optimal points $\bar{z}_e := z_e/y_e$ and $\bar{z}'_e := z'_e/y_e$ with an orange line, labeled in blue with the corresponding flow y_e . Note that, for $y_e > 0$, we have $\tilde{\ell}_e(z_e, z'_e, y_e) = \ell_e(\bar{z}_e, \bar{z}'_e)y_e$, and the vectors \bar{z}_e and \bar{z}'_e are the actual points where the length of the edge e is evaluated. Note also that, by (5.5e), we have $\bar{z}_e \in \mathcal{X}_u$ and $\bar{z}'_e \in \mathcal{X}_v$. The relaxation splits the unit of flow injected in the source into two: half unit is shipped to the target via the top path, the other half via the bottom path. The optimal value of this convex program is 7.0.

The looseness of the relaxation can be explained as follows. If we denote with ρ the flow traversing edge $(1, 3)$, the flow conservation gives $y_{(2,3)} = 1 - \rho$, while the flow through the edge $(3, t)$ is always one. Since the variables $\bar{z}_{(1,3)}$, $\bar{z}_{(2,3)}$, and $\bar{z}'_{(3,t)}$ are forced to match θ_1 , θ_2 , and θ_t , respectively, the cost terms in (5.5a) corresponding to the edges $(1, 3)$, $(2, 3)$, and $(3, t)$ read

$$(9.1) \quad \rho \|\bar{z}'_{(1,3)} - \theta_1\|_2 + (1 - \rho) \|\bar{z}'_{(2,3)} - \theta_2\|_2 + \|\theta_t - \bar{z}_{(3,t)}\|_2.$$

The only constraint that links these variables is (5.5d) for $v = 3$, which gives $\rho \bar{z}'_{(1,3)} + (1 - \rho) \bar{z}'_{(2,3)} = \bar{z}_{(3,t)}$. When $\rho = 1/2$, this constraint asks the mean of $\bar{z}'_{(1,3)}$ and $\bar{z}'_{(2,3)}$ to match $\bar{z}_{(3,t)}$, as opposed to forcing either one of the first two points to match the third, as it would be for $\rho \in \{0, 1\}$. Therefore, while keeping their mean equal to $\bar{z}_{(3,t)}$, the points $\bar{z}'_{(1,3)}$ and $\bar{z}'_{(2,3)}$ can move vertically, and get closer to θ_1 and θ_2 .

This reduces the first two terms in (9.1), and keeps the third term unchanged.

Although this example leads to a relaxation gap of only 5%, a simple variation of it shows that our relaxation can be arbitrarily loose. In particular, if we let $\ell_{(s,1)} := \ell_{(s,2)} := 0$ and we shift the centers of the sets \mathcal{X}_3 and \mathcal{X}_t to the origin, then the cost of the MICP and its relaxation are reduced to 2 and 0, and the relaxation gap becomes 100%. Nevertheless, we emphasize that this is a contrived problem, and the instances we encounter in practice lead to these phenomena very rarely.

10. Conclusions. In this paper we have introduced the SPP in GCS, a versatile generalization of the classical SPP. Our main contribution is a compact MICP formulation for the solution of this NP-hard problem. Numerical experiments show that the convex relaxation of our formulation is typically very tight, and it enables us to quickly solve large problems to global optimality. We have demonstrated the applicability of the proposed framework to control systems: many optimal control problems are interpretable as SPPs in GCS and, in our tests, the proposed formulation outperforms state-of-the-art techniques for their solution.

Acknowledgments. We would like to thank Hongkai Dai for all the time spent improving the solver interface used in the numerical experiments of this paper.

REFERENCES

- [1] W. P. ADAMS AND H. D. SHERALI, *Linearization strategies for a class of zero-one mixed integer programming problems*, Operations Research, 38 (1990), pp. 217–226.
- [2] R. K. AHUJA, T. L. MAGNANTI, AND J. B. ORLIN, *Network Flows: theory, algorithms, and applications*, Prentice-Hall, 1993.
- [3] E. M. ARKIN AND R. HASSIN, *Approximation algorithms for the geometric covering salesman problem*, Discrete Applied Mathematics, 55 (1994), pp. 197–218.
- [4] A. BEMPORAD AND M. MORARI, *Control of systems integrating logic, dynamics, and constraints*, Automatica, 35 (1999), pp. 407–427.
- [5] V. BLANCO, E. FERNÁNDEZ, AND J. PUERTO, *Minimum spanning trees with neighborhoods: Mathematical programming formulations and solution methods*, European Journal of Operational Research, 262 (2017), pp. 863–878.
- [6] F. BORRELLI, A. BEMPORAD, M. FODOR, AND D. HROVAT, *An MPC/hybrid system approach to traction control*, Transactions on Control Systems Technology, 14 (2006), pp. 541–552.
- [7] J. W. BURDICK, A. BOUMAN, AND E. RIMON, *From multi-target sensory coverage to complete sensory coverage: An optimization-based robotic sensory coverage approach*, in International Conference on Robotics and Automation, IEEE, 2021, pp. 10994–11000.
- [8] J. CANNY AND J. REIF, *New lower bound techniques for robot motion planning problems*, in 28th Annual Symposium on Foundations of Computer Science, IEEE, 1987, pp. 49–60.
- [9] S. CERIA AND J. SOARES, *Convex programming for disjunctive convex optimization*, Mathematical Programming, 86 (1999), pp. 595–614.
- [10] M. CONFORTI, G. CORNUÉJOLS, AND G. ZAMBELLI, *Integer programming*, Springer, 2014.
- [11] D. DAVARNIA, J.-P. P. RICHARD, AND M. TAWARMALANI, *Simultaneous convexification of bilinear functions over polytopes with application to network interdiction*, SIAM Journal on Optimization, 27 (2017), pp. 1801–1833.
- [12] S. S. DEY, A. SANTANA, AND Y. WANG, *New SOCP relaxation and branching rule for bipartite bilinear programs*, Optimization and Engineering, 20 (2019), pp. 307–336.
- [13] Y. DISSER, M. MIHALÁK, S. MONTANARI, AND P. WIDMAYER, *Rectilinear shortest path and rectilinear minimum spanning tree with neighborhoods*, in International Symposium on Combinatorial Optimization, Springer, 2014, pp. 208–220.
- [14] M. DROR, A. EFRAT, A. LUBIW, AND J. S. MITCHELL, *Touring a sequence of polygons*, in 35th annual ACM symposium on Theory of computing, 2003, pp. 473–482.
- [15] M. DROR AND M. HAOUARI, *Generalized Steiner problems and other variants*, Journal of Combinatorial Optimization, 4 (2000), pp. 415–436.
- [16] B. EL KHADIR, J. B. LASSERRE, AND V. SINDHWANI, *Piecewise-linear motion planning amidst static, moving, or morphing obstacles*, in International Conference on Robotics and Automation, IEEE, 2021, pp. 7802–7808.

- [17] C. FEREMANS, M. LABBÉ, AND G. LAPORTE, *Generalized network design problems*, European Journal of Operational Research, 148 (2003), pp. 1–13.
- [18] A. FRANGIONI AND C. GENTILE, *Perspective cuts for a class of convex 0–1 mixed integer programs*, Mathematical Programming, 106 (2006), pp. 225–236.
- [19] I. GENTILINI, F. MARGOT, AND K. SHIMADA, *The travelling salesman problem with neighbourhoods: MINLP solution*, Optimization Methods and Software, 28 (2013), pp. 364–378.
- [20] T. GEYER, G. PAPAFOTIU, AND M. MORARI, *Hybrid model predictive control of the step-down DC–DC converter*, Transactions on Control Systems Technology, 16 (2008), pp. 1112–1124.
- [21] M. GRÖTSCHEL, L. LOVÁSZ, AND A. SCHRIJVER, *Geometric algorithms and combinatorial optimization*, Springer Science & Business Media, 2012.
- [22] O. GÜNLÜK AND J. LINDEROTH, *Perspective reformulations of mixed integer nonlinear programs with indicator variables*, Mathematical programming, 124 (2010), pp. 183–205.
- [23] A. GUPTA, S. AHMED, S. S. DEY, AND M. S. CHEON, *Relaxations and discretizations for the pooling problem*, Journal of Global Optimization, 67 (2017), pp. 631–669.
- [24] W. P. HEEMELS, B. DE SCHUTTER, AND A. BEMPORAD, *Equivalence of hybrid dynamical models*, Automatica, 37 (2001), pp. 1085–1091.
- [25] J.-B. HIRIART-URRUTY AND C. LEMARÉCHAL, *Convex analysis and minimization algorithms I: Fundamentals*, vol. 305, Springer science & business media, 2013.
- [26] R. M. KARP, *Reducibility among combinatorial problems*, in Complexity of computer computations, Springer, 1972, pp. 85–103.
- [27] J. KIM AND J. P. HESPANHA, *Discrete approximations to continuous shortest-path: Application to minimum-risk path planning for groups of UAVs*, in International Conference on Decision and Control, vol. 2, IEEE, 2003, pp. 1734–1740.
- [28] J. B. LASSERRE, *Global optimization with polynomials and the problem of moments*, SIAM Journal on Optimization, 11 (2001), pp. 796–817.
- [29] F. LI AND R. KLETTE, *Euclidean shortest paths*, Springer, 2011.
- [30] W.-J. LI, H.-S. J. TSAO, AND O. ULULAR, *The shortest path with at most/nodes in each of the series/parallel clusters*, Networks, 26 (1995), pp. 263–271.
- [31] L. LOVÁSZ AND A. SCHRIJVER, *Cones of matrices and set-functions and 0–1 optimization*, SIAM Journal on Optimization, 1 (1991), pp. 166–190.
- [32] T. LOZANO-PÉREZ AND M. A. WESLEY, *An algorithm for planning collision-free paths among polyhedral obstacles*, Communications of the ACM, 22 (1979), pp. 560–570.
- [33] T. MARCUCCI, R. DEITS, M. GABICINI, A. BICCHI, AND R. TEDRAKE, *Approximate hybrid model predictive control for multi-contact push recovery in complex environments*, in International Conference on Humanoid Robotics, IEEE, 2017, pp. 31–38.
- [34] T. MARCUCCI AND R. TEDRAKE, *Mixed-integer formulations for optimal control of piecewise-affine systems*, in 22nd ACM International Conference on Hybrid Systems: Computation and Control, 2019, pp. 230–239.
- [35] T. MARCUCCI AND R. TEDRAKE, *Warm start of mixed-integer programs for model predictive control of hybrid systems*, Transactions on Automatic Control, 66 (2021), pp. 2433–2448.
- [36] G. P. MCCORMICK, *Computability of global solutions to factorable nonconvex programs: Part i—convex underestimating problems*, Mathematical programming, 10 (1976), pp. 147–175.
- [37] N. MOEHLE AND S. BOYD, *A perspective-based convex relaxation for switched-affine optimal control*, Systems & Control Letters, 86 (2015), pp. 34–40.
- [38] V. V. NAIK AND A. BEMPORAD, *Embedded mixed-integer quadratic optimization using accelerated dual gradient projection*, IFAC-PapersOnLine, 50 (2017), pp. 10723–10728.
- [39] P. A. PARRILO, *Semidefinite programming relaxations for semialgebraic problems*, Mathematical programming, 96 (2003), pp. 293–320.
- [40] P. C. POP, *Generalized network design problems: Modeling and optimization*, vol. 1, Walter de Gruyter, 2012.
- [41] A. P. PUNNEN, P. SRIPRATAK, AND D. KARAPETYAN, *The bipartite unconstrained 0–1 quadratic programming problem: Polynomially solvable cases*, Discrete Applied Mathematics, 193 (2015), pp. 1–10.
- [42] R. T. ROCKAFELLAR, *Convex analysis*, no. 28, Princeton University Press, 1970.
- [43] A. SCHRIJVER, *Combinatorial optimization: polyhedra and efficiency*, vol. 24, Springer Science & Business Media, 2003.
- [44] H. D. SHERALI AND W. P. ADAMS, *A hierarchy of relaxations between the continuous and convex hull representations for zero-one programming problems*, SIAM Journal on Discrete Mathematics, 3 (1990), pp. 411–430.
- [45] P. SRIPRATAK, A. P. PUNNEN, AND T. STEPHEN, *The bipartite boolean quadric polytope*, Discrete Optimization, 44 (2022), p. 100657.
- [46] B. STELLATO, V. V. NAIK, A. BEMPORAD, P. GOULART, AND S. BOYD, *Embedded mixed-*

- integer quadratic optimization using the OSQP solver*, in European Control Conference, IEEE, 2018, pp. 1536–1541.
- [47] R. A. STUBBS AND S. MEHROTRA, *A branch-and-cut method for 0-1 mixed convex programming*, Mathematical Programming, 86 (1999), pp. 515–532.
- [48] L. TACCARI, *Integer programming formulations for the elementary shortest path problem*, European Journal of Operational Research, 252 (2016), pp. 122–130.
- [49] R. TEDRAKE AND THE DRAKE DEVELOPMENT TEAM, *Drake: Model-based design and verification for robotics*, 2019, <https://drake.mit.edu>.
- [50] Y. YANG, M. LIN, J. XU, AND Y. XIE, *Minimum spanning tree with neighborhoods*, in International Conference on Algorithmic Applications in Management, Springer, 2007, pp. 306–316.
- [51] J. ZHEN, D. DE MOOR, AND D. DEN HERTOOG, *An extension of the reformulation-linearization technique to nonlinear optimization*, Available at Optimization Online, (2021).

Appendix A. Other graph problems in GCS. Existing exact algorithms for graph problems with neighborhoods rely on expensive mixed-integer nonconvex optimization [19, 7, 5]. Here we show that, under standard convexity assumptions, the techniques from Section 7 apply beyond the SPP, and yield exact MICP reformulations for a wide variety of graph problems. A thorough numerical evaluation of these novel formulations will be the object of future works.

Given a directed graph $G := (\mathcal{V}, \mathcal{E})$, many combinatorial problems require finding a set of edges $\mathcal{E}^* \subseteq \mathcal{E}$ that is optimal according to a given criterion and given feasibility conditions. Typically, these are formulated as integer linear programs of the form

$$(A.1) \quad \text{minimize} \quad \sum_{e \in \mathcal{E}} l_e y_e \quad \text{subject to} \quad \mathbf{y} \in \mathcal{Y} \cap \{0, 1\}^{|\mathcal{E}|},$$

where $\mathbf{y} := (y_e)_{e \in \mathcal{E}}$. The edge set \mathcal{E}^* is parameterized by the variables y_e as $\mathcal{E}^* = \{e \in \mathcal{E} : y_e = 1\}$, the polyhedron $\mathcal{Y} \subseteq [0, 1]^{|\mathcal{E}|}$ embodies the feasibility conditions, and the cost is a linear function that assigns a weight $l_e \geq 0$ to each edge $e \in \mathcal{E}$.

We extend the graph problem modeled by (A.1) to its version in GCS as done for the SPP. We let the position $\mathbf{x}_v \in \mathbb{R}^n$ of vertex v be a decision variable, constrained in the set \mathcal{X}_v , and we let the length of the edge $e = (u, v)$ be $l_e(\mathbf{x}_u, \mathbf{x}_v)$. The sets \mathcal{X}_v and the functions l_e satisfy the assumptions from Section 2. We define two auxiliary variables $\mathbf{z}_e := y_e \mathbf{x}_u$ and $\mathbf{z}'_e := y_e \mathbf{x}_v$ per edge $e = (u, v)$, and we formulate our graph problem in GCS as in (5.2), with the condition $\mathbf{y} \in \mathcal{Y} \cap \{0, 1\}^{|\mathcal{E}|}$ in place of (5.2b). This yields a mixed-integer program with bilinear constraints. At this point, in the case of the SPP, we grouped the constraints in our problem vertex by vertex, and we applied the relaxation from Lemma 5.4. However, in general, the polyhedron \mathcal{Y} might not enjoy this convenient separability, as it might couple flows that do not share a common vertex. There are two ways around this issue.

One option is just to separate the flow constraints that are vertex-wise separable from the ones that are not. Using only the first to define the polyhedra $\mathcal{Y}_v \subseteq [0, 1]^{|\mathcal{E}_v|}$, we can then proceed as in Remark 7.2. The MICP we get is a valid problem formulation since any point in $\mathcal{Y}_v \cap \{0, 1\}^{|\mathcal{E}_v|}$ is an extreme point of \mathcal{Y}_v and, by Lemma 7.4, our relaxation is exact for those points. The formulation resulting from this approach is compact but it might be weak.

The second option is to introduce new variables that represent the product of each flow y_e and vertex position \mathbf{x}_v , even if edge e is not incident with vertex v . This gives us a total of $n|\mathcal{V}||\mathcal{E}|$ continuous variables $\mathbf{Z} := \mathbf{x}\mathbf{y}^\top$, where $\mathbf{x} := (\mathbf{x}_v)_{v \in \mathcal{V}}$ lives in the Cartesian product $\mathcal{X} := \prod_{v \in \mathcal{V}} \mathcal{X}_v$. Defining the set \mathcal{S} as in (7.1), the constraints of our problem become $\mathbf{y} \in \{0, 1\}^{|\mathcal{E}|}$ and $(\mathbf{x}, \mathbf{y}, \mathbf{Z}) \in \mathcal{S}$. We then use the relaxation \mathcal{S}' of \mathcal{S} from (7.3) to get an MICP whose validity is ensured again by Lemma 7.4. This

second option yields larger but potentially stronger MICPs.

Appendix B. Dual optimization problem. In this appendix we analyze the dual of the convex relaxation of the MICP (5.5), and we draw additional parallels between this problem and the network-flow formulation (5.1) of the SPP.

B.1. Dual of the SPP. As a reference for the discussion below, the dual of the LP (5.1) is

$$\begin{aligned} \text{(B.1a)} \quad & \text{maximize} \quad p_s - p_t \\ \text{(B.1b)} \quad & \text{subject to} \quad p_u - p_v \leq l_e, \quad \forall e = (u, v) \in \mathcal{E}. \end{aligned}$$

Here p_s and p_t are the multipliers of the two constraints in (5.1b), and p_v for $v \neq s, t$ are the multiplier of the flow conservation in (5.1c). These multipliers are interpretable as potentials: the objective asks to maximize the potential jump between source and target, and the constraints ensure that the potential jump along each edge does not exceed the edge length. Since the degree constraints in (5.1c) are redundant (see Remark 5.2), their multipliers do not appear in the dual problem.

For the LPs (5.1) and (B.1), complementary slackness reads $(l_e - p_u + p_v)y_e = 0$ for all edges $e = (u, v)$. Therefore, at optimality, each edge $e \in \mathcal{E}_p$ along the shortest path must have a potential jump equal to its edge length.

B.2. Dual of the SPP in GCS. The convex relaxation of the MICP (5.5) is a conic program, and its dual is derived in the standard way. To make the interpretation of the dual program easier, we assume that the graph G is acyclic, and we remove the degree constraints (5.5c) from the primal. This leads to the following optimization problem:

$$\begin{aligned} \text{(B.2a)} \quad & \text{maximize} \quad p_s - p_t \\ \text{(B.2b)} \quad & \text{subject to} \quad \mathbf{r}_u^\top \mathbf{x}_u + p_u - \mathbf{r}_v^\top \mathbf{x}_v - p_v \leq \ell_e(\mathbf{x}_u, \mathbf{x}_v), \\ & \quad \quad \quad \forall \mathbf{x}_u \in \mathcal{X}_u, \mathbf{x}_v \in \mathcal{X}_v, e = (u, v) \in \mathcal{E}, \\ \text{(B.2c)} \quad & \mathbf{r}_s = \mathbf{r}_t = 0. \end{aligned}$$

The dual variables are p_v and \mathbf{r}_v for all $v \in \mathcal{V}$. The first are paired with the flow constraints as above. The second correspond to the portion of the flow conservation (5.5d) that involves the auxiliary variables \mathbf{z}_e and \mathbf{z}'_e (the additional variables \mathbf{r}_s and \mathbf{r}_t have only the role of simplifying the presentation).

Similarly to the LP (B.1), the dual (B.2) can be interpreted in terms of potentials. For each vertex $v \in \mathcal{V}$, the linear function $\mathbf{r}_v^\top \mathbf{x}_v + p_v$ defines the potential of the point $\mathbf{x}_v \in \mathcal{X}_v$. Because of (B.2c), these functions are constant over the source and target sets, and the objective (B.2a) maximizes the potential jump between s and t as in the classical SPP. Like (B.1b), constraint (B.2b) asks the potential jump along an edge to be smaller than the edge length. By setting all the potential functions to zero, we see that the dual problem is always feasible and has nonnegative optimal value.

For the primal-dual pair (5.5) and (B.2), complementary slackness requires

$$\mathbf{r}_u^\top \mathbf{z}_e + p_u y_e - \mathbf{r}_v^\top \mathbf{z}'_e - p_v y_e = \tilde{\ell}_e(\mathbf{z}_e, \mathbf{z}'_e, y_e)$$

for all edges $e = (u, v)$. As for the classical SPP, this is trivially satisfied if $y_e = 0$. While, for $y_e > 0$, we get $\mathbf{r}_u^\top \bar{\mathbf{z}}_e + p_u - \mathbf{r}_v^\top \bar{\mathbf{z}}'_e - p_v = \ell_e(\bar{\mathbf{z}}_e, \bar{\mathbf{z}}'_e)$, with $\bar{\mathbf{z}}_e := \mathbf{z}_e/y_e$ and $\bar{\mathbf{z}}'_e := \mathbf{z}'_e/y_e$. In words, at optimality, the potential jump along edge e is tight to the edge length ℓ_e at the point $(\bar{\mathbf{z}}_e, \bar{\mathbf{z}}'_e)$.