

Konrad-Zuse-Zentrum
für Informationstechnik Berlin

ZIB

Takustraße 7
D-14195 Berlin-Dahlem
Germany

ANDREAS BLEY

**Inapproximability Results for the Inverse
Shortest Paths Problem with Integer Lengths and
Unique Shortest Paths**

URL: <http://www.zib.de/PaperWeb/abstracts/ZR-05-04>

ZIB-Report 05-04 (2005)

Inapproximability Results for the Inverse Shortest Paths Problem with Integer Lengths and Unique Shortest Paths

Andreas Bley

Konrad-Zuse-Zentrum

Address : Takustr. 7, 14195 Berlin, Germany

Fax: +49-30-84185269

Email: bley@zib.de

April 2005

Abstract

We study the complexity of two Inverse Shortest Paths (ISP) problems with integer arc lengths and the requirement for uniquely determined shortest paths. Given a collection of paths in a directed graph, the task is to find positive integer arc lengths such that the given paths are uniquely determined shortest paths between their respective terminals. The first problem seeks for arc lengths that minimize the length of the longest of the prescribed paths. In the second problem, the length of the longest arc is to be minimized.

We show that it is \mathcal{NP} -hard to approximate the minimal longest path length within a factor less than $8/7$ or the minimal longest arc length within a factor less than $9/8$. This answers the (previously) open question whether these problems are \mathcal{NP} -hard or not. We also present a simple algorithm that achieves an $\mathcal{O}(|V|)$ -approximation guarantee for both variants.

Both ISP problems arise in the planning of telecommunication networks with shortest path routing protocols. Our results imply that it is \mathcal{NP} -hard to decide whether a given path set can be realized with a real shortest path routing protocol such as OSPF, IS-IS, or RIP.

Keywords: Inverse Shortest Paths, Computational Complexity, Approximation, Shortest Path Routing

Mathematical Subject Classification (2000): 68Q25, 90C60, 90C27, 05C38, 90B18

1 Introduction

Computer and communication networks are often based on shortest path routing. Most Internet domains use shortest path protocols such as OSPF, IS-IS, or RIP for internal routing. With these routing protocols, all end-to-end traffic streams are routed along shortest paths with respect to some administrative link lengths (or routing weights).

A major difficulty in planning such networks is to find link lengths that induce a set of globally efficient end-to-end routing paths. The shortest path routing paradigm enforces rather complicated and subtle interdependencies among the paths that comprise a valid routing. Therefore, traditional planning approaches for shortest path networks work with the link lengths as decision variables. Most of these methods rely on local search, simulated annealing, or Lagrangian relaxation techniques that evaluate the effects of routing length modifications [4, 3, 6, 12, 15, 16, 21, 23]. More evolved planning methods for shortest path networks use the end-to-end routing paths as primary decision variables [2, 5, 18, 24]. They decompose the problem of finding an appropriate routing into the task of finding the routing paths and the task of finding routing lengths that induce these paths. In the master problem, efficient end-to-end routing paths are computed

with mixed-integer linear programming or combinatorial optimization techniques. In a subsequent step, linear programming techniques are used to find routing lengths that induce the computed paths. If there are no compatible lengths for some set of end-to-end paths, then the routing length computation returns one or more linear constraints that are valid for all shortest path sets but violated by the given path set. These constraints are then added to the master problem in order to recompute a different set of end-to-end routing paths.

The main advantage of the decomposition approach is that it permits the use of advanced optimization techniques for multicommodity flow and unsplittable flow problems to compute provenly optimal routing paths. However, it is essential for this approach that the INVERSE SHORTEST PATHS (ISP) problem of finding compatible routing lengths (or generating the invalidity certificate) for the computed paths can be solved efficiently: Given a collection of paths in a directed graph, the ISP problem is to find positive integer arc lengths for which the prescribed paths are uniquely determined shortest paths between their respective terminals.

Note that the uniqueness of the prescribed shortest paths as well as the range of admissible length values are very important concerns for shortest path routing in telecommunications. The routers in a network calculate the shortest paths autonomously. Thus, the lengths must be chosen such that for any pair of nodes *only* the prescribed paths are shortest paths. Otherwise, the traffic flows in the real network might differ from those computed during network planning. For a network that is operated in single path routing mode, this means that each prescribed path must be the uniquely determined shortest path between its terminals.¹ Furthermore, the routing lengths must be small integers that fit into the data format of the routing protocol. For link-state protocols like OSPF or IS-IS, this means that each single link length is bounded by some protocol-dependent constant. With a distance-vector protocol such as RIP, the total length of each routing path is bounded. Real network routing domains typically range from 10 to 1000 nodes. Thus, the limited size of the lengths typically is not relevant for OSPF, which admits link lengths between 1 and $2^{16} - 1$. However, it becomes an important issue for IS-IS, which admits values up to 63 only. For RIP, where the link lengths must be chosen such that the total length of the longest routing path does not exceed 15, the restricted range of admissible lengths is one of the major difficulties in practice. Therefore, it is natural to seek for integer lengths that either minimize the largest link length or the length of the longest path.

Computing shortest paths in a weighted digraph is one of the classical combinatorial optimization problems. Various algorithms are known to solve this problem efficiently. The inverse problem of finding arc lengths that induce a prescribed set of paths or node-to-node distances received only little attention in the mathematical literature.

Several groups studied the INVERSE SHORTEST PATHS problem in the context of data engineering where the task is to reconstruct data from inaccurate measurements or observations. Typical applications are the estimation of the average travel times on road segments from total end-to-end travel times, the recovery of the densities of earth crust layers from observed seismic waves, or the reconstruction of relationship degrees in genetic sequencing [7, 22]. In these settings, the goal is to find edge or arc lengths that match as good as possible with the observed distances and shortest paths. Neither the uniqueness of the observed shortest paths nor the integrality of the lengths are required in these ISP variants. Burton and Toint [9, 10] apply a quadratic programming approach to estimate the arc lengths when the perceived shortest paths are known. For the same problem, Tong and Lam [25] propose a conjugate gradient method. The case with additional upper bounds on the shortest paths' lengths is discussed by Burton et al. [8]. Fekete et al. [14] study the complexity of a problem variant where the lengths shall be reconstructed from observed distances in the network, but the shortest paths are not known. Cai and Li [11] show that the more general INVERSE MATROID INTERSECTION problem can be transformed into a minimum cost flow problem and be solved in strongly polynomial time.

The ISP problem with integer lengths and unique shortest paths has been rarely addressed. Farago et al. [13] study a special case where the given paths are shortest paths w.r.t. the number of

¹For simplicity, we consider only single shortest path routing in this article. If traffic is splitting among equally long shortest paths (ECMP routing), then for every node pair all prescribed paths but only those must be shortest paths. The inapproximability results presented in this paper carry over straightforward to this variant.

edges and where the task is to find lengths such that all these paths are unique shortest paths. Ben-Ameur and Gourdin [1] discuss structural properties of (undirected) path sets where all paths are uniquely determined shortest paths for some common edge lengths. Ben-Ameur and Gourdin also consider the problem of finding integer such lengths that minimize the largest link length. They propose several linear programming based heuristics, but they leave open whether this problem is \mathcal{NP} -hard or not.

In this paper, we study the complexity of the ISP problem with integer lengths and unique shortest paths. In the following section, we formally introduce two variants of this problem that arise in the planning of shortest path networks. In Section 3, we prove that both variants are \mathcal{APX} -hard. More precisely, we show that the minimal longest path length cannot be approximated within a factor less than $8/7$ or the minimal longest arc length within a factor less than $9/8$, unless $\mathcal{P} = \mathcal{NP}$. This implies that it is \mathcal{NP} -hard to decide whether a prescribed set of paths can be realized with any real shortest path routing protocol such as OSPF, IS-IS, PNNI or RIP. In Section 4, we present a simple linear programming based algorithm that achieves an $\mathcal{O}(|V|)$ -approximation guarantee for both ISP variants.

2 Notation and Preliminaries

Let $D = (V, A)$ be a directed graph. For any path P in D and any arc lengths $\lambda = (\lambda_a)_{a \in A} \in \mathbb{R}_+^A$, we denote $\lambda(P) := \sum_{a \in P} \lambda_a$.

Definition 2.1 *The non-negative arc lengths $\lambda \in \mathbb{R}_+^A$ are said to be compatible with a given set \mathcal{Q} of paths, if each path $P \in \mathcal{Q}$ is the uniquely determined shortest path between its terminals with respect to λ . A set \mathcal{Q} of paths is called a Shortest Path System (SPS) if there exist compatible arc lengths $\lambda \in \mathbb{R}_+^A$ for \mathcal{Q} .*

The family of SPSs in a digraph D forms an independence system (or hereditary family), because any subset (including the empty set) of an SPS is an SPS as well.

In this paper, we consider two problems of finding small integer lengths that are compatible with a prescribed path set. In the first variant, we want to minimize the length the longest path. This leads to the following combinatorial optimization problem:

<i>Problem:</i>	MINIMUM PATH-LENGTH INVERSE SHORTEST PATHS (MIN-PATH-ISP)
<i>Instance:</i>	A digraph $D = (V, A)$ and a path set $\mathcal{Q} \subseteq \mathcal{P}$.
<i>Solution:</i>	Integer arc lengths $\lambda \in \mathbb{Z}_+^A$ with $\lambda_a \geq 1$ for all $a \in A$ that are compatible with \mathcal{Q} .
<i>Objective:</i>	$\min \max_{P \in \mathcal{Q}} \lambda(P)$.

In the second variant, we wish the maximum arc length to be minimized:

<i>Problem:</i>	MINIMUM ARC-LENGTH INVERSE SHORTEST PATHS (MIN-ARC-ISP)
<i>Instance:</i>	A digraph $D = (V, A)$ and a path set $\mathcal{Q} \subseteq \mathcal{P}$.
<i>Solution:</i>	Integer arc lengths $\lambda \in \mathbb{Z}_+^A$ with $\lambda_a \geq 1$ for all $a \in A$ that are compatible with \mathcal{Q} .
<i>Objective:</i>	$\min \max_{a \in A} \lambda_a$.

The MIN-PATH-ISP problem corresponds to the task of finding lengths that realize the prescribed paths with a distance-vector protocol in practice. The MIN-ARC-ISP corresponds to the analogous task for a link-state protocol.

Without loss of generality, we can assume in both problems that the given path set \mathcal{Q} is an SPS. Otherwise, there are not even real-valued arc lengths compatible with \mathcal{Q} , and none of the two problems has a solution. Whether a path set \mathcal{Q} is an SPS can be decided in polynomial time by solving a linear inequality system as shown in Section 4.

3 Inapproximability Results

In the first part of this section, we show that the problem MIN-PATH-ISP is \mathcal{APX} -hard.




Figure 1: Nodes V_i and arcs A_i introduced for any $i \in S$.

Theorem 3.1 *It is \mathcal{NP} -hard to approximate MIN-PATH-ISP within a factor of $8/7 - \epsilon$, for any $\epsilon > 0$.*

Proof. We construct a polynomial reduction from the \mathcal{NP} -complete decision problem SET PARTITION(3) to the problem of computing a solution of value strictly less than 8 for MIN-PATH-ISP. SET PARTITION(3) is a restricted variant of the SET PARTITION problem introduced by Karp [19] (also called EXACT SET COVER in Garey and Johnson [17]). Given a finite set S and a collection $\mathcal{C} \subset 2^S$ with $|\{C \in \mathcal{C} \mid i \in C\}| = 3$ for all $i \in S$, the task is to decide whether there exists a subcollection $\mathcal{C}' \subseteq \mathcal{C}$ which forms a partition of S . In SET PARTITION(3), each element occurs in exactly three sets of the collection. It is not difficult to verify that this restricted variant of SET PARTITION remains \mathcal{NP} -complete.

Suppose we are given an instance of SET PARTITION(3) that consists of the set S and the collection $\mathcal{C} = \{C_j \subset S \mid j \in J\}$. For each element $i \in S$, we refer with $J(i) := \{j(i), k(i), l(i)\} = \{j \in J \mid i \in C_j\}$ to the indices of the three sets in \mathcal{C} that contain i . We construct a MIN-PATH-ISP instance consisting of a digraph $D = (V, A)$ and a shortest path system \mathcal{Q} as follows.

For each $i \in S$, we introduce 28 nodes

$$V_i := \{v_i^j, w_i^j, b_{i,1}^j, \dots, b_{i,6}^j \mid j \in J(i)\} \cup \{u_{i,1}, u_{i,2}, \bar{v}_i^{l(i)}, \bar{w}_i^{l(i)}\}.$$

Some of these nodes receive multiple names. For every $i \in S$, we also denote $\bar{v}_i^{j(i)} := v_i^{k(i)}$, $\bar{v}_i^{k(i)} := v_i^{l(i)}$, $\bar{w}_i^{j(i)} := w_i^{k(i)}$, and $\bar{w}_i^{k(i)} := w_i^{l(i)}$. These nodes are connected by two types of arcs. For each single element $i \in S$, we add the arcs

$$\begin{aligned} A_i := & \{(v_i^{j(i)}, u_{i,1}), (u_{i,1}, u_{i,2}), (u_{i,2}, \bar{v}_i^{l(i)}), (\bar{v}_i^{l(i)}, \bar{w}_i^{l(i)})\} \\ & \cup \{(v_i^j, \bar{v}_i^j), (w_i^j, \bar{w}_i^j), (v_i^{j(i)}, w_i^{j(i)}) \mid j \in J(i)\} \\ & \cup \{(b_{i,q}^j, b_{i,q+1}^j) \mid j \in J(i), q = 1, \dots, 5\}. \end{aligned}$$

The subgraph induced by these nodes and arcs is illustrated in Figure 1.

Additionally, we introduce two linking arcs $A_{i,i'}^j := \{(v_i^j, w_{i'}^j), (\bar{v}_i^j, \bar{w}_{i'}^j)\}$ for each pair of elements $i, i' \in S$, $i \neq i'$, that are both contained in some set $C_j \in \mathcal{C}$. Figure 2 illustrates these arcs.




Figure 2: All linking arcs $A_{i_1 i_2}^j$, $A_{i_1 i_3}^j$, and $A_{i_2 i_3}^j$ introduced for some set $C_j = \{i_1, i_2, i_3\} \in \mathcal{C}$ with $j = k(i_1) = k(i_2) = l(i_3)$.

Together, these sets form the simple digraph $D = (V, A)$ with

$$V := \bigcup_{i \in S} V_i \quad \text{and} \quad A := \bigcup_{i \in S} A_i \cup \bigcup_{\substack{(i, i') \in \mathcal{C}_j^2 \text{ with} \\ j \in J \text{ and } i \neq i'}} A_{i, i'}^j .$$

In this digraph, we prescribe various shortest paths. For each single element $i \in S$, we prescribe eight paths

$$\begin{aligned} \mathcal{Q}_i := & \{(v_i^{j(i)}, u_{i,1}, u_{i,2}, \bar{v}_i^{l(i)}), (w_i^{j(i)}, w_i^{k(i)}, w_i^{l(i)}, \bar{w}_i^{l(i)})\} \\ & \cup \{(v_i^j, \bar{v}_i^j, \bar{w}_i^j), (v_i^j, w_i^j, b_{i',1}^j, \dots, b_{i',6}^j) \mid j \in J(i)\} . \end{aligned}$$

For each pair $i, i' \in S$, $i \neq i'$, with $i, i' \in C_j$ for some $C_j \in \mathcal{C}$, we prescribe two paths

$$\mathcal{Q}_{i, i'}^j := \{(v_i^j, \bar{v}_i^j, \bar{w}_{i'}^j), (v_i^j, w_{i'}^j, b_{i',1}^j, \dots, b_{i',6}^j)\} .$$

Figures 3 and 4 illustrate these paths. The set of all prescribed paths is

$$\mathcal{Q} := \bigcup_{i \in S} \mathcal{Q}_i \cup \bigcup_{\substack{(i, i') \in \mathcal{C}_j^2 \text{ with} \\ j \in J \text{ and } i \neq i'}} \mathcal{Q}_{i, i'}^j .$$

It is obvious that this construction is polynomial.

In the first part of this proof, we show that the constructed MIN-PATH-ISP instance is solvable for any given SET PARTITION(3) instance, i.e., \mathcal{Q} is an SPS. Furthermore, we show that the value of any solution of the MIN-PATH-ISP instance is at least 7. Consider the arc lengths $\lambda \in \mathbb{Z}_+^A$ defined as

$$\lambda_a := \begin{cases} 3 & , \text{ for all } a = (w_i^j, \bar{w}_i^j) \text{ with } j \in J \text{ and } i \in C_j, \\ 2 & , \text{ for all } a = (v_i^j, \bar{v}_i^j) \text{ with } j \in J \text{ and } i \in C_j, \text{ and} \\ 1 & , \text{ otherwise.} \end{cases} \quad (1)$$




Figure 3: Prescribed shortest paths Q_i for any $i \in S$.

For each $i \in S$, the paths $P = (v_i^{j(i)}, u_{i,1}, u_{i,2}, \bar{v}_i^{l(i)})$ and $P' = (v_i^{j(i)}, v_i^{k(i)}, v_i^{l(i)}, \bar{v}_i^{l(i)})$ are the *only* $(v_i^{j(i)}, \bar{v}_i^{l(i)})$ -paths in D . Since $\lambda(P) = 3 < 6 = \lambda(P')$ by (1), the path P is the uniquely determined shortest $(v_i^{j(i)}, \bar{v}_i^{l(i)})$ -path. Similarly, for each $j \in J$ and each pair $(i, i') \in C_j^2$, the paths $P = (v_i^j, \bar{v}_i^j, \bar{w}_{i'}^j)$ and $P' = (v_i^j, w_{i'}^j, \bar{w}_{i'}^j)$ are the only $(v_i^j, \bar{w}_{i'}^j)$ -paths in D , and $\lambda(P) = 3 < 4 = \lambda(P')$ by (1). The paths $(w_i^{j(i)}, w_i^{k(i)}, w_i^{l(i)}, \bar{w}_i^{l(i)})$, for $i \in S$, and the paths $(v_i^j, w_{i'}^j, b_{i',1}^j, \dots, b_{i',6}^j)$, for $j \in J$ and $(i, i') \in C_j^2$, even are the solitary paths between their respective terminals. Hence, all paths in \mathcal{Q} are uniquely determined shortest paths w.r.t. λ .

Let $P = (v_i^j, w_{i'}^j, b_{i',1}^j, \dots, b_{i',6}^j)$ for some $j \in J$ and $(i, i') \in C_j^2$. As $P \in \mathcal{Q}$ and $|P| = 7$, we have $\max_{P \in \mathcal{Q}} \lambda(P) \geq 7$ for all arc lengths that are compatible to \mathcal{Q} and satisfy $\lambda_a \geq 1$ for all $a \in A$.

In the second part, we show that the constructed MIN-PATH-ISP instance has a solution of value 7 if the given SET PARTITION(3) instance is solvable. Assume there is a subcollection $\mathcal{C}' \subseteq \mathcal{C}$ that forms a partition of S . We define the arc lengths $\lambda = \lambda(\mathcal{C}') \in \mathbb{Z}_+^A$ as

$$\lambda_a := \begin{cases} 3 & , \text{ for all } a = (w_i^j, \bar{w}_i^j) \text{ with } C_j \in \mathcal{C}' \text{ and } i \in C_j, \\ 2 & , \text{ for all } a = (w_i^j, \bar{w}_i^j) \text{ with } C_j \notin \mathcal{C}' \text{ and } i \in C_j, \\ 2 & , \text{ for all } a = (v_i^j, \bar{v}_i^j) \text{ with } C_j \in \mathcal{C}' \text{ and } i \in C_j, \text{ and} \\ 1 & , \text{ otherwise.} \end{cases} \quad (2)$$

Analogous to the first part, we now verify for each path in \mathcal{Q} that it is a unique shortest path between its terminals and its length is at most 7. For this, we consider each of the four path types in \mathcal{Q} individually.

Let $P = (v_i^{j(i)}, u_{i,1}, u_{i,2}, \bar{v}_i^{l(i)})$ for some $i \in S$. By (2), we have $\lambda_a = 1$ for all $a \in P$, and thus $\lambda(P) = 3$. According to our construction, the only other $(v_i^{j(i)}, \bar{v}_i^{l(i)})$ -path in D is $P' = (v_i^{j(i)}, v_i^{k(i)}, v_i^{l(i)}, \bar{v}_i^{l(i)})$. Since \mathcal{C}' defines a partition, the element i is contained in some set $C_j \in \mathcal{C}'$. By (2), this implies $\lambda_{(v_i^j, \bar{v}_i^j)} = 2$ and, therefore, $\lambda(P') \geq 4$.

Let $P = (w_i^{j(i)}, w_i^{k(i)}, w_i^{l(i)}, \bar{w}_i^{l(i)})$ for $i \in S$. As \mathcal{C}' defines a partition of S , i is contained in exactly one set $C_j \in \mathcal{C}'$. W.l.o.g., suppose this set is $C_{j(i)}$. According to (2), we then have

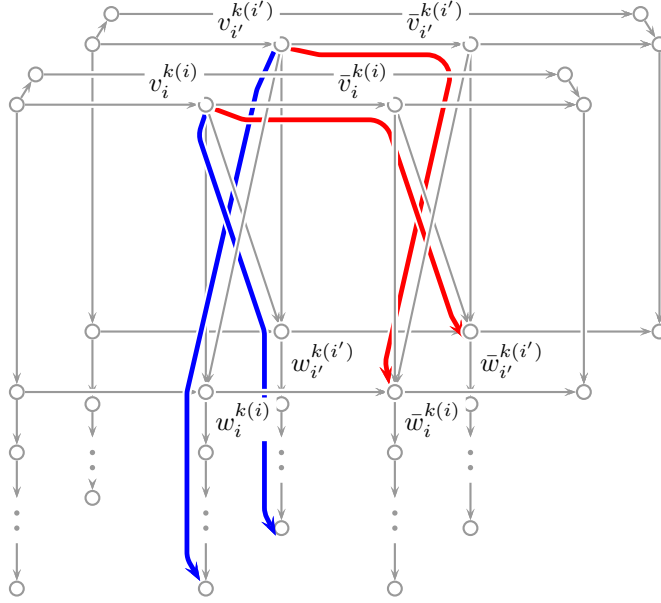


Figure 4: Prescribed shortest paths $Q_{i,i'}^j$ for some pair $i, i' \in S$ with $i \neq i'$ and $j = k(i) = k(i')$.

$\lambda_{(w_i^{j(i)}, \bar{w}_i^{j(i)})} = 3$ and $\lambda_{(w_i^{k(i)}, \bar{w}_i^{k(i)})} = \lambda_{(w_i^{l(i)}, \bar{w}_i^{l(i)})} = 2$. Thus, $\lambda(P) = 7$. Since P is the unique $(w_i^{j(i)}, \bar{w}_i^{l(i)})$ -path in D , it also is the unique shortest such path.

Let $P = (v_i^j, \bar{v}_i^j, \bar{w}_{i'}^j)$ for $j \in J$ and $(i, i') \in C_j^2$. If $C_j \in \mathcal{C}'$, then (2) implies $\lambda_{(v_i^j, \bar{v}_i^j)} = 2$ and $\lambda_{(\bar{v}_i^j, \bar{w}_{i'}^j)} = 1$. If $C_j \notin \mathcal{C}'$, then the length of both arcs in P is one. Consequently, $\lambda(P) \leq 3$. The only other $(v_i^j, \bar{w}_{i'}^j)$ -path in D is $P' = (v_i^j, w_{i'}^j, \bar{w}_{i'}^j)$. If $C_j \in \mathcal{C}'$, then $\lambda_{(v_i^j, w_{i'}^j)} = 1$ and $\lambda_{(w_{i'}^j, \bar{w}_{i'}^j)} = 3$. Otherwise, $\lambda_{(v_i^j, w_{i'}^j)} = 1$ and $\lambda_{(w_{i'}^j, \bar{w}_{i'}^j)} = 2$. Independent of whether $C_j \in \mathcal{C}'$ or not, it follows that $\lambda(P) < \lambda(P')$.

Finally, let $P = (v_i^j, w_{i'}^j, b_{i',1}^j, \dots, b_{i',6}^j)$ for some $j \in J$ and $(i, i') \in C_j^2$. By (2), all arcs in P have length one. Hence, $\lambda(P) = 7$. Furthermore, P is the solitary $(v_i^j, b_{i',6}^j)$ -path.

Summarizing these arguments, it follows that the arc lengths λ are compatible with \mathcal{Q} and that $\max_{P \in \mathcal{Q}} \lambda(P) = 7$. This concludes the second part of the proof.

In the following third part, it remains to show that the constructed MIN-PATH-ISP instance has a solution of value 7 *only* if the given SET PARTITION(3) instance is solvable. For this, we prove that any arc length vector $\lambda \in \mathbb{Z}_+^a$ with $\lambda_a \geq 1$ for all $a \in A$ that is compatible with \mathcal{Q} and satisfies $\max_{P \in \mathcal{Q}} \lambda(P) = 7$ can be transformed into a subcollection $\mathcal{C}' \subseteq \mathcal{C}$ that defines a partition of S . So, suppose we are given such arc lengths and define the subcollection $\mathcal{C}' = \mathcal{C}'(\lambda) \subseteq \mathcal{C}$ as

$$\mathcal{C}' := \{C_j \in \mathcal{C} \mid \lambda_{(v_i^j, \bar{v}_i^j)} \geq 2 \text{ for some } i \in C_j\}. \quad (3)$$

First, we show that every element $i \in S$ is covered by \mathcal{C}' . Let $i \in S$ and consider the two $(v_i^{j(i)}, \bar{v}_i^{l(i)})$ -paths $P = (v_i^{j(i)}, u_{i,1}, u_{i,2}, \bar{v}_i^{l(i)})$ and $P' = (v_i^{j(i)}, v_i^{k(i)}, v_i^{l(i)}, \bar{v}_i^{l(i)})$. Since $P \in \mathcal{Q}$ and the lengths λ are compatible with \mathcal{Q} , the inequality $\lambda(P) \geq \lambda(P') + 1$ must hold. With $\lambda_a \geq 1$ for all $a \in A$, it follows that $\lambda(P') \geq 4$. This, in turn, implies $\lambda_{(v_i^j, \bar{v}_i^j)} \geq 2$ for some $j \in J(i)$. According to (3), C_j then belongs to \mathcal{C}' . Since $i \in C_j$, the element i is covered by \mathcal{C}' .

Secondly, we show that the sets in \mathcal{C}' are pairwise disjoint. Suppose \mathcal{C}' contains two sets C_j and C_k with $j \neq k$ such that $i \in C_j \cap C_k \neq \emptyset$. Since C_j belongs to \mathcal{C}' , there must be some $q \in C_j$ with




Figure 5: Extension of digraph D to digraph D^* .

$\lambda_{(v_q^j, \bar{v}_q^j)} \geq 2$. For this q , consider the two (v_q^j, \bar{w}_i^j) -paths $P_1 = (v_q^j, \bar{v}_q^j, \bar{w}_i^j)$ and $P_2 = (v_q^j, w_i^j, \bar{w}_i^j)$. As λ is compatible with \mathcal{Q} and $P_1 \in \mathcal{Q}$, the lengths λ must satisfy $\lambda(P_2) \geq \lambda(P_1) + 1$, i.e.,

$$\lambda_{(v_q^j, w_i^j)} + \lambda_{(w_i^j, \bar{w}_i^j)} \geq \lambda_{(v_q^j, \bar{v}_q^j)} + \lambda_{(\bar{v}_q^j, \bar{w}_i^j)} + 1. \quad (4)$$

Let $P_3 = (v_q^j, w_i^j, b_{i,1}^j, \dots, b_{i,6}^j)$. Since $P_3 \in \mathcal{Q}$ and $|P_3| = 7$, all arcs in P_3 have length 1. In particular, $\lambda_{(v_q^j, w_i^j)} = 1$. Together with (4) and $\lambda_{(v_q^j, \bar{v}_q^j)} \geq 2$, this implies $\lambda_{(w_i^j, \bar{w}_i^j)} \geq 3$. Analogously, $C_k \in \mathcal{C}'$ implies $\lambda_{(w_i^k, \bar{w}_i^k)} \geq 3$. For the third index $l = J(q) \setminus \{j, k\}$, a similar argument yields $\lambda_{(w_i^l, \bar{w}_i^l)} \geq 2$.

Now consider the path $P = (w_q^{j(a)}, w_q^{k(a)}, w_q^{l(a)}, \bar{w}_q^{l(i)}) \in \mathcal{Q}$. Since all arcs (w_i^j, \bar{w}_i^j) , (w_i^k, \bar{w}_i^k) , and (w_i^l, \bar{w}_i^l) are contained in P , it follows that $\lambda(P) \geq 8$. This, however, conflicts with our assumption $\max_{P \in \mathcal{Q}} \lambda(P) = 7$. Consequently, all sets in \mathcal{C}' are pairwise disjoint and form a partition of S .

Together, the three parts of the proof imply that our construction maps any given SET PARTITION(3) instance to a solvable MIN-PATH-ISP instance. Furthermore, the optimum solution value of the constructed MIN-PATH-ISP instance is 7 if and only if the given SET PARTITION(3) instance is solvable and at least 8 otherwise. As SET PARTITION(3) is \mathcal{NP} -complete and the above construction is polynomial, it is \mathcal{NP} -hard to approximate MIN-PATH-ISP within a factor strictly less than 8/7. \square

An analogous construction proves the \mathcal{APX} -hardness of the corresponding undirected problem variant of MIN-PATH-ISP. The argumentation becomes slightly more complicated, because there are more valid paths between certain terminal pairs in the corresponding undirected graph than in the directed one. Yet, the construction yields the same inapproximability bounds of 8/7.

In the remainder of this section, we show that also the second problem variant MIN-ARC-ISP is \mathcal{APX} -hard. This result follows straightforward from the constant inapproximability threshold for MIN-PATH-ISP.

Theorem 3.2 *It is \mathcal{NP} -hard to approximate MIN-ARC-ISP within a factor of $9/8 - \epsilon$, for any $\epsilon > 0$.*

Proof. Suppose we are given an MIN-PATH-ISP instance consisting of a digraph $D = (V, A)$ and a shortest path system \mathcal{Q} . Consider the MIN-ARC-ISP instance with the same path set \mathcal{Q} , but in an extended digraph $D^* = (V, A \cup A^*)$ with $A^* := \{(s, t) \mid \mathcal{Q} \text{ contains an } (s, t)\text{-path } P_{s,t}\}$, see Figure 5. Note that we introduce a new arc (s, t) in A^* even if the original arc set A contains a parallel arc (s, t) between the terminals of some prescribed (s, t) -path $P_{s,t}$.

Clearly, the MIN-ARC-ISP instance remains solvable. Furthermore, for any arc length vector $\lambda \in \mathbb{Z}_+^A$ that is compatible with \mathcal{Q} in D , the arc lengths $\lambda^* \in \mathbb{Z}_+^{A \cup A^*}$ defined as

$$\lambda_a^* := \begin{cases} \lambda_a & , \text{ for all } a \in A, \text{ and} \\ \lambda(P_{s,t}) + 1 & , \text{ for all } a = (s,t) \in A^*, \end{cases}$$

are compatible with \mathcal{Q} in D^* . Reversely, all arc lengths $\lambda^* \in \mathbb{Z}_+^{A \cup A^*}$ in D^* that are compatible with \mathcal{Q} define arc lengths $\lambda = (\lambda_a^*)_{a \in A}$ in D that are compatible with \mathcal{Q} and satisfy $\lambda(P_{s,t}) \leq \lambda_{(s,t)}^* - 1$ for all $P_{s,t} \in \mathcal{Q}$.

As it is \mathcal{NP} -hard to find lengths λ in D such that the longest path in \mathcal{Q} is strictly less than 8, it is also \mathcal{NP} -hard to find lengths λ^* in D^* such that all arc lengths (in particular those of the arcs in A^*) are strictly less than 9. Therefore, it is \mathcal{NP} -hard to approximate MIN-ARC-ISP within a factor less than $9/8$. \square

Analogously, one can show the same inapproximability threshold of $9/8$ for the corresponding undirected problem.

4 An $\mathcal{O}(|V|)$ -Approximation Algorithm

Real-valued arc lengths that are compatible with a given set of paths can be computed in polynomial time. Let s_P and t_P denote the source and the destination node of a path P , and let $\mathcal{P}(s,t)$ be the set of all (s,t) -paths in D . The following linear program models the problem of finding (unbounded) real-valued arc lengths $\lambda \in \mathbb{R}_+^A$ with $\lambda_a \geq 1$ for all $a \in A$ that are compatible with a given path set \mathcal{Q} .

$$\sum_{a \in P'} \lambda_a - \sum_{a \in P} \lambda_a \geq 1 \quad \text{for all } P \in \mathcal{Q} \text{ and } P' \in \mathcal{P}(s_P, t_P) \setminus \{P\}, \text{ and} \quad (5)$$

$$\lambda_a \geq 1 \quad \text{for all } a \in A. \quad (6)$$

Clearly, this linear system has a solution λ^* if and only if \mathcal{Q} is an SPS.

Many inequalities of type (5) are redundant for a given path set \mathcal{Q} . To characterize the corresponding path, we introduce the notion of weak disjointness. For a path P and two nodes u and v that occur in P (in this order) we denote by $P[u,v]$ the (u,v) -subpath of P .

Definition 4.1 *We say that two (s,t) -paths P_1 and P_2 are weakly disjoint if there are two distinct nodes u and v that occur in this order in both paths P_1 and P_2 and*

(i) $P_1[s,u] = P_2[s,u]$,

(ii) $P_1[v,t] = P_2[v,t]$, and

(iii) $P_1[u,v]$ and $P_2[u,v]$ are (internally) node disjoint.

For an (s,t) -path P , the set of all weakly disjoint paths is denoted by $\mathcal{WD}(P) \subset \mathcal{P}(s,t)$.

Two weakly disjoint paths must be disjoint except for a common subpath at the start and a common subpath at the end. These subpaths may be empty, i.e., a pair of disjoint (s,t) -paths is also weakly disjoint. Figure 6 illustrates the weak disjointness of paths.

It is easy to verify that inequality (5) is redundant for all paths P and P' that are not weakly disjoint. Thus, the linear system (5)–(6) reduces to

$$\sum_{a \in P'} \lambda_a - \sum_{a \in P} \lambda_a \geq 1 \quad \text{for all } P \in \mathcal{Q} \text{ and } P' \in \mathcal{WD}(P), \text{ and} \quad (7)$$

$$\lambda_a \geq 1 \quad \text{for all } a \in A. \quad (8)$$




Figure 6: Examples for weakly disjoint paths (left) and not weakly disjoint paths (right).

Note that there are still exponentially many inequalities (7). However, these inequalities can be separated polynomially, by a 2-shortest paths algorithm [20] for example. Thus, the entire linear system (7)–(8) can be solved in polynomial time.

With an appropriate linear objective function, few additional variables and linear constraints, and the requirement for integer routing length variables λ_a , this linear system can be extended to an integer linear programming formulation for MIN-ARC-ISP:

$$\min \lambda_{\max} \quad (9)$$

$$\sum_{a \in P'} \lambda_a - \sum_{a \in P} \lambda_a \geq 1 \quad \text{for all } P \in \mathcal{Q} \text{ and } P' \in \mathcal{WD}(P), \quad (10)$$

$$1 \leq \lambda_a \leq \lambda_{\max} \quad \text{for all } a \in A, \text{ and} \quad (11)$$

$$\lambda_a \in \mathbb{Z}_+ \quad \text{for all } a \in A. \quad (12)$$

A formulation for MIN-PATH-ISP can be obtained similarly.

The optimal solution λ^* of the linear programming relaxation (9)-(11) corresponds to an optimal 'fractional' solution of MIN-ARC-ISP, i.e., real-valued arc lengths that are compatible with the given paths \mathcal{Q} and minimize the largest arc length. Standard scaling and rounding techniques can be applied to turn the optimal fractional lengths λ^* into an integer solution of MIN-ARC-ISP.

A simple algorithm which first solves the linear programming relaxation and then scales and rounds all fractional lengths simultaneously depending on the length of the longest path is shown below. Ben-Ameur and Gourdin [1] proposed a variant of this algorithm based on two similar linear programming formulations for the undirected MIN-ARC-ISP problem.

MIP-ROUNDING FOR MIN-ARC-ISP
Solve linear program (9)-(11).
If <i>not solvable</i> then
return ' \mathcal{Q} is no SPS'.
($\lambda^*, \lambda_{\max}^*$) := optimal solution of (9)-(11)
$M := \max_{P \in \mathcal{Q}} P $
If $M < V /2$ then
return lengths $\lambda_a := \lceil \lambda_a^* \cdot M \rceil$ for all $a \in A$,
else
return lengths $\lambda_a := \lceil \lambda_a^* \cdot V /2 - 1/2 \rceil$ for all $a \in A$.

Theorem 4.2 *The worst-case approximation ratio of algorithm MIP-ROUNDING for MIN-ARC-ISP is $\rho = \min(|V|/2, \max_{P \in \mathcal{Q}} |P|)$.*

Proof. Consider an MIN-ARC-ISP instance given by the digraph D and the path set \mathcal{Q} . If \mathcal{Q} is not an SPS, then the linear subsystem (10)-(11) has no solution and the algorithm returns infeasibility. So, assume \mathcal{Q} is an SPS and denote by $(\lambda^*, \lambda_{\max}^*)$ the optimal solution of (9)-(11). Clearly, $\lceil \lambda_{\max}^* \rceil$ is a lower bound on the optimal solution value.

First, consider the case where $M := \max_{P \in \mathcal{Q}} |P| < |V|/2$. Let $\bar{\lambda}_a := \lambda_a^* \cdot M$, for all $a \in A$. MIP-ROUNDING returns arc lengths $\lambda_a = \lceil \bar{\lambda}_a \rceil$, for all $a \in A$. Because λ^* satisfies all inequalities (10), we have $\bar{\lambda}(P') - \bar{\lambda}(P) \geq M$ for all $P \in \mathcal{Q}$ and $P' \in \mathcal{WD}(P)$. With $\bar{\lambda}_a \leq \lambda_a < \bar{\lambda}_a + 1$, it

follows that $\lambda(P') - \lambda(P) > 0$ for all $P \in \mathcal{Q}$ and $P' \in \mathcal{WD}(P)$. This implies that all paths in \mathcal{Q} are indeed unique shortest paths w.r.t. λ . Hence, λ is a compatible arc length function for \mathcal{Q} with

$$\lambda_{\max} = \max_{a \in A} \lambda_a \leq \lceil M \cdot \lambda_{\max}^* \rceil \leq M \cdot \lceil \lambda_{\max}^* \rceil.$$

Now, consider the case where $M \geq |V|/2$. Let $\bar{\lambda}_a := \lambda_a^* \cdot |V|/2$, for all $a \in A$. Analogous to the first case, we observe that $\bar{\lambda}(P') - \bar{\lambda}(P) \geq |V|/2$ for all $P \in \mathcal{Q}$ and $P' \in \mathcal{WD}(P)$. Because P and P' are weakly disjoint, there are at most $|V|$ many arcs in *either* P or P' . With $\bar{\lambda}_a - 1/2 \leq \lambda_a < \bar{\lambda}_a + 1/2$, we therefore have $\lambda(P') - \lambda(P) > 0$ for all $P \in \mathcal{Q}$ and $P' \in \mathcal{WD}(P)$. Hence, λ is a compatible with \mathcal{Q} and

$$\lambda_{\max} = \max_{a \in A} \lambda_a \leq \left\lceil \frac{|V|}{2} \cdot \lambda_{\max}^* - \frac{1}{2} \right\rceil \leq \frac{|V|}{2} \cdot \lceil \lambda_{\max}^* \rceil.$$

This concludes the proof. \square

Analogously, it follows that algorithm MIP-ROUNDING achieves an approximation ratio of $\rho = \min(|V|/2, \max_{P \in \mathcal{Q}} |P|)$ for MIN-PATH-ISP if applied to the corresponding linear programming formulation for MIN-PATH-ISP. The same approximation ratio is achieved for the corresponding undirected problem variants, too.

For instances arising in real-world telecommunication network planning, the integer linear programming models for MIN-ARC-ISP and MIN-PATH-ISP usually can be solved efficiently with standard integer linear programming solvers.

5 Concluding remarks

In this paper, we have shown that it is \mathcal{NP} -hard to approximate MIN-PATH-ISP and MIN-ARC-ISP within factors less than $8/7$ and $9/8$, respectively. For any real shortest path routing protocol such as OSPF, IS-IS, PNNI, or RIP, it therefore is \mathcal{NP} -hard to decide whether a given set of paths can be realized. We also presented integer linear programming models and a simple $\mathcal{O}(|V|)$ -approximation algorithm for MIN-ARC-ISP and MIN-PATH-ISP. This algorithm is based on scaling and rounding the optimal solution of the corresponding linear relaxation.

The \mathcal{APX} -hardness of a third problem variant MIN-SUM-ISP, where the objective is to minimize the sum of all arc lengths, can be shown with an approximation-preserving reduction from the MIN VERTEX COVER problem. Because of the smaller practical relevance of the MIN-SUM-ISP problem we omit the proof here.

Theorem 5.1 *It is \mathcal{NP} -hard to approximate MIN-SUM-ISP within a factor of $91/90 - \epsilon$, for any $\epsilon > 0$.*

It remains open whether there are constant factor approximation algorithms for MIN-PATH-ISP, MIN-ARC-ISP, or MIN-SUM-ISP. The best know algorithms are based on linear programming and rounding and achieve a worst-case approximation ratio of $\mathcal{O}(|V|)$ for all three objectives.

References

- [1] W. Ben-Ameur and E. Gourdin. Internet routing and related topology issues. *SIAM Journal on Discrete Mathematics*, 17:18–49, 2003.
- [2] A. Bley. Network optimization for IP-networks. Talk at the International Symposium on Mathematical Programming 2000, Atlanta, 2000.
- [3] A. Bley. A Lagrangian approach for integrated network design and routing in IP networks. In *Proceedings of the First International Network Optimization Conference (INOC 2003), Paris*, pages 107–113, October 2003.

- [4] A. Bley, M. Grötschel, and R. Wessäly. Design of broadband virtual private networks: Model and heuristics for the B-WiN. In *Robust Communication Networks: Interconnection and Survivability*, volume 53 of *DIMACS Series in Discrete Mathematics and Theoretical Computer Science*, pages 1–16. American Mathematical Society, 1998.
- [5] A. Bley and T. Koch. Integer programming approaches to access and backbone IP-network planning. Technical Report ZR-02-41, Konrad-Zuse-Zentrum für Informationstechnik Berlin, 2002.
- [6] L.S. Buriol, M.G.C. Resende, C. Ribeiro, and M. Thorup. A hybrid genetic algorithm for the weight setting problem in OSPF/IS-IS routing. *Optimization Online*, 6 2003.
- [7] D. Burton. *On the Inverse Shortest Path Problem*. PhD thesis, Department of Mathematics, Facultés Universitaires ND de la Paix, Namur, Belgium, 1993.
- [8] D. Burton, B. Pulleyblank, and P.L. Toint. The inverse shortest path problem with upper bounds on shortest path costs. Technical Report 93/03, Department of Mathematics, Facultés Universitaires ND de la Paix, Namur, Belgium, 1993.
- [9] D. Burton and P.L. Toint. The use of an inverse shortest paths algorithm for recovering linearly correlated costs. Technical Report 91/09, Department of Mathematics, Facultés Universitaires ND de la Paix, Namur, Belgium, 1991.
- [10] D. Burton and P.L. Toint. On an instance of the inverse shortest paths problem. *Mathematical Programming*, 53:45–61, 1992.
- [11] M.-C. Cai and Y. Li. Inverse matroid intersection problem. *Mathematical Methods of Operations Research*, 45:235–243, 1997.
- [12] M. Ericsson, M.G.C. Resende, and P.M. Pardalos. A genetic algorithm for the weight setting problem in OSPF routing. Technical report, AT&T Labs Research, 2001.
- [13] A. Farago, A. Szentesi, and B. Szviatovski. Allocation of administrative weights in PNNI. In *Proceedings of Networks 1998*, pages 621–625, 1998.
- [14] S.P. Fekete, W. Hochstättler, S. Kromberg, and C. Moll. The complexity of an inverse shortest path problem. In *Contemporary Trends in Discrete Mathematics: From DIMACS and DIMATIA to the Future*, volume 49. American Mathematical Society, 1999.
- [15] B. Fortz and M. Thorup. Internet traffic engineering by optimizing ospf weights. In *Proceedings of the 19th IEEE Infocom 2000*, Tel-Aviv, Israel, pages 519–528, 2000.
- [16] B. Fortz and M. Thorup. Increasing internet capacity using local search. *Computational Optimization and Applications*, 29:13–48, 2004.
- [17] M.R. Garey and D.S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. Freeman and Company, New York, 1979.
- [18] K. Holmberg and D. Yuan. Optimization of internet protocol network design and routing. *Networks*, 43:39–53, 2004.
- [19] R.M. Karp. *Complexity of Computer Computations*, chapter Reducibility among Combinatorial Problems, pages 85–103. Plenum Press, New York, 1972.
- [20] N. Katoh, T. Ibaraki, and H. Mine. An efficient algorithm for k shortest simple paths. *Networks*, 12:411–427, 1982.
- [21] F.Y.S. Lin and J.L. Wang. Minimax open shortest path first routing algorithms in networks supporting the SMDS service. In *Proceedings of the IEEE International Conference on Communications 1993 (ICC'93)*, Geneva, Suisse, volume 2, pages 666–670, 1993.

- [22] C. Moll. *Das inverse Kürzeste-Wege-Problem*. PhD thesis, Universität zu Köln, 1995.
- [23] M. Pióro, A. Szentési, J. Harmatos, and A. Jüttner. On OSPF related network optimization problems. In *8th IFIP Workshop on Performance Modelling and Evaluation of ATM & IP Networks*, pages 70/1–70/14, Ickley, UK, 2000.
- [24] M. Prytz. *On Optimization in Design of Telecommunications Networks with Multicast and Unicast Traffic*. PhD thesis, Royal Institute of Technology, Stockholm, Sweden, 2002.
- [25] C.W. Tong and K.P. Lam. An embedded connectionist approach for the inverse shortest paths problem. Technical report, Department of Systems Engineering and Engineering Management, The Chinese University of Hong Kong, 1996.