

ANDREAS BLEY

On the Hardness of Finding Small Shortest Path Routing Conflicts

On the Hardness of Finding Small Shortest Path Routing Conflicts

Andreas Bley
Konrad-Zuse-Zentrum
Address : Takustr. 7, 14195 Berlin, Germany
Email: bley@zib.de

April 2009

Abstract

Nowadays most data networks use shortest path protocols such as OSPF or IS-IS to route traffic. Given administrative routing lengths for the links of a network, all data packets are sent along shortest paths with respect to these lengths from their source to their destination. One of the most fundamental problems in planning shortest path networks is to decide whether a given set of routing paths forms a valid routing and, if this is not the case, to find a small subset of the given paths that cannot be shortest paths simultaneously for any routing lengths. In this paper we show that it is \mathcal{NP} -hard to approximate the size of the smallest shortest path conflict by a factor less than $7/6$.

Keywords: shortest path routing, computational complexity

1 Introduction

One of the most fundamental problems in planning networks that employ a shortest path routing protocol such as OSPF or IS-IS is to decide whether a given set of routing paths forms a valid routing and, if this is the case, to find a routing metric for which these paths are shortest paths between their respective terminals. If the given path set does not form a valid shortest path routing, one often wishes to find a small subset of the given paths that form a shortest path conflict, i.e., that cannot be shortest paths simultaneously for any routing metric.

In this paper, we consider the problem of finding a shortest path routing conflict of minimum size in a given path set. This problem naturally arises in integer linear programming approaches for shortest path routing optimization, where invalid routing patterns are cut off the feasible solution space using inequalities based on such shortest path routing conflicts; see [3, 4, 5, 6, 10, 12, 13, 14]. The separation problem for these conflict inequalities is exactly the problem of finding a shortest path routing conflict of minimum weight. Depending on the shortest path routing version and the integer linear programming formulation used to optimize the routing paths, slightly different notions of 'conflict' and of the 'weight' of a shortest path routing conflict have to be used. For the unsplittable shortest path routing version, Bley [3, 4] proposed greedy algorithms to compute conflicts that are inclusion-wise minimal, but not necessarily minimal in terms of size or weight. For the problem version associated with an arc-flow formulation for the shortest multi-path routing version, Tomaszewski et al. [14] proposed an integer programming approach, while Broström and Holmberg [7, 8] derived efficient polynomial time algorithms to optimize (and separate) over a special subclass of the corresponding conflicts.

In this paper, we show that finding minimal shortest path routing conflicts is \mathcal{NP} -hard. We only discuss the unweighted case corresponding to the path flow formulation of the unsplittable shortest path routing variant explicitly. We prove that it is \mathcal{NP} -hard to approximate the size of the smallest path set that is contained in a given set of paths and that does not comprise a

unique shortest path routing by a factor smaller than $7/6$. This result, however, generalizes in a straightforward way to the conflicts associated with the shortest multi-path routing variant and with integer programming formulations based on arc flows or shortest path graphs.

2 Notation and Preliminaries

Throughout this paper, let $D = (V, A)$ be a directed graph, $\mathcal{P}(s, t)$ be the set of all simple (s, t) -paths, and \mathcal{P} be the set of all simple paths in D . We denote the source and the target of a path P by s_P and t_P , respectively. We say that the arc lengths $\lambda \in \mathbb{Z}^A$ are *compatible* with a given path set $\mathcal{S} \subseteq \mathcal{P}$, if each path $P \in \mathcal{S}$ is the uniquely determined shortest path between its terminals with respect to λ . A path set \mathcal{S} is a *Unique Shortest Path System (USPS)* if there exists a vector of compatible arc lengths $\lambda \in \mathbb{Z}^A$ for \mathcal{S} . Otherwise we say that \mathcal{S} is a *non-USPS*. If \mathcal{S} is a non-USPS, but any proper subset of \mathcal{S} is a USPS, then \mathcal{S} is called an *irreducible non-USPS* or a *(unique) shortest path conflict*.

Clearly, any subset of an USPS is an USPS as well. Also the empty path set $\mathcal{S} = \emptyset$ is an USPS. The family of USPSs in a digraph D forms a so-called independence system (or hereditary family) $\mathcal{I} \subseteq 2^{\mathcal{P}}$. The circuits of this independence system are exactly the irreducible non-USPSs. A given path set \mathcal{S} forms a valid unique shortest path routing if and only if it does not fully contain any of these irreducible non-USPSs. Using the linear programming techniques discussed in [1, 2], for example, one can decide in polynomial time whether a given path set $\mathcal{S} \subseteq \mathcal{P}$ is a USPS or not. One easily finds that \mathcal{S} is an USPS if and only if the linear system

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

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

has a solution. As the separation problem for the inequalities (1a) boils down to a two shortest path problem, which can be solved efficiently with the algorithm of Katoh, Ibaraki, and Mine [11] for example, the overall system (1) can be solved in polynomial time. Using scaling and rounding techniques, any fractional solution of (1) can be turned into an integer-valued compatible metric for the given path set \mathcal{S} [1]. If (1) has no solution, then \mathcal{S} is a non-USPS. In this case, we are interested in finding a small non-USPS contained in \mathcal{S} , which can be formally described as follows:

<i>Problem:</i>	MIN-NON-USPS
<i>Instance:</i>	A digraph $D = (V, A)$ and a non-USPS $\mathcal{S} \subseteq \mathcal{P}$.
<i>Solution:</i>	An (irreducible) non-USPS $\mathcal{R} \subseteq \mathcal{S}$.
<i>Objective:</i>	$\min \mathcal{R} $.

The more general problem of finding a non-USPS of minimum weight can be formalized as follows:

<i>Problem:</i>	MIN-WEIGHT-NON-USPS
<i>Instance:</i>	A digraph $D = (V, A)$, a non-USPS $\mathcal{S} \subseteq \mathcal{P}$, and strictly positive weights $w_P \in \mathbb{Z}_+$ for all $P \in \mathcal{S}$.
<i>Solution:</i>	An (irreducible) non-USPS $\mathcal{R} \subseteq \mathcal{S}$.
<i>Objective:</i>	$\min \sum_{P \in \mathcal{R}} w_P$.

This problem arises when we seek for an inequality that separates a given fractional routing from the unique shortest path routing polytope, which is defined by the path sets of all valid unique shortest path routings. Its computational complexity therefore is of great practical importance. If MIN-WEIGHT-NON-USPS cannot be solved in polynomial time, one cannot hope to optimize over the unique shortest path routing polytope in polynomial time.




Figure 1: Subgraph of D corresponding to edges f_1, \dots, f_m , node $w_i \in W$, and incidences $f_1, f_2 \in \delta(w_i)$.

3 Hardness Results

For any fixed $k \in \mathbb{Z}$, we can find a minimum weight non-USPS $\mathcal{R} \subseteq \mathcal{S}$ with $|\mathcal{R}| \leq k$ (or prove that no irreducible non-USPS with $|\mathcal{R}| \leq k$ exists) in polynomial time by solving the linear system (1) for all subsets $\mathcal{R} \subseteq \mathcal{S}$ with $|\mathcal{R}| \leq k$. In special digraphs where the size of irreducible non-USPSs is bounded by some constant, MIN-NON-USPS and MIN-WEIGHT-NON-USPS are therefore solvable in polynomial time. In general, however, it is computationally hard to approximate these problems within a factor strictly smaller than $7/6$.

Theorem 3.1 *For any $\epsilon > 0$, it is \mathcal{NP} -hard to approximate MIN-NON-USPS within a factor $7/6 - \epsilon$. This holds even if each path $P \in \mathcal{S}$ is a shortest (s_P, t_P) -path w.r.t. the number of arcs and $|P| = 2$ for all $P \in \mathcal{S}$.*

Proof. We construct an approximation preserving reduction from the optimization problem MINIMUM VERTEX COVER. The latter problem is defined as follows: Given an undirected graph $H = (W, F)$, find a minimum cardinality set $C \subseteq W$ such that, for each edge $uv \in E$, at least one of the nodes u and v belongs to C . Håstad [9] proved that, unless $\mathcal{P} = \mathcal{NP}$, this problem is not approximable within a factor $7/6 - \epsilon$, for any $\epsilon > 0$.

Suppose we are given an instance $H = (W, F)$ of MINIMUM VERTEX COVER consisting of the nodes w_i with $i \in I := \{1, \dots, n\}$ and the edges f_k with $k \in K = \{1, \dots, m\}$. Let $\alpha \in \mathbb{Z}$ be a large integer number. At the end of the proof, we discuss how to choose α appropriately. We construct a MIN-NON-USPS instance consisting of a digraph $D = (V, A)$ and a path set \mathcal{S} as follows.

For each $i \in I$, we introduce $2\alpha + 2$ nodes u_i^0, \dots, u_i^α and $\bar{u}_i^0, \dots, \bar{u}_i^\alpha$. These nodes are connected by the arcs (u_i^l, \bar{u}_i^l) and (\bar{u}_i^l, u_i^l) , for $l = 0, \dots, \alpha$, and by the arcs (u_i^l, u_i^{l+1}) and $(\bar{u}_i^l, \bar{u}_i^{l+1})$, for $l = 0, \dots, \alpha - 1$. For each $k \in K$, we add four nodes $v_k^1, v_k^2, \bar{v}_k^1$, and \bar{v}_k^2 . These are connected by the arcs (v_k^1, v_k^2) , (v_k^2, v_k^1) , $(\bar{v}_k^1, \bar{v}_k^2)$, $(\bar{v}_k^2, \bar{v}_k^1)$, (v_k^1, \bar{v}_k^1) , (\bar{v}_k^1, v_k^1) , (v_k^2, \bar{v}_k^2) , and (\bar{v}_k^2, v_k^2) for all $k \in K$. Furthermore, we add two arcs (\bar{v}_k^1, v_{k+1}^1) and (\bar{v}_k^2, v_{k+1}^2) for each $k = 1, \dots, m - 1$, and two arcs (\bar{v}_m^1, v_1^1) and (\bar{v}_m^2, v_1^2) for $k = m$. Finally, we introduce arcs for the node-edge incidences in H . For all $i \in I$ and all k with $f_k \in \delta(w_i)$, we add the four arcs (v_k^2, u_i^0) , $(\bar{v}_k^2, \bar{u}_i^0)$, (u_i^α, v_k^1) , and $(\bar{u}_i^\alpha, \bar{v}_k^1)$. The resulting digraph D is illustrated in Figure 1.




Figure 2: Union of the path sets \mathcal{S}_k^2 for all $k \in K$ and \mathcal{S}_i^1 , $\mathcal{S}_{i,k}^3$, and $\mathcal{S}_{i,k}^4$ for some i, k with $f_k \in \delta(w_i)$.

The given path set consists of four different types of paths. For each $i \in I$, we have 2α many paths

$$\mathcal{S}_i^1 := \{(\bar{u}_i^l, u_i^l, u_i^{l+1}), (u_i^l, \bar{u}_i^l, \bar{u}_i^{l+1}) \mid l = 0, \dots, \alpha - 1\}.$$

For each $k \in K$, we are given four paths, namely

$$\begin{aligned} \mathcal{S}_k^2 &:= \{(\bar{v}_k^1, v_k^1, v_k^2), (v_k^2, v_k^1, \bar{v}_k^1), (\bar{v}_k^1, \bar{v}_k^2, v_{k+1}^2), (\bar{v}_k^2, \bar{v}_k^1, v_{k+1}^1)\} && \text{for } k \neq m, \text{ and} \\ \mathcal{S}_m^2 &:= \{(\bar{v}_k^1, v_k^1, v_k^2), (v_k^2, v_k^1, \bar{v}_k^1), (\bar{v}_k^1, \bar{v}_k^2, v_{k+1}^2), (\bar{v}_k^2, \bar{v}_k^1, v_{k+1}^1)\} && \text{for } k = m. \end{aligned}$$

Finally, we are given four paths for each node-edge incidence in H . For each $i \in I$ and each k with $f_k \in \delta(w_i)$, these paths are

$$\mathcal{S}_{i,k}^3 := \{(v_k^2, \bar{v}_k^2, \bar{u}_i^0), (\bar{v}_k^2, v_k^2, u_i^0)\} \quad \text{and} \quad \mathcal{S}_{i,k}^4 := \{(u_i^\alpha, \bar{u}_i^\alpha, \bar{v}_k^1), (\bar{u}_i^\alpha, u_i^\alpha, v_k^1)\}.$$

The path set \mathcal{S} is the union of all those sets, i.e., $\mathcal{S} := \bigcup_i \mathcal{S}_i^1 \cup \bigcup_k \mathcal{S}_k^2 \cup \bigcup_{i,k: f_k \in \delta(w_i)} \mathcal{S}_{i,k}^3 \cup \mathcal{S}_{i,k}^4$. It is illustrated in Figure 2. Note that each path in \mathcal{S} is a shortest path between its terminal nodes and contains exactly two arcs.

In the first part of the proof, we show that any vertex cover $C \subseteq W$ in H can be transformed into a non-USPS $\mathcal{R} := \mathcal{R}(C) \subseteq \mathcal{S}$ in D with $|\mathcal{R}| = 2\alpha |C| + 8m$. As a byproduct, this also proves that the constructed path system \mathcal{S} is indeed a non-USPS.

Let $C \subseteq W$ be a vertex cover in H . For each edge $f = w_i w_j \in F$, we denote $c(f) := w_i$, if $i < j$ and $w_i \in C$, and $c(f) := w_j$, otherwise. Since C is a vertex cover, $c(f) \in C$ for all $f \in F$. We define the path set corresponding to C as

$$\mathcal{R} = \mathcal{R}(C) := \bigcup_{i: w_i \in C} \mathcal{S}_i^1 \cup \bigcup_k \mathcal{S}_k^2 \cup \bigcup_{i,k: w_i = c(f_k)} \mathcal{S}_{i,k}^3 \cup \mathcal{S}_{i,k}^4.$$

The size of this path set is

$$|\mathcal{R}| = 2\alpha \cdot |C| + 8m. \quad (2)$$

The path set \mathcal{R} is a USPS if and only if the linear system (1) has a feasible solution. In order to show that \mathcal{R} is a non-USPS, it is therefore sufficient to show that the following sub-system of

(1) has no solution:

$$\lambda_{(\bar{v}_k^1, v_k^1)} + \lambda_{(v_k^1, v_k^2)} + 1 \leq \lambda_{(\bar{v}_k^1, \bar{v}_k^2)} + \lambda_{(\bar{v}_k^2, v_k^2)} \quad \forall k \quad (3a)$$

$$\lambda_{(v_k^1, \bar{v}_k^1)} + \lambda_{(v_k^2, v_k^1)} + 1 \leq \lambda_{(\bar{v}_k^2, \bar{v}_k^1)} + \lambda_{(v_k^2, \bar{v}_k^2)} \quad \forall k \quad (3b)$$

$$\lambda_{(\bar{v}_k^1, \bar{v}_k^2)} + \lambda_{(\bar{v}_k^2, v_{k+1}^2)} + 1 \leq \lambda_{(\bar{v}_k^1, v_{k+1}^1)} + \lambda_{(v_{k+1}^1, v_{k+1}^2)} \quad \forall k \neq m \quad (3c)$$

$$\lambda_{(\bar{v}_k^2, \bar{v}_k^1)} + \lambda_{(\bar{v}_k^1, v_{k+1}^1)} + 1 \leq \lambda_{(\bar{v}_k^2, v_{k+1}^2)} + \lambda_{(v_{k+1}^2, v_{k+1}^1)} \quad \forall k \neq m \quad (3d)$$

$$\lambda_{(\bar{v}_m^1, \bar{v}_m^2)} + \lambda_{(\bar{v}_m^2, v_1^1)} + 1 \leq \lambda_{(\bar{v}_m^1, v_1^1)} + \lambda_{(v_1^2, v_1^1)} \quad (3e)$$

$$\lambda_{(\bar{v}_m^2, \bar{v}_m^1)} + \lambda_{(\bar{v}_m^1, v_1^2)} + 1 \leq \lambda_{(\bar{v}_m^2, v_1^2)} + \lambda_{(v_1^1, v_1^2)} \quad (3f)$$

$$\lambda_{(\bar{u}_i^l, u_i^l)} + \lambda_{(u_i^l, u_i^{l+1})} + 1 \leq \lambda_{(\bar{u}_i^l, \bar{u}_i^{l+1})} + \lambda_{(\bar{u}_i^{l+1}, u_i^{l+1})} \quad \forall i, l : w_i \in C, l \neq \alpha \quad (3g)$$

$$\lambda_{(u_i^l, \bar{u}_i^l)} + \lambda_{(\bar{u}_i^l, \bar{u}_i^{l+1})} + 1 \leq \lambda_{(u_i^l, u_i^{l+1})} + \lambda_{(u_i^{l+1}, \bar{u}_i^{l+1})} \quad \forall i, l : w_i \in C, l \neq \alpha \quad (3h)$$

$$\lambda_{(\bar{v}_k^2, v_k^2)} + \lambda_{(v_k^2, u_i^0)} + 1 \leq \lambda_{(\bar{v}_k^2, \bar{u}_i^0)} + \lambda_{(\bar{u}_i^0, u_i^0)} \quad \forall k, i : w_i = c(f_k) \quad (3i)$$

$$\lambda_{(v_k^2, \bar{v}_k^2)} + \lambda_{(\bar{v}_k^2, \bar{u}_i^0)} + 1 \leq \lambda_{(v_k^2, u_i^0)} + \lambda_{(u_i^0, \bar{u}_i^0)} \quad \forall k, i : w_i = c(f_k) \quad (3j)$$

$$\lambda_{(u_i^\alpha, \bar{u}_i^\alpha)} + \lambda_{(\bar{u}_i^\alpha, \bar{v}_k^1)} + 1 \leq \lambda_{(u_i^\alpha, v_k^1)} + \lambda_{(v_k^1, \bar{v}_k^1)} \quad \forall k, i : w_i = c(f_k) \quad (3k)$$

$$\lambda_{(\bar{u}_i^\alpha, u_i^\alpha)} + \lambda_{(u_i^\alpha, v_k^1)} + 1 \leq \lambda_{(\bar{u}_i^\alpha, \bar{v}_k^1)} + \lambda_{(\bar{v}_k^1, v_k^1)} \quad \forall k, i : w_i = c(f_k) \quad (3l)$$

Inequalities (3a) ensure that, for each k , the path $(\bar{v}_k^1, v_k^1, v_k^2)$ is strictly shorter than the other two-arc path $(\bar{v}_k^1, \bar{v}_k^2, v_k^2)$ from \bar{v}_k^1 to v_k^2 . Together, (3a)–(3f) express that each path in $\bigcup_k \mathcal{S}_k^2$ is strictly shorter than its alternative other two-arc path. Analogously, inequalities (3g) and (3h) enforce that each path in $\bigcup_{i:w_i \in C} \mathcal{S}_i^1$ is shorter than its respective alternative two-arc path, while inequalities (3i)–(3l) ensure this property for all paths in $\bigcup_{k,i:w_i=c(f_k)} \mathcal{S}_{i,k}^3 \cup \mathcal{S}_{i,k}^4$.

To verify that this linear system has no solution, we apply Farkas' Lemma. For each i, l with $w_i \in C$ and $l \neq \alpha$, we multiply both inequalities (3g) and (3h) with a factor of $\mu(i) := |\{f \in F : w_i = c(f)\}|$, which yields the equivalent inequalities (3g') and (3h'). Adding all inequalities (3a)–(3f), (3g'), (3h'), and (3i)–(3l) then yields an inequality that contains each variable $\lambda_a, a \in A$, with the same coefficient on the left and on the right hand side and a positive constant on the left hand side. As this inequality cannot be satisfied, \mathcal{R} is no USPS.

In the second part of the proof, we now show that any irreducible non-USPS $\mathcal{R} \subseteq \mathcal{S}$ in D can be transformed back into a vertex cover $C := C(\mathcal{R}) \subseteq W$ in H with $2\alpha|C| \geq |\mathcal{R}| - 8m$. It is sufficient to define such a backward transformation only for irreducible non-USPSs, because any non-USPS $\mathcal{R}' \subseteq \mathcal{S}$ in D can be reduced to an irreducible non-USPS $\mathcal{R} \subseteq \mathcal{R}'$ in polynomial time (using the greedy algorithm proposed in [3], for example).

In order to define the backward transformation properly, we first need to show that all irreducible non-USPSs in D have a structure that is similar to that of the non-USPSs $\mathcal{R}(C)$ constructed in the first part of the proof. So, let $\mathcal{R} \subseteq \mathcal{S}$ be an irreducible non-USPS.

First, observe that all paths in $\bigcup_k \mathcal{S}_k^2$ must be contained in \mathcal{R} . Suppose there is some k' such that the path $(v_{k'}^2, v_{k'}^1, \bar{v}_{k'}^1)$ does not belong to \mathcal{R} . Without loss of generality, we may assume that $k' = 1$. Let $M \geq 2|A|$ and consider the metric

$$\lambda_a := \begin{cases} M + l + 1, & \text{if } a \in \{(u_i^l, \bar{u}_i^l), (\bar{u}_i^l, u_i^l)\}, \\ M + \alpha + 2, & \text{if } a \in \{(v_k^1, \bar{v}_k^1), (\bar{v}_k^1, v_k^1)\}, \\ M + 1, & \text{if } a \in \{(\bar{v}_k^1, v_{k+1}^1), (\bar{v}_k^2, v_{k+1}^2), (\bar{v}_m^1, v_1^2), (\bar{v}_m^2, \bar{v}_1^1)\} \\ M + (2m - 1)(\alpha + 4) + 1, & \text{if } a = (v_1^2, v_1^1), \\ M + (k - 2)(\alpha + 4) + 1, & \text{if } a \in \{(v_k^2, v_k^1) : k \neq 1\}, \\ M + (k - 1)(\alpha + 4), & \text{if } a = (\bar{v}_k^2, \bar{v}_k^1), \\ M + (m + k - 2)(\alpha + 4) + 1, & \text{if } a = (v_k^2, v_k^1), \\ M + (m + k - 1)(\alpha + 4), & \text{if } a = (\bar{v}_k^1, \bar{v}_k^2), \text{ and} \\ M, & \text{otherwise.} \end{cases}$$

One easily finds that $M \leq \lambda_a < 3/2M$ for all $a \in A$. Since all paths in \mathcal{S} contain two arcs, no path with three or more arcs in D can be shorter than any path in \mathcal{S} . For each path

$P \in \mathcal{S} \setminus \{(v_1^2, v_1^1, \bar{v}_1^1)\}$, however, there is only one alternative (s_P, t_P) -path with only two arcs, and it is straightforward to verify that each path in $\mathcal{S} \setminus \{(v_1^2, v_1^1, \bar{v}_1^1)\}$ is indeed shorter than the corresponding alternative path. Hence, $\mathcal{S} \setminus \{(v_1^2, v_1^1, \bar{v}_1^1)\}$ is a USPS, which implies that the path $(v_1^2, v_1^1, \bar{v}_1^1)$ must be contained in the (irreducible) non-USPS $\mathcal{R} \subseteq \mathcal{S}$. Analogously, it follows that any other path $P \in \bigcup_k \mathcal{S}_k^2$ is contained in \mathcal{R} .

With the same technique, we can show that, for any $k \in K$, there exists some $i := c(k)$ with $f_k \in \delta(w_i)$ such that $\mathcal{S}_{i,k}^3 \subset \mathcal{R}$. Without loss of generality, let $k = 1$, $f_1 = w_i w_j$, and suppose that at most one of the two paths $\bar{P}_i = (v_1^2, \bar{v}_1^2, \bar{u}_i^0)$ and $P_i = (\bar{v}_1^2, v_1^2, u_i^0)$ and at most one of the two paths $\bar{P}_j = (v_1^2, \bar{v}_1^2, \bar{u}_j^0)$ and $P_j = (\bar{v}_1^2, v_1^2, u_j^0)$ belong to \mathcal{R} . Then, with λ defined as above, the metric

$$\lambda'_a := \begin{cases} \lambda_a + 2m(\alpha + 4) & \text{if } a = (v_1^2, \bar{v}_1^2), \\ \lambda_a + 2m(\alpha + 4) & \text{if } a = (v_1^2, u_i^0) \text{ and } P_i \notin \mathcal{R}, \\ \lambda_a + 2m(\alpha + 4) & \text{if } a = (\bar{v}_1^2, \bar{u}_i^0) \text{ and } \bar{P}_i \notin \mathcal{R}, \\ \lambda_a + 2m(\alpha + 4) & \text{if } a = (v_1^2, u_j^0) \text{ and } P_j \notin \mathcal{R}, \\ \lambda_a + 2m(\alpha + 4) & \text{if } a = (\bar{v}_1^2, \bar{u}_j^0) \text{ and } \bar{P}_j \notin \mathcal{R}, \\ \lambda_a & \text{otherwise,} \end{cases}$$

is compatible with $\mathcal{S} \setminus (\{P_i, P_j, \bar{P}_i, \bar{P}_j\} \setminus \mathcal{R})$. Hence, the non-USPS \mathcal{R} must contain either both paths in $\mathcal{S}_{i,1}^3$ or both paths in $\mathcal{S}_{j,1}^3$.

Analogously, one can show that $\mathcal{S}_i^1 \subset \mathcal{R}$, for each $k \in K$ and $i = c(k)$. Furthermore, one finds that, for any $k \in K$, there exist some $j = c'(k)$ with $f_k \in \delta(w_j)$ such that $\mathcal{S}_{j,k}^4 \subset \mathcal{R}$.¹

Now we can define the vertex set corresponding to the irreducible non-USPS \mathcal{R} as

$$C = C(\mathcal{R}) := \{w_{c(k)} : k \in K\}.$$

Because $c(k)$ is either i or j for any edge $f_k = w_i w_j \in F$, the set C is a vertex cover in H . The above observations imply that

$$|\mathcal{R}| \geq \sum_{k \in K} \left(|\mathcal{S}_k^2| + |\mathcal{S}_{c(k),k}^3| + |\mathcal{S}_{c'(k),k}^4| \right) + \sum_{i \in I: w_i \in C} |\mathcal{S}_i^1| \geq 8m + 2\alpha |C|. \quad (4)$$

It follows in a straightforward way that computing an approximate solution for MIN-NON-USPS is at least as hard as computing an approximate solution for MINIMUM VERTEX COVER. Suppose there is an $(r - \epsilon)$ -approximation algorithm \mathcal{A} for MIN-NON-USPS with $\epsilon > 0$ and $(r - \epsilon) \geq 1$. Then we choose $\alpha := \lceil 8m(r - 1 - \epsilon) / \epsilon \rceil$. With this choice of α , both the construction of the MIN-NON-USPS instance and the backward transformation of an irreducible non-USPS to a vertex cover are polynomial in the encoding size of H .

Due to (2) and (4), we have $|\mathcal{R}^*| = 8m + 2\alpha |C^*|$ for any minimum vertex cover C^* in H and any minimum non-USPS $\mathcal{R}^* \subseteq \mathcal{S}$ in D . Furthermore,

$$\frac{2\alpha |C(\mathcal{R})| + 8m}{2\alpha |C^*| + 8m} \leq \frac{|\mathcal{R}|}{|\mathcal{R}^*|} \leq r - \epsilon$$

implies that $2\alpha |C(\mathcal{R})| \leq 2\alpha |C^*|(r - \epsilon) + 8m(r - 1 - \epsilon)$. Therefore, we have

$$\frac{|C(\mathcal{R})|}{|C^*|} \leq r - \epsilon + \frac{8m(r - 1 - \epsilon)}{2\alpha |C^*|} \leq r - \epsilon + \frac{\epsilon}{2|C^*|} \leq r - \frac{\epsilon}{2}.$$

Consequently, any $(r - \epsilon)$ -approximate solution $\mathcal{R} \subseteq \mathcal{S}$ for the constructed MIN-NON-USPS instance can be transformed back into an $(r - \epsilon/2)$ -approximate solution $C(\mathcal{R})$ of the given MINIMUM VERTEX COVER instance. Thus, any $(r - \epsilon)$ -approximation algorithm for MIN-NON-USPS yields an $(r - \epsilon/2)$ -approximation algorithm for MINIMUM VERTEX COVER.

¹Note that $i = c(k)$ and $j = c'(k)$ may be different. There may exist an irreducible non-USPS $\mathcal{R} \subseteq \mathcal{S}$ in the constructed digraph D that contains only one of the two path sets $\mathcal{S}_{i,k}^3$ and $\mathcal{S}_{i,k}^4$ for each $f_k \in \delta(w_i)$.

Since MINIMUM VERTEX COVER is \mathcal{NP} -hard to approximate within a factor strictly smaller than $7/6$, so is MIN-NON-USPS. \square

The constant inapproximability threshold of $7/6$ carries over directly to MIN-WEIGHT-NON-USPS.

Corollary 3.2 *For any $\epsilon > 0$, it is \mathcal{NP} -hard to approximate MIN-WEIGHT-NON-USPS within a factor $7/6 - \epsilon$.*

4 Concluding remarks

In this paper, we showed that the problems MIN-NON-USPS and MIN-WEIGHT-NON-USPS of finding a minimum size or minimum weight non-USPS contained in a given path set are \mathcal{NP} -hard to approximate within a factor strictly smaller than $7/6$. This implies that the problem of finding an inequality that separates a given fractional path-routing from the unique shortest path routing polytope is \mathcal{NP} -hard as well.

These results and the presented proof carry over in a straightforward way also to the case where we seek for a minimum size or a minimum weight conflict in an invalid shortest multi-path routing, where a conflict is given as a set of shortest paths and a set of non-shortest paths, that cannot be shortest and non-shortest paths simultaneously.

With a slight modification, the presented proofs also carry over to the problem where we seek for a minimum size or minimum weight conflict in a given collection of shortest path graphs (see [7, 3]) for both the unsplittable shortest and the shortest multi-path routing variant. The (extended) shortest path graph for a destination node defines which arcs must be contained in a shortest path towards this destination and which arcs must not be contained in any shortest path. A conflict in this representation is a pair of two sets. One set contains the prescribed destination-arc pairs, where the arc must be contained in a shortest path towards the destination, while the other set contains the forbidden destination-arc pairs, where the arc must not be contained in any shortest path towards the destination. These variants of the problem of finding a minimum weight shortest path routing conflict arise in the separation problem over the polytopes associated with integer linear programming formulations of shortest path routing problems that are based on arc routings or shortest path graphs.

References

- [1] W. Ben-Ameur and E. Gourdin. *Internet routing and related topology issues*, SIAM Journal on Discrete Mathematics **17** (2003) pp. 18–49.
- [2] A. Bley. *Inapproximability results for the inverse shortest paths problem with integer lengths and unique shortest paths*, Networks **50** (2008) pp. 29–36.
- [3] A. Bley. *Routing and Capacity Optimization for IP Networks*, PhD thesis, Technische Universität Berlin, 2007.
- [4] A. Bley. *An integer programming algorithm for routing optimization in IP networks*, In *Proceedings of the 16th Annual European Symposium on Algorithms (ESA 2008)*, Springer, Berlin Heidelberg, 2008, pp. 198–209.
- [5] A. Bley, B. Fortz, E. Gourdin, K. Holmberg, O. Klopfenstein, M. Pióro, A. Tomaszewski, and H. Ümit. *Optimization of OSPF routing in IP networks*, In *Graphs and Algorithms in Communication Networks*, A.M.C.A. Koster and X. Muñoz, eds., Springer, Berlin Heidelberg, to appear in 2009.

- [6] N. Bourquia, W. Ben-Ameur, E. Gourdin, and P. Tolla. *Optimal shortest path routing for Internet networks*, In *Proceedings of the 1st International Network Optimization Conference (INOC 2003)*, Paris, France, 2003, pp. 119–125.
- [7] P. Broström and K. Holmberg. *Determining the non-existence of a compatible OSPF metric*, In *Proceedings of the 2nd International Network Optimization Conference (INOC 2005)*, Lisbon, Portugal, 2005, pp. 702–709.
- [8] P. Broström and K. Holmberg. *Valid cycles: A source of infeasibility in OSPF routing*, *Networks* **52** (2008), pp. 206–215.
- [9] J. Håstad. *Some optimal inapproximability results*, In *Proceedings of the Twenty-ninth Annual ACM Symposium on the Theory of Computing*, El Paso, TX, 1997, pp. 1–10.
- [10] K. Holmberg and D. Yuan. *Optimization of Internet protocol network design and routing*, *Networks* **43** (2004), pp. 39–53.
- [11] N. Katoh, T. Ibaraki, and H. Mine. *An efficient algorithm for k shortest simple paths*, *Networks* **12** (1982), pp. 411–427.
- [12] M. Prytz. *On Optimization in Design of Telecommunications Networks with Multicast and Unicast Traffic*, PhD thesis, Royal Institute of Technology, Stockholm, Sweden, 2002.
- [13] A. Tomaszewski, M. Pióro, M. Dzida, and M. Zagożdżon. *Optimization of administrative weights in IP networks using the branch-and-cut approach*, In *Proceedings of the 2nd International Network Optimization Conference (INOC 2005)*, Lisbon, Portugal, 2005, pp. 393–400.
- [14] A. Tomaszewski, M. Pióro, M. Dzida, M. Mycek, and M. Zagożdżon. *Valid inequalities for a shortest-path routing optimization problem*, In *Proceedings of the 3rd International Network Optimization Conference (INOC 2007)*, Spa, Belgium, 2007.