Konrad-Zuse-Zentrum
für Informationstechnik Berlin

Takustraße 7
D-14195 Berlin-Dahlem
Germany

ANDREAS BLEY

Approximability of Unsplittable Shortest Path Routing Problems

URL: http://www.zib.de/PaperWeb/abstracts/ZR-06-02

ZIB-Report 06-02 (2006)

Approximability of Unsplittable Shortest Path Routing Problems *

Andreas Bley
Konrad-Zuse-Zentrum
Address : Takustr. 7, 14195 Berlin, Germany
Fax: +49-30-84185269
Email: bley@zib.de

March 2006

Abstract

In this paper, we discuss the relation of unsplittable shortest path routing (USPR) to other routing schemes and study the approximability of three USPR network planning problems. Given a digraph $D = (V, A)$ and a set K of directed commodities, an USPR is a set of flow paths $\Phi_{(s,t)}$, $(s,t) \in K$, such that there exists a metric $\lambda = (\lambda_a) \in \mathbb{Z}_+^A$ with respect to which each $\Phi_{(s,t)}$ is the unique shortest (s,t)-path.

In the Min-Con-USPR problem, we seek for an USPR that minimizes the maximum congestion over all arcs. We show that this problem is hard to approximate within a factor of $\mathcal{O}(|V|^{1-\epsilon})$, but easily approximable within $\min(|A|, |K|)$ in general and within $\mathcal{O}(1)$ if the underlying graph is an undirected cycle or a bidirected ring. We also construct examples where the minimum congestion that can be obtained by USPR is a factor of $\Omega(|V|^2)$ larger than that achievable by unsplittable flow routing or by shortest multi-path routing, and a factor of $\Omega(|V|)$ larger than by unsplittable source-invariant routing.

In the Cap-USPR problem, we seek for a minimum cost installation of integer arc capacities that admit an USPR of the given commodities. We prove that this problem is $\mathcal{NP}$-hard to approximate within $2 - \epsilon$ (even in the undirected case), and we devise approximation algorithms for various special cases. The fixed charge network design problem FC-USPR, where the task is to find a minimum cost subgraph of D whose fixed arc capacities admit an USPR of the commodities, is shown to be $\mathcal{NPO}$-complete.

All three problems are of great practical interest in the planning of telecommunication networks that are based on shortest path routing protocols. Our results indicate that they are harder than the corresponding unsplittable flow or shortest multi-path routing problems.

Keywords: Shortest Path Routing, Unsplittable Flow, Computational Complexity, Approximation
Mathematical Subject Classification (2000): 68Q25, 90C60, 90C27, 05C38, 90B18

*A shorter version of this paper appeared as "On the Approximability of the Minimum Congestion Unsplittable Shortest Path Routing Problem" in the Proceedings of 11th Conference on Integer Programming and Combinatorial Optimization (IPCO 2005), Berlin, Germany.

1 Introduction

Many data networks presently employ shortest path routing protocols such as OSPF, IS-IS, or RIP [42, 43, 44]. 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). The simplicity of this policy offers many advantages in practice: It admits the use of decentralized and distributed routing algorithms, it has very good scaling properties with respect to the network size, and it typically leads to less administrative overhead than with classical connection oriented protocols. On the other hand, the shortest path routing policy has an inherent drawback: It is not possible to configure end-to-end routing paths for the communication demands individually. The routing paths can be controlled only jointly and only indirectly by changing the administrative routing lengths of the network links.

Finding a metric of routing lengths that induce a set of globally efficient end-to-end routing paths is a major difficulty in such networks. The shortest path routing paradigm enforces rather complicated and subtle interdependencies among the paths that comprise a valid routing. Additional difficulties arise if the communication demands must be sent unsplit through the network – a requirement that is often imposed in order to ensure tractability of end-to-end traffic flows and to prevent package reordering and other unwanted effects of multi-path routing in practice. In this case, the lengths must be chosen such that the shortest paths are uniquely determined for all communication demands.

The task of finding an efficient such routing in an existing network can be formulated as a minimum congestion unsplittable shortest path routing problem (Min-Con-USPR). The problem input consists of a digraph $D = (V, A)$ with arc capacities $c_a \in \mathbb{Z}_+$ for all $a \in A$, and a set of directed commodities $K \subseteq V \times V$ with demand values $d_{(s,t)} \in \mathbb{Z}_+$, $(s,t) \in K$. A feasible solution is an unsplittable shortest path routing (USPR) of the commodities, i.e., a metric $\lambda = (\lambda_a)_{a \in A} \in \mathbb{Z}_+^A$ that induces a unique shortest (s,t)-path for each commodity $(s,t) \in K$. Each commodity's demand is sent unsplit along its shortest path. The objective is to minimize the maximum congestion (i.e., the flow to capacity ratio) over all arcs, which is a good measure for the overall network service quality.

The task of designing and dimensioning an USPR network may be formulated as a fixed charge network design problem (FC-USPR) or as a capacitated network design problem (Cap-USPR). In both problems, we are given additional arc capacities $c_a \in \mathbb{Z}_+$, $a \in A$, and arc costs $w_a \in \mathbb{Z}_+$, $a \in A$. In FC-USPR, the given capacities are regarded as fix. The task is to find a minimum cost arc set $F \subseteq A$ and an USPR for the given commodities within the subgraph (V, F), such that the induced arc flows do not exceed the capacities c_a. In Cap-USPR, the given values c_a are interpreted as basic capacity units that can be installed in integer multiples. Here, we seek for integer capacity multipliers $z_a \in \mathbb{Z}_+$, $a \in A$, and a corresponding USPR, such that the capacities $c_a z_a$ are not exceeded by the induced flows and the total capacity installation cost $\sum_{a \in A} w_a z_a$ is minimized.

In spite of its long history in practice, USPR has received attention in the mathematical literature only recently. Ben-Ameur and Gourdin [5, 6] and Broström and Holmberg [12, 13] study structural properties of (undirected) path sets where all paths are uniquely determined shortest paths for edge metric. Ben-Ameur and Gourdin also devise integer linear programming models to find a metric that induces a prescribed set of shortest paths (or prove that no such metric exists). Farago et al. [24, 25] study a special case of this Inverse Shortest Paths problem where the given paths are known to be shortest paths with respect to the number of edges and the task is to find lengths such that all these paths are unique shortest paths. Bley [9] shows that finding such a metric is computationally hard if the range of admissible link lengths is bounded. Algorithms based on local search techniques, Lagrangian relaxation, and integer programming methods as well as computational results for real-world network design and congestion minimization problems with USPR and multi-shortest path routing are discussed in [7, 8, 10, 11, 14, 22, 24, 26, 27, 32, 35, 40, 41]. Fortz and Thorup [27] show that it is $\mathcal{NP}$-hard to approximate the minimum congestion that can be obtained with multi-shortest path routing within a factor less than 3/2. Dinitz et al. [19], Kolliopoulos and Stein [34], and Skutella [46] study the approximability of (variants of) the unsplittable flow problem, while Lorenz et al. [36] discuss the relation of source-invariant routing to several other routing schemes. Results concerning the approximability of USPR problems have not been published (to our knowledge).

This paper is organized as follows. In Section 2, we formally define the three USPR problems addressed in this paper and discuss some of their basic properties.

Section 3 contains a comparison of USPR with several other routing schemes. We construct examples where the minimum congestion that can be obtained with USPR is a factor of $\Omega(|V|^2)$ larger than the minimum congestion that is achievable with unsplittable flow, shortest multi-path, or multicommodity flow routing, and a factor of $\Omega(|V|)$ larger than the congestion of an optimal unsplittable source-invariant routing. Furthermore, we show that the so-called no-bottleneck condition, which is typically assumed in unsplittable flow problems, has no effect on the complexity of unsplittable shortest path routing problems. This gives theoretical evidence for the practical experience that routing planning for USPR is harder than for these other routing paradigms.

In Section 4, we present new hardness results for the three USPR problems. We prove that it is $\mathcal{NP}$-hard to approximate the minimum congestion problem Min-Con-USPR within a factor of $\mathcal{O}(|V|^{1-\epsilon})$ (for any $\epsilon > 0$) and that the fixed charge network design problem FC-USPR is $\mathcal{NPO}$-complete. Furthermore, we show that the capacitated network design problem Cap-USPR is $\mathcal{NP}$-hard to approximate within a factor of $\mathcal{O}(2^{log^{1-\epsilon}|V|})$ in the directed and within a factor of $2-\epsilon$ in the undirected case.

In Section 5, we discuss approximation algorithms for Min-Con-USPR and Cap-USPR that work for general underlying (di)graphs. In the first part of this section we devise simple $|A|$- and $|K|$-approximation algorithms for Min-Con-USPR. In the second part, we show how to approximate the uniform and the single-source Cap-USPR problem within a factor $\mathcal{O}(|K|)$ and the undirected uniform Cap-USPR problem within a factor of $\mathcal{O}(\log|V|)$, using techniques that have been proposed for other capacitated network design problems. Finally, in Section 6, we present constant factor approximation algorithms for Min-Con-USPR and Cap-USPR for the special cases where the underlying graph is a bidirected ring or an undirected cycle. FC-USPR remains $\mathcal{NPO}$-complete even in these special cases.

Unless stated otherwise, our hardness results and algorithms hold also for the undirected problem versions, where both the underlying graph and the commodities are undirected. Table 1 summarizes the results of this paper.

2 Unsplittable shortest path routing problems

Let $D = (V, A)$ be a directed graph with arc capacities $c_a \in \mathbb{Z}_+$, $a \in A$, and let $K \subseteq V \times V$ be a set of directed commodities with demand values $d_{(s,t)} \in \mathbb{Z}_+$, $(s,t) \in K$.

For each commodity $(s,t) \in K$, let $\mathcal{P}(s,t)$ denote the set of all (s,t)-paths in D. Further-

Problem	Min-Con-USPR		Cap-USPR		FC-USPR
Hardness		$\Omega(\lvert V\rvert^{1-\epsilon})$	undirected:	$2-\epsilon$	$\mathcal{NPO}$-complete,
			directed:	$\Omega(1^{\log^{1-\epsilon}\lvert V\rvert})$	i.e., $\Omega(2^{\lvert V\rvert^{1-\epsilon}})$
Approx.	general:	$\min\{\lvert A\rvert, \lvert K\rvert\}$	general:	-	-
	undir. cycle:	2	undir. cycle:	2	
	bidir. ring:	3	bidir. ring:	4	
			uniform:	$\lvert K\rvert$	
			single-source:	$\lvert K\rvert$	
			undir. uniform:	$\mathcal{O}(log\lvert V\rvert)$	

Table 1: Approximability of USPR problems.

more, let $\mathcal{P} := \bigcup_{(s,t)\in K} \mathcal{P}(s,t)$. For any path P, we write $a \in P$ or $v \in P$ to indicate that the arc $a \in A$ or the node $v \in V$ occurs in P. The concatenation of two paths $P_1 = (v_0^1, a_1^1, \ldots, v_l^1)$ and $P_2 = (v_0^2, a_1^2, \ldots, v_k^2)$ with $v_l^1 = v_0^2$ is denoted by $P_1 \oplus P_2 := (v_0^1, a_1^1, \ldots, v_l^1 = v_0^2, a_1^2, \ldots, v_k^2)$. For simplicity, we refer to a path $P = (v_0, a_1, v_1, \ldots, a_l, v_l)$ with only its node sequence $P = (v_0, v_1, \ldots, v_l)$ if the underlying digraph D is simple. For any path P and any arc length vector $\lambda = (\lambda_a)_{a\in A} \in \mathbb{R}_+^A$, we denote $\lambda(P) := \sum_{a\in P} \lambda_a$.

Definiton 2.1 *We say that a metric $\lambda = (\lambda_a) \in \mathbb{R}_+^A$ defines an* unsplittable shortest path routing (USPR) *for the commodity set K, if the shortest (s,t)-path with respect to λ is* uniquely determined *for each commodity $(s,t) \in K$. We denote the shortest (s,t)-path with respect to λ with $\Phi_{(s,t)}(\lambda)$.*

The demand of each commodity is routed unsplit along the respective shortest path. For a metric λ that defines an USPR, the total flow through an arc $a \in A$ is therefore

$$f_a(\lambda) := \sum_{(s,t)\in K:\, a\in\Phi_{(s,t)}(\lambda)} d_{(s,t)} \;. \tag{1}$$

The task in the minimum congestion unsplittable shortest path routing problem is to find a metric $\lambda \in \mathbb{Z}_+^A$ that defines an USPR for the given commodity set K. The objective is to minimize the maximum congestion $f_a(\lambda)/c_a$ over all arcs. Formally, this problem is defined as follows:

Problem:	Min-Con-USPR
Instance:	A digraph $D = (V, A)$ with arc capacities $c_a \in \mathbb{Z}_+$, $a \in A$, and a commodity set $K \subseteq V \times V$ with demands $d_{(s,t)} \in \mathbb{Z}_+$, $(s,t) \in K$.
Solution:	A metric $\lambda \in \mathbb{Z}_+^A$, such that the shortest (s,t)-path w.r.t. λ is uniquely determined for each commodity $(s,t) \in K$.
Objective:	$\min(\max_{a\in A} f_a(\lambda)/c_a)$, where $f_a(\lambda)$ is as defined in (1).

In the two network design problems, we are given additional arc costs $w_a \in \mathbb{Z}_+$, $a \in A$. The task in the fixed charge network design problem is to find a minimum cost arc set $F \subseteq A$ and a metric $\lambda \in \mathbb{Z}_+^F$, such that λ defines an USPR for the commodities K in the subgraph (V, F) and such that the induced arc flows f_a do not exceed the original capacities c_a.

Problem:	FC-USPR
Instance:	A digraph $D = (V, A)$ with arc capacities $c_a \in \mathbb{Z}_+$, $a \in A$, and arc costs $w_a \in \mathbb{Z}_+$, $a \in A$, and a commodity set $K \subseteq V \times V$ with demands $d_{(s,t)} \in \mathbb{Z}_+$, $(s,t) \in K$.
Solution:	An arc set $F \subseteq A$ and metric $\lambda \in \mathbb{Z}_+^F$, such that λ induces a unique shortest (s,t)-path $\Phi_{(s,t)}(\lambda)$ in (V, F) for each $(s,t) \in K$, and $f_a(\lambda) \leq c_a$ for all $a \in F$.
Objective:	$\min \sum_{a \in F} w_a$.

In the capacitated network design problem, we seek capacity multipliers $z_a \in \mathbb{Z}_+$, $a \in A$, and a metric $\lambda \in \mathbb{Z}_+^A$ which defines an USPR for the commodity set K, such that $f_a \leq c_a z_a$ for all $a \in A$. The goal is to minimize the total capacity installation cost $\sum_{a \in A} w_a z_a$.

Problem:	Cap-USPR
Instance:	A digraph $D = (V, A)$ with arc capacities $c_a \in \mathbb{Z}_+$, $a \in A$, and arc costs $w_a \in \mathbb{Z}_+$, $a \in A$, and a commodity set $K \subseteq V \times V$ with demands $d_{(s,t)} \in \mathbb{Z}_+$, $(s,t) \in K$.
Solution:	Capacity multipliers $z_a \in \mathbb{Z}_+$, $a \in A$, and a metric $\lambda \in \mathbb{Z}_+^A$, such that the shortest (s,t)-path w.r.t. λ is uniquely determined for each commodity $(s,t) \in K$, and $f_a(\lambda) \leq z_a c_a$ for all $a \in A$.
Objective:	$\min \sum_{a \in A} z_a w_a$.

In all three problems, we may assume without loss of generality that D contains an (s,t)-path for each commodity $(s,t) \in K$ and that D is simple: Loops cannot be contained in a uniquely determined shortest path and, if two parallel arcs were contained in two commodities' routing paths, then these paths would not be unique shortest paths. Furthermore, we may assume that there are no parallel commodities: If there were two or more parallel commodities from s to t, these would have to use the same (uniquely determined shortest) flow path in any unsplittable shortest path routing and, therefore, could be aggregated into one commodity.

Observe that, for any bijection $\text{idx} : A \leftrightarrow \{1, \ldots, |A|\}$, the metric $\lambda_a := 2^{\text{idx}(a)}$ induces unique shortest paths between all node pairs. Hence, any instance of Min-Con-USPR or of Cap-USPR has a feasible solution, provided that the underlying digraph D contains at least one (s,t)-path for each $(s,t) \in K$.

If the underlying graph D contains only one (s,t)-path for each $(s,t) \in K$, then all metrics $\lambda \in \mathbb{Z}_+^A$ define the same USPR. In this case, any metric defines an optimal solution for Min-Con-USPR or Cap-USPR, and FC-USPR is trivially solvable. The simplest non-trivial case is when the underlying digraph contains two paths for each commodity. Yet, already in this case all three USPR problems become (weakly) $\mathcal{NP}$-hard.

Theorem 2.2 Min-Con-USPR*,* FC-USPR*, and* Cap-USPR *are $\mathcal{NP}$-hard, even if the underlying graph is a bidirected ring.*

Proof. We construct a polynomial reduction from the Partition to the problem of solving Min-Con-USPR to optimality. The $\mathcal{NP}$-hardness of FC-USPR and Cap-USPR follow analogously. Given a set of items $i \in \{1, \ldots, k\}$ with sizes $d_i \in \mathbb{Z}_+$, the Partition problem

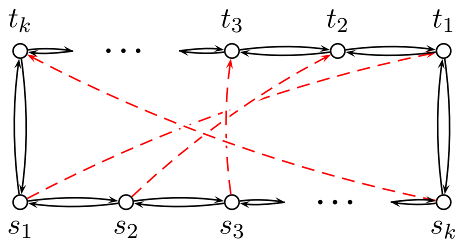

Figure 1: Reduction from PARTITION to MIN-CON-USPR: Solid lines are arcs, dashed lines are commodities.

is to find a subset $S \subseteq \{1, \dots, k\}$ with $\sum_{i \in S} d_i = 1/2 \sum_{i=1}^{k} d_i$ or to prove that no such subset exists. This problem is known to be $\mathcal{NP}$-complete; see Karp [33] or Garey and Johnson [28].

Given a PARTITION instance consisting of the items $i \in \{1, \dots, k\}$ with sizes $d_i \in \mathbb{Z}_+$, the instance of MIN-CON-USPR is built as shown in Figure 1. For each item $i \in \{1, \dots, k\}$, we introduce two nodes s_i, t_i and a commodity (s_i, t_i) with a demand value $d_{(s_i,t_i)} = d_i$. The arc set consists of the arcs (s_i, s_{i+1}), (s_{i+1}, s_i), (t_i, t_{i+1}), and (t_{i+1}, t_i) for all $i = 1, \dots, k-1$, as well as (s_1, t_k), (t_k, s_1) (s_k, t_1), and (t_1, s_k). The arc capacities are set to

$$c_{(s_j,s_{j+1})} = c_{(s_{j+1},s_j)} = c_{(t_j,t_{j+1})} = c_{(t_{j+1},t_j)} = \sum_{i=1}^{k} d_i \quad \text{for all } j = 1, \dots, k-1, \text{ and}$$

$$c_{(s_1,t_k)} = c_{(t_k,s_1)} = c_{(s_k,t_1)} = c_{(t_1,s_k)} = 1/2 \sum_{i=1}^{k} d_i \ .$$

It is not difficult to verify that any unsplittable flow routing of these commodities is also an unsplittable shortest path routing. Therefore, any feasible partition of the items corresponds to an unsplittable shortest path routing of the commodities such that the flows do not exceed the arc capacities, and vise versa. The commodities routed across arc (s_1, t_k) form one set of the partition, those routed across (s_k, t_1) the other set. □

One easily observes that all three USPR problems contain the DISJOINT PATHS problem as a special case. For general directed graphs, these problems therefore are actually $\mathcal{NP}$-hard in the strong sense, even if all demands and capacities are equal to one.

3 Relation to other routing schemes and the no-bottleneck condition

The unsplittable shortest path routing model is very restrictive and inherits structural properties of several other routing models. In this section, we compare USPR to four closely related but less restrictive routing models. We show that the minimal congestion that can be obtained with USPR for a given commodity set may exceed the congestion achievable with the other routing models by an arbitrarily large factor.

The most flexible routing model is (fractional) multicommodity flow routing. With this routing model, the demand of each commodity may be distributed arbitrarily and independent of the other commodities onto several flow paths. It thus admits the best possible use of the available capacities. In order to implement MCF routing in practice, the network must admit the configuration of arbitrary end-to-end routing paths and flow distributions for each commodity individually. This introduces many practical difficulties and complicates the network management. Therefore, many telecommunication network protocols are based on routing models that are less capacity efficient but easier to implement in practice.

With shortest multi-path routing, the traffic that is sent from a node s to a node t is distributed equally to all neighbors of s that are contained in any shortest (s,t)-path with respect to some metric λ. This routing model (adequately) describes so-called equal cost multi-path traffic splitting policies in shortest path routing protocols. Fortz and Thorup [27] show that the minimum congestion achievable with shortest multi-path routing cannot be approximated within a factor less than $3/2$.

The unsplittable flow routing model requires that each commodity is sent unsplit via a single path through the network. In contrast to the USPR model, the commodities' flow paths may be chosen independent of each other. Kolliopoulos and Stein [34] prove that it is $\mathcal{NP}$-hard to approximate the minimum congestion unsplittable flow routing within $2-\epsilon$, for any $\epsilon > 0$.

The fourth routing model that is closely related to USPR is unsplittable source-invariant routing. With this routing model, each commodity is routed on a single flow path. All flow paths with the same destination must form an anti-arborescence directed towards this destination. Once two flows meet on their way to a common destination, they cannot split anymore. This model describes the routing possibilities of packet networks with independently configurable store-and-forward routers. Lorenz et. al. [36] show that finding a minimum congestion unsplittable source-invariant routing is $\mathcal{NP}$-hard. They also show that the minimum congestion may be factor $\Omega(|V|)$ higher for unsplittable source-invariant routing than for unsplittable flow routing.

Given a digraph $D=(V,A)$ with arc capacities $c_a \in \mathbb{Z}_+$, $a \in A$, and commodities $K \subseteq V \times V$ with demand values $d_{(s,t)} \in \mathbb{Z}_+$, $(s,t) \in K$, we denote the optimal solution value of the MIN-CON-USPR problem by L^{USPR}. With L^{MCF}, L^{UFP}, L^{SMPR}, and L^{USIR} we refer to the minimal congestion values that can be obtained with fractional multicommodity flow routing (MCF), an unsplittable flow routing (UFP), a shortest multi-path routing (SMPR), and an unsplittable source-invariant routing (USIR) on the same instance, respectively. It is obvious that

$$\begin{aligned} L^{USPR} &\geq L^{UFP} &&\geq L^{MCF} \;, \\ L^{USPR} &\geq L^{SMPR} &&\geq L^{MCF} \;, \text{ and} \\ L^{USPR} &\geq L^{USIR} &&\geq L^{MCF} \;, \end{aligned}$$

since every unsplittable shortest path routing is also a valid shortest multi-path routing, a valid unsplittable flow routing, and a valid unsplittable source-invariant routing of the given commodities. In the following we construct instances where the gap between USPR and the other routing models becomes arbitrarily large.

Proposition 3.1 *There is a family of instances with*

(i) $L^{USPR} \geq \Omega(|V|^2) \cdot L^{SMPR}$,

(ii) $L^{USPR} \geq \Omega(|V|^2) \cdot L^{UFP}$, *and*

(iii) $L^{USPR} \geq \Omega(|V|^2) \cdot L^{MCF}$.

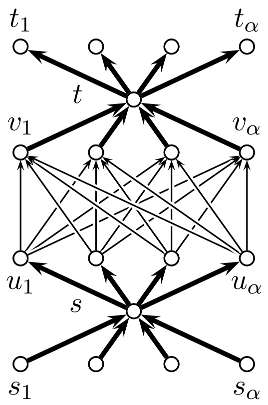

Figure 2: An instance with $L^{USPR} = \alpha^2$ and $L^{UFP} = L^{SMPR} = L^{MCF} = 1$. Arcs with capacity α are bold, arcs with capacity 1 are thin. For each node pair (s_i, t_j) there is a commodity with demand 1.

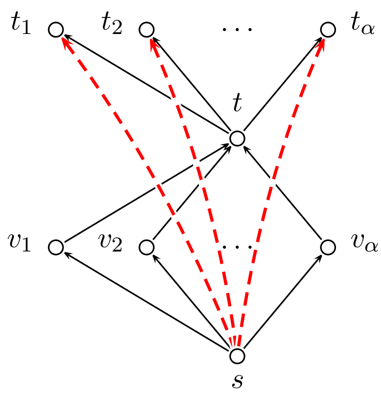

Figure 3: An instance with $L^{USPR} \geq \Omega(|V|)\, L^{USIR}$. The solid lines are arcs, the dashed lines are the commodities. All capacities and demands are one.

Proof. Let $\alpha \in \mathbb{Z}_+$ and consider the digraph $D = (V, A)$ illustrated in Figure 2. It consists of the nodes $V := \{s, t\} \cup \{s_i, t_i, u_i, v_i \; : \; i = 1, \dots, \alpha\}$ and the arcs $A := A_1 \cup A_2$, where $A_1 := \{(u_i, v_j) \; : \; i, j = 1, \dots, \alpha\}$ and $A_2 := \{(s_i, s),\, (s, u_i),\, (v_i, t),\, (t, t_i) \; : \; i = 1, \dots, \alpha\}$. The arc capacities are $c_a = 1$ for all $a \in A_1$ and $c_a = \alpha$ for all $a \in A_2$. In this graph, consider the commodities $K := \{(s_i, t_j) \; : \; i, j = 1, \dots, \alpha\}$ with demands $d_{(s_i,t_j)} = 1$ for all $(s_i, t_j) \in K$.

The congestion of any unsplittable shortest path routing is $L^{USPR} = \alpha^2$, since all commodities' routing paths must follow the same subpath between the nodes s and t in an unsplittable shortest path routing, and therefore share some arc (u_i, v_j) of capacity 1.

On the other hand, the congestion is 1 for an optimal shortest multi-path routing (where all arc lengths are chosen equal), as well as for an optimal unsplittable flow routing or an optimal multicommodity flow routing. For $\alpha \to \infty$, we obtain the claimed relations. $\square$

Proposition 3.2 *There is a family of instances with* $L^{USPR} \geq \Omega(|V|) \cdot L^{USIR}$.

Proof. Let $\alpha \in \mathbb{Z}_+$ and consider the digraph illustrated in Figure 3. It contains the nodes $V := \{s, t\} \cup \{t_i, v_i \; : \; i = 1, \dots, \alpha\}$ and the arcs $A := \{(s, v_i),\, (v_i, t),\, (t, t_i) \; : \; i = 1, \dots, \alpha\}$. All arcs have capacity 1. In this network, consider the commodities $K := \{(s, t_i) \; : \; i = 1, \dots, \alpha\}$ with demands $d_{(s,t_i)} = 1$. In any unsplittable shortest path routing, all commodities are routed via the same subpath between s and t. The minimal congestion value for USPR therefore is $L^{USPR} = \alpha$. With source-invariant routing, the commodities may be routed via different s,t-subpaths, as they have different destinations. The optimal congestion for this routing model therefore is 1, and with $\alpha \to \infty$ we obtain the claim. $\square$

The presented worst-case gaps between the different routing paradigms hold for the corresponding undirected routing variants, too. (In the undirected shortest multi-path routing policy and the undirected unsplittable source invariant routing policy, we arbitrarily choose s as the source and t as the destination for each undirected commodity $(s, t) \in K$.)

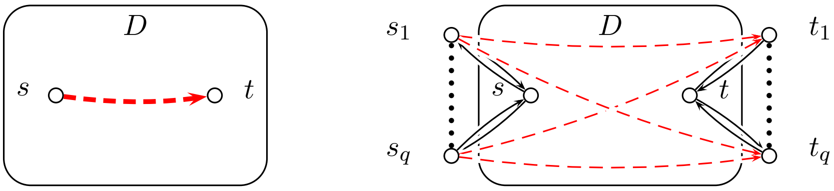

Figure 4: The no-bottleneck condition is irrelevant for USPR problems: The large commodity (s,t) can be replaced by many small commodities (s_i,t_j), which must share the same (s,t)-path.

An assumption commonly made for unsplittable flow problems is that the maximum demand value does not exceed the minimum capacity. Typically, unsplittable flow problems are easier to approximate if this additional condition holds than in the general case; cf. [19, 34, 46]. For unsplittable shortest path routing problems, however, this so-called no-bottleneck condition has no effect on the approximability.

Proposition 3.3 *For any instance* $\mathrm{I} = (D, c, K, d)$ *of* MIN-CON-USPR *with* $d_{\max} > c_{\min}$*, there exists an equivalent instance* $\mathrm{I}' = (D', c', K', d')$ *with* $d'_{\max} \leq c'_{\min}$ *(i.e., any solution for* I *with objective value* L *can be transformed into a solution for* I' *with objective value* L*, and vice versa).*

Proof. Suppose we are given a MIN-CON-USPR instance (D, c, K, d) with $d_{\max} > c_{\min}$. Let $r := d_{\max}/c_{\min}$ and $q := \lceil \sqrt{r}\, \rceil$. For each node $v \in V$, we introduce q additional nodes v_j and $2q$ arcs (v, v_j) and (v_j, v), $j = 1, \ldots, q$, see Figure 4. The capacities c' are given as $c'_a := c_a$ for all $a \in A$, and $c'_{(v,v_j)} = c'_{(v_j,v)} := \sum_{a\in A} c_a$ for all $j = 1, \ldots, q$. Each commodity $(s,t) \in K$ is replaced by q^2 new commodities (s_i,t_j) with demand values $d'_{(s_i,t_j)} \simeq d_{(s,t)}/q^2$. (Or more precisely, with demand values $d'_{(s_i,t_j)} \in \{\lfloor d_{(s,t)}/q^2 \rfloor, \lceil d_{(s,t)}/q^2 \rceil\}$ such that $\sum_{i,j=1}^{q} d'_{(s_i,t_j)} = d_{(s,t)}$.)

Clearly, $d'_{\max} \leq c'_{\min}$ holds.

Now, consider the set of all commodities (s_i,t_j), $i,j = 1, \ldots, q$, for some $(s,t) \in K$. Since all nodes s_i have only one neighbor s and all nodes t_j have only one neighbor t, all these q^2 commodities (s_i,t_j) must be routed via the same (s,t)-subpath in an unsplittable shortest path routing. Therefore, any unsplittable shortest path routing of the commodities K in D corresponds to an unsplittable shortest path routing of the commodities K' in D', and vice versa. As the corresponding routings induce the same flows on the arcs of D, the maximum congestion values are equal for both routings. □

Analogously, any instance of FC-USPR or of CAP-USPR can be transformed into an equivalent instance with $d_{\max} \leq c_{\min}$. However, note that in general this transformation is not polynomial in the strong sense, because the size of the underlying digraph grows by a factor of $\Theta(d_{\max}/c_{\min})$.

4 Inapproximability results

In the following section, we analyze how hard it is to approximate the three unsplittable shortest path routing problems.

4.1 The minimum congestion problem

We begin by showing that MIN-CON-USPR is not approximable within a factor of $\mathcal{O}(|V|^{1-\epsilon})$, unless $\mathcal{P} = \mathcal{NP}$. As a first step, we show that there is no constant factor approximation.

Lemma 4.1 *Let $\alpha \in \mathbb{Z}_+$ be an arbitrary number. It is $\mathcal{NP}$-hard to approximate* MIN-CON-USPR *within a factor less than $\alpha + 1$.*

Proof. We construct a reduction from the $\mathcal{NP}$-complete decision problem FULLY DISJOINT PATHS to MIN-CON-USPR. FULLY DISJOINT PATHS is a restricted variant of the classical DISJOINT PATHS problem [28]. Given a directed graph $H = (W, F)$ and a set of node pairs (s_i, t_i), $i = 1, \ldots, k$, the task is to find (s_i, t_i)-paths P_i in H that are not only internally disjoint but share no nodes at all (including the paths' first and last nodes), i.e., $\{v \in W : v \in P_i \text{ and } v \in P_j\} = \emptyset$ for all $i, j = 1, \ldots, k$ with $i \neq j$. It is easy to verify that the DISJOINT PATHS problem remains $\mathcal{NP}$-complete even with this stronger notion of disjointness. Note that the directed version of FULLY DISJOINT PATHS remains $\mathcal{NP}$-hard even if the number $k \geq 2$ of nodes pairs is not part of the problem input. Yet, we assume that k is part of the input of FULLY DISJOINT PATHS, because then our construction carries over literally to the undirected problem version.

Suppose we are given a FULLY DISJOINT PATHS instance consisting of the digraph $H = (W, F)$ and the node pairs (s_i, t_i), $i = 1, \ldots, k$. W.l.o.g., we may assume that there is an (s_i, t_i)-path in H for each $i = 1, \ldots, k$ and that $\{s_i, t_i\} \cap \{s_j, t_j\} = \emptyset$ for all $i \neq j$.

We construct a MIN-CON-USPR instance (D, c, K, d) as follows. The digraph $D = (V, A)$ contains all nodes and arcs of H. Furthermore, D contains one extra node r and $2k\alpha$ additional nodes u_i^l and v_i^l with $l = 1, \ldots, \alpha$ and $i = 1, \ldots, k$, i.e.,

$$V := W \cup \{r\} \cup \{u_i^l,\, v_i^l \;:\; i = 1, \ldots, k\,,\; l = 1, \ldots, \alpha\}.$$

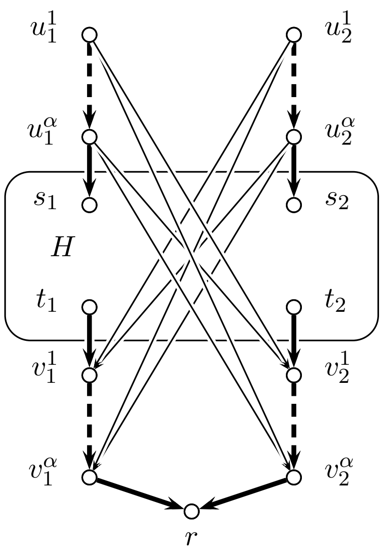

(a) Digraph D with indicated arc capacities: Arcs with capacity α are bold, arcs with capacity 1 are thin.

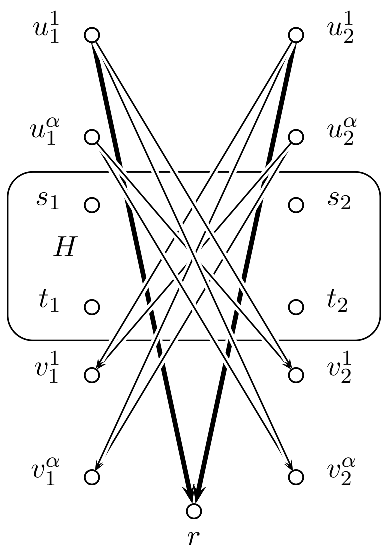

(b) Commodities K with indicated demand values: commodities with demand α are bold, commodities with demand 1 are thin.

Figure 5: Constructed MIN-CON-USPR instance.

For each pair $i, j = 1, \ldots, k$ with $i \neq j$ we add α^2 new arcs. These arcs form the set

$$A_1 := \Big\{(u_i^l, v_j^m) \;:\; i, j = 1, \ldots, k\,,\; i \neq j\,,\; l, m = 1, \ldots, \alpha\Big\}\,.$$

Additionally, we introduce $2\alpha + 1$ arcs for each $i = 1, \ldots, k$, which together comprise the arc set

$$\begin{aligned} A_2 := &\big\{(u_i^\alpha, s_i),\, (t_i, v_i^1),\, (v_i^\alpha, r) \;:\; i = 1, \ldots, k\big\} \\ &\cup \Big\{(u_i^l, u_i^{l+1}),\, (v_i^l, v_i^{l+1}) \;:\; i = 1, \ldots, k\,,\; l = 1, \ldots, \alpha - 1\Big\}\,. \end{aligned}$$

We let $A := F \cup A_1 \cup A_2$. The arc capacities are defined as

$$c_a := \begin{cases} 1, & \text{if } a \in A_1, \text{ and} \\ \alpha, & \text{otherwise.} \end{cases}$$

The commodity set K contains two types of commodities. For each $i = 1, \ldots, k$, there is a commodity (u_i^1, r) with demand value $d_{(u_i^1, r)} = \alpha$. For each pair $i, j = 1, \ldots, k$ with $i \neq j$ and each pair $l, m = 1, \ldots, \alpha$, there is a commodity (u_i^l, v_j^m) with $d_{(u_i^l, v_j^m)} = 1$. Figure 5 illustrates the constructed Min-Con-USPR instance for the case where $k = 2$.

It is obvious that this transformation is polynomial in the encoding size of the given Fully Disjoint Paths instance and α. Furthermore, any metric $\lambda \in \mathbb{Z}_+^A$ that induces unique shortest paths for the commodities in K defines a feasible solution for the constructed Min-Con-USPR instance.

In the first part of the proof, we show that there exists an unsplittable shortest path routing whose induced flows do not exceed the arc capacities if the given Fully Disjoint Paths instance has a feasible solution. Assume there exist fully disjoint (s_i, t_i)-paths P_i in H, $i = 1, \ldots, k$. Then we define the metric λ as

$$\lambda_a := \begin{cases} 1, & \text{if } a \in P_i \text{ for some } i \in \{1, \ldots, k\} \\ & \text{or if } a \in \bigcup_{i=1}^k \big\{(u_i^1, u_i^2), \ldots, (u_i^\alpha, s_i),\, (t_i, v_i^\alpha), \ldots, (v_i^2, v_i^1)\big\}, \text{ and} \\ |A|, & \text{otherwise.} \end{cases}$$

One easily finds that *all* shortest paths in D are unique with respect to λ. In particular, the shortest (u_i^1, r)-path is the path $\Phi_{(u_i^1, r)} = (u_i^1, \ldots, u_i^\alpha, s_i) \oplus P_i \oplus (t_i, v_i^\alpha, \ldots v_i^1, r)$ for each $i = 1, \ldots, k$, and the shortest (u_i^l, v_j^m)-path is $\Phi_{(u_i^l, v_j^m)} = (u_i^l, v_j^m)$ for each $i, j = 1, \ldots, k$ with $i \neq j$ and $l, m = 1, \ldots, \alpha$. Figure 6 illustrates this routing. It is not difficult to verify that the arc flows $f_a(\lambda)$ induced by this routing do not exceed the arc capacities c_a, $a \in A$.

In the second part of the proof, we show that the flows of any unsplittable shortest path routing exceed at least one arc capacity by a factor of at least $\alpha + 1$ if the Fully Disjoint Paths instance has no solution. So, suppose there is no set of fully disjoint (s_i, t_i)-paths in H and let $\lambda \in \mathbb{Z}_+^A$ be an arbitrary metric that defines unique shortest (u, v)-paths $\Phi_{(u,v)}$ for all $(u, v) \in K$.

First, assume that some arc (u_i^l, v_j^m) is contained in the shortest path $\Phi_{(u_h^1, r)}$ for some commodity (u_h^1, r). Since $\Phi_{(u_h^1, r)}$ is the unique shortest (u_h^1, r)-path, its arc (u_i^l, v_j^m) is also

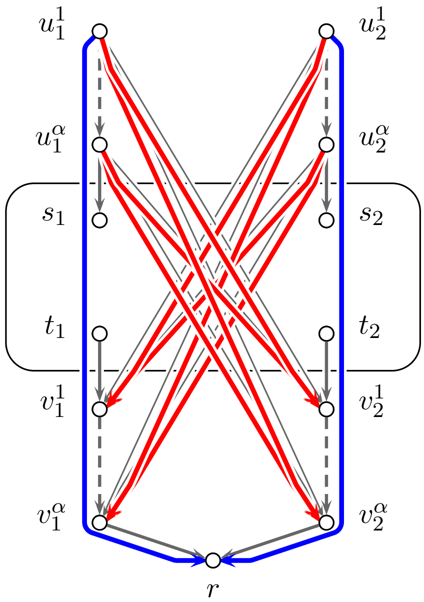

Figure 6: USPR in D if fully disjoint (s_i, t_i)-paths exist in H.

the unique shortest (u_i^l, v_j^m)-path. The total flow across this arc therefore is at least $d_{(u_h^1,r)} + d_{(u_i^l,v_j^m)} = \alpha + 1$, while its capacity is only $c_{(u_i^l,v_j^m)} = 1$. This leads to a congestion of at least $\alpha + 1$ for this routing. In the following, we thus may assume that all commodities (u_i^1, r), $i = 1, \dots, k$, with demands $d_{(u_i^1,r)} = \alpha$ are routed within the subgraph $D' = (V, F \cup A_2)$.

Now, suppose we had $t_i \in \Phi_{(u_i^1,r)}$ for all $i = 1, \dots, k$. Then at least two of these paths, say $\Phi_{(u_1^1,r)}$ and $\Phi_{(u_2^1,r)}$, would have to intersect in some internal node $w \in W$, as illustrated in Figure 7. Otherwise, there would exist fully disjoint (s_i, t_i)-paths in H. However, these two paths $\Phi_{(u_1^1,r)}$ and $\Phi_{(u_2^1,r)}$ cannot be unique shortest paths w.r.t. λ, as they contain different subpaths between w and r.

Consequently, there must be some $i \in \{1, \dots, k\}$ such that $t_i \notin \Phi_{(u_i^1,r)}$. Since $\Phi_{(u_i^1,r)}$ is completely contained in D', there must be some $j \neq i$ such that $t_j \in \Phi_{(u_i^1,r)}$. Furthermore, all nodes u_i^l and v_j^m, $l, m = 1, \dots, \alpha$, are contained in $\Phi_{(u_i^1,r)}$. Hence, all α^2 commodities (u_i^l, v_j^m) are routed along their respective subpath of $\Phi_{(u_i^1,r)}$, see Figure 8. The total flow across arc (t_j, u_j^α) therefore is at least $\alpha + \alpha^2$, while its capacity is only α.

Together, the two parts of the proof imply that it is $\mathcal{NP}$-hard to approximate Min-Con-USPR within a factor less than $\alpha + 1$. □

If we choose $\alpha = 1$ in the above construction, all capacities and demand values of the Min-Con-USPR instance are equal to 1. This yields the following theorem.

Theorem 4.2 *For any $\epsilon > 0$, it is $\mathcal{NP}$-hard to approximate* Min-Con-USPR *within a factor of $2 - \epsilon$, even if all demand values and capacities are equal to one.*

For the general case, we obtain a stronger non-constant inapproximability bound by choosing α depending on the size of the given Fully Disjoint Paths instance.

Theorem 4.3 *For any $\epsilon > 0$, it is $\mathcal{NP}$-hard to approximate* Min-Con-USPR *within a factor of $\mathcal{O}(|V|^{1-\epsilon})$.*

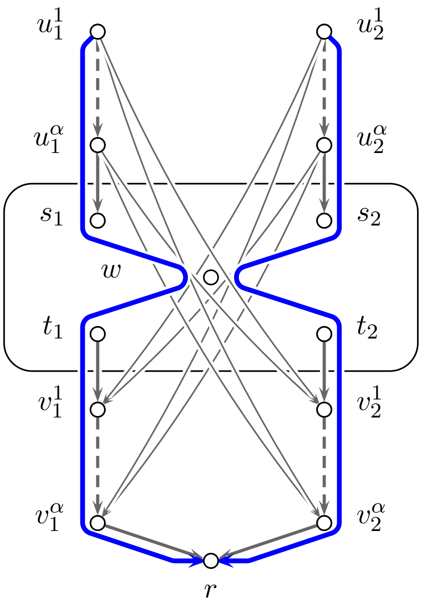

Figure 7: No USPR. If the shortest (u_1^1, r)- and (u_2^1, r)-paths intersect in some internal node w, they are not uniquely determined.

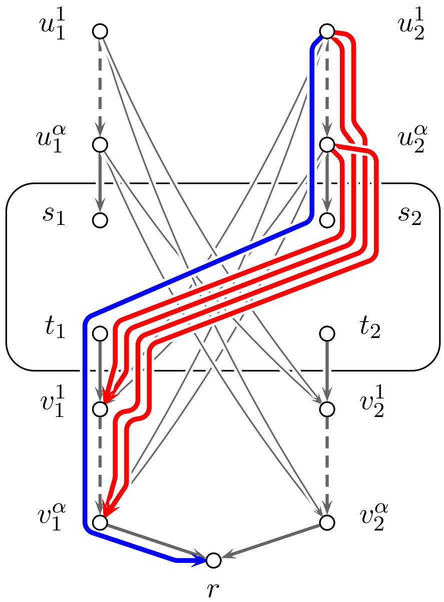

Figure 8: If $t_1 \in \Phi_{(u_2^1, r)}$, then the shortest path property forces all commodities (u_2^l, v_1^m) to follow their respective subpath of $\Phi_{(u_2^1, r)}$.

Proof. The encoding size of the constructed Min-Con-USPR instance is bounded by $\mathcal{O}(\alpha^2 \log \alpha(|W| + |F|))$. With $\alpha = \alpha(H) := |W|^q$, the presented construction thus remains polynomial in $|W| + |F|$ for any fixed $q \in \mathbb{Z}_+$. Because $|V| \in \Omega(\alpha)$, there exists some $q_\epsilon \in \mathbb{Z}_+$ for every $\epsilon > 0$, such that $\alpha \notin \mathcal{O}(|V|^{1-\epsilon})$ for $\alpha := |W|^{q_\epsilon}$. With Lemma 4.1, this implies the claim. $\square$

Analogously, it follows that approximating Min-Con-USPR within a factor of $\mathcal{O}(|A|^{1/2-\epsilon})$ or $\mathcal{O}(\langle \mathrm{I} \rangle^{1/2-\epsilon})$ is $\mathcal{NP}$-hard for any $\epsilon > 0$, where $\langle \mathrm{I} \rangle$ is the encoding size of the Min-Con-USPR instance (including the encoding size of the cost and capacity values).

By adding α many new nodes r^j, $j = 1, \ldots, \alpha$, and replacing each commodity (u_i^1, r) of demand $d_{(u_i^1, r)} = \alpha$ by α many commodities (u_i^1, r^j) with $d_{(u_i^1, r^j)} = 1$, we may transform the Min-Con-USPR instance constructed in the proof of Lemma 4.1 into an instance that satisfies $d_{\max} \leq c_{\min}$. For the special class of Min-Con-USPR instances constructed in the proof of Lemma 4.1, this transformation is strongly polynomial. Therefore, the inapproximability results of Lemma 4.1 and Theorem 4.3 also hold for the case where $d_{\max} \leq c_{\min}$.

4.2 The capacitated network design problem

A problem that is closely related to the capacitated network design problem Cap-USPR is the Generalized Steiner network problem, also known as Point-to-Point Connection problem: Given a (directed) graph $D = (V, A)$ and with arc costs $w_a \in \mathbb{Z}_+$, $a \in A$, and a set of commodities $K \subseteq V \times V$ find a minimum cost arc set $F \subseteq A$ such that the subgraph (V, F) contains an (s, t)-path for each $(s, t) \in K$.

With demand values $d_{(s,t)} := 1$ for all $(s, t) \in K$ and $c_a := |K|$ for all $a \in A$, the Generalized Steiner network problem reduces straightforward to the Cap-USPR problem.

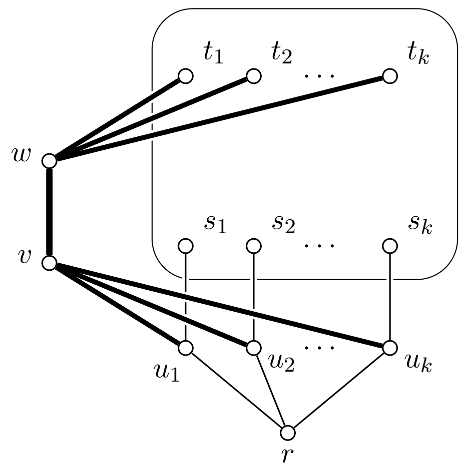

Figure 9: Constructed Cap-USPR instance. Edges with capacity $2M(k-1)$ and $2Mk(k-1)$ are bold, edges with capacity 1 are thin.

Inapproximability results for the directed and undirected Generalized Steiner network problem carry over immediately to the corresponding Cap-USPR problem.

Dodis and Khanna [20] showed that there is no polynomial time $\mathcal{O}(2^{log^{1-\epsilon}n})$ for the directed Generalized Steiner network problem for $\epsilon > 0$, unless $\mathcal{NP} \subseteq \text{DTIME}(n^{polylog(n)})$. This yields the strongest inapproximability threshold for the directed Cap-USPR problem currently known.

Theorem 4.4 ([20]) *For any $\epsilon > 0$, the directed* Cap-USPR *problem is inapproximable within $\mathcal{O}(2^{log^{1-\epsilon}|V|})$, unless $\mathcal{NP} \subseteq \text{DTIME}(n^{polylog(n)})$.*

For the undirected Generalized Steiner network problem, Goemans and Williamson [29] devised a simple primal dual algorithm that achieves a worst case performance guarantee of $2 - 1/|K|$, the best known inapproximability threshold is well below this number. [1] In the following, we prove a slightly stronger result for the undirected Cap-USPR problem.

Theorem 4.5 *For any $\epsilon > 0$, it is $\mathcal{NP}$-hard to approximate the undirected* Cap-USPR *problem within a factor of $2 - \epsilon$.*

Proof. We present a reduction similar to the one used in the proof of Lemma 4.1.

Suppose we are given an undirected Fully Disjoint Paths instance consisting of the graph $H = (W, F)$ and the node pairs (s_i, t_i), $i = 1, \ldots, k$. We may assume w.l.o.g. that there is an (s_i, t_i)-path in H for each $i = 1, \ldots, k$, and that $\{s_i, t_i\} \cap \{s_j, t_j\} = \emptyset$ for all $i \neq j$. Let $M := \lceil (2-\epsilon)k/\epsilon \rceil + 1$.

We construct an undirected Cap-USPR instance (G, c, w, K, d) as shown in Figure 9: The node set V of the graph $G = (V, E)$ contains all nodes in W, the three nodes v, w, and r, and

[1] Andrews [1] shows that there is no $\mathcal{O}(log^{1-\epsilon}|V|)$-approximation algorithm for the Buy-at-Bulk Network Design problem unless $\mathcal{NP} \subseteq \text{ZPTIME}(n^{polylog(n)})$. This proof can be adapted to show the same threshold for the undirected Cap-USPR problem. However, Andrews' construction inherently used randomization and relies on the probabilistic Erdös-Sachs theorem [21]; a deterministic construction that yields the same bound is not known.

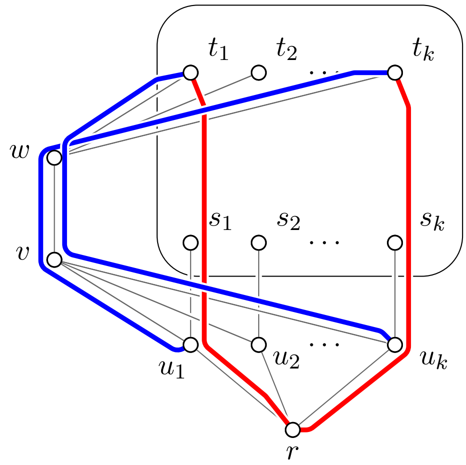

Figure 10: USPR in G if fully disjoint (s_i, t_i)-paths exist in H.

a node u_i for each $i = 1, \dots, k$. The edge set E consists of all edges in F, one edge vw, and the edges ru_i, $u_i s_i$, $u_i v$, and wt_i for each $i = 1, \dots, k$.

The commodity set K contains two types of commodities: For each $i = 1, \dots, k$, we introduce a commodity (r, t_i) with demand value $d_{(r,t_i)} := 1$. For all pairs $i, j = 1, \dots, k$ with $i \neq j$, we introduce a commodity (u_i, t_j) with a demand of $d_{(u_i,t_j)} := 2M$.

The edge capacities and costs are defined as

$$c_e := \begin{cases} 2M\,k(k-1), & \text{if } e = vw, \\ 2M\,(k-1), & \text{if } e \in \{u_i v,\ wt_i \ : \ i = 1, \dots, k\}, \\ 1, & \text{otherwise, and} \end{cases}$$

$$w_e := \begin{cases} M, & \text{if } e = vw, \\ 1, & \text{if } e \in \{u_i s_i \ : \ i = 1, \dots, k\}, \\ 0, & \text{otherwise.} \end{cases}$$

For any fixed $\epsilon > 0$, this construction is polynomial in the size of H.

First, suppose there exist fully disjoint (s_i, t_i)-paths P_i in H. Then we define the metric λ as

$$\lambda_e := \begin{cases} 1, & \text{if } e \in P_i \text{ for some } i \in \{1, \dots, k\} \\ & \text{or if } e \in \{ru_i,\ u_i s_i,\ u_i v,\ wt_i \ : \ i = 1, \dots, k\}, \\ |A|, & \text{if } e = vw, \text{ and} \\ |A| + 2, & \text{otherwise.} \end{cases}$$

It is easy to verify that this metric λ induces an USPR for the commodity set K, as illustrated in Figure 10. For each $i = 1, \dots, k$, the unique shortest (r, t_i)-path is $(r, u_i, s_i) \oplus P_i$, and for pair $i, j = 1, \dots, k$ with $i \neq j$, the unique shortest (u_i, t_j)-path is (u_i, v, w, t_j). If the corresponding commodities are routed along these paths, the induced edge flows f_e do not exceed the given capacities c_e, $e \in E$. Hence, the metric λ and the capacity multipliers $z_e = 1$ for all $e \in E$ form a feasible solution for the constructed Cap-USPR instance. The total capacity installation cost for this solution is

$$\sum_{e \in E} w_e z_e = \sum_{i=1}^{k} w_{u_i s_i} + w_{vw} = k + M \ .$$

Now, suppose there is no set of fully disjoint (s_i, t_i)-paths in H. Let $\lambda \in \mathbb{Z}_+^E$ be an arbitrary metric that defines an USPR for the commodities K. Analogous to the proof of Lemma 4.1, there are two possible cases: Either some commodity (u_i, t_j), $i \neq j$, is routed across an edge $u_l s_l$, or all commodities (u_i, t_j), $i \neq j$, and (at least) one commodity (r, t_i), $i \in \{1, \ldots, k\}$, are routed together across the edge vw.

In the first case, a commodity (u_i, t_j) with demand value $d_{(u_i,t_j)} = 2M$ is routed across an edge $u_l s_l$ with $c_{u_l s_l} = 1$. Then the capacity multiplier $z_{u_l s_l}$ must be at least $2M$ and, therefore, the cost of this solution is no less than $z_{u_l s_l} w_{u_l s_l} = 2M$.

In the second case, all $k(k-1)$ commodities (u_i, t_j), $i \neq j$, with $d_{(u_i,t_j)} = 2M$ and (at least) one commodity (r, t_l) with $d_{(r,t_l)} = 1$ are routed across the edge vw with $c_{vw} = 2Mk(k-1)$. Then the capacity multiplier z_{vw} must be at least two, which also yields a total solution cost of at least $2M$.

Hence, if there are no fully disjoint (s_i, t_i)-paths in H, then any feasible solution of the constructed Cap-USPR instance has a cost of at least $2M$.

Since it is $\mathcal{NP}$-hard to decide whether fully disjoint (s_i, t_i)-paths exist in H, it is also $\mathcal{NP}$-hard to approximate the optimal solution of the constructed undirected Cap-USPR instance within a factor strictly less than $2M/(k + M) > 2 - \epsilon$, as claimed. □

Theorem 4.5 also holds for the directed Cap-USPR problem, but the stronger general inapproximability threshold for the directed case follows from Theorem 4.4. With a construction similar to the one presented in the proof of Theorem 4.5 one can show that the directed Cap-USPR problem remains hard to approximate within $2 - \epsilon$ even if $|K| = 2$. For the special cases where the underlying digraph is a cycle or a bidirected ring, a variant of the reduction used in Theorem 2.2 yields a constant inapproximability threshold of 4/3 for the both the undirected and the directed Cap-USPR problem, respectively.

4.3 The fixed charge network design problem

Intuitively, the fixed charge problem FC-USPR is harder than the Cap-USPR problem, where the installation of arbitrarily large arc capacities is allowed. For any given Cap-USPR instance, there exists a feasible solution which can be easily found (provided that D contains an (s, t)-path for each $(s, t) \in K$). For FC-USPR, on the other hand, already the task of finding some feasible solution is $\mathcal{NP}$-hard, cf. Theorem 2.2. If we were given an FC-USPR instance with cost one for some arcs and prohibitively large costs for all others, then the core of the problem is to find an USPR in the subgraph induced by the edges of cost one. As this is an $\mathcal{NP}$-complete problem, we cannot expect to find a solution of FC-USPR with a reasonable quality guarantee in polynomial time. In the following, we prove this intuition.

Theorem 4.6 FC-USPR *is $\mathcal{NPO}$-complete.*

Proof. We present an approximation preserving reduction, to be more precisely a $\mathcal{PTAS}$-reduction [2, 18], from the Min-Weight-Sat problem.

The Min-Weight-Sat problem is defined as follows: Given a set X of boolean variables, a collection C of disjunctive clauses of at most three literals per clause, and a non-negative integer weight for each variable in X, the aim is to find a truth assignment for X that satisfies all clauses in C and minimizes the sum of the weights of the *true* variables. In Min-Weight-Sat(3), each variable occurs at most three times in total and at least once as a negated and once as an unnegated literal. It was shown by Orponen and Mannila [39] that Min-Weight-Sat is $\mathcal{NPO}$–complete. As for the unweighted satisfiability problem, the restricted problem Min-Weight-Sat(3) remains $\mathcal{NPO}$–complete. The restriction to 3 occurrences of each boolean variable is not necessary to prove Theorem 4.6, but it allows us to use a simpler reduction where all demand values and capacities are either 1 or 2.

Suppose we are given a Min-Weight-Sat(3) instance consisting of the boolean variables $x_1, \ldots, x_n$ with nonnegative weights $w_i \in \mathbb{Z}_+$, $1 \le i \le n$, and the clauses $C_1, \ldots, C_m$. We construct an instance (D, c, w, K, d) of FC-USPR, such that any truth assignment for the Min-Weight-Sat(3) instance corresponds to a solution of the FC-USPR instance whose cost is equal to the weight of the truth assignment.

The digraph $D = (V, A)$ contains the $6n + 2m$ nodes

$$V := \big\{q_i,\, v_i^1,\, v_i^2,\, \bar{v}_i^1,\, \bar{v}_i^2,\, r_i \;:\; i = 1, \ldots, n\big\} \cup \{s_h,\, t_h \;:\; h = 1, \ldots, m\} \;.$$

Among these nodes, we introduce the arcs

$$\begin{aligned} A_x &:= \big\{(q_i, v_i^1),\, (v_i^1, v_i^2),\, (v_i^2, r_i),\, (q_i, \bar{v}_i^1),\, (\bar{v}_i^1, \bar{v}_i^2),\, (\bar{v}_i^2, r_i) \;:\; i = 1, \ldots, n\big\} \text{ , and} \\ A_C &:= \big\{(s_h, \bar{v}_i^1),\, (\bar{v}_i^2, t_h) \;:\; i = 1, \ldots, n \text{ and } h = 1, \ldots, m \text{ s.t. } x_i \text{ appears unnegated in } C_h\big\} \\ &\quad \cup \big\{(s_h, v_i^1),\, (v_i^2, t_h) \;:\; i = 1, \ldots, n \text{ and } h = 1, \ldots, m \text{ s.t. } x_i \text{ appears negated in } C_h\big\} \;. \end{aligned}$$

For each boolean variable x_i, the nodes indexed by i form a *variable subgraph* as shown in Figure 11(a). For each clause C_h, the nodes s_h, t_h, and either the nodes v_i^1, v_i^2 or the nodes $\bar{v}_i^1$, $\bar{v}_i^2$ with index i such that x_i occurs in C_h form a *clause subgraph* as shown in Figure 11(b).

The arc capacities and costs are defined as

$$c_a := \begin{cases} 2, & \text{if } a \in A_x, \\ 1, & \text{if } a \in A_C, \end{cases}$$

$$w_a := \begin{cases} w_i, & \text{if } a = (q_i, v_i^1),\ i \in \{1 \ldots, n\}, \text{ and} \\ 0, & \text{otherwise.} \end{cases}$$

The commodity set K consists of a commodity (q_i, r_i) with a demand value of $d_{(q_i, r_i)} = 2$ for each $i = 1, \ldots, n$, and of a commodity (s_h, t_h) with a demand $d_{(s_h, t_h)} = 1$ for each $h = 1, \ldots, m$.

Clearly, this construction is polynomial in the size of the given Min-Weight-Sat(3) instance.

In the first part of the proof, we show that, for each truth assignment that satisfies all clauses of the Min-Weight-Sat(3) instance, there exists a corresponding feasible solution of the constructed FC-USPR instance. So, let $x \in \{true, false\}^n$ be such a truth assignment.

For each clause C_h, at least one literal evaluates to *true*. The index of the corresponding binary variable is denoted by $i(h) \in \{1, \ldots, n\}$. If more than one literals evaluate to true in

(a) Variable graph for x_i.

(b) Clause graph for $C_h = (\bar{x}_i \vee \bar{x}_j \vee x_k)$.

Figure 11: Constructed FC-USPR instance.

C_h, then $i(h)$ w.l.o.g. denotes the lexicographically first one. We define the corresponding metric $\lambda = \lambda(x) \in \mathbb{Z}_+^A$ as

$$\lambda_a := \begin{cases} 1, & \text{if } a \in \bigcup_{i:\, x_i = true} \{(q_i, v_i^1),\ (v_i^1, v_i^2),\ (v_i^2, r_i)\} \\ & \text{or } a \in \bigcup_{i:\, x_i = false} \{(q_i, \bar{v}_i^1),\ (\bar{v}_i^1, \bar{v}_i^2),\ (\bar{v}_i^2, r_i)\} \\ & \text{or } a \in \bigcup_{h: C_h = (\bar{x}_i(h) \vee \ldots)} \{(s_h, v_{i(h)}^1),\ (v_{i(h)}^2, t_h)\} \\ & \text{or } a \in \bigcup_{h: C_h = (x_i(h) \vee \ldots)} \{(s_h, \bar{v}_{i(h)}^1),\ (\bar{v}_{i(h)}^2, t_h)\}, \text{ and} \\ 2, & \text{otherwise.} \end{cases}$$

This metric λ defines an USPR for the commodities in K. If the boolean variable x_i is *true*, then we route commodity (q_i, r_i) on path $P_i^+ := (q_i, v_i^1, v_i^2, r_i)$, otherwise on path $P_i^- := (q_i, \bar{v}_i^1, \bar{v}_i^2, r_i)$. If the boolean variable $x_{i(h)}$ occurs negated in C_h, then we route commodity (s_h, t_h) along the path $Q_i^h := (s_h, v_i^1, v_i^2, t_h)$. (Note that $\bar{x}_{i(h)}$ then is *true* by definition of $i(h)$.) Otherwise, if the boolean variable $x_{i(h)}$ occurs unnegated (and evaluates to *true*) in clause C_h, then commodity (s_h, t_h) is routed via path $\bar{Q}_i^h := (s_h, \bar{v}_i^1, \bar{v}_i^2, t_h)$. Figure 12 illustrates this routing.

Since each variable x_i occurs in at most two clauses negated and in at most two clauses unnegated in the given Min-Weight-Sat(3) instance, any arc (v_i^1, v_i^2) or $(\bar{v}_i^1, \bar{v}_i^2)$ is contained in at most two commodity routing paths. All other arcs $a \in A$ are contained in at most one shortest path with respect to λ. Hence, the flows that are induced by the corresponding USPR satisfy the given capacities.

Let $F \subseteq A$ be the set of arcs contained in the induced routing paths. Then (F, λ) defines a feasible solution of the FC-USPR instance. Clearly, $(q_i, v_i^1) \in F$ if and only if the boolean variable x_i is *true*. Hence, we have

$$\sum_{a \in F} w_a = \sum_{i:\ x_i = true} w_i \ ,$$

i.e., the weight of the truth assignment x is the same as the cost of its corresponding FC-USPR solution (F, λ). This immediately implies

$$w(\lambda^{opt}) \leq w(x^{opt}) \ . \tag{2}$$

In the second part of the proof, we show that any feasible solution of the constructed FC-USPR instance defines a truth assignment satisfying all clauses. Let (F, λ) be such a

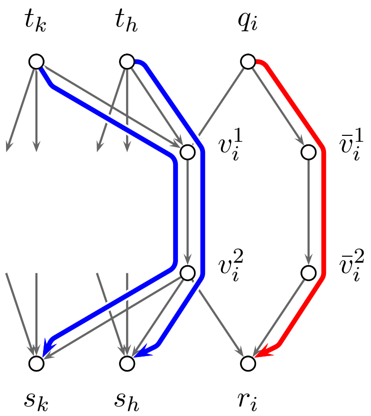

Figure 12: Partial routing in D corresponding to a truth assignment with $x_i = false$ and clauses $C_h = (\bar{x}_i \vee \ldots)$ and $C_k = (\bar{x}_i \vee \ldots)$.

feasible solution. First, observe that each commodity (q_i, r_i), $i \in \{0, \ldots, n\}$, is routed either on path P_i^+ or on path P_i^-. Any other (q_i, r_i)-path contains some arc $a \in A_H$, whose capacity $c_a = 1$ is insufficient to accommodate the demand $d_{(q_i,r_i)} = 2$ of commodity (q_i, r_i) . Thus, we can define the truth assignment $x = x(\lambda) \in \{true, false\}^n$ as

$$x_i := \begin{cases} true, & \text{if } P_i^+ \text{ is the (unique) shortest } (q_i, r_i)\text{-path w.r.t. } \lambda, \text{ and} \\ false, & \text{otherwise.} \end{cases}$$

If the shortest (s_h, t_h)-path w.r.t. λ contains the vertex v_i^1, then commodity (q_i, r_i) must be routed on the path P_i^- in the USPR defined by λ. Otherwise, the capacity of arc (v_i^1, v_i^2) would be violated. Analogously, commodity (q_i, r_i) must be routed via P_i^+ if $\bar{v}_i^1$ is contained in the routing path of some commodity (s_h, t_h). According to our construction, this implies that the corresponding clause C_h evaluates to *true* for the truth assignment x defined by the given routing. Hence, the constructed truth assignment x satisfies all clauses.

The arc set F of the given FC-USPR solution contains all arcs that belong to some of the routing paths. For each boolean variable x_i that was set to *true* in the constructed truth assignment x, the arc (q_i, v_i^1) is contained in the routing path for commodity (q_i, r_i) and therefore also belongs to the set F. Hence, we have

$$\sum_{a \in F} w_a \geq \sum_{i:\ x_i = true} w_i \ . \tag{3}$$

Together (2) and (3) imply that any ϵ-approximate solution of the constructed FC-USPR instance corresponds to an ϵ-approximate solution of the original MIN-WEIGHT-SAT(3) instance. Thus, the given reduction is a $\mathcal{PTAS}$-reduction from MIN-WEIGHT-SAT(3) to FC-USPR. As MIN-WEIGHT-SAT(3) is $\mathcal{NPO}$-complete, so is FC-USPR. $\square$

Theorem 4.6 immediately yields the following corollary:

Corollary 4.7 *For any $\epsilon > 0$, it is $\mathcal{NP}$-hard to approximate* FC-USPR *within a factor of $2^{\langle \mathrm{I} \rangle^{1-\epsilon}}$, where $\langle \mathrm{I} \rangle$ is the encoding size of the* FC-USPR *instance* I.

Note that any $\mathcal{NPO}$ problem can be approximated within a factor of $\mathcal{O}(2^{\langle \mathrm{I} \rangle^{\epsilon}})$ for some $\epsilon > 0$, if at least one feasible solution can be computed in polynomial time. This is due to the polynomial bound on the computation time of the objective function. Such problems belong to the class exp–$\mathcal{APX}$. As for FC-USPR already the problem of finding a feasible solution is $\mathcal{NP}$-hard, this problem is even harder than exp–$\mathcal{APX}$ problems.

A construction analogous to Theorem 2.2 yields that Theorem 4.6 and Corollary 4.7 hold even if the underlying digraph is a bidirected ring. Furthermore, Theorem 4.6 and Corollary 4.7 carry over to the fixed charge network design problem with unsplittable flow routing instead of unsplittable shortest path routing. In the above proof, we only used the requirement that all commodities are routed unsplit. Our construction implicitly guarantees the existence of a compatible metric for every unsplittable flow routing that satisfies the given capacities.

5 Approximation algorithms

In this section, we present polynomial time approximation algorithms for the minimum congestion problem Min-Con-USPR and for the capacitated network design problem Cap-USPR for general underlying digraphs. The fixed charge network design problem FC-USPR is not approximable within any reasonable quality guarantee, unless $\mathcal{P} = \mathcal{NP}$.

5.1 The minimum congestion problem

By Theorem 2.2, Min-Con-USPR is $\mathcal{NP}$-hard to approximate within a factor of $\mathcal{O}(|V|^{1-\epsilon})$. In the following, we show how to compute $\min\{|K|, |A|\}$-approximate solutions.

We begin by showing how to approximate Min-Con-USPR within a factor of $|K|$. Note that, in contrast to the Generalized Steiner Network problem, this is not trivial: We have to ensure that there exist a compatible metric for the chosen routing paths, i.e., a metric such that each routing path is the unique shortest path between its terminals.

Definiton 5.1 *For each path $P \in \mathcal{P}$, let $c_{min}(P) := \min\{c_a : a \in P\}$ be the* thickness *of P.*

An obvious optimal solution for an USPR instance with only one commodity (s,t) is to route (s,t) via an (s,t)-path $\Phi_{(s,t)}$ of maximum thickness, i.e., a path with $c_{min}(\Phi_{(s,t)}) = \max\{c_{min}(P) : P \in \mathcal{P}(s,t)\}$. Choosing a maximum thickness (s,t)-path $\Phi_{(s,t)}$ for each commodity (s,t) therefore yields an unsplittable flow routing whose congestion is at most $|K|$ times the congestion of an optimal unsplittable flow routing in the multicommodity case.

Note that maximum thickness paths are not necessarily unique. Furthermore, it also may be impossible to enforce uniqueness by a small perturbation of the capacities. In order to guarantee the existence of a compatible metric for the chosen paths, we need to consider *all* capacities on the paths instead of only the bottleneck capacity.

Definiton 5.2 *For each path $P \in \mathcal{P}$, the capacity* pattern *$c_{seq}(P)$ of P is the non-decreasingly sorted sequence of its arc capacities, i.e., $c_{seq}(P) := (c_{a_1}, \ldots, c_{a_{|P|}})$ with $a_i \in P$, $a_i \neq a_j$, and $c_{a_i} \leq c_{a_{i+1}}$).*

The lexicographic order on the c_{seq}-sequences defines a prefix-monotone total order on the paths. In contrast to the order by path thickness, ties in the c_{seq}-order can be broken consistently by an arbitrarily small perturbation of the capacities (or via a secondary lexicographical ordering of the paths according to their arc numbers), for example. It is not difficult to verify that, if each path $\Phi_{(s,t)}$ is the unique c_{seq}-maximum (s,t)-path with respect to (a perturbation of) the arc capacities c, then there exists a metric $\lambda \in \mathbb{Z}_+^A$ such that each $\Phi_{(s,t)}$ is a unique shortest (s,t)-path w.r.t. λ. This leads to the following simple algorithm.

Algorithm 5.1 ThickestPath

1. Compute the c_{seq}-maximal (s,t)-path $\Phi_{(s,t)}$ for each commodity $(s,t) \in K$ with respect to (a perturbation) of the capacities c_a, $a \in A$.
2. Find a metric $\lambda \in \mathbb{Z}_+^A$ such that each path $\Phi_{(s,t)}$, $(s,t) \in K$, is the unique shortest (s,t)-path w.r.t. λ.
3. Return λ.

Theorem 5.3 *ThickestPath is an $|K|$-approximation algorithm for* Min-Con-USPR*.*

Proof. Obviously, any c_{seq}-maximal (s,t)-path is a maximum thickness (s,t)-path. Hence, as we route each commodity $(s,t) \in K$ on the c_{seq}-maximal path with respect to (a perturbation of) the capacities c_a, the paths $\Phi_{(s,t)}$ form a routing with congestion at most $|K|$ times the congestion of an optimal solution.

As the lexicographic order on the c_{seq}-sequences defines a prefix-monotone total order on the paths, the c_{seq}-maximal (s,t)-paths can be found in polynomial time using a standard labeling algorithm. If the ties in the c_{seq}-order are broken consistently (by an appropriate perturbation), these paths form an unsplittable shortest path path routing. A compatible metric then can be computed in polynomial time using the linear programming and rounding approaches discussed in [6, 9]. □

Our second approximation algorithm, whose performance guarantee is independent of the number of commodities, is based on the multicommodity flow relaxation of the Min-Con-USPR problem. In a fractional multicommodity flow (MCF) routing, the demand of each commodity (s,t) may be distributed arbitrarily among the paths $\mathcal{P}(s,t)$. Such a routing can be expressed as an assignment $x : \mathcal{P} \to [0,1]$, where each x_P denotes the fraction of the demand $d_{(s,t)}$ that is sent along $P \in \mathcal{P}(s,t)$. The problem of finding an MCF routing of minimal congestion can be formulated as a linear program as follows:

$$\min\ L \qquad \text{(CON-LP)}$$

$$\sum_{P \in \mathcal{P}(s,t)} x_P = 1 \qquad \forall\, (s,t) \in K$$

$$\sum_{(s,t) \in K} \sum_{P \in \mathcal{P}(s,t): a \in P} d_{(s,t)} \cdot x_P \le L \cdot c_a \qquad \forall\, a \in A \tag{4}$$

$$L \ge 1$$

$$0 \le x_P \le 1 \qquad \forall\, (s,t) \in K,\ P \in \mathcal{P}(s,t)$$

Using column generation techniques, (CON-LP) can be solved in polynomial time w.r.t. to the size of the given problem instance, even though it contains exponentially many path variables. Let (L^*, x^*) be an optimal solution of (CON-LP). Clearly, L^* is a lower bound for the minimum congestion that can be obtained with an unsplittable shortest path routing.

The total flow across an arc $a \in A$ in the corresponding multicommodity flow routing is

$$f_a(x^*) := \sum_{(s,t)\in K} \sum_{P\in\mathcal{P}(s,t):a\in P} d_{(s,t)} x_P^* \ .$$

Let $\pi_a^* \in \mathbb{R}_+$, $a \in A$, be the optimal dual variables corresponding to the constraints (4) in (CON-LP). It follows from LP duality that all paths P with $x_P^* > 0$ are shortest paths between their respective terminal nodes with respect to the metric π^*. However, these paths are not necessarily *uniquely determined* shortest paths, and not all shortest paths w.r.t. π^* have a positive flow. The idea behind algorithm PenalizeSmallLinks is to start with the optimal dual variables π_a^* and perturb this metric such that the shortest path is unique for each commodity.

Algorithm 5.2 PenalizeSmallLinks

1. Compute optimal solution (L^*, x^*) of (CON-LP).
   Let π_a^*, $a \in A$, be the optimal dual values for (4).
2. Find *integer* lengths $\lambda'_a \geq 1$ that induce the same shortest paths as π_a^*.
   (In particular, all paths P with $x_P^* > 0$ are shortest paths w.r.t. λ.)
3. Number the arcs $idx : A \to \{1, \ldots, |A|\}$ in order of non-increasing $f_a(x^*)$.
4. Return the metric $\lambda_a := 2^{|A|+1} \cdot \lambda'_a + 2^{idx(a)}$.

Theorem 5.4 *PenalizeSmallLinks is an $|A|$-approximation algorithm for* MIN-CON-USPR.

Proof. It is easy to see that there exist an integer-valued metric $\lambda' \in \mathbb{Z}_+^A$ that induces exactly the same shortest paths as the given fractional metric $\pi^* \in \mathbb{R}_+^A$ (whether or not uniqueness is an issue). Such an integer-valued metric can be computed in polynomial time with linear programming based scaling and rounding methods, as shown in [6, 9]. Furthermore, for any integer valued metric $\lambda' \in \mathbb{Z}_+^A$ and any arc numbering $idx : A \to \{1, \ldots, |A|\}$, all shortest (s,t)-paths in D are unique with respect to the 'perturbed' metric $\lambda_a := 2^{|A|+1} \cdot \lambda'_a + 2^{idx(a)}$. Hence, the metric λ defined in Step 4 of algorithm PenalizeSmallLinks defines an USPR. It remains to show that the congestion induced of this USPR is at most $|A|$ times the congestion of the optimum MCF routing.

Let $f_a(x^*)$ and $f_a(\lambda)$ denote the arc flows induced by the optimal solution (L^*, x^*) of (CON-LP) and by the USPR for metric λ, respectively. Compared to the optimal MCF flow, an arc a receives additional traffic in the USPR routing only from those arcs a' with $idx(a') > idx(a)$. Hence, we have

$$f_a(\lambda) \leq \sum_{a':\ idx(a')\geq idx(a)} f_{a'}(x^*) \leq \sum_{a':\ f_{a'}(x^*)\leq f_a(x^*)} f_{a'}(x^*) \leq |A| \cdot f_a(x^*) \quad \text{for each } a \in A.$$

Consequently, $L^* \leq \max_{a\in A} f_a(\lambda)/c_a \leq \max_{a\in a} |A| \cdot f_a(x^*)/c_a = |A| \cdot L^*$. □

In Proposition 3.1, we have already seen that the ratio of $\Omega(|A|)$ between the optimal congestion values for USPR and MCF routings may be attained. Together, Theorems 5.3 and 5.4 yield the following corollary.

Corollary 5.5 MIN-CON-USPR *is approximable within a factor of* $\min\{|K|, |A|\}$.

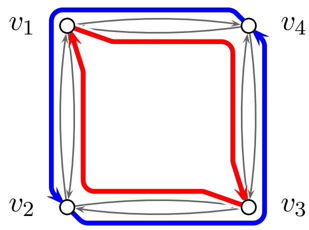

Figure 13: Optimal basic solution of (CON-LP) for Example 5.6.

Note that the routing obtained by rounding the optimal fractional multicommodity flow routing x^* is not necessarily a valid USPR. In the following example, all basic optimal solutions of (CON-LP) are integer and thus form unsplittable flow routings, but none of these routings is an USPR.

Example 5.6 *Consider the bidirected ring $D = (V, A)$ consisting of the four nodes $V = \{v_1, \ldots, v_4\}$ and the eight arcs $A = \{(v_1, v_2), (v_2, v_1), \ldots, (v_1, v_4)\}$. Let $c_a = 1$ for all $a \in A$, and consider the four commodities $K := \{(v_1, v_3), (v_3, v_1), (v_2, v_4), (v_4, v_2)\}$ with demand values $d_{(s,t)} = 1$ for all $(s,t) \in K$.*

For this instance of MIN-CON-USPR*, all basic optimal solutions of* (CON-LP) *are integer and, up to symmetry, correspond to the unsplittable flow routing illustrated in Figure 13. However, there exists no metric such that all paths of such a routing are unique shortest paths. One easily verifies that any USPR of the given commodities induces a flow of two or more on some arc of the bidirected ring.*

Algorithms ThickestPath and PenalizeSmallLinks carry over straightforward to the undirected version of MIN-CON-USPR, where they have worst case approximation ratios of $|K|$ and $|E|$, respectively. In Section 6, we show that PenalizeSmallLinks is a achieves a constant approximation guarantee of 2 in the special case where the underlying graph is an undirected cycle.

5.2 The capacitated network design problem

As many other capacitated network design problems, the CAP-USPR problem seems to be very hard to approximate. No algorithms with non-trivial quality guarantees are known for general arc capacities and arbitrary commodities.[2] However, the uniform and the single-source version of the problem can be approximated within reasonable bounds.

We say that a CAP-USPR problem is *uniform* if all all arc (or edge) capacities are identical. For the uniform CAP-USPR problem with only one commodity, an optimal solution obviously is given by routing this commodity along a shortest (s,t)-path with respect to the arc costs w and installing sufficiently many capacity units on the arcs of this path. Sending each commodity (s,t) along a minimum cost (s,t)-path thus trivially yields a worst case quality guarantee of $|K|$ for the uniform CAP-USPR problem with multiple commodities. As such a routing can be easily realized as an unsplittable shortest path routing by setting

[2] The $2^{\mathcal{O}(\sqrt{\log|V|\log\log|V|})\log d_{max}}$-approximation algorithm of Charikar and Karagiozova [16] for the BUY-AT-BULK NETWORK DESIGN problem does not carry over to the CAP-USPR problem, as the routings computed by this algorithm do not necessarily form unsplittable shortest path routings.

the routing metric λ to an appropriate perturbation of the arc costs w, we get the following result.

Proposition 5.7 *The uniform* CAP-USPR *problem is approximable within a factor of* $|K|$.

Clearly, Proposition 5.7 holds for both the directed and the undirected problem version. Note, however, that the above approach is not applicable for the non-uniform CAP-USPR problem. It is not difficult to construct examples where overlaying the optimal routing paths of the single commodity problems does not yield an unsplittable shortest path routing for the corresponding multicommodity problem.

For the undirected uniform CAP-USPR problem, the trivial $|K|$ bound can be improved to $\mathcal{O}(\log|V|)$ using a probabilistic approximation of the arc costs by dominating tree metrics.[3] Given an undirected graph $G = (V, E)$, the metric $\lambda \in \mathbb{R}_+^E$ is a tree metric if there exists a tree T in G such that all shortest paths with respect to λ are fully contained in T. A metric λ dominates another metric μ if $dist_\lambda(s,t) \geq dist_\mu(s,t)$ for all $(s,t) \in V^2$. The stretch of a dominating metric λ with respect to a metric μ is $stretch(\lambda, \mu) := \max\{dist_\lambda(s,t)/dist_\mu(s,t) \,:, (s,t) \in V^2\}$. Bartal [4] showed that any metric in an undirected graph can be probabilistically approximated by a distribution over dominating tree metrics such that the expected stretch is $\mathcal{O}(\log^2|V|)$. This result was later improved by Fakcharoenphol et al. [23] to an expected stretch of only $\mathcal{O}(\log|V|)$. Charikar et al. [15] showed how to derandomize this probabilistic approximation, i.e., how to approximate a metric with a distribution over only polynomially many tree metrics.

Based on these results, Awerbuch and Azar [3] proposed an approximation algorithm for the undirected uniform BUY-AT-BULK NETWORK DESIGN problem. As all solutions computed by this algorithm are trees, it carries over directly to the undirected uniform CAP-USPR problem. The derandomized version of Awerbuch and Azar's algorithm works as follows: First, we compute polynomially many tree metrics $\lambda_i \in \mathbb{R}_+^E$, $i \in I$, that probabilistically approximate the given arc costs w. For each tree metric, we then compute the cost of the solution that is given by routing each commodity along a shortest path with respect to λ_i and installing sufficient edge capacities. At the end, we return the best of these solutions. The performance guarantee of $\mathcal{O}(\log|V|)$ follows straightforward from [3] and [23].

Theorem 5.8 ([3],[23]) *The undirected uniform* CAP-USPR *problem is approximable within a factor of* $\mathcal{O}(\log|V|)$.

Note that in planar graphs any metric can be probabilistically approximated by tree metrics with constant expected stretch. Hence, Awerbuch and Azar's algorithm yields a constant worst-case guarantee for the undirected uniform CAP-USPR problem on planar graphs.

For the non-uniform and for the directed CAP-USPR problem, the technique of using probabilistic approximations by tree metrics utterly fails.

Another interesting variant of the problem is the *single-source* version, where all commodities share the same source terminal. Single-source network design problems have been considered in the literature for various capacity and routing paradigms. Most proposed solution techniques, however, enforce that the routing paths form a tree and, therefore, can be applied directly for the single-source CAP-USPR problem.

[3] Alternatively, also the technique of approximating the underlying graph by a so-called light-weight distance-preserving spanner with respect to its edge cost function can be applied to obtain a $\mathcal{O}(\log|V|)$-approximation algorithm for the undirected uniform CAP-USPR problem, cf. Mansour and Peleg [37].

A straightforward approach to compute a solution for the single-source problem is to iteratively assign to each commodity the cheapest path such that the new path together with the already assigned paths forms an arborescence. If the commodities are considered in order of decreasing demands, this simple algorithm achieves a worst-case guarantee of $|K|$. Clearly, this approach works for both the directed and the undirected single-source problem.

Proposition 5.9 *The single-source* CAP-USPR *problem is approximable within a factor of $|K|$ in general.*

For the general directed single-source CAP-USPR problem, no better approximation algorithm is known. The undirected problem version can be solved using the algorithms proposed by Guha et al. [30], Gupta et al. [31], Talwar [47], or Meyerson et al. [38] for similar network design problems. These algorithms yield a constant factor approximation for the uniform and an $\mathcal{O}(\log |K|)$-approximation for the non-uniform undirected single-source CAP-USPR problem.

6 Special cases

In this section, we present specialized algorithms that achieve constant factor approximation guarantees for MIN-CON-USPR and CAP-USPR in the special cases where the underlying graph is a bidirected ring or an undirected cycle. As mentioned above, FC-USPR remains $\mathcal{NPO}$-complete even in these special cases.

6.1 Min-Con-USPR on an undirected cycle

Algorithm PenalizeSmallLinks presented in Section 5 carries over straightforward to the undirected version of MIN-CON-USPR. In the special case where the underlying graph $G = (V, E)$ is an undirected cycle, there are only two possible routing paths for each commodity. For any edge $e \in E$, one of these two paths for each commodity contains e. Hence, perturbing the length of one minimum flow edge $e_{\min} := \arg\min f_e(x^*)$ suffices to ensure that all shortest paths are unique. This yields the following theorem.

Theorem 6.1 *Algorithm PenalizeSmallLinks achieves a 2-approximation guarantee for* MIN-CON-USPR *on an undirected cycle.*

Theorem 6.1 extends straightforward to the case where all blocks of the underlying undirected graph are cycles.

Cosares and Saniee [17] and Schrijver et. al. [45] propose algorithms for the undirected ring loading problem (which is equivalent to the minimum congestion unsplittable flow problem on a cycle with unit capacities) that are based on rounding the optimal solution of (CON-LP). This approach also works for the undirected MIN-CON-USPR problem if the rounding procedure is slightly adapted in order to guarantee that the resulting paths form an unsplittable shortest paths routing.

6.2 Cap-USPR on an undirected cycle

Let $G = (V, E)$ be an undirected cycle with edge capacities $c_e \in \mathbb{Z}_+$ and edge costs $w_e \in \mathbb{Z}_+$ for all $e \in E$, and let $K \subseteq V^{(2)}$ be a set of undirected commodities with demand values

$d_{(s,t)} \in \mathbb{Z}_+$ for all $(s,t) \in K$. Consider the following linear programming relaxation of the undirected Cap-USPR problem:

$$\min \sum_{e \in E} w_e z_e \tag{CAP-LP}$$
$$\sum_{P \in \mathcal{P}(s,t)} x_P = 1 \qquad \forall\,(s,t) \in K$$
$$\sum_{(s,t) \in K} \; \sum_{P \in \mathcal{P}(s,t): e \in P} d_{(s,t)} \cdot x_P \le c_e z_e \qquad \forall\, e \in E \tag{5}$$
$$z_e \ge 0 \qquad \forall\, e \in E$$
$$0 \le x_P \le 1 \qquad \forall\,(s,t) \in K,\ P \in \mathcal{P}(s,t)$$

The idea of our approximation algorithm is to round an optimal solution of (z^*, x^*) of (CAP-LP) to obtain an integer solution $([z],[x])$ of (CAP-LP) such that the corresponding routing paths $\mathcal{Q} := \{P : [x]_P = 1\}$ can be realized as an USPR.

It is well known (and follows directly from LP duality), that an optimal solution of (CAP-LP) can be constructed by routing all commodities on shortest paths w.r.t. the edge lengths w_e/c_e and installing exactly the capacities that are consumed by this routing. For an appropriate perturbation of these lengths, the shortest paths are unique and form an USPR. In such a special optimal solution (z^*, x^*) of (CAP-LP), all x_P variables are integer and correspond to an USPR. Only the z_e variables may attain fractional values and need to be rounded up. However, the optimal fractional solution values z^*_e may be arbitrarily small and rounding them all up may increase the cost by an arbitrarily large factor.

If we knew the topology of the optimal Cap-USPR solution in advance, then we could apply the above method in the subgraph defined by this topology. The additional cost of rounding up the capacity multipliers z^*_e then were bounded by the cost of the optimal solution. In the special case where the underlying graph is an undirected cycle, we do not need to now the optimal solution's topology in advance. We can simply enumerate all possible solution topologies, as demonstrated in algorithm EnumerateAndRound.

Algorithm 6.1 EnumerateAndRound

1. Compute solution (λ^0, z^0) as follows:
   1.1 Set $\lambda^0_e := \begin{cases} M w_e/c_e + 1, & \text{if } e = e_0, \\ M w_e/c_e, & \text{otherwise,} \end{cases}$
   for some arbitrary $e_0 \in E$ and $M := 2\Pi_{e \in E} c_e$.
   1.2 Set $z^0_e := \lceil f_e(\lambda^0)/c_e \rceil$ for all $e \in E$.
2. For each $l \in E$, compute solution (λ^l, z^l) as follows:
   2.1 Set $\lambda^l_e := \begin{cases} |E|, & \text{if } e = l, \\ 1, & \text{otherwise.} \end{cases}$
   2.2 Set $z^l_e := \lceil f_e(\lambda^l)/c_e \rceil$ for all $e \in E$.
3. Return minimum cost solution of (λ^0, z^0) and (λ^l, z^l), $l \in E$.

Theorem 6.2 *EnumerateAndRound is a 2-approximation algorithm for* Cap-USPR *on an undirected cycle.*

Proof. The metrics λ^l, $l \in E$, clearly induce uniquely determined shortest paths between all node pairs. The metric λ^0 defined in Step 1.2 is a perturbation of the metric w_e/c_e. Since $\lambda^0_{e_0}$ is odd and all other lengths λ^0_e are even, also λ^0 defines unique shortest paths between all node pairs. Hence, all metrics λ^0 and λ^l, $l \in E$, define valid USPRs, and the corresponding solutions (λ^0, z^0) and (λ^l, z^l) are feasible.

Let (λ^{opt}, z^{opt}) be the optimal solution of the given CAP-USPR instance.

First, assume that $z_e^{opt} \geq 1$ for all $e \in E$. In this case, we have

$$\begin{aligned} w(z^0) &= \sum_{e \in E} w_e z_e^0 = \sum_{e \in E} w_e \lceil f_e(\lambda^0)/c_e \rceil \\ &\leq \sum_{e \in E} w_e \left(f_e(\lambda^0)/c_e + z_e^{opt} \right) = w(z^*) + w(z^{opt}) \leq 2w(z^{opt}) \ , \end{aligned}$$

which implies that (λ^0, z^0) is a 2-approximate solution.

If this is not the case, we have $z_l^{opt} = 0$ for some $l \in E$. As $G - l$ is a path, the routing of all commodities is uniquely determined in this case. Hence, the metric λ^l constructed in algorithm EnumerateAndRound and the optimal solution's metric λ^{opt} induce the same shortest paths, and thus define the same USPR. Clearly, z^l is a minium cost capacity installation for this routing. Therefore, (λ^l, z^l) is an optimal solution in this case.

Consequently, EnumerateAndRound is a 2-approximation algorithm. $\square$

Theorem 6.2 generalizes straightforward to undirected CAP-USPR instances where all blocks of the underlying graph are cycles.

6.3 Min-Con-USPR on a bidirected ring

For MIN-CON-USPR on a bidirected ring, neither the perturbation technique used in algorithm PenalizeSmallLinks nor naive rounding of an optimal solution of the linear programming relaxation leads to a constant factor approximation: The perturbation technique produces $|A|/2$-approximate solutions in the worst case, and the rounding approach may find non-USPR integer routings as illustrated in Example 5.6.

The idea of our algorithm is to remove some arc from the given bidirected ring and round the optimal fractional routing in the residual digraph. The following two lemmas show that any routing obtained this way is a valid USPR.

Definiton 6.3 *Two paths $P_1, P_2 \in \mathcal{P}$ are said to be conflicting, if there are two nodes $s, t \in V$ such that P_1 and P_2 both contain an (s,t)-subpath $P_1[s,t]$ and $P_2[s,t]$, respectively, and $P_1[s,t] \neq P_2[s,t]$.*

Lemma 6.4 *Let $D = (V, A)$ be a bidirected ring. Then there exists an optimal solution (L^*, x^*) of* (CON-LP) *such that $x^*_{P_1} = 0$ or $x^*_{P_2} = 0$ for any pair of conflicting paths $P_1 \in \mathcal{P}(s_1, t_1)$ and $P_2 \in \mathcal{P}(s_2, t_2)$ with $(s_1, t_1) \neq (s_2, t_2)$. Furthermore, such a solution (L^*, x^*) can be found in polynomial time.*

Proof. Suppose we have an optimal solution x^* of (CON-LP) with $x^*_{P_1} > 0$ and $x^*_{P_2} > 0$ for two conflicting paths $P_1 \in \mathcal{P}(s_1, t_1)$ and $P_2 \in \mathcal{P}(s_2, t_2)$ with $(s_1, t_1) \neq (s_2, t_2)$. Let $\bar{P}_1$ be the opposite (s_1, t_1)-path to P_1 and let $\bar{P}_2$ be the opposite (s_2, t_2)-path to P_2. Since P_1 and P_2 conflict, we have $\bar{P}_1 \subsetneq P_2$ and $\bar{P}_2 \subsetneq P_1$. We may assume w.l.o.g. that $x^*_{P_1}$ carries less flow, i.e., $d_{(s_1,t_1)} x^*_{P_1} \leq d_{(s_2,t_2)} x^*_{P_2}$. Let $\alpha := d_{(s_1,t_1)} x^*_{P_1}$.

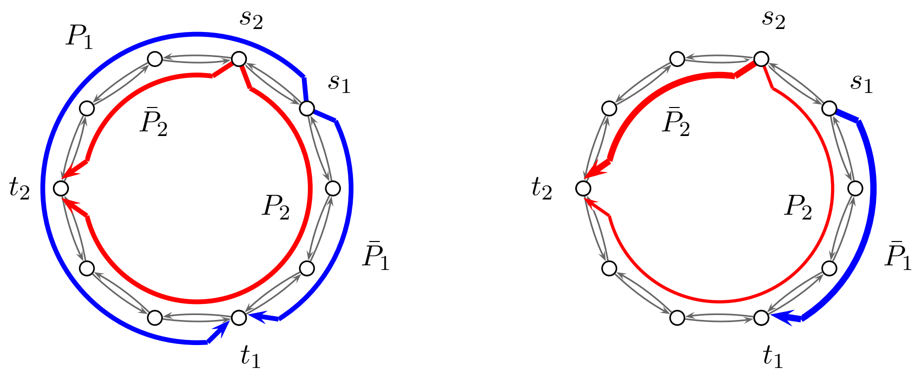

(a) Original routing x^*. (b) Uncrossed routing x'.

Figure 14: Uncrossing the routing of two parallel commodities.

Then we construct another solution x' of (CON-LP) by 'uncrossing' the routing of the two commodities (s_1, t_1) and (s_2, t_2), as shown in Figure 14. For commodity (s_1, t_1), we shift the entire flow of value α from path P_1 to its opposite path $\bar{P}_1$. Simultaneously, we also shift a flow of value α from P_2 to $\bar{P}_2$ for commodity (s_2, t_2). Formally, x' is given as

$$
\begin{aligned}
x'_{P_1} &:= 0 \,, & x'_{\bar{P}_1} &:= 1 \,, \\
x'_{P_2} &:= x^*_{P_2} + \alpha/d_{(s_2,t_2)} \,, & x'_{\bar{P}_2} &:= x^*_{\bar{P}_2} - \alpha/d_{(s_2,t_2)} \quad \text{, and} \\
x'_P &:= x^*_P \quad \text{, for all } P \notin \{P_1, \bar{P}_1, P_2, \bar{P}_2\}.
\end{aligned}
$$

One easily verifies that $f_a(x') \leq f_a(x^*)$ for all $a \in A$, i.e., x' is also an optimal solution of (CON-LP). Furthermore, we have $\sum_{a \in A} f_a(x') < \sum_{a \in A} f_a(x^*)$. Thus, an optimal solution x^* of (CON-LP) which in addition minimizes $\sum_{a \in A} f_a(x)$ (over all optimal solutions of (CON-LP)) has the required properties.

We can find such a solution x^* as follows: First, we solve (CON-LP) to determine the optimal value L^*. Then, we solve (CON-LP) with an additional linear constraint $L \leq L^*$ and the objective function replaced by $\min \sum_{a \in A} f_a(x)$. The optimal solution x^* of this second linear program then has $x^*_{P_1} = 0$ or $x^*_{P_2} = 0$ for any pair of conflicting paths $P_1 \in \mathcal{P}(s_1, t_1)$ and $P_2 \in \mathcal{P}(s_2, t_2)$ with $(s_1, t_1) \neq (s_2, t_2)$. $\square$

Lemma 6.5 *Let $D = (V, A)$ be a bidirected ring and $a_0 \in A$. Let $\mathcal{Q} \subseteq \mathcal{P}$ be a set of paths that contains no pair of conflicting paths and such that $a_0 \notin P$ for all $P \in \mathcal{Q}$. Then there exists a compatible metric $\lambda \in \mathbb{Z}_+^A$ for $\mathcal{Q}$, i.e., a metric λ such that each path $P \in \mathcal{Q}$ is the unique shortest path between its terminals w.r.t. λ.*

Proof. W.l.o.g., we may assume that the nodes of D are labeled v_1 to v_n in a counter-clockwise manner and that $a_0 = (v_n, v_1)$. We denote $a_1 = (v_1, v_n)$.

Suppose $a_1 \notin P$ for all $P \in \mathcal{Q}$. Then all paths $P \in \mathcal{Q}$ are unique shortest paths for the metric

$$
\lambda_a := \begin{cases} |V|, & \text{if } a \in \{a_0, a_1\}, \text{ and} \\ 1, & \text{otherwise.} \end{cases}
$$

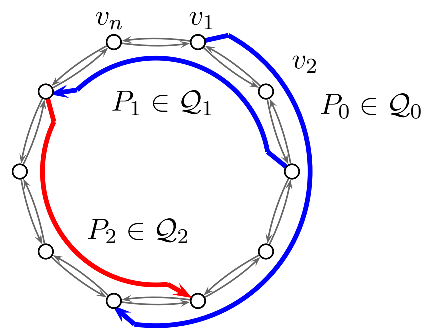

Figure 15: Path sets $\mathcal{Q}_0$, $\mathcal{Q}_1$, and $\mathcal{Q}_2$ in a bidirected ring.

So, we may assume that $a_1 \in P$ for some $P \in \mathcal{Q}$. We distinguish three types of paths in $\mathcal{Q}$:

$$\begin{aligned}
\mathcal{Q}_0 &:= \{P \in \mathcal{Q} : P = (v_i, v_{i+1}, \ldots, v_j) \text{ with } i < j\}\ , \\
\mathcal{Q}_1 &:= \{P \in \mathcal{Q} : a_1 \in P\}\ , \text{ and} \\
\mathcal{Q}_2 &:= \{P \in \mathcal{Q} : P = (v_j, v_{j-1}, \ldots, v_i) \text{ with } i < j\}\ .
\end{aligned}$$

The set $\mathcal{Q}_0$ consists of all clockwise oriented paths of $\mathcal{Q}$, the set $\mathcal{Q}_1$ of all counter-clockwise oriented paths of $\mathcal{Q}$ that contain a_1, and the set $\mathcal{Q}_2$ of all remaining counter-clockwise oriented paths of $\mathcal{Q}$, see Figure 15. As no path of $\mathcal{Q}$ contains a_0, the three sets $\mathcal{Q}_0$, $\mathcal{Q}_1$, and $\mathcal{Q}_2$ form a partition of $\mathcal{Q}$.

First, we show that there is a compatible metric for the smaller path set $\mathcal{Q}_0 \cup \mathcal{Q}_1$. For this, we consider these paths in an undirected setting. Let $G = (V, E)$ be the undirected cycle with $E := \{v_i v_{i+1} : i = 1, \ldots, n\}$, where $v_{n+1} = v_1$. Recall that $a_1 \in P$ for all $P \in \mathcal{Q}_1$, and $a_1 \notin P$ for all $P \in \mathcal{Q}_0$. Hence, there are no two directed paths in $\mathcal{Q}_0 \cup \mathcal{Q}_1$ that correspond to the same undirected path. Furthermore, the set of undirected paths corresponding to $\mathcal{Q}_0 \cup \mathcal{Q}_1$ contains no pair of conflicting undirected paths. It was shown by Ben-Ameur and Gourdin [6] that in this case there exists a metric $\lambda' \in \mathbb{R}^E_+$ in the undirected cycle G that is compatible with the undirected path set corresponding to $\mathcal{Q}_0 \cup \mathcal{Q}_1$. Clearly, the corresponding directed metric $\lambda'' \in \mathbb{R}^A_+$ with $\lambda''_{(u,v)} := \lambda'_{uv}$ for all $(u, v) \in A$ is compatible with the directed path set $\mathcal{Q}_0 \cup \mathcal{Q}_1$ in D.

Now, we modify this metric such that it is compatible with the entire path set $\mathcal{Q}$. Let $M := 1 + \sum_{a \in A \setminus \{a_0\}} \lambda''_a$. Since no path in $\mathcal{Q}_0 \cup \mathcal{Q}_1$ contains the arc a_0, the metric $\lambda \in \mathbb{R}^A_+$ defined as

$$\lambda_a := \begin{cases} M, & \text{if } a = a_0, \text{ and} \\ \lambda''_a, & \text{otherwise,} \end{cases}$$

is compatible with $\mathcal{Q}_0 \cup \mathcal{Q}_1$. Furthermore, any path in $\mathcal{Q}_2$ is shorter than its clockwise counterpart with respect to λ. Thus, all paths $\mathcal{Q}$ are uniquely determined shortest paths w.r.t. λ between their terminals. $\square$

Lemma 6.4 and Lemma 6.5 lead to the constant factor approximation algorithm BidirectedRingRounding: In the first step, we compute an optimal multicommodity flow routing x^* with the additional properties stated in Lemma 6.4. Then we remove the least utilized arc $a_{\min} = \arg\min f_a(x^*)$ from the bidirected ring. Then we 'round' the optimal MCF routing in such a way, that no routing path uses $a_{\min}$ and no pairs of confliction paths are created.

Algorithm 6.2 BidirectedRingRounding

1. Compute an optimal solution (L^*, x^*) of (CON-LP) with $x^*_{P_1} = 0$ or $x^*_{P_2} = 0$ for any pair of conflicting paths $P_1 \in \mathcal{P}(s_1, t_1)$ and $P_2 \in \mathcal{P}(s_2, t_2)$ with $(s_1, t_1) \neq (s_2, t_2)$.
2. Find arc $a_{\min} := \arg\min f_a(x^*)$.
   Let $\bar{a}_{\min}$ be the reverse arc of $a_{\min}$.
3. Define $[x] \in \{0, 1\}^{\mathcal{P}}$ as
$$[x]_P := \begin{cases} 0, & \text{if } a_{\min} \in P \text{ or } x^*_P < 0.5 \\ & \text{or if } x^*_P = 0.5 \text{ and } \bar{a}_{\min} \in P, \text{ and} \\ 1, & \text{otherwise.} \end{cases}$$
4. Compute a compatible metric λ for the path set $\mathcal{Q} := \{P \in \mathcal{P} \ : \ [x]_P = 1\}$.
5. Return λ.

By Lemma 6.5, the routing obtained this way is an USPR of the given commodities. We can compute a compatible metric for this routing in polynomial with the linear programming approaches presented in [6, 9], for example.

Theorem 6.6 *BidirectedRingRounding is a 3-approximation algorithm for* MIN-CON-USPR *on a bidirected ring.*

Proof. It follows immediately from Lemmas 6.4 and 6.5 that algorithm BidirectedRingRounding computes a valid solution for MIN-CON-USPR. It remains to show that this solution has a congestion of at most three times the optimal solution's congestion.

In Step 3 of algorithm BidirectedRingRounding, we shift all flows on paths across arc $a_{\min}$ to the respective opposite flow paths, and we round path variables x_p with $x^*_p \geq 0.5$ to 1. Hence, for any arc $a \in A$, we have $f_a([x]) \leq 2 f_a(x^*) + f_{a_{\min}}(x^*) \leq 3 f_a(x^*)$. □

Algorithm BidirectedRingRounding and Theorem 6.6 straightforward carry over to the case where all strongly connected components of D are bidirected rings (or subgraphs of bidirected rings).

6.4 Cap-USPR on a bidirected ring

In this final section, we show how to approximate CAP-USPR on a bidirected ring within a constant factor. In principle, we use the same approach as for the undirected problem version: We compute one solution (λ^0, z^0) that is a 2-approximation of the optimal solution (λ^{opt}, z^{opt}) if $z^{opt}_a \geq 1$ for all $a \in A$, and $|A|$ many solutions (λ^l, z^l), $l \in A$, to cope with the cases where $z^{opt}_l = 0$ for some $l \in A$.

In contrast to the undirected cycle case, the CAP-USPR problem remains $\mathcal{NP}$-hard on a bidirected ring even with the restriction $z_l = 0$ for some arc $l \in A$. However, it is possible to approximate the restricted problem within a constant factor by rounding the optimal solution of the following linear programming relaxation of CAP-USPR:

$$
\begin{aligned}
\min \sum_{a \in A} w_a z_a & & & \text{(CAP-LP2)}\\
\sum_{P \in \mathcal{P}(s,t)} x_P &= 1 & \forall\, (s,t) \in K &\\
\sum_{(s,t)\in K} \sum_{P \in \mathcal{P}(s,t): a \in P} d_{(s,t)} \cdot x_P &\le c_a z_a & \forall\, a \in A &\\
x_P &\le z_a & \forall\, P \in \mathcal{P},\, a \in P & \quad (6)\\
z_a &\ge 0 & \forall\, a \in A &\\
0 \le x_P &\le 1 & \forall\, (s,t) \in K,\, P \in \mathcal{P}(s,t) &
\end{aligned}
$$

(CAP-LP2) is the directed version of (CAP-LP) strengthened by the inequalities (6). The inequalities (6) are trivially valid for any integer solution of the directed version of (CAP-LP). Hence, any optimal solution of (CAP-LP2) provides a lower bound on the optimal solution value for CAP-USPR. If D is a bidirected ring, these inequalities close a large part of the integrality gap of (CAP-LP) and the strengthened formulation remains polynomially large.

Analogous to the previous section, an optimal solution of (CAP-LP2) can be turned into an USPR by removing one arc and rounding the fractional flows in the residual graph.

Lemma 6.7 *Let $D = (V, A)$ be a bidirected ring. Then there exists an optimal solution (L^*, x^*) of* (CAP-LP2) *such that $x^*_{P_1} + x^*_{P_2} \le 1$ for any pair of conflicting paths $P_1 \in \mathcal{P}(s_1, t_1)$ and $P_2 \in \mathcal{P}(s_2, t_2)$ with $(s_1, t_1) \neq (s_2, t_2)$. Furthermore, such a solution (L^*, x^*) can be found in polynomial time.*

Lemma 6.7 leads straightforward to the 4-approximation algorithm EnumerateAndRound2.

Algorithm 6.3 EnumerateAndRound2

---

1. Compute solution (λ^0, z^0) as follows:
   - 1.1 Set $\lambda^0_a := \begin{cases} Mw_a/c_a + 1, & \text{if } a \in \{(v_0, v_1),\, (v_1, v_0)\}, \\ Mw_a/c_a, & \text{otherwise,} \end{cases}$

     with $M := 2\Pi_{a \in A} c_a$.
   - 1.2 Set $z^0_a := \lceil f_a(\lambda^0)/c_a \rceil$ for all $a \in A$.
2. For each $l \in A$, compute solution (λ^l, z^l) as follows:
   - 2.1 Compute an optimal solution (z^{l*}, x^{l*}) of (CAP-LP2) with the restriction $z^l_l = 0$ such that $x^{l*}_{P_1} + x^{l*}_{P_2} \le 1$ for any pair of conflicting paths $P_1 \in \mathcal{P}(s_1, t_1)$ and $P_2 \in \mathcal{P}(s_2, t_2)$ with $(s_1, t_1) \neq (s_2, t_2)$.
   - 2.2 Let $\bar{l}$ be the reverse arc of l.
   - 2.3 Define $[x^l] \in \{0, 1\}^{\mathcal{P}}$ as

     $$[x]^l_P := \begin{cases} 0, & \text{if } l \in P \text{ or if } x^{l*}_P < 0.5 \text{ or if } x^{l*}_P = 0.5 \text{ and } \bar{l} \in P, \text{ and} \\ 1, & \text{otherwise.} \end{cases}$$
   - 2.4. Compute a compatible metric λ^l for the path set $\mathcal{Q}^l := \{P \in \mathcal{P} \ :\ [x]^l_P = 1\}$.
   - 2.5 Set $z^l_a := \lceil f_a(\lambda^l)/c_a \rceil$ for all $a \in a$.
3. Return minimum cost solution of (λ^0, z^0) and (λ^l, z^l), $l \in A$.

Theorem 6.8 *EnumerateAndRound2 is a* 4*-approximation algorithm for* Cap-USPR *on a bidirected ring.*

Proof. Analogous to the proof of Theorem 6.6, it follows from Lemma 6.5 and Lemma 6.7 that the path sets $\mathcal{Q}^l$ are valid USPRs, i.e., contain a uniquely determined shortest (s,t)-path for each $(s,t) \in K$ w.r.t. the metrics λ^l computed in Step 2.4. The metric λ^0 defined in Step 1.2 is a perturbation of the metric w_a/c_a. As $\lambda^0_{(v_0,v_1)}$ and $\lambda^0_{(v_1,v_0)}$ are odd and all other lengths λ^0_a are even, also λ^0 defines unique shortest paths between all node pairs. Hence, all solutions (λ^0, z^0) and (λ^l, z^l) are feasible.

Let (λ^{opt}, z^{opt}) be the optimal solution of the given Cap-USPR instance. Analogous to the undirected case, it follows that (λ^0, z^0) is a 2-approximate solution if $z^{opt}_a \geq 1$ for all $a \in A$. So, assume that $z^{opt}_l = 0$ for some $l \in A$. Due to inequalities (6), we then have $z^{l*}_a \geq 0.5$ for each arc $a \in A$ with $\lceil z^{l*}_a \rceil \geq 1$. Hence, we have

$$w(z^l) = \sum_{a \in A} w_a z^l_a = \sum_{a \in A} w_a \lceil z^{l*}_a \rceil \leq \sum_{a \in A} w_a \, 4 z^{l*}_a = 4w(z^{l*}) \leq 4w(z^{opt}) \ .$$

Consequently, algorithm EnumerateAndRound2 returns a 4-approximate solution in the worst case. □

Algorithm EnumerateAndRound2 and Theorem 6.8 generalize straightforward to instances where all strongly connected components of the underlying digraph are bidirected rings.

7 Concluding remarks

In this paper, we have shown that it is $\mathcal{NP}$-hard to approximate Min-Con-USPR within a factor of $\mathcal{O}(|V|^{1-\epsilon})$ in general and Cap-USPR within a factor of $\mathcal{O}(2^{\log^{1-\epsilon}|V|})$ in the directed or $2-\epsilon$ in the undirected case. The fixed charge network design problem FC-USPR was proven to be $\mathcal{NPO}$-complete. We presented simple $|A|$- and $|K|$-approximation algorithms for Min-Con-USPR in general networks and we illustrated how known techniques can be used to approximate several special cases of Cap-USPR. For the special cases where the underlying graph is an undirected cycle or a bidirected ring, constant factor approximation algorithms for Min-Con-USPR and Cap-USPR were proposed.

We also constructed examples where the minimum congestion obtainable with unsplittable shortest path routing is a factor of $\Omega(|V|^2)$ larger than the congestion of an optimal unsplittable flow routing or an optimal shortest multi-path routing, and a factor of $\Omega(|V|)$ larger than the congestion of an optimal unsplittable source-invariant routing.

It remains open whether the inapproximability threshold of $\theta(|V|^{1-\epsilon})$ for Min-Con-USPR is tight or whether approximations better than $\min\{|A|, |K|\}$ can be achieved. It is also not known how to compute approximate solutions with reasonable quality guarantees for the general Cap-USPR problem. The methods known for the corresponding Generalized Steiner network or Buy-at-Bulk Network Design problem versions do not necessarily yield feasible solutions for the Cap-USPR problem.

References

[1] M. Andrews, *Hardness of buy-at-bulk network design*, Proceedings of the 45nd Annual IEEE Symposium on Foundations of Computer Science, 2004.

[2] G. Ausiello, P. Crescenzi, G. Gambosi, V. Kann, A. Marchetti-Spaccamela, and M. Protasi, *Complexity and approximation: Combinatorial optimization problems and their approximability properties*, Springer-Verlag, 1999.

[3] B. Awerbuch and Y. Azar, *Buy-at-bulk network design*, Proceedings of the 38th Annual IEEE Symposium on Foundations of Computer Science, Miami Beach FL, 1997, pp. 542–547.

[4] Y. Bartal, *On approximating arbitrary metrics by tree metrics*, Proceedings of the Thirtieth Annual ACM Symposium on the Theory of Computing, Dallas TX, 1998.

[5] W. Ben-Ameur, *Multi-hour design of survivable classical IP networks*, International Journal of Communication Systems **15** (2002), 553–572.

[6] W. Ben-Ameur and E. Gourdin, *Internet routing and related topology issues*, SIAM Journal on Discrete Mathematics **17** (2003), no. 1, 18–49.

[7] W. Ben-Ameur, E. Gourdin, B. Liau, and N. Michel, *Optimizing administrative weights for efficient single-path routing*, Proceedings of Networks 2000, 2000.

[8] A. Bley, *A Lagrangian approach for integrated network design and routing in IP networks*, Proceedings of the First International Network Optimization Conference (INOC 2003), Paris, October 2003, pp. 107–113.

[9] ———, *Inapproximability results for the inverse shortest paths problem with integer lengths and unique shortest paths*, ZIB Preprint ZR-05-04, Konrad-Zuse-Zentrum für Informationstechnik Berlin, 2005.

[10] A. Bley, M. Grötschel, and R. Wessäly, *Design of broadband virtual private networks: Model and heuristics for the B-WiN*, Robust Communication Networks: Interconnection and Survivability (N. Dean, D.F. Hsu, and R. Ravi, eds.), DIMACS Series in Discrete Mathematics and Theoretical Computer Science, vol. 53, American Mathematical Society, 1998, pp. 1–16.

[11] A. Bley and T. Koch, *Integer programming approaches to access and backbone IP-network planning*, ZIB Preprint ZR-02-41, Konrad-Zuse-Zentrum für Informationstechnik Berlin, 2002.

[12] P. Broström and K. Holmberg, *Determining the non-existence of a compatible OSPF metric*, Technical Report LiTH-MAT-R-2004-06, Linköping University, April 2004.

[13] ———, *Stronger necessary conditions for the existence of a compatible OSPF metric*, Technical Report LiTH-MAT-R-2004-08, Linköping University, May 2004.

[14] L. Buriol, M. Resende, C. Ribeiro, and M. Thorup, *A hybrid genetic algorithm for the weight setting problem in OSPF/IS-IS routing*, Networks **46** (2005), 36–56.

[15] M. Charikar, C. Chekuri, A. Goel, S. Guha, and S.A. Plotkin, *Approximating a finite metric by a small number of tree metrics*, Proceedings of the 39th Annual IEEE Symposium on Foundations of Computer Science, Palo Alto CA, 1998, pp. 379–388.

[16] M. Charikar and A. Karagiozova, *On non-uniform multicommodity buy-at-bulk network design*, Proceedings of the Thirty-seventh Annual ACM Symposium on the Theory of Computing, Baltimore, Maryland, USA, 2005.

[17] S. Cosares and I. Saniee, *An optimization problem related to balancing loads on SONET rings*, Telecommunication Systems **3** (1994), 165–182.

[18] P. Crescenzi and V. Kann (eds.), *A compendium of NP optimization problems*, 2003.

[19] Y. Dinitz, N. Garg, and M. Goemans, *On the single source unsplittable flow problem*, Combinatorica **19** (1999), 1–25.

[20] Y. Dodis and S. Khanna, *Designing networks with bounded pairwise distance*, Proceedings of the Thirty-first Annual ACM Symposium on the Theory of Computing, Atlanta GA, 1999, pp. 750–759.

[21] P. Erdös and H. Sachs, *Reguäre Graphen gegebener Taillenweite mit minimaler Knotenzahl*, Wissenschaftliche Zeitung der Universität Halle-Wittenberg **12** (1963), 251–257.

[22] M. Ericsson, M. Resende, and P. Pardalos, *A genetic algorithm for the weight setting problem in OSPF routing*, Journal of Combinatorial Optimization **6** (2002), 299–333.

[23] J. Fakcharoenphol, S. Rao, and K. Talwar, *A tight bound on approximating arbitrary metrics by tree metrics*, Proceedings of the Thirty-fifth Annual ACM Symposium on the Theory of Computing, 2003.

[24] A. Farago, A. Szentesi, and B. Szviatovszki, *Allocation of administrative weights in PNNI*, Proceedings of Networks '98, Sorrento, Italy, 1998, pp. 621–625.

[25] ______, *Inverse optimization in high-speed networks*, Discrete Applied Mathematics **129** (2003), 83–98.

[26] B. Fortz and M. Thorup, *Internet traffic engineering by optimizing OSPF weights*, Proceedings of the 19th IEEE Infocom 2000, Tel-Aviv, Israel, 2000, pp. 519–528.

[27] ______, *Increasing Internet capacity using local search*, Computational Optimization and Applications **29** (2004), no. 1, 13–48.

[28] M. Garey and D. Johnson, *Computers and intractability: A guide to the theory of NP-completeness*, Freeman and Company, New York, 1979.

[29] M.X. Goemans and D.P. Williamson, *A general approximation technique for constrained forest problems*, SIAM Journal on Computing **24** (1995), 296–317.

[30] S. Guha, A. Meyerson, and K. Munagala, *A constant factor approximation for the single sink edge installation problem*, Proceedings of the Thirty-third Annual ACM Symposium on the Theory of Computing, Hersonissos, Crete, Greece, 2001, pp. 383–388.

[31] A. Gupta, A. Kumar, and T. Roughgarden, *Simpler and better approximation algorithms for network design*, Proceedings of the Thirty-fifth Annual ACM Symposium on the Theory of Computing, 2003, pp. 365–372.

[32] K. Holmberg and D. Yuan, *Optimization of Internet protocol network design and routing*, Networks **43** (2004), no. 1, 39–53.

[33] R. Karp, *Reducibility among combinatorial problems*, Complexity of Computer Computations (R. Miller and J. Thatcher, eds.), Plenum Press, New York, 1972, pp. 85–103.

[34] S. Kolliopoulos and C. Stein, *Improved approximation algorithms for unsplittable flow problems*, Proceedings of the 38th Annual IEEE Symposium on Foundations of Computer Science, Miami Beach FL, 1997, pp. 426–435.

[35] F. Lin and J. Wang, *Minimax open shortest path first routing algorithms in networks suporting the SMDS service*, Proceedings of the IEEE International Conference on Communications 1993 (ICC'93), Geneva, Suisse, vol. 2, 1993, pp. 666–670.

[36] D. Lorenz, A. Orda, D. Raz, and Y. Shavitt, *How good can IP routing be?*, Technical Report 2001-17, DIMACS, Center for Discrete Mathematics and Theoretical Computer Science, Rutgers University, Princeton University, AT&T Bell Laboratries and Bellcore, 2001.

[37] Y. Mansur and D. Peleg, *An approximation algorithm for minimum-cost network design*, Tech. Report CS94-22, Weizman Institute of Science, Rehovot, Israel, 1994.

[38] A. Meyerson and K. Munagala S. Plotkin, *Cost-distance: Two-metric network design*, Proceedings of the Thirty-second Annual ACM Symposium on the Theory of Computing, Portland OR, 2000, pp. 624–630.

[39] P. Orponen and H. Mannila, *On approximation preserving reductions: Complete problems and robust measures*, Technical Report C-1987-28, Department of Computer Science, University of Helsinki, 1987.

[40] M. Pióro, A. Szentesi, J. Harmatos, and A. Jüttner, *On OSPF related network optimization problems*, 8th IFIP Workshop on Performance Modelling and Evaluation of ATM & IP Networks (Ilkley, UK), 2000, pp. 70/1–70/14.

[41] M. Prytz, *On optimization in design of telecommunications networks with multicast and unicast traffic*, Ph.D. thesis, Royal Institute of Technology, Stockholm, Sweden, 2002.

[42] *Routing information protocol*, IETF Internet RFC 1058, June 1988.

[43] *Use of OSI IS-IS for routing in TCP/IP and dual environments*, IETF Internet RFC 1195, December 1990.

[44] *OSPF version 2*, IETF Internet RFC 2328, April 1998.

[45] A. Schrijver, P. Seymour, and P. Winkler, *The ring loading problem*, SIAM Journal on Applied Mathematics **11** (1998), no. 1, 1–14.

[46] M. Skutella, *Approximating the single source unsplittable min-cost flow problem*, Mathematical Programming **91** (2002), no. 3, 493–514.

[47] K. Talwar, *Single-sink buy-at-bulk LP has constant integrality gap*, Proceedings of 9th Conference on Integer Programming and Combinatorial Optimization (IPCO 2002), 2002, pp. 475–486.