A comparison of routing sets for robust network design

Michael Poss *

Abstract

Designing a network able to route a set of non-simultaneous demand vectors is an important problem arising in telecommunications. The problem can be seen a two-stage robust program where the recourse function consists in choosing the routing for each demand vector. Allowing the routing to change arbitrarily as the demand varies yields a very difficult optimization problem so that different subsets of admissible routings have been discussed in the literature. In this paper, we compare theoretically the optimal capacity allocation costs for six of these routing sets: affine routing, volume routing and its two simplifications, the routing based on an arbitrary 2-cover of the uncertainty set, and the routing based on a cover delimited by a hyperplane. We show that the two routing sets based on covers of the uncertainty set yield the same optimal costs. We show then that the two simplified volume routings are special cases of affine routings. Finally, assuming that the uncertainty set is the one studied by Bertsimas and Sim (2004), we show that the optimal cost provided by volume routing is not less than the costs provided by the simplified volume routings.

Keywords: Robust optimization; Network design; Routing set; Routing template; Affine routing.

1 Introduction

Given a graph and a set of point-to-point commodities with known demand values, the deterministic network design problem aims at installing enough capacity on the arcs of the graph so that the resulting network is able to route all commodities. In practice it is however very difficult to know with precision the exact values of the demands at the time the design decisions are taken. In the best case, we can estimate a set that contains most likely values for the demand. The introduction of the uncertainty set leads to a robust optimization problem. In this context, a solution is said to be feasible for the problem if it is feasible for all demand vectors that belong to the estimated uncertainty set $\mathcal{D}$, see Soyster [28] and Ben-Tal and Nemirovski [9, 10], among others. This rigid framework is computationally easy but it does not allow the model to react against the uncertainty. To address this drawback, Ben-Tal et al. [8] introduce two-stage robust optimization models that allows to adjust a subset of the problem variables only after observing the actual realization of the data. This adjusting procedure is often called recourse. This two-stage approach applies naturally to network design since first stage capacity design decisions are usually made in the long term while the routing decisions depend on the realization of the demand. Hence, the routing decisions can be seen as the recourse. Free recourse is called dynamic routing in the context of robust network design problems. It has been shown by Chekuri et al. [14] and Gupta et al. [18] that the robust network design with dynamic routing is intractable. Already deciding whether or not a fixed capacity design allows for a dynamic routing of demands in a given polytope is co-$\mathcal{NP}$-complete (on directed graphs).

It is known already that two-stage robust programming with arbitrary recourse is computationally intractable [8]. For this reason, Ben-Tal et al. [8] limit the recourse to affine functions of the uncertainties which makes the problem tractable. Further works by Chen and Zhang [15] and Goh and Sim [17] suggest to extend the second stage to piece-wise linear functions of the uncertainties. In fact, considering special types of recourses had been used already in the context of network design. Ben-Ameur and Kerivin [4, 5] introduce the concept of static routing: after fixing the design, the routing of a commodity is allowed to change but only linearly with the variation of the commodity. Static routing can also be seen as a single stage robust program where the set of routings paths together with the percental splitting among the paths are chosen at the same time the design decisions

*Department of Computer Science, Faculté des Sciences, Université Libre de Bruxelles, Brussels, Belgium, mposs@ulb.ac.be

are made. The resulting set of paths and percental splitting is often called a routing template, which is used by all demand vectors in the uncertainty set. The use of static routing makes the robust network design problem tractable but it yields more expensive capacity allocations than the problem with dynamic routing. Static routing has been used by various authors since its introduction by Ben-Ameur and Kerivin, including Altin et al. [1], Koster et al. [19], Ordóñez and Zhao [22].

Several authors tried to introduce routing schemes that are more flexible than static routing while still being computationally easier than dynamic routing. Ben-Ameur [3] covers the demand uncertainty set by two (or more) subsets using separating hyperplanes and uses specific routings templates for each subset. The resulting optimization problem is $\mathcal{NP}$-hard when no assumptions is made on the hyperplanes. Scutellà [27] generalizes this idea to arbitrary covers of the uncertainty set. She allows a set of routing templates to be used conjointly so that each demand vector can be routed by at least one of the routing templates. She also introduces a procedure that works in two steps. First, an optimal capacity allocation with static routing is computed. Then, she allows allows to reroute part of the demand vectors according to a second routing template. Ben-Ameur and Zotkiewicz [6] introduce volume routing, a framework that shares the demand between two routing templates, according to thresholds. They prove that the resulting optimization problem is $\mathcal{NP}$-hard and introduce two simplifications. Finally, applying the affine recourse from Ben-Tal et al. [8] to robust network design problems, Ouorou and Vial [24] introduce the concept of affine routing. Recently, Poss and Raack [26, 25] study the properties of affine routing, and compare the later to the static and dynamic routings, both theoretically and empirically. They conclude that affine routing tends to yield very good approximations of dynamic routing while being computationally tractable.

In this paper, we compare theoretically the optimal capacity allocation costs provided by the affine routings from Ouorou and Vial [24], the volume routings from Ben-Ameur and Zotkiewicz [6], and the routings based on covers of the uncertainty set in two subsets (Ben-Ameur [3] and Scutellà [27]). In the next section, we introduce the robust network design problem and define a routing set. We model the robust network design problem with the explicit dependency on the routing set and formalize each of the routing frameworks studied herein. Our main results are stated in Section 2.3. In Section 3, we try to understand how good is the cost of the optimal capacity allocation provided by each of the routing sets, and we compare these costs among the different routing sets. We start by comparing the costs obtained with routings based on covers of the uncertainty set in Section 3.2. Then, in Section 3.3, we turn to affine and volume routings. In Section 4, we present examples showing that it is not possible, in general, to compare some of these costs. Finally, we conclude the paper in Section 5.

2 Robust network design

2.1 Problem formulation

The problem is defined below for a directed graph $G = (V, A)$ and a set of commodities K. We formalize first the concept of a routing. Then, we introduce the robust network design problem. Each commodity $k \in K$ has a source $s(k) \in V$, a destination $t(k) \in V$, and a demand value $d^k \geq 0$. A multi-commodity flow is a vector $f \in \mathbb{R}_+^{|A|\times|K|}$ that satisfies the flow conservation constraints at each node of the network:

$$\sum_{a\in\delta^+(v)} f_a^k - \sum_{a\in\delta^-(v)} f_a^k = \begin{cases} d^k & \text{if } v = s(k) \\ -d^k & \text{if } v = t(k) \\ 0 & \text{else} \end{cases} \quad \text{for each } v \in V, \tag{1}$$

where $\delta^+(v)$ and $\delta^-(v)$ respectively denote the set of outgoing arcs and incoming arcs at node v.

In this work the values of the demand vector are uncertain and belong to the closed, convex, and bounded set $\mathcal{D} \subset \mathbb{R}_+^{|K|}$. We call such a set an uncertainty set and any $d \in \mathcal{D}$ is called a realization of the demand. We denote by $(\mathcal{D}, \mathbb{R}^{|A|\times|K|})$ the set of all functions from $\mathcal{D}$ to $\mathbb{R}^{|A|\times|K|}$. Then, a routing is a function $f \in (\mathcal{D}, \mathbb{R}^{|A|\times|K|})$ that satisfies (1) for all realizations of the demand, that is

$$\sum_{a\in\delta^+(v)} f_a^k(d) - \sum_{a\in\delta^-(v)} f_a^k(d) = \begin{cases} d^k & \text{if } v = s(k) \\ -d^k & \text{if } v = t(k) \\ 0 & \text{else} \end{cases} \quad \text{for all } v \in V, d \in \mathcal{D}, \tag{2}$$

and that is non-negative

$$f_a^k(d) \geq 0 \qquad \text{for all } d \in \mathcal{D}. \tag{3}$$

A routing with no further restrictions is called dynamic routing. Hence, the set of all dynamic routings is the set of all functions from $\mathcal{D}$ to $\mathbb{R}^{|A|\times|K|}$ that satisfy (2) and (3):

$$\mathcal{F} \equiv \left\{ f \in (\mathcal{D}, \mathbb{R}^{|A|\times|K|}) \mid f \text{ satisfies (2) and (3)} \right\}. \tag{4}$$

In this paper, we are interested in using special kinds of routings. This corresponds to using specific subsets $\mathcal{F}' \subseteq \mathcal{F}$. These subsets are described in the next section. In what follows, we describe the robust network design problem making explicit the set of admissible routings.

A vector $x \in \mathbb{R}_+^{|A|}$ is called a capacity allocation. A capacity allocation is said to support the set $\mathcal{D}$ if there exists a dynamic routing $f \in \mathcal{F}$ serving $\mathcal{D}$ such that for every $d \in \mathcal{D}$ the corresponding multi-commodity flow $f(d)$ does not exceed the capacities described by x. Similarly, we say that (x, f) supports $\mathcal{D}$ when both the routing f and the capacity allocation x are given. More generally, we say that $(x, \mathcal{F}')$ supports $\mathcal{D}$ when there exists a routing $f \in \mathcal{F}'$ such that (x, f) supports $\mathcal{D}$. Given an uncertainty set $\mathcal{D}$ and a routing set $\mathcal{F}' \subseteq \mathcal{F}$, robust network design now aims at providing the cost minimal capacity allocation x such that $(x, \mathcal{F}')$ supports $\mathcal{D}$:

$$\begin{array}{llll} & \min & \displaystyle\sum_{a\in A} \kappa_a x_a & \\ RND(\mathcal{F}') & \text{s.t.} & f \in \mathcal{F}' & \quad (5) \\ & & \displaystyle\sum_{k\in K} f^k(d) \le x, \qquad d \in \mathcal{D} & \quad (6) \\ & & x \ge 0, & \end{array}$$

where $\kappa_a \in \mathbb{R}$ is the cost for installing one unit of capacity on arc $a \in A$. Notice that in real applications, these costs are usually non-negative. We shall denote the optimal cost of $RND(\mathcal{F}')$ by $opt(\mathcal{F}')$. Problem $RND(\mathcal{F}')$ contains an infinite number of variables $f(d)$ for all $d \in \mathcal{D}$ as well as an infinite number of capacity constraints (6). Moreover, the problem may not even be linear, depending on the constraints defining set $\mathcal{F}'$.

Considering the set of all routings $\mathcal{F}$, $RND(\mathcal{F})$ is a two-stage robust program with recourse following the more general framework described by Ben-Tal et al. [8]. The capacity design has to be fixed in the first stage, and observing a demand realization $d \in \mathcal{D}$, we are allowed to adjust the routing $f(d)$ arbitrarily in the second stage. In that case, (5) is replaced by (2) and (3) so that $RND(\mathcal{F})$ is a linear program, yet infinite. Whenever $\mathcal{D}$ is a polytope, Poss and Raack [26], among others, show how to provide a finite linear programming formulation for $RND(\mathcal{F})$. The formulation is based on enumerating the extreme points of $\mathcal{D}$, so that its size tends to increase exponentially with the number of commodities. In fact, the problem is very difficult to solve given that only deciding whether a given capacity allocation vector x supports $\mathcal{D}$ is co$\mathcal{NP}$-complete for general polytopes $\mathcal{D}$, see Chekuri et al. [14] and Gupta et al. [18]. Moreover, the use of dynamic routings suffers from another drawback. It may be difficult in practice to change arbitrarily the routing according to the demand realization.

For these reasons, various authors study restrictions on the routings that can be used, introducing different subsets of routings $\mathcal{F}' \subset \mathcal{F}$. Their hope is that $opt(\mathcal{F}')$ provides a good approximation of $opt(\mathcal{F})$ while yielding an easier optimization problem $RND(\mathcal{F}')$. For instance, Frangioni et al. [16] and Poss and Raack [26] show under very strong assumptions on $\mathcal{D}$ that the optimal capacity allocations provided by dynamic routings are equivalent to the ones provided by static routings and affine routings, respectively, which are polynomially solvable when $\mathcal{D}$ has a compact formulation. In the next section, we present different choices of $\mathcal{F}'$ discussed in the literature, including static and affine routings. Then, we summarize the main contributions of this paper in Section 2.3.

Note that if there exists only one path from $s(k)$ to $t(k)$ for a commodity $k \in K$, then all routings coincide for that commodity. Unless stated otherwise, in the following we assume that for all $k \in K$ there exist at least two distinct paths p_1, p_2 in G from $s(k)$ to $t(k)$, that is, two paths that differ by one arc at least.

2.2 Routings frameworks

In the next sections, we define formally the set of static routings and the routing sets from Ouorou and Vial [24], Ben-Ameur [3], Scutellà [27] and Ben-Ameur and Zotkiewicz [6].

2.2.1 Static routing

The simplest alternative to dynamic routing has been introduced by Ben-ameur [5] and has been used extensively since then, see Altin et al. [1, 2], Koster et al. [19], Mudchanatongsuk et al. [21], and Ordóñez and Zhao [22]. This framework considers a restriction on the second stage recourse known as *static routing* (also called oblivious routing). Each component $f^k : \mathcal{D} \to \mathbb{R}_+^{|A|}$ is forced to be a linear function of d^k:

$$f_a^k(d) := y_a^k d^k \qquad a \in A, k \in K, d \in \mathcal{D}. \tag{7}$$

Notice that (7) implies that the flow for k is not changing if we perturb the demand for $h \neq k$. By combining (2) and (7) it follows that the multipliers $y \in \mathbb{R}_+^{|A|\times|K|}$ satisfy to

$$\sum_{a\in\delta^+(v)} y_a^k - \sum_{a\in\delta^-(v)} y_a^k = \begin{cases} 1 & \text{if } v = s(k) \\ -1 & \text{if } v = t(k) \\ 0 & \text{else} \end{cases} \quad \text{for each } v \in V. \tag{8}$$

The flow y is called a routing template since it decides, for every commodity, which paths are used to route the demand and what is the percental splitting among these paths. We define formally the the set of all routing templates as

$$\mathcal{Y} \equiv \left\{ y \in \mathbb{R}_+^{|A|\times|K|} \mid y \text{ satisfies } (8) \right\}, \tag{9}$$

and the set of all static routings as

$$\mathcal{F}^{\mathtt{stat}} \equiv \left\{ f \in (\mathcal{D}, \mathbb{R}^{|A|\times|K|}) \mid \exists y \in \mathcal{Y} : f_a^k(d) = y_a^k d^k \quad a \in A, k \in K, d \in \mathcal{D} \right\}.$$

An important result is that a compact linear formulation can be provided for $RND(\mathcal{F}^{\mathtt{stat}})$ as long as the description of $\mathcal{D}$ is compact (see Altin et al. [2] among others). Hence, the resulting optimization problem is polynomially solvable.

In the following, we review alternative routing sets $\mathcal{F}'$ that are less restrictive than static routings while not being as flexible as dynamic routings. Said differently, $\mathcal{F}^{\mathtt{stat}} \subseteq \mathcal{F}' \subseteq \mathcal{F}$.

2.2.2 Covers of the uncertainty set delimited by a hyperplane

Given a set $\mathcal{D}$, a collection of subsets of $\mathcal{D}$ forms a cover of $\mathcal{D}$ if $\mathcal{D}$ is a subset of the union of sets in the collection. Ben-Ameur [3] introduces the idea of covering the uncertainty set by two (or more) subsets using hyperplanes and proposes to use a routing template for each subset. This yields the following set of routings:

$$\mathcal{F}^{2|} \equiv \Big\{ f \in (\mathcal{D}, \mathbb{R}^{|A|\times|K|}) \quad | \quad \exists y^1, y^2 \in \mathcal{Y} \text{ and } \alpha \in \mathbb{R}^K, \beta \in \mathbb{R} : \\ f_a^k(d) \; = \; \begin{cases} y_a^{1k} d^k & d \in \mathcal{D} \cap \{d, \alpha d \le \beta\} \\ y_a^{2k} d^k & d \in \mathcal{D} \cap \{d, \alpha d \ge \beta\} \end{cases} \quad a \in A, k \in K, d \in \mathcal{D} \Big\}.$$

The definition above implies that both routing templates y^1 and y^2 must be able to route demand vectors that lie in the hyperplane $\{d, \alpha d = \beta\}$ without exceeding the capacity. He proves that $RND(\mathcal{F}^{2|})$ is $\mathcal{NP}$-hard in general and describes simplification schemes, where α is given. He further works on the framework in Ben-Ameur and Zotkiewicz [7].

2.2.3 Arbitrary covers of the uncertainty set

Scutellà [27] introduces the idea of using conjointly two routing templates. Formally, she proposes to use two routing templates y^1 and y^2 such that each $d \in \mathcal{D}$ can be served either by y^1 or by y^2 (or both). This yields the following set of routings:

$$\mathcal{F}^{2} \equiv \Big\{ f \in (\mathcal{D}, \mathbb{R}^{|A|\times|K|}) \quad | \quad \exists y^1, y^2 \in \mathcal{Y} \text{ and } \mathcal{D}^1, \mathcal{D}^2 \subseteq \mathcal{D}, \mathcal{D} = \mathcal{D}^1 \cup \mathcal{D}^2 : \\ f_a^k(d) \; = \; \begin{cases} y_a^{1k} d^k & d \in \mathcal{D}^1 \\ y_a^{2k} d^k & d \in \mathcal{D}^2 \end{cases} \quad a \in A, k \in K, d \in \mathcal{D} \Big\}.$$

She mentions that the complexity of $RND(\mathcal{F}^2)$ is unknown. We show in this paper that this optimization problem is $\mathcal{NP}$-hard, because it is a generalization of $RND(\mathcal{F}^{2|})$, proved to be $\mathcal{NP}$-hard by Ben-Ameur [3]. The framework described by $\mathcal{F}^2$ has been independently proposed for general robust programs by Bertsimas and Caramanis [11] (see also Bertsimas et al. [12]) where the authors propose to cover the uncertainty sets with k subsets and devise independent sets of recourse variables for each of these subsets.

2.2.4 Volume routings

More recently, Ben-Ameur and Zotkiewicz [6] introduce a framework that shares the demand between two routing templates, according to thresholds h^k for each $k \in K$. Formally, they use the following set of routings:

$$\begin{aligned}\mathcal{F}^{\mathtt{V}} \equiv \Big\{ f \in (\mathcal{D}, \mathbb{R}^{|A|\times|K|}) \quad | \quad & \exists y^1, y^2 \in \mathcal{Y}, h \in \mathbb{R}^K_+ : \\ f^k_a(d) \quad = \quad & y^{1k}_a \min(d^k, h^k) + y^{2k}_a \max(d^k - h^k, 0) \quad a \in A, k \in K, d \in \mathcal{D} \Big\}.\end{aligned}$$

They prove that $RND(\mathcal{F}^{\mathtt{V}})$ is an $\mathcal{NP}$-hard optimization problem. Hence, they introduce simpler frameworks described below. Defining $d^k_{min} = \min_{d\in\mathcal{D}} d^k$ and $d^k_{max} = \max_{d\in\mathcal{D}} d^k$, the set of routings becomes one of the following

$$\begin{aligned}\mathcal{F}^{\mathtt{VS}} &\equiv \Big\{ f \in (\mathcal{D}, \mathbb{R}^{|A|\times|K|}) \mid \exists y^1, y^2 \in \mathcal{Y} : f^k_a(d) = y^{1k}_a d^k_{min} + y^{2k}_a (d^k - d^k_{min}) \quad a \in A, k \in K, d \in \mathcal{D} \Big\}, \\ \mathcal{F}^{\mathtt{VG}} &\equiv \Big\{ f \in (\mathcal{D}, \mathbb{R}^{|A|\times|K|}) \mid \exists y^1, y^2 \in \mathcal{Y} : \\ & f^k_a(d) = y^{1k}_a d^k_{min} \frac{d^k_{max} - d^k}{d^k_{max} - d^k_{min}} + y^{2k}_a d^k_{max} \frac{d^k - d^k_{min}}{d^k_{max} - d^k_{min}} \; a \in A, k \in K, d \in \mathcal{D} \Big\},\end{aligned}$$

which are both well-defined whenever $d^k_{min} < d^k_{max}$ for each $k \in K$. When $d^k_{min} = d^k_{max}$ for some $k \in K$, the k-th component of $f \in \mathcal{F}^{\mathtt{VG}}$ is defined by $f^k(d) = y^{1k} d^k$.

2.2.5 Affine routings

Ben-Tal et al. [8] introduce Affine Adjustable Robust Counterparts restricting the recourse to be an affine function of the uncertainties. Ourou and Vial [24] apply this framework to robust network design by restricting f^k to be an affine function of all components of d giving

$$\begin{aligned}\mathcal{F}^{\mathtt{aff}} \quad \equiv \quad & \{ f \in (\mathcal{D}, \mathbb{R}^{|A|\times|K|}) | \exists f^0 \in \mathbb{R}^K, y \in \mathbb{R}^{|A|\times|K|} : \\ & f^k_a(d) = f^{0k}_a + \sum_{h\in K} y^{kh}_a d^h \quad a \in A, k \in K, d \in \mathcal{D}, f \text{ satisfies (2) and (3)} \Big\}.\end{aligned}$$

This framework has been compared theoretically and numerically to static and dynamic routings by Poss and Raack [26]. In particular, the authors show that a compact formulation can be described for $RND(\mathcal{F}^{\mathtt{aff}})$ as long as $\mathcal{D}$ has a compact description, generalizing the result obtained for static routing already. We point out that a major difference between $\mathcal{F}^{\mathtt{aff}}$ and the routing described in Section 2.2.1-2.2.4 is that the formers are build up using routing templates, so that it is implicitly assumed that flow conservation constraints (2) and non-negativity constraints (3) are satisfied. In opposition, routings in $\mathcal{F}^{\mathtt{aff}}$ are build up using ordinary vectors so that that satisfaction of (2) and (3) must be stated explicitly.

2.3 Contributions of this paper

The objective of this paper is to compare $opt(\mathcal{F}')$ among the routing sets recalled in previous sections. This comparison is carried out in Section 3. Our main results are stated next.

(a) Let $\mathcal{D}$ be an uncertainty set. It holds that $opt(\mathcal{F}^2) = opt(\mathcal{F}^{2|})$, $opt(\mathcal{F}^{\mathtt{aff}}) \leq opt(\mathcal{F}^{\mathtt{VG}}) \leq opt(\mathcal{F}^{\mathtt{VS}})$, and $opt(\mathcal{F}^{\mathtt{V}}) \leq opt(\mathcal{F}^{\mathtt{VS}})$ for any cost vector $\kappa \in \mathbb{R}^{|A|}$.

(b) Let $\mathcal{D}$ be an uncertainty polytope such that for each $k \in K$, there exists non-negative numbers $0 \leq d^k_{min} \leq d^k_{max}$ such that $d^k \in \{d^k_{min}, d^k_{max}\}$ for each extreme point of $\mathcal{D}$. It holds that $opt(\mathcal{F}^{\mathtt{V}}) = opt(\mathcal{F}^{\mathtt{VS}})$ for any cost vector $\kappa \in \mathbb{R}^{|A|}$.

The polytope introduced by Bertsimas and Sim [13], used for robust network design problems in [6, 23, 24, 20, 26], satisfies the assumption of (b) when the number of deviations allowed is integer.

We present examples in Section 4 showing that it is not possible, in general, to order $opt(\mathcal{F}^2)$, $opt(\mathcal{F}^{\mathtt{V}})$ and $opt(\mathcal{F}^{\mathtt{aff}})$.

3 Optimal costs

The objective of this section is to compare the cost of the optimal capacity allocations obtained for $RND(\mathcal{F}')$ using different routing sets $\mathcal{F}'$. We prove in Section 3.2 that it always holds that $opt(\mathcal{F}^{2|}) = opt(\mathcal{F}^2)$. In Section 3.3, we prove that it always holds that $opt(\mathcal{F}^{\mathtt{aff}}) \le opt(\mathcal{F}^{\mathtt{VG}}) \le opt(\mathcal{F}^{\mathtt{VS}})$. We show also that under additional assumptions on $\mathcal{D}$, it holds that $opt(\mathcal{F}^{\mathtt{VS}}) = opt(\mathcal{F}^{\mathtt{V}})$, $opt(\mathcal{F}^{\mathtt{VG}}) \le opt(\mathcal{F}^{\mathtt{V}})$ and $opt(\mathcal{F}^{\mathtt{aff}}) \le opt(\mathcal{F}^{\mathtt{V}})$. In Section 3.1, we describe the methodology used herein to obtain the desired relations.

3.1 Methodology

Given two routing sets $\mathcal{F}$ and $\mathcal{F}'$, we prove that $opt(\mathcal{F}') \le opt(\mathcal{F}^*)$ using two different approaches. The first approach consists in comparing directly the routing sets themselves, by showing that $\mathcal{F}^* \subseteq \mathcal{F}'$. Proving this inclusion is a very strong result, which holds only for closely related routing sets. In such a situation, we say that $\mathcal{F}^*$ is a special case of $\mathcal{F}'$. Because it is not always possible to compare directly the routing sets themselves, the second approach is based on comparing the sets of all capacity allocations that support $\mathcal{D}$ when considering a specific routing set. These sets are defined formally as

$$\mathcal{X}(\mathcal{F}') \equiv \{x \in \mathbb{R}_+^{|A|} \mid (x, \mathcal{F}') \text{ supports } \mathcal{D}\}, \tag{10}$$

for any routing set $\mathcal{F}'$. To better understand the link between $\mathcal{X}(\mathcal{F}')$ and $opt(\mathcal{F}')$, $RND(\mathcal{F}')$ can be equivalently written as

$$\min \left\{ \sum_{a \in A} \kappa_a x_a \text{ s.t. } x \in \mathcal{X}(\mathcal{F}') \right\}. \tag{11}$$

The second approach is weaker than the first one in the sense that $\mathcal{F}^* \subseteq \mathcal{F}'$ implies that $\mathcal{X}(\mathcal{F}^*) \subseteq \mathcal{X}(\mathcal{F}')$. Hence, it can be applied to more pairs of routing sets.

We prove next a property satisfied by (10). We say that a set $\mathcal{F}' \subset (\mathcal{D}, \mathbb{R}^{|A| \times |K|})$ is convex if the line segment between any two elements of $\mathcal{F}'$ lies in $\mathcal{F}'$, that is, if for any $f, \overline{f} \in \mathcal{F}'$ and any $0 \le \lambda \le 1$ we have that $\lambda f + (1-\lambda)\overline{f} \in \mathcal{F}'$.

Lemma 1. *If $\mathcal{F}'$ is a convex subset of $(\mathcal{D}, \mathbb{R}^{|A| \times |K|})$ then $\mathcal{X}(\mathcal{F}')$ is a convex subset of $\mathbb{R}_+^{|A|}$.*

Proof. Consider $x, \overline{x} \in \mathcal{X}(\mathcal{F}')$. Hence, there exists $f, \overline{f} \in \mathcal{F}'$ such that both (x, f) and $(\overline{x}, \overline{f})$ support $\mathcal{D}$. Because (6) is constituted of linear equations, we see that $(\lambda x + (1-\lambda)\overline{x}, \lambda f + (1-\lambda)\overline{f})$ supports $\mathcal{D}$ for all $0 \le \lambda \le 1$. Therefore, $\lambda f + (1-\lambda)\overline{f} \in \mathcal{F}'$ implies that $\lambda x + (1-\lambda)\overline{x} \in \mathcal{X}(\mathcal{F}')$. □

The two approaches are formalized in the result below.

Proposition 1. *Let $\mathcal{F}'$ and $\mathcal{F}^*$ be two routing sets. The following holds:*

1. *If $\mathcal{F}^* \subseteq \mathcal{F}'$, then $opt(\mathcal{F}') \le opt(\mathcal{F}^*)$ for any cost vector $\kappa \in \mathbb{R}^{|A|}$.*
2. *If $\mathcal{X}(\mathcal{F}^*) \subseteq \mathcal{X}(\mathcal{F}')$ then $opt(\mathcal{F}') \le opt(\mathcal{F}^*)$ for any cost vector $\kappa \in \mathbb{R}^{|A|}$.*
3. *If $opt(\mathcal{F}') \le opt(\mathcal{F}^*)$ for any cost vector $\kappa \in \mathbb{R}^{|A|}$ and $\mathcal{F}'$ is a convex subset of $(\mathcal{D}, \mathbb{R}^{|A| \times |K|})$ then $\mathcal{X}(\mathcal{F}^*) \subseteq \mathcal{X}(\mathcal{F}')$.*

Proof. 1: Follows immediately from the definition of $RND(\mathcal{F}')$.
2: Follows from the fact that $opt(\mathcal{F}')$ is the cost of the optimal solution of (11).
3: Suppose there exists $\overline{x} \in \mathcal{X}(\mathcal{F}^*) \backslash \mathcal{X}(\mathcal{F}')$. Applying Lemma 1, $\mathcal{X}(\mathcal{F}')$ is a convex set so that there exists a hyperplane $H \subset \mathbb{R}^{|A|}$ such that $H \cap \mathcal{X}(\mathcal{F}')$ is empty and $\overline{x} \in H$. Therefore, let κ be the vector in $\mathbb{R}^{|A|}$ orthogonal to H and pointing towards the half-space containing $\mathcal{X}(\mathcal{F}')$. By definition of κ, $\sum_{a \in A} \kappa_a \overline{x}_a < \sum_{a \in A} \kappa_a x_a$ for all $x \in \mathcal{X}(\mathcal{F}')$. Hence, $opt(\mathcal{F}^*) < opt(\mathcal{F}')$. □

In the following sections, we will use Proposition 1 to relate the optimal capacity allocation costs among the routing sets introduced in Section 2.2.

3.2 Routings that cover $\mathcal{D}$

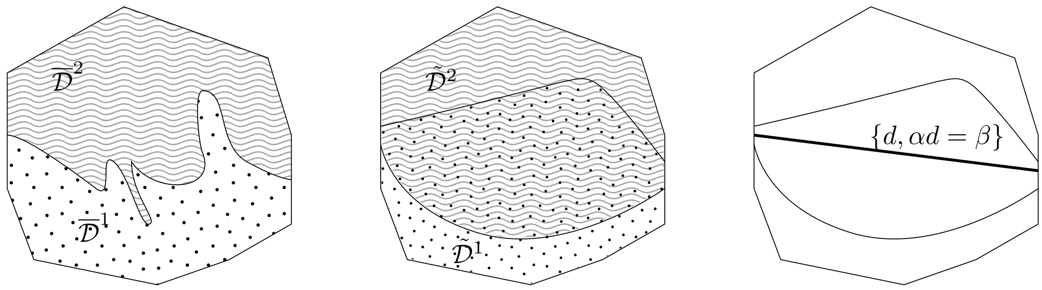

(a) Cover associated with routing $\overline{f} \in \mathcal{F}^2$. (b) Considering sets where $\overline{y}^1$ or $\overline{y}^2$ satisfies (6). (c) The resulting hyperplane.

Figure 1: Construction of a separating hyperplane.

In this section we focus on $\mathcal{F}^2$ and $\mathcal{F}^{2l}$.

Theorem 1. *Let $\mathcal{D}$ be an uncertainty set. It holds that $\mathcal{X}(\mathcal{F}^{2l}) = \mathcal{X}(\mathcal{F}^2)$.*

To prove Theorem 1, we need the following Lemma.

Lemma 2. *Let $\mathcal{D}$ be a convex set in $\mathbb{R}_+^{|K|}$. Let $\mathcal{D}^1$ and $\mathcal{D}^2$ be two closed and convex subsets of $\mathcal{D}$ such that $\mathcal{D}^2 \nsubseteq \mathcal{D}^1$, $\mathcal{D}^1 \nsubseteq \mathcal{D}^2$, and $\mathcal{D} = \mathcal{D}^1 \cup \mathcal{D}^2$. Then there exists a hyperplane $H \subset \mathbb{R}^{|K|}$ such that $H \cap \mathcal{D} \subseteq \mathcal{D}^1 \cap \mathcal{D}^2$.*

Proof. Consider the sets conv$(\mathcal{D}^1 \backslash \mathcal{D}^2)$ and conv$(\mathcal{D}^2 \backslash \mathcal{D}^1)$. There exists a hyperplane $H \subset \mathbb{R}^{|K|}$ such that H separates conv$(\mathcal{D}^1 \backslash \mathcal{D}^2)$ and conv$(\mathcal{D}^2 \backslash \mathcal{D}^1)$. Since $\mathcal{D}^2 \nsubseteq \mathcal{D}^1$ and $\mathcal{D}^1 \nsubseteq \mathcal{D}^2$, $H \cap \mathcal{D}$ is non empty and belongs to cl$(\mathcal{D}^1 \cap \mathcal{D}^2) = \mathcal{D}^1 \cap \mathcal{D}^2$. □

Proof. of Theorem 1. $\subseteq$: Follows from the fact that $\mathcal{F}^{2l}$ is the subset of $\mathcal{F}^2$ where the intersection of $\mathcal{D}^1$ and $\mathcal{D}^2$ must be a hyperplane.
$\supseteq$: Consider a capacity allocation x and a routing $\overline{f} \in \mathcal{F}^2$ such that $(x, \overline{f})$ supports $\mathcal{D}$. Routing $\overline{f}$ is defined by the cover $\mathcal{D} = \overline{\mathcal{D}}^1 \cup \overline{\mathcal{D}}^2$, see Figure 1(a), and the routing templates $\overline{y}^1$ and $\overline{y}^2$. We shall prove that $\overline{f}$ can always be transformed to a routing $\hat{f} \in \mathcal{F}^{2l}$ such that $(x, \hat{f})$ supports $\mathcal{D}$.

Let $\tilde{\mathcal{D}}^1$ (resp. $\tilde{\mathcal{D}}^2$) be defined as the subset of $\mathcal{D}$ where $\overline{y}^1$ (resp. $\overline{y}^2$) satisfies (6) for capacity x. Since we assumed that $(x, \overline{f})$ supports $\mathcal{D}$, $(x, \overline{f})$ satisfies (6) so that $\tilde{\mathcal{D}}^1 \supseteq \overline{\mathcal{D}^1}$ and $\tilde{\mathcal{D}}^2 \supseteq \overline{\mathcal{D}^2}$, see Figure 1(b). Hence, $\mathcal{D} = \tilde{\mathcal{D}}^1 \cup \tilde{\mathcal{D}}^2$. In addition, $\tilde{\mathcal{D}}^1$ and $\tilde{\mathcal{D}}^2$ are convex. To see that $\tilde{\mathcal{D}}^1$ is convex, consider $d_1, d_2 \in \tilde{\mathcal{D}}^1$, so that they satisfy the inequalities $\overline{y}_a^{1k} d_1 \le x_a$ and $\overline{y}_a^{1k} d_2 \le x_a$ for each $a \in A$ and $k \in K$. Then, for any $0 \le \lambda \le 1$,

$$\overline{y}^{1k}(\lambda d_1 + (1-\lambda) d_2) = \lambda d_1 \overline{y}^{1k} + (1-\lambda) d_2 \overline{y}^{1k} \le \lambda x + (1-\lambda) x = x,$$

so that $\lambda d_1 + (1-\lambda) d_2 \in \tilde{\mathcal{D}}^1$. The proof is the same for $\tilde{\mathcal{D}}^2$. Then, since the inequalities in (6) are not strict, we see easily that $\tilde{\mathcal{D}}^1$ and $\tilde{\mathcal{D}}^1$ are closed.

Suppose that $\tilde{\mathcal{D}}^1 \subseteq \tilde{\mathcal{D}}^2$. Then, we can define a routing $\hat{f} \in \mathcal{F}^{2l}$ by considering a hyperplane that does not intersect $\mathcal{D}$ so that the only routing template used is $\overline{y}^2$, which proves the result. We can proceed similarly if $\tilde{\mathcal{D}}^2 \subseteq \tilde{\mathcal{D}}^1$. Hence, we can assume that $\tilde{\mathcal{D}}^2 \nsubseteq \tilde{\mathcal{D}}^1$ and $\tilde{\mathcal{D}}^1 \nsubseteq \tilde{\mathcal{D}}^2$.

Therefore, $\mathcal{D}$, $\tilde{\mathcal{D}}^1$ and $\tilde{\mathcal{D}}^2$ satisfy all the hypothesis of Lemma 2. Hence, there exists a hyperplane $\{d, \alpha d = \beta\}$ such that $\{d, \alpha d = \beta\} \cap \mathcal{D} \subset \tilde{\mathcal{D}}^1 \cap \tilde{\mathcal{D}}^2$, see Figure 1(c). Assume w.l.o.g. that $\mathcal{D} \cap \{d, \alpha d \le \beta\} \subseteq \tilde{\mathcal{D}}^1$ and $\mathcal{D} \cap \{d, \alpha d \ge \beta\} \subseteq \tilde{\mathcal{D}}^2$. Hence, we can construct $\hat{f} \in \mathcal{F}^{2l}$ through the hyperplane $\{d, \alpha d = \beta\}$, that is, $\hat{f}(d) := \overline{y}^{1k} d^k$ for each $d \in \mathcal{D} \cap \{d, \alpha d \le \beta\}$ and $\hat{f}^k(d) := \overline{y}^{2k} d^k$ for each $d \in \mathcal{D} \cap \{d, \alpha d \ge \beta\}$. □

It follows from Proposition 1 that the costs of optimal capacity allocations are always equal.

Corollary 1. *Let $\mathcal{D}$ be an uncertainty set. It holds that $opt(\mathcal{F}^{2l}) = opt(\mathcal{F}^2)$ for any cost vector $\kappa \in \mathbb{R}^{|A|}$.*

Theorem 1 does not imply $\mathcal{F}^2 = \mathcal{F}^{2|}$, since we may face routings in $\mathcal{F}^2$ with domains $\mathcal{D}^1$ and $\mathcal{D}^2$ that do not result from a hyperplane separation of $\mathcal{D}$, see for instance Figure 1(a). We remark that Scutellá [27] mentions that the complexity of $RND(\mathcal{F}^2)$ is unknown. The complexity of $RND(\mathcal{F}^2)$ follows directly from the sufficiency condition of Theorem 1 and the fact that Ben-Ameur [3] proves $RND(\mathcal{F}^{2|})$ to be $\mathcal{NP}$-hard.

Corollary 2. *The optimization problem $RND(\mathcal{F}^2)$ is $\mathcal{NP}$-hard.*

3.3 Volume and affine routings

In this section we compare volume and affine routings. Ben-Ameur and Zotkiewicz [6] mention that $\mathcal{F}^{\mathtt{VS}}$ is a special case of $\mathcal{F}^{\mathtt{V}}$, that is, $\mathcal{F}^{\mathtt{VS}} \subseteq \mathcal{F}^{\mathtt{V}}$. The inclusion is easily verified for $\mathcal{F}^{\mathtt{VS}}$, by choosing $h^k = d^k_{min}$.

Lemma 3. *[6] It holds that $\mathcal{F}^{\mathtt{VS}} \subseteq \mathcal{F}^{\mathtt{V}}$.*

However, we explain in what follows that it is is not true that $\mathcal{F}^{\mathtt{VG}} \subseteq \mathcal{F}^{\mathtt{V}}$. The routings in $\mathcal{F}^{\mathtt{V}}$ can only increase the amount of flow sent on any arc of G for commodity k when d^k rises. Said differently, the flow for any $f \in \mathcal{F}^{\mathtt{V}}$, defined by

$$f^k_a(d) = y^{1k}_a \min(d^k, h^k) + y^{2k}_a \max(d^k - h^k, 0) \qquad a \in A, k \in K,$$

is a non-decreasing function in d. In opposition, routings in $\mathcal{F}^{\mathtt{VG}}$ can also decrease the flow sent on some of the arcs since any routing $f \in \mathcal{F}^{\mathtt{V}}$ is defined by

$$f^k_a(d) = y^{1k}_a d^k_{min} \frac{d^k_{max} - d^k}{d^k_{max} - d^k_{min}} + y^{2k}_a d^k_{max} \frac{d^k - d^k_{min}}{d^k_{max} - d^k_{min}} \qquad a \in A, k \in K,$$

which is a sum of a decreasing term and an increasing term. The advantage of decreasing the flow sent on some arcs when the demand for a commodity rises allows to better combine different commodities within the available capacity. We provide in Section 4.2 an example showing that, in general, it holds that $\mathcal{F}^{\mathtt{VG}} \not\subseteq \mathcal{F}^{\mathtt{V}}$.

Routing sets $\mathcal{F}^{\mathtt{VS}}$ or $\mathcal{F}^{\mathtt{VG}}$ are nevertheless special cases of affine routings. Because any routing in $\mathcal{F}^{\mathtt{VS}}$ or $\mathcal{F}^{\mathtt{VG}}$ is described by an affine functions from $\mathcal{D}$ to $\mathbb{R}^{|A|\times|K|}_+$, it must also belong to $\mathcal{F}^{\mathtt{aff}}$. The theorem below proves, moreover, that $\mathcal{F}^{\mathtt{VS}}$ is a special case of $\mathcal{F}^{\mathtt{VG}}$.

Theorem 2. *Let $\mathcal{D}$ be an uncertainty set. The following holds:*

1. *$\mathcal{F}^{VS} \subseteq \mathcal{F}^{VG}$. The inclusion is strict if and only if $0 < d^k_{min} < d^k_{max}$ for at least one $k \in K$.*
2. *$\mathcal{F}^{VG} \subseteq \mathcal{F}^{aff}$. The inclusion is strict if and only if $\dim(\mathcal{D}) > 1$ or $\dim(\mathcal{D}) = 1$ and $\mathcal{D}$ is orthogonal to one of the coordinate axes of $\mathbb{R}^{|K|}_+$.*

Proof. 1. Let the routing templates $\overline{y}^1$ and $\overline{y}^2$ describe any routing $\overline{f} \in \mathcal{F}^{\mathtt{VS}}$. In what follows, we set up routing templates y^1 and y^2 that yield a routing $f \in \mathcal{F}^{\mathtt{VG}}$ equivalent to $\overline{f}$. We consider independently each component $\overline{f}^k$ of $\overline{f}$. If $d^k_{max} = 0$, then $d^k = 0$ for each $d \in \mathcal{D}$ and we can choose y^{1k} and y^{2k} arbitrarily since the resulting f^k and $\overline{f}^k$ are always null. Similarly, if $d^k_{min} = d^k_{max}$, then any routing in $\mathcal{F}^{\mathtt{VS}}$ or $\mathcal{F}^{\mathtt{VG}}$ is uniquely determined by a unique routing template for k. Suppose that $0 < d^k_{min} < d^k_{max}$. Then, we identify $\overline{f}^k$ and f^k at $d^k = d^k_{min}$ and $d^k = d^k_{max}$, obtaining the following templates for $f \in \mathcal{F}^{\mathtt{VG}}$:

$$\begin{aligned} y^{1k}_a &= \overline{y}^{1k}_a \\ y^{2k}_a &= \overline{y}^{1k}_a \frac{d^k_{min}}{d^k_{max}} + \overline{y}^{2k}_a \frac{d^k_{max} - d^k_{min}}{d^k_{max}} \end{aligned}, \text{ for each } a \in A, k \in K. \tag{12}$$

To see that the inclusion is strict, choose a routing $f \in \mathcal{F}^{\mathtt{VG}}$ defined by routing templates y^1 and y^2 such that $y^{2k}_a - y^{1k}_a \frac{d^k_{min}}{d^k_{max}} < 0$ for some $a \in A$. Using again the identification from (12), we see that $\overline{y}^{2k}_a$ would be strictly less than zero so that $\overline{y}^2$ would not be a routing template.
2. Let $\overline{y}^1$ and $\overline{y}^2$ describe any routing $\overline{f} \in \mathcal{F}^{\mathtt{VG}}$. The components of f^0 and y of the corresponding affine routing are obtained by grouping the terms of $\overline{f}$ according to their degree in d. We obtain

that $y^{kh} = 0$ for each $k \neq h \in K$, and

$$\left.\begin{array}{ll} f_a^{0k} & = \dfrac{d_{min}^k d_{max}^k}{d_{max}^k - d_{min}^k}(\overline{y}_a^{1k} - \overline{y}_a^{2k}) \\ y_a^{kk} & = \dfrac{1}{d_{max}^k - d_{min}^k}(d_{max}^k \overline{y}_a^{2k} - d_{min}^k \overline{y}_a^{1k}) \end{array}\right. , \text{ for each } a \in A, k \in K,$$

which proves the inclusion $\mathcal{F}^{\mathtt{VG}} \subseteq \mathcal{F}^{\mathtt{aff}}$.

Suppose that $\dim(\mathcal{D}) = 1$ and that $\mathcal{D}$ is not orthogonal to any of the coordinate axes, and consider any routing $f \in \mathcal{F}^{\mathtt{aff}}$ and a commodity $k \in K$. The flow for k is given by

$$f^k(d) = f_a^{0k} + \sum_{h \in K} y_a^{kh} d^h. \tag{13}$$

Since $\mathcal{D}$ is not orthogonal to the k-th axis, we can parameterize $\mathcal{D}$ through its orthogonal projection on the axis. Namely, there exists positive reals $\lambda^h \in \mathbb{R}_+$ for each $h \in K \backslash k$ such that $d^h = \lambda^h d^k$ for all $d \in \mathcal{D}$. Hence, (13) becomes

$$f_a^{0k} + \left(y_a^{kk} + \sum_{h \in K \backslash k} \lambda^h y_a^{kh} \right) d^k = f_a^{0k} + y_a'^{kk} d^k.$$

Then, identifying f and $\overline{f} \in \mathcal{F}^{\mathtt{VG}}$ at $d^k = d_{min}^k$ and $d^k = d_{max}^k$, any affine routing f is equivalent to the routing $\overline{f} \in \mathcal{F}^{\mathtt{VG}}$ with

$$\left.\begin{array}{ll} \overline{y}_a^{1k} & = f_a^{0k}/d_{min}^k + y_a'^{kk} \\ \overline{y}_a^{2k} & = f_a^{0k}/d_{max}^k + y_a'^{kk} \end{array}\right. , \text{ for each } a \in A, k \in K, \tag{14}$$

which proves the inclusion $\mathcal{F}^{\mathtt{VG}} \supseteq \mathcal{F}^{\mathtt{aff}}$ when $\dim(\mathcal{D}) = 1$ and $\mathcal{D}$ is not orthogonal to any of the coordinate axes. The quantities in (14) are non-negative because f satisfies non-negativity constraints from (3).

Suppose now that $\dim(\mathcal{D}) = 1$ and that $\mathcal{D}$ is orthogonal to the k-th axis. Therefore, $d_{min}^k = d_{max}^k = d^k$ for any $d \in \mathcal{D}$ so that $\overline{f}^k(d) = y^{1k} d^k$ is constant for any $\overline{f} \in \mathcal{F}^{\mathtt{VG}}$. The problem must contain more than one commodity because $\dim(\mathcal{D}) = 1$ and $\mathcal{D}$ is orthogonal to the k-th axis. Thus, there exists $h \in K \backslash k$ such that $\mathcal{D}$ is not orthogonal to the h-th axis. Therefore, we can define an affine routing $f \in \mathcal{F}^{\mathtt{aff}}$ such that f^k is not constant by choosing a proper $y^{kh} \neq 0$. Finally, if $\dim(\mathcal{D}) > 1$, $\mathcal{D}$ contains a small ball B of dimension at least two. Hence, there exists a pair $\{d_1, d_2\} \subset B$ such that $d_1^k = d_2^k$ and $d_1^h \neq d_2^h$ for some $k, h \in K$. As before, all routings in $\mathcal{F}^{\mathtt{VG}}$ for commodity k yield identical flows for d_1 and d_2, while we can define an affine routing $f \in \mathcal{F}^{\mathtt{aff}}$ yielding different flows by choosing a proper $y^{kh} \neq 0$. $\square$

The strict inclusion of Theorem 2.1. has been verified numerically by Ben-Ameur and Zotkiewicz [6], where it is shown that $opt(\mathcal{F}^{\mathtt{VG}})$ can be strictly smaller than $opt(\mathcal{F}^{\mathtt{VS}})$. We show next that $\mathcal{F}^{\mathtt{VS}}$ is always at least as efficient as $\mathcal{F}^{\mathtt{V}}$ whenever $\mathcal{D}$ satisfies the assumption below. Given a convex set $\mathcal{D} \subset \mathbb{R}_+^{|K|}$, we denote by $\mathrm{ext}(\mathcal{D})$ the set of its extreme points.

Assumption 1. *The uncertainty set $\mathcal{D}$ is a polytope such that for each $k \in K$, there exists non-negative numbers $0 \leq d_{min}^k \leq d_{max}^k$ such that $d^k \in \{d_{min}^k, d_{max}^k\}$ for all $d \in \mathrm{ext}(\mathcal{D})$.*

Assumption 1 is satisfied by a well-known family of uncertainty polytopes, see Example 1.

Example 1. *Bertsimas and Sim [13] consider general linear programs where the coefficients of each linear inequality belong to intervals such that the number of coefficients taking conjointly their maximum value is bounded by a constant Γ. Considering upwards deviations only, their uncertainty set can be formalized in $\mathbb{R}_+^{|K|}$ as follows*

$$\mathcal{D}^\Gamma \equiv \left\{ d \in \mathbb{R}_+^{|K|} \mid d^k \in [d_{min}^k, d_{max}^k] \text{ for each } k \in K, \sum_{k \in K} \frac{d^k - d_{min}^k}{d_{max}^k - d_{min}^k} \leq \Gamma \right\}.$$

When Γ is integer, it is easy to see that $\mathcal{D}^\Gamma$ fulfills Assumption 1. Moreover, $\mathcal{D}^\Gamma$ has been frequently used as the uncertainty set for robust network design problems, see [6, 23, 24, 20, 26], among others.

The proof of the theorem below requires the following simple property. For any $x \in \mathbb{R}_+^{|A|}$ and $f \in \mathcal{F}^{\mathtt{VS}}$,

$$(x, f) \text{ supports } \mathcal{D} \qquad \Leftrightarrow \qquad (x, f) \text{ supports } \operatorname{ext}(\mathcal{D}). \tag{15}$$

Property (15) follows directly from the fact that any routing in $\mathcal{F}^{\mathtt{VS}}$ is a linear function.

Theorem 3. *Let $\mathcal{D}$ be an uncertainty set that fulfills Assumption 1. It holds that $\mathcal{X}(\mathcal{F}^{\mathtt{VS}}) = \mathcal{X}(\mathcal{F}^{\mathtt{V}})$.*

Proof. $\subseteq$: Follows directly from the fact that $\mathcal{F}^{\mathtt{VS}} \subseteq \mathcal{F}^{\mathtt{V}}$ by taking $h^k = d^k_{min}$.
$\supseteq$: Consider a capacity allocation x and a routing $f \in \mathcal{F}^{\mathtt{V}}$ such that (x, f) supports $\mathcal{D}$. We show next that there exists a routing $\overline{f} \in \mathcal{F}^{\mathtt{VS}}$ such that $\overline{f}(d) = f(d)$ for each $d \in \operatorname{ext}(\mathcal{D})$. Therefore, $(x, \overline{f})$ supports $\operatorname{ext}(\mathcal{D})$. Since (15) is satisfied, we have that $(x, \overline{f})$ supports $\mathcal{D}$, proving the result.

Notice that the flow for commodity k of any routing in $\mathcal{F}^{\mathtt{V}}$ or $\mathcal{F}^{\mathtt{VS}}$ is a function that only depends on the k-th component of $d \in \mathcal{D}$. Thus, let us define $\mathcal{D}^k$ as the orthogonal projection of $\mathcal{D}$ into its k-th component so that functions f^k and $\overline{f}^k$ are defined on $\mathcal{D}^k$. In the following, we show how to construct $\overline{f} \in \mathcal{F}^{\mathtt{VS}}$ such that $\overline{f}(d) = f(d)$ for each $d \in \operatorname{ext}(\mathcal{D})$ independently for each commodity $k \in K$.

By Assumption 1, $|\operatorname{ext}(\mathcal{D}^k)| \in \{1, 2\}$. First, suppose that $|\operatorname{ext}(\mathcal{D}^k)| = 1$ so that $\mathcal{D}^k = \{d^k\}$ is a singleton. If $d^k = 0$, then we can choose $\overline{y}^{1k}$ and $\overline{y}^{2k}$ arbitrarily since $\overline{f}^k(0)$ will be null anyway. If $d^k > 0$, then we take $\overline{y}^{1k} = f^k(d^k)/d^k$ and $\overline{y}^{2k} = 0$. Suppose now that $|\operatorname{ext}(\mathcal{D}^k)| = 2$. If $d^k_{min} = 0$, then we can choose $\overline{y}^{1k}$ arbitrarily (since it will be multiplied by $d^k_{min} = 0$) and we set $\overline{y}^{2k} = f^k(d^k_{max})/d^k_{max}$. If $d^k_{min} > 0$, then we set $\overline{y}^{1k} = f^k(d^k_{min})/d^k_{min}$ and $\overline{y}^{2k} = f^k(d^k_{max})/(d^k_{max} - d^k_{min})$ for each $k \in K$. Because of Assumption 1, any $d \in \operatorname{ext}(\mathcal{D})$ is such that $d^k \in \{d^k_{min}, d^k_{max}\}$. Therefore, $f^k(d) = \overline{f}^k(d)$ for each $d \in \operatorname{ext}(\mathcal{D})$, so that $(x, \overline{f})$ supports $\operatorname{ext}(\mathcal{D})$. $\square$

Theorem 3 states that whenever $\mathcal{D}$ satisfies Assumption 1, one should not try to use the complex set of routings $\mathcal{F}^{\mathtt{V}}$, since *opt*$(\mathcal{F}^{\mathtt{V}})$ will never beat *opt*$(\mathcal{F}^{\mathtt{VS}})$. This is of particular interest because $RND(\mathcal{F}^{\mathtt{V}})$ is $\mathcal{NP}$-hard in general while Ben-Ameur and Zotkiewicz [6] show that $RND(\mathcal{F}^{\mathtt{VS}})$ is essentially of the same difficulty as $RND(\mathcal{F}^{\mathtt{stat}})$.

4 Non-comparable routings

In this section, we compare *opt*$(\mathcal{F}^2)$, *opt*$(\mathcal{F}^{\mathtt{V}})$ and *opt*$(\mathcal{F}^{\mathtt{aff}})$ for general uncertainty sets. We show that it is not possible to order these costs by presenting three examples where one of the costs is strictly less than the two others. To devise examples showing that $\mathcal{F}^{\mathtt{aff}}$ may yield more expensive capacity allocations than $\mathcal{F}^2$ and $\mathcal{F}^{\mathtt{V}}$, we shall use the following result. Let e^k be the k-th unit vector in $\mathbb{R}_+^{|K|}$.

Proposition 2. *[26, Proposition 8] Let $\mathcal{D}$ be a demand polytope. If $0 \in \mathcal{D}$ and for each $k \in K$ there is $\epsilon_k > 0$ such that $\epsilon_k e^k \in \mathcal{D}$, then opt$(\mathcal{F}^{aff}) =$ opt$(\mathcal{F}^{stat})$.*

Notice that in our examples some of the commodities have unique paths from their sources to their sinks, so that all routings are equal for these commodities. This enables us to produce simple graphs that present the properties required by our examples. One can easily extend these examples to larger graphs for which each commodity $k \in K$ has at least two different paths from its source $s(k)$ to its sink $t(k)$.

4.1 *opt*$(\mathcal{F}^{\mathtt{V}})$ can be strictly smaller than *opt*$(\mathcal{F}^{\mathtt{aff}})$ and *opt*$(\mathcal{F}^2)$

Consider the network design problem for the graph depicted in Figure 2(a) with two commodities $k_1 : a \to c$ and $k_2 : a \to b$. The uncertainty set $\mathcal{D}$ is defined by the extreme points $d_1 = (2, 1)$, $d_2 = (1, 2)$, $d_3 = (1, 0)$, $d_4 = (0, 1)$, and $d_5 = (0, 0)$, and the capacity unitary costs are the edge labels of Figure 2(a). Edge labels from Figure 2(b) and Figure 2(c) represent optimal capacity allocations with dynamic and static routing, respectively. They have costs of 7 and 8, respectively. A routing $f \in \mathcal{F}$ that satisfies the capacity from Figure 2(b) is depicted on Figure 2(d) and Figure 2(e), for d_1 and d_2, respectively. We show next that the optimal capacity allocations for $\mathcal{F}^{\mathtt{V}}$, $\mathcal{F}^2$, and $\mathcal{F}^{\mathtt{aff}}$ are 7, 8 and 8, respectively. The routing f from Figure 2(d) and Figure 2(e) can be extended to a routing in $\overline{f} \in \mathcal{F}^{\mathtt{V}}$ such that $(x, \overline{f})$ supports $\mathcal{D}$ by fixing $\overline{h}^{k_1} = 1$, $\overline{h}^{k_2} = 2$, $\overline{y}^{1k_1} = f^{k_1}(d_2)$, $\overline{y}^{2k_1} = f^{k_1}(d_1) - f^{k_1}(d_2)$, and $\overline{y}^{1k_2} = \overline{y}^{2k_2} = f^{k_2}(d_1)$.

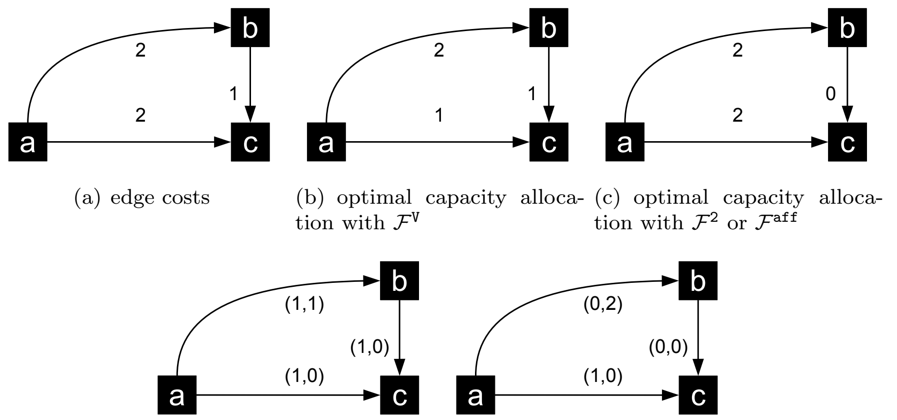

(a) edge costs (b) optimal capacity allocation with $\mathcal{F}^{\mathtt{V}}$ (c) optimal capacity allocation with $\mathcal{F}^2$ or $\mathcal{F}^{\mathtt{aff}}$

(d) $f(d_1)$ (e) $f(d_2)$

Figure 2: Example showing that $opt(\mathcal{F}^{\mathtt{V}})$ can can be strictly smaller than $opt(\mathcal{F}^{\mathtt{aff}})$ and $opt(\mathcal{F}^2)$.

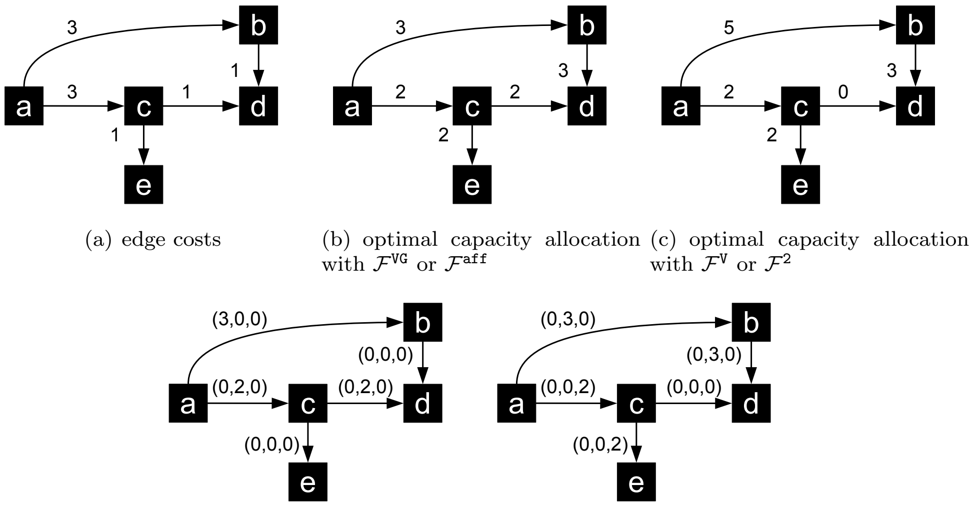

(a) edge costs (b) optimal capacity allocation with $\mathcal{F}^{\mathtt{VG}}$ or $\mathcal{F}^{\mathtt{aff}}$ (c) optimal capacity allocation with $\mathcal{F}^{\mathtt{V}}$ or $\mathcal{F}^2$

(d) $f(d_1)$ (e) $f(d_2)$

Figure 3: Example showing that $opt(\mathcal{F}^{\mathtt{VG}})$ (thus $opt(\mathcal{F}^{\mathtt{aff}})$) can can be strictly smaller than $opt(\mathcal{F}^{\mathtt{V}})$ and $opt(\mathcal{F}^2)$.

However, we explain next why it cannot be extended to a routing in $\mathcal{F}^2$ within the capacity x from Figure 2(b). We restrict our attention to the subset of $\mathcal{D}$ that consists of the line segment $\mathcal{D}' = \text{conv}(d_1, d_2)$ and show that f can already not be extended to a routing in $\mathcal{F}^2$ for $\mathcal{D}'$. Consider flow $f(d_1)$ depicted in Figure 2(d). This flow uses the routing template y^1 defined as $y^{1k} = f^k(d_1^k)/d_1^k$ for $k = k_1, k_2$. Similarly, flow $f(d_2)$ uses the routing template $y^{2k} = f^k(d_2^k)/d_2^k$ for $k = k_1, k_2$. Then, we see that d_1 (resp. d_2) is the unique demand vector in $\mathcal{D}'$ that can be routed within the capacity x from Figure 2(b) using routing template y^1 (resp. y^2). Therefore, defining $\mathcal{D}^1$ (resp. $\mathcal{D}^2$) as the subset of $\mathcal{D}'$ that contains all demand vectors that can be routed along routing template y^1 (resp. y^2), we have that $\mathcal{D}^1 \cup \mathcal{D}^2 \subset \mathcal{D}'$. This shows that it is not possible to extend f to a routing in $\mathcal{F}^2$ for $\mathcal{D}'$, so that it is not possible to do so for $\mathcal{D}$ either.

In fact, we have that the optimal capacity allocation for $\mathcal{F}^2$ is obtained when $\mathcal{D}$ is covered only by itself, yielding $opt(\mathcal{F}^2) = 8$. For $\mathcal{F}^{\mathtt{aff}}$, we can apply Proposition 2 (because $\{(0,0),(1,0),(0,1)\} \subset \mathcal{D}$) so that $opt(\mathcal{F}^{\mathtt{aff}}) = opt(\mathcal{F}^{\mathtt{stat}}) = 8$.

4.2 $opt(\mathcal{F}^{\mathtt{VG}})$ and $opt(\mathcal{F}^{\mathtt{aff}})$ can be strictly smaller than $opt(\mathcal{F}^{\mathtt{V}})$ and $opt(\mathcal{F}^2)$

Consider the network design problem for the graph depicted in Figure 3(a) with three commodities $k_1 : a \to b$, $k_2 : a \to d$ and $k_3 : a \to e$. The uncertainty set $\mathcal{D}$ is defined by the extreme points

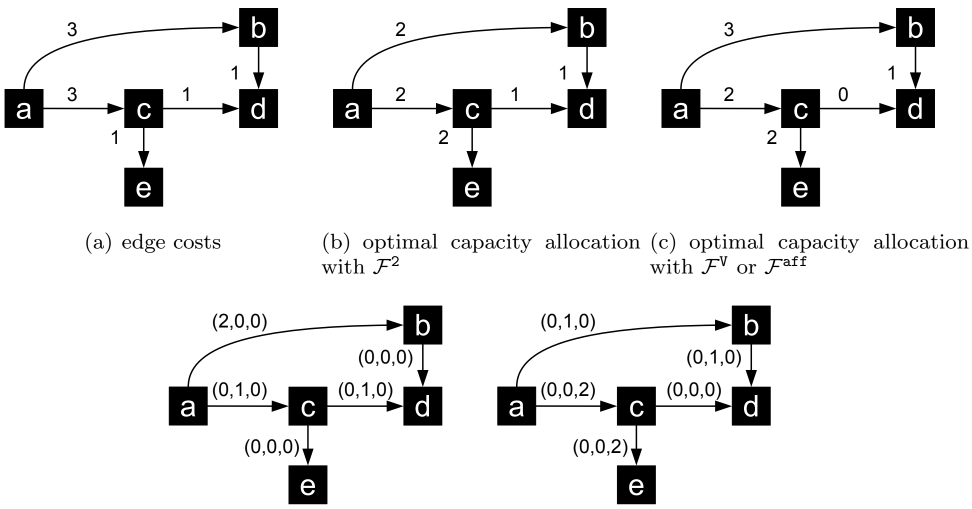

(a) edge costs

(b) optimal capacity allocation with $\mathcal{F}^2$

(c) optimal capacity allocation with $\mathcal{F}^{\mathtt{V}}$ or $\mathcal{F}^{\mathtt{aff}}$

(d) $f(d_1)$

(e) $f(d_2)$

Figure 4: Example showing that $opt(\mathcal{F}^2)$ can can be strictly smaller than $opt(\mathcal{F}^{\mathtt{aff}})$ and $opt(\mathcal{F}^{\mathtt{V}})$.

$d_1 = (3, 2, 0)$ and $d_2 = (0, 3, 2)$, and the capacity unitary costs are the edge labels of Figure 3(a). Edge labels from Figure 3(b) represent an optimal capacity allocation for dynamic routing, with cost 22. A routing $f \in \mathcal{F}$ that satisfies the capacity from Figure 3(b) is depicted on Figure 3(d) and Figure 3(e), for d_1 and d_2, respectively. This routing can be extended to a routing $\overline{f} \in \mathcal{F}^{\mathtt{VG}}$ such that $(x, \overline{f})$ supports $\mathcal{D}$ by setting $\overline{y}^{1k_1} = \overline{y}^{2k_1} = f^{k_1}(d_1)/3$, $\overline{y}^{1k_2} = f^{k_2}(d_1)/2$, $\overline{y}^{2k_2} = f^{k_2}(d_2)/3$ and $\overline{y}^{1k_3} = \overline{y}^{2k_3} = f^{k_3}(d_2)/2$. Applying Theorem 2.2., $\overline{f}$ also belongs to $\mathcal{F}^{\mathtt{aff}}$.

However, f cannot be extended to a routing in $\mathcal{F}^{\mathtt{V}}$ already because $d_2^{k_2} > d_1^{k_2}$ and $f_{ac}^{k_2}(d_2) < f_{ac}^{k_2}(d_1)$, that is, f is not a non-decreasing function. We can show in addition that, using a reasoning similar to the one used in the previous section, f cannot be extended to a routing in $\mathcal{F}^2$ within the existing capacity. We can see that an optimal capacity allocation using $\mathcal{F}^{\mathtt{V}}$ or $\mathcal{F}^2$ is also an optimal capacity allocation using $\mathcal{F}^{\mathtt{stat}}$, and it requires two more units of capacity on ab and no capacity on cd, see Figure 3(c), which yields a total cost of 26.

4.3 $opt(\mathcal{F}^2)$ can be strictly smaller than $opt(\mathcal{F}^{\mathtt{aff}})$ and $opt(\mathcal{F}^{\mathtt{V}})$

Consider the network design problem for the graph depicted in Figure 4(a) with three commodities $k_1 : a \to b$, $k_2 : a \to d$ and $k_3 : a \to e$. The uncertainty set $\mathcal{D}$ is defined by the extreme points $d_1 = (2, 1, 0)$, $d_2 = (0, 1, 2)$, $d_3 = (1, 0, 0)$, $d_4 = (0, 1, 0)$, $d_5 = (0, 0, 1)$, $d_6 = (0, 0, 0)$ and the capacity unitary costs are the edge labels of Figure 4(a) (it is the same as Figure 3(a)). Edge labels from Figure 4(b) and Figure 4(c) represent optimal capacity allocations with dynamic and static routing, respectively. They have costs of 15 and 17, respectively. A routing $f \in \mathcal{F}$ that satisfies the capacity from Figure 4(b) is depicted on Figure 4(d) and Figure 4(e), for d_1 and d_2, respectively. This routing can be extended to a routing $\overline{f} \in \mathcal{F}^2$ such that $(x, \overline{f})$ supports $\mathcal{D}$ by considering the cover through hyperplane $\{d, d^{k_1} = 1\}$ and setting $\overline{y}^{1k_1} = \overline{y}^{2k_1} = f^{k_1}(d_1)/2$, $\overline{y}^{1k_2} = f^{k_2}(d_1)$, $\overline{y}^{2k_2} = f^{k_2}(d_2)$, and $\overline{y}^{1k_3} = \overline{y}^{2k_3} = f^{k_3}(d_2)/2$.

However, f cannot be extended to a routing in $\mathcal{F}^{\mathtt{V}}$ already because $d_2^{k_2} > d_1^{k_2}$ and $f_{ac}^{k_2}(d_2) < f_{ac}^{k_2}(d_1)$, that is, f is not a non-decreasing function. In fact, we have that $opt(\mathcal{F}^{\mathtt{V}}) = opt(\mathcal{F}^{\mathtt{stat}}) = 9$. Then, we can apply Proposition 2 (because $\{(1, 0, 0), (0, 1, 0), (0, 0, 1), (0, 0, 0)\} \subset \mathcal{D}$) to the problem with $\mathcal{F}^{\mathtt{aff}}$, so that $opt(\mathcal{F}^{\mathtt{aff}}) = opt(\mathcal{F}^{\mathtt{stat}}) = 9$.

5 Concluding remarks

This paper studies the optimal capacity allocation cost provided by robust network design models restricted to use specific routing sets. These routing sets are: affine routing, volume routing and its two simplifications, and the routings based on covers of the demand uncertainty set. We show that the routing set based on an arbitrary cover of the uncertainty is equivalent to the routing set that uses a separation hyperplane. We show then that the simplified volume routings are special cases of affine routings. Finally, we show that the general volume routing is no more flexible than its simplifications whenever the uncertainty set is the polytope introduced by Bertsimas and Sim.

An important characteristic of these routing sets is the complexity of the resulting network design problem. In this respect, the general volume routings and the routing sets based on covers of the uncertainty set lead to $\mathcal{NP}$-hard optimization problems. Moreover, while a finite linear programming formulation can be provided for the robust network design problem with dynamic routing under polyhedral uncertainty (by considering only the extreme points of the demand polytope), no such formulations are known for the problems that use the general volume routings or the routings based on covers of the uncertainty set. In this sense, these two routing sets yield optimization problems that are computationally even more difficult than the robust network design with dynamic routing. In opposition, affine routing and the two simplified volume routings lead to polynomially solvable optimization problems, given that the uncertainty polytope has a compact description.

6 Acknowledgements

The author is grateful to Christian Raack for numerous comments that help in improving the presentation of the paper.

References

[1] A. Altin, E. Amaldi, P. Belotti, and M. Ç. Pinar. Provisioning virtual private networks under traffic uncertainty. *Networks*, 49(1):100–115, 2007.

[2] A. Altin, H. Yaman, and M. Ç. Pinar. The robust network loading problem under hose demand uncertainty: Formulation, polyhedral analysis and computations. *INFORMS Journal on Computing*, 23(1):75–89, 2011.

[3] W. Ben-Ameur. Between fully dynamic routing and robust stable routing. In *6th International Workshop on Design and Reliable Communication Networks, 2007. DRCN 2007*, 2007.

[4] W. Ben-Ameur and H. Kerivin. New economical virtual private networks. *Communications of the ACM*, 46(6):69–73, 2003.

[5] W. Ben-Ameur and H. Kerivin. Routing of uncertain demands. *Optimization and Engineering*, 3:283–313, 2005.

[6] W. Ben-Ameur and M. Zotkiewicz. Volume oriented routing. In *14th International Telecommunications Network Strategy and Planning Symposium (NETWORKS)*, pages 1–7, 2010.

[7] W. Ben-Ameur and M. Zotkiewicz. Robust routing and optimal partitioning of a traffic demand polytope. *International Transactions in Operational Research*, 18(3):307–333, 2011.

[8] A. Ben-Tal, A. Goryashko, E. Guslitzer, and A. Nemirovski. Adjustable robust solutions of uncertain linear programs. *Mathematical Programming*, 99(2):351–376, 2004.

[9] A. Ben-Tal and A. Nemirovski. Robust solutions of linear programming problems contaminated with uncertain data. *Mathematical Programming*, 88:411–424, 2000.

[10] A. Ben-Tal and A. Nemirovski. Robust optimization methodology and applications. *Mathematical Programming*, 92:453–480, 2002.

[11] D. Bertsimas and C. Caramanis. Finite adaptability in multistage linear optimization. *Automatic Control, IEEE Transactions on*, 55(12):2751–2766, 2010.

[12] D. Bertsimas, V. Goyal, and A. Sun. Characterization of the power of finite adaptability in multi-stage stochastic and adaptive optimization. To appear in Mathematics of Operations Research., 2011.

[13] D. Bertsimas and M. Sim. The Price of Robustness. *Operations Research*, 52(1):35–53, Jan 2004.

[14] C. Chekuri, G. Oriolo, M. G. Scutellà, and F. B. Shepherd. Hardness of robust network design. *Networks*, 50(1):50–54, 2007.

[15] X. Chen and Y. Zhang. Uncertain Linear Programs: Extended Affinely Adjustable Robust Counterparts. *Operations Research*, 57(6):1469–1482, 2009.

[16] A. Frangioni, F. Pascali, and M. G. Scutellà. Static and dynamic routing under disjoint dominant extreme demands. *Operations Research Letters*, 39(1):36 – 39, 2011.

[17] J. Goh and M. Sim. Distributionally robust optimization and its tractable approximations. *Operations Research*, 58(4-Part-1):902–917, 2010.

[18] A. Gupta, J. Kleinberg, A. Kumar, R. Rastogi, and B. Yener. Provisioning a virtual private network: a network design problem for multicommodity flow. In *STOC '01: Proceedings of the thirty-third annual ACM symposium on Theory of computing*, pages 389–398, 2001.

[19] A. M. C. A. Koster, M. Kutschka, and C. Raack. Towards robust network design using integer linear programming techniques. In *Proceedings of the NGI 2010*, Paris, France, Jun 2010. Next Generation Internet.

[20] A.M.C.A. Koster, M. Kutschka, and C. Raack. Adjustable robust network design: Models, inequalities, and computations, 2010.

[21] S. Mudchanatongsuk, F. Ordonez, and J. Liu. Robust solutions for network design under transportation cost and demand uncertainty. *Journal of the Operations Research Society*, 59:552–562, 2008.

[22] F. Ordóñez and J. Zhao. Robust capacity expansion of network flows. *Networks*, 50(2):136–145, 2007.

[23] A. Ouorou. Affine decision rules for tractable approximations to robust capacity planning in telecommunications. Talk at the International Network Optimization Conference, 2011.

[24] A. Ouorou and J.-P. Vial. A model for robust capacity planning for telecommunications networks under demand uncertainty. In *6th International Workshop on Design and Reliable Communication Networks, 2007. DRCN 2007*, pages 1–4, 2007.

[25] M. Poss and C. Raack. Affine recourse for the robust network design problem: between static and dynamic routing. In *Proceedings of INOC 2011, Hamburg, Germany*, number 6703 in LNCS, pages 150–155. Springer.

[26] M. Poss and C. Raack. Affine recourse for the robust network design problem: between static and dynamic routing. ZIB Report 11-03, Zuse Institute Berlin, February 2011.

[27] M. G. Scutellà. On improving optimal oblivious routing. *Operations Research Letters*, 37(3):197–200, 2009.

[28] A. L. Soyster. Convex programming with set-inclusive constraints and applications to inexact linear programming. *Operations Research*, 21:1154–1157, 1973.