

A comparison of routing sets for robust network design

Michael Poss

Received: date / Accepted: date

Abstract Designing a network able to route a set of non-simultaneous demand vectors is an important problem arising in telecommunications. In this paper, we compare the optimal capacity allocation costs for six routing sets: affine routing, volume routing and its two simplifications, the routing based on an unrestricted 2-cover of the uncertainty set, and the routing based on a cover delimited by a hyperplane.

Keywords Network flows · Robust optimization · Network design · Routing set · Routing template

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. The assumption that the exact demand values are known when planning the network expansion is not realistic. However, using population statistics and traffic measurements, we can estimate a set that contains most of the plausible outcomes for the demand values. 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 the seminal works of [26] and [9]. This rigid framework is computationally easy but it does not allow the model to react against the uncertainty, because optimization variables must take fixed value that are independent of the values taken by the uncertain parameters. This is a well-known drawback of the classical robust optimization framework, and it was the motivation for introducing two-stage robust optimization models where a subset of the optimization variables are fixed after observing the actual realization of the uncertain parameters [8]. 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. Imposing no restriction on the recourse is called dynamic routing in the context of robust network design problems. It has been shown by [19] that the robust network design with dynamic routing is \mathcal{NP} -hard for polyhedral uncertainty.

It is known already that two-stage robust programming with unrestricted recourse is \mathcal{NP} -hard [8]. For this reason, [8] limit the recourse to affine functions of the uncertainties which makes the problem tractable, that is, polynomially solvable. In fact, considering special types of recourse had been used already in the context of network design. Duffield et al. [13] and Fingerhut et al. [14] have independently introduced the concept of static routing for a special case of uncertainty model: after fixing the design, the routing of a commodity is allowed to change but only linearly with the variation of the commodity. Static routing has then been generalized to arbitrary uncertainty polytopes by [4, 5]. 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 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. For instance, [15] have constructed a class of graphs for which static routing can cost, in the worst case, a logarithmic factor more than dynamic routing. Static routing has been used by various authors since its introduction, including [1, 18, 20].

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. Scutellá [24] generalizes this idea to unrestricted 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 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 [8] to robust network design problems, [22] introduce the concept of affine routing. Recently, [23] 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 [22], the volume routings from [6], and the routings based on covers of the uncertainty set in two subsets ([3] and [24]). We introduce in Section 2 the robust network design problem and define formally a routing set. We model the robust network design problem with the explicit de-

pendency on the routing set and formalize each of the routing frameworks studied herein. Then, we study in Section 3 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. In Section 4, we present examples showing that it is not possible, in general, to compare some of these costs. Section 5 provides a numerical comparison of some of these routing set and we conclude the paper in Section 6.

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 $f \in \mathbb{R}_+^{|A| \times |K|}$ with elements f_a^k where for each $k \in K$ and $a \in A$, f_a^k denotes the amount of commodity k going through an arc a . Multi-commodity flow vectors correspond to the vectors of $\mathbb{R}_+^{|A| \times |K|}$ that satisfy the flow conservation constraints at each node of the network:

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

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 vector is a function $f \in (\mathcal{D}, \mathbb{R}^{|A| \times |K|})$ that satisfies (1) for all realizations of the demand, that is

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

and that is non-negative

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

A routing with no further restrictions is called dynamic routing:

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

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 for explicitly defined sets of 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 aims at providing the minimum-cost capacity allocation x and the routing f such that (x, f) supports \mathcal{D} :

$$\begin{aligned}
 \min \quad & \sum_{a \in A} \kappa_a x_a \\
 \text{s.t.} \quad & f \in \mathcal{F}' \\
 & \sum_{k \in K} f^k(d) \leq x, \quad d \in \mathcal{D} \\
 & x \geq 0,
 \end{aligned}
 \tag{5}$$

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 $opt(\mathcal{F}')$. Because of its convexity, set \mathcal{D} is either a singleton or contains an infinite number of vectors. Hence, problem $RND(\mathcal{F}')$ contains in general 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 [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, [2] 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 coNP -complete for general polytopes \mathcal{D} , see [12, 16]. 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 $RND(\mathcal{F}')$. In the next section, we present different choices of \mathcal{F}' discussed in the literature.

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 [22, 3, 24, 6].

2.2.1 Static routing: basic model

The simplest alternative to dynamic routing has been introduced by [13, 14] and has been used extensively since then, see [5, 1, 18, 20]. This framework considers a restriction on the second stage recourse known as *static routing* (also called oblivious routing). Each component $f^k : \mathcal{D} \rightarrow \mathbb{R}_+^{|A|}$ is forced to be a linear function of d^k :

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

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

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

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

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

and the set of all static routings as

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

Recall that a linear formulation is said compact if it involves polynomial numbers of variables and constraints with respect to the problem dimensions (here $|V|$, $|A|$, and $|K|$). An important result is that a compact linear formulation can be provided for $RND(\mathcal{F}^{\text{stat}})$ as long as the description of \mathcal{D} is compact. This result follows directly from the dualization technique introduced by [9] for robust linear optimization. Therefore, 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 i.e., $\mathcal{F}^{\text{stat}} \subseteq \mathcal{F}' \subseteq \mathcal{F}$.

2.2.2 Static routing with hyperplane coverage

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:

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

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 $\{d, \alpha d = \beta\}$. The authors prove that $RND(\mathcal{F}^{21})$ is \mathcal{NP} -hard in general and describes simplification schemes, where α is given. He further works on the framework in [7].

2.2.3 Static routing with unrestricted 2-coverage

Scutellá [24] 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 set of routings

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

She uses a technical proof to show that $RND(\mathcal{F}^2)$ is \mathcal{NP} -hard in [25]. We provide a simpler proof in this paper, based on the fact that $RND(\mathcal{F}^2)$ is a generalization of $RND(\mathcal{F}^{21})$, proved to be \mathcal{NP} -hard by [3]. The framework described by \mathcal{F}^2 has been independently proposed for general robust programs by [10] 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, [6] introduce a framework that shares the demand between two routing templates, according to thresholds h^k for each $k \in K$. Namely, for each $k \in K$, any volume routing sends $\min(d^k, h^k)$ along the first routing template while $\max(d^k - h^k, 0)$ is sent along the second routing template:

$$\mathcal{F}^V \equiv \left\{ f \in (\mathcal{D}, \mathbb{R}^{|A| \times |K|}) \mid \exists y^1, y^2 \in \mathcal{Y}, h \in \mathbb{R}_+^K : \right. \\ \left. f_a^k(d) = y_a^{1k} \min(d^k, h^k) + y_a^{2k} \max(d^k - h^k, 0) \quad a \in A, k \in K, d \in \mathcal{D} \right\}.$$

They prove that $RND(\mathcal{F}^V)$ is an \mathcal{NP} -hard optimization problem. To make the problem more computationally tractable, they introduce simpler frameworks described below. The first simplification (\mathcal{F}^{VS}) supposes that for each commodity, the threshold is equal to the minimal value of the demand over the uncertainty set. Defining $d_{min}^k = \min_{d \in \mathcal{D}} d^k$, any routing in \mathcal{F}^{VS} sends d_{min}^k along the first routing template while $d^k - d_{min}^k$ is sent along the second routing template:

$$\mathcal{F}^{VS} \equiv \left\{ f \in (\mathcal{D}, \mathbb{R}^{|A| \times |K|}) \mid \exists y^1, y^2 \in \mathcal{Y} : \right. \\ \left. f_a^k(d) = y_a^{1k} d_{min}^k + y_a^{2k} (d^k - d_{min}^k) \quad a \in A, k \in K, d \in \mathcal{D} \right\}.$$

Their third routing set (\mathcal{F}^{VG}) also considers pairs of routing templates but without thresholds. Defining $d_{max}^k = \max_{d \in \mathcal{D}} d^k$, any routing in \mathcal{F}^{VG} sends d_{min}^k along the first routing template and d_{max}^k along the second routing template. Then, for any

demand $d^k = \lambda d_{min}^k + (1 - \lambda)d_{max}^k$ with $0 \leq \lambda \leq 1$, the routing sends λd_{min}^k along the first routing template and $(1 - \lambda)d_{max}^k$ along the second routing template:

$$\mathcal{F}^{VG} \equiv \left\{ f \in (\mathcal{D}, \mathbb{R}^{|A| \times |K|}) \mid \exists y^1, y^2 \in \mathcal{Y} : \right. \\ \left. f_a^k(d) = y_a^{1k} d_{min}^k \frac{d_{max}^k - d^k}{d_{max}^k - d_{min}^k} + y_a^{2k} d_{max}^k \frac{d^k - d_{min}^k}{d_{max}^k - d_{min}^k} \quad a \in A, k \in K, d \in \mathcal{D} \right\},$$

which is well-defined whenever $d_{min}^k < d_{max}^k$ for each $k \in K$. When $d_{min}^k = d_{max}^k$ for some $k \in K$, the k -th component of $f \in \mathcal{F}^{VG}$ 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. Ouorou and Vial [22] apply this framework to robust network design by restricting f^k to be an affine function of all components of d giving

$$\mathcal{F}^{aff} \equiv \left\{ f \in (\mathcal{D}, \mathbb{R}^{|A| \times |K|}) \mid \exists f^0 \in \mathbb{R}^K, y \in \mathbb{R}^{|A| \times |K|} : \right. \\ \left. f_a^k(d) = f_a^{0k} + \sum_{h \in K} y_a^{kh} d^h \quad a \in A, k \in K, d \in \mathcal{D}, (2) \text{ and } (3) \right\}.$$

This framework has been compared theoretically and numerically to static and dynamic routings by [23]. In particular, the authors show that a compact formulation can be described for $RND(\mathcal{F}^{aff})$ 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}^{aff} and the routing described in Section 2.2.1-2.2.4 is that the former is built up using routing templates, so that it is implicitly assumed that flow conservation constraints (2) and non-negativity constraints (3) are satisfied. Conversely, routings in \mathcal{F}^{aff} are built up using ordinary vectors so that that satisfaction of (2) and (3) must be stated explicitly.

3 Optimal costs

The objective of this section is to compare the cost of the optimal capacity allocations obtained for $RND(\mathcal{F}')$ by using different routing sets \mathcal{F}' . More specifically, we compare \mathcal{F}^{2l} with \mathcal{F}^2 , \mathcal{F}^{aff} with \mathcal{F}^{VS} and \mathcal{F}^{VG} , and \mathcal{F}^{VS} with \mathcal{F}^V . Our main results are stated next.

- (a) Let \mathcal{D} be an uncertainty set. It holds that $opt(\mathcal{F}^2) \leq opt(\mathcal{F}^{2l})$ and $opt(\mathcal{F}^{aff}) \leq opt(\mathcal{F}^{VG}) \leq opt(\mathcal{F}^{VS})$.
- (b) Let \mathcal{D} be an uncertainty 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 each extreme point of \mathcal{D} . It holds that $opt(\mathcal{F}^V) = opt(\mathcal{F}^{VS})$.

The polytope introduced by [11], used for robust network design problems in [6, 22, 18, 23], 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 $\text{opt}(\mathcal{F}^2)$, $\text{opt}(\mathcal{F}^V)$ and $\text{opt}(\mathcal{F}^{\text{aff}})$.

Given two routing sets \mathcal{F}' and \mathcal{F}^* , we prove that $\text{opt}(\mathcal{F}') \leq \text{opt}(\mathcal{F}^*)$ using two different approaches. The first approach consists in comparing directly the routing sets themselves, by showing that $\mathcal{F}^* \subseteq \mathcal{F}'$. 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

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

for any routing set \mathcal{F}' . To better understand the link between $\mathcal{X}(\mathcal{F}')$ and $\text{opt}(\mathcal{F}')$, $RND(\mathcal{F}')$ can be equivalently written as $\min \{\sum_{a \in A} \kappa_a x_a \text{ s.t. } x \in \mathcal{X}(\mathcal{F}')\}$. The two approaches are formalized in the proposition below stated without proof.

Proposition 1 *Let \mathcal{F}' and \mathcal{F}^* be two routing sets. If $\mathcal{F}^* \subseteq \mathcal{F}'$ or $\mathcal{X}(\mathcal{F}^*) \subseteq \mathcal{X}(\mathcal{F}')$, then $\text{opt}(\mathcal{F}') \leq \text{opt}(\mathcal{F}^*)$.*

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

Proposition 2 *Let \mathcal{D} be an uncertainty set. It holds that $\mathcal{X}(\mathcal{F}^{21}) \subseteq \mathcal{X}(\mathcal{F}^2)$.*

Proof The results directly follows from the fact that \mathcal{F}^{21} is the subset of \mathcal{F}^2 where the intersection of \mathcal{D}^1 and \mathcal{D}^2 must be a hyperplane. \square

The complexity of $RND(\mathcal{F}^2)$ follows directly from the sufficiency condition of Proposition 2 and the fact that [3] proves $RND(\mathcal{F}^{21})$ to be \mathcal{NP} -hard, thus providing a simpler proof than [25].

Corollary 1 *The optimization problem $RND(\mathcal{F}^2)$ is \mathcal{NP} -hard.*

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

Lemma 1 [6] *It holds that $\mathcal{F}^{\text{VS}} \subseteq \mathcal{F}^V$.*

However, it is not true that $\mathcal{F}^{\text{VG}} \subseteq \mathcal{F}^V$ for the following reason. Any routing in \mathcal{F}^V is a non-decreasing functions of d . In opposition, any routing in \mathcal{F}^{VG} is the sum of a non-increasing function of d and a non-decreasing function of d . 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 $\mathcal{F}^{\text{VG}} \not\subseteq \mathcal{F}^V$.

Routing sets \mathcal{F}^{VS} or \mathcal{F}^{VG} are nevertheless special cases of affine routings. Because any routing in \mathcal{F}^{VS} or \mathcal{F}^{VG} is described by an affine functions of d , it must also belong to \mathcal{F}^{aff} . The result below shows, moreover, that \mathcal{F}^{VS} is a special case of \mathcal{F}^{VG} .

Proposition 3 *Let \mathcal{D} be an uncertainty set. The following holds:*

1. $\mathcal{F}^{\text{VS}} \subseteq \mathcal{F}^{\text{VG}}$. The inclusion is strict if and only if there exists $k \in K$ such that $0 < d_{\min}^k < d_{\max}^k$.

2. The inclusion $\mathcal{F}^{VG} \subset \mathcal{F}^{aff}$ holds 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. If $d_{max} = 0$ or $d_{min}^k = d_{max}^k$ for some commodity $k \in K$, then routings in \mathcal{F}^{VS} and \mathcal{F}^{VG} coincide for that commodity. Hence, we suppose in what follows that $0 < d_{min}^k < d_{max}^k$ for each $k \in K$. Let the routing templates \bar{y}^1 and \bar{y}^2 describe any routing $\bar{f} \in \mathcal{F}^{VS}$. Then, routing templates $y_a^{1k} = \bar{y}_a^{1k}$ and $y_a^{2k} = \bar{y}_a^{1k} \frac{d_{min}^k}{d_{max}^k} + \bar{y}_a^{2k} \frac{d_{max}^k - d_{min}^k}{d_{max}^k}$ for each $a \in A$ and $k \in K$ describe a routing $f \in \mathcal{F}^{VG}$ equal to \bar{f} .

Conversely, given a routing $f \in \mathcal{F}^{VG}$ described by routing templates y^1 and y^2 , the equivalent routing $\bar{f} \in \mathcal{F}^{VS}$ is described by routing templates $\bar{y}_a^{1k} = y_a^{1k}$ and $\bar{y}_a^{2k} = y_a^{2k} - y_a^{1k} \frac{d_{min}^k}{d_{max}^k}$ for each $a \in A$ and $k \in K$. Hence, the strict inclusion is verified by choosing a routing $f \in \mathcal{F}^{VG}$ such that $y_a^{2k} - y_a^{1k} \frac{d_{min}^k}{d_{max}^k} < 0$ for some $a \in A$ and $k \in K$.

2. Suppose that $\dim(\mathcal{D}) = 1$ and that \mathcal{D} is not orthogonal to any of the coordinate axes, and consider any routing $f \in \mathcal{F}^{aff}$ and a commodity $k \in K$. We can parameterize \mathcal{D} through its orthogonal projection on the k -th axis. Denoting these projections by $\lambda^h \in \mathbb{R}_+$ for each $h \in K \setminus k$, the flow for k is given by

$$f_a^{0k} + \left(y_a^{kk} + \sum_{h \in K \setminus k} \lambda^h y_a^{kh} \right) d^k = f_a^{0k} + y_a^{kk} d^k. \quad (11)$$

The flow from (11) is equal to the flow prescribed by a routing $\bar{f} \in \mathcal{F}^{VG}$ described by routing templates $\bar{y}^{1k} = f_a^{0k} / d_{min}^k + y_a^{kk}$ and $\bar{y}^{2k} = f_a^{0k} / d_{max}^k + y_a^{kk}$ for each $a \in A$ and $k \in K$.

Suppose that either $\dim(\mathcal{D}) > 1$ or $\dim(\mathcal{D}) = 1$ and \mathcal{D} is orthogonal to one of the coordinate axis. Then, there exists a pair $\{d_1, d_2\} \subset \mathcal{D}$ such that $d_1^k = d_2^k$ and $d_1^h \neq d_2^h$ for some $k, h \in K$. Hence, all routings in \mathcal{F}^{VG} for commodity k yield identical flows for d_1 and d_2 , while we can define an affine routing $f \in \mathcal{F}^{aff}$ yielding different flows by choosing a proper $y^{kh} \neq 0$. \square

Ben-Ameur and Zotkiewicz [6] have provided instances for which the relation $opt(\mathcal{F}^{VG}) < opt(\mathcal{F}^{VS})$ holds. Thus, for these instances, the strict inclusion of Proposition 3.1 holds as well.

We show next that \mathcal{F}^{VS} is always at least as efficient as \mathcal{F}^V whenever \mathcal{D} satisfies the assumption below. Given a convex set $\mathcal{D} \subset \mathbb{R}_+^{|K|}$, we denote by $\text{ext}(\mathcal{D})$ the set of its extreme points.

Assumption 1 *The uncertainty set \mathcal{D} is a polytope such that for each $k \in K$, $d^k \in \{d_{min}^k, d_{max}^k\}$ for all $d \in \text{ext}(\mathcal{D})$ where d_{min}^k and d_{max}^k are defined above (and exist).*

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

Example 1 Bertsimas and Sim [11] 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

$$\mathcal{D}^\Gamma \equiv \left\{ d \in \mathbb{R}_+^{|K|} \mid d^k \in [d_{min}^k, d_{max}^k] \text{ for each } k \in K, \sum_{k \in K} \frac{d^k - d_{min}^k}{d_{max}^k - d_{min}^k} \leq \Gamma \right\}. \quad (12)$$

When Γ is integer, it is easy to see that \mathcal{D}^Γ fulfills Assumption 1. Moreover, \mathcal{D}^Γ has been frequently used as the uncertainty set for robust network design problems, see [6, 22, 18, 23], among others.

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

$$(x, f) \text{ supports } \mathcal{D} \quad \Leftrightarrow \quad (x, f) \text{ supports } \text{ext}(\mathcal{D}). \quad (13)$$

Property (13) follows directly from the fact that any routing in \mathcal{F}^{VS} is a linear function.

Proposition 4 *Let \mathcal{D} be an uncertainty set that fulfills Assumption 1. It holds that $\mathcal{X}(\mathcal{F}^{VS}) = \mathcal{X}(\mathcal{F}^V)$.*

Proof We prove $\mathcal{X}(\mathcal{F}^V) \subseteq \mathcal{X}(\mathcal{F}^{VS})$ since it is already clear that $\mathcal{X}(\mathcal{F}^{VS}) \subseteq \mathcal{X}(\mathcal{F}^V)$ holds. Consider a capacity allocation x and a routing $f \in \mathcal{F}^V$ such that (x, f) supports \mathcal{D} . Because of Property (13), it is enough to show that there exists a routing $\bar{f} \in \mathcal{F}^{VS}$ such that $\bar{f}(d) = f(d)$ only for each $d \in \text{ext}(\mathcal{D})$. If $d_{min}^k = 0$, all routings are equal at that extreme point. Hence, we can suppose that $d_{min}^k > 0$. After simple computations, we see that the flow templates given by $\bar{y}_a^{1k} = f_a^k(d_{min}^k)/d_{min}^k$ and $\bar{y}_a^{2k} = (f_a^k(d_{max}^k) - f_a^k(d_{min}^k))/(d_{max}^k - d_{min}^k)$ for each $a \in A$ and $k \in K$ describe the required routing \bar{f} . \square

Proposition 4 states that whenever \mathcal{D} satisfies Assumption 1, one should not try to use the complex set of routings \mathcal{F}^V , since $\text{opt}(\mathcal{F}^V)$ will never beat $\text{opt}(\mathcal{F}^{VS})$. This is of particular interest because $RND(\mathcal{F}^V)$ is \mathcal{NP} -hard in general while [6] show that $RND(\mathcal{F}^{VS})$ is essentially of the same difficulty as $RND(\mathcal{F}^{stat})$.

4 Non-comparable routings

In this section, we compare $\text{opt}(\mathcal{F}^2)$, $\text{opt}(\mathcal{F}^V)$ and $\text{opt}(\mathcal{F}^{aff})$ 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}^{aff} may yield more expensive capacity allocations than \mathcal{F}^2 and \mathcal{F}^V , we shall use the following result. Let e^k be the k -th unit vector in $\mathbb{R}_+^{|K|}$.

Proposition 5 [23, 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 $\text{opt}(\mathcal{F}^{aff}) = \text{opt}(\mathcal{F}^{stat})$.*

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 $t(k)$.




Fig. 1 Example showing that $\text{opt}(\mathcal{F}^V)$ can be strictly smaller than $\text{opt}(\mathcal{F}^{\text{aff}})$ and $\text{opt}(\mathcal{F}^2)$.

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

Consider the network design problem for the graph depicted in Figure 1(a) with two commodities $k_1 : a \rightarrow c$ and $k_2 : a \rightarrow b$. 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 1(a). Edge labels from Figure 1(b) and Figure 1(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 1(b) is depicted on Figure 1(d) and Figure 1(e), for d_1 and d_2 , respectively. We show next that the optimal capacity allocations for \mathcal{F}^V , \mathcal{F}^2 , and \mathcal{F}^{aff} are 7, 8 and 8, respectively. The routing f from Figure 1(d) and Figure 1(e) can be extended to a routing in $\bar{\mathcal{F}} \in \mathcal{F}^V$ such that (x, \bar{f}) supports \mathcal{D} by fixing $\bar{h}^{k_1} = 1$, $\bar{h}^{k_2} = 2$, $\bar{y}^{1k_1} = f^{k_1}(d_2)$, $\bar{y}^{2k_1} = f^{k_1}(d_1) - f^{k_1}(d_2)$, and $\bar{y}^{1k_2} = \bar{y}^{2k_2} = f^{k_2}(d_1)$.

However, we explain next why it cannot be extended to a routing in \mathcal{F}^2 within the capacity x from Figure 1(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 1(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^2 defined as $y^{2k} = f^k(d_2^k)/d_2^k$ 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 1(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 $\text{opt}(\mathcal{F}^2) = 8$. For \mathcal{F}^{aff} , we can apply Proposition 5 (because $\{(0, 0), (1, 0), (0, 1)\} \subset \mathcal{D}$) so that $\text{opt}(\mathcal{F}^{\text{aff}}) = \text{opt}(\mathcal{F}^{\text{stat}}) = 8$.




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

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

Consider the network design problem for the graph depicted in Figure 2(a) with three commodities $k_1 : a \rightarrow b$, $k_2 : a \rightarrow d$ and $k_3 : a \rightarrow e$. The uncertainty set \mathcal{D} is defined by the extreme points $d_1 = (3, 2, 0)$ and $d_2 = (0, 3, 2)$, and the capacity unitary costs are the edge labels of Figure 2(a). Edge labels from Figure 2(b) represent an optimal capacity allocation for dynamic routing, with cost 22. 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. This routing can be extended to a routing $\bar{f} \in \mathcal{F}^{\text{VG}}$ such that (x, \bar{f}) supports \mathcal{D} by setting $\bar{y}^{1k_1} = \bar{y}^{2k_1} = f^{k_1}(d_1)/3$, $\bar{y}^{1k_2} = f^{k_2}(d_1)/2$, $\bar{y}^{2k_2} = f^{k_2}(d_2)/3$ and $\bar{y}^{1k_3} = \bar{y}^{2k_3} = f^{k_3}(d_2)/2$. Applying Proposition 3.2., \bar{f} also belongs to \mathcal{F}^{aff} .

However, f cannot be extended to a routing in \mathcal{F}^{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}^{V} or \mathcal{F}^2 is also an optimal capacity allocation using $\mathcal{F}^{\text{stat}}$, and it requires two more units of capacity on ab and no capacity on cd , see Figure 2(c), which yields a total cost of 26.

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

Consider the network design problem for the graph depicted in Figure 3(a) with three commodities $k_1 : a \rightarrow b$, $k_2 : a \rightarrow d$ and $k_3 : a \rightarrow e$. 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 3(a) (it is the same as Figure 2(a)). Edge labels from Figure 3(b) and Figure 3(c) represent optimal capacity allocations with dynamic and static routing, respectively. They have costs of 15 and 17, respectively. A




Fig. 3 Example showing that $opt(\mathcal{F}^2)$ can be strictly smaller than $opt(\mathcal{F}^{\text{aff}})$ and $opt(\mathcal{F}^{\text{V}})$.

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 $\bar{f} \in \mathcal{F}^2$ such that (x, \bar{f}) supports \mathcal{D} by considering the cover through hyperplane $\{d, d^{k_1} = 1\}$ and setting $\bar{y}^{1k_1} = \bar{y}^{2k_1} = f^{k_1}(d_1)/2$, $\bar{y}^{1k_2} = f^{k_2}(d_1)$, $\bar{y}^{2k_2} = f^{k_2}(d_2)$, and $\bar{y}^{1k_3} = \bar{y}^{2k_3} = f^{k_3}(d_2)/2$.

However, f cannot be extended to a routing in \mathcal{F}^{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 $opt(\mathcal{F}^{\text{V}}) = opt(\mathcal{F}^{\text{stat}}) = 9$. Then, we can apply Proposition 5 (because $\{(1, 0, 0), (0, 1, 0), (0, 0, 1), (0, 0, 0)\} \subset \mathcal{D}$) to the problem with \mathcal{F}^{aff} , so that $opt(\mathcal{F}^{\text{aff}}) = opt(\mathcal{F}^{\text{stat}}) = 9$.

5 Computational experiments

We report in this section on computational experiments realized for the routing sets that yield polynomially solvable optimization problems: $\mathcal{F}^{\text{stat}}$, \mathcal{F}^{aff} , \mathcal{F}^{VS} , and \mathcal{F}^{VG} . The experiments have been carried out on a computer equipped with a processor Intel Core i7 at 2.90 GHz and 8 GB of RAM memory allowing for 3 threads, and the linear programs were solved by CPLEX 12.4 [17].

The experiments were based on the uncertainty set \mathcal{D}^I defined in (12), for various values of I , and they were realized on three networks from SNDlib [21], *janos-us*, *sun*, and *giul39*. These networks are feasible for a directed formulation and have 26/27/39 nodes and 84/102/172 arcs, respectively. We considered the largest 10 to 50 commodities k with respect to nominal demand value, which allowed us to perform a series of runs for each network. For each $k \in K$, we followed [23] and set d_{min}^k to the nominal demand value and d_{max}^k to $1.4d_{min}^k$.

The results from Table 1 were obtained by solving the compact linear programming formulation obtained when dualizing the robust constraints in $RND(\mathcal{F}')$ for each considered routing set, see for instance [1, 22, 23] for details. Columns ‘‘Cost’’ compare the optimal solution cost of each routing \mathcal{F}' with the cost provided by $\mathcal{F}^{\text{stat}}$. Namely, we report $100(1 - opt(\mathcal{F}')/opt(\mathcal{F}^{\text{stat}}))$ for each considered

\mathcal{F}' . Columns “Time” report solution times in seconds for solving each instance with the barrier solver of CPLEX. We switched off the crossover of CPLEX because our aim is to compute the optimal objective values.

Our experiments show that the two simplified volume routings always yield the same optimal solution costs. This is not surprising because [6] had observed the same situation happening in most of their experiments. We see then that, for many instances, these routings yield optimal solution costs very close to the solution costs provided by affine routing. Moreover, their solution times are order of magnitude smaller than those of affine routings, thus providing an interesting alternative to affine routing when the available computing time is limited.

6 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 simplified volume routings are special cases of affine routings and study when the inclusions are strict. We show then that the general volume routing is no more flexible than its simplifications whenever the uncertainty set is the polytope introduced by Bertsimas and Sim. We complete our comparison by examples and a computational study of some of these routings.

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.

7 Acknowledgements

The author is grateful to Christian Raack and to a referee for numerous comments that helped 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.

Instances			Cost			Time			
	$ K $	Γ	\mathcal{F}^{vs}	\mathcal{F}^{vg}	\mathcal{F}^{aff}	\mathcal{F}^{stat}	\mathcal{F}^{vs}	\mathcal{F}^{vg}	\mathcal{F}^{aff}
janos-us	10	1	5.8	5.8	5.8	0.03	0.06	0.06	3.73
	10	2	0	0	0	0.02	0.05	0.08	1.87
	20	1	7.2	7.2	7.2	0.09	0.14	0.31	43.7
	20	2	6.2	6.2	6.2	0.11	0.17	0.31	64.2
	20	3	0.9	0.9	0.9	0.06	0.19	0.27	51.2
	30	1	7.5	7.5	7.5	0.2	0.37	0.67	160
	30	3	6.6	6.6	6.6	0.16	0.42	0.72	179
	30	5	0	0	0	0.12	0.34	0.72	175
	40	1	8.2	8.2	8.2	0.11	0.61	0.37	528
	40	4	2.9	2.9	2.9	0.08	0.75	0.47	723
	40	6	0	0	0	0.08	0.62	0.33	345
	50	1	8.4	8.4	8.4	0.11	0.95	0.55	1933
	50	3	7.2	7.2	7.3	0.11	1.05	0.66	2010
	50	5	2	2	2.1	0.11	1.08	0.75	3439
	50	7	0	0	0	0.11	0.98	0.5	1360
sun	10	1	8.7	8.7	9.8	0.03	0.06	0.08	1.59
	10	2	1.1	1.1	2.7	0.03	0.08	0.08	1.7
	20	1	8.9	8.9	9.9	0.08	0.14	0.33	25.9
	20	2	8	8	8.9	0.08	0.14	0.31	27.9
	20	3	4.4	4.4	6.3	0.08	0.17	0.34	44.6
	30	1	7.9	7.9	9.2	0.16	0.45	0.76	135
	30	3	6.6	6.6	8.8	0.14	0.44	0.86	146
	30	5	1.7	1.7	2.2	0.14	0.44	0.7	132
	40	1	7.7	7.7	8.5	0.16	0.69	1.31	542
	40	4	6.8	6.8	8.8	0.19	0.72	1.28	563
	40	6	2.8	2.8	4.6	0.16	0.76	1.2	645
	50	1	7	7	8	0.2	1.14	0.34	1956
	50	3	8.6	8.6	10.4	0.25	1.25	0.37	1868
	50	5	6.5	6.5	8.6	0.27	1.29	0.41	1861
	50	7	3.8	3.8	6.2	0.23	1.29	0.33	1691
giul39	10	1	2.6	2.6	2.6	0.05	0.08	0.09	3.6
	10	2	1	1	2.3	0.05	0.11	0.11	3.46
	20	1	4.1	4.1	4.2	0.14	0.39	0.41	62.7
	20	2	5.5	5.5	5.8	0.14	0.41	0.42	65.2
	20	3	4.5	4.5	4.9	0.14	0.36	0.55	65.1
	30	1	7.8	7.8	8.1	0.28	0.72	1.12	422
	30	3	7.4	7.4	8.0	0.25	0.86	1.25	466
	30	5	4	4	5.1	0.27	0.78	1.17	491
	40	1	8.8	8.8	9.2	0.47	1.34	2.03	1292
	40	4	6.8	6.8	7.5	0.45	1.45	2.15	1385
	40	6	3.7	3.7	4.9	0.47	1.37	2	1379
	50	1	8.2	8.2	8.6	0.48	2	3.17	4601
	50	3	9.6	9.6	10.6	0.55	2.23	3.18	4573
	50	5	7.2	7.2	8.1	0.53	2.23	3.48	4377
	50	7	4.2	4.2	5	0.5	2.03	3.65	3774

Table 1 Numerical results.

2. D. Applegate and E. Cohen. Making intra-domain routing robust to changing and uncertain traffic demands: understanding fundamental tradeoffs. In *Proceedings of the 2003 conference on Applications, technologies, architectures, and protocols for computer communications*, SIGCOMM '03, pages 313–324. ACM, 2003.
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 uncertain linear programs. *Operations Research Letters*, 25(1):1–13, 1999.
10. D. Bertsimas and C. Caramanis. Finite adaptability in multistage linear optimization. *Automatic Control, IEEE Transactions on*, 55(12):2751–2766, 2010.
11. D. Bertsimas and M. Sim. The Price of Robustness. *Operations Research*, 52(1):35–53, 2004.
12. C. Chekuri, G. Oriolo, M. G. Scutellà, and F. B. Shepherd. Hardness of robust network design. *Networks*, 50(1):50–54, 2007.
13. N. G. Duffield, P. Goyal, A. Greenberg, P. Mishra, K. K. Ramakrishnan, and J. E. van der Merive. A flexible model for resource management in virtual private networks. *SIGCOMM Comput. Commun. Rev.*, 29(4):95–108, 1999.
14. J. A. Fingerhut, S. Suri, and J. S. Turner. Designing least-cost nonblocking broadband networks. *Journal of Algorithms*, 24(2):287 – 309, 1997.
15. N. Goyal, N. Olver, and F. B. Shepherd. Dynamic vs. oblivious routing in network design. In *Proceedings of the ESA 2009*, pages 277–288, 2009.
16. 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.
17. ILOG CPLEX Division, Gentilly, France. *ILOG. ILOG CPLEX 12.4 Reference Manual.*, 2007.
18. A. M.C.A. Koster, M. Kutschka, and C. Raack. Robust network design: Formulations, valid inequalities, and computations. Technical Report 11-34, ZIB, 2011.
19. M. Minoux. Robust network optimization under polyhedral demand uncertainty is np-hard. *Discrete Applied Mathematics*, 158(5):597–603, 2010.
20. F. Ordóñez and J. Zhao. Robust capacity expansion of network flows. *Networks*, 50(2):136–145, 2007.
21. S. Orłowski, M. Pióro, A. Tomaszewski, and R. Wessäly. SNDlib 1.0—Survivable Network Design Library. *Networks*, 55(3):276–286, 2010. <http://sndlib.zib.de>.
22. 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.
23. M. Poss and C. Raack. Affine recourse for the robust network design problem: between static and dynamic routing. *Networks*, 61(2):180–198, 2013.
24. M. G. Scutellà. On improving optimal oblivious routing. *Operations Research Letters*, 37(3):197–200, 2009.
25. M. G. Scutellà. Hardness of some optimal oblivious routing generalizations. Technical Report TR-10-05, Università di Pisa, 2010.
26. A. L. Soyster. Convex programming with set-inclusive constraints and applications to inexact linear programming. *Operations Research*, 21:1154–1157, 1973.