

UNIVERSITÀ DI PISA
DIPARTIMENTO DI INFORMATICA

TECHNICAL REPORT: TR-09-23

Static and dynamic routing under the single-source Hose model

Antonio Frangioni

Fausto Pascali

Maria Grazia Scutellà

ADDRESS: Largo B. Pontecorvo 3, 56127 Pisa, Italy. TEL: +39 050 2212700 FAX: +39 050 2212726

Static and dynamic routing under the single-source Hose model

Antonio Frangioni *

Fausto Pascali †

Maria Grazia Scutellà ‡

Abstract

This paper addresses special cases of the robust network design problem under the single-source Hose model. We show that, in the case of unitary bounds, the static and the dynamic routing approaches lead to the same optimal solution, and this is true for both the splittable and the unsplittable scenarios. As a consequence, in such a special case, the robust network design problem with (splittable) dynamic routing is polynomially solvable, whereas the problem is coNP -Hard under the general single-source Hose model. The results are based on the fact that the single-source Hose polyhedron with unitary bounds is dominated by a polynomial number of demand vectors. A feasible static routing can then be constructed as a convex combination of a set of routing templates which are feasible for the dominant demand vectors. The equivalence between static and dynamic routing is a consequence of those results, and it can also be generalized to some single-source Hose cases with non unitary bounds.

Keywords: *Robust Optimization, Network Design, Routing, Single-source, Hose Model.*

1 Introduction

Let $\mathcal{G} = (\mathcal{V}, \mathcal{A})$ be a directed network, with $|\mathcal{V}| = n$ and $|\mathcal{A}| = m$. Let \mathcal{K} be a set of k origin-destination pairs which represent users that wish to communicate, and c_{ij} denote the non-negative cost of installing a unit of capacity along arc $(i, j) \in \mathcal{A}$. Let \mathcal{D} be a bounded non-empty polyhedron describing the possible *non-simultaneous* demands between the given origin-destination pairs. The *robust network design problem* (RND) on \mathcal{G} consists of determining a minimum cost capacity allocation for the arcs of \mathcal{G} such that the network is able to support each demand in \mathcal{D} .

Several variants and generalizations of RND have been proposed in the literature in the last decade. Concerning the routing constraints, each origin-destination pair may be required to communicate through a single path (*unsplittable* routing), or the traffic can be split among different paths (*splittable* routing). In addition, the routing can be *dynamic*, that is, it can change as the traffic demand varies in \mathcal{D} , or *static*, that is, the same routing template must be used for each traffic demand in \mathcal{D} . Observe that static routing—also referred to as “oblivious” [2] or “stable” [4]—can be preferable in network applications where migrating from one routing to another one is costly [3]. Clearly, even if the objective function only depends on the capacity allocation costs, the optimal solution value depends on the considered routing constraints; in general, splittable routing leads to a cheaper solution than unsplittable routing, and dynamic routing leads to a cheaper solution than static routing.

Concerning the shape of the demand polyhedron \mathcal{D} , the most widely studied case in the literature is the so-called *Hose* model, which simply specifies a bound on the maximum total traffic that each node can receive (considering the destination nodes) or send out (for the origins). In the particular case where a single node can send traffic along the network while all the other nodes are potential receivers, the model is referred to as the *single-source Hose* model.

*Dipartimento di Informatica, Università di Pisa, Polo Universitario della Spezia, Via dei Colli 90, 19121 La Spezia, Italy. E-mail: frangio@di.unipi.it

†Dipartimento di Informatica, Università di Pisa, Largo B. Pontecorvo 3, 56127 Pisa, Italy. E-mail: pascali@di.unipi.it

‡Dipartimento di Informatica, Università di Pisa, Largo B. Pontecorvo 3, 56127 Pisa, Italy. E-mail: scut@di.unipi.it

From a time complexity perspective, the very special case where \mathcal{D} is a singleton (the so-called *nominal case*) is clearly polynomially solvable both in the unsplittable and in the splittable scenario, as it can be easily solved by computing shortest paths between all origin-destination pairs [1]. In addition, it is useful to remark that, when \mathcal{D} is a singleton, then the static and the dynamic routing leads to the same optimal solution, and this implies that the problem is polynomially solvable also in the dynamic scenario. The splittable static case is still polynomially solvable for any separable convex \mathcal{D} by using standard results in robust optimization [4]. On the other hand, in the unsplittable case RND is \mathcal{NP} -Hard, since the Steiner tree problem can be reduced to the single-source Hose model [10]. As a consequence, RND is \mathcal{NP} -Hard in the unsplittable dynamic case, too. RND is also difficult in the splittable dynamic case, as it is $\text{co}\mathcal{NP}$ -Hard even in the special case of the single-source Hose model both in directed [10] and in undirected networks [6]. Thus, dynamic routing is, in general, substantially more difficult than static routing. This has motivated the study of “intermediate scenarios” such as the one where the demands in \mathcal{D} can be served by *two alternative* routing templates [14], which allows one to obtain cheaper solutions than static routing while being computationally tractable in some cases. Another possible approach is to study special cases of RND that are solvable in polynomial time due to the special structure of the demand polyhedron \mathcal{D} . For example, when \mathcal{D} is built upon a discrete number of scenarios whose number is polynomial in n and m , then RND with splittable dynamic routing is polynomially solvable since a compact linear programming formulation can be devised [12]. In addition, in the case of symmetric Hose polyhedra on undirected networks, RND with static unsplittable routing is polynomially solvable [9]. The fact that the time complexity of the problem depends on \mathcal{D} has motivated approaches [13] where the demand polyhedron can be reduced by discarding some *dominated* demand vectors. For a more detailed survey on time complexity results concerning RND, its variants and its generalizations, the interested reader is referred to [5] where several interesting variants of RND are discussed, e.g. cases where the optimum routing support is required to be a tree, and where the objective function involves congestion aspects.

In this paper, we partially answer to one of the open questions in [5]: can RND with splittable dynamic routing be solved in polynomial time in other cases than the mentioned ones where \mathcal{D} is a singleton or \mathcal{D} is built upon a polynomial number of scenarios? We prove that this is true for the single-source Hose model with unitary bounds. As for the nominal case, the result follows from a stronger result. We prove in fact that imposing a static or a dynamic routing leads to the same optimum solution, and this is true for both the splittable and the unsplittable case. Such properties still hold true when the source bound is smaller than any receiver bound, and when the source bound is greater than or equal to the sum of all receiver bounds. In both cases, by using the idea of domination between demand vectors [13], we reduce the demand polyhedron to a polynomial number of dominant points. Although all results are proved for the case of directed networks, they easily generalize to the undirected case.

The paper is organized as follows. Firstly we introduce the unitary single-source Hose model, and recall some preliminary results and notions about RND. The main result is then proved. The more general cases follow. We end with two examples. The first shows that dynamic routing may lead to a cheaper solution than static routing for single-source Hose models where the source bound is greater than some receiver bounds, but less than the sum of all receiver bounds. The latter example shows that, also in the unitary single-source Hose case, splittable routing may lead to a cheaper solution than unsplittable routing.

2 The unitary single-source Hose model

Let y denote a vector of routing variables, i.e., y_{ij}^{st} be the fraction of the demand of $(s, t) \in \mathcal{K}$ to be routed along the arc $(i, j) \in \mathcal{A}$. Then $y: \mathcal{A} \times \mathcal{K} \rightarrow [0, 1]$ is a *routing template* if it satisfies the following flow conservation constraints

$$\sum_{(j,i) \in BS(i)} y_{ji}^{st} - \sum_{(i,j) \in FS(i)} y_{ij}^{st} = \phi_i^{st} = \begin{cases} -1 & \text{if } i = s, \\ 1 & \text{if } i = t, \\ 0 & \text{otherwise,} \end{cases}$$

for all $i \in \mathcal{V}$ and $(s, t) \in \mathcal{K}$, where, as customary, $FS(i)$ and $BS(i)$ denote the set of arcs leaving node i and entering it, respectively. This can be equivalently restated in compact form as $Ey^{st} = \phi^{st}$ for all $(s, t) \in \mathcal{K}$, where E denotes the node-arc incidence matrix of \mathcal{G} and $\phi^{st} = [\phi_i^{st}]_{i \in \mathcal{V}}$. Hereafter we shall denote the set

of all routing templates by \mathcal{Y} , and by x a vector of design variables such that x_{ij} denotes the amount of capacity to be allocated to the arc $(i, j) \in \mathcal{A}$.

Definition 1 Given a routing template $y \in \mathcal{Y}$ and a capacity allocation $x \in \mathbb{R}_+^m$, the pair (y, x) supports \mathcal{D} if

$$\sum_{(s,t) \in \mathcal{K}} y_{ij}^{st} d_{st} \leq x_{ij} \quad (i, j) \in \mathcal{A}, d \in \mathcal{D} . \quad (1)$$

In this case, y is said to be a feasible static routing with respect to x and \mathcal{D} .

When the routing can change dynamically with $d \in \mathcal{D}$, then y must become a *routing function* $y : \mathcal{A} \times \mathcal{K} \times \mathcal{D} \rightarrow [0, 1]$, where $y(d)_{ij}^{st}$ denotes the fraction of the demand of (s, t) to be routed along (i, j) when the demand vector is d . Of course, the notion of feasibility has to be changed accordingly.

Definition 2 The capacity allocation $x \in \mathbb{R}_+^m$ supports \mathcal{D} if, for each $d \in \mathcal{D}$, there exists a routing template $y(d) \in \mathcal{Y}$ such that

$$\sum_{(s,t) \in \mathcal{K}} y(d)_{ij}^{st} d_{st} \leq x_{ij} \quad (i, j) \in \mathcal{A} . \quad (2)$$

The family $y(d)$, $d \in \mathcal{D}$, is said to be a feasible dynamic routing with respect to x .

Of course, (1) implies (2), since one can use the same y for all $d \in \mathcal{D}$, while the converse, in general, is not true. Hence, dynamic routing may allow to accommodate cheaper capacity allocations than static routing. In fact, let $OPT(\cdot)$ denote the optimal value of a problem, F and G denote respectively the robust network design problem with splittable and unsplittable routing, and dyn denotes the case of dynamic routing (hence no further qualification denotes static routing). We extend a result of [11] by comparing the optimum solution values of instances with the same underlying data (network \mathcal{G} and demand polyhedron \mathcal{D}) also in the dynamic setting.

Theorem 1 Both in the directed and undirected case,

$$OPT(F \text{ dyn}) \leq OPT(F) \leq OPT(G) \quad (3)$$

Proof. The first inequality has already been commented upon: any static routing is a dynamic one. The second inequality comes from the fact that (F) is the relaxation of (G) obtained by removing the requirement to use a single path for each origin-destination pair. ■

By the same token, $OPT(G \text{ dyn}) \leq OPT(G)$; however, nothing is known in general about the relationship between $OPT(G \text{ dyn})$ and $OPT(F)$. In Section 5, some examples will be provided where the inequalities in (3) hold as strict inequalities.

We are interested in the so-called asymmetric Hose model [7, 8], where each terminal $v \in \mathcal{V}$ has an upper bound b_v^{out} on the cumulative amount of traffic that can be sent by v , as well as an upper bound b_v^{in} on the cumulative amount of traffic that can be received by v . Formally, the *asymmetric Hose polyhedron* can be defined as follows:

$$\mathcal{D}_{Asym} = \left\{ d \in \mathbb{R}_+^k : \sum_{t : (v,t) \in \mathcal{K}} d_{vt} \leq b_v^{out}, \sum_{s : (s,v) \in \mathcal{K}} d_{sv} \leq b_v^{in} \quad v \in \mathcal{V} \right\} .$$

If the traffic can only be sent by a single *source* node, say r , whereas only *destination* nodes $t \in \mathcal{T} \subseteq \mathcal{V} \setminus \{r\}$ can receive traffic, one obtains the *single-source Hose polyhedron*

$$\mathcal{D}_{Ss} = \left\{ d \in \mathbb{R}_+^{|\mathcal{T}|} : \sum_{t \in \mathcal{T}} d_{rt} \leq b_r^{out}, d_{rt} \leq b_t^{in} \quad t \in \mathcal{T} \right\} .$$

In this special case, the number of origin-destination pairs (or commodities) is equal to the number of destination nodes, i.e., $k = |\mathcal{T}|$. Since the source r belongs to all pairs, hereafter y^{rt} will be denoted simply by y^t .

In this paper we consider the further specialization of the single-source Hose model where all bounds are unitary, i.e., the *unitary single-source Hose polyhedron*

$$\mathcal{D}_{Uss} = \left\{ d \in \mathbb{R}_+^{|\mathcal{T}|} : \sum_{t \in \mathcal{T}} d_{rt} \leq 1 \right\}.$$

Clearly, the inequalities $d_{rt} \leq 1$ relative to the destination nodes $t \in \mathcal{T}$ are redundant and can be dropped. \mathcal{D}_{Uss} is therefore a very simple polyhedron, the k -dimensional tetrahedron Δ^k . It is well-known that its vertices are 0 and all the unitary vectors e^t , $t \in \mathcal{T}$, of the orthogonal basis of the demand space, i.e., $e_h^t = 1$ if $h = t$, and $e_h^t = 0$ otherwise; in other words, $\Delta^k = \text{conv}(0, \Lambda^k)$ where $\Lambda^k = \{ \lambda \geq 0 : \sum_h \lambda_h = 1 \}$ is the unitary simplex. Thus, in this case RND with (splittable) dynamic routing is polynomially solvable.

The result follows from the next lemma which states that, given a routing template for each of a discrete set of demand vectors, it is possible to build up (in linear time) a feasible routing for each demand vector belonging to the convex hull of the demand vectors¹.

Lemma 1 *Any capacity allocation $x \in \mathbb{R}^m$ that supports a finite set $D = \{ d^h \in \mathbb{R}^k \}_{h \in H}$ of demand vectors also supports any demand $d \in \text{conv}(D)$.*

Proof. By hypothesis, for each $h \in H$ there exists a routing template $y(d^h)$ which is feasible w.r.t. x . Consider a vector of convex multipliers $\lambda \in \Lambda^{|H|}$ and the corresponding demand vector $d^\lambda = \sum_{h \in H} \lambda_h d^h \in \text{conv}(D)$; we claim that

$$z(d^\lambda)_{ij}^{st} = \sum_{h \in H} \lambda_h \frac{y(d^h)_{ij}^{st} d_{st}^h}{d_{st}^\lambda} \quad (4)$$

defines a feasible routing template with respect to x . Of course, (4) is well-defined only if $d_{st}^\lambda > 0$, but $d_{st}^\lambda = 0$ implies $d_{st}^h = 0$ for all $h \in N$, and therefore $y(d^h)_{ij}^{st} = 0$ for all $(i, j) \in A$. Thus, in the (unlikely) case that $d_{st}^\lambda = 0$, one can take $z(d^\lambda)_{ij}^{st} = 0$ for all $(i, j) \in A$. Now, for all $v \in \mathcal{V}$ and $(s, t) \in \mathcal{K}$

$$\begin{aligned} \sum_{(i,j) \in A} E_{ij}^v z(d^\lambda)_{ij}^{st} &= \sum_{(i,j) \in A} E_{ij}^v \frac{\sum_{h \in H} \lambda_h y(d^h)_{ij}^{st} d_{st}^h}{d_{st}^\lambda} = \\ &= \frac{1}{d_{st}^\lambda} \sum_{h \in H} \lambda_h d_{st}^h \sum_{(i,j) \in A} E_{ij}^v y(d^h)_{ij}^{st} = \end{aligned} \quad (5)$$

$$= \frac{1}{d_{st}^\lambda} \sum_{h \in H} \lambda_h d_{st}^h \phi_v^{st} = \phi_v^{st} \quad (6)$$

where (5) is due to algebraic manipulations and (6) follows by the hypothesis that $y(d^h)$ are routing templates for all $h \in H$. Moreover,

$$\sum_{(s,t) \in \mathcal{K}} d_{st}^\lambda z(d^\lambda)_{ij}^{st} = \sum_{(s,t) \in \mathcal{K}} d_{st}^\lambda \frac{\sum_{h \in H} \lambda_h y(d^h)_{ij}^{st} d_{st}^h}{d_{st}^\lambda} = \quad (7)$$

$$= \sum_{h \in H} \lambda_h \sum_{(s,t) \in \mathcal{K}} y(d^h)_{ij}^{st} d_{st}^h \leq \sum_{h \in H} \lambda_h x_{ij} = x_{ij} \quad (8)$$

where (7) follows from (4) and (8) follows by algebraic manipulations and the hypothesis that $y(d^h)$ is feasible w.r.t. x for each $h \in H$ (as well as by the definition of λ). ■

This result will be used in the next section to show that, for the unitary single-source Hose model, the dynamic and the static routing approaches lead to the same optimal solution; clearly, this equivalence also implies the polynomial solvability of the dynamic case. In addition, the result is true for both the splittable and the unsplittable case. An instrumental concept in our analysis is that of domination between demand vectors [13]. A demand vector d^1 *dominates* d^2 if any capacity allocation $x : \mathcal{A} \rightarrow \mathbb{R}_+$ supporting d^1 also supports d^2 . Moreover, d^1 *totally dominates* d^2 if any pair (y, x) supporting d^1 also supports d^2 . Clearly, total domination implies domination. A nice characterization of total domination is the following:

¹Indeed, the result above also follows from the results in [12], since a polynomial number of traffic demands has to be taken into account in order to route all demands in \mathcal{D}_{Uss}

Theorem 2 [13, Theorem 2.5] d^1 totally dominates d^2 if and only if $d_{st}^1 \geq d_{st}^2$ for all $(s, t) \in \mathcal{K}$.

3 Static versus dynamic routing for \mathcal{D}_{Uss}

We first identify the dominant extreme points of the polyhedron \mathcal{D}_{Uss} .

Lemma 2 For each $d \in \mathcal{D}_{Uss}$ there exists a demand vector $\bar{d} \in \Lambda^k$ such that \bar{d} totally dominates d .

Proof. The property follows immediately from Theorem 2 and the trivial observation that for each $d \in \mathcal{D}_{Uss}$ there exists $\bar{d} \in \Lambda^k$ such that $\bar{d} \geq d$. ■

As a consequence:

Corollary 1 A capacity allocation x supports e^t for each $t \in \mathcal{T}$ if and only if x supports each demand in \mathcal{D}_{Uss} .

Proof. $[\Rightarrow]$ If x supports e^t for each $t \in \mathcal{T}$, then from Lemma 1 there exists a feasible routing template $y(d)$ for each $d \in \Lambda^k$. So, the family $y(d)$, $d \in \Lambda^k$, is a feasible dynamic routing supported by x . From Lemma 2, the family $y(d)$, $d \in \Lambda^k$, is able to route each demand in \mathcal{D}_{Uss} . Therefore, x supports \mathcal{D}_{Uss} . $[\Leftarrow]$ is trivial since $e^t \in \mathcal{D}_{Uss}$ for each $t \in \mathcal{T}$. ■

We are now ready to state the main result.

Theorem 3 The following statements are equivalent:

- x supports \mathcal{D}_{Uss} (9)

- there exists $z \in [0, 1]^{m \times k}$ such that (z, x) supports \mathcal{D}_{Uss} (10)

Proof. (10) \Rightarrow (9) is trivial. In fact, as already remarked, a static routing is also a dynamic routing. In order to prove the implication (9) \Rightarrow (10) observe that, from Corollary 1, x supports \mathcal{D}_{Uss} if and only if x supports a routing template $y(e^t)$ for each $t \in \mathcal{T}$. The relevant characteristic of each vector $y(e^t)$, $t \in \mathcal{T}$, is that the only non zero demand concerns commodity t . Without loss of generality we can then assume that $y(e^t)_{ij}^q = 0$ for all $(i, j) \in \mathcal{A}$ and $q \in \mathcal{T} \setminus \{t\}$. Now, pick any $d \in \mathcal{D}_{Uss}$. We have $d = \sum_{t \in \mathcal{T}} \lambda_t e^t$ for $\lambda_t = d_{rt}$. Hence, by exploiting the fact that $y(e^t)_{ij}^q = 0$ for $q \neq t$, the feasible dynamic routing $z(d)$ determined according to (4) satisfies

$$z(d)_{ij}^t = \lambda_t y(e^t)_{ij}^t e_t^t / d_{rt}$$

which (using the relations $\lambda_t = d_{rt}$ and $e_t^t = 1$) yields

$$z(d)^t = y(e^t)^t \quad \forall d \in \mathcal{D}_{Uss} . \quad (11)$$

Therefore, $z(d)$ is actually *independent on d* , i.e., it is a *routing template* which can be simply denoted by z . We claim that z is a feasible static routing, and that (z, x) supports \mathcal{D}_{Uss} . Firstly, z satisfies the flow conservation constraints by construction. Moreover, let $y_{ij}^{max} = \max_{t \in \mathcal{T}} \{ y(e^t)_{ij}^t \}$ for all $(i, j) \in \mathcal{A}$. Then

$$\sum_{t \in \mathcal{T}} z_{ij}^t d_{rt} = \sum_{t \in \mathcal{T}} y(e^t)_{ij}^t d_{rt} \leq y_{ij}^{max} \sum_{t \in \mathcal{T}} d_{rt} \leq y_{ij}^{max} \leq x_{ij}$$

where the first inequality follows from the definition of y_{ij}^{max} , the second inequality follows from the definition of \mathcal{D}_{Uss} , and the third inequality follows from the hypothesis that $y(e^t)$ is a feasible routing with respect to x for each $t \in \mathcal{T}$. Hence, (z, x) supports \mathcal{D}_{Uss} . ■

Corollary 2 If the routing templates corresponding to the dominant demands in \mathcal{D}_{Uss} are integer valued, i.e., $y(e^t) \in \{0, 1\}^{m \times k}$ for each $t \in \mathcal{T}$, then the static routing z in Theorem 3 is also unspittable, that is, $z \in \{0, 1\}^{m \times k}$.

Proof. Just consider the construction rule specified in equation (11): if all $y(e^t)^t$ are integer vectors, then z is an integer vector, too. ■

Theorem 3 and Corollary 2 allow to strengthen the results of Theorem 1 in the case of the unitary single-source Hose polyhedron, for which one has

$$OPT(F \text{ dyn}) = OPT(F) \leq OPT(G \text{ dyn}) = OPT(G) .$$

These properties hold true also in more general cases, as shown in the following section.

4 Generalizations

4.1 Unlimited source bound

When the source node has a unlimited bound b_r^{out} (equivalently, r can push an amount of traffic greater than or equal to the overall requirement of the destination nodes), then the dominant demand vectors in \mathcal{D}_{Uss} are determined only by the destination node bounds. Let d^b denote the demand vector such that $d_{rt}^b = b_t^{in}$, $t \in \mathcal{T}$. The following property then holds true:

Theorem 4 *If $\sum_{t \in \mathcal{T}} b_t^{in} \leq b_r^{out}$, then RND under \mathcal{D}_{Ss} is equivalent to RND under $\mathcal{D} = \{ d^b \}$, which is the nominal case where all destinations require the maximum amount of traffic they can.*

Proof. Since $d_{rt} \leq b_{in}^t$ for all $t \in \mathcal{T}$, one has

$$\sum_{t \in \mathcal{T}} d_{rt} \leq \sum_{t \in \mathcal{T}} b_{in}^t \Rightarrow \sum_{t \in \mathcal{T}} d_{rt} \leq b_r^{out} \quad \forall d \in \mathcal{D}_{Ss} .$$

The Hose constraint related to the source r is so redundant. According to the characterization of totally domination previously reviewed, it follows that each demand in \mathcal{D}_{Ss} is totally dominated by the single demand d^b , which belongs to \mathcal{D}_{Ss} . ■

As a consequence, the single-source Hose model with unlimited source bound can be solved by solving a nominal network design problem w.r.t. the single demand vector d^b . Therefore, also in this case $OPT(F \text{ dyn}) = OPT(F)$, and this is true also for the unsplittable case.

4.2 Source bound limited by each receiver bound

Let us now turn to the special case of \mathcal{D}_{Ss} where $b_r^{out} \leq b_t^{in}$, $t \in \mathcal{T}$. We will show that this special case can be reduced to the single-source Hose model with unitary bounds. We will need a technical result from [11], with the following notation: (\mathcal{G}, b, c) will be used to denote an instance of RND, under the Hose model, specified by a network \mathcal{G} , an upper bound vector b and a cost vector c .

Theorem 5 [11, Lemma 3.4] *For each $\beta \in \mathbb{R}_+$, the instance $(\mathcal{G}, \beta b, c)$ has a feasible solution of value βC if and only if the instance (\mathcal{G}, b, c) has a feasible solution of value C .*

Theorem 6 *If $b_r^{out} \leq b_t^{in}$ for all $t \in \mathcal{T}$, then RND with \mathcal{D}_{Ss} is equivalent to RND with \mathcal{D}_{Uss} .*

Proof. Consider the RND instance where the upper bound vector is $b = (b_r^{out}, b_1^{in}, \dots, b_k^{in})$; from Theorem 5 we can scale b by a factor b_r^{out} , thereby obtaining an equivalent RND problem (up to the chosen scaling factor) with upper bound vector $(1, \bar{b}_1^{in}, \dots, \bar{b}_k^{in})$, where $\bar{b}_i^{in} = b_i^{in} / b_r^{out} \geq 1$. Clearly, the inequalities $d_{rt} \leq \bar{b}_i^{in}$ can be removed since they are redundant ($\sum_{t \in \mathcal{T}} d_{rt} \leq 1 \Rightarrow d_{rt} \leq 1$); therefore, the corresponding instance of RND under the unitary single-source Hose model is equivalent (up to the scaling factor) to the original RND instance. ■

Therefore, also in this case $OPT(F \text{ dyn}) = OPT(F)$, and this is true also in the unsplittable case.

5 Some examples

Let us first provide an example showing that the dynamic routing may lead to a cheaper solution, with respect to the static routing, for single-source Hose models where the source bound is greater than some receiver bounds, but less than the sum of all receiver bounds (i.e., $OPT(F_{dyn}) < OPT(F)$). Consider the network in Fig. 1(a) where the arcs (a,b) and (a,c) cost 1.6 and all the other arcs cost 1. Let $\mathcal{K} = \{(a,b), (a,c)\}$, $b_a^{out} = 1.5$, $b_b^{in} = 1$, $b_c^{in} = 1$, while the other bounds are 0. These bounds define a single-source Hose model with respect to the source node a , where the source bound is greater than the destination bounds, but less than the sum of all receiver bounds.

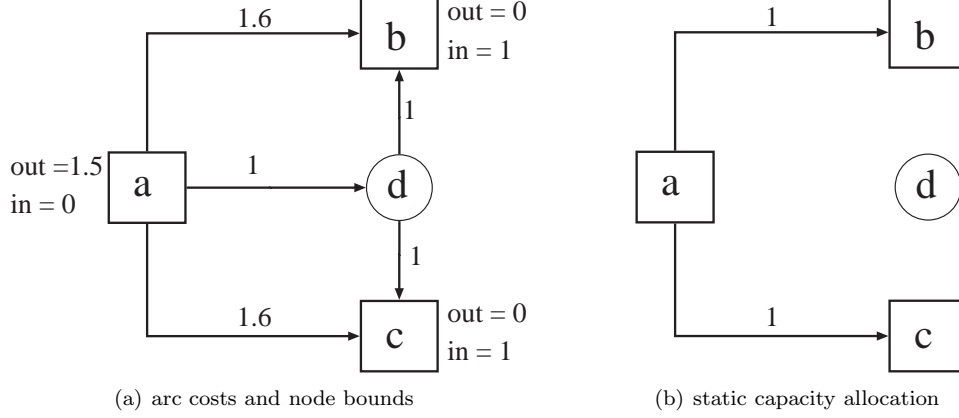


Figure 1: *Single-source Hose model*

The dominant demand vectors of this Hose polyhedron are $d^1 = (1, 0.5)$ and $d^2 = (0.5, 1)$. Consider the following routing template: a sends one unit of flow to b along (a,b) and a sends one unit of flow to c along (a,c) . In order to satisfy each demand in the Hose polyhedron, a capacity 1 has therefore to be installed on (a,b) and a capacity 1 has to be installed on (a,c) . Such a capacity allocation costs 3.2 (it is reported in Fig. 1(b)), and it is optimal for the considered RND instance in the case of static routing.

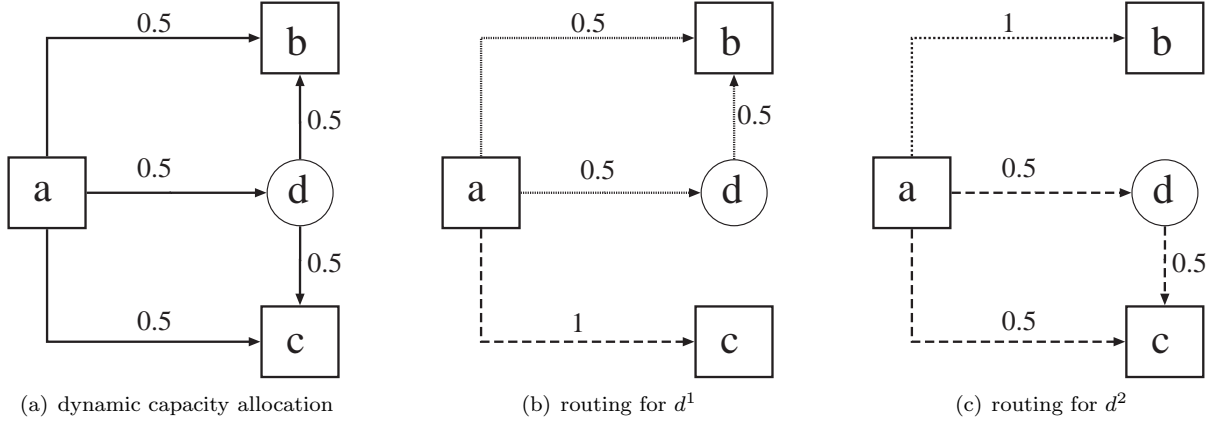


Figure 2: *Single-source Hose model*

Let us show now that a cheaper solution does exist in the case of dynamic routing. Consider the capacity allocation, of cost 3.1, reported in Fig 2(a), where we installed capacity 0.5 on each arc of the network. The routing template in Fig 2(b), where a sends 0.5 to b along (a,b) , a sends 0.5 to b along (a,d,b) and a sends 1 to c along (a,c) , is feasible for the demand vector $d^1 = (1, 0.5)$, while the routing template in Fig 2(c), where a sends a unitary flow to b along (a,b) , a sends 0.5 to c along (a,c) and a sends 0.5 to c along (a,d,c) ,

is feasible for the demand vector $d^2 = (0.5, 1)$. Therefore, since d^1 and d^2 are dominant for the considered single-source polyhedron, the stated capacity allocation, of cost 3.1, supports the Hose polyhedron, as follows from Lemma 1 and from the concept of domination between demand vectors.

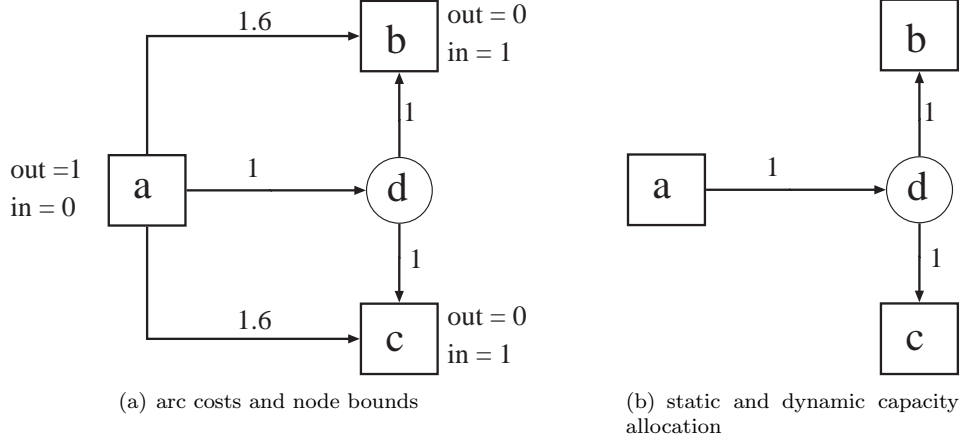


Figure 3: Unit single-source Hose model

Note that, if the source bound would be $b_a^{out} = 1$ (Fig 3(a)), and therefore the instance satisfies the hypothesis of Theorem 3, then the optimal solution, of cost 3, would consist on installing a capacity 1 on the arcs (a, d) , (d, b) , (d, c) (Fig 3(b)), and this is true for both dynamic and static routing. Observe that, since the induced subgraph is a tree, the splittable and the unsplittable scenarios coincide in the considered example.




Figure 4:

However, this is not necessarily true. The example below shows in fact that, also in the unitary single-source Hose case, the splittable scenario may lead to a cheaper solution than the unsplittable scenario, (i.e., $OPT(F) < OPT(G)$). Consider the network in Fig. 4, where all arcs cost 1. Let $\mathcal{K} = \{(r, d), (r, e), (r, f)\}$, $b_r^{out} = b_d^{in} = b_e^{in} = b_f^{in} = 1$, while all the other bounds are 0. These bounds define a unitary single-source Hose polyhedron with respect to the source node r . Recall that, for the unitary single-source Hose case, $OPT(F \text{ dyn}) = OPT(F)$ and $OPT(G \text{ dyn}) = OPT(G)$, as proved in Section 3. Therefore, there is no need to specify whether the static or the dynamic routing scenario has to be addressed. In order to satisfy each demand in the Hose polyhedron, a capacity 1 has to be installed on all the arcs of the subset $\{(r, b), (r, c), (b, d), (b, e), (c, f)\}$. Such a capacity allocation costs 5 (it is reported in Fig. 5(a)), and it is optimal for the considered RND instance in the case of unsplittable routing. However, a cheaper solution exists in the case of splittable routing. Consider in fact the capacity allocation reported in Fig 5(b), of cost

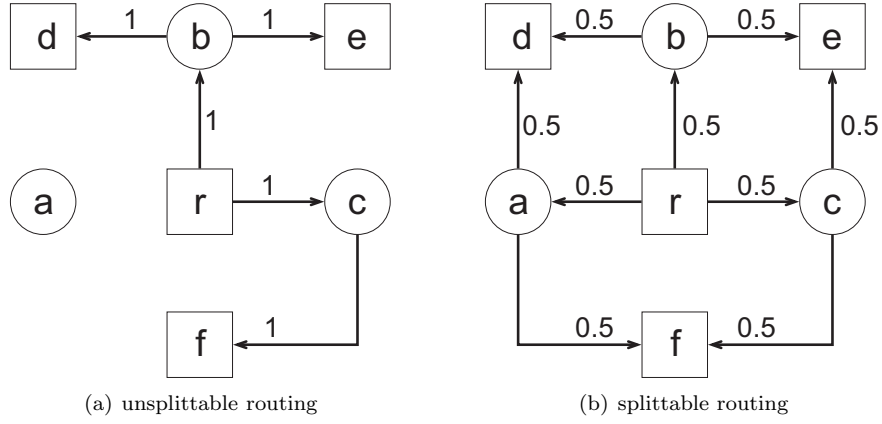


Figure 5: Capacity allocations

4.5, where we install a capacity 0.5 on each arc of the network. Such a capacity allocation is supported by the (splittable) routing template which halves the demand of each commodity (r, t) between the two paths in the network linking r to t . Therefore, splittable may be cheaper than unsplittable also for the unitary single-source Hose case.

6 Conclusions

In this paper we have investigated how the particular shape of the demand polyhedron \mathcal{D} may affect the time complexity of related robust network design problems. We have addressed special cases of RND under the single-source Hose model and shown that, in the case of unitary bounds, the static and the dynamic routing approaches lead to the same optimal solution. This property is true for both the splittable and the unsplittable case, and it can be generalized to some single-source Hose cases with non unitary bounds.

An interesting line of research would be to provide a full characterization of those polyhedra \mathcal{D} such that the $(F \text{ dyn})$ version of RND is polynomially solvable. Two sufficient conditions that guarantee the polynomial solvability of the splittable dynamic version are:

- the number of the dominant traffic vectors in \mathcal{D} is polynomial in the input size;
- $OPT(F \text{ dyn}) = OPT(F)$.

Observe that the unitary single-source Hose polyhedron satisfies both conditions. Observe also that the latter is not a necessary condition in the general case, as revealed by the example in Fig 1(a). It is an open question whether polyhedra do exist whose number of dominant traffic vectors is not polynomial, but the dynamic and the static routing lead to the same optimal solution. In addition, since $OPT(F) < OPT(G \text{ dyn})$ may hold true, as illustrated by the example in Fig 4, we think that the relationship between such two versions of RND would be worth investigating.

References

- [1] R. K. Ahuja, T. L. Magnanti, and J. B. Orlin. *Network flows: theory, algorithms, and applications*. Prentice-Hall, Inc., Upper Saddle River, NJ, USA, 1993.
- [2] D. Applegate and E. Cohen. Making intra-domain routing robust to changing and uncertain traffic demands: understanding fundamental tradeoffs. In *SIGCOMM '03: Proceedings of the 2003 conference on Applications, technologies, architectures, and protocols for computer communications*, pages 313–324, New York, NY, USA, 2003. ACM.

- [3] W. Ben-Ameur and H. Kerivin. Networks new economical virtual private. *Communication of the ACM*, 46(6):69–73, 2003.
- [4] W. Ben-Ameur and H. Kerivin. Routing of uncertain traffic demands. *Optimization and Engineering*, 6(3):283–313, 2005.
- [5] C. Chekuri. Routing and network design with robustness to changing or uncertain traffic demands. *SIGACT News*, 38(3):106–129, 2007.
- [6] C. Chekuri, F. B. Shepherd, G. Oriolo, and M. G. Scutellà. Hardness of robust network design. *Networks*, 50(1):50–54, 2007.
- [7] 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. In *SIGCOMM '99: Proceedings of the conference on Applications, technologies, architectures, and protocols for computer communication*, pages 95–108, New York, NY, USA, 1999. ACM Press.
- [8] J. A. Fingerhut, S. Suri, and J. S. Turner. Designing least-cost nonblocking broadband networks. *Journal Algorithms*, 24(2):287–309, 1997.
- [9] N. Goyal, N. Olver, and F. B. Shepherd. The vpn conjecture is true. In *STOC '08: Proceedings of the 40th annual ACM symposium on Theory of computing*, pages 443–450, New York, NY, USA, 2008. ACM.
- [10] 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, New York, NY, USA, 2001. ACM.
- [11] C. A. J. Hurkens, J. C. M. Keijsper, and L. Stougie. Virtual private network design: A proof of the tree routing conjecture on ring networks. *SIAM Journal of Discrete Mathematics*, 21(2):482–503, 2007.
- [12] M. Labbé, R. Séguin, P. Soriano, and C. Wynants. Network synthesis with non-simultaneous multi-commodity flow requirements: Bounds and heuristics. Technical report, ULB, 1999.
- [13] G. Oriolo. Domination between traffic matrices. *Mathematics of Operations Research*, 33(1):91–96, February 2008.
- [14] M. G. Scutellà. On improving optimal oblivious routing. *Operations Research Letters*, 37(3):197–200, May 2009.