Approximability of Robust Network Design: The Directed Case

Yacine Al-Najjar[1]
Huawei Technologies, Paris Research Center, France
Samovar, Telecom SudParis, Institut Polytechnique de Paris, France

Walid Ben-Ameur
Samovar, Telecom SudParis, Institut Polytechnique de Paris, France

Jérémie Leguay
Huawei Technologies, Paris Research Center, France

Abstract

We consider robust network design problems where an uncertain traffic vector belonging to a polytope has to be dynamically routed to minimize either the network congestion or some linear reservation cost. We focus on the variant in which the underlying graph is directed. We prove that an $O(\sqrt{k}) = O(n)$-approximation can be obtained by solving the problem under static routing, where k is the number of commodities and n is the number of nodes. This improves previous results of Hajiaghayi et al. [SODA'2005] and matches the $\Omega(n)$ lower bound of Ene et al. [STOC'2016] and the $\Omega(\sqrt{k})$ lower bound of Azar et al. [STOC'2003]. Finally, we introduce a slightly more general problem version where some flow restrictions can be added. We show that it cannot be approximated within a ratio of $k^{\frac{c}{\log \log k}}$ (resp. $n^{\frac{c}{\log \log n}}$) for some constant c. Making use of a weaker complexity assumption, we prove that there is no approximation within a factor of $2^{\log^{1-\epsilon} k}$ (resp. $2^{\log^{1-\epsilon} n}$) for any $\epsilon > 0$.

2012 ACM Subject Classification Theory of computation; Theory of computation → Data structures design and analysis; Theory of computation → Approximation algorithms analysis; Theory of computation → Routing and network design problems

Keywords and phrases Robust Optimization, Network Design, Approximation, Inapproximability, Competitive Ratio of Oblivious Routing

Digital Object Identifier 10.4230/LIPIcs.STACS.2022.6

1 Introduction

Network optimization [38, 28] plays a crucial role for telecommunication operators since it permits to carefully invest in infrastructures. As the traffic is continuously increasing, the network's capacity needs to be expanded through careful investments every year. However, the dynamic nature of the traffic due to ordinary daily fluctuations, long term evolution and unpredictable events requires to consider uncertainty on the traffic demand when dimensioning network resources. In this context, we provide new approximability results on two tightly related variants of the robust network design problem, the minimization of either the congestion or a linear cost.

Let us consider a directed graph $G = (V(G), E(G))$ representing a communication network. The traffic is characterized by a set of commodities $h \in \mathcal{H}$ associated to different node pairs and traffic values d_h. The demand vector $d = (d_h)_{h \in \mathcal{H}}$ is assumed to be uncertain and more precisely to belong to a polyhedral set $\mathcal{D}$. The *polyhedral model* was introduced in [6, 7] as an extension of the *hose model* [14, 17], where limits on the total traffic going into (resp. out of) a node are considered.

[1] corresponding author

The routing of a commodity h can be represented by a unit flow (also called routing template) $f_{h,\cdot} = (f_{h,e})_{e \in E(G)}$ from the source $s(h)$ to the sink $t(h)$.

When solving a robust network design problem, several objective functions can be considered. Given a capacity c_e for each edge e, one might be interested in minimizing the congestion given by $\max_{e \in E(G)} \frac{u_e}{c_e}$ where u_e is the reserved capacity on edge e. Another common objective function is given by the linear reservation cost $\sum_{e \in E(G)} \lambda_e u_e$. This can also represent the average congestion by taking $\lambda_e = \frac{1}{c_e}$. The goal is to choose a reservation vector u so that the network is able to support any demand vector $d \in \mathcal{D}$, i.e., there exists a (fractional) routing serving every commodity such that the total flow on each edge e is less than the reservation u_e.

The robust network design problem that we are focusing on in this paper, is referred to as *dynamic routing* in the literature since the network is optimized such that any realization of traffic vector in the uncertainty set has its own routing (i.e., $f_{h,\cdot}$ depends on d). The robust network design problem where a linear reservation cost is minimized was proved to be co-NP hard in [21] when the graph is directed. A stronger co-NP hardness result is given in [12] where the graph is undirected (this implies the directed case result). Some exact solution methods for robust network design have been considered in [13, 30]. Some special cases where dynamic routing is easy to compute have been described in [8, 18, 31]. For each of the two problems of congestion minimization and linear reservation cost minimization under dynamic routing, it is proved in [1] that the optimal value cannot be approximated within any constant (unless $P = NP$) and within $\Omega(\frac{\log n}{\log \log n})$ (under ETH assumption) for an undirected graph having n vertices. This leads again to the same inapproximability result for the directed case.

Routing with uncertain demands has received a significant interest from the community. As opposed to dynamic routing, *static routing* or *stable routing* was introduced in [6]: it consists in choosing a fixed flow $f_{h,\cdot}$ of value 1 for each commodity h. Static routing is also called *oblivious routing* in [2, 3]. In this case, polynomial-time algorithms to compute optimal static routing (with respect to either congestion or linear reservation cost) have been proposed [2, 3, 6, 7] based on either duality or cutting-plane algorithms.

To further improve solutions of static routing and overcome complexity issues related to dynamic routing, a number of restrictions on routing have been considered to design polynomial-time algorithms. This includes, for example, the approaches proposed in [5, 9, 27, 34, 35, 39].

Most of the literature studied the undirected case of the robust network design problem while only a few papers, such as [3, 7, 21, 24], address the directed case. In this work, we mainly focus on the approximability of robust network design problems under dynamic routing in directed networks, while minimizing either congestion or some linear reservation cost.

In the rest of this section, we summarize the main contributions of the paper and their positioning with regard to prior work. Then, we review some related state of the art.

1.1 Our results

- We prove that compared to dynamic routing, when static routing is considered, congestion is multiplied by a factor less than or equal to $\sqrt{8k}$ where k is the number of commodities. This implies that the gap between static routing and dynamic routing for the congestion minimization problem is $O(\sqrt{k}) = O(n)$ where n is the number of nodes. The best-known previous bound is $O(\sqrt{k} n^{\frac{1}{4}} \log n)$ and was given by [24]. The same $\sqrt{8k}$ bound applies to the linear reservation cost problem. The new upper bound matches the $\Omega(\sqrt{k})$ lower bound of [3] and the $\Omega(n)$ lower bound of [16].

- We introduce a more general version of the two robust network design problems (related to congestion and linear cost) by considering some flow restrictions (each commodity h can only be routed through edges inside a given subset E_h). The upper bound $\sqrt{8k} = O(n)$ is still valid and the static versions of the problems can still be solved in polynomial-time. We show some strong inapproximability results for this problem. More precisely, we prove that unless $NP \subseteq SUBEXP$, neither minimum dynamic congestion nor optimal linear cost can be approximated within a ratio of $k^{\frac{c}{\log\log k}}$ (resp. $n^{\frac{c}{\log\log n}}$) for some constant c. Making use of a weaker assumption, we get that unless $NP \subseteq QP$, there is no approximation within a factor of $2^{\log^{1-\epsilon} k}$ (resp. $2^{\log^{1-\epsilon} n}$) for any $\epsilon > 0$. This result improves the $\Omega(\frac{\log n}{\log\log n})$ inapproximability bound of [1] for the undirected case that also applies to the directed one.

1.2 Related work

Let us first consider that the graph is undirected and a linear cost is minimized. A result attributed to A. Gupta ([11], see also [19] for a more detailed presentation) leads to an $O(\log n)$ approximation algorithm for linear cost under dynamic fractional routing. Furthermore, this approximation is achieved by a routing on a (fixed) single tree. In particular, this shows that the ratios between the dynamic and the static solutions under fractional routing ($\frac{\text{lin}_{\text{sta-frac}}}{\text{lin}_{\text{dyn-frac}}}$) (lin denotes here the optimal linear cost of the solution) and between single path and fractional routing under the static model ($\frac{\text{lin}_{\text{sta-sing}}}{\text{lin}_{\text{sta-frac}}}$) is in $O(\log n)$ and provides an $O(\log n)$ approximation for static single path routing $\text{lin}_{\text{sta-sing}}$. On the other hand [33] shows that the static single path problem cannot be approximated within a $\Omega(\log^{\frac{1}{4}-\epsilon} n)$ ratio unless $NP \not\subset ZPTIME(n^{polylog(n)})$. As noticed in [19], this implies (assuming this complexity conjecture) that the gap $\frac{\text{lin}_{\text{sta-sing}}}{\text{lin}_{\text{sta-frac}}}$ is in $\Omega(\log^{\frac{1}{4}-\epsilon} n)$. [19] has shown that the gap $\frac{\text{lin}_{\text{sta-frac}}}{\text{lin}_{\text{dyn-frac}}}$ is $\Omega(\log n)$.

For the linear cost objective function and undirected graphs, an extensively studied polyhedron is the *symmetric* hose model. The demand vector is here not oriented (i.e, there is no distinction between a demand from i to j and a demand from j to i), and uncertainty is defined by considering an upper-bound limit b_i for the sum of demands related to node i. A 2-approximation has been found for the dynamic fractional case [17, 21] based on tree routing (where we route through a static tree that should be found) showing that $\frac{\text{lin}_{\text{sta-tree}}}{\text{lin}_{\text{dyn-frac}}} \leq 2$. It has been conjectured that this solution resulted in an optimal solution for the static single path routing. This question has been open for some time and has become known as the *VPN conjecture*. It was finally answered by the affirmative in [20]. The *asymmetric* hose polytope was also considered in many papers. An approximation algorithm is proposed to compute $\text{lin}_{\text{sta-sing}}$ within a ratio of 3.39 [15] (or more precisely 2 plus the best approximation ratio for the Steiner tree problem). If $\mathcal{D}$ is a *balanced* asymmetric hose polytope, i.e., $\sum_{v\in V} b_v^{out} = \sum_{v\in V} b_v^{in}$ where b_v^{in} (resp. b_v^{out}) is the upper bound for the traffic entering into (resp. going out of) v, then the best approximation factor becomes 2 [15]. Moreover, if we assume that $b_v^{out} = b_b^{in}$, then $\text{lin}_{\text{sta-sing}}$ is easy to compute and we get that $\text{lin}_{\text{sta-tree}} = \text{lin}_{\text{sta-sing}}$ [32]. In other words, there is some similarity with the case where $\mathcal{D}$ is a symmetric hose polytope.

When congestion is considered, [36] proved the existence of an oblivious (or static) routing with a competitive ratio of $O(\log^3 n)$ with respect to optimum routing of any traffic matrix. Then, [25] improved the bound to $O(\log^2 n \log\log n)$ and gave a polynomial-time algorithm to find such a static routing. Finally, [37] described an $O(\log n)$ approximation algorithm for static routing with minimum congestion. Notice that the bound given by static routing cannot provide a better bound than $O(\log n)$ since a lower bound of $\Omega(\log n)$ is achieved by static routing for planar graphs [29, 4]. It has also been shown in [23] that the gap between the dynamic fractional routing and a dynamic fractional routing restricted to a polynomial number of paths can be $\Omega(\frac{\log n}{\log\log n})$.

In a recent study on the approximability of robust network design [1] for the undirected case, it was proved that minimum dynamic congestion and the optimal linear cost cannot be approximated within any constant factor. Then using the ETH conjecture, it is shown there that they cannot be approximated within $\Omega(\frac{\log n}{\log\log n})$. This implies that the well-known $O(\log n)$ approximation ratio established in [37] is tight. Using a Lagrange relaxation approach, it is also shown in [1] that any α-approximation algorithm for the robust network design problem with linear reservation costs directly leads to an α-approximation for the problem of minimum congestion. This is used there to prove in a different way the $O(\log n)$ result of [37] starting from the one of [22] (attributed to A. Gupta and related to the linear cost minimization).

The closest papers to ours are [3, 16, 24]. When a directed graph is considered and congestion is minimized, [3] has shown that the gap between static fractional routing and dynamic fractional routing can be $\Omega(\sqrt{k})$ while [24] proves that the gap is upper-bounded by $O(\sqrt{k}n^{\frac{1}{4}}\log n)$. Since the instance provided in [3] contains vertices with large degree, [24] studied the version where the degree is less than some constant and all commodities have the same sink. An instance with a $\Omega(\sqrt{n})$ gap was then provided in [24], while the upper bound becomes $O(\sqrt{n}\log n)$. [24] considered also the case of symmetric demands (in that paper, symmetry means that for any two nodes u and v, the demand from u to v is equal to the demand from v to u) and shows that the upper bound of the static to dynamic ratio becomes $O(\sqrt{k}\log^{5/2} n)$. A general $\Omega(n)$ lower bound was later proposed in [16]. They also introduced the notion of balance for directed graphs. A weighted directed graph is α-balanced if for every subset $S \subseteq V$, the total weight of edges going from S to $V\backslash S$ is within a factor α of the total weight of edges directed from $V\backslash S$ to S. Using this new parameter, they show that for single source instances an upper bound of $O(\alpha\frac{\log^3 n}{\log\log n})$ holds for the competitive ratio of static routing.

2 Preliminaries

In this section, we give more formal definitions of the *robust network design* problems considered in this paper. Some notation and basic results are also recalled. The *congestion minimization* variant takes as input a graph $G = (V, E)$, a vector of link capacities $c \in \mathbb{R}_+^E$ and a set of commodities $\mathcal{H}$. Each commodity $h \in \mathcal{H}$ has a source $s(h)$ and destination $t(h)$ in V. We also have as input a polytope $\mathcal{D}$ of all possible demand vectors $d \in \mathbb{R}_+^{\mathcal{H}}$ specifying the demand d_h that needs to be sent from $s(h)$ to $t(h)$. An instance $\mathcal{I}$ of the congestion minimization problem might be denoted by $\mathcal{I} = (G, c, \mathcal{H}, \mathcal{D})$. We use n to denote the number of nodes ($n = |V|$), while k denotes the number of commodities ($k = |\mathcal{H}|$). Given two nodes $s, t \in V$, a routing template (also called a unit flow) from s to t is a vector $f \in \mathbb{R}_+^E$ satisfying the standard flow conservation constraints. For each vertex v, $\sum\limits_{e\in\delta^+(v)} f_e - \sum\limits_{e\in\delta^-(v)} f_e$ is required to be equal to $1, -1$ or 0 when v is respectively the source s, the destination t or any other node, where $\delta^+(v)$ (resp. $\delta^-(v)$) denotes the set of edges going out of (resp. entering into) v.

A vector $f \in \mathbb{R}_+^{\mathcal{H}\times E}$ is a routing if for each commodity $h \in \mathcal{H}$, $f_{h,.} = (f_{h,e})_{e\in E}$ is a routing template from $s(h)$ to $t(h)$. The set of all possible routing schemes is denoted by $\mathcal{F} \subseteq \mathbb{R}_+^{\mathcal{H}\times E}$. The total flow on each link $e \in E$ is $\sum\limits_{h\in\mathcal{H}} f_{h,e}d_h$ and its congestion is the total flow on the link e divided by its capacity c_e. Let $\text{cong}(f,d)$ denote the maximum congestion over all links $e \in E$, i.e. $\text{cong}(f,d) = \max_{e\in E} \sum_{h\in\mathcal{H}} \frac{f_{h,e}d_h}{c_e}$. Two problems can be considered depending on whether the routing can be adapted to each demand vector d in $\mathcal{D}$ or if only one fixed routing $f \in \mathcal{F}$ can be used. In the first case, the routing is said to be *dynamic*. The dynamic congestion is formally defined as: $\text{cong}_{\text{dyn}}(\mathcal{I}) = \max_{d\in\mathcal{D}} \min_{f\in\mathcal{F}} \text{cong}(f,d)$. In the second case, the routing is said to be *static* (or oblivious). This static congestion is formally defined as: $\text{cong}_{\text{sta}}(\mathcal{I}) = \min_{f\in\mathcal{F}} \max_{d\in\mathcal{D}} \text{cong}(f,d)$. Notice that when clear from the context, we might use notation $\text{cong}_{\text{dyn}}(\mathcal{D})$ and $\text{cong}_{\text{sta}}(\mathcal{D})$ to insist on the dependency on $\mathcal{D}$ when all other parameters of the instance $\mathcal{I}$ are fixed.

In the same way, we can define the robust linear reservation problem. As already said in Section 1, given a positive cost vector $(\lambda_e)_{e\in E}$, we aim to reserve a capacity $u_e \geq 0$ on each link e such that $\sum_{e\in E} \lambda_e u_e$ is minimized and $\sum_{h\in\mathcal{H}} f_{h,e}d_h \leq u_e$ holds for any demand vector d. An instance can then be denoted by $(G,\lambda,\mathcal{H},\mathcal{D})$. We also have two variants depending on routing. The optimal cost is then denoted by $\text{lin}_{\text{dyn}}(\mathcal{I})$ (or $\text{lin}_{\text{dyn}}(\mathcal{D})$) and $\text{lin}_{\text{sta}}(\mathcal{I})$ (or $\text{lin}_{\text{sta}}(\mathcal{D})$). Notice that only fractional routing is considered in this paper (this is why the subscript *frac* used in Section 1.2 is omitted in the rest of the paper).

For concise notation, the four variants of the robust optimization problems considered in this paper will simply be denoted by lin_{sta}, lin_{dyn}, cong_{sta} and cong_{dyn}.

All previous definitions still make sense even when $\mathcal{D}$ is not a polytope. However, the next lemma tells us that the optimal objective value does not increase when the uncertainty set $\mathcal{S}$ is replaced by its convex-hull (this lemma can be considered as a folklore result that is implicitly used in many robust optimization papers).

▶ **Lemma 1.** *Let $\mathcal{S} \subset \mathbf{R}_+^{\mathcal{H}}$ be a compact set. Then $cong_{sta}(\mathcal{S}) = cong_{sta}(conv(\mathcal{S}))$, $cong_{dyn}(\mathcal{S}) = cong_{dyn}(conv(\mathcal{S}))$, $lin_{sta}(\mathcal{S}) = lin_{sta}(conv(\mathcal{S}))$, and $lin_{dyn}(\mathcal{S}) = lin_{dyn}(conv(\mathcal{S}))$.*

Proof. Since $\mathcal{S} \subseteq \text{conv}(\mathcal{S})$, we have $\text{cong}_{\text{sta}}(\mathcal{S}) \leq \text{cong}_{\text{sta}}(\text{conv}(\mathcal{S}))$ and $\text{cong}_{\text{dyn}}(\mathcal{S}) \leq \text{cong}_{\text{dyn}}(\text{conv}(\mathcal{S}))$. The same holds for the robust linear cost problem. Moreover, given a static routing solution f and the corresponding reservation vector u, we have $\sum_{h\in\mathcal{H}} f_{h,e}d_h \leq u_e$ for any $d \in \mathcal{S}$. Consider any point d' of $\text{conv}(\mathcal{S})$ written as $d' = \sum_{d\in\mathcal{S}} \alpha_d d$ ($\alpha_d \geq 0, \sum_{d\in\mathcal{S}} \alpha_d = 1$). By multiplying the previous inequalities by α_d and summing them all, we get that $\sum_{h\in\mathcal{H}} f_{h,e}d'_h \leq u_e$ implying that f and u are feasible. Therefore, we have $\text{lin}_{\text{sta}}(\mathcal{S}) = \text{lin}_{\text{sta}}(\text{conv}(\mathcal{S}))$. The proof can be easily extended to the dynamic routing version and to the congestion objective function. ◀

Let us now focus on the connection between the congestion problem and the linear cost problem. The first proposition is from [19] and states that if the static to dynamic ratio is less than or equal to α for the congestion problem, then the same applies to the robust linear reservation problem.

▶ **Proposition 2** ([19])**.** *Let $\mathcal{I} = (G,c,\mathcal{H},\mathcal{D})$ and assume that $cong_{sta}(\mathcal{I}) \leq \alpha\ cong_{dyn}(\mathcal{I})$ for some $\alpha \geq 1$ and for **any** vector $c \in \mathbb{R}_+^{\mathcal{H}}$. Then $lin_{sta}(\mathcal{I}') \leq \alpha\ lin_{dyn}(\mathcal{I}')$ where $\mathcal{I}' = (G,\lambda,\mathcal{H},\mathcal{D})$ for **any** cost vector $\lambda \in \mathbb{R}_+^{\mathcal{H}}$.*

Proof. Given a cost vector λ, let $c^*_{\text{dyn}} \in \mathbb{R}_+^E$ be the reservation vector (i.e., u) obtained when the linear cost is minimized under dynamic routing. Let then $\mathcal{I} = (G, c^*_{\text{dyn}}, \mathcal{H}, \mathcal{D})$. We clearly have $\text{cong}_{\text{dyn}}(\mathcal{I}) \leq 1$ and $\text{cong}_{\text{sta}}(\mathcal{I}) \leq \alpha\ \text{cong}_{\text{dyn}}(\mathcal{I}) \leq \alpha$. Therefore, $\alpha\ c^*_{\text{dyn}}$ is a feasible reservation vector for the lin_{sta} problem related to instance $\mathcal{I}' = (G,\lambda,\mathcal{H},\mathcal{D})$ and its cost is α times the cost of c^*_{dyn}. ◀

A converse result is presented in [1]. While the proof in [1] was given in the context of undirected graphs, it can be repeated verbatim for the directed case (the proof is based on a Lagrange relaxation approach and a careful application of the ellipsoid method).

▶ **Proposition 3** ([1]). *Let $\mathcal{I}' = (G, \lambda, \mathcal{H}, \mathcal{D})$ and assume that $lin_{sta}(\mathcal{I}') \le \alpha\ lin_{dyn}(\mathcal{I}')$ for some $\alpha \ge 1$ and for* ***any*** *cost vector $\lambda \in \mathbb{R}^{\mathcal{H}}_+$. Then $cong_{sta}(\mathcal{I}) \le \alpha\ cong_{dyn}(\mathcal{I})$ where $\mathcal{I} = (G, c, \mathcal{H}, \mathcal{D})$ for* ***any*** *capacity vector $c \in \mathbb{R}^{\mathcal{H}}_+$. Moreover, any β-approximation ($\beta \ge 1$) for lin_{dyn} leads to a β-approximation for $cong_{dyn}$.*

To close this section, let us recall some notation and assumptions that will be used in the rest of the paper. The uncertainty set (i.e., the set of demand vectors) $\mathcal{D}$ is assumed to be polyhedral and down monotone (i.e., if $d \in \mathcal{D}$, then $d' \in \mathcal{D}$ for any $0 \le d' \le d$). Let $d^{max}(\mathcal{D})$ be the vector representing the maximum commodity values (i.e., $d^{max}_h(\mathcal{D}) = \max_{d \in \mathcal{D}} d_h$). We will naturally assume that $d^{max}_h > 0$ for any $h \in \mathcal{H}$ since otherwise the commodity can just be ignored. When the polytope $\mathcal{D}$ is clear from the context, we just write d^{max} (instead of $d^{max}(\mathcal{D})$).

Let I, J be some set of indices. For a vector $v \in \mathbb{R}^{I \times J}$ and $i \in I$ we denote by $v_{i,\cdot}$ the vector $w \in \mathbb{R}^J$ defined by $w_j = v_{i,j}$. Given a set $X \in \mathbb{R}^I$ and $\lambda \ge 0$, we denote by λX the set $\{\lambda x | x \in X\}$.

3 Approximation of dynamic congestion by static congestion

We are going to prove Theorem 4 stating that compared to dynamic routing, when static routing is considered, congestion is multiplied by a factor less than or equal to $\sqrt{8k}$. This result improves the upper bound $O(\sqrt{k} n^{1/4} \log n)$ from [24]. It implies that the gap between static and dynamic congestion is $O(\sqrt{k}) = O(n)$. By combining Proposition 2 with Theorem 4, we also obtain similar results for the minimization of a linear reservation cost, i.e., that $\mathrm{lin}_{\mathrm{sta}}(\mathcal{D}) \le \sqrt{8k}.\mathrm{lin}_{\mathrm{dyn}}(\mathcal{D})$.

▶ **Theorem 4.** *$cong_{sta}(\mathcal{D}) \le \sqrt{8k}.cong_{dyn}(\mathcal{D})$. Therefore $\frac{cong_{sta}(\mathcal{D})}{cong_{dyn}(\mathcal{D})} = O(n)$.*

To derive an upper bound for the ratio $\mathrm{cong}_{\mathrm{sta}}(\mathcal{D})/\mathrm{cong}_{\mathrm{dyn}}(\mathcal{D})$, our strategy first consists in approximating the uncertainty set either by a box or a simplex where $\mathrm{cong}_{\mathrm{sta}}(\mathcal{D}) = \mathrm{cong}_{\mathrm{dyn}}(\mathcal{D})$. While this method yields an $O(k)$ upper bound, we obtain further improvement by partitioning the set of commodities into two sets $\mathcal{H}_1$, $\mathcal{H}_2$ and considering a box approximation for $\mathcal{D}_1$ and a simplex approximation for $\mathcal{D}_2$, where $\mathcal{D}_1$ and $\mathcal{D}_2$ are respectively the projections of $\mathcal{D}$ on $\mathbb{R}^{\mathcal{H}_1}$ and $\mathbb{R}^{\mathcal{H}_2}$.

To prove Theorem 4, we first present some preliminary lemmas.

Lemma 5 states that if the uncertainty set $\mathcal{D}$ can be well approximated by another set $\mathcal{D}'$ for which $\mathrm{cong}_{\mathrm{sta}}(\mathcal{D}') = \mathrm{cong}_{\mathrm{dyn}}(\mathcal{D}')$, then $\mathrm{cong}_{\mathrm{sta}}(\mathcal{D})$ gives a good approximation of $\mathrm{cong}_{\mathrm{dyn}}(\mathcal{D})$.

▶ **Lemma 5.** *Let $\mathcal{D}$ and $\mathcal{D}'$ be two compact subsets of $\mathbb{R}^{\mathcal{H}}_+$ and $\alpha \in \mathbb{R}_+$ such that $\mathcal{D}' \subseteq \mathcal{D} \subseteq \alpha\mathcal{D}'$ and $cong_{sta}(\mathcal{D}') = cong_{dyn}(\mathcal{D}')$. Then $cong_{sta}(\mathcal{D}) \le \alpha \cdot cong_{dyn}(\mathcal{D})$.*

Proof. The proof of this lemma relies on two simple facts. The first one is that if we scale the demand values by a factor α, then the congestion (either static or dynamic) is also scaled by the same factor α. The second fact is that $\mathrm{cong}_{\mathrm{dyn}}$ and $\mathrm{cong}_{\mathrm{sta}}$ are increasing in $\mathcal{D}$. In other words, if $\mathcal{D}_1$ and $\mathcal{D}_2$ are two subsets of $\mathbb{R}^{\mathcal{H}}_+$ such that $\mathcal{D}_1 \subseteq \mathcal{D}_2$, then

$\text{cong}_{\text{dyn}}(\mathcal{D}_1) \leq \text{cong}_{\text{dyn}}(\mathcal{D}_2)$ and $\text{cong}_{\text{sta}}(\mathcal{D}_1) \leq \text{cong}_{\text{sta}}(\mathcal{D}_2)$. Combining the two facts, we can write the following:

$$\text{cong}_{\text{sta}}(\mathcal{D}) \leq \text{cong}_{\text{sta}}(\alpha\mathcal{D}') = \alpha \cdot \text{cong}_{\text{sta}}(\mathcal{D}') = \alpha \cdot \text{cong}_{\text{dyn}}(\mathcal{D}') \leq \alpha \cdot \text{cong}_{\text{dyn}}(\mathcal{D}) \tag{1}$$

◀

We now provide in Lemmas 6 and 7 two classes of polytopes, based on box and simplex sets, for which $\text{cong}_{\text{sta}}(\mathcal{D}) = \text{cong}_{\text{dyn}}(\mathcal{D})$.

For a vector $d^{max} \in \mathbb{R}_+^{\mathcal{H}}$, let $\mathcal{B}(d^{max})$ be the box set defined by $\{d \in \mathbb{R}^{\mathcal{H}} |\ 0 \leq d \leq d^{max}\}$.

▶ **Lemma 6.** *Let $\mathcal{D} = \mathcal{B}(d^{max})$ for some $d^{max} \in \mathbb{R}_+^{\mathcal{H}}$. Then $cong_{dyn}(\mathcal{D}) = cong_{sta}(\mathcal{D})$.*

Proof. For a routing $f \in \mathcal{F}$ and a demand vector $d \in \mathcal{D}$, we have $\text{cong}(f, d) \leq \text{cong}(f, d^{max})$. Since $d^{max} \in \mathcal{D}$, it implies that $\max\limits_{d\in\mathcal{D}} \text{cong}(f, d) = \text{cong}(f, d^{max})$. Minimizing both sides of the equality over $f \in \mathcal{F}$, we get that $\text{cong}_{\text{sta}}(\mathcal{D}) = \min\limits_{f\in\mathcal{F}} \text{cong}(f, d^{max})$. We can also write that $\min\limits_{f\in\mathcal{F}} \text{cong}(f, d) \leq \min\limits_{f\in\mathcal{F}} \text{cong}(f, d^{max})$. Taking the maximum over all $d \in \mathcal{D}$ leads to $\text{cong}_{\text{dyn}}(\mathcal{D}) = \max\limits_{d\in\mathcal{D}} \min\limits_{f\in\mathcal{F}} \text{cong}(f, d) \leq \min\limits_{f\in\mathcal{F}} \text{cong}(f, d^{max})$. Since $d^{max} \in \mathcal{D}$, the previous inequality becomes $\text{cong}_{\text{dyn}}(\mathcal{D}) = \min\limits_{f\in\mathcal{F}} \text{cong}(f, d^{max})$. ◀

For a vector $d \in \mathbb{R}_+^{\mathcal{H}}$, let $\Delta(d)$ be the simplex set whose vertices are the zero vector and the k vectors $d_h e_h$ where e_h denotes the vector in $\mathbb{R}_+^{\mathcal{H}}$ with a component of 1 for commodity h and 0 otherwise. Formally, we have $\Delta(d) = \text{conv}\left(\{d_h e_h | h \in \mathcal{D}\} \cup \{0\}\right)$.

▶ **Lemma 7.** *Let $\mathcal{D} = \Delta(d^{max})$ where $d^{max} \in \mathbb{R}_+^{\mathcal{H}}$. Then $cong_{dyn}(\mathcal{D}) = cong_{sta}(\mathcal{D})$.*

Proof. Assume that $\text{cong}_{\text{dyn}}(\mathcal{D})$ has been computed and consider the obtained dynamic routing. The extreme points of $\mathcal{D}$ are the demand vectors $\{d_h^{max} e_h | h \in \mathcal{H}\} \cup \{0\}$. For each demand vector $d_h^{max} e_h$, we consider the flow $f_{h,.}$ representing its routing. Let us build a static routing f just by routing each commodity h in accordance to $f_{h,.}$. By construction, taking the extreme points of $\mathcal{D}$, we have $\text{cong}_{\text{sta}}(\{d_h^{max} e_h | h \in \mathcal{H}\} \cup \{0\}) = \text{cong}_{\text{dyn}}(\{d_h^{max} e_h | h \in \mathcal{H}\} \cup \{0\})$. By considering the convex-hulls and applying Lemma 1, we get that $\text{cong}_{\text{dyn}}(\mathcal{D}) = \text{cong}_{\text{sta}}(\mathcal{D})$. ◀

Let $\alpha_1(\mathcal{D}) = \max\limits_{d\in\mathcal{D}} \sum\limits_{h\in\mathcal{H}} \frac{d_h}{d_h^{max}}$ (remember that $d_h^{max} = \max_{d\in\mathcal{D}} d_h$). It is then clear that $\Delta(d^{max}) \subseteq \mathcal{D} \subseteq \alpha_1(\mathcal{D})\Delta(d^{max})$. Consider the box $\mathcal{B}(d^{max})$ and let $\alpha_2(\mathcal{D})$ be the smallest factor α such that d^{max}/α belongs to $\mathcal{D}$. In other words, $\alpha_2(\mathcal{D})$ represents the best approximation ratio that can be obtained through boxes. We obviously have $\frac{1}{\alpha_2(\mathcal{D})}\mathcal{B}(d^{max}) \subseteq \mathcal{D} \subseteq \mathcal{B}(d^{max})$. Figure 1 illustrates the approximations by boxes and simplices for a 2-dimensional demand polytope $\mathcal{D}$.

Since $\frac{1}{k}\mathcal{B}(d^{max}) \subseteq \Delta(d^{max}) \subseteq \mathcal{D} \subseteq \mathcal{B}(d^{max})$, $\alpha_2(\mathcal{D})$ is always less than or equal to k. And by definition, $\alpha_1(\mathcal{D})$ is also less than or equal to k.

It is easy to check that the upper bound k is reached since $\alpha_1(\mathcal{B}(d^{max})) = k$ and $\alpha_2(\Delta(d^{max})) = k$. In other words, using box and simplex approximations with the approach above, we cannot expect to prove a better upper bound for the ratio $\text{cong}_{\text{sta}}(\mathcal{D})/\text{cong}_{\text{dyn}}(\mathcal{D})$ for arbitrary uncertainty sets.

A more refined strategy is to take the best of the two bounds $\alpha_1(\mathcal{D})$, $\alpha_2(\mathcal{D})$. The next proposition states that a better bound is obtained if $\mathcal{D}$ is permutation-invariant (i.e., by permuting the components of any vector d of $\mathcal{D}$ we always get a vector inside $\mathcal{D}$). The proof is provided in Appendix.

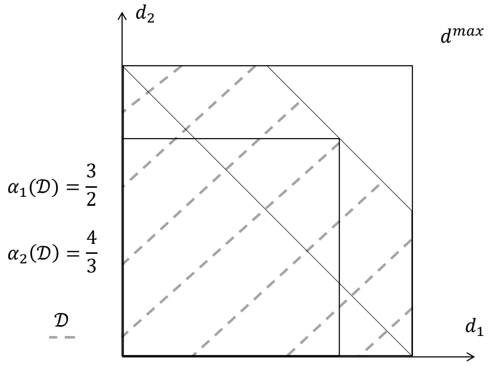

Figure 1 Approximation using boxes and simplices: example of a 2-dimensional demand polytope $\mathcal{D}$.

▶ **Proposition 8.** *If $\mathcal{D}$ is permutation-invariant then* $\min\{\alpha_1(\mathcal{D}), \alpha_2(\mathcal{D})\} \leq \sqrt{k}$.

One can wonder whether a general $O(\sqrt{k})$ bound can be obtained by trying to find a better upper bound for $\min\{\alpha_1(\mathcal{D}), \alpha_2(\mathcal{D})\}$. The following example, on a specific polytope $\mathcal{D}$, shows that this is not possible. Let $\mathcal{D}$ be the product of a box $\mathcal{B}(d^1)$ of dimension $k/2$ and a simplex $\Delta(d^2)$ of the same dimension. Using the remark above we know that $\alpha_1(\mathcal{B}(d^1)) = k/2$ and $\alpha_2(\Delta(d^2)) = k/2$ implying that $\alpha_1(\mathcal{D}) \geq k/2$ and $\alpha_2(\mathcal{D}) \geq k/2$.

To overcome this difficulty, we are going to partition the set of commodities $\mathcal{H}$ into two well-chosen subsets $\mathcal{H}_1$ and $\mathcal{H}_2$, then we approximate $\mathcal{D}_1$ (resp. $\mathcal{D}_2$) defined as the projection of $\mathcal{D}$ on $\mathbb{R}^{\mathcal{H}_1}$ (resp. $\mathbb{R}^{\mathcal{H}_2}$) using a simplex (resp. a box). The algorithm used to partition the set of commodities is an adaptation of an algorithm of [10] proposed in a different context. We will also slightly improve the analysis of this algorithm ($\sqrt{8k}$ instead of $3\sqrt{k}$).

Let us start with Lemma 9 where we show how an approximation of cong_{dyn} in $\mathcal{D}$ can be obtained from cong_{sta} using the approximations related to $\mathcal{D}_1$ and $\mathcal{D}_2$.

▶ **Lemma 9.** *Let $\mathcal{H}_1, \mathcal{H}_2$ be a partition of $\mathcal{H}$ and $\mathcal{D}_1, \mathcal{D}_2$ be the projection of $\mathcal{D}$ on $\mathbb{R}^{\mathcal{H}_1}$ and $\mathbb{R}^{\mathcal{H}_2}$. Suppose that for some $\alpha_1, \alpha_2 \geq 1$ we have $cong_{sta}(\mathcal{D}_1) \leq \alpha_1 cong_{dyn}(\mathcal{D}_1)$ and $cong_{sta}(\mathcal{D}_2) \leq \alpha_2 cong_{dyn}(\mathcal{D}_2)$, then $cong_{sta}(\mathcal{D}) \leq (\alpha_1 + \alpha_2) cong_{dyn}(\mathcal{D})$.*

Proof. We first show that we have $\text{cong}_{\text{sta}}(\mathcal{D}) \leq \text{cong}_{\text{sta}}(\mathcal{D}_1) + \text{cong}_{\text{sta}}(\mathcal{D}_2)$.

$$\begin{aligned}
\text{cong}_{\text{sta}}(\mathcal{D}) &= \max_{d \in \mathcal{D}} \min_{f \in \mathcal{F}} \text{cong}(f, d) \\
&\leq \max_{d^1 \in \mathcal{D}_1, d^2 \in \mathcal{D}_2} \min_{f \in \mathcal{F}} \text{cong}(f, d^1) + \text{cong}(f, d^2) \\
&= \max_{d^1 \in \mathcal{D}_1} \min_{f \in \mathcal{F}} \text{cong}(f, d^1) + \max_{d^2 \in \mathcal{D}_2} \min_{f \in \mathcal{F}} \text{cong}(f, d^2) \\
&= \text{cong}_{\text{sta}}(\mathcal{D}_1) + \text{cong}_{\text{sta}}(\mathcal{D}_2)
\end{aligned}$$

We now prove the lemma: $\text{cong}_{\text{sta}}(\mathcal{D}) \leq \text{cong}_{\text{sta}}(\mathcal{D}_1) + \text{cong}_{\text{sta}}(\mathcal{D}_2) \leq \alpha_1 \text{cong}_{\text{dyn}}(\mathcal{D}_1) + \alpha_2 \text{cong}_{\text{dyn}}(\mathcal{D}_2) \leq (\alpha_1 + \alpha_2) \text{cong}_{\text{dyn}}(\mathcal{D})$. ◀

Let us now present Algorithm 1 that can be seen as a direct adaptation of the partitioning algorithm of [10] (Algorithm $\mathcal{A}$, Fig. 1) for our dynamic routing problem. It has initially been introduced for the analysis of affine policies in a class of two-stage adaptive linear optimization problems. The main idea of Algorithm 1 is to partition the set of commodities into two sets $\mathcal{H}_1$ and $\mathcal{H}_2$ and to produce a vector $\beta \in \mathbb{R}_+^{\mathcal{H}}$ such that $\max_{d \in \mathcal{D}} \sum_{h \in \mathcal{H}_1} \frac{d_h}{d_h^{max}} \leq \gamma \sqrt{k}$

(i.e., $\alpha_1(\mathcal{D}_1) \leq \gamma\sqrt{k}$ for $\gamma > 0$) and $\beta_h \geq d_h^{max}$ for any $h \in \mathcal{H}_2$. The returned vector β is built as a sum of at most Z points of $\mathcal{D}$ where Z is the number of iterations of the algorithm. Since the vector $\frac{1}{Z}\beta$ belongs to $\mathcal{D}$, we deduce that $\alpha_2(\mathcal{D}_2) \leq Z$. We will show in Lemma 10 that Z is less than or equal to $2\frac{\sqrt{k}}{\gamma}$ leading to $\alpha_2(\mathcal{D}_2) \leq 2\frac{\sqrt{k}}{\gamma}$. Notice that γ is equal to 1 in the original algorithm of [10]. Let us describe more precisely the different steps of Algorithm 1. At iteration i, $\mathcal{H}_1^i, \mathcal{H}_2^i$ denote the current partitions of commodities while $\mathcal{D}_1^i, \mathcal{D}_2^i$ denote the projections of $\mathcal{D}$ on $\mathbb{R}^{\mathcal{H}_1^i}$ and $\mathbb{R}^{\mathcal{H}_2^i}$. A vector b^i is also defined and used to update $\mathcal{H}_1^i, \mathcal{H}_2^i$. We start with $\mathcal{H}_1^0 = \mathcal{H}$, $\mathcal{H}_2^0 = \emptyset$ and $b^0 = 0$.

If $\alpha_1(\mathcal{D}_1^i) > \gamma\sqrt{k}$ then we consider a traffic vector u^i maximizing $\sum_{h \in \mathcal{H}_1^{i-1}} \frac{d_h}{d_h^{max}}$, otherwise a partition is returned. The vector u^i is then used to update b^i (lines 5-7). Observe that only the components related to commodities inside $\mathcal{H}_1^{i-1}$ are updated while the others do not change. This means that the returned vector $\beta = \sum_{1 \leq i \leq Z} u^i$ (line 19) is such that $\beta \geq b^Z$. The sets $\mathcal{H}_1^i$ and $\mathcal{H}_2^i$ are updated by moving each commodity $h \in \mathcal{H}_1^{i-1}$ to $\mathcal{H}_2^i$ if $b_h^i \geq d_h^{max}$ (lines 8-15). Notice that we always have $\mathcal{H}_1^i \subseteq \mathcal{H}_1^{i-1}$. It is then clear that when the algorithm stops, the obtained partition satisfies what is announced above. The only fact that remains to be proved is that the number of iterations Z is bounded by $2\frac{\sqrt{k}}{\gamma}$.

Algorithm 1 Commodity partitioning algorithm (adapted from [10]).

1: Initialize $i \leftarrow 0$, $\mathcal{H}_1^0 \leftarrow \mathcal{H}, \mathcal{H}_2^0 \leftarrow \emptyset$, $b^0 \leftarrow 0$
2: **while** $\alpha_1(\mathcal{D}_1^i) > \gamma\sqrt{k}$ **do**
3: $\quad i \leftarrow i+1$
4: $\quad u^i \in \arg\max_{d \in \mathcal{D}} \sum_{h \in \mathcal{H}_1^{i-1}} \frac{d_h}{d_h^{max}}$
5: $\quad$ **for all** $h \in \mathcal{H}$ **do**
6: $\qquad b_h^i = \begin{cases} b_h^{i-1} + u_h^i & \text{if } h \in \mathcal{H}_1^{i-1} \\ b_h^{i-1} & \text{otherwise} \end{cases}$
7: $\quad$ **end for**
8: $\quad$ **for all** $h \in \mathcal{H}_1^{i-1}$ **do**
9: $\qquad$ **if** $b_h^i \geq d_h^{max}$ **then**
10: $\qquad\quad \mathcal{H}_1^i \leftarrow \mathcal{H}_1^{i-1} \backslash \{h\}$
11: $\qquad\quad \mathcal{H}_2^i \leftarrow \mathcal{H}_2^{i-1} \cup \{h\}$
12: $\qquad$ **else**
13: $\qquad\quad \mathcal{H}_1^i \leftarrow \mathcal{H}_1^{i-1}$, $\mathcal{H}_2^i \leftarrow \mathcal{H}_2^{i-1}$
14: $\qquad$ **end if**
15: $\quad$ **end for**
16: **end while**
17: $Z \leftarrow i$, $\mathcal{H}_1 \leftarrow \mathcal{H}_1^Z, \mathcal{H}_2 \leftarrow \mathcal{H}_2^Z$
18: $\beta \leftarrow \sum_{1 \leq i \leq Z} u^i$

▶ **Lemma 10.** *For any $\gamma > 0$, the commodity set $\mathcal{H}$ can be partitioned in two subsets $\mathcal{H}_1, \mathcal{H}_2$ such that $\alpha_1(\mathcal{D}_1) \leq \gamma\sqrt{k}$ and $\alpha_2(\mathcal{D}_2) \leq \frac{2\sqrt{k}}{\gamma}$ where $\mathcal{D}_1, \mathcal{D}_2$ are the projections of $\mathcal{D}$ on $\mathbb{R}^{\mathcal{H}_1}$ and $\mathbb{R}^{\mathcal{H}_2}$.*

Proof. We only have to prove that $Z \leq \frac{2\sqrt{k}}{\gamma}$. This can be done by slightly modifying the proof of Lemma 10 of [10].

We first argue that $b_h^Z \leq 2d_h^{max}$ for all $h \in \mathcal{H}$. For $h \in \mathcal{H}$, let $i(h)$ be the last iteration number when $h \in \mathcal{H}_1^i$. Therefore we have $b_h^{i(h)-1} \leq d_h^{max}$. Also $u^{i(h)} \leq d_h^{max}$ leading to $b_h^{i(h)} \leq 2d_h^{max}$. Now for $i \geq i(h)$ we have $b_h^Z = b_h^i = b_h^{i(h)}$ implying that, $\sum_{h \in \mathcal{H}} \frac{b_h^Z}{d_h^{max}} \leq 2k$.

Alternatively, $\sum_{h\in\mathcal{H}} \frac{b_h^Z}{d_h^{max}} = \sum_{h\in\mathcal{H}} \sum_{i=1}^{Z} \frac{b_h^i - b_h^{i-1}}{d_h^{max}} = \sum_{i=1}^{Z} \sum_{h\in\mathcal{H}} \frac{b_h^i - b_h^{i-1}}{d_h^{max}} = \sum_{i=1}^{Z} \sum_{h\in\mathcal{H}_1^{i-1}} \frac{u_h^i}{d_h^{max}} \geq \sum_{i=1}^{Z} \gamma\sqrt{k} = Z\gamma\sqrt{k}$. Therefore we have that $Z\gamma\sqrt{k} \leq \sum_{h\in\mathcal{H}} \frac{b_h^Z}{d_h^{max}} \leq 2k$ which implies that $Z \leq \frac{2\sqrt{k}}{\gamma}$. Since β is the sum of Z points in $\mathcal{D}$, we have $\mathcal{B}(\frac{1}{Z}\beta) \subseteq \mathcal{D}$. Moreover, the projection β^2 of β on $\mathbb{R}^{\mathcal{H}_2}$ satisfies $\mathcal{D}_2 \subseteq \mathcal{B}(\beta^2)$ and thus $\alpha_2(\mathcal{D}_2) \leq \frac{2\sqrt{k}}{\gamma}$. ◀

To prove Theorem 4, we only have to take $\gamma = \sqrt{2}$, use Lemma 10, and then invoke Lemma 9 to conclude. Using $k = O(n^2)$, we get that the ratio $\frac{\text{cong}_{\text{sta}}}{\text{cong}_{\text{dyn}}}$ is $O(n)$.

4 Inapproximability with flow restrictions

Let us consider a more general variant of the robust congestion problem where each commodity can only be routed on a subset of allowed edges $E_h \subseteq E$. These restrictions seem to be quite natural to ensure quality of service requirements such as delay constraints.

Observe first that cong_{sta} can still be computed in polynomial-time for this variant. Moreover, the $O(\sqrt{k})$ bound of Section 3 still holds here since all the proofs presented there do not change if we assume that each commodity h can only be routed using edges inside E_h.

The $\Omega(\frac{\log n}{\log\log n})$ inapproximability bound shown for the undirected case [1] (under ETH assumption) still applies to the directed case (with and without flow restrictions). It is however quite far from the $O(\sqrt{k})$ approximation ratio deduced from Section 3. We will prove stronger inapproximability results for the generalisation of cong_{dyn} with flow restrictions under some classical complexity conjectures.

A standard way to prove this kind of results is to first prove that the problem is inapproximable under some constant and then to amplify this constant, see for example [26].

Let us first introduce some additional notations. Taking into account the flow restrictions and given a subset of edges $C \subseteq E$, let $\mathcal{H}_C \subseteq \mathcal{H}$ be the set commodities such that each valid path related to any commodity $h \in \mathcal{H}_C$ intersects C. Even if C is not necessarily a cut in the standard sense of graph theory, C is called a cut in what follows. Given a demand vector $d \in \mathcal{D}$ and a cut C, $\sum_{h\in\mathcal{H}_C} d_h / \sum_{e\in C} c_e$, is obviously a lower bound of $\text{cong}_{\text{dyn}}(\mathcal{D})$. The maximum over all demand vectors $d \in \mathcal{D}$ and all cuts C of the ratio $\sum_{h\in\mathcal{H}_C} d_h / \sum_{e\in C} c_e$ is called cut congestion and denoted by $\text{cong}_{\text{cut}}(\mathcal{D})$. We also use $E_{\mathcal{H}}$ to denote the set of all flow restrictions: $E_{\mathcal{H}} = (E_h)_{h\in\mathcal{H}}$. An instance of cong_{dyn} with flow restrictions is then defined by $(G, c, \mathcal{H}, \mathcal{D}, E_{\mathcal{H}})$.

In Lemma 11, we will prove that it is *NP*-hard to distinguish between instances where $\text{cong}_{\text{dyn}}(\mathcal{D})$ is less than or equal to 1 and those where the cut congestion $\text{cong}_{\text{cut}}(\mathcal{D})$ is greater than or equal to $1+\rho$ for some constant $\rho > 0$. Then, in Lemma 12, we will show that given two instances of this problem, it is possible to build some kind of product instance whose dynamic congestion is less than or equal to the product of the dynamic congestion of the two instances and the cut congestion is greater than or equal to the product of the cut congestion of the two initial instances. Finally, by repetitively using the product of Lemma 12 on the instance of Lemma 11, we can amplify the gap leading to some strong inapproximability results.

Given a 3-SAT instance φ, $\text{val}(\varphi)$ denotes the maximum proportion of clauses that can be simultaneously satisfied (thus φ is satisfiable when $\text{val}(\varphi) = 1$). We will consider polytopes $\mathcal{D}$ that can be expressed through linear constraints and auxiliary variables ξ, i.e., $\mathcal{D} = \{d \in \mathbb{R}^{\mathcal{H}} | Ad + B\xi \leq b\}$ where A and B are matrices of polynomial size (the maximum of the number of columns and the number of rows is polynomially bounded). Notice that it is important to consider polytopes that can be easily described (otherwise the difficulty of solving $\mathrm{cong}_{\mathrm{dyn}}$ would be a consequence of the difficulty of describing the polytope).

▶ **Lemma 11.** *For $0 < \rho < 1$, there is a polynomial-time mapping from a 3-SAT instance φ to an instance $\mathcal{I} = (G, c, \mathcal{H}, \mathcal{D})$ of $cong_{dyn}$ where $\mathcal{D} = \{d \in \mathbb{R}^{\mathcal{H}} | Ad + B\xi \le b\}$ such that:*
- *If $val(\varphi) \le 1 - \rho$ then $cong_{dyn}(\mathcal{I}) \le 1$*
- *If φ is satisfiable then $cong_{cut}(\mathcal{I}) \ge 1 + \rho$.*

Furthermore, $|V(G)|$, $|E(G)|$, $|\mathcal{H}|$ and the size of the matrices A and B defining $\mathcal{D}$ are all $O(m)$ where m is the number of clauses of φ.

Proof. Given a 3-SAT instance φ with m clauses, we build an instance of $\mathrm{cong}_{\mathrm{dyn}}$ as follows. We consider two nodes: a source s and a destination t. Then, for each $i = 1, ..., m$ we create a path form s to t containing three directed edges $e_{i,j}$ of capacity 1 for $j = 1, 2, 3$ corresponding to the $i - th$ clause of φ. For each $i = 1, ..., m$ and $j = 1, 2, 3$, $\mathcal{H}$ contains a commodity $h_{i,j}$ with the same source and destination as edge $e_{i,j}$. We also add a commodity $h_{s,t}$ from s to t. The polytope $\mathcal{D}$ is defined as follows. We set $d_{h_{s,t}} = \rho \cdot m$. For each literal l (i.e. a variable or its negation) of the 3-SAT instance φ we add an auxiliary variable ξ_l. Intuitively $\xi_l = 1$ will correspond to setting the literal l to true. For each variable v, we add the constraint $\xi_v + \xi_{\neg v} = 1$ in addition to non-negativity constraints $\xi_v \ge 0$ and $\xi_{\neg v} \ge 0$. For each $i = 1, ..., m$ and $j = 1, 2, 3$, we consider the constraint $d_{h_{i,j}} = \xi_{l_{i,j}}$ where $l_{i,j}$ is the literal appearing in the $i - th$ clause in the $j - th$ position. Observe that the size of $\mathcal{D}$ is $O(m)$. The numbers of nodes, edges and commodities are also $O(m)$.

Consider first the case $\mathrm{val}(\varphi) \le 1 - \rho$. The set of extreme points of $\mathcal{D}$ is such that the ξ_l variables take their values in $\{0, 1\}$. The maximum dynamic congestion is attained for a demand vector of this form (see Lemma 1). Let d be such a demand vector and consider the corresponding solution of the 3-SAT instance φ. Notice that demand $d_{h_{i,j}}$ can only be routed on $e_{i,j}$. If for some $i = 1, ..., m$ the $i - th$ clause is false, then the demands $d_{h_{i,1}}, d_{h_{i,2}}, d_{h_{i,3}}$ are equal to 0 and therefore one unit of flow of the commodity $h_{s,t}$ can be routed on the path $(e_{i,1}, e_{i,2}, e_{i,3})$. Since $\mathrm{val}(\varphi) \le 1 - \rho$, there are at least $m \cdot \rho$ such indices i (i.e., false clauses) and therefore the demand $d_{h_{s,t}}$ can be routed with a congestion less than or equal to 1.

We now consider the case where φ is satisfiable. Let d be the demand vector corresponding to a truth assignment satisfying φ. For each $i = 1, ..., m$, let $j(i)$ be the position of a literal set to true in the $i - th$ clause. Therefore we have $d_{h_{i,j(i)}} = 1$ for all $i = 1, ..., m$. Consider the cut $C = \{e_{i,j(i)} | i = 1, ..., m\}$. C intersects the paths related to the m demands $d_{h_{i,j(i)}}$ of value 1 in addition to demand $d_{h_{s,t}}$ of value $m \cdot \rho$. The total capacity of this cut is m while the sum of demands belonging to C is $m + m \cdot \rho$. Therefore the congestion of this cut is $\frac{m + m \cdot \rho}{m} = 1 + \rho$. ◀

Next Lemma (whose proof is provided in Appendix) shows how to build a product instance leading to some gap amplification.

▶ **Lemma 12.** *Given two instances of $cong_{dyn}$ with flow restrictions $\mathcal{I}_1 = (G_1, c_1, \mathcal{H}_1, \mathcal{D}_1, E_{\mathcal{H}_1})$ and $\mathcal{I}_2 = (G_2, c_2, \mathcal{H}_2, \mathcal{D}_2, E_{\mathcal{H}_2})$, we can build a new instance $\mathcal{I} = \mathcal{I}_1 \times \mathcal{I}_2 = (G, c, \mathcal{H}, \mathcal{D}, E_{\mathcal{H}})$ such that:*
- *$cong_{dyn}(\mathcal{I}) \le cong_{dyn}(\mathcal{I}_1) \cdot cong_{dyn}(\mathcal{I}_2)$*
- *$cong_{cut}(\mathcal{I}) \ge cong_{cut}(\mathcal{I}_1) \cdot cong_{cut}(\mathcal{I}_2)$.*

Furthermore, we have $|E(G)| = |E(G_1)| \cdot (|E(G_2)| + 2|V(G_2)|)$, $|V(G)| = |V(G_1)| + |V(G_2)| \cdot |E(G_1)|$, $|\mathcal{H}| = |\mathcal{H}_1| \cdot |\mathcal{H}_2|$ and the size of $\mathcal{D}$ is less than or equal to the product of the sizes of $\mathcal{D}_1$ and $\mathcal{D}_2$.

Combining the two previous lemmas, one can amplify the gap as follows.

▶ **Lemma 13.** *For some $0 < \rho < 1$ and each $p \in \mathbb{N}$, each 3-SAT instance φ can be mapped to an instance $\mathcal{I}^p = (G^p, c^p, \mathcal{H}^p, \mathcal{D}^p, E_{\mathcal{H}^p})$ of $cong_{dyn}$ with flow restrictions where $\mathcal{D}^p = \{d \in \mathbb{R}^{\mathcal{H}^p} | A^p d + B^p \xi \leq b^p\}$ such that:*

- *If $val(\varphi) \leq 1 - \rho$ then $cong_{dyn}(\mathcal{I}^p) \leq 1$*
- *If φ is satisfiable then $cong_{cut}(\mathcal{I}^p) \geq (1 + \rho)^p$.*

Furthermore, there exists a positive constant θ such that $|V(G^p)|$, $|E(G^p)|$, $|\mathcal{H}^p|$ and the size of the matrices A^p and B^p defining $\mathcal{D}^p$ are all less than or equal to $(\theta m)^p$ where m is the number of clauses of φ.

Proof. Let $\mathcal{I}^1$ be the instance defined in Lemma 11. We recursively build $\mathcal{I}^p$ as the product of $\mathcal{I}^{p-1}$ and $\mathcal{I}^1$. Using notation of Lemma 12, we take $\mathcal{I}_1 = \mathcal{I}^{p-1}$, $\mathcal{I}_2 = \mathcal{I}^1$ and $\mathcal{I}^p = \mathcal{I} = \mathcal{I}_1 \times \mathcal{I}_2$. Using what is already known about the size of the instance $\mathcal{I}^1$ of Lemma 11 and the results of Lemma 12, a simple induction proves the existence of a constant θ such that $(\theta m)^p$ is an upper bound of the number of vertices, number of edges, number of commodities and the size of the matrices defining the polytope $\mathcal{D}^p$. ◀

By making use of some standard complexity assumptions, inapproxiambility resuls can be directly deduced from the previous lemma.

▶ **Proposition 14.** *Unless $NP \subseteq SUBEXP$, $cong_{dyn}$ with flow restrictions cannot be approximated within a factor of $k^{\frac{c'}{\log \log k}}$ (resp. $n^{\frac{c'}{\log \log n}}$) for some constant c'.*

Proof. $SUBEXP$ is the class of problems that can be solved in $2^{n^{\epsilon}}$ time for all $\epsilon > 0$. Therefore, if $NP \nsubseteq SUBEXP$ then there is a constant $\epsilon_0 > 0$ such that no algorithm can solve the Gap-3-SAT problem in time $O(2^{m^{\epsilon_0}})$ where m is the number of clauses of the 3-SAT instance.
Let $\epsilon_1 < \epsilon_0$ and let $p(m) = \frac{m^{\epsilon_1}}{\log m}$. The size of the instance $\mathcal{I}^{p(m)}$ is polynomial in $m^{p(m)}$. Therefore if we run a polynomial approximation algorithm on the instance $\mathcal{I}^{p(m)}$, the running time will be $m^{c_1 p(m)}$ for some constant c_1. Furthermore, $m^{c_1 p(m)} = m^{c_1 \frac{m^{\epsilon_1}}{\log m}} = 2^{c_1 m^{\epsilon_1}} < 2^{m^{\epsilon_0}}$ for big enough m.

The number of commodities k in the instance $\mathcal{I}^{p(m)}$ is bounded by $(\theta m)^{p(m)}$. We consequently have $\log k \leq \log(\theta m) \frac{m^{\epsilon_1}}{\log m}$ implying that $m > a \log^{\frac{1}{\epsilon_1}} k$ for some constant a and big enough m.

The gap between the congestion of the instances $\mathcal{I}^{p(m)}$ corresponding to a 3-SAT instance for which $val(\varphi) < 1 - \rho$ and those for which $val(\varphi) = 1$ is:

$$(1+\rho)^{p(m)} > (1+\rho)^{p(a \log^{\frac{1}{\epsilon_1}} k)} = (1+\rho)^{\frac{a^{\epsilon_1} \log k}{\frac{1}{\epsilon_1} \log a \log k}} > k^{\frac{c'}{\log \log k}} \quad \text{for some constant c'.}$$

Hence, if a polynomial-time algorithm could solve cong_{dyn} with flow restrictions within an approximation ratio of $O(k^{\frac{c'}{\log \log k}})$, we could use it to solve the Gap-3-SAT problem in $O(2^{m^{\epsilon_0}})$ time. The same proof applies if parameter n (the number of vertices) is considered instead of k. ◀

A slightly weaker inapproximability result is obtained using a weaker complexity assumption (the proof is provided in Appendix).

▶ **Proposition 15.** *Unless $NP \subseteq QP$, $cong_{dyn}$ with flow restrictions cannot be approximated within a factor of $2^{\log^{1-\epsilon} k}$ (resp. $2^{\log^{1-\epsilon} n}$) for any $\epsilon > 0$.*

Using the last part of Proposition 3, all inapproximability results stated for the congestion problem cong_{dyn} are also valid for lin_{dyn}.

References

1 Y. Al-Najjar, W. Ben-Ameur, and J. Leguay. On the approximability of robust network design. *Theoretical Computer Science*, 860:41–50, 2021.

2 D. Applegate and E. Cohen. Making intra-domain routing robust to changing and uncertain traffic demands: Understanding fundamental tradeoffs. In *Proc. ACM SIGCOMM*, 2003.

3 Y. Azar, E. Cohen, A. Fiat, H. Kaplan, and H. Racke. Optimal oblivious routing in polynomial time. In *Proc. ACM STOC*, 2003.

4 Y. Bartal and S. Leonardi. Ondashline routing in all-optical networks. *Theoretical Computer Science*, 221(1):19–39, 1999.

5 W. Ben-Ameur. Between fully dynamic routing and robust stable routing. In *International Workshop on Design and Reliable Communication Networks*, 2007.

6 W. Ben-Ameur and H. Kerivin. New economical virtual private networks. *Commun. ACM*, 46(6):69–73, June 2003.

7 W. Ben-Ameur and H. Kerivin. Routing of uncertain traffic demands. *Optimization and Engineering*, 6(3):283–313, 2005.

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

9 W. Ben-Ameur and M. Żotkiewicz. Multipolar routing: where dynamic and static routing meet. *Electronic Notes in Discrete Mathematics*, 41:61–68, 2013.

10 D. Bertsimas and V. Goyal. On the power and limitations of affine policies in two-stage adaptive optimization. *Mathematical programming*, 134(2):491–531, 2012.

11 C. Chekuri. Routing and network design with robustness to changing or uncertain traffic demands. *SIGACT News*, 38(3):106–129, September 2007.

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

13 G. Claßen, A. M. C. A. Koster, M. Kutschka, and I. Tahiri. Robust metric inequalities for network loading under demand uncertainty. *APJOR*, 32(5), 2015.

14 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, August 1999.

15 F. Eisenbrand, F. Grandoni, G. Oriolo, and M. Skutella. New approaches for virtual private network design. *SIAM Journal on Computing*, 37(3), 2007.

16 A. Ene, G. Miller, J. Pachocki, and A. Sidford. Routing under balance. In *Proceedings of the Forty-Eighth Annual ACM Symposium on Theory of Computing*, STOC '16, pages 598–611, New York, NY, USA, 2016. Association for Computing Machinery.

17 J.A. Fingerhut, S. Suri, and J. S. Turner. Designing least-cost nonblocking broadband networks. *Journal of Algorithms*, 24(2):287–309, 1997.

18 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.

19 N. Goyal, N. Olver, and F. Shepherd. Dynamic vs. oblivious routing in network design. In *European Symposium on Algorithms*, pages 277–288. Springer, 2009.

20 N. Goyal, N. Olver, and F. B. Shepherd. The vpn conjecture is true. *Journal of the ACM (JACM)*, 60(3):1–17, 2013.

21 A. Gupta, J. Kleinberg, A. Kumar, R. Rastogi, and B. Yener. Provisioning a virtual private network: A network design problem for multicommodity flow. In *Proc. ACM STOC*, 2001.

22 A. Gupta and J. Könemann. Approximation algorithms for network design: A survey. *Surveys in Operations Research and Management Science*, 16(1):3–20, 2011.

23 M. Hajiaghayi, R. Kleinberg, and T. Leighton. Semi-oblivious routing: lower bounds. In *Proceedings of ACM-SIAM symposium on Discrete algorithms*, pages 929–938, 2007.

24 M. T. Hajiaghayi, R. D. Kleinberg, H. Räcke, and T. Leighton. Oblivious routing on node-capacitated and directed graphs. *ACM Trans. Algorithms*, 3(4), November 2007. doi:10.1145/1290672.1290688.

25 C. Harrelson, K. Hildrum, and S. Rao. A polynomial-time tree decomposition to minimize congestion. In *Proc. SPAA*, 2003.

26 I. Haviv and O. Regev. Tensor-based hardness of the shortest vector problem to within almost polynomial factors. In *Proceedings of the thirty-ninth annual ACM symposium on Theory of computing*, pages 469–477, 2007.

27 P. Kumar, Y. Yuan, C. Yu, N. Foster, R. Kleinberg, P. Lapukhov, C. Lin Lim, and R. Soulé. Semi-oblivious traffic engineering: The road not taken. In *15th USENIX Symposium on Networked Systems Design and Implementation (NSDI 18)*, pages 157–170. USENIX Association, 2018.

28 H. P. L. Luna. Network planning problems in telecommunications. In *Handbook of Optimization in Telecommunications*, pages 213–240. Springer, 2006.

29 B. M. Maggs, F. Meyer auf der Heide, B. Vöcking, and M. Westermann. Exploiting locality for data management in systems of limited bandwidth. In *Proc. FOCS*, 1997.

30 S. Mattia. The robust network loading problem with dynamic routing. *Comp. Opt. and Appl.*, 54(3):619–643, 2013.

31 M. Minoux. Robust network optimization under polyhedral demand uncertainty is np-hard. *Discrete Applied Mathematics*, 158(5):597–603, 2010.

32 N. Olver. *Robust network design*. PhD thesis, McGill University Library, 2010.

33 N. Olver and F. B. Shepherd. Approximability of robust network design. *Mathematics of Operations Research*, 39(2):561–572, 2014.

34 A. Ouorou and J.P. Vial. A model for robust capacity planning for telecommunications networks under demand uncertainty. In *Workshop on Design and Reliable Communication Networks*, 2007.

35 M. Poss and C. Raack. Affine recourse for the robust network design problem: Between static and dynamic routing. *Networks*, 61(2):180–198, 2013.

36 H. Räcke. Minimizing congestion in general networks. In *The 43rd Annual IEEE Symposium on Foundations of Computer Science, 2002. Proceedings.*, pages 43–52, 2002.

37 H. Räcke. Optimal hierarchical decompositions for congestion minimization in networks. In *40th annual ACM symposium on Theory of Computing*, pages 255–264, 2008.

38 N. Wang, Kin H. Ho, G. Pavlou, and M. Howarth. An overview of routing optimization for internet traffic engineering. *IEEE Communications Surveys & Tutorials*, 10(1):36–56, 2008.

39 M. Żotkiewicz and W. Ben-Ameur. More adaptive robust stable routing. In *IEEE Global Telecommunications Conference*, 2009.

A Appendix

Proof of Proposition 8. Let d^* be the demand maximizing $\max_{d\in\mathcal{D}} \sum_{h\in\mathcal{H}} \frac{d_h}{d_h^{max}}$. Since $\mathcal{D}$ is permutation-invariant, $d_h^{max} = d_{h'}^{max}$ for all $h, h' \in \mathcal{H}$ and d^* can be chosen such that $d_h^* = d_{h'}^*$ for all $h, h' \in \mathcal{H}$. Consequently, we have $\alpha_1(\mathcal{D}) \leq k\frac{d_{h_0}^*}{d_{h_0}^{max}}$. Moreover, since $\frac{d_{h_0}^*}{d_{h_0}^{max}}\mathcal{B}(d^{max}) \subseteq \mathcal{D} \subseteq \mathcal{B}(d^{max})$ we also have $\alpha_2(\mathcal{D}) \leq \frac{d_{h_0}^{max}}{d_{h_0}^*}$. Therefore, using notation $x = \frac{d_{h_0}^*}{d_{h_0}^{max}}$, we get that $\min\{\alpha_1(\mathcal{D}), \alpha_2(\mathcal{D})\} \leq \min\{kx, \frac{1}{x}\}$ and x is such that $0 \leq x \leq 1$. To conclude, observe that $\max_{0\leq x\leq 1} \min\{kx, \frac{1}{x}\} = \sqrt{k}$. ◀

Proof of Lemma 12. Let $\mathcal{I}_1 = (G_1, c_1, \mathcal{H}_1, \mathcal{D}_1, E_{\mathcal{H}_1})$ and $\mathcal{I}_2 = (G_2, c_2, \mathcal{H}_2, \mathcal{D}_2, E_{\mathcal{H}_2})$ be two instances of cong_{dyn} with flow restrictions. We denote by G_2' the graph obtained from G_2 by adding two nodes $s(G_2)$ and $t(G_2)$ to G_2, an edge from $s(G_2)$ to each node of G_2 having an infinite capacity (i.e., $|V(G_2)|$ edges), and an edge from each node of G_2 to $t(G_2)$ having also an infinite capacity (i.e., $|V(G_2)|$ edges). We build a graph G by replacing each edge e of G_1 by a copy of G_2' while identifying the node $s(e)$ (resp. $t(e)$) with the node $s(G_2)$ (resp.

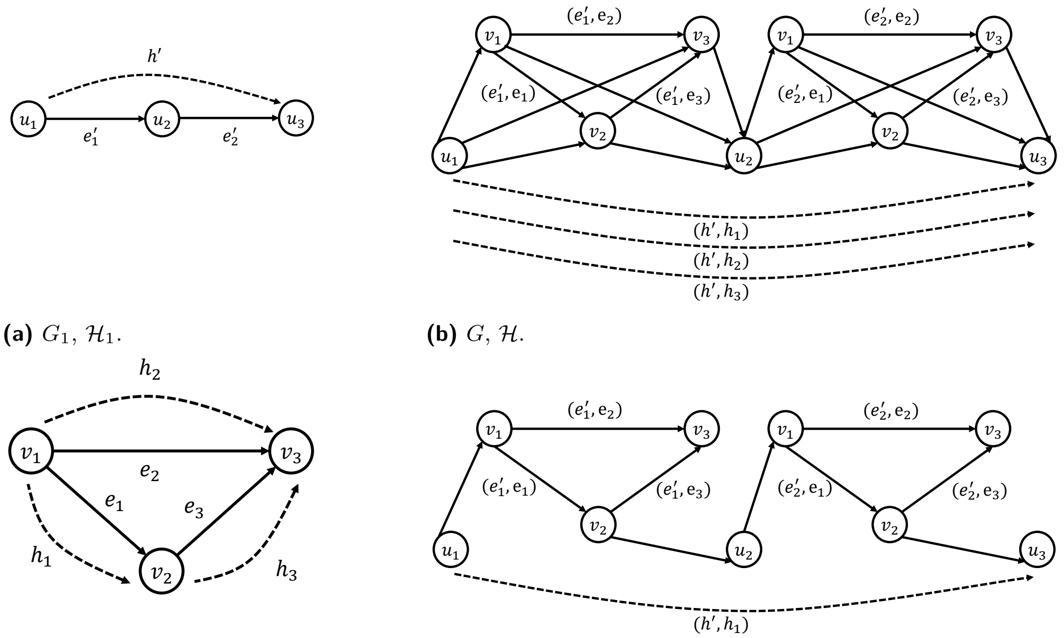

(a) G_1, $\mathcal{H}_1$.

(b) G, $\mathcal{H}$.

(c) G_2, $\mathcal{H}_2$.

(d) $E_{(h',h_1)}$: the set of edges allowed for commodity (h', h_1).

Figure 2 Illustration of the construction of the product instance.

$t(G_2)$). Figure 2 illustrates the construction of the product instance. We denote by (e_1, e_2) the edge e_2 in G'_2 corresponding to the copy of G'_2 related to $e_1 \in E(G_1)$. The capacity of the edge (e_1, e_2) is the product of the capacity of edges e_1 and e_2: $c_{(e_1,e_2)} = c_{1e_1} \cdot c_{2e_2}$.

We create a set of commodities $\mathcal{H}$ in G by taking $\mathcal{H} = \mathcal{H}_1 \times \mathcal{H}_2$ and assuming that $s(h_1, h_2) = s_{h_1}$ and $t(h_1, h_2) = t_{h_1}$ for $(h_1, h_2) \in \mathcal{H}$. We also assume that edges of type $(s(G_2) = s(e), v)$ can only be used by a commodity $(h_1, h_2) \in \mathcal{H}$ if $s(h_2) = v$. Similarly, edges of type $(v, t(G_2) = t(e))$ can only be used by (h_1, h_2) if $t(h_2) = v$. In other words, when a commodity (h_1, h_2) is routed through the copy of G_2 related to an edge $e \in E(G_1)$, then it should enter from $s(h_2)$ and leave at $t(h_2)$ (cf. Figure 2). Other flow restrictions are added by considering the restrictions related to $\mathcal{I}_1$ and $\mathcal{I}_2$. If $h' \in \mathcal{H}_1$ is not allowed to use edge $e' \in E(G_1)$, then all commodities (h', h_2) are not allowed to be routed through the e' copy of G'_2. Moreover, if $e_2 \in E(G_2)$ does not belong to E_{h_2} for some $h_2 \in \mathcal{H}_2$, then for each $e_1 \in E(G_1)$ and each $h_1 \in \mathcal{H}_1$, (e_1, e_2) cannot be used to route commodity (h_1, h_2).
Observe that $|E(G)| = |E(G_1)| \cdot (|E(G_2)| + 2|V(G_2)|)$, $|V(G)| = |V(G_1)| + |V(G_2)| \cdot |E(G_1)|$.

We define $\mathcal{D}$ as the set of vectors $d \in \mathbb{R}_+^{\mathcal{H}_1 \times \mathcal{H}_2}$ such that there is a vector $d^1 \in \mathcal{D}_1$ satisfying $d_{h_1,.} \in d^1_{h_1} \mathcal{D}_2$ for all $h_1 \in \mathcal{H}_1$. The constraint $d_{h_1,.} \in d^1_{h_1} \mathcal{D}_2$ can be enforced with linear inequalities as follows. Suppose that $\mathcal{D}_2 = \{d^2 \in \mathbb{R}^{\mathcal{H}_2} | A_2 d^2 + B_2 \xi \leq b_2\}$ for some matrices A_2, B_2. We also assume that this description contains the constraints $d^2_{h_2}/d^{2\,max}_{h_2} \leq 1$ for all $h_2 \in \mathcal{H}_2$ in addition to the non-negativity constraints of demand values $d^2_{h_2}$. Then we can write the constraint $d_{h_1,.} \in d^1_{h_1} \mathcal{D}_2$ as $A_2 d_{h_1,.} + B_2 \xi' - d^1_{h_1} b_2 \leq 0$. Indeed, $d^1_{h_1} = 0$ implies $d_{h_1,.} = 0$ while for $d^1_{h_1} > 0$ we have $A_2 d_{h_1,.} + B_2 \xi' - d^1_{h_1} b_2 \leq 0$ if and only if $d_{h_1,.}/d^1_{h_1} \in \mathcal{D}_2$. Polytope $\mathcal{D}$ is then defined by constraints $A_2 d_{h_1,.} + B_2 \xi^{h_1} - d^1_{h_1} b_2 \leq 0$ for each $h_1 \in \mathcal{H}_1$ in addition to $A_1 d^1 + B_1 \xi \leq b_1$. Observe that a subscript h_1 is added to express the fact that the auxiliary variables ξ^{h_1} depend on $h_1 \in \mathcal{H}_1$. Notice also that the size of the matrices defining $\mathcal{D}$ is less than or equal to the product of the sizes of the matrices defining $\mathcal{D}_1$ and $\mathcal{D}_2$.

We will now prove that $\text{cong}_{\text{dyn}}(\mathcal{I}) \le \text{cong}_{\text{dyn}}(\mathcal{I}_1) \cdot \text{cong}_{\text{dyn}}(\mathcal{I}_2)$. Let d be a vector in $\mathcal{D}$ and let $d^1 \in \mathcal{D}_1$ be a vector such that $d_{h_1,.} \in d^1_{h_1}\mathcal{D}_2$. For h_1 in $\mathcal{H}_1$, we define $d^{2,h_1} \in \mathcal{D}_2$ by $d^{2,h_1}_{h_2} = \frac{d_{h_1 h_2}}{d^1_{h_1}}$ if $d^1_{h_1} \neq 0$ and $d^{2,h_1} = 0$ if $d^1_{h_1} = 0$. We clearly have $d_{h_1,h_2} = d^1_{h_1} \cdot d^{2,h_1}_{h_2}$ for all $h_1 \in \mathcal{H}_1, h_2 \in \mathcal{H}_2$.
Let f^1, f^{2,h_1} be the optimal routing schemes for $d^1 \in \mathbb{R}^{\mathcal{H}_1}$ and $d^{2,h_1} \in \mathbb{R}^{\mathcal{H}_2}$ for $h_1 \in \mathcal{H}_1$. To route commodity (h_1, h_2), we consider the following multi-commodity flow in G defined by $f_{(h_1,h_2),(e_1,e_2)} = f^1_{h_1,e_1} f^{2,h_1}_{h_2,e_2}$. The total flow on the edge (e_1, e_2) is then given by:

$$\begin{aligned}
\sum_{(h_1,h_2)\in\mathcal{H}_1\times\mathcal{H}_2} d^1_{h_1} d^{2,h_1}_{h_2} f_{(h_1,h_2),(e_1,e_2)} &= \sum_{h_1\in\mathcal{H}_1} d^1_{h_1} f^1_{h_1,e_1} \sum_{h_2\in\mathcal{H}_1} d^{2,h_1}_{h_2} f^{2,h_1}_{h_2,e_2} \\
&\le \sum_{h_1\in\mathcal{H}_1} d^1_{h_1} f^1_{h_1,e_1} \text{cong}_{\text{dyn}}(\mathcal{I}_2) c_{2e_2} \\
&\le \text{cong}_{\text{dyn}}(\mathcal{I}_1) \cdot \text{cong}_{\text{dyn}}(\mathcal{I}_2) \cdot c_{1e_1} \cdot c_{2e_2} \\
&= \text{cong}_{\text{dyn}}(\mathcal{I}_1) \cdot \text{cong}_{\text{dyn}}(\mathcal{I}_2) \cdot c_{(e_1,e_2)}.
\end{aligned}$$

Since this holds for any edge (e_1, e_2) of G (the other edges of G have an infinite capacity), we deduce that $\text{cong}_{\text{dyn}}(\mathcal{I}) \le \text{cong}_{\text{dyn}}(\mathcal{I}_1) \cdot \text{cong}_{\text{dyn}}(\mathcal{I}_2)$.

Let us now show that $\text{cong}_{\text{cut}}(\mathcal{I}) \ge \text{cong}_{\text{cut}}(\mathcal{I}_1) \cdot \text{cong}_{\text{cut}}(\mathcal{I}_2)$. Let C_1 (resp. C_2) be a cut of G_1 (resp. G_2) achieving the maximal congestion $\text{cong}_{\text{cut}}(\mathcal{I}_1)$ (resp. $\text{cong}_{\text{cut}}(\mathcal{I}_2)$), and let $d^1 \in \mathcal{D}_1$ (resp. $d^2 \in \mathcal{D}_2$) be a demand vector for which the maximal cut congestion is obtained. In other words, we have $\sum_{h_1\in\mathcal{H}_{C_1}} d^1_{h_1} / \sum_{e_1\in C_1} c_{e_1} = \text{cong}_{\text{cut}}(\mathcal{I}_1)$ and $\sum_{h_1\in\mathcal{H}_{C_2}} d^2_{h_2} / \sum_{e_2\in C_2} c_{e_2} = \text{cong}_{\text{cut}}(\mathcal{I}_2)$.
Observe that the set of edges $C_1 \times C_2$ is a cut of G that is intersecting all demands of $\mathcal{H}_{C_1} \times \mathcal{H}_{C_2}$. Notice that the flow restrictions that have been considered are crucial here to guarantee the previous fact. Let $d \in \mathbb{R}^{\mathcal{H}}$ be the demand defined by $d_{(h_1,h_2)} = d^1_{h_1} \cdot d^2_{h_2}$. Since $d^1 \in \mathcal{D}_1$ and $d^2 \in \mathcal{D}_2$, we also have $d \in \mathcal{D}$. The congestion on the cut $C_1 \times C_2$ is given by:

$$\begin{aligned}
\frac{\sum_{(h_1,h_2)\in\mathcal{H}_{C_1}\times\mathcal{H}_{C_2}} d_{(h_1,h_2)}}{\sum_{(e_1,e_2)\in C_1\times C_2} c_{(e_1,e_2)}} &= \frac{\sum_{h_1\in\mathcal{H}_{C_1}} d^1_{h_1}}{\sum_{e_1\in C_1} c_{1e_1}} \cdot \frac{\sum_{h_2\in\mathcal{H}_{C_2}} d^2_{h_2}}{\sum_{e_2\in C_2} c_{2e_2}} \\
&= \text{cong}_{\text{cut}}(\mathcal{I}_1) \cdot \text{cong}_{\text{cut}}(\mathcal{I}_2).
\end{aligned}$$

This clearly implies that $\text{cong}_{\text{cut}}(\mathcal{I}) \ge \text{cong}_{\text{cut}}(\mathcal{I}_1) \cdot \text{cong}_{\text{cut}}(\mathcal{I}_2)$. ◀

Proof of Proposition 15. Let us take $p(m) = \log^{c_1}(m)$ for an arbitrary constant c_1. If we run a polynomial-time algorithm on instance the instance $\mathcal{I}^{p(m)}$, we get an algorithm running in poly-logarithmic time. The number of commodities k is bounded by $(\theta m)^{p(m)}$. Thus $\log k \le \log^{c_1} m \log \theta m < \log^{c_1+2} m$ for big enough m and therefore $m > exp(\log^{\frac{1}{c_1+2}} k)$.

The gap between the congestion of the instances $\mathcal{I}^{p(m)}$ corresponding to 3-SAT instances such that $val(\varphi) < 1 - \rho$ and those such that $val(\varphi) = 1$ is:

$$\begin{aligned}
(1+\rho)^{p(m)} &> (1+\rho)^{p(\log^{\frac{1}{c_1+2}} k))} \\
&= (1+\rho)^{\log^{\frac{c_1}{c_1+2}} k} \\
&> (1+\rho)^{\log^{1-\epsilon} k}
\end{aligned}$$

for any $\epsilon > 0$ if we take c_1 such that $\frac{c_1}{c_1+2} > 1 - \epsilon$. The $(1 + \rho)$ term can be replaced by 2 by observing that $2^{\log^{1-\epsilon'} k} = o((1+\rho)^{\log^{1-\epsilon} k})$ for any $\epsilon' < \epsilon$. The same proof applies if parameter n is considered instead of k. ◀