



HAL
open science

On the approximability of robust network design

Yacine Al-Najjar, Walid Ben-Ameur, Jérémie Leguay

► **To cite this version:**

Yacine Al-Najjar, Walid Ben-Ameur, Jérémie Leguay. On the approximability of robust network design. Theoretical Computer Science, 2021, 860, pp.41-50. 10.1016/j.tcs.2021.01.026 . hal-03557332

HAL Id: hal-03557332

<https://hal.science/hal-03557332>

Submitted on 9 Mar 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution - NonCommercial 4.0 International License

On the Approximability of Robust Network Design

Yacine Al-Najjar^{a,b,*}, Walid Ben-Ameur^b, Jérémie Leguay^a

^a*Huawei Technologies, Paris Research Center, France.*

^b*Samovar, Telecom SudParis, Institut Polytechnique de Paris, France.*

Abstract

Given the dynamic nature of traffic, we investigate the variant of robust network design where we have to determine the capacity to reserve on each link so that each demand vector belonging to a polyhedral set can be routed. The objective is either to minimize congestion or a linear cost. Routing is assumed to be fractional and dynamic (i.e., dependent on the current traffic vector). We first prove that the robust network design problem with minimum congestion cannot be approximated within any constant factor. Then, using the ETH conjecture, we get a $\Omega(\frac{\log n}{\log \log n})$ lower bound for the approximability of this problem. This implies that the well-known $O(\log n)$ approximation ratio established by Räcke in 2008 is tight. Using Lagrange relaxation, we obtain a new proof of the $O(\log n)$ approximation. An important consequence of the Lagrange-based reduction and our inapproximability results is that the robust network design problem with linear reservation cost cannot be approximated within any constant ratio. This answers a long-standing open question of Chekuri (2007). We also give another proof of the result of Goyal&al (2009) stating that the optimal linear cost under static routing can be $\Omega(\log n)$ more expensive than the cost obtained under dynamic routing. Finally, we show that even if only two given paths are allowed for each commodity, the robust network design problem with minimum congestion or linear cost is hard to approximate within some constant.

Keywords: Approximability, PCP, ETH, Robust Network Design.

*Corresponding author. Full postal address:
18 quai du point du jour, 92100 Boulogne-Billancourt, France.

Email addresses: yacine.alnajjar@huawei.com (Yacine Al-Najjar),
walid.benameur@telecom-sudparis.eu (Walid Ben-Ameur),
jeremie.leguay@huawei.com (Jérémie Leguay)

1. Introduction

Network optimization [1, 2] plays a crucial role for telecommunication operators since it permits to carefully invest in infrastructures, i.e. reduce capital expenditures. As Internet traffic is ever increasing, the network's capacity needs to be expanded through careful investments every year or even half-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.

Ideally, the network capacity should follow the demand. When the traffic demand can be precisely known, several approaches have been proposed to solve the capacitated network design problem using for instance decomposition methods and cutting planes [3, 4, 5]. But in practice, perfect knowledge of future traffic is not available at the time the decision needs to be taken. The dynamic nature of the traffic due to ordinary daily fluctuations, long term evolution and unpredictable events requires to consider uncertainty on traffic demands when dimensioning network resources. While overestimated traffic forecasts could be used to solve a deterministic optimization problem, it is likely to yield to a costly over-provisioning of the network capacities, which is not acceptable. Therefore, robust optimization under uncertainty sets is a must for the design of network capacities. In this context, our paper presents 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's consider an undirected 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 the routing of a commodity can be represented by a flow $f^h \in \mathbb{R}^{E(G)}$ of intensity d_h . To take into account the changing nature of the demand, d 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* [8, 9], where limits on the total traffic going into (resp. out of) a node are considered.

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

36 reserved capacity on edge e . Another common objective function is given by
 37 the linear reservation cost $\sum_{e \in E(G)} \lambda_e u_e$. This can also represent the average
 38 congestion by taking $\lambda_e = \frac{1}{c_e} \cdot \sum_{h \in \mathcal{H}} f_e^h(d)$. A popular approach is to minimize
 39 the maximum edge utilization given by $\max_{e \in E(G)} \frac{u_e}{c_e}$. The goal is to choose a
 40 reservation vector u so that the network is able to support any demand vec-
 41 tor $d \in \mathcal{D}$, i.e., there exists a (fractional) routing serving every commodity
 42 such that the total flow on each edge e is less than the reservation u_e .

43 The robust network design variant that we are focusing on in this paper, is
 44 referred to as *dynamic routing* in the literature since the network is optimized
 45 such that any realization of traffic matrix in the uncertainty set has its own
 46 routing. The robust network design problem where a linear reservation cost is
 47 minimized was proved to be co-NP hard in [10] when the graph is directed. A
 48 stronger co-NP hardness result is given in [11] where the graph is undirected
 49 (this implies the directed case result). Some exact solution methods for
 50 robust network design have been considered in [12, 13]. Some special cases
 51 where dynamic routing is easy to compute have been described in [14, 15, 16].

52 Routing with uncertain demands has received a significant interest from
 53 the community. As opposed to dynamic routing, *static routing* or *stable*
 54 *routing* was introduced in [6]: it consists in choosing a fixed flow x^h of value
 55 1 for each commodity h . The actual flow $f^h(d)$ for the demand scenario d
 56 will then be scaled by the actual demand d_h of commodity h , i.e. $f^h(d) =$
 57 $d_h x^h$. Static routing is also called *oblivious routing* in [17, 18]. In this case,
 58 polynomial-time algorithms to compute optimal static routing (with respect
 59 to either congestion or linear reservation cost) have been proposed [6, 7, 17,
 60 18] based on either duality or cutting-plane algorithms. When the routing
 61 for each commodity is restricted to be on a single path, the problem becomes
 62 *NP-hard* even in the deterministic case, this follows from the *NP-hardness*
 63 of the edge-disjoint path computation problem [19].

64 To further improve solutions of static routing and overcome complexity is-
 65 sues related to *dynamic routing*, a number of restrictions on routing have been
 66 considered to design polynomial-time algorithms (see [20, 21] for a complete
 67 survey). This includes, for example, the multi-static approach, introduced
 68 in [22], where the uncertainty set is partitioned using an hyperplane and
 69 routing is restricted to be static over each partition. This idea has been gen-
 70 eralized in [23] to unrestricted covers of the uncertainty set and an extension
 71 to share the demand between routing templates, called volume routing, has
 72 been proposed in [24]. [25] applied affine routing for robust network design,

73 based on affine adjustable robust counterparts introduced in [26], restricting
 74 the recourse to be an affine function of the uncertainties. The performance
 75 of this framework has been extensively compared to the static and dynamic
 76 routing, both theoretically and empirically [27, 21]. In practice, affine rout-
 77 ing provides a good approximation of the dynamic routing while it can be
 78 solved in reasonable time thanks to polynomial-time algorithms. Finally, an
 79 approach encompassing the previous approaches is the multipolar approach
 80 proposed in [28, 29].

81 In this work, we will only focus on the complexity of the robust network
 82 design problem under dynamic routing, while minimizing either congestion
 83 or some linear cost. To close this section, let us summarize the main contri-
 84 butions of the paper and review some related work.

85 1.1. Our contributions

- 86 • We first prove that the robust network design problem with minimum
 87 congestion cannot be approximated within any constant factor. The
 88 reduction is based on the PCP theorem and some connections with
 89 the Gap-3-SAT problem [30]. The same reduction also allows to show
 90 inapproximability within $\Omega(\log \frac{n}{\Delta})$ where Δ is the maximum degree in
 91 the graph and n is the number of vertices.
- 92 • Using the ETH conjecture [31, 32], we prove a $\Omega(\frac{\log n}{\log \log n})$ lower bound
 93 for the approximability of the robust network design problem with min-
 94 imum congestion. This implies that the well-known $O(\log n)$ approxi-
 95 mation ratio that can be obtained using the result in [33] is tight.
- 96 • We show that any α -approximation algorithm for the robust network
 97 design problem with linear costs directly leads to an α -approximation
 98 for the problem with minimum congestion. The proof is based on La-
 99 grange relaxation. We obtain that robust network design with mini-
 100 mum congestion can be approximated within $O(\log n)$. This was al-
 101 ready proved in [33] in a different way.
- 102 • An important consequence of the Lagrange-based reduction and our in-
 103 approximability results is that the robust network design problem with
 104 linear reservation cost cannot be approximated within any constant
 105 ratio. This answers a long-standing open question stated in [34].

106 Another consequence is a new proof for the existence of instances for
 107 which the optimal static solution can be $\Omega(\log n)$ more expensive than

108 a solution based on dynamic routing, when a linear cost is minimized.
109 This was already proved in [35] in a different way.

- 110 • We show that even if only two given paths are allowed for each com-
111 modity, there is a constant k such that the robust network design prob-
112 lem with minimum congestion or linear costs cannot be approximated
113 within k .

114 1.2. Related work

115 Let us first assume that the graph is undirected and a linear cost is
116 minimized. A result attributed to A. Gupta ([34], see also [35] for a more
117 detailed presentation) leads to an $O(\log n)$ approximation algorithm for lin-
118 ear cost under dynamic fractional routing. Furthermore, this approximation
119 is achieved by a routing on a (fixed) single tree. In particular, this shows
120 that the ratios between the dynamic and the static solutions under frac-
121 tional routing ($\frac{Lin_{stat-frac}}{Lin_{dyn-frac}}$ (Lin denotes here the optimal linear cost of the
122 solution) and between single path and fractional routing under the static
123 model ($\frac{Lin_{stat-sing}}{Lin_{stat-frac}}$) is $O(\log n)$ and provides an $O(\log n)$ approximation for
124 static single path routing $Lin_{stat-sing}$. On the other hand [36] shows that
125 the static single path problem cannot be approximated within a $\Omega(\log^{\frac{1}{4}-\epsilon} n)$
126 ratio unless $NP \not\subseteq ZPTIME(n^{\text{polylog}(n)})$. As noticed in [35], this implies
127 (assuming this complexity conjecture) that the gap $\frac{Lin_{stat-sing}}{Lin_{stat-frac}}$ is $\Omega(\log^{\frac{1}{4}-\epsilon} n)$.
128 [35] has shown that the gap $\frac{Lin_{stat-frac}}{Lin_{dyn-frac}}$ is $\Omega(\log n)$.

129 For the linear cost and undirected graphs, an extensively studied polyhe-
130 dron is the *symmetric* hose model. The demand vector is here not oriented
131 (i.e, there is no distinction between a demand from i to j and a demand from
132 j to i), and uncertainty is defined by considering an upper-bound limit b_i for
133 the sum of demands related to node i . A 2-approximation has been found
134 for the dynamic fractional case [9, 10] based on tree routing (where we route
135 through a static tree that should be found) showing that $\frac{Lin_{stat-tree}}{Lin_{dyn-frac}} \leq 2$. It
136 has been conjectured that this solution resulted in an optimal solution for the
137 static single path routing. This question has been open for some time and
138 has become known as the *VPN conjecture*. It was finally answered by the af-
139 firmative in [37]. The *asymmetric* hose polytope was also considered in many
140 papers. An approximation algorithm is proposed to compute $Lin_{stat-sing}$
141 within a ratio of 3.39 [38] (or more precisely 2 plus the best approximation

142 ratio for the Steiner tree problem). If \mathcal{D} is a *balanced* asymmetric hose poly-
 143 tope, 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
 144 for the traffic entering into (resp. going out of) v , then the best approxi-
 145 mation factor becomes 2 [38]. Moreover, if we assume that $b_v^{out} = b_v^{in}$, then
 146 $Lin_{stat-sing}$ is easy to compute and we get that $Lin_{stat-tree} = Lin_{stat-sing}$
 147 [39]. In other words, there is some similarity with the case where \mathcal{D} is a
 148 symmetric hose polytope.

149 When congestion is considered, [40] proved the existence of an obli-
 150 vious (or static) routing with a competitive ratio of $O(\log^3 n)$ with respect
 151 to optimum routing of any traffic matrix. Then, [41] improved the bound
 152 to $O(\log^2 n \log \log n)$ and gave a polynomial-time algorithm to find such a
 153 static routing. Finally, [33] described an $O(\log n)$ approximation algorithm
 154 for static routing with minimum congestion. Notice that the bound given
 155 by static routing cannot provide a better bound than $O(\log n)$ since a lower
 156 bound of $\Omega(\log n)$ is achieved by static routing for planar graphs [42, 43].

157 It has also been shown in [44] that the gap between the dynamic fractional
 158 routing and a dynamic fractional routing restricted to a polynomial number
 159 of paths can be $\Omega(\frac{\log n}{\log \log n})$.

160 When a directed graph is considered and congestion is minimized, [18] has
 161 shown that the gap between static fractional routing and dynamic fractional
 162 routing can be $\Omega(\sqrt{n})$ while [45] proves that the gap is upper-bounded by
 163 $O(\sqrt{kn}^{\frac{1}{4}} \log n)$ (where $k = |\mathcal{H}|$ is the number of commodities). More results
 164 can be found in [45] and the references therein.

165 Using an approximate separation oracle for the dual problem to obtain
 166 an approximate solution of the primal is a well-known technique already
 167 used in [46, 47, 48] at least in the context of packing-covering problems.
 168 Lagrangian relaxations are also used in [49, 50, 51] to produce dual solutions
 169 that are near-optimal.

170 **2. From Gap-3-SAT to robust network design with minimum con-** 171 **gestion**

172 Given an edge e , let $s(e)$ and $t(e)$ be the extremities of e . Similarly to
 173 edges, for a commodity $h \in \mathcal{H}$, let $s(h)$ and $t(h)$ denote the endpoints of h .
 174 And let $\mathcal{U}(\mathcal{D})$ be the set of $u \in \mathbb{R}^{E(G)}$ such that each traffic vector $d \in \mathcal{D}$
 175 can be routed on the network when a capacity u_e is assigned to edge e . Since \mathcal{D}
 176 is polyhedral, $\mathcal{U}(\mathcal{D})$ is also polyhedral (see, e.g, [34]).

177 We are interested in minimizing the congestion under polyhedral uncer-
 178 tainty and dynamic routing: $\min_{u \in \mathcal{U}(\mathcal{D})} \max_{e \in E(G)} \frac{u_e}{c_e}$.

179 Given a polytope represented by $Ax \leq b$, the size of the polytope denotes
 180 the total encoding size of the entries in A and b .

181 Our first main result is related to the inapproximability of the minimum
 182 congestion problem within a constant factor.

183 **Theorem 2.1.** *Unless $P = NP$, the minimum congestion problem cannot*
 184 *be approximated with a polynomial-time algorithm within any constant factor*
 185 *even if \mathcal{D} is given by $\{d : Ad + B\psi \leq b\}$ whose size is polynomially bounded*
 186 *by $|V(G)|$.*

187 Notice that it is important to consider polyhedral uncertainty sets that are
 188 easy to describe (otherwise the inapproximability results would be a direct
 189 consequence of the difficulty to separate from the uncertainty set).

190 To prove Theorem 2.1, we will need the PCP (Probabilistically Checkable
 191 Proof) theorem [30] and an intermediate lemma. For a 3-SAT formula φ we
 192 note $val(\varphi)$ the maximum fraction of the clauses which are satisfiable at
 193 the same time. In particular, $val(\varphi) = 1$ means that φ is satisfiable. The
 194 problem where we have to decide if $val(\varphi) < \rho$ or $val(\varphi) = 1$ for a 3-SAT
 195 formula φ is called Gap-3-SAT.

196 The instances such that $\rho \leq val(\varphi) < 1$ do not need to be considered.
 197 One way to state PCP theorem is to say that there exists a constant $0 <$
 198 $\rho < 1$ for which Gap-3-SAT is NP-complete. In other words, it is NP-hard to
 199 distinguish between satisfiable 3-SAT formulas and those for which strictly
 200 less than a fraction ρ of clauses can be simultaneously satisfied.

201 To prove the theorem 2.1, we will use the following lemma (where $cong$
 202 denotes the optimal congestion of the corresponding instance).

203 **Lemma 2.1.** *For every $\gamma \in \mathbb{N}$ there is a mapping f_γ computable in poly-*
 204 *nomial time from 3-SAT instances to minimum congestion instances de-*
 205 *finied by an undirected graph G_γ , a set of commodities \mathcal{H}_γ and a polytope*
 206 *$\mathcal{D}_\gamma = \{d : A_\gamma d + B_\gamma \psi_\gamma \leq b_\gamma\}$ such that $|V(G_\gamma)| = O(m^\gamma)$, $|E(G_\gamma)| = O(m^\gamma)$*
 207 *and the size of \mathcal{D}_γ is $O(m^{c\gamma})$ where c is some positive constant and m is the*
 208 *number of clauses. The mapping satisfies the following:*

- 209 • $val(\varphi) = 1 \implies cong(f_\gamma(\varphi)) \geq 1 + \gamma(1 - \rho)$
- 210 • $val(\varphi) < \rho \implies cong(f_\gamma(\varphi)) \leq 1$.

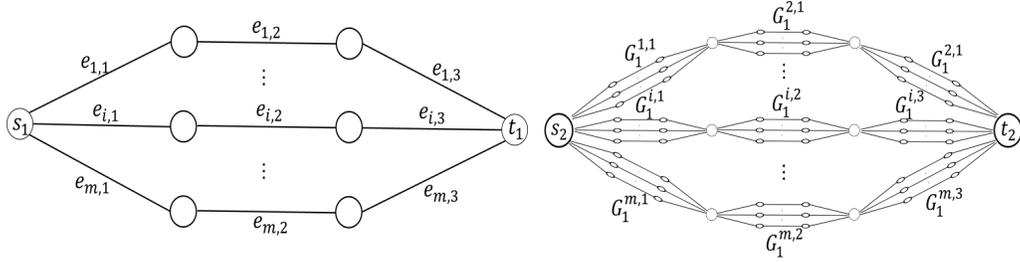


Figure 1: G_1 and G_2

211 *Proof. of Theorem 2.1* We are going to use Lemma 2.1 and PCP Theorem
 212 for the proof. Suppose that congestion can be approximated in polynomial
 213 time within a constant approximation factor α . We first choose γ such that
 214 $\alpha < 1 + \gamma(1 - \rho)$.

215 Starting from a 3-SAT formula such that either $val(\varphi) < \rho$ or $val(\varphi) = 1$,
 216 we construct $f_\gamma(\varphi)$ in polynomial time. The optimal congestion will satisfy
 217 either $cong(f_\gamma(\varphi)) \geq 1 + \gamma(1 - \rho)$ or $cong(f_\gamma(\varphi)) \leq 1$. Applying the α -
 218 approximation to $f_\gamma(\varphi)$ provides an approximate value $\tilde{\beta}$ for congestion. If
 219 $\tilde{\beta} < 1 + \gamma(1 - \rho)$ holds, then we can deduce that $cong(f_\gamma(\varphi)) \leq \tilde{\beta} < 1 + \gamma(1 - \rho)$.
 220 This implies that $cong(f_\gamma(\varphi)) \leq 1$ and hence $val(\varphi) < \rho$. Otherwise, we have
 221 $\tilde{\beta} \geq 1 + \gamma(1 - \rho)$ and $\alpha \times cong(f_\gamma(\varphi)) \geq \tilde{\beta}$ (since $\tilde{\beta}$ is an α -approximation),
 222 leading to $cong(f_\gamma(\varphi)) \geq \frac{1 + \gamma(1 - \rho)}{\alpha} > 1$. We consequently have $cong(f_\gamma(\varphi)) \geq$
 223 $1 + \gamma(1 - \rho)$ and $val(\varphi) = 1$. This proves that a constant α -approximation
 224 for the congestion problem allows the solution of Gap-3-SAT.

225 Furthermore, as the size of the polytope used in Lemma 2.1 is $O(m^{c\gamma})$
 226 while $|V(G_\gamma)| = O(m^\gamma)$, its size is polynomially bounded in the number of
 227 vertices as announced in Theorem 2.1 \square

228 We are now going to prove Lemma 2.1 by first constructing instances of
 229 the congestion problem leading to some inapproximability factor. Then, this
 230 factor is increased by recursively building larger instances with higher values
 231 of γ .

232 *Proof. of Lemma 2.1, case $\gamma = 1$*

233 We start with a 3-SAT formula φ , with m clauses and r variables. We
 234 note $\mathcal{L} = \{l_1, \dots, l_r, \neg l_1, \dots, \neg l_r\}$ the set of the literals appearing in formula
 235 φ and $l_{i,j}$ the literal appearing in the i -th clause C_i at the j -th position for
 236 $i = 1, \dots, m$ and $j = 1, 2, 3$ (it is not restrictive to assume that each clause
 237 contains exactly 3 literals).

238 We build as follows a graph G_1 and a set of commodities \mathcal{H}_1 . For each
 239 $i = 1, \dots, m$, $j = 1, 2, 3$ we add 3 consecutive edges $e_{i,j}$ (i.e. such that $t(e_{i,1}) =$
 240 $s(e_{i,2})$ and $t(e_{i,2}) = s(e_{i,3})$) and 3 commodities $h_{i,j}$ with $s(h_{i,j}) = s(e_{i,j})$ and
 241 $t(h_{i,j}) = t(e_{i,j})$. We impose that all nodes $s(e_{i,1})$ (resp. $t(e_{i,3})$) for $i = 1, \dots, m$
 242 are equal to a single node noted s_1 (resp. t_1) (see Figure 1). We consider an
 243 additional commodity h_0 between s_1 and t_1 .

244 We create a polyhedron \mathcal{D}_1 of the form $\mathcal{D}_1 = \{d : A_1 d + B_1 \psi_1 \leq b_1\}$ as
 245 follows. We consider for each literal $l \in \mathcal{L}$ a non-negative variable ξ_l and add
 246 for $k = 1, \dots, r$ the constraint $\xi_{l_k} + \xi_{-l_k} = 1$. We also consider, for $i = 1, \dots, m$
 247 and $j = 1, 2, 3$, the constraint $d_{h_{i,j}} = \xi_{l_{i,j}}$. A constraint related to d_{h_0} is also
 248 integrated: $d_{h_0} \leq m(1 - \rho)$. Finally, the capacity c_e of each edge e is here
 249 equal to 1 ($c_e = 1$).

250 If $val(\varphi) = 1$, then there is a demand vector such that for each path
 251 between s_1 and t_1 (there is one path corresponding to each clause), at least
 252 one commodity whose endpoints are on the path is equal to 1 (a commodity
 253 corresponding to a true literal). This implies that all paths are blocked and
 254 thus the optimal routing for commodity h_0 is to equally spread $m(1 - \rho)$
 255 between the m paths leading to a congestion of $1 + (1 - \rho)$.

256 Let us now assume that $val(\varphi) < \rho$. Notice first that the components of
 257 the extreme points of the polyhedron \mathcal{D}_1 are integers (except d_{h_0}). This is
 258 due to the fact that \mathcal{D}_1 can be seen as a coordinate projection of the higher
 259 dimensional polytope $\{(d, \psi) | A_1 d + B_1 \psi\}$ whose extreme points are obviously
 260 integers (except d_{h_0}). For such an extreme demand vector $d \in \mathcal{D}_1$ there are
 261 at least $m(1 - \rho)$ free paths to route the demand d_{h_0} allowing a congestion
 262 less than or equal to 1. This implies that all demands in \mathcal{D}_1 can also be
 263 routed with a congestion less than or equal to 1

264 Observe that $|V(G_1)| = O(m)$, $|E(G_1)| = O(m)$, \mathcal{D}_1 has the appropriate
 265 form ($\mathcal{D}_1 = \{d : A_1 d + B_1 \psi_1 \leq b_1\}$) and the size of \mathcal{D}_1 is $O(m^c)$ for some
 266 constant c . \square

267 *Proof. of Lemma 2.1, case $\gamma \geq 2$*

268 For $\gamma \geq 2$, having constructed $G_{\gamma-1}, \mathcal{H}_{\gamma-1}, \mathcal{D}_{\gamma-1}$, we build $G_\gamma, \mathcal{H}_\gamma, \mathcal{D}_\gamma$
 269 as follows. We will construct the graph G_γ , by taking the graph G_1 and
 270 replacing each edge by a copy of the graph $G_{\gamma-1}$ denoted by $G_{\gamma-1}^{i,j}$. Each
 271 copy $G_{\gamma-1}^{i,j}$ contains a node $s_{\gamma-1}$ that is identified with $s(e_{i,j})$ and a node
 272 $t_{\gamma-1}$ identified with $t(e_{i,j})$ (see Figure 1). All commodities related to $G_{\gamma-1}^{i,j}$
 273 (belonging to $\mathcal{H}_{\gamma-1}$) are also considered as commodities of \mathcal{H}_γ . Let us use
 274 $d^{i,j} \in \mathbb{R}^{\mathcal{H}_{\gamma-1}}$ to denote the related demand vector. \mathcal{H}_γ also contains a non-

275 negative commodity $h_{0,\gamma}$ constrained by $d_{h_{0,\gamma}} \leq m^\gamma(1 - \rho)$. Thus $|\mathcal{H}_\gamma| =$
 276 $1 + 3m \times |\mathcal{H}_{\gamma-1}|$.

277 We are going to build an uncertainty set \mathcal{D}_γ as a coordinate projection of
 278 a higher-dimensional polyhedron Ξ_γ , involving demand variables in addition
 279 to auxiliary non-negative variables ξ_l related to literals, and also auxiliary
 280 variables $\psi_{\gamma-1}^{i,j}$ related to $G_{\gamma-1}^{i,j}$ and the description of $\mathcal{H}_{\gamma-1}$. We gradually
 281 explain the construction. For $k = 1, \dots, r$, we add the constraint $\xi_{l_k} + \xi_{-l_k} = 1$.
 282 And for $e_{i,j} \in E(G_1)$, we impose that $d^{i,j} \in \xi_{l_{i,j}} \mathcal{D}_{\gamma-1} := \{\xi_{l_{i,j}} d_0 \mid d_0 \in \mathcal{D}_{\gamma-1}\}$.
 283 Let us explain how this can be done. By induction, we know that $\mathcal{D}_{\gamma-1} = \{d :$
 284 $A_{\gamma-1}d + B_{\gamma-1}\psi_{\gamma-1} \leq b_{\gamma-1}\}$ and this representation includes (among others)
 285 non-negativity constraints of all variables in addition to constraints implying
 286 that all variables are upper-bounded. Then by writing $A_{\gamma-1}d^{i,j} + B_{\gamma-1}\psi_{\gamma-1}^{i,j} \leq$
 287 $\xi_{l_{i,j}} b_{\gamma-1}$, we can ensure that $\xi_{l_{i,j}} = 0$ implies $d^{i,j} = 0$, while $\xi_{l_{i,j}} > 0$ leads
 288 to $\frac{1}{\xi_{l_{i,j}}} d^{i,j} \in \mathcal{D}_{\gamma-1}$. In particular when $\xi_{l_{i,j}} = 0$, from outside, the whole
 289 subgraph corresponding to $G_{\gamma-1}^{i,j}$ acts like a single edge of capacity $m^{\gamma-1}$.

\mathcal{D}_γ can be seen as the projection of a polytope $\Xi_\gamma = \{(d, \psi_\gamma) \mid A_\gamma d + B_\gamma \psi_\gamma \leq$
 $b_\gamma\}$ where ψ_γ contains the auxiliary variables appearing in all levels. More
 precisely, Ξ_γ is defined by:

$$\begin{aligned}
 d_{h_{0,\gamma}} &\leq m^\gamma(1 - \rho) \\
 -d_{h_{0,\gamma}} &\leq 0, \\
 -\xi_l &\leq 0, \quad \forall l \in \mathcal{L} \\
 \xi_{l_k} + \xi_{-l_k} &\leq 1, \quad \forall k = 1, \dots, r \\
 -\xi_{l_k} - \xi_{-l_k} &\leq -1, \quad \forall k = 1, \dots, r \\
 A_{\gamma-1}d^{i,j} + B_{\gamma-1}\psi_{\gamma-1}^{i,j} - \xi_{l_{i,j}}b_{\gamma-1} &\leq 0, \quad \forall i = 1, \dots, m, j = 1, 2, 3. \quad (1)
 \end{aligned}$$

290 By simple induction, we have $|V(G_\gamma)| = O(m^\gamma)$, $|E(G_\gamma)| = O(m^\gamma)$ and
 291 the size of \mathcal{D}_γ is $O(m^{c\gamma})$ where c is some positive constant.

292 We observe that all extreme points of Ξ_γ are such that $\xi_l \in \{0, 1\}$ for
 293 $l \in \mathcal{L}$. To verify that, we first recall that constraints (1) are equivalent to
 294 $d^{i,j} \in \xi_{l_{i,j}} \mathcal{D}_{\gamma-1}$ (in this way, the vectors $\psi_{\gamma-1}^{i,j}$ can be ignored). Second, let \mathcal{L}_+
 295 be the set of literals appearing in positive form. We observe that variables ξ_l
 296 for $l \in \mathcal{L}_+$ are pairwise independent. Only variables $d^{i,j}$ such that either $l_{i,j} =$
 297 l or $l_{i,j} = -l$ depend on ξ_l since $d^{i,j} \in \xi_l \mathcal{D}_{\gamma-1}$ in the first case and $d^{i,j} \in (1 -$
 298 $\xi_l) \mathcal{D}_{\gamma-1}$ in the second case. This immediately implies that given some arbitrary
 299 real vectors $q_{i,j}$ and f , minimizing $\sum_{i=1, \dots, m; j=1, 2, 3} q_{i,j}^T d^{i,j} + \sum_{l \in \mathcal{L}_+} f_l \xi_l$ is equiv-

300 alent to minimizing $\sum_{l \in \mathcal{L}_+} \xi_l \left(f_l + \sum_{i,j:l_{i,j}=l} \min_{d^{i,j} \in \mathcal{D}_{\gamma-1}} q_{i,j}^T d^{i,j} - \sum_{i,j:l_{i,j}=-l} \min_{d^{i,j} \in \mathcal{D}_{\gamma-1}} q_{i,j}^T d^{i,j} \right)$.

301 It is then clear that optimal ξ_l values will be either 0 or 1. Since this holds
 302 for an arbitrary linear objective function, we get the wanted result about
 303 extreme points.

304

305 Let us now show that $val(\varphi) < \rho \implies cong(f_\gamma(\varphi)) \leq 1$. Assume that
 306 $val(\varphi) < \rho$. We prove by induction that the congestion of $(G_\gamma, \mathcal{H}_\gamma, \mathcal{D}_\gamma)$ is
 307 1. Suppose that this is true for some $\gamma - 1$. If $\xi_{l_{i,1}} = \xi_{l_{i,2}} = \xi_{l_{i,3}} = 0$ for
 308 some i , a flow of value $m^{\gamma-1}$ can be routed between s_γ and t_γ by sending a
 309 flow of value 1 on each edge of $G_{\gamma-1}^{i,j}$ for $j = 1, 2, 3$. Since $val(\varphi) < \rho$, there
 310 are necessarily at least $m(1 - \rho)$ such i , thus we can send the whole demand
 311 $m^{\gamma-1}m(1 - \rho) = m^\gamma(1 - \rho)$ this way. For the indices i, j such that $\xi_{l_{i,j}} = 1$,
 312 by the induction hypothesis ($cong(f_{\gamma-1}(\varphi)) \leq 1$), the demands inside $G_{\gamma-1}^{i,j}$
 313 can be routed without sending more than one unit of flow on each edge of
 314 $G_{\gamma-1}^{i,j}$.

315 Notice that to show that all traffic vectors of \mathcal{D}_γ can be routed with con-
 316 gestion 1, we considered demand vectors corresponding with $\{0, 1\}$ ξ vari-
 317 ables. The result shown above about extreme points is useful here since it
 318 allows us to say that each extreme point of \mathcal{D}_γ can be routed with congestion
 319 less than or equal to 1 implying that each demand vector inside \mathcal{D}_γ can also
 320 be routed with congestion less than or equal to 1.

321

322 Let us now show that $val(\varphi) = 1 \implies cong(f_\gamma(\varphi)) \geq 1 + \gamma(1 - \rho)$. We
 323 are going to use induction to build a cut $\delta(C_\gamma)$ where C_γ is set of vertices
 324 of $V(G_\gamma)$ containing s_γ and not containing t_γ . The number of edges of the
 325 cut will be m^γ and each edge has a capacity equal to 1. We also show the
 326 existence of a demand vector $d \in \mathcal{D}_\gamma$ such that the sum of the demands
 327 traversing the cut is greater than or equal to $m^\gamma(1 + \gamma(1 - \rho))$. This would
 328 show that there is at least one edge that carries at least $1 + \gamma(1 - \rho)$ units of
 329 flow.

330 Since φ is satisfiable, there is a truth assignment represented by ξ variables
 331 (the auxiliary variables) such that for each $i = 1, \dots, m$ there is a $j(i)$ such
 332 that $\xi_{l_{i,j(i)}} = 1$. By considering the graph $G_{\gamma-1}^{i,j(i)}$ and using the induction
 333 hypothesis, we can build a cut $\delta(C_{\gamma-1}^i)$ separating the node $s(e_{i,j(i)})$ and
 334 $t(e_{i,j(i)})$ and containing $m^{\gamma-1}$ edges. We also build a demand vector $d^{i,j(i)} \in$
 335 $\mathcal{D}_{\gamma-1}$ such that the sum of demands traversing the cut is greater than or equal

336 to $m^{\gamma-1}(1 + (\gamma - 1)(1 - \rho))$ (still possible by induction). By taking the union
337 of these m disjoint cuts we get a cut $\delta(C_\gamma)$ that is separating s_γ and t_γ having
338 the required number of edges. A demand vector d can be built by combining
339 the vectors $d^{i,j(i)}$ and the demand $d_{h_0,\gamma}$ taken equal to $m^\gamma(1 - \rho)$. Since the
340 demand from s_γ to t_γ is also traversing the cut, the total demand through
341 $\delta(C_\gamma)$ is greater than or equal to $m^\gamma(1 - \rho) + m \cdot m^{\gamma-1}(1 + (\gamma - 1)(1 - \rho)) =$
342 $m^\gamma(1 + \gamma(1 - \rho))$. \square

343 Lemma 2.1 can be further exploited in different ways since there are many
344 possible connections between the value $1 + \gamma(1 - \rho)$ and the characteristics of
345 the undirected graph built in the proof of the lemma. Observe, for example,
346 that by a simple induction we get that the number of vertices $|V(G_\gamma)| =$
347 $2 + 2m \frac{(3m)^\gamma - 1}{3m - 1}$ leading to $|V(G_\gamma)| \simeq 2 \times 3^{\gamma-1} m^\gamma$ (when m goes to infinity).
348 We also have $\Delta(G_\gamma)$ equal to m^γ where $\Delta(\cdot)$ denotes the maximum degree
349 in the graph. Consequently, $\log\left(\frac{|V(G_\gamma)|}{\Delta(G_\gamma)}\right) \simeq \gamma \log 3 + \log 2/3$. Then by taking
350 any constant k such that $k \times \log 3 < (1 - \rho)$ where ρ is the constant in the
351 PCP Theorem we get a lower bound of the approximability ratio. This is
352 stated in the following corollary.

353 **Corollary 2.1.** *Under conditions of Theorem 2.1, for any constant $k <$
354 $\frac{1-\rho}{\log 3}$, it is not possible to approximate the minimum congestion problem in
355 polynomial time within a ratio of $k \log \frac{|V(G)|}{\Delta(G)}$.*

356 3. A $\Omega\left(\frac{\log n}{\log \log n}\right)$ approximability lower bound

357 To get an approximability lower bound, we will use the well-known ETH
358 conjecture that is recalled below.

359 **Conjecture 3.1** (Exponential Time Hypothesis). *[31, 32] There is a con-*
360 *stant δ such that no algorithm can solve 3-SAT instances in time $O(2^{\delta m})$,*
361 *where m is the number of clauses.*

362 Let us use n to denote the number of vertices of the graph.

363 **Theorem 3.1.** *Under Conjecture 3.1, there exists a constant k such that no*
364 *polynomial-time algorithm can solve the minimum congestion problem with*
365 *the approximation ratio $k \frac{\log n}{\log \log n}$.*

366 *Proof.* The combination of PCP Theorem and ETH Conjecture 3.1 implies
 367 that distinguishing between 3-SAT instances such that $val(\varphi) < \rho$ and
 368 $val(\varphi) = 1$ cannot be done in time $O(2^{m^\beta})$ for some constant $\beta > 0$ (a
 369 slightly better bound is $O(2^{m/\log^c m})$ for some constant c , but this will not
 370 help us to improve the lower bound of Theorem 3.1).

371 Suppose that there is an algorithm that solves the minimum congestion
 372 problem with an approximation factor $\alpha(n)$ and a running time $O(n^{c_1})$.
 373 Given a 3-SAT instance and a function $\gamma : \mathbb{N} \rightarrow \mathbb{N}$ we can construct a
 374 minimum congestion instance $f_{\gamma(m)}(\varphi)$ as in Lemma 2.1 in time $O(m^{c_2\gamma(m)})$
 375 and where the number of vertices of the instance is $m^{\gamma(m)}$. Then by running
 376 the approximation algorithm for minimum congestion we get a total time of
 377 $O(m^{c_3\gamma(m)})$ where $c_3 = \max\{c_1, c_2\}$. Thus by choosing $\gamma(m) = \frac{m^\beta}{c_3 \log m}$ we get
 378 an algorithm that runs in time $O(2^{m^\beta})$. And if the approximation factor $\alpha(n)$
 379 is small enough, that is if $\alpha(m^{\gamma(m)}) < 1 + (1 - \rho)\gamma(m)$ for a big enough m ,
 380 we get an algorithm solving Gap-3-SAT and thus contradicting Conjecture
 381 3.1. This is the case for $k \frac{\log n}{\log \log n}$ for some constant k . To see this, we can
 382 observe that:

383
$$\frac{1+(1-\rho)\gamma(m)}{\alpha(m^{\gamma(m)})} = \frac{1+(1-\rho)\frac{m^\beta}{c_3 \log m}}{k \frac{m^\beta/c_3}{\beta \log m - \log c_3}} \simeq \frac{\beta(1-\rho)}{k}$$
. By taking $k < \beta(1 - \rho)$ we get the
 384 wanted inapproximability result. \square

385 4. From minimum congestion to linear costs

386 Given any $\lambda \geq 0$, the robust network design problem with linear costs
 387 is simply the following where $\mathcal{U}(\mathcal{D})$ is the set of possible capacity vectors
 388 defined in Section 2:

$$\min_{u \in \mathcal{U}(\mathcal{D})} \lambda^T u. \quad (2)$$

389 Assume that there exists a number $\alpha \geq 1$ such that Problem (2) can be
 390 solved in polynomial time within an approximation ratio α . More precisely,
 391 we have a polynomial-time oracle that takes as input a non-negative linear
 392 cost $\lambda \in \mathbb{R}^{E(G)}$ and outputs a $u^{ap}(\lambda) \in \mathcal{U}(\mathcal{D})$ such that $\lambda^T u(\lambda) \leq \lambda^T u^{ap}(\lambda) \leq$
 393 $\alpha \lambda^T u(\lambda)$ where $u(\lambda) \in \mathcal{U}(\mathcal{D})$ is the optimal solution of (2).

394 Recall that the congestion problem is given by

$$\begin{aligned}
& \min_{\beta, u} \beta & (3) \\
& u_e \leq c_e \beta, \forall e \in E(G) \\
& u \in \mathcal{U}(\mathcal{D})
\end{aligned}$$

395 where β and u are optimization variables.

396 Let us consider a Lagrange relaxation of (3) by dualizing the capacity
397 constraints and using λ for the dual multipliers. The dual problem is then
398 given by $\max_{\lambda \geq 0} \min_{\beta, u \in \mathcal{U}(\mathcal{D})} \beta + \sum_{e \in E(G)} \lambda_e (u_e - \beta c_e)$ (where β is an optimization
399 variable). If λ is chosen such $\sum_e \lambda_e c_e \neq 1$, then the value of the inner minimum
400 would be $-\infty$. Thus in an optimal solution, we will always have $\sum_e \lambda_e c_e = 1$.

401 The problem is then equivalent to:

$$\max_{\substack{\lambda \geq 0 \\ \sum_{e \in E(G)} \lambda_e c_e = 1}} \min_{u \in \mathcal{U}(\mathcal{D})} \sum_{e \in E(G)} \lambda_e u_e = \max_{\substack{\lambda \geq 0 \\ \sum_{e \in E(G)} \lambda_e c_e = 1}} \lambda^T u(\lambda). \quad (4)$$

402 Since $\mathcal{U}(\mathcal{D})$ is polyhedral and all constraints and the objective function are
403 linear, there is no duality gap between (3) and (4).

Observe that (4) can be expressed as follows:

$$\max_{\beta, \lambda \geq 0} \beta \quad (5a)$$

$$\beta \leq \sum_{e \in E(G)} \lambda_e u_e, \forall u \in \mathcal{U}(\mathcal{D}) \quad (5b)$$

$$1 = \sum_{e \in E(G)} \lambda_e c_e \quad (5c)$$

404 We are going to approximately solve (5) using a cutting-plane algorithm
405 where inequalities (5b) are iteratively added by using the α -approximation or-
406 acle. Let (β', λ') be a potential solution of (5), we can run the α -approximation
407 of robust network design problem (2) with the cost vector λ' to get a solution
408 $u^{ap}(\lambda')$. If $\beta' > \sum_{e \in E(G)} \lambda'_e u_e^{ap}(\lambda')$ we return the inequality $\beta \leq \sum_{e \in E(G)} \lambda_e u_e^{ap}(\lambda')$,
409 otherwise the algorithm stops and returns (β', λ') . We know from the separation-
410 optimization equivalence theorem [52] that (5) can be solved by making
411 a polynomial number of calls to the separation oracle leading a globally

412 polynomial-time algorithm. Notice that this happens if the separation oracle
 413 is exact. In our case, the oracle is only an approximate one, implying that the
 414 cutting plane algorithm might be prematurely interrupted before obtaining
 415 the true optimum of (5). Observe however that this implies that the com-
 416 puting time is polynomially bounded. Let $(\tilde{\beta}, \tilde{\lambda})$ be the solution returned by
 417 the cutting-plane algorithm. Let (β^*, λ^*) be the true optimal solution of (5).
 418 The next lemma states that the returned solution is an α -approximation of
 419 the optimal solution.

420 **Lemma 4.1.** *The cutting-plane algorithm computes in polynomial time a*
 421 *solution $\tilde{\beta}$ satisfying:*

$$\beta^* \leq \tilde{\beta} \leq \alpha\beta^*. \quad (6)$$

422

423 *Proof.* Observe that $\beta^* = \lambda^{*T}u(\lambda^*)$. Moreover, since (5) is equivalent to
 424 (4), we get that $\lambda^{*T}u(\lambda^*) = \beta^* \geq \tilde{\lambda}^T u(\tilde{\lambda})$. From the approximation factor
 425 of the oracle, one can write that $\tilde{\lambda}^T u^{ap}(\tilde{\lambda}) \leq \alpha\tilde{\lambda}^T u(\tilde{\lambda})$. Using the fact that
 426 no inequalities can be added for $(\tilde{\beta}, \tilde{\lambda})$, we get that $\tilde{\beta} \leq \tilde{\lambda}^T u^{ap}(\tilde{\lambda})$. Finally,
 427 since (β^*, λ^*) is feasible for (5), we obviously have $\tilde{\beta} \geq \beta^*$. Combining the 4
 428 previous inequalities leads to (6). \square

429 The above lemma has many consequences.

430 **Theorem 4.1.** *Unless $P = NP$, the robust network design problem with*
 431 *linear costs cannot be approximated in polynomial time within any constant*
 432 *ratio. Unless the ETH conjecture is false, the robust network design problem*
 433 *with linear costs cannot be approximated within $\Omega(\frac{\log n}{\log \log n})$.*

434 *Proof.* The result is an immediate consequence of Theorems 2.1, 3.1 and
 435 Lemma 4.1. \square

436 The theorem above answers a long-standing open question of [34]. All
 437 other inapproximability results proved for the congestion problem directly
 438 hold for the robust network design problem with linear cost.

439 Another important consequence is that the congestion problem can be
 440 approximated within $O(\log n)$. This result was already proved in [33] using
 441 other techniques. In our case, the result is an immediate consequence of
 442 the $O(\log n)$ -approximation algorithm for the robust network design problem
 443 with linear cost provided by [53, 54] and fully described in [34, 35].

444 **Theorem 4.2.** [33] *Congestion can be approximated within $O(\log n)$.*

445 Notice that Theorem 3.1 tells us that the ratio $O(\log n)$ is tight.

446 Starting from the results of [42, 43] showing the existence of instances for
 447 which the ratio $\frac{cong_{stat-frac}}{cong_{dyn-frac}}$ is $\Omega(\log n)$, one can also use the reduction above
 448 to prove, differently from [35], the existence of instances for which the ratio
 449 $\frac{Lin_{stat-frac}}{Lin_{dyn-frac}}$ is $\Omega(\log n)$ where a linear cost is minimized.

450 **Theorem 4.3.** [35] *There are instances for which $\frac{Lin_{stat-frac}}{Lin_{dyn-frac}}$ is $\Omega(\log n)$.*

451 *Proof.* Similarly to $\mathcal{U}(\mathcal{D})$ defined when dynamic routing is considered, let
 452 $\mathcal{U}_{stat}(\mathcal{D})$ be the set of capacity vectors for which there exists a static frac-
 453 tional routing satisfying all demand vectors of \mathcal{D} . $\mathcal{U}_{stat}(\mathcal{D})$ is obviously a
 454 polyhedral set. The mathematical programs (3), (4) and (5) can be consid-
 455 ered in the same way: we only have to replace $\mathcal{U}(\mathcal{D})$ by $\mathcal{U}_{stat}(\mathcal{D})$. All results
 456 stated above about the equivalence of (3), (4) and (5) still hold in the static
 457 case. Consider an instance from [42, 43] for which $\frac{cong_{stat-frac}}{cong_{dyn-frac}}$ is $\Omega(\log n)$.
 458 $cong_{stat-frac}$ is computed from (5). Then there is at least one vector $\lambda^{stat} \geq 0$
 459 and one vector $u^{stat} \in \mathcal{U}_{stat}(D)$ such that $cong_{stat-frac} = \sum_{e \in E(G)} \lambda_e^{stat} u_e^{stat}$
 460 and $\sum_{e \in E(G)} \lambda_e^{stat} c_e = 1$. This implies that u^{stat} is an optimal solution of
 461 the linear problem where we minimize $\sum_{e \in E(G)} \lambda_e^{stat} u_e$ under the condition
 462 $u \in \mathcal{U}_{stat}(D)$. We consequently have $cong_{stat-frac} = Lin_{stat-frac}$ for the con-
 463 sidered instance.

464 Let $u' \in \mathcal{U}(\mathcal{D})$ be an optimal solution minimizing the linear cost $\sum_{e \in E(G)} \lambda_e^{stat} u_e$
 465 under dynamic routing. In other words, $Lin_{dyn-frac} = \sum_{e \in E(G)} \lambda_e^{stat} u'_e$ when
 466 the coefficients of the objective function are λ^{stat} . Moreover, we know from
 467 (5) that $cong_{dyn-frac}$ is obtained by maximizing through λ , implying that
 468 $cong_{dyn-frac} \geq \sum_{e \in E(G)} \lambda_e^{stat} u'_e = Lin_{dyn-frac}$. Using that $\frac{cong_{stat-frac}}{cong_{dyn-frac}}$ is
 469 $\Omega(\log n)$, we get that $\frac{Lin_{stat-frac}}{Lin_{dyn-frac}}$ is $\Omega(\log n)$ for the same instance where the
 470 linear objective function is defined through λ^{stat} . \square

471 5. Restriction to a constant number of given paths per commodity

472 First, observe that in the proof of Lemma 2.1, the minimum congestion
 473 instances built there are such that some commodities can be routed along
 474 many paths. For example, in graph G_1 (Figure 1), commodity h_0 (between
 475 s and t) can use up to m paths. Second, consider an instance of the mini-
 476 mum congestion problem where only one path is given for each commodity.

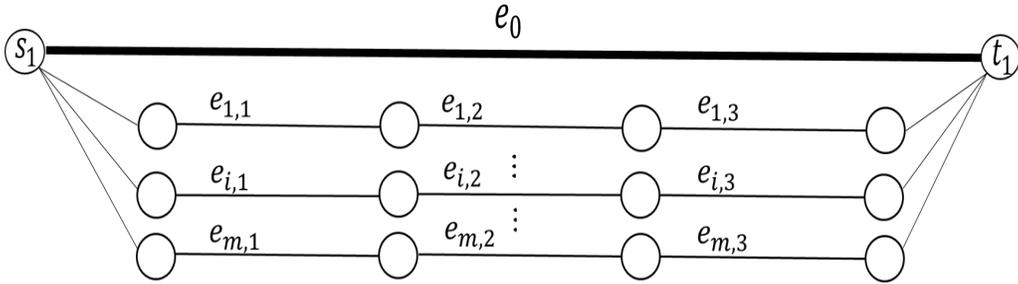


Figure 2: G'

477 Then computing the minimum congestion is easy since we only have to compute
 478 $\max_{d \in \mathcal{D}} \sum_{h \in \mathcal{H}_e} d_h$ where \mathcal{H}_e denotes the set of commodities routed through
 479 e . The congestion is just given by $\max_{e \in E(G)} \frac{1}{c_e} \max_{d \in \mathcal{D}} \sum_{h \in \mathcal{H}_e} d_h$. Combining these
 480 two observations, one can wonder whether the difficulty of the congestion
 481 problem is simply due to the number of possible paths that can be used by
 482 each commodity. We will show that the problem is still difficult even if each
 483 commodity can be routed along at most two fixed given paths.

484 **Theorem 5.1.** *Unless $P = NP$, for some positive constant k , minimum*
 485 *congestion cannot be approximated within a ratio k even if each commodity*
 486 *can be routed along at most two given paths.*

487 *Proof.* The proof is a simple modification of the proof of Lemma 2.1 (case
 488 $\gamma = 1$). We are going to slightly modify graph G_1 in such a way that at most
 489 2 paths are allowed for each commodity. Given a 3-SAT formula φ with m
 490 clauses, we construct $G', \mathcal{H}', \mathcal{D}'$ as follows. We first create two nodes s_1 and
 491 t_1 and an edge e_0 between s_1 and t_1 of capacity $m\rho$ (ρ is the constant in
 492 PCP theorem). Then for each clause index $i = 1, \dots, m$, as in Lemma 2.1,
 493 we create 3 consecutive edges $e_{i,j}$ ($j = 1, 2, 3$) such that $t(e_{i,j}) = s(e_{i,j+1})$
 494 and a commodity $h_{i,j}$ between $s(e_{i,j})$ and $t(e_{i,j})$ that is allowed to be routed
 495 only through $e_{i,j}$. We also add one edge between $s(e_{i,1})$ and s_1 and one edge
 496 connecting t_1 and $t(e_{i,3})$ of infinite capacity and a commodity $h_{i,0}$ between
 497 $s(e_{i,1})$ and $t(e_{i,3})$ with a demand $d_{h_{i,0}} = 1$. $h_{i,0}$ is allowed to be routed only
 498 through the path P_i containing the edges $(e_{i,1}, e_{i,2}, e_{i,3})$ and the path going
 499 through s_1, e_0 and t_1 (see Figure 2). We consider auxiliary variables ξ_l
 500 for each literal l . We add constraints $\xi_l + \xi_{-l} = 1$ and $d_{h_{i,j}} = \xi_{l_{i,j}}$.

501 If $val(\varphi) < \rho$ there are at least $m(1 - \rho)$ commodities $h_{i,0}$ that can be
502 routed on the paths P_i and the remaining $m\rho$ can be routed on the edge e_0 .
503 This implies that each extreme point of \mathcal{D}' can be routed with congestion
504 ≤ 1 . Notice that the observation made in the proof of Lemma 2.1 about
505 extreme points is still valid here: extreme points corresponds to 0 – 1 values
506 of the variables ξ_l .

507 If $val(\varphi) = 1$, then there is a cut and a demand vector d (corresponding to
508 the truth assignment satisfying φ) such that the capacity of the cut is $m\rho + m$
509 and the demand that needs to cross the cut is $2m$. There is consequently
510 at least one edge of congestion greater than or equal to $\frac{2m}{(1+\rho)m} = \frac{2}{1+\rho}$. By
511 taking $k < \frac{2}{1+\rho}$ we get the wanted result. \square

512 Finally, observe that the result above can also be stated for the linear
513 cost case using again the Lagrange based reduction of the previous section.

514 **Corollary 5.1.** *Unless $P = NP$, for some positive constant k , robust network*
515 *design with linear costs is difficult to approximate within a ratio k even if each*
516 *commodity can be routed along at most two given paths.*

517 Acknowledgments

518 The authors are grateful to anonymous referees for their helpful sugges-
519 tions that greatly improved the quality of the paper. We also point out that
520 Theorem 4.3 was added following a suggestion of one of the reviewers who
521 was wondering whether the reduction of Section 4 could be used to prove the
522 theorem.

523 **References**

- 524 [1] N. Wang, K. H. Ho, G. Pavlou, M. Howarth, An overview of routing opti-
525 mization for internet traffic engineering, *IEEE Communications Surveys
526 & Tutorials* 10 (1) (2008) 36–56.
- 527 [2] H. P. L. Luna, Network planning problems in telecommunications, in:
528 *Handbook of Optimization in Telecommunications*, Springer, 2006, pp.
529 213–240.
- 530 [3] A. M. Costa, A survey on benders decomposition applied to fixed-charge
531 network design problems, *Computers & operations research* 32 (6) (2005)
532 1429–1450.
- 533 [4] A. Frangioni, B. Gendron, 0–1 reformulations of the multicommod-
534 ity capacitated network design problem, *Discrete Applied Mathematics*
535 157 (6) (2009) 1229–1241.
- 536 [5] C. Raack, A. M. Koster, S. Orlowski, R. Wessály, On cut-based inequal-
537 ities for capacitated network design polyhedra, *Networks* 57 (2) (2011)
538 141–156.
- 539 [6] W. Ben-Ameur, H. Kerivin, New economical virtual private networks,
540 *Association for Computing Machinery. Communications of the ACM*
541 46 (6) (2003) 69–69.
- 542 [7] W. Ben-Ameur, H. Kerivin, Routing of uncertain traffic demands, *Op-
543 timization and Engineering* 6 (3) (2005) 283–313.
- 544 [8] N. G. Duffield, P. Goyal, A. Greenberg, P. Mishra, K. K. Ramakrishnan,
545 J. E. van der Merive, A flexible model for resource management in virtual
546 private networks, in: *Proceedings of the conference on Applications,
547 technologies, architectures, and protocols for computer communication*,
548 1999, pp. 95–108.
- 549 [9] J. A. Fingerhut, S. Suri, J. S. Turner, Designing least-cost nonblocking
550 broadband networks, *Journal of Algorithms* 24 (2) (1997) 287–309.
- 551 [10] A. Gupta, J. Kleinberg, A. Kumar, R. Rastogi, B. Yener, Provisioning a
552 virtual private network: a network design problem for multicommodity
553 flow, in: *Proceedings of the thirty-third annual ACM symposium on
554 Theory of computing*, 2001, pp. 389–398.

- 555 [11] C. Chekuri, F. B. Shepherd, G. Oriolo, M. G. Scutella, Hardness of
556 robust network design, *Networks: An International Journal* 50 (1) (2007)
557 50–54.
- 558 [12] G. Claßen, A. M. Koster, M. Kutschka, I. Tahiri, Robust metric in-
559 equalities for network loading under demand uncertainty, *Asia-Pacific*
560 *Journal of Operational Research* 32 (05) (2015).
- 561 [13] S. Mattia, The robust network loading problem with dynamic routing,
562 *Computational Optimization and Applications* 54 (3) (2013) 619–643.
- 563 [14] W. Ben-Ameur, M. Żotkiewicz, Robust routing and optimal partitioning
564 of a traffic demand polytope, *International Transactions in Operational*
565 *Research* 18 (3) (2011) 307–333.
- 566 [15] M. Minoux, Robust network optimization under polyhedral demand un-
567 certainty is NP-hard, *Discrete Applied Mathematics* 158 (5) (2010) 597–
568 603.
- 569 [16] A. Frangioni, F. Pascali, M. G. Scutellà, Static and dynamic routing
570 under disjoint dominant extreme demands, *Operations research letters*
571 39 (1) (2011) 36–39.
- 572 [17] D. Applegate, E. Cohen, Making intra-domain routing robust to chang-
573 ing and uncertain traffic demands: Understanding fundamental trade-
574 offs, in: *Proceedings of the 2003 conference on Applications, technolo-*
575 *gies, architectures, and protocols for computer communications, 2003,*
576 pp. 313–324.
- 577 [18] Y. Azar, E. Cohen, A. Fiat, H. Kaplan, H. Räcke, Optimal oblivious
578 routing in polynomial time, *Journal of Computer and System Sciences*
579 69 (3) (2004) 383–394.
- 580 [19] S. Fortune, J. Hopcroft, J. Wyllie, The directed subgraph homeomor-
581 phism problem, *Theoretical Computer Science* 10 (2) (1980) 111–121.
- 582 [20] W. Ben-Ameur, A. Ouorou, M. Żotkiewicz, Robust routing in commu-
583 nication networks, *Progress in Combinatorial optimization*, Iste, Wiley,
584 Mahjoub Eds. (2012) 353–389.

- 585 [21] M. Poss, A comparison of routing sets for robust network design, *Opti-*
586 *mization Letters* 8 (5) (2014) 1619–1635.
- 587 [22] W. Ben-Ameur, Between fully dynamic routing and robust stable rout-
588 *ing*, in: 2007 6th International Workshop on Design and Reliable Com-
589 *munication Networks*, IEEE, 2007, pp. 1–6.
- 590 [23] M. G. Scutellà, On improving optimal oblivious routing, *Operations*
591 *Research Letters* 37 (3) (2009) 197–200.
- 592 [24] M. Zotkiewicz, W. Ben-Ameur, More adaptive robust stable routing, in:
593 *GLOBECOM IEEE Global Telecommunications Conference*, 2009, pp.
594 1–6.
- 595 [25] A. Ouorou, J.-P. Vial, A model for robust capacity planning for telecom-
596 *munications networks under demand uncertainty*, in: *International*
597 *workshop on design and reliable communication networks*, IEEE, 2007,
598 pp. 1–4.
- 599 [26] A. Ben-Tal, A. Goryashko, E. Guslitzer, A. Nemirovski, Adjustable ro-
600 *bust solutions of uncertain linear programs*, *Mathematical programming*
601 99 (2) (2004) 351–376.
- 602 [27] M. Poss, C. Raack, Affine recourse for the robust network design prob-
603 *lem: Between static and dynamic routing*, *Networks* 61 (2) (2013) 180–
604 198.
- 605 [28] W. Ben-Ameur, M. Żotkiewicz, Multipolar routing: where dynamic and
606 *static routing meet*, *Electronic Notes in Discrete Mathematics* 41 (2013)
607 61–68.
- 608 [29] W. Ben-Ameur, A. Ouorou, G. Wang, M. Żotkiewicz, Multipolar ro-
609 *bust optimization*, *EURO Journal on Computational Optimization* 6 (4)
610 (2018) 395–434.
- 611 [30] S. Arora, B. Barak, *Computational Complexity: A Modern Approach*,
612 1st Edition, Cambridge University Press, New York, NY, USA, 2009.
- 613 [31] R. Impagliazzo, R. Paturi, On the complexity of k-sat, *Journal of Com-*
614 *puter and System Sciences* 62 (2) (2001) 367–375.

- 615 [32] R. Impagliazzo, R. Paturi, F. Zane, Which problems have strongly ex-
616 ponential complexity?, *Journal of Computer and System Sciences* 63 (4)
617 (2001) 512–530.
- 618 [33] H. Räcke, Optimal hierarchical decompositions for congestion minimiza-
619 tion in networks, in: *Proceedings of the fortieth annual ACM symposium*
620 *on Theory of computing*, 2008, pp. 255–264.
- 621 [34] C. Chekuri, Routing and network design with robustness to changing or
622 uncertain traffic demands, *ACM SIGACT News* 38 (3) (2007) 106–129.
- 623 [35] N. Goyal, N. Olver, F. B. Shepherd, Dynamic vs. oblivious routing
624 in network design, in: *European Symposium on Algorithms*, Springer,
625 2009, pp. 277–288.
- 626 [36] N. Olver, F. B. Shepherd, Approximability of robust network design,
627 *Mathematics of Operations Research* 39 (2) (2014) 561–572.
- 628 [37] N. Goyal, N. Olver, F. B. Shepherd, The vpn conjecture is true, *Journal*
629 *of the ACM (JACM)* 60 (3) (2013) 1–17.
- 630 [38] F. Eisenbrand, F. Grandoni, G. Oriolo, M. Skutella, New approaches
631 for virtual private network design, *SIAM Journal on Computing* 37 (3)
632 (2007).
- 633 [39] N. K. Olver, Robust network design, Ph.D. thesis, McGill University
634 Library (2010).
- 635 [40] H. Racke, Minimizing congestion in general networks, in: *The 43rd An-*
636 *annual IEEE Symposium on Foundations of Computer Science*, 2002. *Pro-*
637 *ceedings.*, IEEE, 2002, pp. 43–52.
- 638 [41] C. Harrelson, K. Hildrum, S. Rao, A polynomial-time tree decompo-
639 sition to minimize congestion, in: *Proceedings of the fifteenth annual*
640 *ACM symposium on Parallel algorithms and architectures*, 2003, pp.
641 34–43.
- 642 [42] B. M. Maggs, F. M. auf der Heide, B. Vocking, M. Westermann, Ex-
643 ploiting locality for data management in systems of limited bandwidth,
644 in: *Proceedings 38th Annual Symposium on Foundations of Computer*
645 *Science*, IEEE, 1997, pp. 284–293.

- 646 [43] Y. Bartal, S. Leonardi, Ondashline routing in all-optical networks, The-
647 oretical Computer Science 221 (1-2) (1999) 19–39.
- 648 [44] M. Hajiaghayi, R. Kleinberg, T. Leighton, Semi-oblivious routing: lower
649 bounds, in: Proceedings of ACM-SIAM symposium on Discrete algo-
650 rithms, Citeseer, 2007, pp. 929–938.
- 651 [45] M. T. Hajiaghayi, R. D. Kleinberg, H. Räcke, T. Leighton, Oblivious
652 routing on node-capacitated and directed graphs, ACM Transactions on
653 Algorithms (TALG) 3 (4) (2007) 51.
- 654 [46] L. Fleischer, M. X. Goemans, V. S. Mirrokni, M. Sviridenko, Tight
655 approximation algorithms for maximum general assignment problems,
656 in: Proceedings of ACM-SIAM symposium on Discrete algorithms, 2006,
657 p. 611–620.
- 658 [47] R. Carr, S. Vempala, Randomized metarounding, Random Structures &
659 Algorithms 20 (3) (2002) 343–352.
- 660 [48] K. Jain, M. Mahdian, M. R. Salavatipour, Packing steiner trees, in:
661 Proceedings of ACM-SIAM symposium on Discrete algorithms, Vol. 3,
662 2003, pp. 266–274.
- 663 [49] N. Garg, J. Koenemann, Faster and simpler algorithms for multicom-
664 modity flow and other fractional packing problems, SIAM Journal on
665 Computing 37 (2) (2007) 630–652.
- 666 [50] S. A. Plotkin, D. B. Shmoys, É. Tardos, Fast approximation algorithms
667 for fractional packing and covering problems, Mathematics of Opera-
668 tions Research 20 (2) (1995) 257–301.
- 669 [51] E. Neal, Randomized rounding without solving the linear program, in:
670 Proceedings of the sixth annual ACM-SIAM symposium on Discrete
671 algorithms, Vol. 76, 1995, p. 170.
- 672 [52] M. Grötschel, L. Lovász, A. Schrijver, Geometric Algorithms and
673 Combinatorial Optimization, Vol. 2 of Algorithms and Combinatorics,
674 Springer, 1988.
- 675 [53] A. Gupta, Personal communication (2004).

- 676 [54] A. Gupta, J. Könemann, Approximation algorithms for network de-
677 sign: A survey, *Surveys in Operations Research and Management Sci-*
678 *ence* 16 (1) (2011) 3–20.

Declaration of interests

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

The authors declare the following financial interests/personal relationships which may be considered as potential competing interests: