European Journal of Operational Research 000 (2016) 1–15

Contents lists available at ScienceDirect

European Journal of Operational Research

journal homepage: www.elsevier.com/locate/ejor

Discrete Optimization

Design of survivable networks with vulnerability constraints

Luis Gouveia[a], Markus Leitner[b,*]

[a] Departamento de Estatistica e Investigaçao Operational, Faculdade de Ciênçias, Universidade de Lisboa, CMAFCIO, Lisbon, Portugal
[b] Department of Statistics and Operations Research, Faculty of Business, Economics and Statistics, University of Vienna, Vienna, Austria

ARTICLE INFO

Article history:
Received 18 November 2015
Accepted 2 September 2016
Available online xxx

Keywords:
Networks
Integer programming
Survivable network design
Hop-constraints
OR in telecommunications

ABSTRACT

We consider the Network Design Problem with Vulnerability Constraints (NDPVC) which simultaneously addresses resilience against failures (network survivability) and bounds on the lengths of each communication path (hop constraints). Solutions to the NDPVC are subgraphs containing a path of length at most H_{st} for each commodity $\{s, t\}$ and a path of length at most H'_{st} between s and t after at most $k-1$ edge failures. We first show that a related and well known problem from the literature, the Hop-Constrained Survivable Network Design Problem (kHSNDP), that addresses the same two measures produces solutions that are too conservative in the sense that they might be too expensive in practice or may even fail to provide feasible solutions. We also explain that the reason for this difference is that Mengerian-like theorems not hold in general when considering hop-constraints. Three graph theoretical characterizations of feasible solutions to the NDPVC are derived and used to propose integer linear programming formulations. In a computational study we compare these alternatives with respect to the lower bounds obtained from the corresponding linear programming relaxations and their capability of solving instances to proven optimality. In addition, we show that in many cases, the solutions produced by solving the NDPVC are cheaper than those obtained by the related kHSNDP.

1. Introduction

For Internet service providers, it is essential to provide stable and reliable communications between any two points in their supporting telecommunication networks. However, it is not easy to provide a precise definition of reliability. A vast body of literature, both in the engineering and operations research community, suggests various concepts for network reliability. This is because the reliability of a network depends on several factors. On the one side, it depends on the technical equipment installed along the links and nodes of the network. On the other side, even with the best available equipment, reliability may be easily destroyed, if the underlying network *topology* is vulnerable to failures. Therefore, resistance to network failures (also known as network *survivability*) has been used in the network optimization literature as one of the main criteria for designing reliable communication networks (see, e.g., Kerivin & Mahjoub, 2005). A network is said to be survivable, if communications between nodes can be established, even after failures of a pre-defined number of nodes or links. Starting with the seminal work by Grötschel, Monma, and Stoer (1995), a large body of mathematical models and algorithmic approaches for designing survivable networks has been proposed.

Another important issue for Internet service providers is quality of service, see e.g., Klincewicz (2006). Each packet of a data flow traveling through a path from its source node to its destination node suffers a total delay that is given by the propagation delay on each link and the queuing and transmission delays on each intermediate node. Jitter, defined as the time difference between the maximum delay and the minimum delay among all packets of a data flow, is an important quality of service parameter that should be bounded to guarantee a given quality of service (Sheikh & Ghafoor, 2011). This parameter is of particular importance for multimedia services (see, Roychoudhuri, Al-Shaer, & Brewster, 2006) but also for data service running over mobile networks (see, Scharf, Necker, & Gloss, 2004). Note that the dominant factor on jitter is the queuing delay since propagation introduces a constant delay on each packet and transmission delay is only dependent on packet size statistics. A simple way of bounding jitter is to bound the number of packet queues, which is equivalent to bound the number of hops of each routing path. Hence, the quality of service can be ensured by imposing so-called hop-constraints.

Recent literature suggests to combine survivability and quality of service by additionally imposing hop-constraints when designing survivable networks. Thus, one guarantees that for every distinct pair of nodes, there exists a pre-defined number of edge/node

* Corresponding author.
E-mail addresses: legouveia@fc.ul.pt (L. Gouveia), markus.leitner@univie.ac.at (M. Leitner).

http://dx.doi.org/10.1016/j.ejor.2016.09.003
0377-2217/

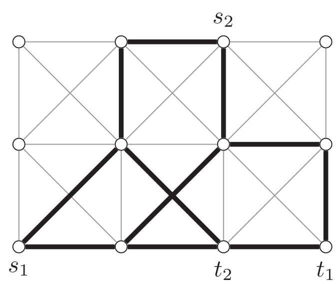

Fig. 1. Example of an instance of the NDPVC with $\mathcal{R} = \{\{s_1, t_1\}, \{s_2, t_2\}\}$, $H_{st} = 3$, $H'_{st} = 4$, $\forall\{s, t\} \in \mathcal{R}$, and a feasible solution to this instance (bold edges).

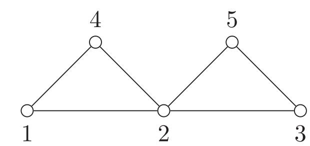

Fig. 2. A feasible solution for the edge disjoint variant of NDPVC with $k = 2$, $\mathcal{R} = \{\{1, 3\}\}$, $H_{13} = 2$, $H'_{13} = 3$, that does not contain two edge disjoint paths of length 2 and 3 between nodes 1 and 3, cf. Exoo (1982).

disjoint paths, such that each such path does not exceed the given hop limit (Botton, Fortz, Gouveia, & Poss, 2013). In this article we show that solutions to this problem variant are too conservative and too expensive, from the perspective of a network provider. We therefore propose to study a new (and related) problem, that ensures both survivability, and the maintenance of the hop-limits after a failure of a pre-defined number of nodes or links, but with significantly less restrictions on the underlying network topology. We call the problem the *Network Design Problem with Vulnerability Constraints* (NDPVC). The term "vulnerability" is coined from graph theory (see, e.g., Bermond, Bond, Paoli, & Peyrat, 1983) and is usually associated to the study of changes in the distance between pairs of nodes due to graph alterations (e.g., edge or node removals). We will have more to say about this in Section 2. As in many related problems, one may consider protection against edge or node failures. We focus on the case of edge-connectivity for which the formal problem definition is given below. The node-connectivity case will be briefly addressed in Section 5.

Problem definition and motivation. We are given an undirected graph $G = (V, E)$, with nonnegative edge costs $c_e \geq 0$, for all $e \in E$, and a parameter $k \in \mathbb{N}$ specifying the network survivability. In addition, we are given a set of commodities $\mathcal{R} \subseteq V \times V \backslash \{(v, v) \mid v \in V\}$ and two hop limits, $H_{st} \leq H'_{st} \in \mathbb{N}$, for each pair $\{s, t\} \in \mathcal{R}$. The goal is to find a minimum cost subgraph of G, such that for each pair $\{s, t\} \in \mathcal{R}$, it contains a path of length at most H_{st}, and after removal of any $k - 1$ edges from it, the resulting graph contains a path of length at most H'_{st}.

Fig. 1 illustrates an input graph G with commodities $\mathcal{R} = \{\{s_1, t_1\}, \{s_2, t_2\}\}$, $k = 2$, and hop limits $H_{st} = 3$, $H'_{st} = 4$ for all $\{s, t\} \in \mathcal{R}$, together with a feasible solution. Notice, that the well known, NP-hard, survivable network design problem (see, e.g., Grötschel et al., 1995) is a special case of the NDPVC when the hop-limits are redundant. Thus the NDPVC is NP-hard as well.

A related problem that has already been studied in the literature is the hop-constrained survivable network design problem (kHSNDP). In this problem, one searches for a minimum-cost subgraph, such that between each pair of commodities there exist k edge- (node-) disjoint paths, each containing at most H edges. Various integer programming formulations and solution algorithms have been proposed recently for the kHSNDP, see, e.g., Botton et al. (2013); Gouveia, Patricio, and de Sousa (2006); Mahjoub, Simonetti, and Uchoa (2013). At first one may assume that the NDPVC and the kHSNDP are equivalent, at least when $H = H'$. The reason for this is that if one would ignore the hop constraints, the two problems, kHSNDP and NDPVC, become equivalent. This equivalence follows immediately from Menger's theorem (Menger, 1927). As a matter of fact, quite often in the literature, mathematical formulations for modeling survivable networks are derived using the results of the Menger's theorem (see, e.g, Grötschel et al., 1995; Kerivin & Mahjoub, 2005; Ljubić., 2010).

Unfortunately, once the hop-constraints are imposed, the two problems are no longer equivalent, since hop-constrained *Mengerian-like theorems* (see Section 2 for a more formal definition and discussion) are valid only for small or large hop-limits. To see that the two problems, NDPVC and kHSNDP, are different in general case, consider the example given in Fig. 2. Assume that a network provider wants to make this network survivable against single edge failures (i.e., $k = 2$), and to protect the vulnerability of the network by assuming that for a commodity pair $\mathcal{R} = \{\{1, 3\}\}$, the distance between 1 and 3 should be at most $H_{13} = 2$ and, after a single edge failure, this distance should not be greater than $H'_{13} = 3$. The path $P = (1, 2, 3)$ is the unique (1, 3)-path of length less than or equal to H_{13}. If an arbitrary edge in this graph fails, the solution will still contain a feasible (1, 3)-path of length at most three. Hence, the graph depicted in Fig. 2 is a feasible NDPVC solution. On the other hand, if the network provider would try solve the related kHSNDP on this graph, instead, then one can easily observe that no feasible solution exist. A feasible kHSNDP solution needs to contain two *edge-disjoint* paths between 1 and 3, such that one path contains at most H_{13} edges and the other at most H'_{13} edges. This is, however impossible, since the only (1, 3)-path P' that is edge disjoint to P, is given by $P' = (1, 4, 2, 5, 3)$ and has length four. This example illustrates that for network providers it could be more attractive to consider the NDPVC to protect vulnerability of a network, rather than the kHSNDP.

This example can be easily generalized for $k \geq 3$ and for other values of H_{st} and H'_{st}, $\{s, t\} \in \mathcal{R}$. The optimal value of kHSNDP always gives an upper bound on the optimal value of the NDPVC. It is easy to find examples where solutions to both problems exist but the optimal solution to the kHSNDP is more expensive than of the NDPVC. Observe that not only the gap between the cost of an optimal solution of NDPVC and kHSDNP can be arbitrarily large, but, as demonstrated above, there exist networks which are feasible for NDPVC but infeasible for kHSNDP. Since these relations motivate this work we summarize them in Observation 1.

Observation 1. Let I be an arbitrary, feasible instance of the NDPVC and $v(I)$ be its optimal cost. Then, exactly one of the following holds:

(i) There does not exist a feasible solution of the kHSNDP for I.
(ii) $v(I) \leq v'(I)$, where $v'(I)$ is the optimal cost of the kHSNDP for I.

Furthermore, there exist instances such that $v(I) < v'(I)$ and $\frac{v'(I)}{v(I)}$ can be arbitrary large.

Note also that, cases in which $\mathcal{R}$ is induced by all node pairs $\{s, t\}$ from a given set of "terminals" T, i.e., $\mathcal{R} = \{\{s, t\} \mid s, t \in T\}$, give rise to several interesting diameter variants of the problem. As one example, we mention the case when $H_{st} = D$ and $H'_{st} = D'$ for all $\{s, t\} \in \mathcal{R}$, in which case we aim to identify a minimum cost Steiner subgraph with diameter at most D such that after the removal of k edges (nodes), the graph is connected and has diameter at most D'.

The remainder of this article in which we mainly focus on the case of single edge failures (i.e. $k = 2$) is organized as follows. We