European Journal of Operational Research 225 (2013) 223–235

Contents lists available at SciVerse ScienceDirect

European Journal of Operational Research

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

Discrete Optimization

Hierarchical approach for survivable network design

Anantaram Balakrishnan [a], Mihai Banciu [b], Karolina Glowacka [c], Prakash Mirchandani [d,*]

[a] McCombs School of Business, The University of Texas at Austin, Austin, TX 78712, USA
[b] School of Management, Bucknell University, Lewisburg, PA 17837, USA
[c] Department of Supply Chain Management, Hang Seng Management College, New Territories, Hong Kong
[d] Katz Graduate School of Business, University of Pittsburgh, Pittsburgh, PA 15260, USA

ARTICLE INFO

Article history:
Received 1 January 2012
Accepted 26 September 2012
Available online 17 October 2012

Keywords:
Metaheuristics
Survivable network design
Heuristic analysis
Tabu search

ABSTRACT

A central design challenge facing network planners is how to select a cost-effective network configuration that can provide uninterrupted service despite edge failures. In this paper, we study the Survivable Network Design (SND) problem, a core model underlying the design of such resilient networks that incorporates complex cost and connectivity trade-offs. Given an undirected graph with specified edge costs and (integer) connectivity requirements between pairs of nodes, the SND problem seeks the minimum cost set of edges that interconnects each node pair with at least as many edge-disjoint paths as the connectivity requirement of the nodes. We develop a hierarchical approach for solving the problem that integrates ideas from decomposition, tabu search, randomization, and optimization. The approach decomposes the SND problem into two subproblems, Backbone design and Access design, and uses an iterative multi-stage method for solving the SND problem in a hierarchical fashion. Since both subproblems are NP-hard, we develop effective optimization-based tabu search strategies that balance intensification and diversification to identify near-optimal solutions. To initiate this method, we develop two heuristic procedures that can yield good starting points. We test the combined approach on large-scale SND instances, and empirically assess the quality of the solutions vis-à-vis optimal values or lower bounds. On average, our hierarchical solution approach generates solutions within 2.7% of optimality even for very large problems (that cannot be solved using exact methods), and our results demonstrate that the performance of the method is robust for a variety of problems with different size and connectivity characteristics.

1. Introduction

As businesses and individuals become increasingly dependent on round-the-clock use of telecommunication services to interact and access information (including data and multi-media content), their expectations about the performance of the networks that provide these services has also increased. In particular, customers seek uninterrupted service despite failures in one or more edges of the network. To meet these expectations, the underlying telecommunications network must be configured so that it contains redundancies in terms alternate paths. A network is said to be *survivable* if it continues to function (allows its customers to communicate and access online services) even after some of the edges fail. Undirected networks that have a tree configuration (with no redundancy) fall on one end of the survivability spectrum, whereas those that have a complete (i.e., fully connected) topology lie at the other extreme. In a tree network, the failure of just a single edge disrupts the required communication between at least one pair of nodes; in contrast, a complete network on n nodes is highly survivable since every pair of network nodes can continue to communicate even if any $n-2$ edges of the network fail. Complete networks provide the highest level of reliability, but they are also very expensive because they assume that all pairs of nodes require the same (high) level of protection against edge failures even though only a subset of nodes may have stringent connectivity requirements. For instance, nodes representing hospitals, emergency call-centers, airports, and financial institutions are very important, and require a high level of protection via backup communication paths that can be used when the primary paths connecting these facilities fail. On the other hand, less important nodes such as individual customers may not need this level of protection since these users may be willing to tolerate temporary disruptions of their services. To develop cost-effective topologies, network planners must address complex cost-connectivity trade-offs to take advantage of the differential connectivity requirements.

In this paper, we address the *Survivable Network Design* (SND) problem, a core network design problem to identify the least cost network that provides the required connectivity, expressed in terms of number of edge-disjoint paths, between pairs of nodes. The SND

* Corresponding author. Tel.: +1 412 648 1652; fax: +1 412 648 1693.
E-mail addresses: anantb@mail.utexas.edu (A. Balakrishnan), mihai.banciu@bucknell.edu (M. Banciu), karolinaglowacka@hsmc.edu.hk (K. Glowacka), pmirchan@katz.pitt.edu (P. Mirchandani).

0377-2217/$ - see front matter © 2012 Elsevier B.V. All rights reserved.
http://dx.doi.org/10.1016/j.ejor.2012.09.045

problem is *NP*-hard (Garey and Johnson, 1979) as are many of its special cases; indeed, the SND problem generalizes several well-known, but intractable, problems including the Traveling Salesman problem (TSP) and the Steiner tree problem. Researchers have also studied variants of the SND problem. For instance, Baldacci et al. (2007) and Naji-Azimi et al. (2010) study the *m*-ring-star problem that requires installing rings, each with size no more than *m* nodes and intersecting only at a central node, to span some or all customer nodes, and attaching other customers to the rings. Terblanche et al. (2011) consider network configuration and equipment installation to satisfy multiple non-simultaneous demand scenarios using alternate routes between demand nodes. Since solving SND problems using standard models and methods has proved to be very challenging, researchers have focused on strengthening the problem formulation via strong valid inequalities to accelerate exact solution methods, characterizing the worst-case performance of heuristic procedures, and designing efficient algorithms for special cases. Grötschel et al. (1995), Raghavan and Magnanti (1997), and Kerivin and Mahjoub (2005) provide excellent surveys of the research on SND problems; Contreras and Fernández (2011) provide a recent review and classification of network design problems. Recent work includes linear-time algorithms for network design on series–parallel graphs (Raghavan, 2004), strong formulations based on bidirectional flow (Magnanti and Raghavan, 2005), analysis of the tightness of connectivity-splitting models (Balakrishnan et al., 2004), and optimization-based heuristics using connectivity-upgrading models (Balakrishnan et al., 2009). Yet, effectively solving large-scale instances of SND problems (e.g., with more than 50 nodes) remains elusive.

In this paper, we develop and test a composite algorithm for solving large SND problems. The approach, combining decomposition, metaheuristics, randomization, and optimization, has several distinctive characteristics. At its core, the method employs a *hierarchical* (multi-stage) decomposition framework, consisting of *backbone network design* and *access network design*, to improve SND solutions. For each stage, we develop a tailored *tabu search* method, combined with *exact algorithms* for embedded subproblems, to explore the appropriate solution neighborhood. To initiate the algorithm, we apply two optimization-based heuristic methods that can generate good starting solutions. Subsequent iterations of the multi-stage procedure are driven by a randomization step to augment the current best solution and permit fuller exploration of the solution space. The various components of our algorithm exploit the problem's underlying special structures and solution characteristics, and incorporate tradeoffs between diversification and intensification of the neighborhood search procedures taking into account both solution speed and quality.

We implemented the composite solution approach and tested it on 78 SND problem instances, varying in size and in the structure of connectivity requirements (both the maximum number of edge-disjoint paths needed and the proportion of nodes at each connectivity level). The test problems contain up to 100 nodes and 400 edges, and nodes have connectivity requirements of up to four edge-disjoint paths. The compact flow-based integer programming formulations for these instances have up to 79,600 variables (of which 400 are binary) and 128,600 constraints. CPLEX could not solve some large problems (with more than 80 nodes and 320 edges) even after 72 hours of computation, whereas our method finds near-optimal solutions relatively quickly. On average, the solutions we obtain have an optimality gap of 2.7% relative to the optimal value or a lower bound on this value. Thus, our solutions are provably close (less than 3% on average) to the optimal SND solutions. Comparison of the method's performance across different problem scenarios demonstrates that its effectiveness is robust to variations in problem size and connectivity requirements.

The rest of this paper is organized as follows. Section 2 defines and formulates the SND problem and discusses the structural features of SND solutions that motivate our hierarchical approach. Section 3 provides a detailed description of our multi-stage method including the tabu search procedure, and Section 4 discusses our computational design and summarizes the results. Section 5 offers concluding remarks.

2. SND problem definition and solution structure

2.1. Problem formulation

Given an undirected network G: (N,E), with N and E respectively representing the set of nodes and available edges of the network, nonnegative costs c_{ij} for each edge in $\{i,j\}$ in E, and nonnegative integer connectivity parameters ρ_i for each node i, the SND problem seeks the minimum-cost set of edges that meets all the connectivity requirements. More important nodes that require greater level of protection of their communication paths have higher values of ρ_i. The connectivity parameters translate to the requirement that, for any pair of nodes i and j, the chosen network must contain at least $r_{ij} = \min(\rho_i, \rho_j)$ edge-disjoint paths between these two nodes. Nodes i with connectivity requirement $\rho_i = 0$ represent intermediate points that the network may optionally use to reduce the total cost of the network; however, we are not required to necessarily span these nodes. We refer to these locations as *Steiner* nodes. Nodes with connectivity parameter $\rho_i = 1$, called *regular* nodes, represent customers or locations that have minimal connectivity requirements. That is, the network design must span these nodes, but we only require one path between a regular node and every other node with positive connectivity parameter. Finally, we refer to nodes with $\rho_i \geqslant 2$ as *critical* nodes; each of these nodes requires protection in the form of two or more paths to every other critical node. By permitting nodes to have different connectivity parameters, the SND problem differentiates nodes in terms of their importance and protection requirements, and can exploit these differences to reduce the cost of the network.

To formulate the SND problem, we define a binary (0 or 1) variable u_{ij} for each edge $\{i, j\} \in E$ to indicate whether or not the SND solution includes this edge. One way to enforce the connectivity requirements is using cutset constraints (see, for example, Grötschel et al., 1995), one for each cutset defined by node partitions $\{T, N\backslash T\}$, $T \subset N$, that separate at least one node pair i, j with positive connectivity requirement r_{ij}. The constraint corresponding to this cutset specifies that the solution must select at least $max\{r_{ij}$: $i \in T$, $j \in N\backslash T\}$ edges from this cutset. Since the network has an exponential number of cutsets, the number of cutset constraints is also exponential; therefore, solving this model may require using a cutting plane procedure that dynamically adds violated cutset constraints. Alternatively, we can use the max-flow min-cut theorem to develop an equivalent flow-based formulation that has a polynomial number of constraints (see, for example, Magnanti and Raghavan, 2005). For this flow formulation, we select a node with the highest connectivity, say node 1, as the *root* node, and let H denote the set of non-Steiner nodes excluding the root node. For each $h \in H$, we define a commodity h with the root node as the origin node, node h as the destination node, and demand of ρ_h units. The flow formulation includes additional continuous flow variables f_{ij}^h to denote the flow of each commodity $h \in H$ on edge $\{i,j\}$ from node i to node j, and uses flow conservation constraints together with forcing constraints to impose the connectivity requirements. The two formulations, [*CUT*] and [*FLOW*], are presented below.

$$[CUT] \quad \min \sum_{\{i,j\}\in E} c_{ij}u_{ij}$$

$$\text{subject to} \quad \sum_{\{i,j\}\in\{T,N\backslash T\}} u_{ij} \geqslant \max r_{ij_{i\in T, j\in N\backslash T}}, \quad \forall \text{ cutsets } T$$

$$u_{ij} \in \{0,1\}, \quad \forall \{i,j\} \in E.$$