Designing Hierarchical Survivable Networks

A. Balakrishnan, T. L. Magnanti and P. Mirchandani

OR 291-94 January 1994

Designing Hierarchical Survivable Networks

Anantaram Balakrishnan*
Thomas L. Magnanti

Sloan School of Management
Massachusetts Institute of Technology
Cambridge, MA

Prakash Mirchandani†

Katz Graduate School of Business
University of Pittsburgh
Pittsburgh, PA

January 1994

* Supported in part by a grant from the AT&T Research Fund
† Supported in part by a Faculty Grant from the Katz Graduate School of Business, University of Pittsburgh

Abstract

As the computer, communication, and entertainment industries begin to integrate phone, cable, and video services and to invest in new technologies such as fiber optic cables, interruptions in service can cause considerable customer dissatisfaction and even be catastrophic. In this environment, network providers want to offer high levels of service—in both serviceability (e.g., high bandwidth) and survivability (failure protection)—and to segment their markets, providing better technology and more robust configurations to certain key customers. We study core models with three types of customers (critical, primary, and secondary) and two types of services/technologies (primary and secondary). The network must connect primary customers using primary (high bandwidth) services and, additionally, contain a back-up path connecting certain critical primary customers. Secondary customers require only single connectivity to other customers and can use either primary or secondary facilities. We propose a general multi-tier survivable network design model to configure cost effective networks for this type of market segmentation. When costs are triangular, we show how to optimally solve single-tier subproblems with two critical customers as a matroid intersection problem. We also propose and analyze the worst-case performance of tailored heuristics for several special cases of the two-tier model. Depending upon the particular problem setting, the heuristics have worst-case performance ratios ranging between 1.25 and 2.6. We also provide examples to show that the performance ratios for these heuristics are the best possible.

Introduction

Increasingly, survivability is becoming an important criterion in the design of telecommunication networks. Several recent developments have prompted this change. The first is technological: fiber-optic and opto-electronic cables are replacing traditional copper cables as a telecommunication medium. Because these newer technologies can carry substantially more traffic (both more channels and at a higher frequency) than traditional copper cables, telecommunication networks designed solely to minimize costs will tend to be sparse. In this case, the failure of a single edge can create significant system-wide disruptions, disabling traffic between many customer locations if the network does not provide alternate paths for routing. Second, customers, individual as well as industrial, are increasingly using telecommunication networks not only for transmitting voice, but also to transmit video and data. For example, in their logistics operations, many companies are now using Electronic Data Interchange (EDI) systems to connect suppliers and customers throughout the supply chain. EDI not only permits the immediate transmittal of sales and demand information between the different links in the supply chain, but also provides up-to-date inventory status throughout the chain. In addition, because EDI also provides automatic billing, monitoring of key marketing variables, and other advantages, companies have become quite dependent on their inter-organizational telecommunication networks for day-to-day operations. As yet another motivating factor, recently merged telecommunication and cable companies will be offering new entertainment services to their customers; this change has increased the reliance on communication networks connected to individual households. For all these reasons, and in all these contexts, network providers need to offer services that are highly reliable and that are robust to localized equipment (edge and node) failures.

Recent developments have brought about yet another change in telecommunications: network designers now have a choice of multiple transmission and switching technologies. For example, they can use twisted pair (copper), fiber optics, or opto-electonic transmission media, and add/drop multiplexers or digital cross-connect switches. Moreover, a particular physical technology such as fiber optic cables might be able to provide different types of service (such as DS1 or DS3). These technologies and services differ in their cost, reliability, and capacity. As a result, networks need to connect important customers using higher cost, but also more reliable and higher capacity switches and transmission media, while connecting less critical customers using less expensive, but also lower capacity equipment. This technology choice adds a new, and as yet only partially studied, dimension to the design of survivable networks.

The prevailing literature on network survivability (see, for example, Cornuéjols, Fonlupt, and Naddef [1985], Grötschel, Monma, and Stoer [1992], and Monma, Munson, and Pulleyblank [1990]) considers a single interconnection technology. These models represent survivability through node-connectivity requirements specifying the number of edge or node-disjoint paths required between every pair of nodes. The network must provide a larger number of edge-disjoint paths connecting more important node pairs.

Node-connectivity requirements of two or more provides one form of network reliability. Another recent stream of research in the network design literature attempts to provide reliable designs by using multiple interconnection technologies. Examples are the Hierarchical Network Design Problem (Current, Revelle, and Cohon [1986]) and the more general Multi-Level Network Design Problem (Balakrishnan, Magnanti, and Mirchandani [1992a]); this "serviceability" approach to network design provides higher grade (more reliable and more costly) service between certain "important" pairs of nodes, and lower grade service between other nodes. This approach does not incorporate multiple paths.

This paper aims to bring together these two disparate streams of research by viewing network reliability/survivability as a function of both node-connectivity and of the technology choices. We propose a multi-tier, multi-connected network design model that incorporates differential technologies as well as multiple connectivity requirements between certain node pairs in the network. The single-tier, multi-connected as well as the multi-tier, single-connected network design problems in the literature are special cases of this model.

In Section 1, we introduce a general model and describe various specializations and alternative modeling assumptions. We then recast the problem as an "overlay optimization problem," a class of models introduced by Balakrishnan, Magnanti, and Mirchandani [1994a] which has a "base" subproblem and an "overlay" subproblem(s); these subproblems are linked by the requirement that the overlay solution is "contained in" the base solution. Since multi-tier survivable network design problems can be modeled as special cases of the overlay optimization problem, as we show in Section 2, the heuristic worst-case results in Balakrishnan et al. apply directly. However, Sections 4 and 5 demonstrate that we can strengthen these results by using idiosyncratic problem characteristics.

The results in Sections 4 and 5 build upon heuristic and optimal methods for solving single-tier, multi-connected versions of the general multi-tier problem. We first examine the single-tier models in Section 3. In this discussion, we consider two basic problems: a dual path tree problem and a dual path Steiner tree problem. In the dual path tree (DPT) problem we seek a cost-minimizing network that connects all the nodes and has two edge-disjoint paths between two specified nodes. The dual path Steiner tree (DPST) problem is a Steiner tree version of the DPT problem; it contains a set of additional Steiner nodes that can (but need not) be used as intermediate nodes in the optimal design. We describe a heuristic method with a worst-case performance guarantee of 2 for both these problems. When the costs satisfy the triangle inequality, we can do better: using a matroid intersection algorithm we can optimally solve the DPT problem. We also provide an easily implemented "1-tree" heuristic with a worst-case performance guarantee of 3/2 for the DPT problem. We then consider a more general cost structure, called μ-direct, and show that in this case the 1-tree heuristic has a worst-case performance guarantee of $1+\mu/2$ ($\mu = 1$ for problems with triangular costs).

Sections 4 and 5 address various two-level, two-connected survivability models. In these problem settings, we can use either high grade or low grade transmission facilities. We need to connect certain primary nodes using only high grade paths, we can use any type of path to connect other, secondary nodes. In addition, the network design must include an alternative back-up transmission path between certain of the primary nodes. By making alternate assumptions concerning the nature of the back-up path (high grade or general) and by making assumptions about the number of primary nodes and their connectivity requirements, we obtain four different types of models. For each of these models, we develop two or more heuristic solution procedures and design a composite heuristic solution procedure that chooses the best of the individual heuristic solution values. We analyze the performance of this procedure for various cost structures. Our analysis shows that, depending upon the specific problem setting, the heuristic performance guarantees for the composite heuristic range from 1.25 to 2.6.

As we note in the conclusions (Section 6), the analysis in this paper extends to more general multi-tier, multi-connected problems. For example, we could require K instead of 2 paths between the special nodes, or we could consider models with K special nodes that must all lie on a common ring (and so have connectivity two). Recent SONET networks use this type of ring topology.

1. The Multi-tier Survivable Network Design Problem

Let G = (N, E) denote an undirected graph with node set N and edge set E. Let L denote the number of different technology (service) types, indexed from 1 to L; level $l = 1$ refers to the highest grade technology (e.g., fiber optic cables) and level $l = L$ corresponds to the lowest grade. A grade l facility on edge {i,j} costs c^l_{ij}, with $c^l_{ij} \geq c^{l'}_{ij}$ if $l < l'$. The Multi-tier Survivable Network Design (MTS) model represents survivability through L nonnegative connectivity parameters r^l_{ij}, for $l = 1, 2, ..., L$, defined for each pair i and j of nodes. The integer connectivity value r^l_{ij} ($= r^l_{ji}$) specifies the minimum required number of edge-disjoint paths connecting node i to node j containing facilities of service grade l or higher. Therefore, $r^l_{ij} \geq r^{l'}_{ij}$ if $l' < l$. Whenever each connectivity value r^l_{ij} equals 0, 1, or 2, we will say that the problem has low connectivity requirement; for most of this paper, we consider only low connectivity problems. As Grötschel et al. [1992] have noted, these models are relevant for designing contemporary telecommunication networks.

1.1 Multi-tier problem formulation

To formulate the multi-tier survivable design problem as an integer program, for any subset of nodes $S \subset N$ and $T = N\backslash S$, let {S,T} denote the edge-cutset defined by S and T, i.e., {S,T} includes all edges $\{i,j\} \in E$ with $i \in S$ and $j \in T$. Let u^l_{ij} equal 1 if we install a level-l facility on edge {i,j}, and equal 0 otherwise. Define $U^l_{S,T} = \sum_{\{i,j\} \in \{S,T\}} u^l_{ij}$, i.e., $U^l_{S,T}$ denotes the aggregate number of level-l facilities across the {S,T} cutset. Let $R^l_{S,T}$ denote the maximum level-l connectivity requirement across the {S,T} cutset, i.e.,

$$R^l_{S,T} = \max_{i \in S, j \in T} r^l_{ij}.$$

Using the facility design variables u, we can formulate the multi-tier, survivable network design problem as follows.

Problem [MTS]:

$$\text{minimize} \quad \sum_{1 \leq l \leq L} \sum_{\{i,j\} \in E} c^l_{ij}\, u^l_{ij} \tag{1.1}$$

subject to

$$\sum_{1 \leq l' \leq l} U^{l'}_{S,T} \;\geq\; R^l_{S,T} \qquad \text{for all } S \subset N,\ T = N\backslash S,\ 1 \leq l \leq L, \tag{1.2}$$

$$\sum_{1 \le l \le L} u_{ij}^{l} \le 1 \qquad \text{for all } \{i,j\} \in E, \text{ and} \tag{1.3}$$

$$u_{ij}^{l} = 0 \text{ or } 1 \qquad \text{for all } \{i,j\} \in E,\ 1 \le l \le L. \tag{1.4}$$

By Menger's theorem (Ford and Fulkerson [1962]), constraints (1.2) establish the connectivity requirement for each level of service. Constraints (1.3) ensure that we can install at most one facility on each edge. Formulation [MTS] uses only design variables u. Alternatively, we could introduce auxiliary flow variables and use these variables to establish the connectivity requirements. The flow-based formulation has more variables, but far fewer constraints. A directed version of this alternative model has proven to be very effective computationally for solving multi-tier, single-connected network design problems (Balakrishnan, Magnanti, and Mirchandani [1992b]).

Rather than using the intuitive formulation [MTS], the heuristic analysis presented in this paper is based on an alternative "overlay optimization" model (Balakrishnan et al. [1994a]) that has the following generic formulation. Let b^l for $1 \le l \le L$ be m-dimensional cost vectors with nonnegative elements b_{ij}^{l}. Let v^l for $1 \le l \le L$ be m-dimensional decision vectors with components v_{ij}^{l}. For all l, V^l denotes a set in Z_+^m satisfying the property that $V^{l+1} \subseteq V^l$. Consider the following L-level overlay optimization problem.

Problem [OOP]:

$$\text{Minimize} \sum_{1 \le l \le L} b^l v^l \tag{1.5}$$

subject to

$$v^l \in V^l \qquad \text{for all } 1 \le l \le L, \text{ and} \tag{1.6}$$

$$v^l \le v^{l+1} \qquad \text{for all } 1 \le l \le L-1. \tag{1.7}$$

Observe that the overlay optimization problem consists of L subproblems $v^l \in V^l$ along with *linking* constraints (1.7). These constraints specify that the solution to the l^{th} subproblem must be "overlayed" or embedded in the $(l+1)$-level solution. An alternative version of overlay model requires embedding the higher grade facilities on a common base (level L) design, i.e., this model replaces constraints (1.7) with $v^l \le v^L$ for all $1 \le l \le L-1$. Balakrishnan et al. [1994a] use this latter model to analyze the multicommodity uncapacitated network design problem.

To interpret [MTS] as an overlay optimization problem, let $b^{l}_{ij} = c^{l}_{ij} - c^{l+1}_{ij}$, with $c^{L+1}_{ij} = 0$, denote the **incremental** cost of installing a level l facility on edge {i,j}. The reformulation represents the decision to install a level l facility on edge {i,j} as the decision to first install a level L facility on {i,j} and then to successively upgrade this facility to level l', for $l' = L-1, L-2, ..., l$. The level l facility upgrading variable v^{l}_{ij} takes the value 1 if we upgrade the level $(l+1)$ facility on edge {i,j} to level l, and is 0 otherwise. For all l, let $v^{l}_{ij} = \sum_{1 \le l' \le l} u^{l'}_{ij}$ and $V^{l}_{S,T} = \sum_{1 \le l' \le l} U^{l'}_{S,T}$, and define the set

$$V^{l} = \{v = (v^{l}_{ij}): \text{each } v^{l}_{ij} \in Z_{+},\ v^{l}_{ij} \le 1, \text{ and } V^{l}_{S,T} \ge R^{l}_{S,T}, \text{ for all } S \subset N,\ T = N\backslash S, \}.$$

With this variable redefinition, formulation [MTS] is equivalent to formulation [OOP]. Note that $u^{l}_{ij} = v^{l}_{ij} - v^{l-1}_{ij}$ for $l = 2, 3, ..., L$; therefore, the nonnegativity restriction on the u variables becomes the linking constraints for the v variables.

The model [MTS] is deceptively simple; however, as shown in Figure 1, it includes as special cases many network design problems. The model can permit single or multiple grades for the transmission facilities and it allows single or multiple connectivities (for some or all nodes). We can further categorize multi-tier models depending upon the number of 1-connected and multi-connected nodes (i.e., nodes with connectivity requirement greater than 1) at each level and whether all the paths between the multi-connected nodes need to use the same grade paths. For two level problems, we refer to multi-connected nodes at the higher level as *critical* nodes. If the edge-disjoint paths connecting these nodes must all use the same (or higher) level paths, we say that the problem requires *full back-up*; otherwise, we say that it requires *partial back-up*. Providing full back-up between two nodes i and j ensures that the network can accommodate all traffic from i to j if a single link on the regular i-to-j path fails; networks with full back-up are expensive and in periods of normal operation have considerable underutilized high-grade capacity. To reduce network cost, planners might be satisfied with providing minimal communication capability (for critical traffic) when a link on the regular path fails. In this case, we provide partial back-up by permitting lower grade facilities on the back-up path.

Yet another distinction between multi-connected models concerns assumptions regarding *edge duplication*. Edge duplication permits us to install parallel transmission facilities between pairs of nodes, and to treat these parallel facilities as edge-disjoint for purposes of establishing back-up paths. Our model [MTS], unlike some other models in the network survivability literature, does not allow duplicated edges: it permits at most one facility to be installed on any edge (constraints (1.3)). As we will see and as might be expected, allowing duplicate edges simplifies the heuristic solution methods and improves their worst-case performance. Yet, survivability issues often dictate that we do not allow duplicated edges (for instance, often a single conduit carries parallel transmission lines and so if the conduit breaks, then so do all the lines in that conduit). Finally, while we consider undirected facilities in this paper, we could define multi-tier survivability problems for the directed case as well.

Figure 1: Hierarchy of Multi-tier, Multi-connected Network Design Problems

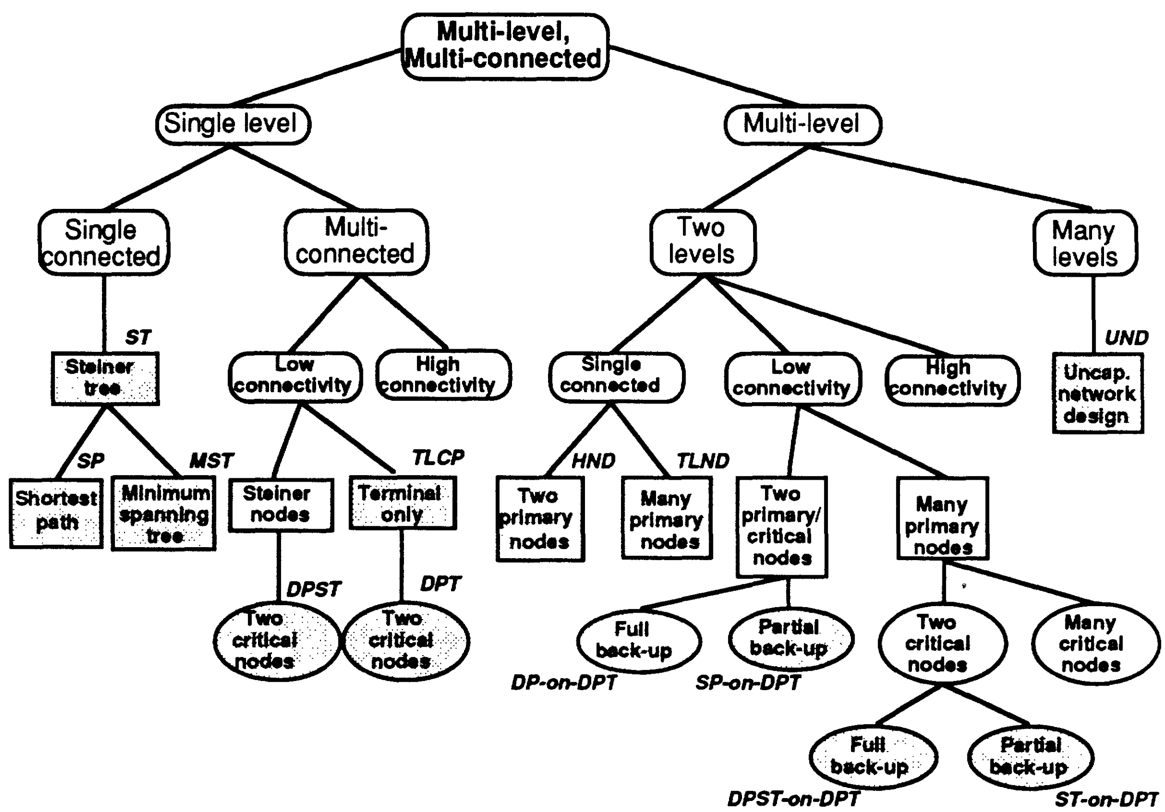

To summarize, the multi-tier, multi-connected network design framework covers a very broad range of models. Rather than applying a general solution method or developing general worst-case bounds that apply to all models, we might wish to exploit the structure of specialized models within this framework to sharpen the bounds and develop more effective solution methods. A taxonomy of multi-tier, multi-connected models might include the following items:

- the number of node levels and facility grades;
- the number of higher level nodes (for example, two or more than two);
- the number of multi-connected nodes at each level;

• the maximum connectivity level and type of redundancy (full or partial back-up);
• the use or non-use of duplicate edges; and,
• undirected or directed facilities.

In this paper, we focus on undirected, two-tier, low connectivity models without edge duplication, making distinctions between (i) two and many nodes at the high grade level, and (ii) full or partial back-up. At various points in our discussion, we comment on how to adapt our results to situations permitting edge duplication. To the best of our knowledge, this paper represents the first study of multi-tier, multi-connected network design optimization problems.

2. A Composite Heuristic for Two-tier Overlay Optimization Models

We first briefly review prior heuristic analysis for the two-tier overlay optimization problem (Balakrishnan et al. [1994a]). For notational simplicity, we use X instead of V^2 and Y instead of V^1. Similarly, we let x denote v^2 and y denote v^1, and use a and b respectively to denote the cost vectors b^2 and b^1. We let c denote the total cost (a+b). With these variable changes, the overlay version of the two-tier survivable network design problem has the following form:

Problem [TTS]:

$$Z^* = \text{minimize} \quad ax + by \tag{2.1}$$

subject to

Overlay constraints: $$y \in Y \tag{2.2}$$

Base constraints: $$x \in X \tag{2.3}$$

Linking constraints: $$y \leq x \tag{2.4}$$

We begin by noting that if we ignore the linking constraints (2.4) in formulation [TTS], then we obtain two subproblems, $Z_B(a) = \min\{ax: x \in X\}$, and $Z_O(b) = \min\{by: y \in Y\}$. We refer to these problems as the *base* and the *overlay subproblems.*

Since our model assumes that $X \subseteq Y$, we can generate feasible solutions to problem [TTS] by finding feasible solutions $x \in X$ to the base subproblem, and then setting y = x. If we choose x as a solution (approximate or optimal) of the base subproblem BP(c), using the total costs c, we refer to this method as the *Base Upgrading* (BU) heuristic. A complementary heuristic, which we call the *Overlay Completion* (OC) heuristic, first generates a feasible solution $\hat{y}$ to the overlay subproblem OP(c), using total costs, and then "completes" this overlay solution by solving the following *completion subproblem*: $Z_B(a,\hat{y}) = \min\{ax: x \geq \hat{y}, x \in X\}$. Since $x \geq \hat{y}$ and the cost vector a is nonnegative, the optimal value $Z_B(a,\hat{y})$ of the completion problem must be at least $a\hat{y}$. We refer to the difference $\delta(\hat{y}) = Z_B(a,\hat{y}) - a\hat{y}$ as the optimal completion cost. Our analysis applies to problem classes that satisfy the following condition: for any feasible problem instance and a given overlay solution $\hat{y}$, the optimal completion cost does not exceed $\lambda\, Z_B(a)$ for some finite known constant λ. We refer to this condition as the *feasible completion property*, and to λ as the *completion cost multiplier*. For three out of the four models that we analyze in this paper, $\lambda = 1$. Note that any problem that permits duplicate edges satisfies the feasible completion property with $\lambda = 1$ since we can get a feasible solution to the completion problem by setting $x_{ij} = \hat{y}_{ij} + 1$ for all the edges $\{i,j\}$ in the optimal solution to the base subproblem, and $x_{ij} = \hat{y}_{ij}$ for all other edges.

We analyze the worst-case performance of a *composite* heuristic that applies both the BU and OC heuristics to any given problem instance, and selects the solution with the smaller total cost. Let Z^{Comp} denote the cost of this solution. If we solve the base and overlay subproblems using heuristic methods with worst-case performance guarantees of ρ_B and ρ_O respectively, then the BU heuristic solution value is bounded from above by $\rho_B Z_B(c)$, and the OC heuristic solution value is bounded from above by $\rho_O Z_O(c)+\lambda\rho_B Z_B(a)$. Therefore, assuming $\lambda = 1$, the cost of the composite heuristic solution is bounded from above by

$$Z^{Comp} \leq \min\{\rho_B Z_B(c), \rho_O Z_O(c)+\rho_B Z_B(a)\}. \tag{2.5}$$

In the subsequent analysis, we let $\rho = \rho_O/\rho_B$.

2.1 General worst-case results for two-tier models

For the generic two-tier overlay optimization problem [TTS] satisfying the feasible completion property with $\lambda = 1$, Balakrishnan et al. [1994a] have characterized the worst-case performance ratio of the composite heuristic, that is, the maximum possible ratio between the objective value Z^{Comp} of the solution generated by the composite heuristic and the optimal value Z^* of problem [TTS]. They consider two cases: (i) problems for which total and base costs are *proportional*, i.e., $c_{ij}/a_{ij} = r$, a constant for all edges $\{i,j\}$, and (ii) the general case with *unrelated* total-to-base costs. The following two theorems summarize these prior results.

Theorem 1:

For overlay optimization problems with $\lambda = 1$ and proportional costs, the performance ratio ω_{prop} of the composite heuristic is bounded from above by

$$\omega_{prop} \leq \rho_B \frac{4}{4-\rho} \quad \text{if } \rho \leq 2, \tag{2.6a}$$

$$\leq \rho_B \, \rho \quad \text{if } \rho > 2. \tag{2.6b}$$

Theorem 2:

For overlay optimization problems with $\lambda = 1$ and unrelated costs, the worst-case performance ω_{unrel} of the composite heuristic is

$$\omega_{unrel} \leq \rho_O + \rho_B \quad \text{if } Z_B(a) > 0, \text{ and} \tag{2.7a}$$

$$\leq \rho_O \quad \text{if } Z_B(a) = 0. \tag{2.7b}$$

2.2 Applications of general worst-case results

To illustrate the use of these theorems, we consider two special cases of the two-tier network design problem: the *Hierarchical Network Design (HND)* and the *Two-Level Network Design (TLND)* problems. In both these problems, (i) every pair of nodes has a connectivity requirement of 1, and (ii) there are two service levels, primary and secondary, corresponding, for example, to fiber-optic cables and copper cables. In the HND problem, we designate two nodes of G, say nodes 1 and 2, as primary nodes, and refer to a path containing only primary facilities, as a primary path. The HND problem seeks a cost minimizing spanning tree that contains a primary path connecting nodes 1 and 2; the remaining edges of the tree have secondary facilities. The TLND problem generalizes the HND problem by designating more than two nodes as primary nodes; the solution must connect all the primary nodes to each other using primary paths. The optimal TLND solution is a cost minimizing spanning tree that contains a primary subtree (a Steiner tree) connecting all the primary nodes (as well, perhaps, as some secondary nodes); the remaining edges of the tree have secondary facilities.

For the HND problem, the base subproblem is a minimum spanning tree problem, and the overlay subproblem is a shortest path problem. Therefore, $\rho_O = \rho_B = \rho = 1$. So for problems with proportional costs, by Theorem 1 the cost of the composite heuristic solution is at most 4/3 rds the optimal cost. For the TLND problem, $\rho_B = 1$ and the minimum spanning tree (MST) heuristic (Takahashi and Matsuyama [1980]) solves the overlay (Steiner tree) subproblem with a worst-case ratio $\rho_O = 2$. Therefore, Theorem 1 implies that the worst-case ratio of the composite heuristic for TLND problems with proportional costs must not exceed 2. These worst-case bounds for the composite heuristic for HND and TLND problems are tight (Balakrishnan et al. [1992a]).

In Sections 4 and 5, we show that by exploiting special problem structure we can improve upon the worst-case bounds of Theorems 1 and 2 for several two-tier, two-connected network design models. For example, for one proportional cost model that we consider in Section 4.1, $\rho_O = 3/2$ and $\rho_B = 1$ and so Theorem 1 provides the bound $\omega_{prop} \leq 9/5$ whereas the bound we obtain has an improved worst-case performance guarantee of 8/5. In another instance, we are able to reduce the bound from 4/3 to 5/4.

2.3 Heuristic analysis strategy

Theorems 1 and 2 and our worst-case analysis in Sections 4 and 5 use the following general approach. The analysis begins with the upper bound (2.5) on the cost of the composite heuristic. This bound depends on the costs of the BU and OC heuristic solutions. For each specialized model that we consider, we attempt to improve the BU and OC heuristics and obtain sharper estimates of their costs. We also determine a lower bound on the optimal value Z^* as follows. If we ignore the linking constraints (2.4) in formulation [TTS], as we noted previously, the problem decomposes into the overlay subproblem with costs b and the base subproblem with costs a. Consequently, the sum of the optimal values for these two subproblems is a valid lower bound on Z^*. We obtain another lower bound by ignoring the base constraints (2.3). Since all costs are nonnegative, setting x = y is optimal for this relaxation, and so the optimal value of the relaxation is $Z_O(c)$. Combining these two lower bounds shows that

$$Z^* \geq \max\{Z_O(b) + Z_B(a), Z_O(c)\}. \qquad (2.8)$$

Dividing the heuristic upper bound (2.5) by the lower bound (2.8) gives an upper bound on the heuristic worst-case performance ratio. For the proportional costs case, we express this ratio in terms of two parameters—the cost ratio r and the unknown ratio s = $Z_O(a)/Z_B(a)$ (we assume $Z_B(a) > 0$). To obtain a data-independent performance characterization, we maximize the performance ratio with respect to s and r.

3. Solution Methods and Analysis for Underlying Single-tier Models

Sections 4 and 5 analyze two-tier versions of low connectivity network design models. These models have two new single-tier models—the dual path tree problem and the dual path Steiner tree problem—as their base and overlay subproblems. In this section, we study solution methods for these two single-level problems. This analysis will provide the values of the worst-case parameters ρ_O and ρ_B that we require for our subsequent two-level analysis.

Before beginning our analysis, let us introduce some terminology and briefly review relevant prior results. By *triangularizing* an undirected graph G = (N,E) with costs a_{ij} for all edges $\{i,j\} \in E$ we mean constructing a complete graph G' = (N,E') with edge costs a'_{ij} for all i, j ∈ N equal to the shortest path distance from node i to node j in G. We refer to G' as the triangularized graph and the costs a'_{ij} as triangularized costs. When we consider edge duplication, we will rely on the following property proved by Goemans and Bertsimas [1993] for single-tier survivable network design (SND) problems: the optimal value of the SND problem defined over the triangularized graph G' (with edge duplication permitted in this graph as well) is the same as the optimal value over the original graph G. We will refer to this property as the *duplication equivalence property.* To construct a feasible SND solution over the original graph G from a feasible solution over G', we replace each edge {i,j} in the latter solution with the edges of the shortest i-to-j path in G (with replications if an edge in G appears in more than one such shortest path). We refer to the resulting solution to the original problem as the *recovered* solution.

Dual Path Steiner Tree (DPST) problem:

Given an undirected graph G=(N,E) with nonnegative edge costs a_{ij}, and a subset $P \subseteq N$ of primary nodes containing two critical nodes 1 and 2, find the minimum cost subgraph that spans all the nodes of P via optional "Steiner" nodes from N\P, and that connects nodes 1 and 2 via two edge-disjoint paths.

In terms of the terminology we introduced for the general MTS model, the DPST problem has L = 1, and $r^1_{ij} = 1$ for all node pairs i and j ∈ P except $r^1_{12} = r^1_{21} = 2$, and $r^1_{ij} = 0$ if i or j $\notin$ P. The **Dual Path Tree (DPT) problem** is a special case of the DPST model with P = N, i.e., the solution must span all the nodes of graph G.

The DPST problem is NP-hard since it generalizes the Steiner network problem. As we will show later, if we assume triangular costs then the DPT problem is polynomially solvable. For DPT and DPST problems with arbitrary edge costs a_{ij}, Balakrishnan, Magnanti and Mirchandani [1994b] propose the following efficient *dual path greedy completion (DPGC)* heuristic. Using a graph doubling argument, they show that the DPGC method solves the DPST and DPT problems with a worst-case performance guarantee of 2. This bound holds for the problems with or without edge duplication.

Dual Path Greedy Completion (DPGC) heuristic:

Step 1: Find the minimum cost pair of edge-disjoint paths from node 1 to node 2. Let E_1 and N_1 be the subset of edges and nodes belonging to these paths.

Step 2: Contract the subgraph $G_1=(N_1,E_1)$ into a single node 0, triangularize the resulting graph, and eliminate all the Steiner nodes not in N_1 and their incident edges, creating a graph G*. Find the minimum spanning tree of G*. Recover the original edges corresponding to the edges of this spanning tree and add one copy of each recovered edge to E_1 to obtain a feasible DPST solution.

The method derives its name from the operations of first finding the optimal "dual paths" (in Step 1) and then completing this solution in a greedy fashion (Step 2).

If we do not permit edge duplication, then as is well-known we can find the optimal dual paths in Step 1 by solving a minimum cost network flow problem defined on the following network. The network contains all the nodes and edges of G. Node 1 has a supply of 2 units, node 2 has a demand of 2 units, and all other nodes are transshipment nodes. The flow cost on each edge $\{i,j\}$ is the original edge cost a_{ij}, and every edge has a capacity of 1 unit. The minimum cost flow solution routes 1 unit of flow on each of the two required edge-disjoint 1-to-2 paths. When we permit edge duplication, the optimal dual path solution consists of two copies of the shortest 1-to-2 path.

When the edge costs have special properties, can we develop alternative solution methods that have better worst-case performance than the DPGC method? For the DPT problem, we can indeed develop more effective methods. In particular, when the edge costs satisfy the triangle inequality, as we show in Section 3.1, the DPT problem is polynomially solvable using a matroid intersection algorithm. For a broader class of cost structures that we call μ-direct costs, Section 3.2 describes and analyzes the worst-case performance of a simple 1-tree heuristic that is more effective than the DPGC method for a range of μ values. The models considered in both Sections 3.1 and 3.2 prohibit edge duplication; Section 3.3 discusses algorithmic and worst-case implications for models that permit edge duplication.

3.1 Dual path trees for graphs with triangular costs

DPT problems with triangular costs are polynomially solvable. To establish this result, we use the following property.

Proposition 3:

If the edge costs satisfy the triangle inequality, then the DPT problem has an optimal solution containing exactly $|N|$ edges.

Proof:

The optimal solution to the DPT problem spans all the nodes in the graph and contains two edge-disjoint paths, say P_1 and P_2, connecting the primary nodes 1 and 2. Because the costs are nonnegative, we can choose both P_1 and P_2 as simple paths (they do not revisit nodes). If the paths P_1 and P_2 intersect only at nodes 1 and 2, then the optimal solution spans all nodes and contains exactly one cycle, and thus contains exactly $|N|$ edges.

Next suppose that the paths P_1 and P_2 intersect at some intermediate node(s) other than nodes 1 and 2. Let us orient these paths from node 1 to node 2; that is, node 1 is their first node and node 2 their last node. If paths P_1 and P_2 intersect at more than one intermediate node, let a be the first intersection point (after node 1) on P_1, and let b be the first intersection point on P_2. (Nodes a and b might be the same node). First, observe that nodes b and a cannot simultaneously be (immediate) successors of each other on paths P_1 and P_2, since then both paths would contain the edge $\{a,b\}$, contradicting the fact that P_1 and P_2 are edge-disjoint. So, suppose that the node b is not the successor of node a on path P_1. Let i and j denote the predecessor and successor of node a on path P_1. In path P_1 replace the edges $\{i,a\}$ and $\{a,j\}$ with the edge $\{i,j\}$; the triangle inequality implies that the cost of the resulting path P_3 does not exceed the cost of path P_1.

Now note that if node a's predecessor is node $i \neq 1$, then since node a is the first intersection node on path P_1, $i \notin P_2$ and so $\{i,j\} \notin P_2$. If $i = 1$, the definition of node b as the first intersection node on path P_2 and the fact that $b \neq j$ implies that $(i,j) \notin P_2$. In either case, the paths P_2 and P_3 are edge disjoint. Moreover, by our previous observation these two paths cost no more than the two paths P_1 and P_2. Therefore, we have found another optimal solution to the DPT problem with one less node in common to the two paths.

Repeatedly identifying nodes a and b allows us to short-circuit one of the two paths. Since each path contains a finite number of nodes, this constructive procedure terminates when the two paths intersect at only nodes 1 and 2. ♦

Matroid Intersection algorithm:

We now show that the dual path tree is the intersection of two matroids. A **1-tree** of a graph G is the union of a spanning tree and one edge not in the spanning tree. Clearly, a 1-tree contains exactly one cycle. A **q-restricted 1-tree** is a 1-tree with the property that the unique cycle formed by the additional edge contains a particular node q of the graph. We can interpret a dual path tree with exactly $|N|$ edges as the intersection of a 1-restricted 1-tree and a 2-restricted 1-tree. Subsets of q-restricted 1-trees form a matroid (see Exercise 13.39 in Ahuja, Magnanti, and Orlin [1993]). Since the weighted matroid intersection problem is solvable in polynomial time (Edmonds [1979]), Proposition 3 implies that we can optimally solve the DPT problem with triangular secondary costs in polynomial time. We have thus established the following result.

Theorem 4:

If the edge costs satisfy the triangle inequality, then a weighted matroid intersection algorithm solves the DPT problem optimally in polynomial time.

For DPST problems with triangular costs, suppose we use the corresponding optimal DPT solution over the primary nodes as a heuristic solution. Can we characterize the worst-case performance of this DPST heuristic method? Balakrishnan et al. [1994b] have shown that for any low connectivity Steiner problem with triangular costs, the heuristic solution obtained by optimally solving the corresponding low connectivity problem over the terminal nodes costs at most twice the original optimal value, and this bound is tight. This result implies that the matroid intersection-based heuristic for triangular cost DPST problems has a worst-case performance of 2, which is the same as the worst-case performance of the more general and simpler DPGC heuristic.

Although polynomial, the generic matroid intersection algorithm is complex and is typically difficult to implement (its specialization for the dual path tree problem might be much easier though). As an alternative, we might wish to use a simple heuristic method for solving the DPT problem even when the costs are triangular. In the Section 3.2, we develop one such heuristic in the context of a broader class of graphs than those with triangular costs.

3.2 Dual path trees for μ-direct graphs

Whenever the graph G contains the edge {1,2}, this edge can potentially serve as one of the two edge-disjoint 1-to-2 paths. Therefore, if we do not permit edge duplication and

G has a feasible dual path tree, then if we start with edge {1,2}, the problem must have a feasible completion, i.e., the residual graph obtained by deleting edge {1,2} must contain a 1-to-2 path. (When we permit edge duplication and G is connected, we can always complete any given 1-to-2 path.) This observation motivates the following 1-tree heuristic. We state and analyze this $O(|E| + |N| \log|N|)$ heuristic in its general form, which is capable of solving both DPT and DPST problems.

1-Tree Heuristic:

<u>Step 1:</u> Remove the direct edge {1,2} from G and find an approximate or optimal solution STREE to the Steiner tree problem STP spanning all the primary nodes (with optional intermediate Steiner nodes) on the resulting residual graph G_{12}.

<u>Step 2:</u> Add edge {1,2} to STREE to obtain the 1-tree heuristic solution to the DPST problem.

When applied to the DPT problem, Step 1 merely requires finding the minimum spanning tree of G_{12}. In order to bound the worst-case performance of this heuristic, we need to be able to bound (i) the cost of the Steiner tree it produces, and (ii) bound the cost of edge {1,2} relative to the rest of the network. For (i), we let ρ_{ST} denote the worst-case ratio of the method we use to find the Steiner tree STREE. For (ii), we restrict our attention to a special class of graphs that we call **μ-direct**: these are graphs that: (a) have nonnegative edge costs, (b) contain edge {1,2}, (c) contain a path connecting nodes 1 and 2 without edge {1,2}, and (d) satisfy the property that the *cost a_{12} of the edge {1,2} is no more than μ (≥ 0) times the cost A_{12} of the shortest 1-to-2 path when we remove edge {1,2}*. This assumption implies that any DPT and DPST solution that does not contain the edge {1,2} costs at least $2A_{12} \geq 2a_{12}/\mu$. Note that triangular graphs are μ-direct graphs with $\mu = 1$. In stating the following worst-case result for the 1-tree heuristic, we let $\hat{\mu} = \max\{\mu, 1\}$.

Proposition 5:

Let $s = \min\{a_{12}, A_{12}\}/Z_{DPST} \leq 1/2$ be the cost of the shortest path from node 1 to node 2 relative to the optimal cost of the DPST problem. For μ-direct graphs, the 1-tree heuristic generates a solution to the DPST problem with a worst-case bound of at most $\rho_{ST} + \min\{\mu/2, \hat{\mu}s\}$. If any optimal DPST solution contains edge {1,2}, then the 1-tree solution has a worst-case bound of at most ρ_{ST}.

Proof:

We claim that the optimal value Z_{DPST} of the DPST problem is no less than the optimal value Z_{STP} of the Steiner tree problem STP that we solve in Step 1 of the 1-tree heuristic. Let Q* be an optimal DPST solution. Let Q' be any Steiner tree formed by dropping an edge from the cycle in Q* containing nodes 1 and 2, choosing edge {1,2} if Q* contains this edge. Since Q' is a feasible solution to the Steiner tree problem STP, its cost Z(Q') is greater than or equal to Z_{STP}. Therefore, the cost Z(STREE) of the exact or approximate Steiner tree STREE satisfies the following inequalities

$$Z_{DPST} \geq Z(Q') \geq Z_{STP} \geq \frac{Z(STREE)}{\rho_{ST}}. \tag{3.1}$$

If some optimal DPST solution OS contains edge {1,2}, then this edge plus the Steiner tree solution on G^{12} obtained by removing this edge from OS solves the DSPT problem. Therefore, since $\rho_{ST} \geq 1$ and $a_{12} \geq 0$,

$$Z_{DPST} \geq Z_{STP} + a_{12} \geq \frac{Z(STREE)}{\rho_{ST}} + a_{12} \geq \frac{Z(STREE)}{\rho_{ST}} + \frac{a_{12}}{\rho_{ST}}.$$

But , if $Z^{1\text{-tree}}$ denotes the cost of the 1-tree heuristic solution, then

$$Z^{1\text{-tree}} = Z(STREE) + a_{12} \leq \rho_{ST} Z_{DPST},$$

which is the last conclusion of the proposition.

If no optimal DPST solution contains edge {1,2}, then since the graph is μ-direct and the costs are nonnegative, $Z_{DPST} \geq 2A_{12} \geq 2a_{12}/\mu$. Combined with expression (3.1), this inequality implies that the cost $Z^{1\text{-tree}}$ of the 1-Steiner tree heuristic solution is bounded as follows:

$$Z^{1\text{-tree}} = Z(STREE) + a_{12} \leq \rho_{ST} Z_{DPST} + \frac{\mu}{2} Z_{DPST} = (\rho_{ST} + \frac{\mu}{2}) Z_{DPST}. \tag{3.2}$$

If edge {1,2} isn't a shortest 1-to-2 path (i..e., $\mu > 1$), then $sZ_{DPST} = A_{12} \geq a_{12}/\mu$. In this case, the previous inequality becomes

$$Z^{1\text{-tree}} = Z(STREE) + a_{12} \leq \rho_{ST} Z_{DPST} + \mu s Z_{DPST} = (\rho_{ST} + \mu s) Z_{DPST}. \tag{3.3}$$

If edge {1,2} is the shortest 1-to-2 path (i.e. $\mu \leq 1$), then

$$Z^{1\text{-tree}} = Z(STREE) + a_{12} \leq \rho_{ST} Z_{DPST} + s Z_{DPST} = (\rho_{ST} + s) Z_{DPST}. \tag{3.4}$$

The inequalities (3.2), (3.3), and (3.4) imply that

$$Z^{1\text{-tree}} \leq \rho_{ST} + \min \{\frac{\mu}{2}, \max\{\mu s, s\}\} Z_{DPST} = \rho_{ST} + \min \{\frac{\mu}{2}, s \max\{\mu, 1\}\} Z_{DPST}. \blacklozenge$$

When we apply the 1-tree heuristic to the DPT problem, $\rho_{ST} = 1$ and so we obtain the following corollary.

Corollary 6:

For DPT problems defined on μ-direct graphs, the 1-tree heuristic generates a solution with a worst-case bound of at most $1 + \min\{\mu/2, \hat{\mu}s\} \leq 1 + \mu/2$.

The following example shows that the bound in Corollary 6 is tight. Consider a μ-direct network containing 3 paths from node 1 to node 2: a direct path costing μ, and two alternate unit-cost paths, each containing q "short" edges (every edge on these two paths has a cost of $1/q$). If $1/q < \mu$, then the optimal solution is the two non-direct paths at a cost of 2; the 1-tree heuristic solution uses all but one of the short edges and costs $2 + \mu - 1/q$. Therefore, the ratio of the 1-tree solution's cost to the optimal value is $1 + \mu/2 - 1/2q$, which approaches $1+\mu/2$ as q approaches infinity.

Note that for solving DPST problems, the 1-tree heuristic is not competitive (in terms of worst-case performance) with the DPGC method unless we solve the Steiner problem optimally, or we know that the optimal DPST solution contains edge $\{1,2\}$, or we use Berman and Ramaiyer's [1992] heuristic (with $\rho_{ST} = 16/9$) to solve the Steiner tree problem and $\mu \leq 4/9$. Also, if $\mu \geq 2$, then the DPGC method is superior even for DPT problems defined on μ-direct graphs. In subsequent sections we assume, for convenience, that we always apply the DPGC method to approximately solve the DPST problem.

3.3 Dual path trees with edge duplication

When we permit edge duplication, we can solve the DPT problem with arbitrary costs in polynomial time using a modified matroid intersection algorithm. The following proposition proves a special property of an optimal DPST solution with edge duplication in triangularized graphs, enabling us to extend the matroid intersection algorithm to the edge duplication case as well.

Proposition 7:

The DPST problem with edge duplication has an optimal solution that either chooses edge $\{1,2\}$ twice or contains only unduplicated edges.

Proof:

By the duplication equivalence property, assume the costs are triangular. Consider an optimal solution that duplicates some edge $\{i,j\} \neq \{1,2\}$. If edge $\{i,j\}$ does not belong to both the 1-to-2 paths in the dual path tree, we can delete one copy of this edge from the solution. Otherwise, we can short-circuit either node i or $j \neq 1, 2$ on one of the paths, obtaining a feasible solution with equal or lower costs and using fewer edges in the 1-to-2 paths. Repeating this procedure for all duplicated edges $\{i,j\} \neq \{1,2\}$ provides the required optimal DPST solution. ♦

For the DPT problem with edge duplication, consider a solution satisfying the conditions of Proposition 7. If this solution contains edge $\{1,2\}$ twice, then the remaining edges must be spanning tree edges, and so the solution contains exactly n edges. Otherwise, the solution must be optimal for the "unduplicated" version of the problem, i.e., for the triangularized DPT problem without edge duplication. Proposition 3 shows that this unduplicated problem has an optimal solution containing exactly n edges. So we have the following corollary to Proposition 7.

Corollary 8:

The DPT problem with edge duplication has an optimal solution containing exactly n edges.

We can exploit this property to optimally solve the DPT problem with *general* (nonnegative) costs and with edge duplication as follows:

Edge-duplicating Matroid Intersection algorithm:

Triangularize the given graph G and add one parallel copy of edge $\{1,2\}$ to form a new augmented graph G". Apply the matroid intersection algorithm (without edge duplication) to this augmented graph. Recover the solution to the original graph G.

As we noted before, solving the DPT problem optimally over the primary nodes provides a heuristic DPST solution that costs at most twice the optimal DPST value. Consequently, for DPST problems with arbitrary costs but edge duplication, we can apply the Edge-duplicating Matroid Intersection algorithm to the corresponding DPT problem to obtain a solution with the same worst-case guarantee of 2 as the DPGC heuristic.

Notice that we could replace the use of the matroid intersection algorithm for solving the triangularized version of the problem in the augmented graph G" by any heuristic method that applies to triangular cost DPT problems with edge duplication. In particular, we could apply the 1-tree heuristic with the following adaptation: in Step 1 of the method, we delete only one copy of edge {1,2} before finding the optimal or approximate Steiner tree. Equivalently, in Step 1 we find the optimal or approximate Steiner tree for the triangularized graph G', and then add edge {1,2} (a parallel copy if this edge already exists in the Steiner tree) in Step 2. It is easy to adapt our prior analysis to show that this *Edge-duplicating 1-tree* heuristic has the same worst-case bounds as the original version of the 1-tree heuristic (see Proposition 5 and Corollary 6). In particular, for DPT problems with edge duplication, the edge-duplicating 1-tree heuristic produces a solution that costs at most 1.5 times the optimal value.

Goemans and Bertsimas [1993] proposed a "tree heuristic" to solve a large class of single-level, survivable network design problems with edge duplication. When specialized to the DPT or DPST problems, this heuristic is the same as the Edge-duplicating 1-tree heuristic (assuming that to solve the DPST problem this heuristic applies the MST heuristic to solve the Steiner subproblem in Step 1). Goemans and Bertsimas showed that the tree heuristic generates a solution that costs no more than twice the optimal value of the survivable network design problem. Relative to this bound, Corollary 6 provides a tighter bound of 3/2 for DPT problems, but for DPST problems Proposition 5 implies a weaker bound of 5/2. Our DPGC heuristic achieves the bound of 2 for both DPT and DPST problems with general costs, with or without edge duplication.

This section has examined the DPT and DPST single-tier subproblems of the two-tier models that we study next. The general purpose DPGC algorithm finds a heuristic DPST or DPT solution within a factor of 2 of the optimal value. For DPT problems defined over triangular and μ-direct graphs, the matroid intersection method and the 1-tree heuristic solve the problem optimally or within a factor of $1+\mu/2$. As we have shown, for situations that permit edge duplication, we can optimally solve DPT problems with arbitrary costs using matroid intersection.

4. Heuristic Analysis of Cycle+Tree Problems with Full Back-up

This section and Section 5 examine full and partial back-up versions of two-connected generalizations of the hierarchical network design and two-level network design problems.

In these two-tier models, the undirected graph G=(N,E) has a subset P of primary or high-level nodes that must be interconnected via primary paths (with optional intermediate secondary nodes). Furthermore, two *critical* nodes in P, nodes 1 and 2, require mutual two-connectivity, and the design must span all the remaining secondary nodes using secondary facilities. We refer to this class of models as Cycle+Tree models since the required network configuration consists of a tree plus edges of a cycle.

In the full back-up versions of these two-tier survivable network design problems, the two edge-disjoint paths connecting the critical nodes must both contain only primary facilities. Thus, the connectivity requirements of this model are: (i) at the primary service level: $r_{ij}^1 = 1$ if i and j $\in$ P, except $r_{12}^1 = r_{21}^1 = 2$; and (ii) at the secondary service level: $r_{12}^2 = r_{21}^2 = 2$ and $r_{ij}^2 = 1$ otherwise. For the general version of this problem, the overlay subproblem is the dual path Steiner tree (DPST) model and the base subproblem is a dual dath tree (DPT) problem; we therefore refer to this model as the *DPST-on-DPT* model. We also consider the *DP-on-DPT* special case containing only two primary nodes, both of which are critical. This model has the polynomially solvable dual path (DP) problem as its overlay subproblem.

Using Section 3's algorithms and worst-case results for DPT and DPST problems, we develop specialized worst-case bounds for the composite heuristic for the DPST-on-DPT and DP-on-DPT problems; these results improve upon the bounds we would obtain by applying general overlay results in Section 2 .

4.1 The DPST-on-DPT problem

For the DPST-on-DPT problem, the BU heuristic that we described in Section 2 solves the DPT problem using primary edge costs, and installs primary facilities on all the edges of this solution. From our discussions in Section 3, we note that the worst-case ratio ρ_B of the embedded procedure to solve the base (DPT) subproblem depends on the problem's cost structure (triangular, μ-direct, or general) and on the method that we apply (matroid intersection, 1-tree heuristic, or DPGC heuristic).

The OC heuristic solves the DPST problem using primary costs, and completes this solution by adding edges (containing secondary facilities) in order of increasing secondary costs to connect the remaining secondary nodes. As for the base subproblem, the worst-case ratio ρ_O of the embedded procedure to solve the overlay subproblem depends on the cost structure and the method that we apply. The total secondary cost of all the edges that the OC heuristic adds to complete any overlay (DPST) solution does not exceed the cost of the minimum spanning tree of G using secondary costs, which does not exceed the optimal value $Z_B(a)$ of the base (DPT) subproblem. Therefore, the DPST-on-DPT problem satisfies the feasible completion property with a completion cost multiplier $\lambda = 1$.

Since the OC heuristic's greedy completion step (adding edges in increasing order of secondary cost to span the nodes not in the DPST solution) incurs a completion cost of at most $Z_B(a)$, the OC heuristic solution costs at most $Z_O(c)+Z_B(a)$. In contrast, the analysis of general overlay optimization problems in Section 2 assumes that we solve the completion subproblem using the same heuristic that we use to solve the base subproblem. So, the general completion procedure can add a completion cost of up to $\rho_B Z_B(a)$, giving the looser OC upper bound $Z_O(c)+\rho_B Z_B(a)$ (see inequality (2.5)). Since the BU heuristic costs at most $\rho_B Z_B(c)$, the composite solution to the DPST-on-DPT problem costs no more than $\min\{\rho_B Z_B(c), Z_O(c)+Z_B(a)\}$. This observation and the lower bound (2.8) imply the following bound on the composite heuristic's worst-case ratio for DPST-on-DPT problems:

$$\frac{Z^{Comp}}{Z^*} \leq \frac{\min\{\rho_B Z_B(c), \rho_O Z_O(c)+Z_B(a)\}}{\max\{Z_O(c), Z_O(b)+Z_B(a)\}}. \tag{4.1}$$

For the *proportional costs* case, if s denotes the optimal secondary cost of the overlay subproblem relative to the optimal base subproblem value $Z_B(a)$, then inequality (4.1) reduces to

$$\omega_{prop} \leq \frac{\min\{\rho_B r, \rho_O rs+1\}}{\{(r-1)s+1\}}. \tag{4.2}$$

We must select values of s and r that maximize the right-hand side of inequality (4.2) subject to the constraints that $0 \leq s \leq 1$ and $r \geq 1$. If we ignore the restrictions on s, the right-hand side of (4.2) achieves its maximum when

$$s^* = \frac{\rho_B r - 1}{\rho_O r}. \tag{4.3}$$

Notice that since ρ_B and r are both greater than or equal to 1, s* is nonnegative. Furthermore, for both triangular and arbitrary costs, if we use our single-level heuristics from Section 3 or if we optimally solve the overlay and base subproblems, then $1 \leq \rho_B \leq \rho_O$. Therefore, the value of s* given by equation (4.3) is always less than or equal to 1 whenever the value of r is at least 1.

Substituting this value of s* in (4.2), we obtain

$$\omega_{prop} \leq \frac{\rho_B \, \rho_O \, r^2}{(1 - r[\rho_B+1-\rho_O] + \rho_B r^2)}. \tag{4.4}$$

Let us now consider two cases:

Case 1: $0 < \rho_B+1-\rho_O \leq 2$.

Differentiating the right-hand side of (4.4) with respect to r, we find that this expression achieves its maximum when

$$r^* = \frac{2}{(\rho_B+1-\rho_O)}. \tag{4.5}$$

Since $0 < \rho_B+1-\rho_O \leq 2$, this value of r* satisfies the requirement $r \geq 1$. Substituting for r* in (4.4), we obtain

$$\omega_{prop} \leq \frac{4 \, \rho_B \, \rho_O}{\rho_B \, (2+2\rho_O-\rho_B) - (\rho_O-1)^2} \; . \tag{4.6a}$$

Case 2: $\rho_B+1-\rho_O \leq 0$

In this case, the right-hand side of inequality (4.4) increases with r, and achieves its maximum value when $r^* = \infty$. Therefore,

$$\omega_{prop} \leq \rho_O. \tag{4.6b}$$

We do not consider the third case when $\rho_B+1-\rho_O > 2$ because, as we discuss next, this case does not apply when we use our single-level heuristics to solve the overlay and base subproblems.

Let us now consider various possible combinations of overlay and base solution methods. We can either solve the DPT base subproblem optimally (using the matroid intersection algorithm if costs are triangular) or apply the 1-tree or DPGC heuristics for problems with triangular or arbitrary costs. For the DPST overlay subproblem, we consider the options of solving it optimally (using, say, branch-and-bound) or approximately using the DPGC heuristic. Table I lists the resulting combinations of overlay and base solution methods, the corresponding base and overlay heuristic worst-case ratios, and the composite heuristic's performance ratio for proportional cost problems. To keep the discussions simple, we do not consider DPST-on-DPT problems defined on general μ-direct graphs, but instead limit our attention to the special case of triangular costs.

Table I: Proportional costs DPST-on-DPT solution options

Solution strategy identifier	Base (DPT) solution method	Overlay (DPST) solution method	ρ_B	ρ_O	ω_{prop}
DPST-on-DPT with triangular proportional costs					
T1	Optimal (Matroid Intersection)	Optimal	1	1	4/3
T2	Matroid Intersection	Dual Path Greedy Completion (DPGC)	1	2	2
T3	1-tree	DPGC	3/2	2	48/23
DPST-on-DPT with arbitrary proportional costs					
A1	Optimal	Optimal	1	1	4/3
A2	DPGC	DPGC	2	2	16/7

For solution strategies T1 and A1, $(\rho_B+1-\rho_O) = 1$, and so the bound (4.6a) applies. Since $\rho_B = \rho_O = 1$ for both these strategies, the composite heuristic has a worst-case ratio of 4/3, which is the same bound as in Theorem 1. For strategy T2 as well, Theorem 1 gives the same worst-case bound of 2 as (4.6b). However, for strategies T3 and A2, Theorem 1's bounds are 9/4 and 8/3 while inequality (4.6b) gives better bounds of 48/23 and 16/7. Note that the bounds in Theorem 1 apply to problems without edge duplication. If we permit edge duplication, then the DPT subproblem can be solved optimally even when costs do not satisfy the triangle inequality. Therefore, we might be interested in solving such problems using strategy T2 or T3 instead of strategy A2.

Consider the *unrelated costs* case. Using the OC heuristic, we obtain,

$$\omega_{unrel} \leq \frac{\rho_O Z_O(a)+Z_B(a)}{\max\{Z_O(c), Z_O(b)+Z_B(a)\}}$$

$$\leq \rho_O + 1. \qquad (4.7)$$

Notice that, unlike Theorem 2, the worst-case ratio for the unrelated costs case does not depend on the performance of the base heuristic.

Theorem 9:

For the DPST-on-DPT problem, the worst-case performance ratios ω_{prop} and ω_{unrel} corresponding to problems with proportional and unrelated costs are bounded from above as follows:

$$\omega_{prop} \le \frac{4\,\rho_B\,\rho_O}{\rho_B\,(2+2\rho_O-\rho_B) - (\rho_O-1)^2} \qquad \text{if } 0 < \rho_B+1-\rho_O \le 2, \quad (4.8a)$$

$$\le \rho_O \qquad \text{if } \rho_B+1-\rho_O \le 0; \text{ and}, \quad (4.8b)$$

$$\omega_{unrel} \le \rho_O + 1. \quad (4.8c)$$

4.2 The DP-on-DPT problem

For the DP-on-DPT special case with only two primary (and critical) nodes, the overlay subproblem seeks the minimum cost pair of edge-disjoint paths connecting nodes 1 and 2. As we noted in Section 3, this dual path problem is a minimum cost network flow problem, and so $\rho_O = 1$. The completion cost multiplier λ is 1 since the overlay completion procedure of adding edges (to the dual path) in increasing secondary cost order to span the remaining secondary nodes incurs a cost no more than the optimal value of the base subproblem. The BU heuristic is the same for both the DP-on-DPT and DPST-on-DPT problems: we find an approximate or optimal DPT solution and install primary facilities on all the edges of this design. Therefore, the results of Theorem 9 apply. Substituting ρ_O = 1 in expressions (4.8a) and (4.8c) gives

Corollary 10:

For the DP-on-DPT model, the worst-case performance ratios ω_{prop} and ω_{unrel} for problems with proportional costs and unrelated costs are bounded from above as follows:

$$\omega_{prop} \le \frac{4}{4-\rho_B}, \text{ and} \quad (4.9a)$$

$$\omega_{unrel} \le 2. \quad (4.9b)$$

For proportional cost problems, when we solve the DPT subproblem optimally (e.g., if costs are triangular or edge duplication is permitted, and we apply the matroid intersection algorithm), Corollary 10 gives the same worst-case bound of 4/3 as Theorem 1. However, when we use the dual path greedy completion (DPGC) heuristic with $\rho_B = 2$ to approximately solve the DPT base subproblem, Corollary 10 reduces the bound on ω_{prop} from 16/7 (in Theorem 1) to 2. Similarly, when we apply the 1-tree heuristic (with ρ_B = 3/2) to approximately solve the triangular cost DPT base problem, Corollary 10 gives a proportional costs worst-case ratio of 8/5 while Theorem 1 implies a ratio of 9/5.

4.2.1 DP-on-DPT worst-case examples

Since we will present DP-on-DPT worst-case examples for several cases, we first provide a brief overview of these examples. Figures 2 and 3 describe worst-case examples for the proportional cost DP-on-DPT problem. These examples achieve the bounds of 4/3 and 2 corresponding to situations when we either (i) solve the DPT subproblem optimally, or (ii) use the DPGC heuristic with a worst-case performance ratio of 2 to solve the DPT subproblem. For triangular, proportional cost problems, we have not been able to construct an example that achieves the bound of 8/5 when we use 1-tree heuristic as the embedded base subproblem solution method. Figure 4 describes an example with a heuristic performance ratio of 3/2.

Figures 5 and 6 describe worst-case examples for the unrelated costs DP-on-DPT problem. Although we considered only the OC heuristic in order to develop the worst-case bound of 2 (Theorem 9) for problems with unrelated costs, our examples demonstrate that the bound is tight even when we apply the BU heuristic and choose the better of the BU and OC heuristic solutions. Figure 5 assumes that we solve the DPT subproblem optimally, while Figure 6 assumes that we apply the DPGC heuristic to approximately solve the DPT subproblem.

Let us now discuss these examples in more detail. Figure 2 contains a worst-case example for the *proportional cost DP-on-DPT problem* to show that, when we solve the DPT subproblem optimally, the bound of 4/3 is tight. Figure 2(a) shows the network configuration and the secondary costs; the primary-to-secondary cost ratio r is 2 for all edges. Edges {1,a} and {2,b} each have secondary costs of 1/q: q is a sufficiently large multiple of 4. Edges {a,2} and {b,1} each have a cost of 1/4. The network contains two parallel paths, each containing q/4 nodes, connecting nodes a and 2; every edge on these paths has a secondary cost of 1/q. Each intermediate node on these two parallel paths is connected to the node vertically adjacent to it with an edge of cost 1/q. A similar configuration of parallel paths connects node b to node 1. The OC heuristic solution, shown in Figure 2(b), costs $2\{1/2 + 2(1/q)\} + (4q/4)(1/q) = 2 + 4/q$. The BU heuristic solution (Figure 2(c)) costs $2\{2(q/4 + 1)(1/q) + 2(q/4)(1/q) + 2(1/q)\} = 2 + 8/q$. Finally, the optimal solution (Figure 2(d)) costs $2\{(q/4 + 1)(1/q) + 2(1/q)\} + 2(q/4)(1/q) = 3/2 + 8/q$. Therefore, the heuristic performance ratio for this example approaches 4/3 as q approaches infinity.

Figure 3 describes a proportional cost worst-case example that achieves the bound of 2 when we solve the DPT subproblem using the DPGC heuristic. Figure 3(a) shows the network configuration and secondary costs; in this example, the cost ratio r is 1. The network has four alternate paths connecting the primary nodes 1 and 2: (i) a direct path of secondary cost 1, (ii) a two-edge path with edge costs 1/q and (1–1/q), and (iii) two q-edge paths with total secondary cost 1. The OC heuristic, shown in Figure 3(b), costs 4–2/q. The BU heuristic, shown in Figure 3(c), also costs 4–2/q. Figure 3(d) shows the optimal solution, which costs 2 + 1/q. Thus, the performance ratio of the composite heuristic approaches 2 as q approaches infinity.

Figure 4 presents a DP-on-DPT example with *triangular*, proportional costs. The given graph G is the triangularized version of the graph shown in the figure. Unlike the previous example, the cost ratio r is 2 instead of 1. The OC heuristic solution shown in Figure 4(b) costs 6–2/q, the BU heuristic solution (using the embedded 1-tree heuristic) shown in Figure 4(c) costs 6, while the optimal solution (Figure 4(d)) costs 4+1/q. As q approaches infinity, the performance ratio for this example approaches 3/2.

Figures 5 and 6 contain examples for the *unrelated costs DP-on-DPT problem*. Assuming that we can solve the DPT subproblem optimally, our worst-case example has the same network configuration as Figure 2(a), but uses the cost parameters shown in Figure 5. Figures 2(b), 2(c), and 2(d) depict the structure of the OC heuristic solution, BU heuristic solution (assuming that we solve the DPT subproblem optimally), and optimal solution for this example. For large values of q, the performance ratio approaches the worst-case performance bound of 2.

We can similarly modify the costs of Figure 3(a) to show that our worst-case bound is tight even when we use the DPGC heuristic to solve the DPT subproblem. Figure 6 shows the cost parameters for this example. Figures 3(b), 3(c), and 3(d) depict the structure of the OC heuristic solution, BU heuristic solution, and optimal solution for this example. For large values of q, the performance ratio approaches the worst-case bound of 2.

In closing, we note that whenever we use the DPGC heuristic to solve the DPT base subproblem of the DP-on-DPT model, then the composite heuristic achieves the tight performance ratio bound of 2 for both proportional and unrelated costs. Thus, this model provides one example for which the worst-case performance ratio for the base subproblem equals the worst-case bound for the two-tier, two-connected overlay optimization problem.

5. Heuristic Analysis of Cycle+Tree Problems with Partial Back-up

The partial back-up Cycle+Tree model permits secondary facilities on one of the two edge-disjoint paths connecting the two critical nodes 1 and 2. This section studies the partial back-up counterparts—the SP-on-DPT and ST-on-DPT problems—of the full-back-up DP-on-DPT and DPST-on-DPT problems that we considered in Section 4. In the SP-on-DPT model, the graph contains two primary nodes, both of which are critical. We must find the minimum cost spanning subgraph that contains a primary path connecting nodes 1 and 2, and an alternate edge-disjoint 1-to-2 path containing either primary or secondary facilities. Its overlay subproblem is the shortest path (SP) problem, and its base subproblem is the dual path tree (DPT) problem. The more general ST-on-DPT problem contains more than two primary nodes. The design must contain (i) a primary subgraph spanning all the primary nodes, (ii) an alternate 1-to-2 path containing primary or facilities, and (iii) spanning tree edges connecting all the remaining secondary nodes. The connectivity requirements for the ST-on-DPT model are: (i) at the primary service level: $r^1_{ij} = 1$ if i and $j \in P$; and (ii) at the secondary service level: $r^2_{12} = r^2_{21} = 2$ and $r^2_{ij} = 1$ otherwise. For this two-tier, partial back-up model, the overlay subproblem is the Steiner tree (ST) problem.

Although the distinction between full and partial back-up models appear to be relatively minor, the analysis for the partial back-up versions is quite different because, in general, the overlay subproblem might generate a solution that cannot be feasibly completed. However, as we show in Section 5.1, problems defined on μ-direct graphs satisfy the feasible completion property. In this case, we obtain improvements on the general overlay results of Section 2 by modifying and obtaining sharper worst-case bounds on both the BU and OC heuristic procedures. Also, unlike Section 4 where we directly applied the worst-case results for the DPST-on-DPT problemto the DP-on-DPT problem, in this section we exploit the special features of the SP-on-DPT problem to obtain better bounds than the general ST-on-DPT model. We study the SP-on-DPT problem in Section 5.1, and the ST-on-DPT model in Section 5.2.

5.1 The SP-on-DPT problem

Let us first address the potential difficulties in overlay completion for SP-on-DPT problems with general costs. The SP-on-DPT problem's overlay subproblem is a shortest path problem (from node 1 to 2). The OC heuristic first finds the shortest 1-to-2 path using primary costs, and then attempts to "complete" this path by adding secondary edges to (i) create an alternate edge-disjoint path from node 1 to node 2, and (ii) span all the nodes. Unfortunately, as the simple proportional costs example in Figure 7(a) demonstrates, the shortest overlay path might not have a feasible completion. For this example, the shortest overlay path is 1-3-4-2. However, when we attempt to complete this solution, we find that the network does not contain an alternate edge-disjoint 1-to-2 path, i.e., the completion subproblem is infeasible, even though the original problem is feasible. The optimal SP-on-DPT solution uses 1-3-2 as the primary path with a total cost of 202; path 1-4-2 is an edge-disjoint secondary path with a cost of 101. Observe that if we permit edge duplication, then the counterexample of Figure 7(a) is not valid and we can always find a feasible completion for any overlay path.

Notice that the graph in Figure 7(a) does not contain edge $\{1,2\}$. If we add this edge, then the problem instance has a feasible completion for any overlay solution. Furthermore, for μ-direct problems, we can obtain a bound on the completion cost multiplier λ.

Proposition 11:

SP-on-DPT problems with μ-direct costs satisfy the feasible completion property with completion cost multiplier λ of at most $1 + \mu/2$.

Proof:

Consider any overlay solution to the SP-on-DPT problem, i.e., any simple path PATH from node 1 to node 2. We need to show that we can complete this solution into a feasible SP-on-DPT solution, incurring an incremental cost of no more than $\{1 + \mu/2\}Z_B(a)$. If PATH is the edge $\{1,2\}$, then adding the edges in any optimal base subproblem solution (except edge $\{1,2\}$) produces a feasible completion that costs no more than $Z_B(a)$. If PATH is not edge $\{1,2\}$, then we can complete it by adding the edges (except edges in PATH), from the feasible solution found by the 1-tree heuristic. By Corollary 6, these added edges cost at most $\{1 + \mu/2\}Z_B(a)$. Therefore, we obtain the desired result. ♦

As the example in Figure 7(b) shows, this bound of $1+\mu/2$ on λ is tight for arbitrary μ. In this figure, the curved paths 1 to 4 and 3 to 2 have each have $q-1$ intermediate nodes, and the total length of these paths is 1. The optimal DPT solution for this example, shown in bold, has a total secondary cost of 2 if $\mu > 1/q$. Suppose we first select a shortest 1-to-2 path, 1-3-4-2, as the overlay path. Then, to complete this path, we must select all but the last edge on each of the curved paths from 1 to 4 and 2 to 3 in order to span these nodes.

Since we require a back-up 1-to-2 path that is edge-disjoint with respect to the overlay path 1-3-4-2, we must select edge {1,2} to complete the solution. The total completion cost is, therefore, $\mu + (2-2/q)$, and λ approaches $1+\mu/2$ as q tends to infinity.

5.1.1 SP-on-DPT problems with μ-direct, proportional costs

To develop improved worst-case bounds for SP-on-DPT problems with μ-direct proportional costs, we will modify the general OC and BU procedures to exploit the problem's special structure.

Modified OC heuristic ModOC1:

Step 1: Install a primary facility on edge {1,2}.

Step 2: Delete edge {1,2} from G, and install secondary facilities on all edges of the minimum spanning tree T* of the residual graph G*.

This heuristic differs from the general OC heuristic because it does not choose the optimal overlay solution in the first step even though the overlay subproblem (which is a shortest path problem) is easy to solve. This procedure is an adaptation of the 1-tree heuristic for the single-level DPT problem to the two-level SP-on-DPT problem. (In Section 5.2 we consider an overlay completion heuristic ModOC2 for the ST-on-DPT problem that also applies to the SP-on-DPT problem and does begin by choosing the optimal overlay solution.)

To develop an upper bound on the cost of the ModOC1 solution, note that if $\mu \leq 1$, then $a_{12} = s\, Z_B(a)$ since edge {1,2} is the shortest 1-to-2 path. Otherwise, if A_{12} denotes the length of the shortest 1-to-2 path (without edge {1,2}),

$$a_{12} \;\leq\; \mu\, A_{12} \;=\; \mu s Z_B(a). \qquad (5.1)$$

Therefore, if we define $\hat{\mu} = \max\,\{\mu, 1\}$, the heuristic solution produced by ModOC1 costs at most

$$Z^{\text{ModOC1}} \;\leq\; \hat{\mu} rs\, Z_B(a) + Z_B(a) \;=\; \{\hat{\mu} rs + 1\}\, Z_B(a). \qquad (5.2)$$

Note that ModOC1 has a completion cost multiplier $\lambda = 1$.

Note that when $\mu = \hat{\mu} = 1$ (e.g., when the costs are triangular), the right-hand side of inequality (5.2) is the same as the OC heuristic upper bound of $\{rs+1\}Z_B(a)$ that we used for the DP-on-DPT problem. So, if we apply the standard BU heuristic, then for SP-on-DPT problems with proportional costs, the composite heuristic has the same worst-case bound of $4/(4-\rho_B)$ (see Corollary 10). In particular, if the edge costs are triangular and we solve the DPT base subproblem optimally using the matroid intersection algorithm, the resulting composite heuristic has a worst-case bound of 4/3. However, we can reduce this bound by using the following improved BU heuristic for SP-on-DPT problems.

The standard BU heuristic solves the DPT base subproblem optimally or approximately, and installs primary facilities on all the edges in this solution. Instead, we will consider the following modified procedure:

Modified BU heuristic ModBU:

Step 1: Solve the DPT subproblem optimally or approximately using secondary costs.

Step 2: Install primary facilities on either of the two paths connecting nodes 1 and 2.

Choosing the shorter 1-to-2 path as the primary path in the second step obviously produces a superior SP-on-DPT solution; however, our worst-case bound applies even if we install primary facilities on the longer 1-to-2 path. The modified BU heuristic outperforms our original BU heuristic since it avoids installing non-essential primary facilities on edges of the DPT solution that do not lie on the chosen 1-to-2 path. (As shown by Balakrishnan et al. [1994a], for the most general overlay optimization problems, this strategy does not improve the worst-case performance of the composite heuristic. However, for the SP-on-DPT problem, we will show that it does.)

Let ρ_B be the worst-case ratio for the DPT solution method that we use in Step 1. Let s' be the secondary cost of the shorter 1-to-2 path in the (optimal or approximate) base solution divided by the optimal base value $Z_B(a)$; s' must be at least as large as s, the relative cost of the shortest 1-to-2 path. Installing primary facilities on either of the two edge-disjoint 1-to-2 paths in the DPT solution produces a SP-on-DPT solution with a maximum total cost of

$$Z^{ModBU} \quad \leq \quad Z_B(a)\,\{r\,(\rho_B - s') + s'\}. \tag{5.3}$$

This upper bound is valid because the DPT solution costs at most $\rho_B Z_B(a)$ of which the shorter 1-to-2 path accounts for a cost of $s'\,Z_B(a)$. Consequently, even if we install primary facilities on the longer 1-to-2 path, we incur a total primary cost of no more than $(\rho_B - s')\,Z_B(a)$. Since $s' \geq s$, inequality (5.3) implies that

$$Z^{ModBU} \quad \leq \quad Z_B(a)\,\{r\rho_B - (r-1)s\}. \tag{5.4}$$

Let us now analyze the worst-case performance for SP-on-DPT problems with μ-direct, proportional costs assuming we use the modified OC and BU heuristics as the embedded solution methods in the composite procedure. The lower bound obtained by relaxing the linking constraint (2.4) is $Z_B(a)\{(r-1)s + 1\}$. Since the composite heuristic selects the better of the modified OC and BU heuristic solutions, the bounds (5.2) and (5.4) imply that

$$\omega_{prop} \quad \leq \quad \frac{\min\{r\rho_B-(r-1)s,\ \hat{\mu}rs+1\}}{\{(r-1)s+1\}}. \tag{5.5}$$

As before, the relative performance ratios of the modified BU and OC heuristics decrease and increase with s. For the SP-on-DPT problem, s must be less than or equal to 1/2 since the optimal DPT solution contains two 1-to-2 paths and therefore costs at least twice the optimal overlay (shortest path) solution using secondary costs. If we ignore this upper bound on s, the right-hand side of (5.5) attains its maximum value when

$$s^* \;=\; \frac{r\rho_B - 1}{(\hat{\mu}+1)r - 1}. \tag{5.6}$$

Note that since r and ρ_B have values at least 1, s* is nonnegative. However, for certain $\hat{\mu}$ and ρ_B values the value of s* given by (5.6) can exceed 1/2. We separately analyze one such case later. Note that even if the s* value computed using (5.6) exceeds 1/2, substituting this value in the right-hand side of (5.5) and maximizing the expression with respect to r gives a valid upper bound on the performance ratio. In this case, the maximizing value of r is

$$r^* \;=\; \frac{1 + \sqrt{1 + (1+\rho_B\hat{\mu}-\hat{\mu}^2)(\hat{\mu}-\rho_B)/\rho_B}}{1 + \rho_B\hat{\mu} - \hat{\mu}^2}. \tag{5.7}$$

As an illustration, let us consider the triangular costs case. Substituting $\hat{\mu} = 1$ in (5.7) gives

$$r^* \;=\; \frac{1 + \sqrt{2-\rho_B}}{\rho_B}. \tag{5.8}$$

Again, by definition, the cost ratio r must be ≥ 1. Nevertheless, even if the r* value given by (5.8) is < 1, substituting this value in the right-hand side of (5.5) gives a valid upper bound on ω_{prop}. Substituting for s* and r* from (5.6) and (5.8) in (5.5) gives

$$\omega_{prop} \leq \frac{4 - 2\rho_B + 3\sqrt{2-\rho_B}}{4 - 2\rho_B + (3-\rho_B)\sqrt{2-\rho_B}} \tag{5.9}$$

for SP-on-DPT problems with triangular, proportional costs.

Let us apply this bound to two scenarios. Suppose we solve the DPT base problem optimally using the matroid intersection algorithm: then (5.8) implies $r^* = 2$ and (5.9) implies $\omega_{prop} \leq 5/4$. Instead, suppose we use the 1-tree heuristic to approximately solve the DPT subproblem. In this case, $\rho_B = 3/2$ and so $r^* = 2(\sqrt{2}+1)/3\sqrt{2} = 1.1381$, $s^* = 3/\{4+\sqrt{2}\} = 0.5541$ and $\omega_{prop} = 1.5147$. Note that in this case s^* exceeds the upper bound of 1/2 on s. Therefore, the bound of 1.5147 is not likely to be tight. The following arguments enable us to improve upon this bound.

Notice that when $\hat{\mu} = 1$, for values of $\rho_B \geq 1.5$, the second term $(rs + 1)$ in the numerator of the right-hand side in inequality (5.5) is less than or equal to the first term $(r\rho_B-(r-1)s)$ since $r \geq 1$ and $s \leq 1/2$. That is, for μ-direct SP-on-DPT problems with $\mu \leq 1$, if we use a base heuristic with a worst-case ratio $\rho_B \geq 1.5$, the modified OC heuristic solution has a smaller upper bound than the modified BU heuristic. Therefore, inequality (5.5) reduces to

$$\omega_{prop} \leq \frac{rs+1}{\{(r-1)s+1\}}. \tag{5.10}$$

Since we require $s \leq 1/2$, the right-hand side of (5.10) achieves its maximum at $s^* = 1/2$ for all values of $r \geq 1$. At this value of s^*, the maximizing value of r is $r^* = 1$. Substituting $r^* = 1$ and $s^* = 1/2$ in (5.10) we obtain

$$\omega_{prop} \leq 3/2. \tag{5.11}$$

This bound applies, for instance, to triangular, proportional cost DPT problems when we solve the base subproblem using the 1-tree heuristic.

Theorem 12:

For SP-on-DPT problems with triangular, proportional costs, the worst-case performance ratio of the composite heuristic is bounded from above by

$$\begin{aligned} \omega_{prop} &\leq \frac{5}{4} \quad \text{if we solve the DPT base problem optimally, or} \\ &\leq \frac{3}{2} \quad \text{if we apply the 1-tree DPT heuristic for the base subproblem.} \end{aligned}$$

Thus, by modifying the BU and OC heuristics for the triangular, proportional cost SP-on-DPT problem we have reduced the worst-case bounds (i) from 4/3 (Theorem 1 and Corollary 10) to 5/4 when we solve the base subproblem optimally, and (ii) from 9/5 (Theorem 1) to 3/2 when we use the 1-tree heuristic to approximately solve the DPT base subproblem.

5.1.2 SP-on-DPT worst-case examples

Figure 8(a) shows a worst-case example to prove that the bound of 5/4 is tight for triangular, proportional cost SP-on-DPT problems when we solve the DPT subproblem optimally. The actual graph G for this problem is the triangularized version of the graph shown in Figure 8(a). Figures 8(b), 8(c), and 8(d) show the modified OC heuristic solution, the modified BU heuristic solution (this solution installs primary facilities on the longer 1-to-2 path), and the optimal solution, assuming a cost ratio $r = 2$. The OC heuristic solution costs $5/3 - 2/3q$, the BU heuristic solution costs 5/3, but the optimal value is 4/3, thus achieving the worst-case performance ratio of 5/4 for large q.

In the SP-on-DPT example in Figure 9, the composite heuristic achieves Theorem 12's bound of 3/2 when we use the 1-tree method as the embedded DPT heuristic. Figure 9(a) shows the secondary costs for select edges. The actual graph G is the triangularized version of this graph, and the cost ratio is 1. The network contains three alternate 1-to-2 paths, each with a total secondary length of 1. The solutions constructed by the modified OC and BU heuristics, shown in Figures 9(b) and 9(c), cost $3-1/q$ each. On the other hand, the optimal solution (Figure 9(d)) costs 2. Therefore, by choosing a large value of q, we obtain a performance ratio that is arbitrary close to 3/2.

5.2 The ST-on-DPT problem

The ST-on-DPT model generalizes the SP-on-DPT problem by permitting more than two primary nodes, but only two of these primary nodes are critical. The overlay problem is, therefore, the minimum Steiner tree problem with the primary nodes as terminals, and the base problem is the DPT problem. As in the SP-on-DPT model, for ST-on-DPT problems with general costs the overlay solution might not always have a feasible completion. However, problems with μ-direct costs satisfy the feasible completion property. For our worst-case analysis, we will focus on the triangular costs special case (with $\mu = 1$); extensions of these results apply to problems with arbitrary (but prespecified) values of μ. As we show below, we can modify the OC heuristic to guarantee feasible completion with $\lambda = 1$ for any feasible overlay solution.

For our worst-case analysis of the ST-on-DPT problem, we assume the standard BU heuristic, i.e., the method solves the base DPT subproblem and installs primary facilities on all the edges of this solution. If r is the common primary-to-secondary cost ratio for all edges, the cost of the BU solution is at most $r\, Z_B(a)$. Unlike the SP-on-DPT problem for which we exploited the superior performance of the Modified BU heuristic, we cannot obtain a tighter upper bound than r $Z_B(a)$ even if we modify the BU heuristic to upgrade only essential facilities (i.e., only on the edges of the minimal subtree in the base solution spanning all primary nodes).

We can, however, improve the bound on the OC solution if we replace the standard OC heuristic with a different optimal completion heuristic; in Step 1 of this heuristic we find an approximate or optimal Steiner tree spanning all the terminal nodes and install primary facilities on the edges of this tree.

Modified OC heuristic ModOC2 for ST-on-DPT problems:

<u>Step 1</u>: Find the optimal or approximate Steiner tree STREE in G spanning all the primary nodes, and install primary facilities on all the edges of this tree.

<u>Step 2</u>: If STREE contains edge {1,2}, then in the graph obtained by removing edge {1,2}, complete the solution by adding edges in order of increasing costs (without adding any new cycles) until the solution spans all the nodes; install secondary facilities on the added edges. Otherwise, use the 1-tree heuristic to find a feasible solution to the DPT problem with secondary costs. Install secondary facilities on the edges in this solution that aren't already in STREE.

This procedure differs from the standard OC heuristic in the following way. Instead of completing the overlay solution obtained in Step 1 by solving (optimally or approximately) the base subproblem, we apply a greedy completion procedure. Not only is this procedure efficient, but it also guarantees that the completion cost will not exceed 1.5 times the optimal base value. Thus, unlike the general overlay analysis, our upper bound for the cost of the OC heuristic will not contain the factor ρ_B in the overlay completion cost.

To develop an upper bound on the total cost of the heuristic solution generated by ModOC2, let s be the (secondary) cost of the minimum Steiner tree spanning all the terminal nodes expressed as a proportion of the optimal DPT value $Z_B(a)$. Note that, unlike the SP-on-DPT problem, s can now assume any value between 0 and 1. If we use a heuristic with worst-case ratio ρ_O to find the Steiner tree in Step 1, then the total cost of the primary facilities installed in this step is no greater than ρ_O r s $Z_B(a)$. Since the costs are triangular, $\hat{\mu} = \max\{\mu, 1\} = 1$, and since the cost of the Steiner tree STREE is at least the cost of the shortest 1-to-2 path, Corollary 6 implies that the 1-tree solution found in Step 2 costs no more than $(1 + \min\{(1/2), s\} Z_B(a)$. When STREE contains edge {1,2}, the total secondary cost of all the edges chosen by the greedy completion procedure in Step 2 does not exceed the optimal DPT value $Z_B(a) \leq (1 + \min\{(1/2),s\}Z_B(a)$. These observations provide the following upper bound on the total cost of the ModOC2 heuristic solution:

$$Z^{\mathrm{ModOC2}} \leq \rho_O rs\, Z_B(a) + (1 + \min\{\tfrac{1}{2}, s\})\, Z_B(a)$$
$$= \min\{(\rho_O r+1)s + 1, \rho_O rs + \tfrac{3}{2}\}\, Z_B(a). \quad (5.12)$$

We note parenthetically that the arguments we used to develop the upper bound (5.12) also show that for the triangular cost ST-on-DPT problem, the completion cost multiplier λ is 1.5. (In contrast, our previous three models had $\lambda = 1$.) Since Theorems 1 and 2 assume $\lambda = 1$, those results do not apply directly to the ST-on-DPT problem; we need to develop worst-case expressions in terms of arbitrary (but prespecified) values of the completion cost multiplier λ.

Note that when specialized, this ModOC2 procedure provides an alternative to ModOC1 for the SP-on-DPT problem. For the SP-on-DPT problem, Step 1 becomes a shortest path computation and $\rho_O = 1$. For problems with triangular costs (and so $\mu \leq 1$), the heuristic ModOC1 never provides a higher cost solution than this specialized version of ModOC2. For this reason, we considered only ModOC1 in Section 5.1. Similarly, for the ST-on-DPT problem, although we could develop a heuristic analogous to ModOC1, this heuristic does not provide a tighter upper bound than (5.12).

To analyze the worst-case performance of the composite procedure (using the standard BU heuristic and the modified OC heuristic ModOC2), we will first consider only the first term (i.e., the expression $(\rho_O r+1)s + 1$) in the minimand in the right-hand side of the OC heuristic bound (5.12). We will refer to this bound as the *direct-relative-to-overlay* bound since it uses the overlay cost as an upper bound on the cost of the direct edge {1,2}. Subsequently, we will develop an alternate bound for the composite procedure using the second expression (i.e., $\rho_O rs + \frac{3}{2}$) in (5.12). We refer to this bound as the *direct-relative-to-base* bound. The composite heuristic has the smaller of these two worst-case bounds. As we find later, neither bound dominates the other, i.e., for two out of the three solution options that we consider, the first bound is smaller, while the second bound is smaller for the third option.

In the following discussion, we consider only those cases that apply to values of the parameters ρ_O and ρ_B that are of interest to us. Since the costs are triangular, we can solve the DPT base problem either optimally using the matroid intersection algorithm, or approximately using the 1-tree heuristic. Correspondingly, ρ_B is either 1 or 3/2. Similarly, for the overlay (Steiner tree) subproblem, we will either solve it optimally, or solve it approximately using the MST heuristic (with $\rho_O = 2$).

Furthermore, we will not consider the case when we solve the overlay subproblem optimally but solve the base subproblem only approximately. So, we will consider the following three combinations of embedded base–overlay solution methods: optimal–optimal, optimal–MST, and 1-tree–MST. We will refer to these combinations as *relevant combinations.*

Direct-relative-to-overlay bound:

For ST-on-DPT problems with triangular, proportional costs, we obtain the worst-case ratio for the modified composite heuristic by choosing the minimum of the BU upper bound and the modified OC upper bound, and dividing by the lower bound (2.8). Using the direct-relative-to-overlay bound for the modified OC heuristic solution, we obtain

$$\omega_{prop} \leq \frac{\min\{\rho_B r, (\rho_O r+1)s + 1\}}{\{(r-1)s + 1\}}. \tag{5.13}$$

The BU heuristic's ratio (the first term in the numerator divided by the denominator) decreases as s increases and since $(\rho_O r+1) > (r-1)$, the OC heuristic's ratio increases. Thus, the value of s that maximizes the right-hand side of (5.13) is

$$s^* = \frac{\rho_B r - 1}{\rho_O r + 1}. \tag{5.14}$$

For the relevant combinations, this value of s* satisfies $0 \leq s^* \leq 1$. Substituting for s* in (5.13), we obtain

$$\omega_{prop} \leq \frac{\rho_B r(\rho_O r+1)}{\rho_B r^2+(\rho_O-\rho_B-1)r+2}. \tag{5.15}$$

Note that if $(\rho_O-\rho_B-1) \geq 0$, then the right-hand side of (5.15) achieves its maximum value of ρ_O when $r = \infty$. However, for the relevant combinations, $(\rho_O-\rho_B-1) < 0$. So, we maximize the right-hand side of (5.15) with respect to r. The maximizing value of r is

$$r^* = \frac{2\rho_O + \sqrt{2(\rho_O+1)(\rho_O+\rho_B)}}{\rho_B + \rho_O\rho_B + \rho_O - \rho_O^2}. \tag{5.16}$$

For all the relevant combinations, this value of r* is at least 1 as required. Substituting this value of r* in (5.15) gives an upper bound on ω_{prop} in terms of the base and overlay worst-case ratios ρ_B and ρ_O. The following theorem evaluates the numerical values of this ratio for the relevant combinations.

Theorem 13:

For ST-on-DPT problems with triangular, proportional costs, the composite method's (using the modified OC heuristic ModOC2) worst-case ratio has the following upper bounds for various combinations of the embedded base–overlay solution methods:

$$
\begin{aligned}
\omega_{prop} &\le 1.522 && \text{for the optimal–optimal combination} \\
&\le 2.061 && \text{for the optimal–MST combination, and} \\
&\le 2.255 && \text{for the 1-tree–MST combination.}
\end{aligned}
$$

Direct-relative-to-base bound:

Using the direct-relative-to-base bound (in the right-hand side of inequality (5.12)) for the Modified OC heuristic, we obtain the following worst-case ratio for the modified composite heuristic:

$$\omega_{prop} \le \frac{\min\{\rho_B r,\ \rho_O rs + 3/2\}}{\{(r-1)s + 1\}}. \tag{5.17}$$

The BU heuristic's ratio decreases as s increases and if $\rho_O r > 3(r-1)/2$, then the OC heuristic's ratio increases. If $\rho_O \ge 3/2$, this inequality is always valid; if $\rho_O = 1$, we require r to be less than 3. The optimal value r* that we compute later for the relevant combination $\rho_O = \rho_B = 1$ is 2, and so satisfies this condition. In this case, the value of s that maximizes the right-hand side of (5.17) is

$$s^* = \frac{2\rho_B r - 3}{2\rho_O r}. \tag{5.18}$$

For our relevant combinations $\rho_O \ge \rho_B$, and so s* does not exceed 1. When $\rho_B = 3/2$, s* ≥ 0, and for all combinations with $\rho_B = 1$, r* = 2 and so s* ≥ 0. Substituting for s* in (5.17), we obtain

$$\omega_{prop} \le \frac{2\rho_B\rho_O r^2}{2\rho_B r^2 + (2\rho_O - 2\rho_B - 3)r + 3}. \tag{5.19}$$

If $(2\rho_O - 2\rho_B - 3) < 0$, a valid condition for our combinations, the right-hand side of (5.19) first increases and then decreases with r. The maximizing value of r is

$$r^* = \frac{6}{3 - 2\rho_O + 2\rho_B}. \tag{5.20}$$

For all the relevant combinations, this value of r* is greater than or equal to 1 as required. Substituting this value of r* in (5.19) gives the following upper bound on ω_{prop}:

$$\omega_{prop} \le \frac{24\rho_B\rho_O}{24\rho_B-(3+2\rho_B-2\rho_O)^2}. \tag{5.21}$$

Theorem 14:

For ST-on-DPT problems with triangular, proportional costs, the composite method's (using the modified OC heuristic ModOC2) worst-case ratio has the following upper bounds for various combinations of the embedded base–overlay solution methods:

$$\begin{aligned} \omega_{prop} &\le 1.6 && \text{for the optimal–optimal combination} \\ &\le \frac{48}{23} = 2.087 && \text{for the optimal–MST combination, and} \\ &\le 2.25 && \text{for the 1-tree–MST combination.} \end{aligned}$$

Notice that for the first two combinations of base-overlay solution methods, the direct-relative-to-overlay bound is superior, whereas the direct-relative-to-base bound is smaller for the 1-tree–MST combination. Notice also that we can improve the worst-case bounds in Theorems 13 and 14 if we use a better heuristic, say Berman and Ramaiyer's [1992] heuristic (which has a worst-case bound of 16/9) to solve the overlay (Steiner tree) subproblem.

As we remarked earlier, the results of Theorem 1 do not apply directly to ST-on-DPT problems since the completion cost multiplier λ for this class of problems is 1.5 whereas Theorem 1 assumes that $\lambda = 1$. Since we have exploited the ST-on-DPT problem's special structure, we expect the bounds of Theorems 13 and 14 to be superior to the general bounds. To gauge the improvement, we can derive results analogous to those in Theorem 1 but for general values of λ (in this case expression (2.5) becomes $Z^{Comp} \le \min \{\rho_B Z_B(c), \rho_O Z_O(c) + \rho_B \lambda Z_B(a)\}$. This exercise gives the following *general overlay bound* applicable to our relevant combinations:

$$\omega_{prop} \le \rho_B \frac{4\,\rho\,\lambda}{4\lambda - (1+\lambda-\rho)^2}. \tag{5.22}$$

Notice that if we substitute $\lambda = 1$, then the right-hand side of (5.22) reduces to the right-hand side of the bound (2.6a) in Theorem 1. For each of our relevant combinations, substituting the corresponding values of ρ_B and ρ_O, and setting $\lambda = 1.5$ in (5.22) gives the following *general* overlay bounds:

$$\begin{aligned} \omega_{prop} &\le 8/5 = 1.6 && \text{for the optimal–optimal combination} \\ &\le \frac{48}{23} = 2.087 && \text{for the optimal–MST combination, and} \end{aligned}$$

$$\leq \frac{432}{167} = 2.587 \quad \text{for the 1-tree–MST combination,}$$

which are higher than the specialized bounds of Theorem 13 and 14.

To conclude this section, we note that bounds in this section for the ST-on-DPT model are higher than those in Section 5.1 for the SP-on-DPT model. These latter bounds exploited the special properties of the subproblem solutions (e.g., the optimal overlay cost is no more than 1/2 the optimal base cost) when the network contains only two primary nodes.

6. Conclusion

In this paper, we have initiated the analysis of general multi-tier, multi-connected network design models. We have provided an integer programming formulation of a rather general model and shown how to interpret it as a special case of the so-called overlay optimization problem. This interpretation permits us to apply our prior analysis concerning overlay optimization to develop and assess the worst-case performance of heuristic solution methods for general multi-tier, multi-connected network design models.

After reviewing some general results for two-tier overlay optimization problems, and applying them to the well-known hierarchical and two-level network design models, we have described two core single-tier survivability problems: the Dual Path Tree (DPT) and the Dual Path Steiner Tree (DPST) problems. The DPT problem seeks a minimum cost subgraph consisting of a spanning tree plus additional edges so that the solution also contains a cycle connecting two specified (critical) nodes. In this case, we showed how to solve the problem optimally as a matroid intersection problem when the edge costs are triangular. We also provided an easily implemented minimum spanning tree based heuristic method, the 1-tree heuristic, for the DPT problem; this method has a worst-case performance guarantee of 3/2. For the DPST problem we used a dual path greedy completion heuristic with a worst-case performance guarantee of two for problems with general (nonnegative) costs.

Building upon these results, we then studied four versions of the general multi-tier, multi-connected network design model—all have two service levels, require that the design be a dual path tree, and have two specially designated critical nodes that must be on the cycle in the dual path tree. The set of edges with the higher level of service (the overlay edges) need to be: (i) a path between the critical nodes, (ii) a cycle containing the critical nodes, (iii) a Steiner tree containing the critical nodes and other designated primary nodes, or (iv) a dual path Steiner tree that contains all the primary nodes and that has the critical nodes on its cycle.

For all four of these models, we have developed heuristics with worst-case bounds that depend upon the cost structure—proportional or unrelated, and arbitrary, μ-direct, or triangular—and how accurately we solve the dual path tree problem and the overlay problem. Table II illustrates these results by summarizing the bounds we have obtained for versions of the proportional cost model.

Table II: Comparison of general overlay bound and specialized bounds for selected versions of proportional cost problems

	Solve DPT optimally *(arbitrary nonnegative costs)*		1-tree heuristic for DPT *(triangular costs)*	
Method → / Problem	General Overlay Bound	Tailored Heuristic	General Overlay Bound	Tailored Heuristic
DP-on-DPT	4/3	4/3	9/5	8/5
DPST-on-DPT*	8/3	16/7	9/4	48/23
SP-on-DPT	4/3	5/4	9/5	3/2
ST-on-DPT*	2.087	2.061	2.587	2.25

* Table assumes we solve the DPST and ST problems by heuristics with worst-case bounds of 2.

These results improve upon the worst-case bounds for the general overlay problem (as specified in Section 2) by exploiting the problems' special structure.

Our discussions in Sections 1 and 2 suggest several opportunities for modeling, analysis, and algorithmic development in the new arena of multi-tier, multi-connected problems, a class of models that is likely to gain increasing importance as the telecommunications industry emphasizes cost effective investments to upgrade the switching and transmission facilities while providing adequate levels of reliable service to different customer classes. Decomposition algorithms and optimization-based heuristics offer considerable promise to effectively solve these difficult problems. For the HND and

TLND problems, Balakrishnan et al. [1992b] developed and tested a dual ascent technique that generates linear programming-based heuristic solutions as well as lower bounds to verify the quality of these solutions. Although this method can have arbitrarily poor worst-case performance, extensive computational testing confirmed that the method generates very good heuristic solutions that are within 1% of the lower bounds for a variety of cost structures. These results also show that the gap between the objective values of these problems and their linear programming relaxations tends to be very small. A similar approach might prove to be fruitful for solving the two-tier, two-connected network design problem, and other overlay optimization models. For instance, we might explore the possibility of building upon Magnanti and Raghavan's [1992] dual ascent method for the single-level network design problem with connectivity constraints to solve our multi-level, multi-connected model (the single-level model is the base subproblem for the multi-level model). Or, we might adopt the type of polyhedral results for network survivability developed by Grötschel et al. [1992] for solving any single level subproblem in a decomposition or polyhedral approach.

Our worst-case analysis for the four two-tier, two-connected models can extend directly or motivate similar approaches for various more complex versions of multi-tier survivable network design problems. We briefly discuss three possible extensions: higher connectivity requirements for critical nodes, rings on trees, and multiple groups of primary and critical nodes. For simplicity, we will describe these extensions to the DP-on-DPT model; many of the principles apply to analogous extensions of the other three models as well (for the partial back-up models, we might need to modify the OC and BU heuristics and make appropriate assumptions to satisfy the feasible completion property).

K-connected critical nodes:

Suppose the two critical nodes in the HND model require K edge-disjoint full-backup paths for some specified value of $K > 2$. The overlay problem is then the "K-path" problem of finding K edge-disjoint 1-to-2 paths, and the base subproblem is the corresponding K-path Tree (KPT) problem. The K-path (KP) problem is easy to solve with or without edge duplication. With edge duplication, the optimal solution is K replications of the shortest 1-to-2 path, and if we do not permit edge duplication a minimum cost network flow problem finds the optimal solution. Consequently, $\rho_O = 1$ for the KP-on-KPT problem. Furthermore, we can show that a greedy completion of the optimal K-path overlay solution incurs a completion cost of no more than $Z_B(a)$. Although we do not know if the KPT problem is solvable in polynomial time when costs are triangular, the K-path Greedy Completion (KPGC) heuristic (i.e., find the optimal K edge-disjoint 1-to-2 paths, and add edges in increasing cost order to span all remaining nodes) finds a heuristic solution to the KPT problem (or even the K-path Steiner tree problem) that costs no more than twice the optimal value (see Balakrishnan et al. [1994b]). This result holds even for non-triangular costs. Therefore, the worst-case bounds in Corollary 10 apply to the KP-on-KPT problem with $\rho_O = 1$ and $\rho_B = 2$, i.e., the composite heuristic generates a feasible solution with a performance guarantee of 2 for both proportional cost and unrelated cost problems. Similar arguments can establish the validity of a bound of 16/7 (Theorem 9) for KPST-on-KPT problems if we apply the KPGC heuristic to approximately solve the overlay and base subproblems.

Rings on trees:

Consider a generalization of the DP-on-DPT problem containing more than two primary nodes that are all critical and must be connected via a primary *Steiner ring*, i.e., a hamiltonian tour containing primary facilities that might optionally contain secondary nodes. The network design must connect the secondary nodes to this ring via subtrees containing secondary facilities. This type of ring-on-tree topology is a core configuration in SONET networks. The overlay problem is the Steiner ring (SR) problem which is NP-hard (since the traveling salesman problem is a special case) and the base subproblem is the Ring+Tree (R+T) problem of finding a spanning network containing a (edge-disjoint) ring that visits each primary node exactly once. Balakrishnan et al. [1994b] have shown that if we find the optimal Steiner ring, then adding spanning tree edges to this ring in order of increasing costs to span the remaining secondary nodes produces an R+T heuristic solution that costs no more than twice the optimal value. Therefore, if we solve the SR problem optimally (both for the overlay solution and for the base heuristic), then Corollary 10 applies with $\rho_O = 1$ and $\rho_B = 2$, giving a worst-case performance bound of 2 for the SR-on-R+T problem. We can similarly apply the bound in Theorem 9 to a generalization that includes non-critical primary nodes (that must be connected to each other and to the critical nodes via at least one primary path). Furthermore, the analysis extends to problems that do not impose the requirement that the paths connecting critical nodes must be node-disjoint, i.e., the overlay subproblem is the 2-connected Steiner network design problem. For this problem, Monma et al. [1990] have shown that if the costs are triangular, the optimal traveling salesman tour that visits just the critical nodes costs no more than 4/3rds the optimal 2-connected solution. If the network does not contain any non-critical primary nodes, then this result implies that $\rho_O = 4/3$ if we find the optimal TSP. We can then incorporate this result in our analysis of the two-tier problem.

Multiple groups of primary nodes:

As a final extension, consider the generalization of the DP-on-DPT model with K disjoint pairs of critical nodes; the two nodes in each pair must be connected via two edge-disjoint primary paths that might optionally include secondary nodes as well as other primary nodes. In this case, the overlay solution might contain more than one component. Every two-connected pair of critical nodes must belong to a single component, and each component provides two-connectivity for all the critical node pairs that it spans. One heuristic solution for the overlay subproblem for this model is the union of optimal dual paths connecting every pair of critical nodes. We have not analyzed the worst-case performance of this method. Suppose the worst-case ratio is α. Adding spanning tree edges to this solution in order of increasing costs generates a heuristic base solution with a performance guarantee of 2α (see Balakrishnan et al. [1994b]). Furthermore, the completion cost multiplier λ for the overlay completion method is 1 using the greedy completion method. Therefore, our strategy for developing the worst-case results for Corollary 10 applies to this problem setting as well.

We can consider model enhancements to incorporate non-critical primary nodes. Again, instead of having a single group of primary nodes, we are given K clusters, each containing two critical nodes. The primary subnetwork must connect pairs of nodes within each cluster by primary facilities, but can use paths containing nodes from other clusters and/or secondary nodes. Similarly, the two critical nodes within each cluster must have two edge-disjoint paths that might optionally contain primary nodes from other clusters. We might consider full and partial back-up versions of these problems. Consider for instance the partial back-up version. The overlay problem is the Steiner forest (SF) problem (see, for instance, Balakrishnan et al. [1994a]), and the base problem is the following generalization of the dual path problem: find the minimum cost spanning network that contains two edge-disjoint paths connecting every pair of critical nodes. In this "multi dual path tree problem (MDPT)," the partial back-up version is the SF-on-MDPT problem, and the full back-up version is the DPSF-on-MDPT problem since its overlay problem is the "Dual Path Steiner Forest problem (DPSF)." Just as the single-tier DPT and DPST worst-case results of Section 3 shed light on the analysis of the models analyzed in this paper, we could develop multi-path, single tier results to serve as a starting point for developing heuristics with worst-case bounds for these survivability problems.

As this discussion suggests, there are plenty of opportunities for studying new versions of multi-tier, multi-connected network design problems and to develop effective new practical algorithms for these important problems. Hopefully, the perspective of overlay optimization discussed in this paper can play a role in these developments

Acknowledgments:
We are grateful to Professor James Orlin for many a discerning discussions on the general topic of overlay optimization and, particularly, for his insights concerning the matroid intersection algorithm for the Dual Path Tree problem. We are also grateful to the referees for thoughtful suggestions that have led to a number of improvements and enhancements to this work.

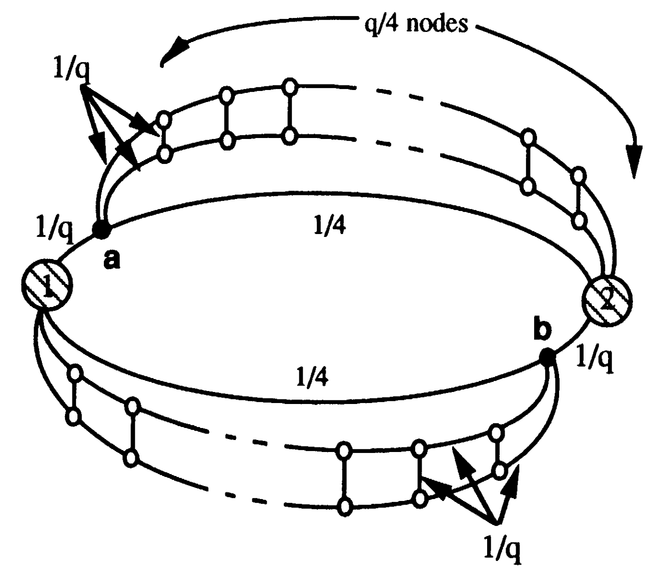

Figure 2(a): DP-on-DPT Worst-case example
secondary cost shown on edges; cost ratio = 2

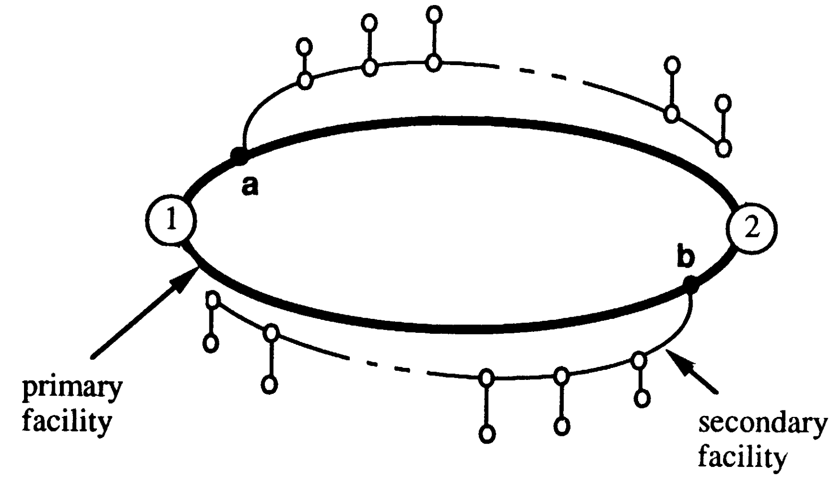

Figure 2(b): OC heuristic solution

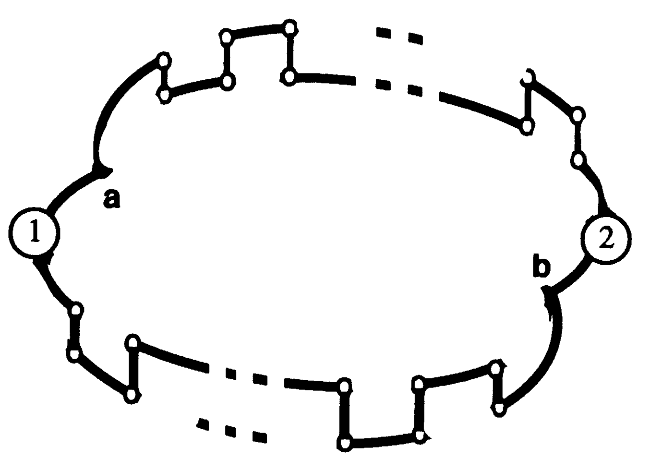

Figure 2(c): BU heuristic solution

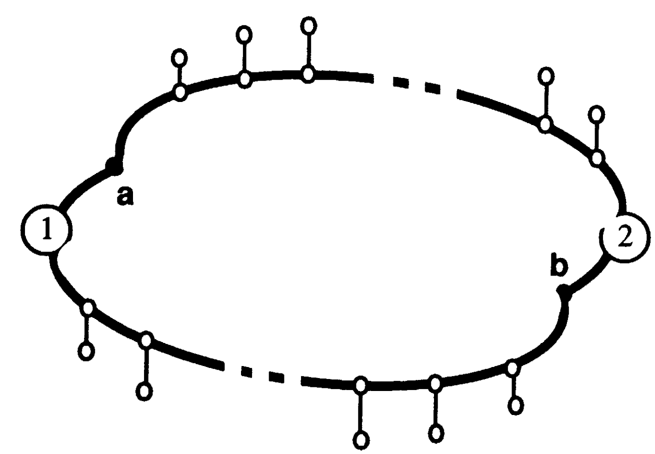

Figure 2(d): Optimal solution

Figure 2: Worst-case example for DP-on-DPT problem with proportional costs
DPT base subproblem solved optimally using Matroid Intersection algorithm

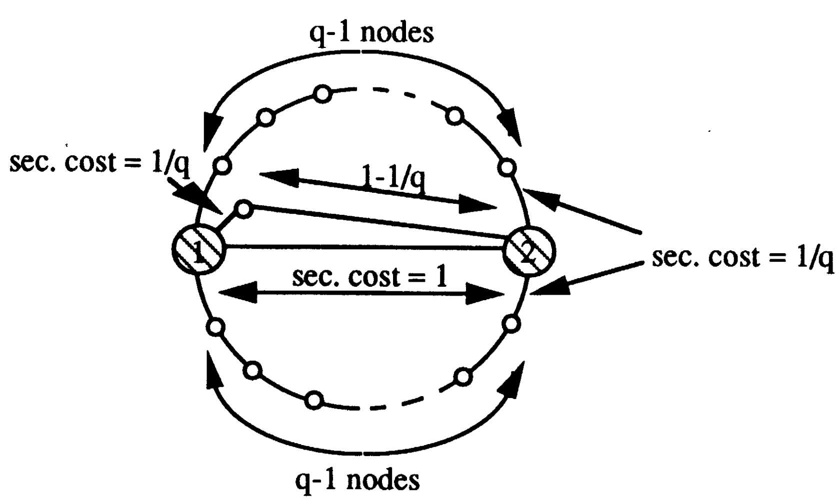

Figure 3(a): DP-on-DPT example with proportional costs
cost ratio = 1

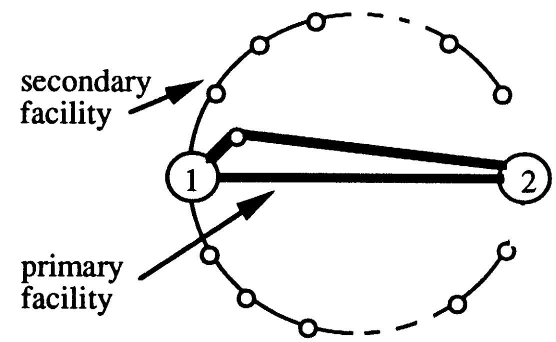

Figure 3(b): OC heuristic solution

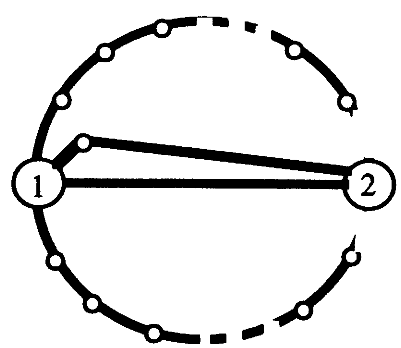

Figure 3(c): BU heuristic solution

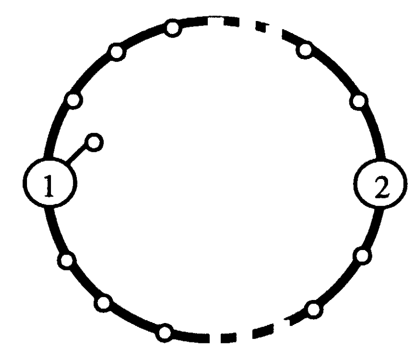

Figure 3(d): Optimal solution

Figure 3: Worst-case example for DP-on-DPT problem with proportional costs
Using Dual Path Greedy Completion heuristic to solve the DPT base subproblem

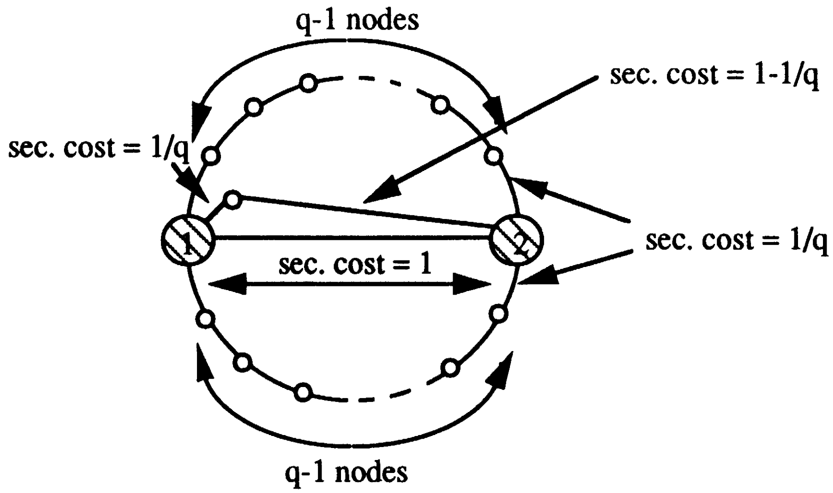

Figure 4(a): DP-on-DPT example with triangular costs
cost ratio = 2

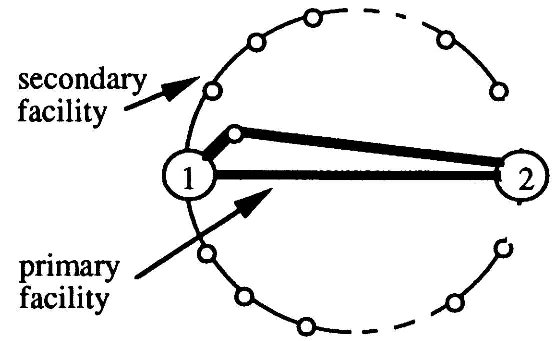

Figure 4(b): OC heuristic solution

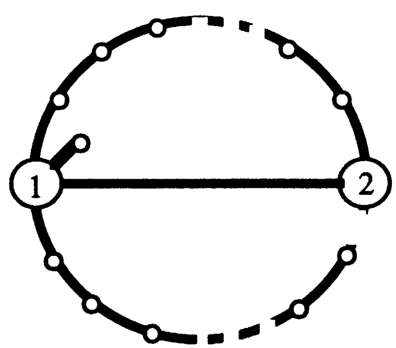

Figure 4(c): BU heuristic solution

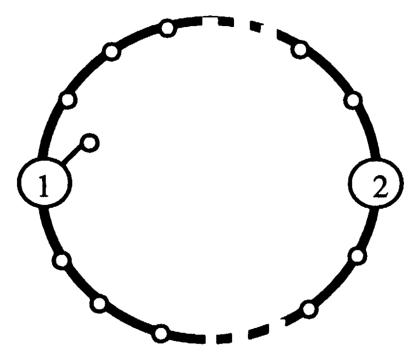

Figure 4(d): Optimal solution

Figure 4: Worst-case example for DP-on-DPT problem with proportional, triangular costs
Using 1-tree heuristic to solve the DPT base subproblem

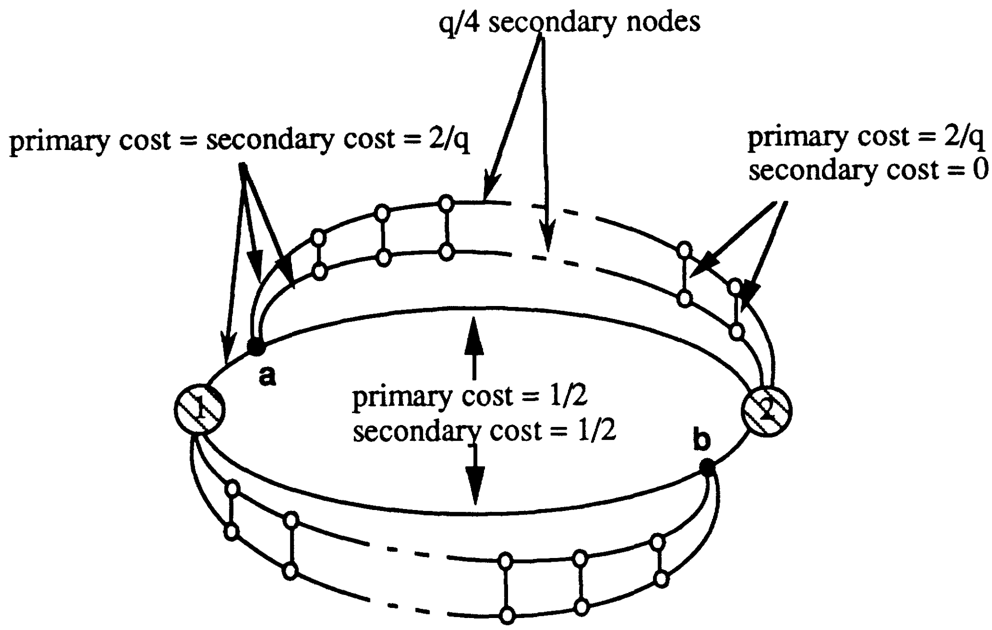

Figure 5: Worst-case example for DP-on-DPT problem with unrelated costs
DPT base subproblem solved optimally

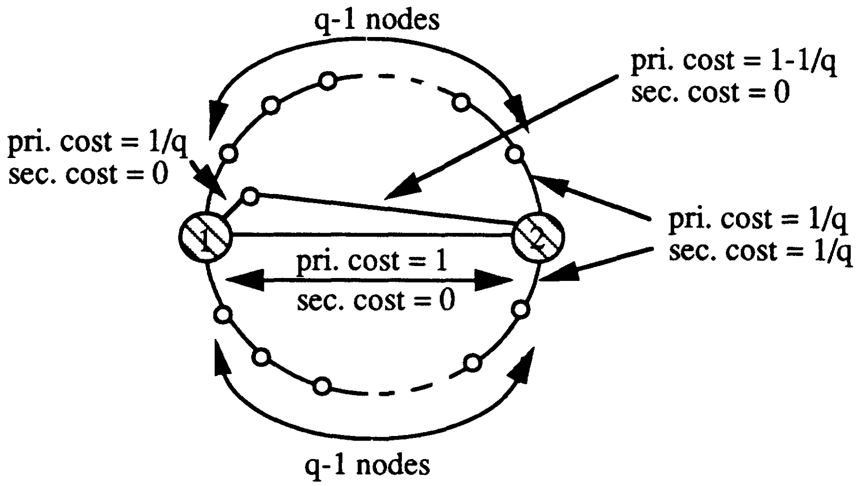

Figure 6: Worst-case example for DP-on-DPT problem with unrelated costs

Using Dual Path Greedy Completion heuristic to solve the DPT base subproblem

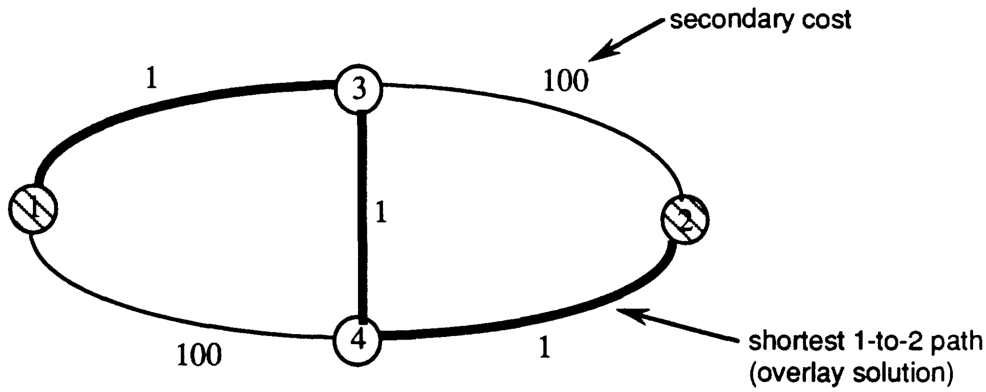

Figure 7(a): Shortest 1-to-2 path has no feasible overlay completion
primary-to-secondary cost ratio r = 2 for all edges

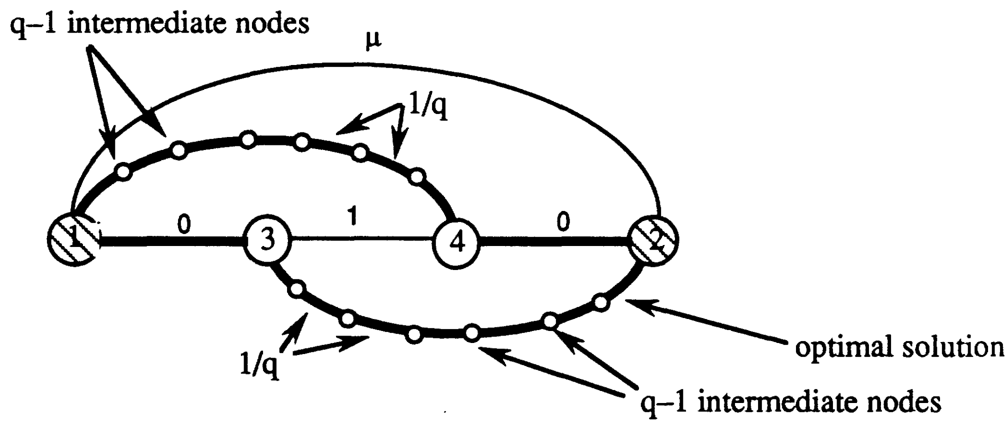

Figure 7(b): Completion cost multiplier for μ-direct graphs
secondary costs shown above
primary-to-secondary cost ratio = 1 for all edges

Figure 7: Overlay Completion Examples for SP-on-DPT problems

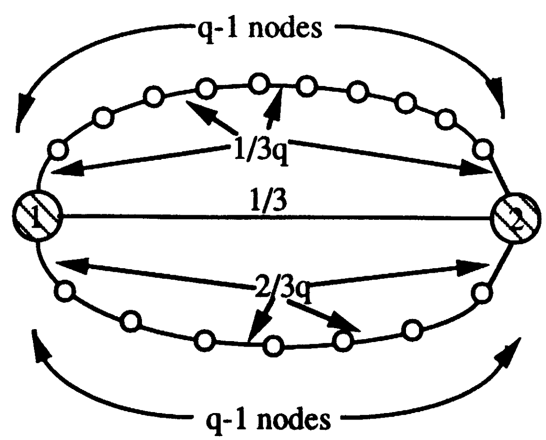

Figure 8(a): SP-on-DPT example with triangular costs
cost ratio = 2; secondary cost shown on edges

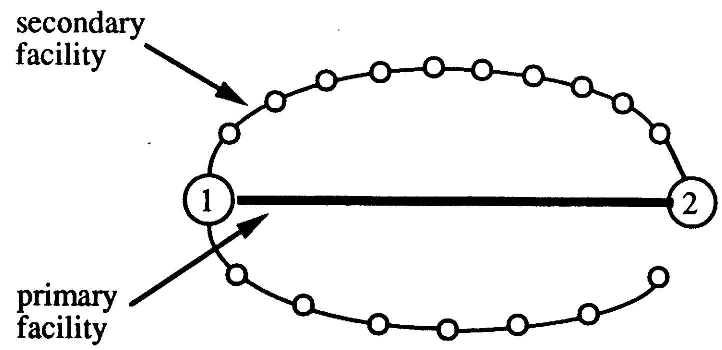

Figure 8(b): OC heuristic solution

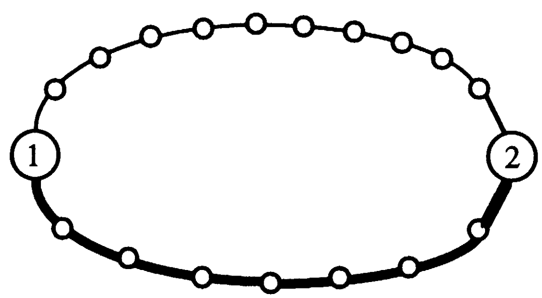

Figure 8(c): Modified BU heuristic solution

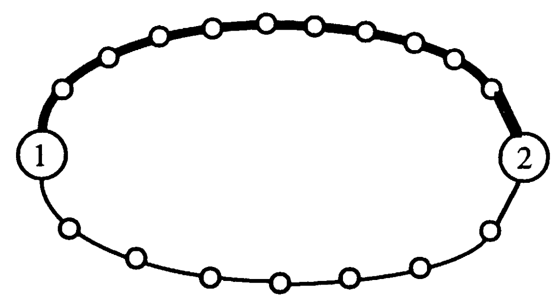

Figure 8(d): Optimal solution

Figure 8: Worst-case example for SP-on-DPT problem with proportional, triangular costs
DPT base subproblem solved optimally

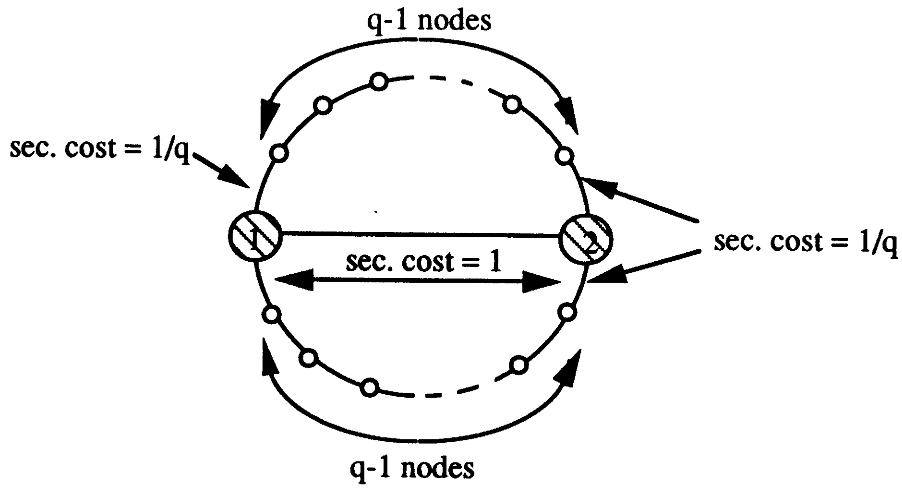

Figure 9(a): SP-on-DPT example with triangular costs
cost ratio = 1;

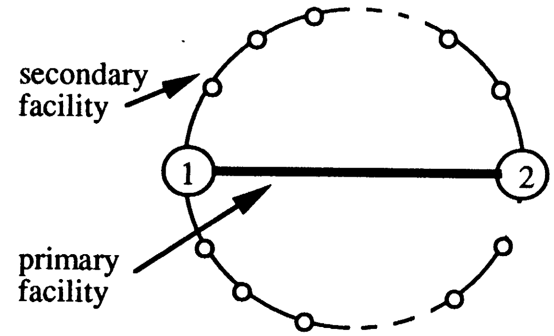

Figure 9(b): OC heuristic solution

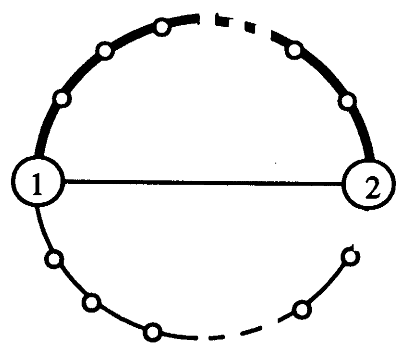

Figure 9(c): BU heuristic solution

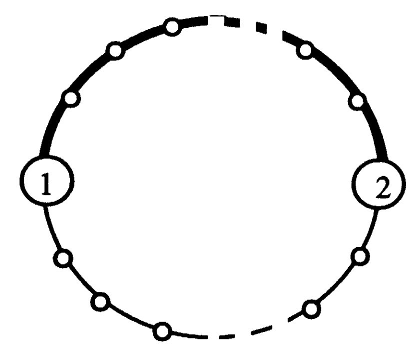

Figure 9(d): Optimal solution

Figure 9: Worst-case example for SP-on-DPT problem with proportional, triangular costs
Using 1-tree heuristic to solve the DPT base problem

References

AHUJA, R., T. L. MAGNANTI, and J. B. ORLIN. 1993. *Network Flows: Theory, Algorithms, and Applications*, Prentice-Hall, Englewood Cliffs, New Jersey.

BALAKRISHNAN, A., T. L. MAGNANTI, and P. MIRCHANDANI. 1992a. Modeling and Worst-case Performance Analysis of the Two-level Network Design Problem. Working Paper # 3498-92-MSA, Sloan School of Management, Massachusetts Institute of Technology, Cambridge, to appear in *Management Science.*

BALAKRISHNAN, A., T. L. MAGNANTI, and P. MIRCHANDANI. 1992b. A Dual-based Algorithm for Multi-level Network Design. Working Paper # 3365-91-MSA, Sloan School of Management, Massachusetts Institute of Technology, Cambridge, to appear in *Management Science.*

BALAKRISHNAN, A., T. L. MAGNANTI, and P. MIRCHANDANI. 1994a. Heuristics, LPs, and Trees on Trees: Network Design Analyses. Working Paper, Sloan School of Management, Massachusetts Institute of Technology, Cambridge.

BALAKRISHNAN, A., T. L. MAGNANTI, and P. MIRCHANDANI. 1994b. Designing Low Connectivity Steiner Networks. In preparation.

BERMAN, S., and V. RAMAIYER. 1992. An Approximation Algorithm for the Steiner Tree Problem. *Proc. of the Third ACM-SIAM Symposium on Discrete Algorithms.*

CORNUÉJOLS, G., J. FONLUPT, and D. NADDEF. 1985. The Traveling Salesman Problem on a Graph and Some Related Polyhedra. *Math. Prog.,* **33**, 1-27.

CURRENT, J. R., C. S. REVELLE, and J. L. COHON. 1986. The Hierarchichal Network Design Problem. *Eur. J. Oper. Res.,* **27**, 57-66.

EDMONDS, J. 1979. Matroid Intersection. *Annals of Discrete Mathematics*, **4**, 39-49.

FORD, L. R., and D. R. FULKERSON. 1962. *Flows in Networks*. Princeton University Press, Princeton, New Jersey.

GOEMANS, M. X., and D. J. BERTSIMAS. 1993. Survivable Networks, Linear Programming Relaxations and the Parsimonious Property. *Math. Prog*., **60**, 145-166.

GRÖTSCHEL, M., C. L. MONMA, and M. STOER. 1992. Computational Results with Cutting Plane Algorithm for Designing Communication Networks with Low-connectivity Constraints. *Oper. Res.,* **40**, 309-330.

MAGNANTI, T. L., and S. RAGHAVAN. 1992. A Dual-Ascent Procedure for Network Design with Connectivity Requirements. ORSA/TIMS Joint National Meeting, San Francisco, November 1992.

MONMA, C. L., B. S. MUNSON, and W. R. PULLEYBLANK. 1990. Minimum-weight Two-connected Spanning Networks. *Math. Prog,* **46**, 153-171.

TAKAHASHI, H., and A. MATSUYAMA. 1980. An Approximation Solution for the Steiner Problem in Graphs. *Mathematica Japonica*, **24**, 573-577.