

Designing Least Cost Nonblocking Broadband Networks

J. Andrew Fingerhut Subhash Suri Jonathan S. Turner

Department of Computer Science
Washington University
St. Louis, MO 63130
{jaf, jst, suri}@cs.wustl.edu

May 5, 1997

Abstract

Integrated network technologies, such as ATM, support multimedia applications with vastly different bandwidth needs, connection request rates, and holding patterns. Due to their high level of flexibility and communication rates approaching several gigabits per second, the classical network planning techniques, which rely heavily on statistical analysis, are less relevant to this new generation of networks. In this paper, we propose a new model for broadband networks and investigate the question of their optimal topology from a worst-case performance point of view. Our model is more flexible and realistic than others in the literature, and our worst-case bounds are among the first in this area. Our results include a proof of intractability for some simple versions of the network design problem, and efficient approximation algorithms for designing nonblocking networks of provably small cost. More specifically, assuming some mild global traffic constraints, we show that a minimum-cost nonblocking *star* network achieves near-optimal cost; the cost ratio is at most 2 if switch source and sink capacities are symmetric, and at most 3 when the total source and sink capacities are balanced. In the special case of unit link costs, we can show that a star network is indeed the cheapest nonblocking network.

1 Introduction

We consider optimization and combinatorial issues arising in the planning and design of modern broadband digital networks. In order to address these questions from a complexity-theoretic viewpoint, we first propose a new traffic model for designing these networks. Our model is less restrictive and more flexible than other models in several key aspects; it emphasizes the parameters that are most reliably available and eliminates

the need for overly specific traffic description. Our model is inspired by and builds upon the classical theory of nonblocking switching networks developed by Beneš [1], Clos [3], Pippenger [12], among others, and generalized to multirate switching networks by Melen and Turner [10]. The present paper differs from the work in switching networks in that it addresses the design of networks with irregular topologies and traffic characteristics, and it takes into account the costs of transmission links spanning substantial geographical distances. It also differs from prior work in topological design of networks in allowing a much less constrained and detailed specification of traffic requirements [8, 9, 11]. One can think of our model as implicitly allowing the specification of a very large number of traffic matrices.

Our research is motivated by the new generation of packet-based, virtual-circuit networks, such as ATM (Asynchronous Transfer Mode), capable of supporting high speed multimedia applications. ATM networks differ from telephone and wide area data networks in several ways. First, they are multirate networks, meaning that their virtual circuits can operate at any bandwidth, ranging from a few bits per second to over one hundred megabits per second. These networks also promise to support a wide range of applications with different bandwidth needs, different connection request rates and different holding times. Furthermore, unlike traditional data networks, ATM networks need to provide connections with a guaranteed quality of service, requiring allocation of bandwidth to individual virtual circuits and raising the possibility of virtual circuit blocking. Second, ATM networks support *multipoint* connections, not just point-to-point virtual circuits. Multipoint virtual circuits are essential for applications like video distribution or multimedia conferencing and include both one-to-many and many-to-many transmission patterns. The design of optimal networks supporting multipoint virtual circuits is largely uncharted territory, although some point-to-point results can be usefully generalized to the multipoint environment. Finally, ATM networks are much less predictable than telephone networks or traditional low speed data networks. There is no reliable statistical data on application characteristics and connection request patterns. Indeed, the very flexibility that is ATM's greatest strength also makes it highly unpredictable, and so classical network planning techniques, which rely heavily on statistical analysis, become less relevant. In ATM networks, the whole notion of blocking probability for virtual circuit setup must be called into question, since there is no reasonable possibility of validating the probabilistic assumptions that must go into any analysis of blocking probability.

Therefore, we believe that the network models that rely on specifying demands between all node-pairs are overly restrictive for designers of these networks and, given all the uncertainties in the usage of these networks, may completely miss the mark. We propose that a more reliable and meaningful parameter is the total amount of traffic *entering* or *leaving* a switch, without regard to the destination or source of this traffic. These numbers can often be estimated by looking only at the *devices* that use a particular switch as their network gateway. Within this global traffic model, we study the network design problem with the *linear link costs*: the cost of a link grows as a linear function of its capacity. In addition, we set the switch costs to zero; this is not a major limitation, since the switch costs can be distributed among its incident links. (See Concluding Remarks for

more on this issue.) Under this assumption, we are able to prove *worst-case* results for a variety of network design cases, where almost no such results were known previously. On the theoretical side, our results develop some mathematical techniques that might be useful in establishing lower bounds on the costs of optimal nonblocking networks in other contexts as well. On the practical side, our theorems can also be viewed as lending mathematical support to some commonly employed network topologies. We postpone a precise statement of our results until after we have formally introduced our model and formulated the problem (Section 2.2).

This paper has six sections. Section 2 introduces the necessary definitions, formalizes the network design problem considered in this paper, and briefly summarizes our main results. Section 3 addresses the computational complexity of the problem, and shows that the problem is NP-Complete. Section 4 describes our approximation techniques. Section 5 describes an optimal network for the special case of unit link costs. Section 6 provides some closing remarks and discussion of some practical network design issues.

2 Our Network Model

Our formulation of the network design problem consists of a complete digraph, $G = (V, E)$, where each vertex represents a *switch* and each directed edge represents a *link group*, comprising one or more physical transmission links. The vertices and edges of G have the following parameters associated with them:

- Each vertex u has an integer *source capacity* $\alpha(u)$, and an integer *sink capacity* $\omega(u)$, representing the maximum traffic rate that can originate or terminate at u .
- Each vertex pair (u, v) has a function $\gamma_\ell(u, v, x)$ representing the cost of constructing a link of capacity x from u to v .¹

We also have a switch cost function $\gamma_s(x)$ giving the cost of a switch of total capacity x . If we assign a capacity $\kappa_\ell(u, v)$ to every edge (u, v) the resulting network cost is defined as

$$\sum_{(u,v) \in E} \gamma_\ell(u, v, \kappa_\ell(u, v)) + \sum_{u \in V} \gamma_s(\kappa_s(u)), \quad (1)$$

where $\kappa_s(u)$ is the capacity of switch u in the network and it equals

$$\alpha(u) + \omega(u) + \sum_{v \in V, v \neq u} (\kappa_\ell(u, v) + \kappa_\ell(v, u)).$$

Thus, our model does not constrain traffic on a switch-pair basis, rather only at individual switches. The latter data is not only available more reliably, but it also gives the network designers more flexibility.

In order to define the notion of nonblocking networks, we first need to define connection requests and their routing in the network. A *connection request* $R = (S, D, w)$ comprises

¹We require that the costs satisfy the *triangle inequality*, meaning that the direct path of any given capacity between two vertices is never more expensive than an indirect path with the same capacity.

a non-empty set of *sources* S , a non-empty set of *destinations* D and an integer *weight* $w \leq B$, where B is a maximum connection weight. A *route* T for a request R is a subgraph of G for which the underlying undirected graph is a tree and in which there is a directed path from every vertex in S to every vertex in D . A collection of routes C places a connection weight $\lambda_C(u, v)$ on an edge (u, v) , which is defined as the sum of the weights of all routes that include the edge (u, v) . $\lambda_C(u)$ denotes the weight on a switch u , which is equal to the sum of the weights of its incident edges.

A set of connection requests is *valid* if, for every vertex u , the sum of the weights of the requests containing u in their source and sink sets, respectively, does not exceed $\alpha(u)$ and $\omega(u)$. A collection of routes C is *valid* if it satisfies a set of valid connection requests, and if $\lambda_C(u, v) \leq \kappa_\ell(u, v)$, for every edge (u, v) , and if $\lambda_C(u) \leq \kappa_s(u)$, for every vertex u .

A *state* of a network is a valid set of routes. A *routing algorithm* is a procedure that maintains a valid set of routes under the following four operations: (1) add a new route satisfying a specified connection request; (2) remove an existing route; (3) add a new vertex to either the source set, the destination set, or both for some route in the current state; (4) remove some vertex from either the source set, the destination set, or both for some route in the current state.² We are only concerned with routing algorithms that are *incremental*, meaning that they only add, delete or modify a single route when carrying out a requested operation and that they cannot both add and remove edges from an existing route in a single operation.

The *reachable states* for a routing algorithm on a network with specified link capacities is the set of all states that can be reached by sequences of the four operations given above, starting from the empty state. We say that a network is *nonblocking* under a given routing algorithm if for every reachable state and every operation request whose completion would not exceed the source or sink capacity of any vertex in that state, the algorithm produces a new state satisfying the operation request.

2.1 The Nonblocking Network Design Problem

The *nonblocking network design problem* is to *determine a set of link capacities* that will yield a nonblocking network of least cost under either a specified routing algorithm or some routing algorithm from a specified class of routing algorithms. In the latter case, the design problem is to produce both the link capacities and a specific routing algorithm from the given class, for which the network is nonblocking. Figure 1 shows an instance of the network design problem on the left and a solution on the right. On the left, the numbers next to each vertex denote the switch capacities, $\alpha(v), \omega(v)$; the number next to each edge denotes the link cost per unit capacity (assuming symmetric link costs). The solution on the right shows directional capacities on links. This network is nonblocking if connections are always routed using shortest available paths.

In many situations, some special cases of the network design problem are of interest. In the *linear cost* version, switch costs are zero and all link costs satisfy $\gamma_\ell(u, v, x) =$

²A routing algorithm may fail to carry out operations of type (1) or (3), but will always carry out operations of type (2) or (4).

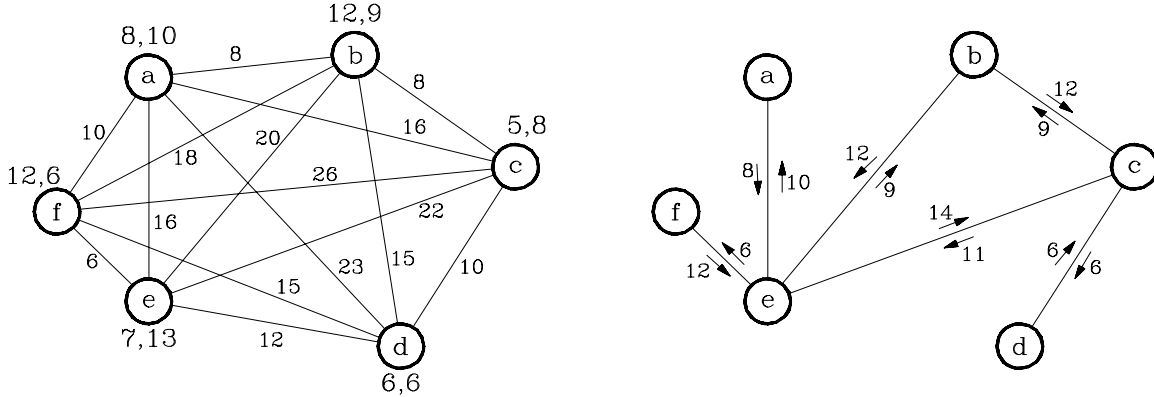


Figure 1: An example of the design problem and a suboptimal solution

$x \times \gamma(u, v)$, where $\gamma(u, v)$ is a constant that depends only on u and v . The *symmetric* version of the problem has $\alpha(u) = \omega(u)$ for all vertices u , $\gamma_\ell(u, v, x) = \gamma_\ell(v, u, x)$ for all pairs u, v and restricts the choice of link capacities so that $\kappa_\ell(u, v) = \kappa_\ell(v, u)$. In the *balanced* version of the problem, we have $\sum_{u \in V} \alpha(u) = \sum_{u \in V} \omega(u)$.

2.2 A Summary of Our Results

In this paper, we focus on the linear link cost model, that is, $\gamma_\ell(u, v, x) = x \times \gamma(u, v)$ and switch costs are zero. In this model, we prove that networks of star topology achieve near-optimal cost. In particular, for the symmetric case, we prove that the least cost nonblocking network of an *arbitrary* topology has cost at least half the cost of the cheapest nonblocking star network. The ratio becomes 1/3 when the source and sink traffic capacities are asymmetric, but balanced. For arbitrary traffic capacities, the performance ratio of the star networks degrades gracefully (cf. Theorem 4.11). Finally, we show that in the special case of *unit* link cost function, meaning $\gamma(u, v, x) = c \cdot x$ for some absolute constant c , a star network is indeed optimal.

Even in our simplified linear link cost model, the problem of computing a least-cost nonblocking network turns out to be NP-Complete, meaning that approximation algorithms are the only recourse for designing nonblocking networks of provably good cost ratios. We show several hardness results. In the usual RAM model of computation, we can compute a least-cost nonblocking star network in $O(n^2)$ time, where n is the number of switches. We start by addressing the computational complexity question first.

3 Computational Complexity of the Problem

A solution to the network design problem asks for a cheapest set of link capacities as well as an incremental strategy for setting up valid connections. In general, the routing problem in itself is a hard problem. In networking literature, a variety of routing strategies are used: (1) fixed path routing, (2) alternate path routing, and (3) shortest available path routing. In the fixed path routing, a precomputed routing table stores a directed

path between each pair of nodes (u, v) , and if this path has insufficient bandwidth to add a connection request from u to v , the connection is refused. In alternate path routing, each node pair has two paths, a primary path and a secondary path, and the secondary path is tried in case the primary path is unable to route the connection. If the secondary one cannot route the connection either, the connection blocks.

Both the fixed path and alternate path routing algorithms can block a connection even when there exists a path of sufficient bandwidth between u and v . The shortest available path algorithm is able to route a connection as long as there is *some* path in the network of sufficient bandwidth. As its name suggests, the shortest available path algorithm uses, among all valid paths, a least cost path. It is not difficult to see that even with shortest available path routing, connections can block when a simple *rearrangement* of routes will free up enough bandwidth to accept the new connection. The following theorem, proved in [4], shows that determining whether a given network blocks some sequence of switch capacity-compliant connection requests is NP-Hard.

Theorem 3.1 [4] *Let V be a set of switches, let $\alpha(u), \omega(u)$ be their source and sink capacities, and let $\kappa_\ell(u, v)$ be the capacity of link (u, v) . Suppose that all connection requests have weights that are multiples of a minimum weight b and connections are routed using the shortest available path. Then the problem of deciding whether a sequence of point-to-point requests compatible with the switch capacities blocks is NP-Hard in the strong sense.*

The intractability of *checking* whether a network is nonblocking does not imply that *designing* one is also hard. However, we show below that several simple versions of the design problem are indeed intractable. In the first theorem, we let the source and sink capacities be arbitrary, with no constraints of symmetry or balance. In this version, the well-known Steiner tree problem in graphs with edge costs in $\{1, 2\}$ turns out to be special case of the network design problem. (Let $G = (V, E)$ be a graph, $w(e) \in \{1, 2\}$ be weights on edges, $R \subset V$ be a subset, and B a positive integer bound. The Steiner tree problem asks if there is a subtree of G spanning all nodes of R with a total cost of at most B ? This version of the Steiner tree problem was proved MAXSNP-Hard by Bern and Plassman [2]; see also the book by Garey and Johnson [7].)

Theorem 3.2 *Given a set of switches V , their source and sink capacities $\alpha(v), \omega(v)$, and a linear link cost function $\gamma(u, v)$ for each switch-pair in V , the problem of finding a minimum cost, nonblocking network for $(V, \alpha, \omega, \gamma)$ is MAXSNP-Hard.*

PROOF. The set of switches V is the set of nodes V . The link costs are the same as the edge costs in G , namely, $\gamma(u, v) = w(u, v)$. Observe that the link costs satisfy triangle inequality. We pick an arbitrary “root” node $r \in R$, from the Steiner subset $R \subset V$. Set $\alpha(r) = 1$ and $\omega(r) = 0$. For the remaining Steiner nodes $u \in R$, we set $\alpha(u) = 0$ and $\omega(u) = 1$. All other nodes of V have $\alpha(v) = \omega(v) = 0$, where $v \in V - R$. Thus, the root node can originate one unit of traffic, but it has no termination capacity. Every other node of the Steiner subset R has one unit of termination capacity and no origination capacity. The nodes in $V - R$ have no origination/termination capacity at all.

Let \mathcal{N}^* be a minimum cost nonblocking network for the above instance. It is easily seen that V admits a Steiner tree of cost B on R if and only if $\text{cost}(\mathcal{N}^*) \leq B$. In particular, every Steiner tree spanning R can be turned into a nonblocking network, by directing all edges away from the root node and assigning unit capacity to each link. Conversely, every nonblocking network is a Steiner tree of R . \square

Next, we show that the network design problem even with symmetric switch capacities is hard, for a slight variation of the linear link cost model. In particular, assume that setting up a link from u to v of capacity $\kappa_\ell(u, v)$ has cost

$$c(u, v) + \gamma(u, v) \times \kappa_\ell(u, v),$$

where $c(u, v)$ is a fixed installation cost, independent of the link capacity. In this case, we show a polynomial time reduction from the following well-known *set cover problem* to the network design problem:

Given a finite set X and a family $\mathcal{F} = \{S_1, S_2, \dots, S_m\}$ of subsets of X , find a *minimum cardinality* subset $J \subseteq \{1, 2, \dots, m\}$ such that $\cup_{j \in J} S_j = X$.

We construct a bipartite graph with elements of X and \mathcal{F} as node classes, and put an edge between x and S if $x \in S$. Thus, we have n nodes labeled x_1, x_2, \dots, x_n , and m nodes labeled S_1, S_2, \dots, S_m . There is an edge between x_i and S_j if and only if $x_i \in S_j$. Finally, we add a new node x_0 that is joined to each set S_j , $j = 1, 2, \dots, m$. We assign capacities as follows:

$$\begin{aligned} \alpha(x_i) &= 1, & \omega(x_i) &= 0 & i &= 1, 2, \dots, n \\ \alpha(S_j) &= 0, & \omega(S_j) &= 0 & j &= 1, 2, \dots, m \\ \alpha(x_0) &= 0, & \omega(x_0) &= m \end{aligned}$$

The link costs are defined as follows:

$$\begin{aligned} c(x_i, S_j) &= 1, & \gamma(x_i, S_j) &= 1 & \text{if } x_i \in S_j \\ c(S_j, x_0) &= 1, & \gamma(S_j, x_0) &= 0 & j = 1, 2, \dots, m \end{aligned}$$

All other link costs are defined by using the shortest path metric in this graph. (With an appropriate interpretation of what it means to set up a link along some shortest path, these costs obey the triangle inequality.) Figure 2 illustrates the construction.

Lemma 3.3 *The network design problem formulated above has a solution of cost $2n + |J^*|$ if and only if (X, \mathcal{F}) admits a set cover of cardinality $|J^*|$.*

PROOF. In order to prove the “if” part, consider a set cover of size J^* , and let $S^{*}_1, S^{*}_2, \dots, S^{*}_{J^*}$ denote the member sets of this cover. We can construct a nonblocking network of cost $2n + |J^*|$ as follows: for each x_i , assign capacity one to the edge (x_i, S^{*}_j) , where j is the lower-indexed set containing x_i . (Since S^{*}_j ’s form a set cover, each x_i is joined to

some S_j^* by this rule.) Next, for each S_j^* , assign capacity $|S_j^*|$ to the edge (S_j^*, x_0) . The total cost of this network is $2n + \sum_{j \in J^*} 1 = m + |J^*|$, completing the forward implication.

In order to prove the “only if” part, consider a nonblocking network \mathcal{N}^* of cost at most $2n + |J^*|$. We will exhibit a set cover of size $|J^*|$. All the connections of the form (x_i, S_j) must cost at least $2n$, since each x_i is connected to some S_j , and each such link has a fixed cost of 1 and a per unit cost of 1. Thus, the number of connections of the form (S_j, x_0) is no more than $|J^*|$. Since all x_i must be able to reach x_0 via some S_j , these S_j 's must be joined to all x_i 's, and hence form a set cover of X . This completes the proof. \square




Figure 2: The proof of NP-Completeness.

We have the following theorem.

Theorem 3.4 *Let V be a set of switches, let $\alpha(v)$, $\omega(v)$ be their source and sink capacities, $\gamma(u, v)$ be the linear link cost for each switch-pair (u, v) , and let c be a fixed cost per link. Then, the problem of finding a minimum cost, nonblocking network is NP-Complete even with $\alpha(u) = \omega(u)$ for each node and $c \in \{0, 1\}$.*

We point out that our approximation results hold for the modified link cost function of the preceding theorem. In view of these hardness results, we focus our attention on efficient algorithms for designing nonblocking networks of provably small cost.

4 Designing Low-Cost Nonblocking Networks

We show that *star networks* produce nearly optimal results. In particular, we prove that there exists a star network, rooted at one of the nodes of V , that is nonblocking and has a cost at most twice the minimum cost in the symmetric case (i.e., $\alpha(v) = \omega(v)$). In the balanced case, the same network is also shown to be within a factor 3 of optimal. As the balance condition worsens, the quality of approximation degrades gracefully: we prove that there is a star network with cost no more than $2 + \frac{\sum \alpha(u)}{\sum \omega(u)}$ times the optimal, where we assume without loss of generality that $\sum \alpha(u) \geq \sum \omega(u)$. An optimal nonblocking star can

be found algorithmically in $O(n^2)$ time, where n is the number of switches. (The routing strategy for star networks is obvious: use the unique path between two communicating nodes.)

We will bound the cost of an optimal star network in terms of a quantity \mathcal{D} defined below, and then derive a lower bound on the cost of a cheapest nonblocking network also in terms of \mathcal{D} to establish our results. We will frequently need to refer to the total source and sink capacities. For convenience, let us introduce the following shorthand notation:

$$\mathcal{A} = \sum_{v \in V} \alpha(v) \quad \text{and} \quad \mathcal{Z} = \sum_{v \in V} \omega(v).$$

Throughout the following discussion, we assume without loss of generality that $\mathcal{A} \geq \mathcal{Z}$. The quantity \mathcal{D} is defined as follows:

$$\mathcal{D} = \sum_u \sum_v \alpha(u) \times \omega(v) \times \gamma(u, v). \quad (2)$$

We are now ready to proceed with our proof of the approximation bound; we first establish the general upper bound, and then sharpen it further for the symmetric case of switch capacities.

4.1 General Switch Capacities

In establishing the upper bound, we use an intermediate network that has the form of a double star. The double star $S(v_k, v_l)$ corresponding to an ordered pair (v_k, v_l) is defined by the following link capacities:

1. $\kappa(v_i, v_l) = \alpha(v_i)$, for $i \neq l$;
2. $\kappa(v_k, v_i) = \omega(v_i)$, for $i \neq k$;
3. $\kappa(v_l, v_k) = \mathcal{Z} - \omega(v_l)$.

All other links in $S(v_k, v_l)$ have zero capacity. See Figure 3 for an illustration. We will show that the cheapest double star achieves the desired cost. But first let us show that the double star described above is indeed a nonblocking network.

Lemma 4.5 *The double star $S(v_k, v_l)$ is a nonblocking network for $(V, \alpha, \omega, \gamma)$.*

PROOF. The link (v_l, v_k) clearly has sufficient bandwidth to route all valid connections, since the maximum traffic to all receiving switches, other than v_l itself, cannot exceed $\mathcal{Z} - \omega(v_l)$. Since each v_i has outgoing link capacity $\alpha(v_i)$ and each v_j has incoming link capacity $\omega(v_j)$, it is easily seen that no valid connection request is blocked. \square

In order to complete our proof of the approximation bound, we show below that there exists a double-star in $B(V)$ whose cost is within a factor $2 + \frac{\mathcal{A}}{\mathcal{Z}}$ of the cost of an optimal network.




Figure 3: Illustrating a double star.

Lemma 4.6 *A minimum cost double star of V has cost no greater than*

$$\frac{(\mathcal{A} + 2\mathcal{Z})\mathcal{D}}{\mathcal{A} \cdot \mathcal{Z}}.$$

PROOF. We prove the lemma by considering a multiset of double stars of V and arguing that the cost of an average double star in this multiset has the claimed bound. Since the minimum of a set cannot exceed its average, the lemma follows. So, let \mathcal{M} denote the multiset of double stars, in which $S(v_k, v_l)$ appears $\alpha(v_k) \times \omega(v_l)$ times. The family \mathcal{M} has size

$$\begin{aligned} |\mathcal{M}| &= \sum_{k=1}^{|V|} \sum_{l=1}^{|V|} \alpha(v_k) \times \omega(v_l) \\ &= \mathcal{A} \times \mathcal{Z}. \end{aligned} \tag{3}$$

Let us now count the total cost of all the double stars in this multiset. We do this by counting the contribution of each edge (v_k, v_l) , and summing over all pairs. An edge (v_k, v_l) contributes costs in three ways:

1. In the double star $S(v_k, v_l)$, the oppositely directed edge (v_l, v_k) has capacity $\mathcal{Z} - \omega(v_l)$. This double star appears $\alpha(v_k) \times \omega(v_l)$ times, and, by symmetry of the link cost, $\gamma(v_k, v_l) = \gamma(v_l, v_k)$. Thus, the the total contribution is

$$\alpha(v_k)\omega(v_l)\gamma(v_l, v_k) (\mathcal{Z} - \omega(v_l)).$$

2. In each of the double stars $S(v_k, v_j)$, it appears with capacity $\omega(v_l)$. The total number of these double stars in \mathcal{M} is $\alpha(v_k) \times \sum_{j=1}^{|V|} \omega(v_j)$, which implies the total contribution at most

$$\alpha(v_k)\omega(v_l)\gamma(v_k, v_l) \times \mathcal{Z}$$

3. In each of the double stars $S(v_i, v_l)$, it appears with capacity $\alpha(v_k)$. The total number of these double stars in \mathcal{M} is $\omega(v_l) \times \sum_{i=1}^{|V|} \alpha(v_i)$, which implies the total contribution at most

$$\alpha(v_k)\omega(v_l)\gamma(v_k, v_l) \times \mathcal{A}.$$

Recalling that $\mathcal{D} = \sum_k \sum_l \alpha(v_k)\omega(v_l)\gamma(v_k, v_l)$, we obtain that the total cost of all the double star in \mathcal{M} is at most

$$(\mathcal{A} + 2\mathcal{Z})\mathcal{D}. \quad (4)$$

Thus, we get an upper bound on the cost of an average double star in \mathcal{M} by dividing the quantity in Eq. (4) by the quantity in Eq. (3), which gives the bound claimed in the lemma. This completes the proof. \square

Finally, we show that triangle inequality implies that the cost of a cheapest nonblocking star cannot exceed the cost of a cheapest double star. In particular, we show that the double star $S(v_k, v_l)$ can be converted to a star rooted at v_l with no increase in cost. In the double star $S(v_k, v_l)$, we leave all incoming links of v_l the same, but transfer all outgoing links of v_k to v_l . Clearly, this yields a nonblocking star rooted at v_l . The following lemma proves the bound on the cost.

Lemma 4.7 *The cost the cheapest nonblocking star rooted at v_l does not exceed the cost of $S(v_k, v_l)$.*

PROOF. In modifying the double star into the star network, we effectively replace the path (v_l, v_k, v_i) with the direct path (v_l, v_i) . The total capacity of all outgoing links at v_l is $\mathcal{Z} - \omega(v_l)$. By triangle inequality,

$$(\mathcal{Z} - \omega(v_l)) \times \gamma(v_l, v_k) + \sum_{i \neq l} \omega(v_i) \times \gamma(v_k, v_i) \leq \sum_{i \neq l} \omega(v_i) \times \gamma(v_l, v_i).$$

The leftover term $\omega(v_l)\gamma(v_l, v_k)$ is used to charge the link from v_l to v_k of capacity $\omega(v_l)$. This completes the proof. \square

4.2 An Improved Bound for Symmetric Switch Capacities

In this case, we can directly bound the cost of an optimal star network. Consider the least cost nonblocking star rooted at node u . It has cost

$$\sum_{v \neq u} (\alpha(v) + \omega(v))\gamma(u, v) = \sum_{v \neq u} 2\alpha(v)\gamma(u, v),$$

where we use the fact that $\alpha(v) = \omega(v)$. Let \mathcal{M} denote the multiset of stars in which the star with root u appears exactly $\alpha(u)$ times. An edge (u, v) contributes a total cost of $2\gamma(u, v)$ in each of the $\alpha(u) + \alpha(v)$ stars. Thus, the total cost of all the stars in \mathcal{M} is

$$\sum_u \sum_v 2\gamma(u, v) \times \alpha(u) \times \omega(v) = 2\mathcal{D}.$$

Since $|\mathcal{M}| = \mathcal{A}$, it follows that the cheapest star in \mathcal{M} has cost no more than $2\mathcal{D}/\mathcal{A}$, giving the following lemma.

Lemma 4.8 *Let V be a set of switches, with symmetric source and sink capacities, $\alpha(v) = \omega(v)$, and link costs $\gamma(u, v)$ for all switch-pairs u, v . Then, the cost of a cheapest nonblocking star network for $(V, \alpha, \omega, \gamma)$ is at most $\frac{2\mathcal{D}}{\mathcal{A}}$.*

In order to show that these star networks are near optimal, we need to establish a lower bound on the cost of any nonblocking network. We do this in the following subsection.

4.3 A Lower Bound on the Cost of an Optimal Network

Suppose \mathcal{N}^* is a nonblocking network for the switch capacities $\alpha(v)$, $\omega(v)$ and link costs $\gamma(u, v)$, where $u, v \in V$. Being a nonblocking network, \mathcal{N}^* is able to route any set of switch capacity-compliant connections. Consider a feasible connection between u and v at data rate $f(u, v)$, where feasibility dictates that $f(u, v) \leq \min\{\alpha(u), \omega(v)\}$. Then, by triangle inequality, the route(s) used by \mathcal{N}^* to set up this connection must cost at least $\gamma(u, v) \times f(u, v)$. Now, if there are two simultaneously feasible connections, one from u to v at rate $f(u, v)$ and another from x and y at rate $f(x, y)$, then the *linearity* of link costs implies that the network has cost at least

$$\gamma(u, v) \times f(u, v) + \gamma(x, y) \times f(x, y). \quad (5)$$

Thus, any set of simultaneously feasible connections implies a lower bound of the form Eq. (5) on the cost of \mathcal{N}^* . In order to get the best lower bound, we seek connections of maximum cost.

The problem of finding a set of simultaneous connections maximizing the cost is essentially a *maximum-cost multi-commodity* flow problem. However, for our purpose, we are interested in a quantitative, and not numerical, estimate of the cost. In particular, we would like a lower bound in terms of the quantity \mathcal{D} , so as to relate it to the upper bound of the preceding subsections. One possibility to derive such a lower bound is to use a maximum cost *matching* in the network. But, due to varying switch capacities, a valid connection between u and v has rate at most $\min\{\alpha(u), \omega(v)\}$. We, therefore, may need to set up multiple connections from u to exhaust its capacity. In order to find these multiple connection conveniently, we carry out a node-splitting transformation, which splits a node u into $\alpha(u)$ source nodes and $\omega(u)$ sink nodes, each with unit capacity.

More formally, let $v_i \in V$ be a switch with source capacity $\alpha_i = \alpha(v_i)$ and sink capacity $\omega_i = \omega(v_i)$. We replace v_i with α_i copies of itself labeled *source nodes* $a_{i1}, a_{i2}, \dots, a_{i\alpha_i}$, and with ω_i copies labeled *sink nodes* $z_{i1}, z_{i2}, \dots, z_{i\omega_i}$. Assign $\alpha(a_{ij}) = 1$ and $\omega(a_{ij}) = 0$, and $\alpha(z_{ij}) = 0$ and $\omega(z_{ij}) = 1$. Thus, each source node has send capacity of one and receive capacity of zero, while each sink node has the send capacity of zero and receive capacity of one. Construct a bipartite graph by joining each a -node to each z -node and “inheriting” the link cost from the original problem. Specifically, we assign

$$\gamma(a_{ij}, z_{kl}) = \gamma(v_i, v_k), \quad \text{for } j = 1, 2, \dots, \alpha_i, \text{ and } l = 1, 2, \dots, \omega_i.$$

An example of our graph transformation is shown in Figure 4. We call this bipartite graph $B(V)$. Observe that $B(V)$ has $\mathcal{A} + \mathcal{Z}$ nodes and $\mathcal{A} \times \mathcal{Z}$ edges, where recall that

$\mathcal{A} = \sum_v \alpha(v)$ and $\mathcal{Z} = \sum_v \omega(v)$, and we assume that $\mathcal{A} \geq \mathcal{Z}$. In order to simplify the notation, let us renumber the nodes so that the source nodes are labeled $a_1, a_2, \dots, a_{\mathcal{A}}$, and the sink nodes are labeled $z_1, z_2, \dots, z_{\mathcal{Z}}$.




Figure 4: Illustrating the graph transformation. In the figure, $\alpha(v_1) = 2$, $\omega(v_1) = 1$, and $\alpha(v_2) = 1$, $\omega(v_2) = 2$. Edges (a_{11}, z_{11}) , (a_{12}, z_{11}) , (a_{21}, z_{11}) , and (a_{21}, z_{22}) have link costs zero; others have cost d .

Let M denote an arbitrary matching in $B(V)$; recall that a matching is a collection of vertex disjoint edges. We claim that a maximum-weight matching in $B(V)$ has weight at least \mathcal{D}/\mathcal{A} , where the weight of the matching is the total cost of its edges.

Lemma 4.9 *Let M be a maximum-weight matching in $B(V)$. Then, $\text{cost}(M) \geq \mathcal{D}/\mathcal{A}$.*

PROOF. First, observe that

$$\mathcal{D} = \sum_{i=1}^{\mathcal{A}} \sum_{j=1}^{\mathcal{Z}} \gamma(a_i, z_j);$$

this follows because the node-splitting transformation makes $\alpha(u) \times \omega(v)$ copies of the edge (u, v) . The number of different matchings in $B(V)$ is $\binom{\mathcal{A}}{\mathcal{Z}} \cdot \mathcal{Z}!$. (The first term counts the number of ways to pick which \mathcal{Z} source nodes to match with the sink nodes, and the second term counts the number of ways to do this matching.) Every edge of $B(V)$ gets counted $\binom{\mathcal{A}-1}{\mathcal{Z}-1} \cdot (\mathcal{Z} - 1)!$ times over all the matchings. Thus, the total weight of all the matchings is $\binom{\mathcal{A}-1}{\mathcal{Z}-1} \cdot (\mathcal{Z} - 1)! \cdot \mathcal{D}$. Since the maximum of a set is at least as large as its average, the maximum-weight matching satisfies

$$\begin{aligned} \text{cost}(M) &\geq \frac{\binom{\mathcal{A}-1}{\mathcal{Z}-1} \cdot (\mathcal{Z} - 1)! \cdot \mathcal{D}}{\binom{\mathcal{A}}{\mathcal{Z}} \cdot \mathcal{Z}!} \\ &= \frac{\mathcal{D}}{\mathcal{A}}, \end{aligned}$$

which completes the proof. □

A matching in $B(V)$ corresponds (uniquely) to a set of valid connections, with the same total cost as the matching, giving the following corollary.

Corollary 4.10 *Let V be a set of switches with source and sink capacities $\alpha(v), \omega(v)$, for $v \in V$, and assume that the link cost $\gamma(u, v)$, for all switch-pairs $(u, v) \in V \times V$, satisfies the triangle inequality. Then, a minimum-cost nonblocking network for $(V, \alpha, \omega, \gamma)$ has cost at least $\frac{\mathcal{D}}{\mathcal{A}}$.*

4.4 Approximation Ratios for Star Networks

Comparing the cost of a cheapest star network to the lower bound of Corollary 4.10, we can bound the approximation factor of our star network. The approximation factor is given by

$$\begin{aligned} \frac{\text{cost}(\text{cheapest star})}{\text{cost}(\mathcal{N}^*)} &\leq \frac{\frac{(\mathcal{A}+2\mathcal{Z})\mathcal{D}}{\mathcal{A}\mathcal{Z}}}{\frac{\mathcal{D}}{\mathcal{A}}} \\ &\leq \frac{\mathcal{A} + 2\mathcal{Z}}{\mathcal{Z}} \\ &\leq 2 + \frac{\mathcal{A}}{\mathcal{Z}} \\ &\leq 3 \quad \text{if } \mathcal{A} = \mathcal{Z}. \end{aligned}$$

Thus, in the balanced case, namely $\mathcal{A} = \mathcal{Z}$, there exists a nonblocking star network for the network design problem $(V, \alpha, \omega, \gamma)$ whose cost does not exceed three times the cost of an optimal network. Without any balance condition, the cost of the best star network is within $2 + \frac{\mathcal{A}}{\mathcal{Z}}$ times of the optimal. For the symmetric capacity case, the ratio of the star to optimal network is 2 (cf. Lemma 4.8). We conclude with the following theorem.

Theorem 4.11 *Let V be a set of switches, with source and sink capacities $\alpha(v)$ and $\omega(v)$, and link costs $\gamma(u, v)$ for all switch pairs u, v . Then, the ratio between the cost of a cheapest nonblocking star and an optimal network is at most 2 if the switch capacities are symmetric, at most 3 if the switch capacities are balanced, and at most $2 + \frac{\mathcal{A}}{\mathcal{Z}}$ in general, where $\mathcal{A} \geq \mathcal{Z}$.*

4.5 How tight is the lower bound?

We have shown that a nonblocking network of cost at most twice (resp. three times) the optimal can be found in polynomial time for the symmetric (resp. balanced) case of switch capacities. Whether these approximation factors can be improved, remains an open problem. We can exhibit examples, however, showing that the *ratio* between the cheapest star and the lower bound in Corollary 4.10 is tight.

Figure 5 shows an example where the ratio of maximum-cost matching to minimum-cost nonblocking star is $2 + \frac{\mathcal{A}-2}{\mathcal{Z}}$, which comes arbitrarily close to the bound stated in Theorem 4.11. Similarly, Figure 6 shows an example where the ratio is tight even for the symmetric case.

Consider the example shown in Figure 5. We have \mathcal{A} switches with $\alpha(v) = 1$ and $\omega = 0$, and \mathcal{Z} switches with $\alpha(v) = 0$ and $\omega = 1$. Call the switches in the former

group *source* nodes and the ones in the latter group *sink* nodes. We label the source nodes $u_1, u_2, \dots, u_{\mathcal{A}}$ and sink nodes $v_1, v_2, \dots, v_{\mathcal{Z}}$, and let us assume that $\mathcal{A} \geq \mathcal{Z} + 2$. To complete the description of the problem, we specify the link costs as follows:

$$\begin{aligned} \gamma(u_i, v_j) &= 1 & \forall i, j \\ \gamma(u_i, u_j) &= 2 & \forall i, j \\ \gamma(v_i, v_j) &= 2 & \forall i, j \end{aligned}$$




Figure 5: An example showing that the ratio between the cost of a cheapest star and a maximum-weight matching is tight. The links costs are 1 for a u - v pair, and 2 otherwise. A minimum-cost nonblocking star is also shown, having a cost of $\mathcal{A} + 2(\mathcal{Z} - 1)$.

Thus, the links joining a source node to a sink node have costs one; all others have cost two. Clearly, the link costs satisfy the triangle inequality. It is easy to see that maximum-weight matching in this graph has cost \mathcal{Z} —each sink node can have one edge incident to it from a source node, at the cost of one. Thus,

$$\text{cost}(M) = \mathcal{Z}.$$

Next, the cheapest nonblocking star has cost $\mathcal{A} + 2(\mathcal{Z} - 1)$; such a star is obtained by picking one of the sink nodes and connecting it to all others, with edges directed appropriately. Such a star is illustrated in Figure 5. It follows that the ratio between the cheapest star and the maximum-weight matching is

$$\begin{aligned} R &\geq \frac{\mathcal{A} + 2\mathcal{Z} - 2}{\mathcal{Z}} \\ &= 2 + \frac{\mathcal{A} - 2}{\mathcal{Z}}, \end{aligned}$$

which comes arbitrarily close to the approximation bound stated in Theorem 4.11.

To show that the bound in Lemma 4.8 is also tight, consider the example in Figure 6, which has n nodes, each with $\alpha(v) = \omega(v) = 1$, and $\gamma(u, v) = 1$, for all $u, v \in V$. In this

case, it is easily seen that the maximum-weight matching has cost n : every node can send and receive one unit of data at the cost of one. On the other hand, every nonblocking star has cost $2(n - 1)$: connect a root node to all others, with two unit-capacity edges directed oppositely. See Figure 6. It follows that the ratio between the cheapest star and maximum-weight matching is $2 - \frac{2}{n}$, which approaches the bound of Lemma 4.8.




Figure 6: Approximation bound is tight even for symmetric case. The links costs are 1 for all pairs. A maximum-weight matching has cost at most \mathcal{A} , while a minimum-cost nonblocking star, such as the one shown in the figure, has cost $2(\mathcal{A} - 1)$.

5 Unit Link Costs

We now consider the case when all link costs are the same, and show that a star network is optimal when the switch capacities are balanced. Despite being a specialized case, it applies to practical situations where the link costs are dominated by the cost of the terminating electronics, or where there is a single type of link from which larger link groups must be constructed.

Since all links have the same cost, without loss of generality, we assume that $\gamma(u, v) = 1$, for all u, v . In this case, the problem can be specified with three parameters: (V, α, ω) . We first prove the following lemma, which is useful in the proof of the main theorem.

Lemma 5.12 *Suppose V is a set of switches, with balanced source and sink capacities $\alpha(v)$ and $\omega(v)$, and unit link cost function between pairs of switches. Let \mathcal{N} be a nonblocking network for (V, α, ω) such that $\kappa_{\ell}(u, v) \geq \min\{\alpha(u), \omega(v)\}$, for all $(u, v) \in V \times V$. Then, the following holds:*

$$\text{cost}(\mathcal{N}) \geq \sum_v (\alpha(v) + \omega(v)) - \max_v (\alpha(v) + \omega(v)).$$

PROOF. We note that the bound on the right hand side is the cost of a nonblocking star, rooted at the node with a maximum capacity; unit link costs imply that the cost of a network equals its total link capacity. In counting the link capacities in \mathcal{N} , we charge

each link to its *destination* node. Let v_m denote the switch with the maximum capacity (source or sink) of all switches, and without loss of generality assume that

$$\omega(v_m) = \max_{v \in V} \{\alpha(v), \omega(v)\}.$$

Considering any other node v_i , where $i \neq m$, we get

$$\kappa_\ell(v_i, v_m) \geq \alpha(v_i), \tag{6}$$

since $\alpha(v_i) \leq \omega(v_m)$. All these links are charged to v_m , and they sum to $\sum_v \alpha(v) - \alpha(v_m)$.

Next, if the total incoming link capacity at each v_i , for $i \neq m$, is at least $\omega(v_i)$, then we get the desired bound on the overall cost of the network, completing the proof. So, assume that the incoming link capacity falls short at some node, say, v_i . Since we must have $\kappa_\ell(v_j, v_i) \geq \min\{\alpha(v_j), \omega(v_i)\}$, for all $v_j \neq v_i$, the total incoming link capacity at v_i fails to add up to $\omega(v_i)$ only if the following holds:

$$\omega(v_i) > \sum_v \alpha(v) - \alpha(v_i);$$

that is, the sink capacity of v_i exceeds the combined source capacity of all other nodes. When this happens, we conclude that

$$\alpha(v_i) + \omega(v_i) > \sum_{v \in V} \alpha(v). \tag{7}$$

We now re-apply the argument, using v_i in place of v_m as the purported root of the star. Since $\omega(v_i) > \sum_{v \in V} \alpha(v) - \alpha(v_i)$, it follows that each *incoming* link (v_k, v_i) at v_i has capacity $\kappa_\ell(v_k, v_i) = \alpha(v_k)$. We charge these links to v_i , and consider the incoming links at any other node v_k . Can it happen again that at some node v_k , for $k \neq i$, we find

$$\omega(v_k) > \sum_v \alpha(v) - \alpha(v_k)? \tag{8}$$

Suppose it did. Then, inequalities (7) and (8) together imply that

$$\begin{aligned} (\alpha(v_i) + \omega(v_i)) + (\alpha(v_k) + \omega(v_k)) &> 2 \sum_{v \in V} \alpha(v) \\ &= \sum_{v \in V} (\alpha(v) + \omega(v)), \end{aligned} \tag{9}$$

which is clearly not possible. Thus, the incoming links at each node v_k , for $i \neq k$, sum to $\omega(v_k)$, and thus the total link capacity of \mathcal{N} is at least

$$\sum_v (\alpha(v) + \omega(v)) - \max_v (\alpha(v) + \omega(v)),$$

and the proof is completed. □

We can now prove the result that, for unit link costs, a star network is optimal.

Theorem 5.13 *Let V be a set of switches, with source and sink capacities $\alpha(v)$ and $\omega(v)$, and assume unit link cost function between pairs of switches. Then, for balanced switch capacities, a minimum cost nonblocking star network is an optimal network.*

PROOF. We show that any nonblocking network must have a total link capacity at least

$$\sum_{v \in V} (\alpha(v) + \omega(v)) - \max_{v \in V} (\alpha(v) + \omega(v)). \quad (10)$$

It is easy to see that this matches the cost of a cheapest nonblocking star network, obtained by choosing as root the switch with the maximum source plus sink capacity. Let \mathcal{N}^* be an optimal nonblocking network, and let $\kappa_\ell(u, v)$ denote the capacity of the link (u, v) ; if there is no link between u and v , this capacity is zero. Consider any pair of nodes $(u, v) \in V \times V$ for which the following inequality holds:

$$\kappa_\ell(u, v) < \min\{\alpha(u), \omega(v)\}. \quad (11)$$

We set up two connections from u to v , first at the rate of $\kappa_\ell(u, v)$, and second at the rate $f(u, v) = \min\{\alpha(u), \omega(v)\} - \kappa_\ell(u, v)$. Due to the capacity constraint, the second connection must use an indirect path, requiring at least two links. We now tear-down the *first* connection, freeing up switch capacities $\kappa_\ell(u, v)$ at both u and v .

Since connection rerouting is not permitted in nonblocking networks, the second connection continues to be routed along the indirect path. This connection consumes $f(u, v)$ units of source (resp. sink) capacity of u (resp. v). It also consumes at least $2f(u, v)$ link capacities in \mathcal{N}^* , by virtue of being an indirect path. Subtract $f(u, v)$ from the switch capacities of u and v , and link capacities of all the links in the indirect path used by the connection. Observe that this modification keeps the switch capacities balanced. (In order not to introduce extra notation, we continue to use $\alpha(v)$, $\omega(v)$, and $\kappa_\ell(u, v)$ for the *residual* capacities of switches and links.)

We now repeat the connection setup procedure at any other link for which the condition in Ineq. (11) holds, until no such link exists. Suppose that the total source capacity consumed by the indirect connections is \mathcal{A}_1 ; an equal amount of sink capacity is also consumed. By the simultaneous connection argument used in Eq. (5), these (indirect) connections saturate at least

$$2\mathcal{A}_1 \quad (12)$$

units of link capacity in \mathcal{N}^* .

When the condition in (11) no longer holds, every node-pair $(u, v) \in V \times V$ satisfies:

$$\kappa_\ell(u, v) \geq \min\{\alpha(u), \omega(v)\},$$

and the total residual source capacity is $\mathcal{A} - \mathcal{A}_1$. We now invoke Lemma 5.12 on the residual network, which must be nonblocking for the residual (balanced) capacities. Combining the lower bound of Lemma 5.12 with the bound in (12), we conclude that

$$\begin{aligned} \text{cost}(\mathcal{N}^*) &\geq 2\mathcal{A}_1 + 2(\mathcal{A} - \mathcal{A}_1) - \max_v (\alpha(v) + \omega(v)) \\ &\geq \sum_v (\alpha(v) + \omega(v)) - \max_v (\alpha(v) + \omega(v)), \end{aligned}$$

which completes the proof. \square

6 Discussion and Future Research Directions

Our approximation algorithms have focused exclusively on star networks. These networks have a tremendous practical and theoretical appeal: they are extremely simple to build and maintain, and require very little overhead in setting up or tearing down connections. One potential disadvantage of the star networks is the huge transit capacity needed at the root switch. There are known switch architectures, however, whose cost grows as a function $c_1L + c_2L \log L$, where L is the total capacity of all links incident to the switch. The constants c_1 and c_2 are technology-dependent, but currently $c_1 \gg c_2$, and so the majority of the switch cost can be effectively combined with the link costs whenever L is too large—see Turner [16]. Thus, having a provable cost guarantee for star networks has a lot of appeal, and it remains a tantalizing problem to determine the best possible approximation bound for a star network.

The examples presented in Section 4.5 show that we cannot hope to improve the approximation bound using the maximum-weight matching as a lower bound. Interestingly enough, a star network is indeed *optimal* for the examples in Figures 5 and 6, and so the weakness is on the lower bound side.

The linear link cost model ignores the fact that a communication link of an arbitrary capacity must be constructed by combining links from a limited set of types, each with its own fixed capacity. In theory, the problem of even determining the cheapest combinations of links to achieve a particular bandwidth is equivalent to the well-known knapsack problem, and therefore intractable. In practice, however, the small set of available choices permit an efficient dynamic programming solution. We are currently working on extending our results to multiple (fixed number of) links types. A major source of difficulty is that, without the linearity of link costs, the combination equation for the lower bound Eq. (5) does not hold. Details on some practical heuristics and their performance can be found in [5, 6].

Finally, while global traffic constraints are clearly appealing to network designers, they can lead to overly expensive network designs. There is a trade-off here between minimizing the amount of information required of the network manager and providing enough information to yield the most cost-effective designs. We are currently exploring additional constraints that can improve the quality of network designs without placing an undue burden on the network manager. Some examples include *hierarchical clustering*, *distance-bounded clusters*, and *node-set pair constraints*.

References

- [1] V. E. Beneš. *Mathematical Theory of Connecting Networks and Telephone Traffic*. Academic Press, NY, 1965.
- [2] M. Bern and P. Plassman. “The Steiner Problem with Edge Lengths 1 and 2.” *Info. Proc. Letters*, 32, 1989, 171–176.
- [3] C. Clos. “A Study of Nonblocking Switching Networks.” *Bell Systems Technical Journal*, 32, 1953.

- [4] J. A. Fingerhut. "Approximation Algorithms for Configuring Nonblocking Communication Networks." Washington Univ. Dept of CS Doctoral Dissertation, 5/94.
- [5] J. A. Fingerhut, R. Jackson, S. Suri, J. S. Turner. "Design of Nonblocking ATM Networks." In preparation.
- [6] J. A. Fingerhut, R. Jackson, S. Suri, J. S. Turner. "New Models for Network Design." In preparation.
- [7] M. R. Garey and D. S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman, 1979.
- [8] A. Gersht and R. Weismayer. "Joint Optimization of Data Network Design and Facility Selection." *IEEE J. on Selected Areas in Communications*, 12/90.
- [9] A. Kershenbaum P. Kermani and G. Grover. "MENTOR: An Algorithm for Mesh Network Topological Optimization and Routing." *IEEE Transactions on Communications*, April 1991.
- [10] R. Melen and J. S. Turner. "Nonblocking Multirate Distribution Networks." *IEEE Transactions on Communications*, vol. 41(2), 1993, pp. 362–369.
- [11] M. Minoux. "Network Synthesis and Optimum Network Design Problems: Models, Solution Methods and Applications." *Networks*, 1989.
- [12] N. Pippenger. "Telephone Switching Networks." *Proc. Symp. on Applied Mathematics*, 1982.
- [13] M. De Prycker. "ATM Switching on Demand." *IEEE Net. Magazine*, 6(2), 1992.
- [14] F. Tobagi, T. Kwok and F. Chiussi. "Architectures, Performance, and Implementation of the Tanden Banyan Fast Packet Switch." *IEEE J. on Selected Areas of Communications*, 9(8), 1991.
- [15] J. S. Turner. "Design of a Broadcast Packet Network." *IEEE Transactions on Communications*, 6/88.
- [16] J. S. Turner. "An Optimal Nonblocking Multicast Virtual Circuit Switch." *INFOCOM*, 1994, pp. 298–305.
- [17] B. Waxman. "Routing of Multipoint Connections." *IEEE J. on Selected Areas of Communications*, 12/88.
- [18] Y. S. Yeh, M. G. Hluchyj and A. S. Acampora. "The Knockout Switch: A Simple Modular Architecture for High Performance Packet Switching." *IEEE J. on Selected Areas of Communications*, 5(8), 1987.