Directed Network Design with Orientation Constraints*

Sanjeev Khanna [†] Joseph (Seffi) Naor[‡] F. Bruce Shepherd[§]

Abstract

We study directed network design problems with *orientation constraints*. An orientation constraint on a pair of nodes u and v states that a feasible solution may include at most one of the arcs (u, v) and (v, u). Such constraints arise naturally in many network design problems, since link or edge resources such as fibre can be used to support traffic in one of two possible directions only. Our first result is that the directed network design problem with orientation constraints can be solved in polynomial time in the case where the requirement function f is positively intersecting supermodular. (The case where there are no orientation constraints follows from work of Frank [6].) The second main result of the paper is a 4-approximation algorithm for the minimum cost strongly connected subgraph problem with orientation constraints. Our algorithm shows that the problem of enforcing orientation constraints can be reduced to the minimum cost 2-edge connected subgraph problem on undirected graphs. Finally, we study the problem for general crossing supermodular functions and show the following bi-criteria approximation result. Let k denote the maximum requirement of any set under the given requirement function f. We give a $2k$-approximation algorithm to construct a solution that satisfies a slightly weaker requirement function, namely, $f'(S) = \max\{f(S) - 1, 0\}$.

*A preliminary version of this paper appeared in [16].

†Department of Computer and Information Science, University of Pennsylvania, Philadelphia, PA 19104. Most of this work was done when the author was at Bell Laboratories, Lucent Technologies, 700 Mountain Avenue, Murray Hill, NJ 07974. Email : `sanjeev@cis.upenn.edu`.

‡Computer Science Department, Technion, Haifa 32000, Israel. Most of this work was done when the author was at Bell Laboratories, Lucent Technologies, 700 Mountain Avenue, Murray Hill, NJ 07974. E-mail: `naor@cs.technion.ac.il`.

§Bell Laboratories, Lucent Technologies, 700 Mountain Avenue, Murray Hill, NJ 07974. E-mail: `bshep@research.bell-labs.com`.

1 Introduction

We study directed network design problems in mixed networks with *orientation constraints.* We are given a directed graph $D = (V, A)$ and an undirected graph $G = (V, E)$. Let $M = (V, E \cup A)$ denote the *mixed graph* obtained by taking their union. An *orientation* of an undirected edge is obtained by replacing it by one of two possible directed arcs parallel to it. We consider problems which amalgamate two previously studied models for network design. We wish, at minimum cost, to fulfill a given connectivity requirement in a network by (a) selecting a subgraph, and (b) orienting its undirected edges. We start by first considering a concrete special case of such orientation and subgraph selection problems.

Let $r \in V$ be a special node, and consider the following two problems, the first concerning D, the second concerning M. Call a digraph *k-connected from r* if it contains k arc-disjoint paths from r to each other node.

(A) Given a cost function on A, find a minimum cost subdigraph (if there is one) which is k-connected from r.

(B) Find an orientation (if there is one) of the undirected edges of M which results in a minimum cost digraph which is k-connected from r.

The polyhedron for problem (A) is described in [6] while that for (B) is derived from the integrality of submodular flow polyhedra (see, e.g., [9, 23]). One aim of the present paper is to describe the polyhedron for the following common generalization.

(AB) Find a mixed subgraph M' of M and orient the undirected edges of M' so that the resulting digraph is k-connected from r, and so that the cost of the directed edges plus the cost of the orienting the undirected edges of M' is minimized.

We study the problem (AB) in the context of the following more general framework.

INPUT: We are given a directed graph $D = (V, A)$ with a cost function $c : A \to \mathcal{Z}$, a requirement function f defined over all subsets of V, and a collection $\mathcal{E}$ of disjoint *constrained* arc pairs, each of which induces a *digon* (i.e., two arcs directed in opposite directions). Let $A_p \subseteq A$ denote the set of constrained arcs, and for any arc $a \in A_p$, let a^- denote the arc that a is paired with. We are also given lower and upper bound vectors l, u for each pair in $\mathcal{E}$ and each arc a.

GOAL: Find an optimal solution to the integer program below; here $\delta^+(S)$ denotes the set of arcs leaving $S \subseteq V$.

$$(I) \qquad \min \quad \sum_{a \in A} c_a x_a$$

$$\sum_{a \in \delta^+(S)} x_a \geq f(S) \qquad \text{for each } S \subset V \qquad (1)$$

$$l_{a,a^-} \leq x_a + x_{a^-} \leq u_{a,a^-} \quad \text{for each } \{a, a^-\} \in \mathcal{E} \qquad (2)$$

$$x_a \in \{l_a, l_a + 1, \ldots, u_a\} \qquad \text{for each } a \in A \qquad (3)$$

We refer to constraints (1) as the *cut* constraints, constraints (2) as the *orientation* constraints, and constraints (3) as the *integrality* constraints. The term orientation here "arises" from the consideration of constrained pairs with $l_{a,a^-} = u_{a,a^-} = 1$. In this case, the choice of the variables x_a, x_{a^-} amounts to determining the orientation of an associated undirected edge.

The above framework, without the orientation constraints, already captures a large number of fundamental combinatorial optimization problems. Some representative examples include minimum cost branchings, minimum cost k-strongly connected subgraphs, and the directed steiner network problem. Many of these problems are NP-hard and thus research is often focused on the design of approximation algorithms for these problems.

In the present paper, we restrict attention to *crossing supermodular* requirement functions f. That is, for every $X, Y \subseteq V$ such that $X \cap Y \neq \emptyset$ and $X \cup Y \neq V$ we have that

$$f(X) + f(Y) \leq f(X \cap Y) + f(X \cup Y).$$

Directed network design problems with a crossing supermodular requirement function remain NP-hard. An example of a crossing supermodular function is the function $f(S) = k$ for all subsets $S \subset V$; the associated problem is known as the minimum cost k-strongly connected subgraph problem. For this case, a simple 2-approximation algorithm is obtained by solving two minimum cost k-disjoint arborescence problems (one into and one out from) at an arbitrary node v [5].

Frank [6] showed that in the special case where the requirement function is also *intersecting supermodular*, i.e., the above inequality holds whenever X and Y intersect, the network design problem can be solved optimally in polynomial time. Melkonian and Tardos [21] have recently shown that his result can be used to obtain a 2-approximation algorithm for any requirement function which is crossing supermodular.

In undirected graphs, so-called weakly supermodular functions model a broad class of network design problems, including for instance, the generalized Steiner tree problem. Following a line of work ([1, 12, 13, 24]), Jain [15] devised an ingenious 2-approximation algorithm for weakly supermodular functions. He proved that every basic feasible solution to the linear programming (LP) relaxation of the problem contains a variable of value at least a half. Jain's algorithm finds and rounds such a *large component* iteratively until a final integral solution is obtained.

Network design problems in undirected graphs are often much better understood than their directed counterparts. In particular, techniques for network design problems on undirected graphs, e.g., the widely used primal-dual approach [1, 12], do not seem to be easily amenable to directed network design problems. In the work of Melkonian and Tardos [21] a directed analog of Jain's result is given which requires significant further insight into the combinatorial structure of basic solutions. They show that every basic solution to an LP relaxation of the design problem contains a variable of value at least a quarter, whenever the requirement function f is crossing supermodular. We note that the LP relaxation for (I) is polynomially solvable by the ellipsoid method [14].

1.1 Our Problems

We study directed network design problems with orientation constraints (as specified in (I)). We can view these as two-phase problems: finding a subgraph of an undirected graph, or a mixed graph, and then orienting its edges so as to satisfy the cut constraints. The cost function associated with the orientation may in general be *asymmetric*, i.e., the cost of orienting an edge $e = uv$ from v to u is different from orienting it from u to v. (An edge can only be oriented in one direction.) The cost of an orientation is defined to be the sum of the costs of the orientations of the edges.

Thus, our network design problems combine constraints of two types, subgraph constraints, and orientation constraints. While each type of constraint has been well-studied separately, much less is known for design problems that combine these two types of constraints simultaneously.

Perhaps, the most basic orientation problem with asymmetric costs that involves both subgraph constraints and orientation constraints is finding among all subgraphs of G that admit a strong orientation, one that has a strong orientation of minimum cost. This problem generalizes two well known NP-hard problems. If the orientation cost function is symmetric, then the problem reduces to finding a minimum cost 2-edge connected subgraph of G. On the other hand, if there are no orientation constraints, then the problem reduces to finding a minimum cost strongly connected subgraph of a directed graph. We note that for both problems, 2-approximation algorithms are known.

One case we study for which a complete solution is given is that where the requirement function is positively intersecting supermodular. This is a generalization of the problem (AB) defined at the outset of the paper. These results appear in Section 3.

Another important case of special interest in our study is the design of strongly connected directed graphs with orientation constraints. In this case the requirement function f satisfies $f(X) = 1$ for every proper nonempty subset X, and $f(V) = f(\emptyset) = 0$.

An interesting special case of the asymmetric orientation problem is when the constraints (2) in (I) become $x_a + x_{a^-} = 1$ for each arc $a \in A$. Equivalently, we are given an undirected graph and we need to find a minimum cost orientation that satisfies the requirement function f. A good characterization for the special case of strong connectivity requirements, follows from the classical min-max theorem of Lucchesi and Younger [20]. A *dijoin* is a set of arcs in a digraph which intersects every directed cut $\delta^+(S)$, i.e., S such that $\delta^-(S) = \emptyset$. For an integer cost vector c, the Lucchesi-Younger Theorem states that the minimum c-cost of a dijoin is equal to the maximum packing of directed cuts where each arc a is allowed in at most c_a cuts in this packing. This theorem also implies the existence of a polynomial time algorithm to find the minimum dijoin via the Ellipsoid Method - see [14]. We can find a minimum cost strongly connected orientation as follows. First, orient each constrained pair in the cheaper of its two directions to obtain a digraph D'. We must now reverse some of these orientations in order to make D strongly connected. If A' is the subset of arcs of D' which are flipped, then evidently A' must be a dijoin. So we could do no better than to choose a cheapest such dijoin; but does reversing the arcs of a dijoin result in a strongly connected digraph? In general no, but Lovász [19, Exercise 6.11] and independently Younger

[25] proved that any minimal dijoin does indeed have this property. The orientation problem for general crossing supermodular requirement functions can also be solved in polynomial time via reductions to submodular flows, see [9, 23].

Orientation constraints arise in many network design problems, since link/edge resources such as fibre, are commonly unidirectional (i.e., they support traffic in only one of the two possible directions at a given time). Asymmetric costs may arise in many network routing problems. For instance, consider the setting where traffic demand is being incrementally introduced in an existing network. Load balancing constraints may favor forcing traffic in opposite directions between a given pair of switches. Hence when routing new demands, costs on the directed links may increase proportionately to the amount of existing traffic. Asymmetric costs may also arise in network planning due to assorted line termination equipment; these are the costs associated with terminating the two ends of a link.

1.2 Our Results

Our first result is that the directed network design problem with orientation constraints can be solved in polynomial time using the Ellipsoid Method, in the case where the requirement function f is positively intersecting supermodular. We show that any basic solution for the relaxation of (I) in this case has integral components. This generalizes the work of Frank [6] who proved the same result for the variant with no orientation constraints. In fact, we establish this result for the more general formulation given in (I).

Our second result is that the minimum cost strongly connected subgraph problem with orientation constraints has a 4-approximation algorithm. We give a combinatorial approximation algorithm based on the idea that the problem of enforcing orientation constraints can be reduced to the minimum cost 2-edge connected subgraph problem. We start with any feasible solution to the minimum cost strongly connected subgraph problem and use the above reduction to modify the solution so as to satisfy the violated orientation constraints.

Finally, we study our problem for general crossing supermodular functions and show the following bi-criteria approximation result. Let k denote the maximum requirement of any set under the given requirement function f. We give a $2k$-approximation algorithm to construct a solution that satisfies a slightly weaker requirement function, namely, $f'(S) = \max\{f(S) - 1, 0\}$.

2 Preliminaries

We denote a directed graph by $D = (V, A)$. For $S \subseteq V$, denote by $\delta^+(S)$ (respectively $\delta^-(S)$) the set of arcs with tail in S (respectively $V - S$) and head in $V - S$ (respectively S).

A pair of subsets X, Y, of a ground set V, is *intersecting* if $X \cap Y, X - Y, Y - X \neq \emptyset$. An intersecting pair X, Y of sets is *crossing* if $X \cup Y \neq V$ and X, Y are noncomparable. A family $\mathcal{F}$ of nonempty subsets of V is *intersecting* if we have $X \cap Y, X \cup Y \in \mathcal{F}$ for each intersecting pair $X, Y \in \mathcal{F}$. A family is *crossing* if $X \cap Y, X \cup Y \in \mathcal{F}$ for each crossing pair

$X, Y \in \mathcal{F}$. A function $f : 2^V \to \mathcal{Z}_+$ is *positively crossing* (respectively *positively intersecting supermodular*) on a crossing (respectively intersecting) family $\mathcal{F}$, if:

1. $f(V) = f(\emptyset) = 0$.
2. For each crossing (respectively intersecting) pair $X, Y \in \mathcal{F}$ such that $f(X), f(Y) > 0$, $f(X) + f(Y) \le f(X \cap Y) + f(X \cup Y)$.

We emphasize that we only require the inequality to hold for X, Y in the support of f.[1]

For any ordered pair (u, v) of nodes of V, we define an operator Ψ_{uv} as follows. Given any $f : 2^V \to \mathcal{Q}_+$, $\Psi_{uv}(f)$ is a new function such that $\Psi_{uv}(f)(S) = f(S) - 1$ if $u \in S, v \notin S$ and $f(S) > 0$. Otherwise $\Psi_{uv}(f)(S) = f(S)$. The following is proved in [6].

Lemma 1 *If f is positively crossing (respectively positively intersecting) supermodular on $\mathcal{F}$, then $\Psi_{uv}(f)$ is also positively crossing (respectively positively intersecting) supermodular.*

Note that this result does not hold for crossing (intersecting) supermodular functions (i.e., without the positive requirement). When f is clear from the context, we denote by $\mathcal{F}(uv)$ the family obtained from $\mathcal{F}$ by removing all sets $S \neq V, \emptyset$ for which $\Psi_{uv}(f)(S) = 0$. One easily checks that if f is positively crossing (intersecting) supermodular on $\mathcal{F}$, then it is also positively crossing (intersecting) supermodular on $\mathcal{F}(uv)$.

A family $\mathcal{S} = \{S_i\}_{i=1}^m$ of proper, nonempty subsets of a finite ground set V is *cross-free* if no pair of sets in $\mathcal{S}$ cross. The family is *laminar* if for each pair S_i, S_j of distinct sets in $\mathcal{S}$, we have either $S_i \subseteq S_j, S_j \subseteq S_i$, or $S_i \cap S_j = \emptyset$.

3 Intersecting Supermodularity

In this section we study the polyhedron obtained by relaxaing the integrality constraints (3) in (I), i.e., for each $a \in A$, we now require only $l_a \le x_a \le u_a$. In the following, $\mathcal{F}$ denotes an intersecting family of subsets of V, and f is a positively intersecting supermodular function on $\mathcal{F}$.

Let D be a digraph and $\mathcal{E}$ be a disjoint collection of arc pairs, a and a^-, each of which forms a *digon*, i.e., directed circuit of length two. A quadruple $(D, f, \mathcal{F}, \mathcal{E})$ will be called simply an *f-network* (or positively intersecting supermodular network according to f). A *capacitated* f-network is a sextuple $(D, f, \mathcal{F}, \mathcal{E}, l, u)$ where $(D, f, \mathcal{F}, \mathcal{E})$ is an f-network and $l, u : A \cup \mathcal{E} \to \mathcal{Z}_+$ are assignments of capacities to the arcs and digons of $\mathcal{E}$.

For any such capacitated f-network and cost vector c, the *f-connectivity problem* is to find an optimal solution to (I). Denote by $P(D, f, \mathcal{F}, \mathcal{E}, l, u)$ the polyhedron defined by the relaxation of (I), that is replacing (3) by $l_a \le x_a \le u_a$ for each arc a. We assume that the polyhedron in consideration is nonempty; in particular, note also that it is pointed

[1]This property has also been referred to as **weakly** crossing (intersecting) supermodularity in [6], whereas supermodularity had been referred to functions which satisfy condition (2) for **all** crossing (intersecting) pairs X, Y in $\mathcal{F}$.

and hence has vertices. Each extreme point is thus defined by a system of m linearly independent tight inequalities. We are interested primarily in integer solutions to such capacitated f-connectivity problem and so our goal is to show:

Theorem 2 *For any capacitated positively intersecting f-network $(D, f, \mathcal{F}, \mathcal{E}, l, u)$, the extreme points of $P(D, f, \mathcal{F}, \mathcal{E}, l, u)$ are integral.*

The special case where there are no orientation constraints follows from Frank [6]. We note that it is easy to construct examples such that there is an unbounded gap between the cost of optimal solutions with and without orientation constraints.

This theorem first appeared in [16]; we follow the same proof which is based on a primal analysis of the extreme points of $P(D, f, \mathcal{F}, \mathcal{E}, l, u)$. Subsequently and independently, it was communicated to us (by J. Cheriyan as well as A. Frank) that the result can be proved by showing that the system of inequalities associated with $P(D, f, \mathcal{F}, \mathcal{E}, l, u)$ is *totally dual integral* (TDI). That is, for every integer vector c, the dual linear program has an integer optimum. In the meantime, another strengthening of the original result has appeared in [11]. In particular, they extend the result to a hypergraph setting and to orientation constraints over larger sets of possible orientations. They also show that Theorem 2 for (not positively) intersecting supermodularity, follows from the theory of submodular flows using a reduction of Schrijver [23]. They indicate, however, that they do not know whether the polyhedron $P(D, f, \mathcal{F}, \mathcal{E}, l, u)$ for positively intersecting supermodular functions arises from a submodular flow polyhedron.

We now proceed with a proof of Theorem 2. Henceforth we let $P(D, f, \mathcal{F}, \mathcal{E}, l, u)$, be a positively intersecting supermodular f-connectivity polyhedron and x^* be an extreme point. Note that any extreme point has a *defining system* determined by a triple $\mathcal{S}, \mathcal{R}^+, \mathcal{R}^-$ where $\mathcal{S} \subseteq \mathcal{F}$, $\mathcal{R} = \mathcal{R}^+ \cup \mathcal{R}^- \subseteq A \cup \mathcal{E}$ and $m = |A| = |\mathcal{S}| + |\mathcal{R}|$. That is, x^* is the unique solution (in R^A) to the system of equalities:

1. $x_a = u_a$, for each $a \in \mathcal{R}^+$
2. $x_a = l_a$, for each $a \in \mathcal{R}^-$
3. $x_a + x_{a^-} = u_{a,a^-}$, for each $\{a, a^-\} \in \mathcal{R}^+$
4. $x_a + x_{a^-} = l_{a,a^-}$, for each $\{a, a^-\} \in \mathcal{R}^-$
5. $x(\delta^+(S)) = f(S)$, for each $S \in \mathcal{S}$

and in particular, $\mathcal{S}, \mathcal{R}$ identify a set of linearly independent rows in the constraint matrix for the f-connectivity problem.

We now analyze the structure of such an extreme point x^* and this (not necessarily unique) defining system.

Lemma 3 *If $(D, f, \mathcal{F}, \mathcal{E}, l, u)$ is a positively intersecting supermodular problem, then any extreme point x^* has a defining system such that $\mathcal{S}$ is laminar.*

Proof: We first produce a system which is cross-free. The proof of this part is by the well-known uncrossing technique; its essence is identical to Lemma 2 in [21], and so we only sketch the proof. If $X, Y \in \mathcal{S}$, then we may replace these inequalities by those for $X \cap Y, X \cup Y$ (which are again tight) without destroying linear independence of the system.

So we assume that $\mathcal{S}$ is cross-free and suppose that $X, Y \in \mathcal{S}$ such that $X \cap Y, X - Y, Y - X$ are all nonempty but $X \cup Y = V$. Let $x_{ab} = x^*([X-Y, Y-X])$, $x_{ba} = x^*([Y-X, X-Y])$, $x_{ia} = x^*([X \cap Y, X-Y])$, $x_{ib} = x^*([X \cap Y, Y-X])$. For example, $x_{ia} = \sum_{a \in \delta^+(X \cap Y) - \delta^+(X)} x_a$; in particular, $x^*(X) = x_{ab} + x_{ib}, x^*(Y) = x_{ba} + x_{ia}$. Thus we have

$$
\begin{aligned}
&(x_{ab} + x_{ib}) + (x_{ba} + x_{ia}) \\
&\quad = \quad f(X) + f(Y) \\
&\quad \leq \quad f(X \cap Y) + f(X \cup Y) \\
&\quad = \quad f(X \cap Y) \\
&\quad = \quad (x_{ib} + x_{ia})
\end{aligned}
$$

from which we deduce that $x_{ab} = x_{ba} = 0$ and that $X \cap Y$ is again tight for x^*. In this case, we may replace the set Y in $\mathcal{S}$ by $X \cap Y$. Let a, b, c be the $0, 1$ incidence vectors of $\delta^+(X), \delta^+(Y), \delta^+(X \cap Y)$ respectively. Note that we have $a + c = 2a + b$ and hence the resulting system is again defining (i.e., induce a nonsingular matrix) for the vector x^*.

Applying this procedure decreases the value $\sum_{S \in \mathcal{S}} |S|$ and does not create any new intersecting pairs. Thus we may repeatedly apply this operation until we obtain the desired laminar system. □

The following lemma guarantees an integral component in extreme points of a positively intersecting supermodular connectivity polyhedron.

Lemma 4 *Let x^* be an extreme point which has a defining system $\mathcal{S}, \mathcal{R}$ for which $\mathcal{S}$ is laminar. Then there is some arc a such that x^*_a is an integer.*

Proof: Suppose that $x^*, \mathcal{S}, \mathcal{R}$ is a counterexample for a digraph D with a minimum number of arcs and subject to this minimizes $\sum_a x^*_a$. Clearly, each member of $\mathcal{R}$ must define a constrained pair, or we would have an integral component immediately. We first claim that without loss of generality

$$\text{Each arc appears in some cut } \delta^+(S) \text{ where } S \in \mathcal{S}. \tag{4}$$

Suppose this is not the case for some a, then clearly $a \in \mathcal{R}$ and the constraint $x_a + x_{a^-} \leq 1$ is the only constraint from the defining system which involves a. For any value z, let x^z be the vector obtained by replacing the value x^*_a by the value z. Let z_0 be the smallest value such that x^{z_0} is feasible. Note that if $z_0 > 0$, then a is in some tight cut $\delta^+(S)$ for x^{z_0}, and so we may obtain a linearly independent defining system for x^z by adding this cut to $\mathcal{S}$ and removing the orientation constraint for $\{a, a^-\}$. If $x^z = x^*$, then we have a desired system for x^*. Otherwise, $x^z < x^*$ is another extreme point which contradicts the minimality of x^*. If $z_0 = 0$, then let x' be obtained from x^* by restricting to the components of $A - a$; also set $\mathcal{E}' = \mathcal{E} - \{a, a^-\}$. Then x' is an extreme point for the f-connectivity polyhedron of $D - a$ and $\mathcal{E}'$, contradicting minimality.

In the following, let R denote those arcs which lie in come constrained pair in $\mathcal{R}$; also let $F = A - R$. For a set S, we let $F(S) = F \cap \delta^+(S)$ and $R(S) = R \cap \delta^+(S)$. For any laminar family $\mathcal{S}$ we may define an acyclic *laminar directed graph* $H(\mathcal{S})$ with a node v_S associated with each $S \in \mathcal{S}$ and where there is an arc (v_{S_i}, v_{S_j}) if S_i is a minimal set containing S_j. Let $\Phi : A \to V(H(\mathcal{S}))$ be a mapping. The function Φ is a *legal mapping* for a subgraph $H' \subseteq H$ if each arc of H' is mapped to a node v_S such that S contains its tail, i.e., the tail of each arc in $\Phi^-(v_S)$ lies in S. The *contribution* to v_S by Φ, denoted by $C(v)$, is

$$\sum_{a \in F, \Phi(a) = v_S} 1 \; + \sum_{a \in R, \Phi(a) = v_S} \frac{1}{2} \, .$$

We next prove the following claim.

Claim 5 *For each subtree T rooted at a node v_S in $H(\mathcal{S})$ we may find a legal mapping such that $C(v_S) \geq 1 + |F(S)|/2$ and for each other node $v \in T$ we have $C(v) \geq 1$.*

Proof of Claim: Note that since $x^* > 0$ (there are no integral components), we must have that $|\delta^+(S)| \geq 2$ for each $S \in \mathcal{S}$. Hence the result holds in the case where $|V(H)| = 1$. Now suppose that $v_{S_1}, v_{S_2}, \ldots, v_{S_q}$ are the out-neighbours, or children, of v_S. By induction, for each i, there is a legal mapping Φ_i for the subgraph induced by the descendants of v_{S_i} such that the contribution to v_{S_i} is at least $1 + |F(S_i)|/2$. For each i, we first partition the arcs in $\delta^+(S_i)$ as follows. Let $F_i^+ = \delta^+(S) \cap F(S_i)$ and $R_i^+ = \delta^+(S) \cap R(S_i)$. Let $F_i^- = F(S_i) - F_i^+$ and $R_i^- = R(S_i) - R_i^+$. Let $F' = F(S) - (\cup_i F_i^+)$ and $R' = R(S) - (\cup_i R_i^+)$. The cardinality of each of these sets we denote by changing the large capital letter to small, e.g., $f_i^+ = |F_i^+|$.

We create a legal map Φ from the individual maps Φ_i as follows. First, every arc from F', R' is mapped to v_S. Next, for each i, select either one arc of $F(S_i)$, or two arcs of $R(S_i)$ from $\Phi^-(v_{S_i})$ and then reassign any remaining arcs of $\Phi^-(v_{S_i})$ to v_S. Thus v_{S_i} still has a a contribution of 1 but v_S inherents its excess contribution of $|F(S_i)|/2$. Note that if the resulting map does not contribute $1 + |F(S)|/2$ to v_S, then we must have

$$f' + \frac{r'}{2} + \frac{1}{2}\sum_{i=1}^{q}(f_i^+ + f_i^-) < 1 + \frac{f'}{2} + \frac{1}{2}\sum_{i=1}^{q} f_i^+ .$$

This implies that $f' + r' + \sum_{i=1}^{q} f_i^- \leq 1$. Consider adding the constraint for S together with all orientation constraints $\{a, a^-\}$ for which both ends of a are contained in S. To this combination, subtract the equalities associated with the cuts S_i. We are left with the following equation:

$$f' + r' - \sum_{i=1}^{q} f_i^- \sum_{a \in F' \cup R'} x_a^* - \sum_{i=1}^{q} \sum_{a \in F_i^-} x_a^* = \gamma,$$

where γ is some integer. (Note that we require (4) in order to deduce this fact.) Thus if $f' = r' = \sum_{i=1}^{q} f_i^- = 0$, then we have a nontrivial combination of tight constraints which yields the zero vector, a contradiction. Thus exactly one of $f', r', f_1^-, f_2^-, \ldots, f_q^-$ is 1. But, in this case, the equality above shows that x^* has an integral component, a contradiction. □

We now continue with the proof of the lemma. Suppose that $S_1, S_2 \ldots S_p$ are the maximal sets in $\mathcal{S}$ and apply the claim to the subtree rooted at each v_{S_i}; let Φ_i be the legal mapping thus found. Note first that the subtrees are disjoint for if this were not the case, then there would exist some node, say v_S of in-degree two or more. But in that case, any pair of parents S', S'' would be incomparable (by definition of an arc in H) and hence $S' - S'', S'' - S'$ and $S'' \cap S'$ would all be nonempty, contradicting laminarity.

Thus we can combine the mappings Φ_i into an overall mapping Φ such that the contribution to each node is at least 1 and $C(v_{S_i}) \geq 1 + |F(S_i)|/2$ for each $i = 1, 2, \ldots, p$. We thus have

$$|\mathcal{S}| \leq \sum_{v_S \in H} C(v_S) \leq \sum_{a \in A \cap F} 1 + \sum_{a \in A \cap R} \frac{1}{2} = |A| - |\mathcal{R}|.$$

But we know that $|\mathcal{S}| = |A| - |\mathcal{R}|$ and so all inequalities must be tight above. In particular, we must have $C(v_{S_i}) = 1$ for each $i = 1, 2, \ldots, p$. Hence $|F(S_i)| = 0$ for each i. But then we have a nontrivial combination of the zero vector by adding the cut inequalities for the maximal sets S_i and subtracting the orientation constraints corresponding to arcs appearing in some $\delta^+(S_i)$. This final contradiction completes the proof of Lemma 4. □

Proof of Theorem 2 Suppose the statement is false and x^* is a non-integral extreme point for some $P(D, f, \mathcal{F}, \mathcal{E}, l, u)$. Moreover, choose such a counterexample for which $|A|$ is minimized and subject to this, minimizes $\sum_a x_a^*$. Suppose x_a^* is an integer, say k, for some arc a. Then we claim to find a smaller counterexample by reducing a's component. If $k > 0$, then let x' be obtained from x^* by decreasing by 1 the component for a; we may also reduce by 1 any u_e, l_e corresponding to an arc or constrained pair involving a. Let l', u' be the vectors obtained. The vector x' is an extreme point for $P(D, \Psi(f), \mathcal{F}(wv), \mathcal{E}, l', u')$, where $a = (w, v)$. To see this, note any inequality which was tight for x^* is still tight for x' in the new system. Thus suppose that $k = 0$. Let x' be obtained by restricting x to the variables $A - \{a\}$. Also let $\mathcal{E}'$ be obtained from $\mathcal{E}$ by removing any digon of the form $\{a, a^-\}$ if it exists. Also, let l', u' be obtained by restricting to arcs and constrained pairs which do not involve a. The non-integral vector x' is then an extreme point of $P(D - a, f, \mathcal{F}, \mathcal{E}', l', u')$, contradicting minimality of x^*. Thus we may assume that each component of x^* is fractional. However, this now contradicts Lemmas 3 and 4. □

4 Strong Connectivity

We present in this section a combinatorial 4-approximation algorithm for the problem of strong connectivity with orientation constraints. This is an important special case of (I) that is not captured by our study in Section 3 — the requirement function for this problem is crossing supermodular. For clarity of presentation, we assume that any parallel arcs form a digon, and any such pair is involved in an orientation constraint, that is, the pair appears in $\mathcal{E}$. However, our algorithm extends trivially to the general case. In what follows, we assume that the input directed graph is $D = (V, A)$ and the optimal solution is a directed graph $D^* = (V, A^*)$. We use OPT to denote the cost of the optimal directed graph D^*. We say that an arc (u, v) in a directed graph $D = (V, A)$ is *simple* if $(v, u) \notin A$, and we say that it is *non-simple* otherwise. A directed cycle is called *non-trivial* if it is a simple cycle of length at least three.

At the center of our approach is a procedure that takes as input any strongly connected subgraph of D, possibly violating some orientation constraints, and reduces the problem of "amending" its violated orientation constraints to that of finding a minimum cost 2-edge connected subgraph in an undirected graph. In fact, the precise problem we reduce it to is a minimum cost augmentation of a spanning tree to a 2-edge connected subgraph. For the latter problem, a 2-approximation algorithm [5, 17] is known. We now describe our algorithm in detail:

1. Pick any node r and compute a minimum cost in-branching to r, say T_1, as well as a minimum cost out-branching from r, say T_2. Consider the directed graph $D_1 = (V, A_1)$ induced by $T_1 \cup T_2$. Clearly, D_1 is strongly connected and its cost is at most $2 \cdot \text{OPT}$. Assume without loss of generality that D_1 is minimal.

2. The set of simple arcs $A' \subseteq A_1$ induces a collection of strongly connected components $C_1, C_2, ..., C_k$ (see Lemma 6). Shrink each component C_i to a single node x_i and construct a directed graph $D_2 = (X, A_2)$, where $X = \{x_1, ..., x_k\}$. An arc $(x_i, x_j) \in A_2$, if and only if D_1 contained some arc (u, v) with $u \in X_i, v \in X_j$. The minimality of D_1 implies that D_2 is minimally strongly connected as well.

3. Replacing each non-simple pair of arcs by an undirected edge, evidently results in a tree, by minimality. Thus, D_2 has $k-1$ non-simple arc pairs $(a_1, b_1), ..., (a_{k-1}, b_{k-1})$. It is convenient to view D_2 as an undirected tree $T = (X, E_T)$, such that T contains an edge e_i for each pair (a_i, b_i). Let $(X_i, \bar{X}_i)$ be the partition of node set X induced by removal of a pair (a_i, b_i). We associate with any such partition a cut $(S_i, \bar{S}_i)$ in D_1, where $S_i = \bigcup_{x_j \in X_i} V(C_j)$. We refer to these cuts as the *fundamental cuts* of D_1. We say that an arc a *hits* a fundamental cut if $a \in \delta^+(S_i) \cup \delta^+(\bar{S}_i)$. For each fundamental cut, $A^* \setminus A_1$ contains an arc $a \in \delta^+(S_i) \cup \delta^+(\bar{S}_i)$ (see Lemma 7). Thus, the minimum cost Z of a set of arcs in $A \setminus A_2$, that hits all fundamental cuts of D_1 (call it the *fundamental directed cut cover* of D_1) is no more than OPT. Finding an optimal fundamental directed cut cover is NP-hard, but a 2-approximation can be easily obtained through the undirected version of the problem (see Lemma 8). Let A_3 be a set of arcs obtained in this fashion. Since Z is at most OPT, the cost of A_3 is at most $2 \cdot \text{OPT}$. Let $D_3 = (V, A_1 \cup A_3)$ be the directed graph obtained by adding arcs in A_3 to the directed graph D_1. The total cost of arcs in D_3 is at most $4 \cdot \text{OPT}$.

4. The final step is to show that the directed graph D_3 above can be modified into another directed graph $D_4 = (V, A_4)$ such that (i) D_4 is strongly connected, (ii) $A_4 \subseteq (A_2 \cup A_3)$, and (iii) all arcs in D_4 are simple. To achieve this we use the

necessary and sufficient conditions for the existence of a strong orientation of a mixed graph given by Boesch and Tindell [2] (see Lemma 9). The costs of steps (1) and (3) are no more than $2 \cdot \mathrm{OPT}$ each. Therefore we now have a 4-approximation to our problem.

Lemma 6 *In any minimally strongly connected directed graph H, the set of simple arcs induces a collection of strongly connected components.*

Proof: It suffices to show that every simple arc lies on a directed cycle that consists only of simple arcs. Consider any simple arc (u, v) and a path $P(v, u)$ from v to u in H. By the minimality of H, every arc on $P(v, u)$ must be simple. The lemma follows. □

Lemma 7 *$A^* \setminus A_1$ hits every fundamental cut of D_1.*

Proof: Suppose not. Then there is a cut $(S_i, \bar{S}_i)$ such that A^* has at most one arc (from the pair $\{a_i, b_i\}$) that crosses the cut. This contradicts that D^* is strongly connected. □

Lemma 8 *There is a 2-approximation algorithm for the minimum cost fundamental directed cut cover problem.*

Proof: We solve this problem by a reduction to the minimum cost 2-edge connected subgraph problem. Let B be the set of arcs $A \setminus A_2$. Consider the undirected graph $H = (X, E)$ obtained as follows. There is an edge $x_i x_j \in E$ if and only if there is an arc in B that connects some node in X_i to X_j or vice versa. Moreover, the cost of this edge is equal to the minimum cost such arc. Also, for each edge in E_T, we include an edge of cost 0 in H.

We now claim that the problem of finding a minimum cost fundamental directed cut cover of D_1 is equivalent to finding a minimum cost 2-edge connected subgraph $H' = (X, E')$ of H. To see this, consider any fundamental directed cut cover $\hat{A}$. Then, the set of undirected edges in H, obtained from the arcs in $\hat{A}$ along with the edges in E_0, has the property that every cut in H has at least two edges crossing it. Moreover, the cost of this collection of edges is no more than the cost of the fundamental directed cut cover $\hat{A}$. In the other direction, let us consider a 2-edge connected subgraph $\hat{E}$ of H. Since exactly one edge in E_0 crosses a cut in H, there must be at least one additional edge in $\hat{E}$ crossing any cut. Thus, the directed arcs corresponding to the edges in $\hat{E}$ form a fundamental cut cover of D_1.

Since the zero cost edges induce a spanning tree, the minimum cost 2-edge connected subgraph problem we need to solve is essentially a minimum cost augmentation of a spanning tree to a 2-edge connected subgraph. For this problem, a 2-approximation algorithm is known [5, 17]. It then suffices to use this algorithm to get a 2-approximation algorithm for the minimum cost fundamental directed cut cover problem. □

Lemma 9 *The directed graph D_3 obtained at the end of Step (3) of the algorithm, can be modified into a directed graph D_4 such that D_4 is strongly connected, contained in D_3, and does not contain any non-simple arcs.*

Proof: The lemma essentially follows from the necessary and sufficient conditions of Boesch and Tindell [2] for the existence of a strong orientation of a mixed graph. The BT conditions state that a mixed graph whose underlying graph is 2-edge connected has a strong orientation if and only if it does not contain a directed cut, i.e., a cut where all edges are directed in the same direction. Note that the underlying graph of D_3 is 2-edge connected, and since A_2 spans D_3, every cut contains an undirected edge and therefore there are no directed cuts in D_3.

For completeness, we sketch here a short proof that does not use the BT conditions directly. The only non-simple arcs (or undirected edges) in D_3 are the ones corresponding to the pairs (a_i, b_i). We remove from D_3, one by one, exactly one arc from each such pair while keeping it strongly connected. Consider a leaf x_i in T. There must be an arc in A_3 of the form (x_i, x_j) or (x_j, x_i) that hits the fundamental cut $(x_i, X \setminus \{x_i\})$. Assume without loss of generality that it is of the form (x_i, x_j). Consider the directed path from x_j to x_i in D_2, consisting only of non-simple arcs, and remove every non-simple arc that is oriented in a direction opposite to this path. Contract the resulting cycle and repeat this procedure on a leaf of the resulting tree T'. It is easy to verify that the procedure continues until the resulting tree reduces to a single node, corresponding to the graph D_4 above. □

5 Crossing Supermodularity

We now focus our attention on general crossing supermodular functions. While simple constant factor approximation algorithms are known in the absence of orientation constraints [21], the problem seems to become much harder in the presence of orientation constraints. Although we do not resolve this question here, we make some progress towards solving our original problem (I) for crossing supermodular functions. We establish the following result.

Theorem 10 *Let* OPT *denote the optimal cost of a fractional solution to the problem* (I) *with a crossing supermodular function* f *defined on a crossing family* $\mathcal{F}$. *Let* $p = \max_S f(S)$ *denote the maximum requirement of any set. Then, there is a polynomial time algorithms that finds an integral solution of cost* $2p \cdot \mathrm{OPT}$ *satisfying the weaker requirement function* f' *defined as* $f'(S) = \max\{f(S) - 1, 0\}$.

An immediate corollary of the above theorem is that we can find a solution to a $(k-1)$-strong connectivity problem with orientation constraints at a cost that is no more than $2k$ times the optimal cost for k-strong connectivity. We devote the rest of this section to the proof of Theorem 10.

A key step in our algorithm finds a minimum cost orientation of a graph for general crossing supermodular requirement functions. That is, it solves problem (I) in the case where the orientation constraints are equalities. As mentioned earlier, this problem can be solved in polynomial time via reductions to submodular flows, see [9, 23]. (In [10] an orientation theorem is given for the more general class of so-called crossing G-supermodular functions.) We briefly explain how this is done, following [9].

First, choose the cheaper arc from each pair of arcs, a and a^-, appearing in an orientation constraint $x_a + x_{a^-} = 1$, yielding a directed graph denoted by $D_1 = (V, A_1)$. Clearly, the cost of A_1 is a lower bound on the cost of an optimal solution. Then, find a minimum cost collection of arcs $(u, v) \in A_1$, such that if they are *flipped*, i.e., (u, v) is replaced by (v, u), then a directed graph satisfying the requirement function f is obtained. Call such a set of arcs an *f-flip set*.

Mininum cost f-flip sets can be optimally computed by a reduction to the minimum cost submodular flow problem, which is itself solvable in polynomial time. It is easily seen that a set A_2 is a f-flip set if and only if its incidence vector x satisfies for every proper subset S:

$$|\delta^+(S)| + x(\delta^-(S)) - x(\delta^+(S)) \geq f(S)$$

Define a new function g where

$$g(S) = -f(S) + |\delta^+(S)|.$$

Since f is crossing supermodular, we get that g is submodular. Therefore, finding a minimum cost f-flip set is equivalent to solving the submodular flow problem $\min\{cx : x \geq 0,\ x(\delta^+(S)) - x(\delta^-(S)) \leq g(S)$ for each proper subset $S\}$.

We now go back to the proof of Theorem 10. Let D be our digraph and f be a crossing supermodular function on the family $\mathcal{F}$. Let $p = \max_S(f(S) - 1)$. Define $f'(S) = \lfloor \frac{f(S)}{1+1/p} \rfloor$ for each set S. It is easy to verify that $f'(S) = f(S) - 1$ if $f(S) > 0$, and that f' is also crossing supermodular on $\mathcal{F}$. Let $\mathrm{OPT}(f)$, or OPT if the context is clear, denote the cost of an optimal solution x^* to problem (I) for a requirement function f.

1. Define u_a to be $\lceil px^*(a) \rceil$. Clearly, $u_a + u_{a^-} \leq p + 1$. Also, define $f_p(S) = p \cdot f(S)$ to be a new crossing supermodular function.

2. We now solve two separate intersecting supermodular LPs, corresponding to f_p, with upper bounds just defined. These LP's are obtained by splitting the "requiring" sets into $\mathcal{F}^1 = \{F \in \mathcal{F} : v \in F\}$ and $\mathcal{F}^2 = \mathcal{F} - \mathcal{F}^1$ where v is an arbitrarily chosen node. The first we may solve directly; the latter is not actually intersecting and so we work with the intersecting family $\{V - S : S \in \mathcal{F}^2\}$. For the second family we must also work with the function f^* defined by $f^*(S) = f(V - S)$ and use the digraph with the arcs reversed (see [21]). By Theorem 2, any basic solution for these LP's is integral. We may thus find two such vectors z^1, z^2 in polynomial time. Actually, since these LP's do not have orientation constraints, we can also find integral optimal solutions using Frank's approach ([6], see also [23]). Now define a new vector z by setting $z_a = \max\{z_a^1, z_a^2\}$ for each $a \in A$. Clearly z is integral, satisfies f_p, and costs no more than $2 \cdot \mathrm{OPT}(f_p) \leq 2p \cdot \mathrm{OPT}(f)$.

3. By setting $y = \frac{1}{p}z$ we obtain a solution for the original function f which is $(1/p)$-integral and has cost at most 2OPT. The solution y violates any orientation constraint by at most a factor of $1 + 1/p$. We scale down all violating arc pairs to satisfy the constraint $x_a + x_{a^-} = 1$. We also uniformly scale up any arc pairs with $1/p \leq x_a + x_{a^-} < 1$ to satisfy $x_a + x_{a^-} = 1$. The resulting solution clearly satisfies the function f' and has a cost that is at most $2p \cdot \mathrm{OPT}$.

4. At this point, since all orientation constraints are tight, we can find an integral solution of no greater cost.

We can view the above algorithm as a process in which we tighten the orientation constraints, while bounding the increase in cost. However, we can only guarantee that the cut constraints are "almost" satisfied. We conjecture that the following approach, similar to that of Jain [15], and Melkonian and Tardos [21], can yield a constant factor approximation algorithm for crossing supermodular requirement functions. Let x^* be an optimal (basic) solution to the linear relaxation of (I) in the case where f is crossing supermodular. We conjecture that there exists a pair of arcs $a, a^- \in A$, for which the orientation constraint is not tight, yet $x^*(a) + x^*(a^-)$ is greater than a constant, say 1/4. If this conjecture is correct, then we can make the orientation constraint on a and a^- tight, and resolve the problem. We repeat this until we obtain a solution in which all orientation constraints are tight. The cost of this solution increases by only a constant factor with respect to x^*. Once all orientation constraints are tight, as before, we can find an integral solution at no greater cost.

Acknowledgements: The authors are grateful to the detailed comments of one referee which greatly improved the paper. We also thank Joseph Cheriyan and Andras Frank for insightful remarks on the contents of the paper.

References

[1] A. Agrawal, P. Klein, and R. Ravi, An approximation algorithm for generalized Steiner tree problems in networks, *Proceedings of the 23rd ACM Symposium on Theory of Computing*, (1991) 134-144.

[2] F. Boesch and R. Tindell, Robbins's theorem for mixed multigraphs, *American Mathematical Monthly*, **87** (1980), pp. 716–719.

[3] J. Edmonds, R. Giles, A min-max relation for submodular functions on graphs, *Annals Discrete Math.,* **1** (1977), pp. 185–204.

[4] K.P. Eswaran, R.E. Tarjan, Augmentation problems, *SIAM Journal on Computing,* **5** (1976), pp. 653–665.

[5] G. N. Frederickson and J. JáJá, Approximation algorithms for several graph augmentation problems, *SIAM Journal on Computing*, **10**(2) (1981), pp. 270–283.

[6] A. Frank, Kernel systems of directed graphs, *Acta Sci. Math (Szeged),* **41** (1979), pp. 63–76.

[7] A. Frank, How to make a digraph strongly connected, *Combinatorica,* **1**(2) (1981), pp. 145–153.

[8] A. Frank, Augmenting graphs to meet edge-connectivity requirements, *SIAM Journal on Discrete Math.,* **5**(1) (1992), pp. 25–53.

[9] A. Frank, Applications of submodular functions, *Surveys in Combinatorics, London Mathematical society Lecture Notes Series 187*, Cambridge University Press (Ed. K. Walker) (1993), pp. 85–136.

[10] A. Frank, Orientations of graphs and submodular flows, *Congressus Numerantium (Ed: A. J. W. Hilton)*, **113** (1996), pp. 111–142.

[11] A. Frank, T. Király, Z. Király, On the orientation of graphs and hypergraphs, Technical Report TR-2001-06 of the Egerváry Research Group, Budapest, Hungary, (2001).

[12] M. Goemans and D. Williamson, A general approximation technique for constrained forest problems, *SIAM Journal on Computing*, **24** (1995), pp. 296–317.

[13] M. Goemans, A. Goldberg, S. Plotkin, D. Shmoys, É. Tardos, and D. Williamson, Approximation algorithms for network design problems, *Proceedings of the 5th Annual ACM-SIAM Symposium on Discrete Algorithms*, (1994), pp. 223–232.

[14] M. Grötschel, L. Lovász, and A. Schrijver, The ellipsoid method and its consequences in Combinatorial Optimization, *Combinatorica*, **1** (1981), 169-197.

[15] K. Jain, A factor 2 approximation algorithm for the generalized steiner network problem, *Combinatorica*, **21** (2001), 39-60.

[16] S. Khanna, J. Naor, and F. B. Shepherd, Directed network design with orientation constraints, *Proceedings of the 11th Annual ACM-SIAM Symposium on Discrete Algorithms*, (2000), pp. 663–671.

[17] S. Khuller, R. Thurimella, Approximation algorithms for graph augmentation *Journal of Algorithms*, **14** (1993), 214–225.

[18] L. Lovász, On two min-max theorems in graph theory, *JCT Series B,* **21** (1976) 96-103.

[19] L. Lovász, Combinatorial problems and exercises, *North Holland Press*, Amsterdam, 1979.

[20] C.L. Lucchesi, D.H. Younger, A minimax relation for directed graphs, *J. London Math. Soc.,* **17**(2) (1978), pp. 369–374.

[21] V. Melkonian and É. Tardos, Approximation algorithms for a directed network design problem, *IPCO '99.*

[22] A. Schrijver, Packing and covering crossing families of cuts, *JCT Series B*, **35** (1983), pp. 104–128.

[23] A. Schrijver, Total dual integrality from directed graphs, crossing families and sub- and supermodular functions, in: W.R. Pulleyblank, ed., *Progress in Combinatorial Optimization*, Academic Press, Toronto, 1984, 315-362.

[24] D. Williamson, M. Goemans, M. Mihail, and V. Vazirani, A Primal-dual approximation algorithm for generalized steiner network problems, *Combinatorica*, **15** (1995), pp. 435–454.

[25] D. Younger, Private Communication, (1985).