Tight Guarantees for Cut-Relative Survivable Network Design via a Decomposition Technique

Nikhil Kumar
University of Waterloo, Canada

J. J. Nan
University of Waterloo, Canada

Chaitanya Swamy
University of Waterloo, Canada

Abstract

In the classical *survivable-network-design problem* (SNDP), we are given an undirected graph $G = (V, E)$, non-negative edge costs, and some k tuples (s_i, t_i, r_i), where $s_i, t_i \in V$ and $r_i \in \mathbb{Z}_+$. The objective is to find a minimum-cost subset $H \subseteq E$ such that each s_i-t_i pair remains connected even after the failure of any $r_i - 1$ edges. It is well-known that SNDP can be equivalently modeled using a weakly-supermodular *cut-requirement function* f, where the objective is to find the minimum-cost subset of edges that picks at least $f(S)$ edges across every cut $S \subseteq V$.

Recently, motivated by fault-tolerance in graph spanners, Dinitz, Koranteng, and Kortsartz proposed a variant of SNDP that enforces a *relative* level of fault tolerance with respect to G. Even if a feasible SNDP-solution may not exist due to G lacking the required fault-tolerance, the goal is to find a solution H that is at least as fault-tolerant as G itself. They formalize the latter condition in terms of paths and fault-sets, which gives rise to *path-relative* SNDP (which they call relative SNDP). Along these lines, we introduce a new model of relative network design, called *cut-relative* SNDP (CR-SNDP), where the goal is to select a minimum-cost subset of edges that satisfies the given (weakly-supermodular) cut-requirement function to the maximum extent possible, i.e., by picking $\min\{f(S), |\delta_G(S)|\}$ edges across every cut $S \subseteq V$.

Unlike SNDP, the cut-relative and path-relative versions of SNDP are not equivalent. The resulting cut-requirement function for CR-SNDP (as also path-relative SNDP) is not weakly supermodular, and extreme-point solutions to the natural LP-relaxation need not correspond to a laminar family of tight cut constraints. Consequently, standard techniques cannot be used directly to design approximation algorithms for this problem. We develop a *novel decomposition technique* to circumvent this difficulty and use it to give a *tight* 2*-approximation algorithm for* CR-SNDP. We also show some new hardness results for these relative-SNDP problems.

2012 ACM Subject Classification Theory of computation → Approximation algorithms analysis; Mathematics of computing → Discrete optimization

Keywords and phrases Approximation algorithms, Network Design, Cut-requirement functions, Weak Supermodularity, Iterative rounding, LP rounding algorithms

Digital Object Identifier 10.4230/LIPIcs.ESA.2025.38

Related Version *Full Version*: https://arxiv.org/abs/2507.04473

Funding Supported in part by NSERC grant 327620-09.

1 Introduction

In the classical *survivable-network-design problem* (SNDP), we are given an undirected graph $G = (V, E)$, non-negative edge costs $\{c_e\}_{e \in E}$, and some k source-sink pairs $s_i, t_i \in V$ with requirements $r_i \in \mathbb{Z}_+$, for $i = 1, \ldots, k$. The goal is to find a minimum-cost set $H \subseteq E$ of edges such that there are at least r_i edge-disjoint s_i-t_i paths in H, for all $i \in [k]$. We overload notation and use H to denote both the edge-set, and the corresponding subgraph of G.

SNDP imposes an absolute level of fault-tolerance in a solution subgraph H, by ensuring that for any bounded-size *fault-set* $F \subseteq E$ of edges that may fail, the graph $H - F$ retains some connectivity properties. More precisely, due to Menger's theorem, or the max-flow min-cut theorem, the feasibility condition on H can be stated equivalently as follows. For every $i \in [k]$, and every fault-set $F \subseteq E$ with $|F| < r_i$:

(S1) (Path-version) there is an s_i-t_i path in $H - F$;

(S2) (Cut-version) for every s_i-t_i cut $S \subseteq V$ (i.e., $|S \cap \{s_i, t_i\}| = 1$), we have $\delta_{H-F}(S) \neq \emptyset$, where $\delta_{H-F}(S) := \delta(S) \cap (H - F)$.

Recently, motivated by work on fault-tolerance in graph spanners, Dinitz, Koranteng, and Kortsartz [8], proposed a variant of SNDP that aims to impose a *relative* level of fault-tolerance with respect to the graph G. The idea is that even if the SNDP instance in infeasible, because G itself does not possess the required level of fault-tolerance, one should not have to completely abandon the goal of fault-tolerance: it is still meaningful and reasonable to seek a solution that is "as fault-tolerant as G", and in this sense is fault-tolerant relative to G.

To formalize this, it is useful to consider the definition of SNDP in terms of fault-sets. Roughly speaking, we would like to capture that H is a feasible solution if

for every fault-set F (valid for SNDP), $H - F$ and $G - F$ have similar connectivity. (*)

To elaborate, in SNDP, if there is *even one* fault-set F' under which $G - F'$ fails to have the required connectivity, i.e., $|F'| < r_i$ and $F' \supseteq \delta_G(S)$ for some s_i-t_i cut S, then "all bets are off"; we declare that the instance in infeasible. But declaring infeasibility here feels unsatisfactory because we allow one specific "problematic" fault-set F' to block us from obtaining any kind of fault-tolerance. In contrast, (*) aims to provide a *per-fault-set guarantee*, by asking for a solution H that functions as well as G in terms of connectivity under the failure of any fault-set F (considered for SNDP). We can formalize "$H - F$ and $G - F$ have similar connectivity" in two ways, via paths or cuts, and this gives rise to the following two problem definitions.

► **Definition 1.1.** *Let* $\big(G = (V, E), \{c_e\}_{e \in E}, \{s_i, t_i, r_i\}_{i \in [k]}\big)$ *be an* SNDP *instance.*

(R1) Path-relative SNDP*:* $H \subseteq E$ *is a feasible solution if for every* $i \in [k]$ *and* $F \subseteq E$ *with* $|F| < r_i$, $G - F$ *has an* s_i-t_i *path* $\implies$ $H - F$ *has an* s_i-t_i *path.*

(R2) Cut-relative SNDP *(*CR-SNDP*):* $H \subseteq E$ *is a feasible solution if for every* $i \in [k]$ *and* $F \subseteq E$ *with* $|F| < r_i$*, for every* s_i-t_i *cut* S, $\delta_{G-F}(S) \neq \emptyset \implies \delta_{H-F}(S) \neq \emptyset$.

The goal in both problems is to find a minimum-cost feasible solution.

Perhaps surprisingly, path-relative SNDP and cut-relative SNDP are *not equivalent*, even when the SNDP instance involves a single s-t pair. It is not hard to show that if H is feasible for CR-SNDP, it is also feasible for path-relative SNDP[1] but the converse fails to hold; see Section 7 for an example.

Path-relative SNDP was defined by Dinitz et al. [8], who referred to this problem simply as relative SNDP. They considered (among other problems) the path-relative version of k-edge connected subgraph (k-ECSS), which is the special case of SNDP where we have an s_i-t_i pair with $r_i = k$ for every pair of nodes. (In this case, the path-relative and cut-relative versions do turn out to be equivalent [8].) They called the resulting path-relative problem, *k-edge*

[1] Fix any $i \in [k]$ and $F \subseteq E$ with $|F| < r_i$. If $G - F$ has an s_i-t_i path, then $\delta_{G-F}(S) \neq \emptyset$ for every s_i-t_i cut. So since H is feasible for cut-relative SNDP, we have $\delta_{H-F}(S) \neq \emptyset$ for every s_i-t_i cut, which implies that $H - F$ has an s_i-t_i path.

fault-tolerant subgraph (k-EFTS), and observed that the feasibility condition for k-EFTS can be equivalently stated as: H is feasible iff $|\delta_H(S)| \geq g(S) := \min\{k, |\delta_G(S)|\}$ for all $S \subseteq V$. The function $g : 2^V \mapsto \mathbb{Z}$ is called a *cut-requirement function*, as it stipulates the (minimum) number of edges across any cut in any feasible solution. Cut-requirement functions constitute a very versatile framework for specifying network-design problems, where given a cut-requirement function $f : 2^V \mapsto \mathbb{Z}$, the corresponding *$f$-network-design problem* (f-NDP), also called the *f-connectivity problem*, is to find a minimum-cost set $H \subseteq E$ such that $|\delta_H(S)| \geq f(S)$ for all $S \subseteq V$. For instance, it is easy to see that SNDP corresponds to f-NDP for the cut-requirement function f^{SNDP}, defined by $f^{\mathsf{SNDP}}(S) := \max\{r_i : |S \cap \{s_i, t_i\}| = 1\}$ (and k-ECSS is $f^{k\text{-}\mathsf{ECSS}}$-NDP where $f^{k\text{-}\mathsf{ECSS}}(S) := k$ for all $\emptyset \neq S \subsetneq V$); various other network-design problems, such as the T-join problem, point-to-point connection problem etc., can be modeled using suitable cut-requirement functions (see, e.g., [10, 11]).

Cut-relative SNDP is a problem that we introduce in this paper. Its formulation in terms of cuts is the problem one obtains when we replace $f^{k\text{-}\mathsf{ECSS}}$ in the cut-based formulation of k-EFTS by the cut-requirement function f^{SNDP} for (general) SNDP: that is, (similar to k-EFTS), we can say that $H \subseteq E$ is feasible for CR-SNDP iff $|\delta_H(S)| \geq g^{\mathsf{CR\text{-}SNDP}}(S) := \min\{f^{\mathsf{SNDP}}(S), |\delta_G(S)|\}$ for all $S \subseteq V$ (we show this easy equivalence in Section 2); in other words, CR-SNDP corresponds to $g^{\mathsf{CR\text{-}SNDP}}$-NDP. To gain some intuition for cut-relative SNDP, and contrast it with path-relative SNDP, consider SNDP with a single s-t pair and requirement r, for simplicity. A feasible solution H to path-relative SNDP offers a per-fault-set guarantee as encapsulated by (R1). But, for a given fault-set F with $|F| < r$, we get a weak fault-tolerance guarantee in terms of s-t cuts: if $\delta_{G-F}(S) = \emptyset$ for *even one* s-t cut S then there is no requirement on H for this fault-set F. Cut-relative SNDP offers a per-fault-set *and per-cut guarantee*, since for (every F with $|F| < r$ and) every s-t cut S, we require that $\delta_{H-F}(S) \neq \emptyset$ if $\delta_{G-F}(S) \neq \emptyset$.

Along the lines of CR-SNDP, we can easily extend the cut-relative model to capture, more broadly, the cut-relative version of any network-design problem specified by a cut-requirement function: given a cut-requirement function f, in the corresponding *cut-relative f-network-design problem* (CR-f-NDP), we seek a min-cost $H \subseteq E$ such that $|\delta_H(S)| \geq g^{\mathsf{CR}\text{-}f\text{-}\mathsf{NDP}}(S) := \min\{f(S), |\delta_G(S)|\}$ for all $S \subseteq V$; that is, we seek to satisfy the cut-requirement function f to the maximum extent possible. We refer to f as the *base cut-requirement function*, to distinguish it from the cut-requirement function $g^{\mathsf{CR}\text{-}f\text{-}\mathsf{NDP}}$ that defines CR-f-NDP.

Modeling power. We arrived at CR-f-NDP, which can be seen as a fault-tolerant cut-covering problem, as a technically natural extension of k-EFTS. We show below that CR-f-NDP offers a surprising amount of modeling power, and, in particular, allows one to capture stronger forms of relative fault-tolerance compared to k-EFTS. In k-EFTS, there is a sharp relative-fault-tolerance threshold at k: if H is feasible, then $G - F$ and $H - F$ have the same components whenever $|F| < k$, but there are no guarantees for larger fault-sets. With CR-f-NDP, one can capture a weaker guarantee also for other fault-sets, which allows for a more *graceful degradation* of relative fault-tolerance as $|F|$ increases.

▶ **Example 1.2.** One way of modeling graceful relative-fault-tolerance degradation is as follows: given a *non-increasing function* $\tau : \mathbb{Z}_+ \mapsto \mathbb{Z}_+$, we seek a min-cost subgraph H such that for *every* fault-set F, $G - F$ and $H - F$ have exactly the same components with at most $\tau(|F|)$ nodes. If G models a communication network where connected nodes can communicate with each other, then this yields the desirable guarantee that, *post-faults*, if a node can communicate with more than $\tau(|F|)$ nodes in G then the same holds for H; if this number is at most $\tau(|F|)$, then post-faults, it can communicate with the same nodes in G and H. Note that if $\tau(k-1) \geq n = |V|$, then H is feasible for k-EFTS; by adding more "levels" to τ, we can obtain fault-tolerance guarantees for larger fault-sets.

In Section 3, we show that this problem can be modeled as CR-f-NDP with the weakly-supermodular function $f = f^{\mathsf{grace}}$, where $f^{\mathsf{grace}}(S) := \min\{\ell : \tau(\ell) < |S|\}$, for $S \neq \emptyset$.

▶ **Example 1.3.** Generalizing Example 1.2 substantially, suppose along with τ, we have a *monotone* function $\pi : 2^V \mapsto \mathbb{R}_+$, i.e., $\pi(T) \leq \pi(S)$ if $T \subseteq S$. We now seek a subgraph H such that for every fault-set F and $S \subseteq V$ with $\pi(S) \leq \tau(|F|)$, we have that S is (the node-set of) a component of $G - F$ iff it is a component of $H - F$. Similar to Example 1.2, we can model this as CR-f-NDP by defining $f(S) := \min\{\ell : \tau(\ell) < \pi(S)\}$, for $S \neq \emptyset$ (Theorem 3.2).

This setup yields fault-tolerance degradation depending on monotone properties of components other than just their size, which creates a rich space of problems. For example, suppose $\pi(S) = \max_{u,v \in S} \min_{u\text{-}v \text{ paths } P \text{ in } G} |P|$ is the weak-diameter of S. Setting $\tau(\ell) = n$ if $\ell < k$, and some $t < n$ for larger ℓ, we seek a k-EFTS solution H such that for $F \subseteq E$ with $|F| \geq k$, $G - F$ and $H - F$ have the same components of weak-diameter at most t.

1.1 Our contributions

We introduce cut-relative network-design problems and develop strong approximation guarantees for these problems. *We obtain an approximation guarantee of* 2 *for cut-relative* SNDP, which, notably, matches the best-known approximation factor for SNDP. Our guarantee applies more generally to CR-f-NDP, whenever the base cut-requirement function f satisfies certain properties and the natural LP-relaxation of the problem ($\text{CRLP}_{f,G}$) can be solved efficiently. Our guarantee is relative to the optimal value of this LP, *and is tight in that it matches the integrality gap of this LP.*

We also show that even in the simplest SNDP setting with only one s-t pair (wherein SNDP is polytime solvable), cut-relative SNDP and path-relative SNDP capture k-ECSS as a special case, and are thus *APX*-hard (Section 6). Previously, even *NP*-hardness of path-relative SNDP in the s-t case (which is studied by [8, 9]) was not known.

Technical contributions and overview. Technically, our main contribution is to show that we can obtain such a strong guarantee *despite the fact that the cut-requirement function defining cut-relative network-design (i.e.,* $g^{\mathsf{CR}\text{-}f\text{-}\mathsf{NDP}}$*) is not weakly supermodular*, which is the key property that drives the 2-approximation algorithm for SNDP. To elaborate, there is a celebrated 2-approximation algorithm for SNDP due to Jain [15] based on iterative rounding. This *crucially* exploits the fact that the underlying cut-requirement function f^{SNDP} is *weakly supermodular*: for any two node-sets $A, B \subseteq V$, we have:

$$f^{\mathsf{SNDP}}(A)+f^{\mathsf{SNDP}}(B) \leq \max\{f^{\mathsf{SNDP}}(A\cap B)+f^{\mathsf{SNDP}}(A\cup B),\ f^{\mathsf{SNDP}}(A-B)+f^{\mathsf{SNDP}}(B-A)\}.$$

This property allows one to argue that given an LP-solution x, any two *tight* sets that cross – i.e., sets A, B for which $x(S) = f^{\mathsf{SNDP}}(S)$ holds for $S \in \{A, B\}$ and $A \cap B, A - B, B - A$ are all non-empty – can be "uncrossed", and thereby an extreme-point solution $\widehat{x}$ to the LP can be defined via a *laminar family* of tight cut constraints. By a laminar family, we mean that any two sets A, B in the family satisfy $A \cap B = \emptyset$ or $A \subseteq B$, and uncrossing two crossing tight sets means that they can be replaced by an equivalent laminar family of tight sets. Jain's seminal contribution was to show that given such a laminar family defining an extreme point, there always exists an edge e for which $\widehat{x}_e \geq \frac{1}{2}$; picking such an edge and iterating then yields the 2-approximation. This technique of uncrossing tight cut-constraints to obtain a laminar family of tight cut-constraints has proved to be quite versatile, becoming a staple technique that has been leveraged to obtain strong guarantees in various network-design settings such as, most notably, network-design problems with degree constraints [13, 17, 19, 5, 4, 16]. More generally, this technique of uncrossing to obtain a structured family of tight constraints has been utilized for various combinatorial-optimization problems; see, e.g., [18].

As mentioned earlier, the main complication with CR-f-NDP is that the underlying cut-requirement function $g^{\text{CR-}f\text{-NDP}}$ is usually *not* weakly-supermodular, even if the base cut-requirement function f is. In particular, the cut-requirement function $g^{\text{CR-SNDP}}$ defining cut-relative SNDP is not weakly-supermodular. This difficulty was also noted by Dinitz et al. [8] when they considered k-edge fault-tolerant subgraph (k-EFTS, or path- or cut-relative k-ECSS in our terminology). For this special case, they identify a notion called local weak-supermodularity that is satisfied by the cut-requirement function $\min\{k, |\delta_G(S)|\}$ for k-EFTS, and observed that Jain's machinery can be used as is if this property holds; that is, one can uncross tight sets to obtain a laminar family, and hence obtain a 2-approximation for k-EFTS proceeding as in Jain's algorithm for SNDP.

However, for general CR-SNDP, such an approach breaks down, because, as we show in Section 2.1, one can construct quite simple CR-SNDP instances where an (optimal) extreme point *cannot be defined by any laminar family of tight cut constraints*. In particular, this implies that the cut-requirement function $g^{\text{CR-}f\text{-NDP}}$ is (in general) not locally weakly-supermodular (even when f is weakly-supermodular); moreover, it does not satisfy any property that allows for the uncrossing of tight sets. The upshot here is that the standard technique of uncrossing tight sets to obtain a laminar family does not work, and one needs new ideas to deal with cut-relative network design.

The chief technical novelty of our work is to show how to overcome the impediment posed by the lack of any such nice structural property for $g^{\text{CR-}f\text{-NDP}}$, via a novel *decomposition technique* that allows one to *reduce* cut-relative f-NDP with a weakly supermodular *base cut-requirement function* to (standard) f-NDP with (different) weakly supermodular cut-requirement functions (see Theorem 4.1).

We introduce a little notation to describe this. Let CR-(h, D)-NDP denote CR-f-NDP on graph D with base cut-requirement function h. For $S \subseteq V$, define its *deficiency* to be $f(S) - |\delta_G(S)|$; we say that S is a *small cut* if its deficiency is positive (i.e., $|\delta_G(S)| < f(S)$). We use $G[S]$ to denote the subgraph induced by S. Let $\overline{S}$ denote $V - S$.

Assume that the underlying base cut-requirement function f is weakly supermodular (as is the case with SNDP). Observe first that if there is no small cut then CR-f-NDP is the same as f-NDP. Our key technical insight is that, otherwise, we can utilize a novel splitting operation to simplify our problem. We show (see Theorem 4.5) that if we pick a small cut $A \subseteq V$ with *maximum deficiency*, then one can define suitable weakly-supermodular functions $f_A : 2^A \mapsto \mathbb{Z}$ and $f_{\overline{A}} : 2^{\overline{A}} \mapsto \mathbb{Z}$ such that *any solution to the original* CR-f-NDP*-instance yields solutions to* CR-$(f_A, G[A])$-NDP *and* CR-$(f_{\overline{A}}, G[\overline{A}])$-NDP*, and vice-versa.* The intuition here is as follows. Since A is a small cut, any feasible CR-f-NDP solution must include all edges in $\delta_G(A)$. The functions f_A and $f_{\overline{A}}$ are essentially the restrictions of f to A and $\overline{A}$ respectively, taking this into account. Consequently, moving to CR-$(f_A, G[A])$-NDP and CR-$(f_{\overline{A}}, G[\overline{A}])$-NDP does not really impact the constraints for cuts that do not cross A; but we do eliminate the constraints for crossing cuts. This is the simplification that we obtain via the splitting operation, and one needs to argue that this does not hurt feasibility; that is, the constraints for such cuts are still satisfied when we take $\delta_G(A)$ along with the CR-$(f_A, G[A])$-NDP and CR-$(f_{\overline{A}}, G[\overline{A}])$-NDP solutions. Here we exploit the weak-supermodularity of f, which implies that the deficiency function $\mathrm{def}(S) := f(S) - |\delta_G(S)|$ is also weakly supermodular. One can use this, and the fact that A has maximum deficiency, to bound the deficiency of any crossing set T in terms of the deficiency of non-crossing sets, and hence show that the cut-constraint for T is satisfied.

By recursively applying the above splitting operation to A and $\overline{A}$, we end up with a partition $V_1, \dots, V_k$ of V and weakly-supermodular functions $f_i : 2^{V_i} \mapsto \mathbb{Z}$ for all $i \in [k]$ such that, for each $i \in [k]$, the graph $G[V_i]$ contains no small cut with respect to f_i. It follows that solving f_i-NDP on the graph $G[V_i]$ for all $i \in [k]$ yields a solution to the original CR-f-NDP instance (Theorem 4.1).

While we can obtain this decomposition efficiently, assuming that one has suitable algorithmic primitives involving the base cut-requirement function f (shown in the full version), in fact we only need this decomposition in the *analysis*: in order to obtain the stated LP-relative 2-approximation algorithm, we only need to be able to find an extreme-point optimal solution to the natural LP-relaxation ($\mathrm{CRLP}_{f,G}$) for CR-f-NDP (also after fixing some variables to 1). This is because the translation between CR-f-NDP-solutions for G, and f_i-NDP-solutions on $G[V_i]$ for all $i \in [k]$, also applies to *fractional solutions*, and implies that any extreme-point solution to ($\mathrm{CRLP}_{f,G}$) yields extreme-point solutions to the LP-relaxations of f_i-NDP on $G[V_i]$ for all $i \in [k]$ (parts (b) and (c) of Theorem 4.1); hence, by Jain's result, we can always find some fractional edge of value at least $\frac{1}{2}$, round its value to 1, and iterate. For SNDP, we can find an extreme-point optimal solution efficiently because we argue that one can devise an efficient separation oracle for ($\mathrm{CRLP}_{f,G}$) (Lemma 2.3). This yields a 2-approximation algorithm for CR-SNDP, and more generally CR-f-NDP, whenever f is weakly supermodular and ($\mathrm{CRLP}_{f,G}$) can be solved efficiently (see Algorithm CRNDP-Alg in Section 5).

Other related work. Network design is a fairly broad research topic, with a vast amount of literature. We limit ourselves to the work that is most closely connected to our work.

As mentioned earlier, path-relative SNDP was introduced by Dinitz et al. [8]. In addition to k-EFTS, which is the path-relative version of k-ECSS, they study the path-relative versions of two other special cases of SNDP, for which they developed constant-factor approximation algorithms: (1) SNDP with $r_i \le 2$ for all i; (2) one s_i-t_i pair with $r_i \le 3$. Both results were subsequently extended by [9] who gave a 2-approximation for SNDP with $r_i \le 3$ for all i, and a $2^{O(k^2)}$-approximation when there is one s_i-t_i pair with $r_i = k$. The only result known for general path-relative SNDP is due to [7], who (among other results) devise an $O(k^2 \log^2 n \log\log n)$-approximation algorithm with $n^{O(k)}$ running time. As noted earlier, on the hardness side, prior to our work, it was not known if the setting with one s_i-t_i pair is even *NP*-hard.

An influential paper by Goemans and Williamson [10] (see also the survey [11]), which built upon the work of [1] for the Steiner forest problem, developed the primal-dual method for network-design problems. Their work also popularized the use of cut-requirement functions to specify network-design problems, by showing how this framework can be used to capture a slew of problems, and how their primal-dual framework leads to a 2-approximation algorithm for $\{0, 1\}$-cut-requirement functions satisfying some properties. This work was extended to handle general integer-valued cut-requirement functions, which captures SNDP, in [20, 12], which led to an $O(\log \max_i r_i)$-approximation for SNDP. The seminal work of Jain [15] later improved this to a 2-approximation, and in doing so, introduced the powerful technique of *iterative rounding*. Later work of [17, 19] on degree-bounded network design added another ingredient to iterative rounding, namely iteratively dropping some constraints. This paradigm of *iterative rounding and relaxation* has proved to be an extremely versatile and powerful tool in developing algorithms (and structural results) for network-design problems, and more generally in combinatorial optimization; see [18] for an extensive study.

Iterative rounding and relaxation derives its power from the fact that an extreme-point solution to an LP-relaxation of the problem can be defined using a structured family of tight constraints, such as, often a laminar family of cut constraints in the case of network-design problems. As noted earlier, this key property does not hold for cut-relative SNDP (and path-relative SNDP). Recently, various works have considered network-design problems that give rise to cut constraints that are less structured and not quite uncrossable; see [2, 3, 6] and the references therein. Our work can be seen as furthering this research direction.

2 Preliminaries and notation

Recall that we are given an undirected graph $G = (V, E)$ with nonnegative edge costs $\{c_e\}_{e \in E}$. For a subset $S \subseteq V$ and subset $Z \subseteq E$, which we will interchangeably view as the subgraph (V, Z) of G, we use $\delta_Z(S)$ to denote $\delta_G(S) \cap Z$, i.e., the edges of Z on the boundary of S. Recall that $\overline{S}$ denotes $V - S$, and $G[S]$ denotes the subgraph induced by S. A cut-requirement function on G is a function $f : 2^V \mapsto \mathbb{Z}$. (We allow $f(S)$ to be negative chiefly for notational simplicity; this does not affect anything, as we will only examine sets for which $f(S) \geq 0$.) Recall the following network-design problems.

- *f-network-design problem* (f-NDP): find a min-cost edge-set H such that $|\delta_H(S)| \geq f(S)$ for all $S \subseteq V$.
- *Cut-relative network-design problem* (CR-f-NDP): find a min-cost edge-set H such that $|\delta_H(S)| \geq g^{\text{CR-}f\text{-NDP}}(S) := \min\{f(S), |\delta_G(S)|\}$ for all $S \subseteq V$. This corresponds to $g^{\text{CR-}f\text{-NDP}}$-NDP. We call f the base cut-requirement function.
- *Survivable network design problem* (SNDP) is a special case of f-NDP. The input also specifies k tuples $\{(s_i, t_i, r_i)\}_{i \in [k]}$, where $s_i, t_i \in V$ and $r_i \in \mathbb{Z}_+$ for all $i \in [k]$. Defining $f^{\text{SNDP}}(S) := \max\{r_i : |S \cap \{s_i, t_i\}| = 1\}$ for all $S \subseteq V$, we obtain that SNDP is the f^{SNDP}-network-design problem.
- *Cut-relative* SNDP (CR-SNDP): The input here is the same as in SNDP. In Definition 1.1 (R2), CR-SNDP was defined as: find a min-cost $H \subseteq E$ such that for every $i \in [k]$, every $F \subseteq E$ with $|F| < r_i$, and every s_i-t_i cut S, we have $\delta_{H-F}(S) \neq \emptyset$ whenever $\delta_{G-F}(S) \neq \emptyset$. Later in Section 1, we stated that this can be equivalently formulated as CR-f-NDP with base cut-requirement function f^{SNDP}. We now justify this statement.
  Suppose that H is feasible under (R2) but $|\delta_H(S)| < \min\{f^{\text{SNDP}}(S), |\delta_G(S)|\}$ for some $S \subseteq V$. Suppose $f^{\text{SNDP}}(S) = r_i$. Then taking $F = \delta_H(S)$, we have $|F| < r_i$, $\delta_{G-F}(S) \neq \emptyset$, but $\delta_{H-F}(S) = \emptyset$, yielding a contradiction. So H is feasible for CR-f^{SNDP}-NDP. Conversely, suppose H is feasible for CR-f^{SNDP}-NDP. Consider any $i \in [k]$, any $F \subseteq E$ with $|F| < r_i$, and any s_i-t_i cut S with $\delta_{G-F}(S) \neq \emptyset$. Then, $|\delta_F(S)| < \min\{r_i, |\delta_G(S)|\} \leq \min\{f(S), |\delta_G(S)|\} \leq |\delta_H(S)|$; so $\delta_{H-F}(S) \neq \emptyset$, showing that H satisfies (R2).

We will consider network-design problems defined on various graphs and with various cut-requirement functions. Given a graph $D = (V_D, E_D)$ and cut-requirement function $h : 2^{V_D} \mapsto \mathbb{Z}$, we use (h, D)-NDP to denote f-NDP on the graph D with cut-requirement function h, and CR-(h, D)-NDP to denote CR-f-NDP on the graph D with base cut-requirement function h.

Properties of cut-requirement functions. Let $f : 2^V \mapsto \mathbb{R}$. We say that f is *weakly supermodular* if for any two node-sets $A, B \subseteq V$, we have $f(A) + f(B) \leq \max\{f(A \cap B) + f(A \cup B), f(A - B) + f(B - A)\}$. We say that f is *symmetric* if $f(S) = f(V - S)$ for all $S \subseteq V$, and is *normalized* if $f(\emptyset) = f(V) = 0$.

Since we are working on undirected graphs, we can always replace any cut-requirement function f by its symmetric, normalized version f^{sym}, where $f^{\mathsf{sym}}(S) := \max\{f(S), f(V-S)\}$ for all $\emptyset \neq S \subsetneq V$ and $f^{\mathsf{sym}}(\emptyset) = f^{\mathsf{sym}}(V) := 0$, without changing the resulting network-design problem(s), i.e., we have f-NDP $\equiv$ f^{sym}-NDP and CR-f-NDP $\equiv$ CR-f^{sym}-NDP.

Lemma 2.1 shows that weak-supermodularity is preserved under various operations. This will be quite useful in the analysis of our decomposition technique. Part a of the lemma shows that weak supermodularity is preserved under symmetrization and normalization, so for any f-NDP or CR-f-NDP instance involving a weakly-supermodular base cut-requirement function f, we may assume that f is symmetric and normalized.

All omitted proofs and details can be found in the full version of the paper.

▶ **Lemma 2.1.** *Let $f : 2^V \mapsto \mathbb{R}$ be weakly supermodular. Then*

(a) *f^{sym} is weakly supermodular.*

(b) *Let $w \in \mathbb{R}_+^E$. Define $g_1(S) := f(S) - w\big(\delta_G(S)\big)$ for all $S \subseteq V$. Then g_1 is weakly supermodular, and g_1 is symmetric if f is symmetric. In particular, for any $Z \subseteq E$, the function $f(S) - |\delta_Z(S)|$ is weakly supermodular, and is symmetric if f is symmetric.*

(c) *Let $X \subseteq V$, and $\mathcal{C}$ be a collection of subsets of $V-X$ closed under taking set intersections, unions, differences, and complements with respect to $V-X$, i.e., if $S \in \mathcal{C}$ then $V-X-S \in \mathcal{C}$, and $S, T \in \mathcal{C}$ implies that $\{S\cap T, S\cup T, S-T, T-S\} \subseteq \mathcal{C}$.*
Define the projection of f onto X with respect to $\mathcal{C}$ *as the function $g_2 : 2^X \mapsto \mathbb{R}$ given by $g_2(\emptyset) = g_2(X) := 0$, and $g_2(S) := \max_{T\in\mathcal{C}} f(S \cup T)$ for all $\emptyset \neq S \subsetneq X$. If f is symmetric, then g_2 is weakly supermodular and symmetric.*

Proof. Part (a) follows from a simple case analysis; part (b) follows because $w\big(\delta_G(S)\big)$ is symmetric, submodular. We focus on proving part (c). Essentially, the closure properties of $\mathcal{C}$ enable us to transfer the weak-supermodularity of f to the function g_2.

We first argue that g_2 is symmetric. If $A \in \{\emptyset, X\}$, we have $g_2(A) = g_2(X - A)$ by definition. Otherwise, if $g_2(A) = f(A \cup T_A)$, we have $f(A \cup T_A) = f\big(V - (A \cup T_A)\big) = f\big((X-A)\cup(V-X-T_A)\big) \leq g_2(X-A)$ where the last inequality follows since $V-X-T_A \in \mathcal{C}$. We also have $g_2(X - A) \leq g_2(A)$ by the same type of reasoning, and so $g_2(A) = g_2(X - A)$ for all $A \subseteq X$.

Now we show that g_2 is weakly supermodular. Consider $A, B \subseteq X$. If any of the sets $A-B, B-A, A\cap B, X-(A\cup B)$ are empty, then $g_2(A)+g_2(B) = g_2(A\cap B)+g_2(A\cup B)$ or $g_2(A) + g_2(B) = g_2(A - B) + g_2(B - A)$. The former holds when $A - B$ or $B - A$ is empty; the latter holds when $A \cap B$ or $X - (A \cup B)$ is empty, where we also utilize the fact that g_2 is symmetric. So assume otherwise. Then, for every $S \in \{A, B, A\cap B, A\cup B, A-B, B-A\}$, we have $g_2(S) = \max_{T\in\mathcal{C}} f(S \cup T)$.

Let $g_2(A) = f(A \cup T_A)$ and $g_2(B) = f(B \cup T_B)$, where $T_A, T_B \in \mathcal{C}$. Let $A' = A \cup T_A$ and $B' = B\cup T_B$. For any $S \in \{A'\cap B', A'\cup B', A'-B', B'-A'\}$, we can write $S = (S\cap X)\cup T$, where $T \in \mathcal{C}$ due to the closure properties of $\mathcal{C}$. It follows that $g_2(S \cap X) \geq f(S)$ for all $S \in \{A'\cap B', A'\cup B', A'-B', B'-A'\}$. So if $f(A')+f(B') \leq f(A'\cap B')+f(A'\cup B')$, we obtain that $g_2(A)+g_2(B) \leq g_2(A\cap B)+g_2(A\cup B)$, and if $f(A')+f(B') \leq f(A'-B')+f(B'-A')$, we obtain that $g_2(A) + g_{@}(B) \leq g_2(A - B) + g_2(B - A)$. ◀

The intuition behind the projection operation is as follows. Suppose f is a cut-requirement function and we know that no edges are picked from $E_G(V - X) \cup \delta_G(X)$. Then, roughly speaking, $g_2(S)$ captures the constraints that arise on $S \subseteq X$ due to f and $\mathcal{C}$.

2.1 LP-relaxation for CR-f-NDP

We consider the following natural LP-relaxation for CR-f-NDP. From now on, we assume that the cut-requirement function $f : 2^V \mapsto \mathbb{Z}$ is weakly supermodular, symmetric, and normalized. Given $\alpha \in \mathbb{R}^E$ and $F \subseteq E$, let $\alpha(F)$ denote $\sum_{e \in F} \alpha_e$.

$$\min \sum_{e \in E} c_e x_e \tag{CRLP$_{f,G}$}$$

$$\text{s.t.} \quad x\big(\delta_G(S)\big) \geq \min\big\{f(S), |\delta_G(S)|\big\} \qquad \forall S \subseteq V \tag{1}$$

$$0 \leq x_e \leq 1 \qquad \forall e \in E. \tag{2}$$

The LP-relaxation for f-NDP is similar, with the cut-constraints taking the simpler form $x\big(\delta_G(S)\big) \geq f(S)$ for all $S \subseteq V$. We denote this LP by (NDP$_{f,G}$). Our algorithm is based on iterative rounding, where we iteratively fix some x_e variables to 1. Suppose that we have fixed $x_e = 1$ for all $e \in Z_1$. Then, in that iteration, we are considering the residual problem on the subgraph $G' = (V, E' = E - Z_1)$, which leads to the following variant of (CRLP$_{f,G}$):

$$\min \sum_{e \in E'} c_e x_e \tag{P'}$$

$$\text{s.t.} \quad x\big(\delta_{G'}(S)\big) \geq \min\big\{f(S), |\delta_G(S)|\big\} - |\delta_{Z_1}(S)| \qquad \forall S \subseteq V \tag{3}$$

$$0 \leq x_e \leq 1 \qquad \forall e \in E'.$$

Defining $f'(S) = f(S) - |\delta_{Z_1}(S)|$ for all $S \subseteq V$, we can rephrase constraints (3) as $x\big(\delta_{G'}(S)\big) \geq \min\big\{f'(S), |\delta_{G'}(S)|\big\}$ for all $S \subseteq V$. Since f' is weakly supermodular, symmetric (Lemma 2.1 (b)), and normalized, (P') is of the same form as (CRLP$_{f,G}$): we have (P') $\equiv$ (CRLP$_{f',G'}$). This will be convenient for iterative rounding, as the residual problem is of the same form as the original problem. (CRLP$_{f,G}$) and (P') involve an exponential number of constraints, but using the ellipsoid method, we can find an extreme-point optimal solution provided we have a separation oracle for them. Recall that a separation oracle is a procedure that given a candidate point determines if it is feasible, or returns a violated inequality. Observe that a separation oracle for (CRLP$_{f,G}$) also yields a separation oracle for (P'), since (P') is simply (CRLP$_{f,G}$) after fixing some variables to 1.

▶ **Fact 2.2** (Ellipsoid method [14]). *Given a polytime separation oracle for* (CRLP$_{f,G}$)*, one can find an extreme-point optimal solution to LPs* (CRLP$_{f,G}$) *and* (P') *in polytime.*

For CR-SNDP, i.e., when $f = f^{\mathsf{SNDP}}$, a polytime separation oracle for (CRLP$_{f,G}$) can be obtained using min-cut computations. This was shown by [8] for path-relative k-ECSS, which is equivalent to cut-relative k-ECSS; refining their ideas yields a separation oracle for (CRLP$_{f,G}$) for general SNDP.

▶ **Lemma 2.3.** *When $f = f^{\mathsf{SNDP}}$, there is a polytime separation oracle for* (CRLP$_{f,G}$)*. Hence, one can obtain extreme-point optimal solutions to* (CRLP$_{f,G}$) *and* (P') *in polytime.*

Proof. The second statement follows from Fact 2.2. Let $(s_i, t_i, r_i)_{i \in [k]}$ be the k tuples defining the SNDP instance, which recall leads to the cut-requirement function $f^{\mathsf{SNDP}}(S) := \max\{r_i : |S \cap \{s_i, t_i\}| = 1\}$ for all $S \subseteq V$. Say that $S \subseteq V$ is a small cut if $f^{\mathsf{SNDP}}(S) > |\delta_G(S)|$, and it is a small cut for i if $|S \cap \{s_i, t_i\}| = 1$ and $r_i > |\delta_G(S)|$. Note that if $e \in \delta_G(S)$ for any small cut S, then we must have $x_e = 1$ in any feasible solution.

Now we can equivalently phrase the constraints of (CRLP$_{f,G}$) as follows: for any edge $a = uv \in E$, we must have $x_a = 1$ or, for every $i \in [k]$, the minimum s_i-t_i cut also separating u and v under the $\{x_e\}_{e \in E}$-capacities should have capacity at least r_i. So to detect if x is a feasible solution, we simply need to consider every edge $a = uv$ with $x_a < 1$ and every $i \in [k]$, and find the minimum s_i-t_i cut also separating u and v under the $\{x_e\}_{e \in E}$ capacities. If an s_i-t_i cut separates u, v, then u lies on the s_i-side and v lies on the t_i-side, or the other way around. So to find the minimum s_i-t_i cut also separating u and v, we find the minimum $\{s_i, u\}$-$\{t_i, v\}$ cut – by which we mean the minimum s_i-t_i cut where s_i, u are on the same side, and t_i, v are on the other side–and the minimum $\{s_i, v\}$-$\{t_i, u\}$-cut and take the one with smaller capacity. ◀

Extreme-point solution to $(\text{CRLP}_{f,G})$ not defined by a laminar family

We show that for the CR-SNDP instance shown in Fig. 1, the extreme-point solution $\widehat{x}$ shown in the figure, cannot be expressed as the solution to *any laminar family of tight cut constraints* (1) (along with tight edge constraints (2)). The SNDP instance is s-t 2-edge connectivity: that is, the base cut-requirement function is given by $f(S) = 2$ if S is an s-t cut, and $f(S) = 0$ otherwise.

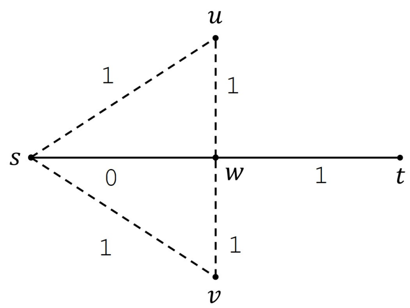

Figure 1 CR-SNDP instance, and an extreme-point solution $\widehat{x}$ to $(\text{CRLP}_{f,G})$. The solid edges have $\widehat{x}_e = 1$ and the dashed edges have $\widehat{x}_e = 0.5$, so $\widehat{x}_e = 0.5$ for $e \in \{su, uw, sv, vw\}$ and $\widehat{x}_e = 1$ for $e \in \{sw, wt\}$.

Note that the only small cut above containing s is $S = \{s, u, v, w\}$. One can verify by inspection that $\widehat{x}$ is feasible. It is not hard to see that $\widehat{x}$ is the unique solution to the following set of equations, which come from constraints that are tight at $\widehat{x}$.

$$x\big(\delta(S)\big) = 2 \qquad \forall S \in \big\{\{s\}, \{s, u, v\}, \{s, u, w\}, \{s, v, w\}\big\}, \tag{4}$$

$$x_e = 1 \qquad \forall e \in \{sw, wt\}. \tag{5}$$

This shows that $\widehat{x}$ is an extreme-point solution to $(\text{CRLP}_{f,G})$. With edge costs $c_{sw} = 0$ and $c_e = 1$ for all other edges, $\widehat{x}$ is also an optimal solution. The constraints $x\big(\delta(S)\big) \geq 2$ for all $S \in \big\{\{s\}, \{s, u, v\}, \{s, u, w\}, \{s, v, w\}\big\}$ imply that (since $x_e \leq 1$ for all e)

$$x_{su} + x_{sv} \geq 1, \quad x\big(\delta(u)\big) \geq 1, \quad x\big(\delta(v)\big) \geq 1, \quad x_{uw} + x_{vw} \geq 1.$$

Adding these constraints, we obtain that $x_{su} + x_{sv} + x_{uv} + x_{vw} \geq 2$. Therefore, the cost of any feasible solution is at least 3, and we have $c^T\widehat{x} = 3$.

$\widehat{x}$ cannot be defined via a laminar family. We now argue that $\widehat{x}$ cannot be defined by any laminar family of tight cut constraints and tight edge constraints. Note that any tight cut constraint corresponds to an s-t cut. So a laminar family $\mathcal{L}$ of tight cut constraints consists of a union of two chains $\mathcal{L}_1 \cup \mathcal{L}_2$, where a chain is a nested collection of sets: $\mathcal{L}_1$

comprising sets containing s but not t, and $\mathcal{L}_2$ comprising sets containing t but not s. Then $\mathcal{L}' = \mathcal{L}_1 \cup \{V - S : S \in \mathcal{L}_2\}$ is also a laminar family, yielding the same cut-constraints as $\mathcal{L}$, and and all sets in $\mathcal{L}'$contain s but not t. So we only need to consider laminar families of tight cut constraints corresponding to a chain of sets containing s and not t.

A simple argument now rules out the existence of any such chain that can uniquely define $\widehat{x}$ along with the tight edge constraints (5). Any such chain of sets $\mathcal{C}$ is obtained by ordering the elements u, v, w, and taking sets of the form $\{s\} \cup T$, where T is some prefix (including the empty prefix) of this ordered sequence; so $\mathcal{C}$ contains at most 4 sets. Since $\widehat{x}$ has 4 fractional edges, $E' = \{su, sv, uw, vw\}$, it must be that $\{\chi^{\delta_{E'}(S)}\}_{S \in \mathcal{C}}$ contains 4 linearly independent vectors. So $\mathcal{C}$ contains exactly 4 sets, and in particular the set $\{s, u, v, w\}$, and all vectors in $\{\chi^{\delta_{E'}(S)}\}_{S \in \mathcal{C}}$ are linearly independent. This yields a contradiction, since $\chi^{\delta_{E'}(\{s,u,v,w\})} = \vec{0}$.

3 Modeling power of CR-f-NDP

We prove that the CR-f-NDP Examples 1.2 and 1.3 listed under "Modeling power" in Section 1 model the stated problems. Recall that in both examples, we have a non-increasing function $\tau : \mathbb{Z}_+ \mapsto \mathbb{Z}_+$.

▶ **Theorem 3.1.** *Consider the base cut-requirement function f^{grace} in Example 1.2, where $f^{\mathsf{grace}}(S) := \min\{\ell : \tau(\ell) < |S|\}$, for $S \neq \emptyset$.*

(a) *The function f^{grace} is weakly-supermodular.*

(b) *H is feasible for CR-f^{grace}-NDP iff for every fault-set $F \subseteq E$, and every $S \subseteq V$ with $|S| \leq \tau(|F|)$, S is (the node-set of) a component of $G - F$ iff it is a component of $H - F$.*

Proof. The function f^{grace} is *downwards-monotone*: if $\emptyset \neq T \subseteq S$, then $f^{\mathsf{grace}}(T) \geq f^{\mathsf{grace}}(S)$. This follows from the definition, since τ is a non-increasing function. Weak-supermodularity of f^{grace} is now immediate since we have $f^{\mathsf{grace}}(A) + f^{\mathsf{grace}}(B) \leq f^{\mathsf{grace}}(A - B) + f^{\mathsf{grace}}(B - A)$, if $A - B, B - A \neq \emptyset$, and otherwise $f^{\mathsf{grace}}(A) + f^{\mathsf{grace}}(B) = f^{\mathsf{grace}}(A \cap B) + f^{\mathsf{grace}}(A \cup B)$.

Note that f^{grace} is not symmetric or normalized, but as indicated by Lemma 2.1, we can consider CR-f-NDP with the symmetric, normalized version of f^{grace} as the base cut-requirement function, without changing the problem.

To prove the feasibility characterization in part (b), we note that the function f^{grace} is constructed so that for any $S \subseteq V$ and $F \subseteq E$, we have $|S| \leq \tau(|F|)$ iff $f^{\mathsf{grace}}(S) > |F|$, so that $|\delta_H(S)| \geq f^{\mathsf{grace}}(S)$ implies that S is not a component of $H - F$. Suppose H is such that $|\delta_H(S)| \geq \min\{f^{\mathsf{grace}}(S), |\delta_G(S)|\}$ for all $S \subseteq V$. Consider any fault-set $F \subseteq E$ and any $S \subseteq V$ such that $|S| \leq \tau(|F|)$. By the above, if S is a component of $H - F$, then $|\delta_H(S)| < f^{\mathsf{grace}}(S)$, and so S must also be a component of $G - F$. But this also implies that if S is a component of $G - F$, then it is a component of $H - F$; otherwise, there is some $A \subsetneq S$ that is a component of $H - F$ with $|A| < |S| \leq \tau(|F|)$, and so A must be a component of $G - F$, which is not the case.

Conversely, suppose that for every $F \subseteq E$, and $S \subseteq V$ with $|S| \leq \tau(|F|)$, we have that S is a component of $G - F$ iff it is a component of $H - F$. If $|\delta_H(S)| < \min\{f^{\mathsf{grace}}(S), |\delta_G(S)|\}$ for some $S \subseteq V$, then take $F = \delta_H(S)$. Since $f^{\mathsf{grace}}(S) > |F|$, we have $|S| \leq \tau(|F|)$. But S is a component of $H - F$ and not $G - F$, which yields a contradiction. ◀

▶ **Theorem 3.2.** *Let $\pi : 2^V \mapsto \mathbb{R}_+$ be a monotone function, i.e., $\pi(T) \leq \pi(S)$ if $T \subseteq S$, as in Example 1.3. Consider CR-f-NDP, where the base cut-requirement function f is given by $f(S) := \min\{\ell : \tau(\ell) < \pi(S)\}$, for all $S \neq \emptyset$. Then*

(a) *f is weakly-supermodular;*

(b) *H is feasible for* CR-f-NDP *iff for every fault-set $F \subseteq E$, and every $S \subseteq V$ with $\pi(S) \leq \tau(|F|)$, S is (the node-set of) a component of $G - F$ iff it is a component of $H - F$.*

Proof. We mimic the proof of Theorem 3.1. Since π is monotone, we obtain that f is downwards-monotone, and hence weakly-supermodular.

For part (b), similar to before, we have that $\pi(S) \leq \tau(|F|)$ iff $f(S) > |F|$. Suppose H is feasible for CR-f-NDP. Consider any fault-set $F \subseteq E$ and any $S \subseteq V$ be such that $\pi(S) \leq \tau(|F|)$. By the above, $f(S) > |F|$, so if S is a component of $H - F$, then it must also be a component of $G - F$. Again, this also implies that if S is a component of $G - F$, then it is a component of $H - F$: otherwise, there is some $A \subsetneq S$ that is a component of $H - F$, and $\pi(A) \leq \pi(S) \leq \tau(|F|)$, which means that A is a component of $G - F$.

Conversely, suppose that for every $F \subseteq E$, and $S \subseteq V$ with $\pi(S) \leq \tau(|F|)$, we have that S is a component of $G - F$ iff it is a component of $H - F$. If $|\delta_H(S)| < \min\{f(S), |\delta_G(S)|\}$ for some $S \subseteq V$, then take $F = \delta_H(S)$. Since $f(S) > |F|$, we have $\pi(S) \leq \tau(|F|)$. But S is a component of $H - F$ and not $G - F$, which yields a contradiction. ◀

Our intent with Examples 1.2 and 1.3 is to illustrate that CR-f-NDP can be used to model some interesting problems. The choice of the τ and/or π functions can of course impact our ability to solve the LP-relaxation ($\text{CRLP}_{f,G}$) efficiently.

4 Structure of feasible CR-f-NDP solutions: a decomposition result

We now prove our decomposition result, showing that CR-f-NDP can be reduced to f-NDP with suitably-defined weakly-supermodular cut-requirement functions. Recall that the base cut-requirement function f is weakly supermodular (and symmetric, normalized), but the cut-requirement function $g^{\text{CR-}f\text{-NDP}}$ underlying CR-f-NDP need not be. Given $x \in \mathbb{R}^E$ and $S \subseteq V$, we use x^S to denote the restriction of x to edges in S, i.e., the vector $(x_e)_{e \in E(S)} \in \mathbb{R}^{E(S)}$.

▶ **Theorem 4.1 (Decomposition result).** *There is a partition $V_1, V_2, \ldots, V_k$ of V and weakly-supermodular, symmetric, and normalized functions $f_i : 2^{V_i} \mapsto \mathbb{Z}$ such that the following hold.*

(a) *$F \subseteq E$ is a feasible solution to* CR-(f, G)-NDP *iff $F \supseteq \delta_G(V_i)$ and $F(V_i)$ is a feasible solution to* $(f_i, G[V_i])$-NDP *for all $i \in [k]$.*

(b) *$x \in \mathbb{R}^E$ is feasible to ($\text{CRLP}_{f,G}$) iff for all $i \in [k]$, we have $x_e = 1$ for all $e \in \delta_G(V_i)$, and x^{V_i} is feasible to ($\text{NDP}_{f_i,G[V_i]}$).*

(c) *x is an extreme-point solution to ($\text{CRLP}_{f,G}$) iff for all $i \in [k]$, we have $x_e = 1$ for all $e \in \delta_G(V_i)$, and x^{V_i} is an extreme-point solution to ($\text{NDP}_{f_i,G[V_i]}$).*

Before delving into the proof of Theorem 4.1, we state the following immediate corollary of Theorem 4.1 (c), which directly leads to a 2-approximation algorithm for CR-f-NDP whenever ($\text{CRLP}_{f,G}$) can be solved efficiently.

▶ **Corollary 4.2.** *Let $\widehat{x}$ be an extreme-point solution to ($\text{CRLP}_{f,G}$). There exists some $e \in E$ for which $\widehat{x}_e \geq 1/2$.*

Proof. By Theorem 4.1 (c), there is a partition of V and f-NDP-instances defined on each part such that $\widehat{x}$ induces an extreme-point solution in each f-NDP-instance. By Jain's result, this implies that in each of these instances, there is some edge e with $\widehat{x}_e \geq 1/2$. ◀

The rest of this section is devoted to the proof of Theorem 4.1. For a set $S \subsetneq V$, define its *deficiency* $\mathrm{def}(S) := f(S) - |\delta_G(S)|$. We use $\mathrm{def}_{f,G}(S)$ (or $\mathrm{def}_{f,E}(S)$) when we want to make the cut-requirement function and graph (or edge-set) explicit. By Lemma 2.1 (b), $\mathrm{def}(S)$ is weakly supermodular; also, it is symmetric (and normalized). We say that S is a *small cut* (or more explicitly a *small-(f,G)-cut*) if $\mathrm{def}(S) > 0$. Clearly, if there is no small cut, then CR-f-NDP becomes equivalent to f-NDP. Our chief insight and key structural result is that if there is a small cut, then we can simplify the CR-f-NDP instance and move to smaller CR-f-NDP instances, by "splitting" the instance along a suitable small cut (Theorem 4.5). By repeating this operation, we eventually obtain an instance with no small cuts, which leads to Theorem 4.1. The splitting operation relies on the following modification of f.

▶ **Definition 4.3.** *Let $\emptyset \neq S \subsetneq V$. Recall that $\overline{S} := V - S$. Let $Z = \delta_G(S)$. Define the function $f_S : 2^S \mapsto \mathbb{Z}$, which we call the* restriction of f to S, *as follows.*

$$f_S(T) = \begin{cases} 0; & \text{if } T = \emptyset \text{ or } T = S \\ \max\{f(T) - |\delta_Z(T)|, f(T \cup \overline{S}) - |\delta_Z(T \cup \overline{S})|\} & \text{otherwise} \end{cases} \qquad \forall T \subseteq S.$$

The intuition here is that for $T \subseteq S$, $f_S(T)$ is a lower bound on the residual requirement that needs to be met from $\delta_{G[S]}(T)$ given that we pick all edges of Z. Note that since f and $|\delta_Z(.)|$ are symmetric, $f(T \cup \overline{S}) - |\delta_Z(T \cup \overline{S})| = f(S - T) - |\delta_Z(S - T)|$, and so f_S is the symmetric, normalized version of of the residual requirement function $f - |\delta_Z(\cdot)|$.

In the splitting operation, we pick a small cut A of *maximum* deficiency, and move to CR-f-NDP instances on the smaller graphs $G[A]$, and $G[\overline{A}]$ obtained by restricting f to A, and restricting f to $\overline{A}$ respectively (Theorem 4.5). Lemma 4.4 states some basic properties of the restriction of f. Part (b) will be quite useful in proving that the splitting operation maintains feasibility, and part (a) allows us to repeat the splitting operation on the smaller instances. For a set $S \subseteq V$, and set $F \subseteq E$ of edges, we use $F(S)$ to denote the edges of F with both ends in S.

▶ **Lemma 4.4.** *Let $\emptyset \neq S \subsetneq V$. Let f_S be the restriction of f to S. Let $\delta_G(S) \subseteq F \subseteq E$.*

(a) *f_S is weakly supermodular, symmetric, and normalized.*

(b) *Let $\emptyset \neq T \subsetneq S$. Then, $\mathrm{def}_{f_S,F(S)}(T) = \max\{\mathrm{def}_{f,F}(T), \mathrm{def}_{f,F}(T \cup \overline{S})\}$.*

Proof. Let $Z = \delta_G(S)$. Define $f' : 2^V \mapsto \mathbb{R}$ by $f'(T) := f(T) - |\delta_Z(T)|$. Then, f' is weakly supermodular and symmetric (Lemma 2.1 (b)).

For part (a), we can proceed in two ways. As noted earlier, f_S is the symmetric, normalized version f', so part (a) follows from Lemma 2.1 (a). Alternatively, observe that f_S is the projection of f' onto S with respect to $\mathcal{C} := \{\emptyset, \overline{S}\}$, which is closed under set intersections, unions, differences, and complements with respect to $V - S$. So by Lemma 2.1 (c), we obtain that f_S is weakly supermodular and symmetric. Also, f_S is normalized by definition.

For part (b), the stated equality follows by simply plugging in f_S, and noting that for $Y \in \{T, T \cup \overline{S}\}$, we have $|\delta_F(Y)| = |\delta_{F(S)}(T)| + |\delta_Z(Y)|$. ◀

▶ **Theorem 4.5** (Splitting along a small cut)**.** *Suppose that there is some small cut. Let $A \subseteq V$ be a maximum-deficiency cut. Note that $\mathrm{def}(A) > 0$ and $\emptyset \neq A \subsetneq V$.*

(a) *$F \subseteq E$ is a feasible solution to CR-(f,G)-NDP iff $F \supseteq \delta_G(A)$, $F(A)$ is a feasible solution to CR-$(f_A, G[A])$-NDP, and $F(\overline{A})$ is a feasible solution to CR-$(f_{\overline{A}}, G[\overline{A}])$-NDP.*

(b) *$x \in \mathbb{R}^E_+$ is feasible to $(\mathrm{CRLP}_{f,G})$ iff $x_e = 1$ for all $e \in \delta_G(A)$, x^A is feasible to $(\mathrm{CRLP}_{f_A,G[A]})$, and $x^{\overline{A}}$ is feasible to $(\mathrm{CRLP}_{f_{\overline{A}},G[\overline{A}]})$.*

Proof. Part (a) can be seen as the special case of part (b), where x is integral. We prove part (a) here, as the proof is somewhat (notational) simpler. For notational simplicity, let $(f_1, G_1, F_1) = \big(f_A, G[A], F(A)\big)$ and $(f_2, G_2, F_2) = \big(f_{\overline{A}}, G[\overline{A}], F(\overline{A})\big)$. Let $Z = \delta_G(A)$.

The "only if" direction follows easily from the definition of the restriction of f. Let F be feasible to CR-(f, G)-NDP. Since A is a small cut, we must have $F \supseteq Z$. Consider any $\emptyset \neq T \subsetneq A$. Since F is feasible, we have $|\delta_F(T)| \geq \min\{f(T), |\delta_G(T)|\}$, which implies that $|\delta_{F_1}(T)| \geq \min\{f(T) - |\delta_Z(T)|, |\delta_{G_1}(T)|\}$. Similarly, we have $|\delta_F(T \cup \overline{A})| \geq \min\{f(T \cup \overline{A}), |\delta_G(T \cup \overline{A})|\}$, so $|\delta_{F_1}(T)| \geq \min\{f(T \cup \overline{A}) - |\delta_Z(T \cup \overline{A})|, |\delta_{G_1}(T)|\}$. Combining the inequalities, and using the definition of restriction, we obtain that $|\delta_{F_1}(T)| \geq \min\{f_1(T), |\delta_{G_1}(T)|\}$. Since T was arbitrary, this shows that F_1 is feasible to CR-(f_1, G_1)-NDP.

A symmetric argument shows that F_2 is a feasible solution to CR-(f_2, G_2)-NDP.

Conversely, suppose that $F \supseteq \delta_G(A)$, $F_1 = F(A)$ is a feasible solution to CR-(f_1, G_1)-NDP, and $F_2 = F(\overline{A})$ is a feasible solution to CR-(f_2, G_2)-NDP. Recall that $Z = \delta_G(A)$. Consider any $\emptyset \neq T \subsetneq V$. Lemma 4.4 (b) can be used to readily argue that if $T \subseteq A$ or $T \subseteq \overline{A}$, then the constraint for set T in CR-(f, G)-NDP is satisfied. We show this for $T \subseteq A$; a symmetric argument applies for $T \subseteq \overline{A}$. If T is a small-(f_1, G_1)-cut, then $F_1 \supseteq \delta_{G_1}(T)$. Since $F \supseteq Z$, this implies that $F \supseteq \delta_G(T)$. Otherwise, we have $\mathrm{def}_{f_1,F_1}(T) \leq 0$, so by Lemma 4.4 (b), we have that $\mathrm{def}_{f,F}(T) \leq \mathrm{def}_{f_1,F_1}(T) \leq 0$. Thus, we always have $|\delta_F(T)| \geq \min\{f(T), |\delta_G(T)|\}$.

Now consider T such that $T \cap A, T \cap \overline{A} \neq \emptyset$. We exploit the weak supermodularity of $\mathrm{def}_{f,F}$, and that A has maximum deficiency, to bound $\mathrm{def}_{f,F}(T)$ in terms of the deficiency of sets that do not cross A, and thus show that the constraint for T is satisfied. We have

$$\mathrm{def}_{f,F}(T) + \mathrm{def}_{f,F}(A) \leq \mathrm{def}_{f,F}(T \cap A) + \mathrm{def}_{f,F}(T \cup A) = \mathrm{def}_{f,F}(T \cap A) + \mathrm{def}_{f,F}(\overline{A} - T) \quad \text{or}$$
$$\mathrm{def}_{f,F}(T) + \mathrm{def}_{f,F}(A) \leq \mathrm{def}_{f,F}(A - T) + \mathrm{def}_{f,F}(T - A)$$

where in the first inequality we also use the fact that $\mathrm{def}_{f,F}$ is symmetric. In both cases, we have $\mathrm{def}_{f,F}(T) + \mathrm{def}_{f,F}(A) \leq \mathrm{def}_{f,F}(X) + \mathrm{def}_{f,F}(Y)$ where $X \subseteq A$, $Y \subseteq \overline{A}$ and $\delta_G(T) \subseteq \delta_G(X) \cup \delta_G(Y) \cup Z$. As shown previously, for $S \in \{X, Y\}$, we have $|\delta_F(S)| \geq \min\{f(S), |\delta_G(S)|\}$, or equivalently, $\mathrm{def}_{f,F}(S) \leq \max\{0, \mathrm{def}_{f,G}(S)\}$. Note that $\mathrm{def}_{f,F}(A) = \mathrm{def}_{f,G}(A) > 0$ since $F \supseteq \delta_G(A)$. Since A is a maximum-deficiency small cut, this also implies that $\mathrm{def}_{f,F}(A) \geq \mathrm{def}_{f,F}(S)$ for $S \in \{X, Y\}$: we have $\mathrm{def}_{f,F}(S) \leq \max\{0, \mathrm{def}_{f,G}(S)\} \leq \max\{0, \mathrm{def}_{f,G}(A)\} = \mathrm{def}_{f,F}(A)$.

If both X and Y are small-(f, G)-cuts, then we are done since then we have that $F \supseteq \delta_G(X) \cup \delta_G(Y) \cup Z \supseteq \delta_G(T)$. Otherwise, since $\mathrm{def}_{f,F}(A) \geq \max\{\mathrm{def}_{f,F}(X), \mathrm{def}_{f,F}(Y)\}$, we have $\mathrm{def}_{f,F}(T) \leq \min\{\mathrm{def}_{f,F}(X), \mathrm{def}_{f,F}(Y)\} \leq 0$, since at least one of X, Y is not a small-(f, G)-cut. ◀

Proof of Theorem 4.1. Parts (a) and (b) follow easily by induction on the number of nodes, using parts (a) and (b) of Theorem 4.5. Consider part (a). The base case is when there is no small-(f, G)-cut, in which case, we have $V_1 = V$, and $f_1 = f$. Otherwise, let $A = \mathrm{argmax}_{S \subseteq V} \mathrm{def}_{f,G}(S)$. Since f_A and $f_{\overline{A}}$ are weakly-supermodular, we can recurse and induct on CR-$(f_A, G[A])$-NDP and CR-$(f_{\overline{A}}, G[\overline{A}])$-NDP instances, both of which have fewer nodes, to obtain partitions of A and $\overline{A}$; we combine these to obtain $V_1, \ldots, V_k$. Theorem 4.5 (a) and the induction hypothesis then readily yield part (a) here. The only thing to note is that since the induction hypothesis yields $F \supseteq \delta_{G[A]}(V_i)$ for a part $V_i \subseteq A$, and $F \supseteq \delta_G(A)$ by Theorem 4.5 (a), we have $F \supseteq \delta_G(V_i)$. Similarly, $F \supseteq \delta_G(V_i)$ for any part $V_i \subseteq \overline{A}$.

Part (b) follows similarly from Theorem 4.5 (b). Part (c) follows from part (b), because it is not hard to see, using part (b), that any convex combination of feasible solutions to $(\mathrm{CRLP}_{f,G})$ yields a convex combination of feasible $(\mathrm{NDP}_{f_i, G[V_i]})$-solutions, and vice versa. ◀

One may wonder if the $(f_i, G[V_i])$-NDP-instances in the decomposition of Theorem 4.1 can be specified more directly or succinctly. In the full version, we prove some properties of the decomposition given by its proof, which as a *consequence*, enable one to do so. But we need our splitting approach to argue why this succinct description fulfills the properties in Theorem 4.1. The properties we establish also imply that the decomposition in Theorem 4.1 can be computed efficiently given suitable algorithmic access to the base cut-requirement function f.

5 Algorithmic implication: LP-relative 2-approximation algorithm

We now exploit the decomposition result stated in Theorem 4.1, and Corollary 4.2, to obtain a 2-approximation algorithm. We assume that we have a separation oracle for $(\text{CRLP}_{f,G})$, which implies that we can efficiently find an extreme-point optimal solution to $(\text{CRLP}_{f,G})$ and any LP that arises in a subsequent iteration after fixing some x_e variables to 1 (Fact 2.2).

Algorithm 1 CRNDP-Alg.

The input is $\big(G = (V, E), \{c_e \geq 0\}, f : 2^V \mapsto \mathbb{Z}\big)$, where f is weakly-supermodular (and symmetric, normalized). We assume there is a separation oracle for $(\text{CRLP}_{f,G})$.

A1. Initialize $F \leftarrow \emptyset$, $G' = (V, E') \leftarrow G$, and $f' \leftarrow f$.

A2. While F is not a feasible solution (which can be detected using the separation oracle), repeat:

- **A2.1** Find an extreme-point optimal solution $\widehat{x}$ to $(\text{CRLP}_{f',G'})$.
- **A2.2** Let $Z_1 = \{e \in E' : \widehat{x}_e \geq 1/2\}$. Update $F \leftarrow F \cup Z_1$, $E' \leftarrow E' - Z_1$, and $f'(S) = f'(S) - |\delta_{Z_1}(S)|$ for all $S \subseteq V$.

A3. Return F.

▶ Remark 5.1. The collection of small cuts *does not change across iterations*; that is, S is a small-(f, G)-cut iff it is a small-(f', G')-cut in some other iteration. This is because, for any $S \subseteq V$, we change $|\delta_{G'}(S)|$ and $f'(S)$ by the same amount in every iteration. In the first iteration, we have $\widehat{x}_e = 1$ for every $e \in \delta_G(S)$ and every small cut S, and so $\delta_{G'}(S) = \emptyset$ in any subsequent iteration. So we do not have any constraints for small cuts in subsequent iterations, and this causes problems with uncrossing tight constraints. Nevertheless, we are able to argue, due to our decomposition result, that there is always some edge with $\widehat{x}_e \geq 1/2$.

▶ **Theorem 5.2.** *Let OPT denote the optimal value of* $(\text{CRLP}_{f,G})$*. Algorithm* CRNDP-Alg *returns a feasible* CR-f-NDP*-solution $F \subseteq E$ such that $c(F) \leq 2 \cdot OPT$.*

Proof. By Corollary 4.2, at least one edge gets added to F in every iteration, so the algorithm terminates in at most $|E|$ iterations. The cost bound follows by induction on the number of iterations to termination. The base case is trivial. Suppose $F = Z_1 \cup F'$, where F' is the solution found in subsequent iterations for the instance CR-(f', G')-NDP. Let OPT' be the optimal solution to $(\text{CRLP}_{G',f'})$. Then, by induction, we have $c(F_1) \leq 2 \cdot OPT'$ and $OPT' \leq \sum_{e \in E'} c_e \widehat{x}_e$, since $(\widehat{x}_e)_{e \in E'}$ is a feasible solution to $(\text{CRLP}_{f',G'})$. We also have $c(Z_1) \leq 2 \cdot \sum_{e \in Z_1} c_e \widehat{x}_e$. Therefore, $c(F) \leq 2c^T \widehat{x} = 2 \cdot OPT$. ◀

Since $(\text{CRLP}_{f^{\text{SNDP}},G})$ admits a polytime separation oracle (Lemma 2.3), we obtain the following result as a corollary.

▶ **Theorem 5.3.** *There is a 2-approximation algorithm for* CR-SNDP*.*

6 Hardness result

We now show that cut-relative SNDP and path-relative SNDP are APX-hard, even when the base SNDP instance involves only one s-t pair, by showing that the k-edge connected subgraph (k-ECSS) problem can be cast as a special case of these problems. (This hints at the surprising amount of modeling power of {cut, path}- relative SNDP.) Previously, even NP-hardness of path-relative SNDP in the s-t case was not known.

Let $\big(G = (V, E), \{c_e\}_{e\in E}\big)$ be an instance of k-ECSS. We may assume that G is at least k-edge connected, i.e., $|\delta_G(S)| \geq k$ for all $\emptyset \neq S \subsetneq V$, as otherwise there is no feasible solution (and this can be efficiently detected).

Consider the following path-relative SNDP instance. We add a source s and sink t, and edges sv, vt of cost 0, for all $v \in V$. Let $G' = (V', E')$ denote this graph. The base-SNDP instance asks for $(k+n)$-edge connectivity between s and t, where $n = |V|$. That is, the base cut-requirement function f is given by $f(S') = k + n$ if $S' \subseteq V'$ is an s-t cut, and $f(S') = 0$ otherwise. We show that path-relative SNDP and cut-relative SNDP are actually equivalent in this case, and that they encode k-ECSS on the graph G.

▶ **Theorem 6.1.** *$H \subseteq E'$ is a feasible solution to the above {cut, path}-relative* SNDP *instance $\iff$ $H \supseteq \delta_{G'}(s) \cup \delta_{G'}(t)$ and $H(V)$ is a feasible k-*ECSS*-solution for G. Hence, this {cut,path}-relative* SNDP *instance is the same as k-*ECSS* on the graph G.*

Proof. Recall that we are assuming that G is (at least) k-edge connected. The following observation will be handy. The only small cuts in G' are $\{s\}$ and $\{t\}$ (and their complements). For any other s-t cut $S' \subseteq V'$, taking $S = S' \cap V$, we have that $\emptyset \neq S \subsetneq V$. So we have $|\delta_{G'}(S')| = |\delta_G(S)| + |S| + |V - S| \geq k + n$.

$\Leftarrow$ direction. We argue that H is feasible for cut-relative SNDP, which also implies that it is feasible for path-relative SNDP. Clearly, H covers both small cuts. For any other s-t cut $S' \subseteq V'$, taking $S = S' \cap V$, we have $|\delta_H(S')| \geq |\delta_{H(V)}(S)| + |S| + |V - S| \geq k + n$, since $H(V)$ is a feasible k-ECSS-solution for G.

$\Rightarrow$ direction. Suppose that H is feasible for path-relative SNDP. If there is some edge $sv \notin H$, then consider the fault-set $F = \delta_{G'}(s) - \{sv\}$. Then $H - F$ clearly has no s-t path, since $\delta_{H-F}(s) = \emptyset$. But $G' - F$ has an s-t path sv, vt. So we must have $H \supseteq \delta_{G'}(s)$. A similar argument shows that $H \supseteq \delta_{G'}(t)$.

We next argue that $|\delta_H(S')| \geq n + k$ for every s-t cut S' for which $\emptyset \neq S := S' \cap V \subsetneq V$. This is equivalent to showing that $|\delta_{H(V)}(S)| \geq k$, since $|\delta_H(S') \cap \delta_{G'}(s)| + |\delta_H(S') \cap \delta_{G'}(t)| = n$. Assume that $s \in S'$ without loss of generality. Suppose $|\delta_{H(V)}(S)| < k$. Since $|\delta_G(S)| \geq k$, there is some edge $uv \in \delta_G(S) - \delta_{H(V)}(S)$. Suppose $u \in S$. Consider the fault-set $F = \delta_{H(V)}(S) \cup \{sw : w \in V - S\} \cup \{wt : w \in S\}$. Then $|F| < k + n$ and $H - F$ has no s-t path, since we have constructed F to ensure that $\delta_{H-F}(S') = \emptyset$. However, $G - F$ has an s-t path: su, uv, vt. This contradicts that H is feasible for path-relative SNDP.

So we have shown both that $H(V)$ is feasible k-ECSS solution for G, and that H is feasible for cut-relative SNDP.

Since the cost of the edges in $E' - E$ is 0, the above {cut, path}-relative SNDP instance is exactly the same as k-ECSS on the graph G. ◀

7 Non-equivalence of path-relative SNDP and cut-relative SNDP

We consider the same graph and base SNDP instance as in Fig. 1. The graph is reproduced below for convenience, and recall that the base cut-requirement function models s-t 2-edge connectivity: so $f(S) = 2$ if S is an s-t cut, and $f(S) = 0$ otherwise.

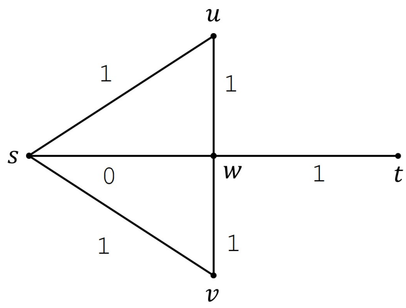

Figure 2 An instance where path-relative SNDP and cut-relative SNDP are not equivalent. The number labeling an edge gives the cost of the edge.

A feasible solution to path-relative SNDP is given by the edge-set $H_1 = \{su, uw, sw, wt\}$, which has cost 3. (This is in fact an optimal solution to path-relative SNDP.)

However, H_1 *is not feasible for cut-relative* SNDP. This is because, we require $\delta_H(\{s, u, w\}) \geq \min\{2, 3\} = 2$ in any feasible CR-SNDP solution H. Moreover, one can infer that any feasible CR-SNDP-solution must include all but one of the edges from $\{su, uw, sw, sv, vw\}$. This is because for any $F \subseteq \{su, uw, sw, sv, vw\}$ with $|F| = 2$, one can verify that there is an s-t cut S such that $|\delta_G(S)| = 3$ and $F \subseteq \delta_G(S)$; so $|\delta_{G-F}(S)| \leq 1$, which implies that there is no feasible solution that excludes F. A feasible CR-SNDP-solution must also include the edge wt, so the optimal value for cut-relative SNDP is 4.

While cut-relative SNDP and path-relative SNDP are not in general equivalent, they are closely related. Dinitz et al. [9] showed (see Lemma A.1 in [9]) that one can give a cut-based formulation for path-relative SNDP, where H is feasible iff $|\delta_H(S)| \geq \min\{f(S), |\delta_G(S)|\}$ for every $S \subseteq V$ *such that* $G[S]$ *and* $G[V' - S]$ *are connected graphs*, where $V' \subseteq V$ is the vertex-set of the component containing S. That is, we require the cut-constraints of CR-SNDP to hold for a suitable collection of node-sets (as opposed to all node-sets in CR-SNDP). This immediately implies that CR-SNDP and path-relative SNDP are equivalent when the underlying graph is a (capacitated) complete graph. Also, as mentioned earlier, path-relative k-ECSS and cut-relative k-ECSS are equivalent [8].

References

1 A. Agrawal, Phil Klein, and R. Ravi. When trees collide: An approximation algorithm for the generalized Steiner problem on networks. *SIAM Journal on Computing*, 24:440–456, 1995. doi:10.1137/S0097539792236237.

2 Ishan Bansal. A global analysis of the primal-dual method for edge-augmentation problems. In *Proceedings of the 26th IPCO*, pages 58–71, 2025. doi:10.1007/978-3-031-93112-3_5.

3 Ishan Bansal, Joseph Cheriyan, Logan Grout, and Sharat Ibrahimpur. Improved approximation algorithms by generalizing the primal-dual method beyond uncrossable functions. *Algorithmica*, 86(8):2575–2604, 2024. doi:10.1007/S00453-024-01235-2.

4 Nikhil Bansal, Rohit Khandekar, Jochen Könemann, Viswanath Nagarajan, and Britta Peis. On generalizations of network design problems with degree bounds. *Mathematical Programming*, 141(1-2):479–506, 2013. doi:10.1007/S10107-012-0537-8.

5 Nikhil Bansal, Rohit Khandekar, and Viswanath Nagarajan. Additive guarantees for degree-bounded directed network design. *SIAM Journal on Computing*, 39(4):1413–1431, 2009. doi:10.1137/080734340.

6 Sylvia C. Boyd, Joseph Cheriyan, Arash Haddadan, and Sharat Ibrahimpur. Approximation algorithms for flexible graph connectivity. *Mathematical Programming*, 204:493–516, 2024. doi:10.1007/S10107-023-01961-5.

7 Chandra Chekuri and Rhea Jain. Approximation algorithms for network design in non-uniform fault models. In *Proceedings of the 50th ICALP*, pages 36:1–36:20, 2023. doi:10.4230/LIPICS.ICALP.2023.36.

8 Michael Dinitz, Ama Koranteng, and Guy Kortsarz. Relative survivable network design. In *Proceedings of the 25th APPROX*, pages 41:1–41:19, 2022. doi:10.4230/LIPICS.APPROX/RANDOM.2022.41.

9 Michael Dinitz, Ama Koranteng, Guy Kortsarz, and Zeev Nutov. Improved approximations for relative survivable network design. In *Proceedings of the 21st WAOA*, pages 190–204, 2023. doi:10.1007/978-3-031-49815-2_14.

10 M. Goemans and D. Williamson. A general approximation technique for constrained forest problems. *SIAM J. on Computing*, 24:296–317, 1995. doi:10.1137/S0097539793242618.

11 M. Goemans and D. Williamson. The primal-dual method for approximation algorithms and its application to network design problems. In Dorit Hochbaum, editor, *Approximation algorithms for NP-hard problems*, chapter 4. PWS Publishing Company, 1996.

12 Michel X. Goemans, Andrew V. Goldberg, Serge A. Plotkin, David B. Shmoys, Éva Tardos, and David P. Williamson. Improved approximation algorithms for network design problems. In *Proceedings of the 5th SODA*, pages 223–232, 1994. URL: http://dl.acm.org/citation.cfm?id=314464.314497.

13 M.X. Goemans. Minimum bounded-degree spanning trees. In *Proceedings of the 47th FOCS*, pages 273–282, 2006.

14 Martin Grötschel, Lászlo Lovász, and Lex Schrijver. *Geometric algorithms and combinatorial optimization*. Algorithms and Combinatorics 2. Springer, 1993.

15 Kamal Jain. A factor 2 approximation algorithm for the generalized Steiner network problem. *Combinatorica*, 21(1):39–60, 2001. doi:10.1007/S004930170004.

16 N. Klein and N. Olver. Thin trees for laminar families. In *Proceedings of the 64th FOCS*, 2023.

17 Lap Chi Lau, Seffi Naur, Mohammad Salavatipour, and Mohit Singh. Survivable network design with degree or order constraints. *SIAM Journal on Computing*, 39(3):1062–1087, 2009. doi:10.1137/070700620.

18 Lap Chi Lau, R. Ravi, and Mohit Singh. *Iterative methods in combinatorial optimization*. Cambridge University Press, 2011.

19 Mohit Singh and Lap Chi Lau. Approximating minimum bounded degree spanning trees to within one of optimal. In *Proceedings of the 39th STOC*, pages 661–670. ACM, 2007. doi:10.1145/1250790.1250887.

20 David P. Williamson, Michel X. Goemans, Milena Mihail, and Vijay V. Vazirani. A primal-dual approximation algorithm for generalized steiner network problems. *Combinatorica*, 15(3):435–454, 1995. doi:10.1007/BF01299747.