Improved Approximations for Relative Survivable Network Design

Michael Dinitz[*]
Johns Hopkins University
mdinitz@cs.jhu.edu

Ama Koranteng[†]
Johns Hopkins University
akorant1@jhu.edu

Guy Kortsarz
Rutgers University, Camden
guyk@camden.rutgers.edu

Zeev Nutov
The Open University, Israel
nutov@openu.ac.il

October 4, 2023

Abstract

One of the most important and well-studied settings for network design is edge-connectivity requirements. This encompasses uniform demands such as the Minimum k-Edge-Connected Spanning Subgraph problem, as well as nonuniform demands such as the Survivable Network Design problem (SND). In a recent paper by [Dinitz, Koranteng, Kortsarz APPROX '22], the authors observed that a weakness of these formulations is that we cannot consider fault-tolerance in graphs that have small cuts but where *some* large fault sets can still be accommodated. To remedy this, they introduced new variants of these problems under the notion *relative* fault-tolerance. Informally, this requires not that two nodes are connected if there are a bounded number of faults (as in the classical setting), but that two nodes are connected if there are a bounded number of faults *and the two nodes are connected in the underlying graph post-faults.*

This seemingly minor change in definition makes the problem dramatically more complex, and so the results in [Dinitz, Koranteng, Kortsarz APPROX '22] are quite limited. For the Relative Survivable Network Design problem (RSND) with non-uniform demands, they are only able to give a nontrivial result when there is a single demand with connectivity requirement 3—a non-optimal 27/4-approximation. We strengthen this result in two significant ways: We give a 2-approximation for RSND where *all requirements are at most* 3, and a $2^{O(k^2)}$-approximation for RSND with a single demand of *arbitrary value k*. To achieve these results, we first use the "cactus representation" of minimum cuts to give a lossless reduction to normal SND. Second, we extend the techniques of [Dinitz, Koranteng, Kortsarz APPROX '22] to prove a generalized and more complex version of their structure theorem, which we then use to design a recursive approximation algorithm.

1 Introduction

Fault-tolerance has been a central object of study in approximation algorithms, particularly for network design problems where the graphs we study represent physical objects which might fail (communication links, transportation links, etc.). In these settings it is natural to ask for whatever object we build to be fault-tolerant. The precise definition of "fault-tolerance" varies in different settings, but a common formulation is edge fault-tolerance, which typically takes the form of edge connectivity. Informally, these look like guarantees of the form "if up to k edges fail, then the nodes I want to be connected are still connected." For example, consider the following classical fault-tolerance problem.

Definition 1. In the Survivable Network Design problem (SND, sometimes referred to as Generalized Steiner Network) we are given an edge-weighted graph G and demands $\{(s_i, t_i, k_i)\}_{i \in [\ell]}$, and we are supposed to find the minimum-weight subgraph H of G so that there are at least k_i edge-disjoint paths between s_i and t_i for every $i \in [\ell]$. In other words, for every $i \in [\ell]$, if fewer than k_i edges fail then s_i and t_i will still be connected in H even after failures.

The Survivable Network Design problem is well-studied (see [16, 24, 26, 31] for a sample); notably, Jain gives a 2-approximation algorithm for the problem in a seminal paper [26]. Beyond SND, edge fault-tolerance has been studied in many related network design contexts, with the k-Edge Connected Spanning Subgraph, Fault-Tolerant Group Steiner Tree, Fault-Tolerant Spanner, and Fault-Tolerant Shortest Paths problems being just a few examples (see [4, 16, 19, 27] for examples). These and other classical fault-tolerance problems, including the Survivable Network Design problem, are *absolute* fault-tolerance problems—if up to k objects fail, the remaining graph should function as desired. This differs from the stronger notion of fault-tolerance introduced in [17], called *relative* fault-tolerance. Relative fault-tolerance makes guarantees that rather than being absolute ("if at most k edges fail the network still functions") are relative to an underlying graph or system ("if at most k edges fail, the subgraph functions just as well as the original graph post-failures").

Relative fault-tolerance is therefore a natural generalization of absolute fault-tolerance: If the input graph has the desired connectivity, then the relative fault-tolerance and absolute fault-tolerance definitions are equivalent. However, if the input graph does not have the requested connectivity, then relative fault-tolerance allows us to return a solution with interesting and nontrivial guarantees while absolute fault-tolerance forces us to return nothing. In this way, relative fault-tolerance overcomes a significant weakness of absolute fault-tolerance.

This relative fault-tolerance definition was inspired by a recent line of work on relative notions of fault-tolerance for graph spanners and emulators [5–10, 12, 18, 19]. In these settings, the goal is generally to find existential bounds and algorithms to achieve them, rather than to do optimization. In [17], by contrast, their approach takes the point of view of optimization and approximation algorithms. With this notion of fault-tolerance in network design, the authors of [17] define the relative version of the Survivable Network Design problem.

Definition 2. In the Relative Survivable Network Design problem (RSND), we are given a graph $G = (V, E)$ with edge weights $w : E \to \mathbb{R}_{\geq 0}$ and demands $\{(s_i, t_i, k_i)\}_{i \in [\ell]}$. A feasible solution is a subgraph H of G where for all $i \in [\ell]$ and $F \subseteq E$ with $|F| < k_i$, if there is a path in $G \setminus F$ from s_i to t_i then there is also a path in $H \setminus F$ from s_i to t_i. Our goal is to find the minimum weight feasible solution.

Note that if s_i and t_i are k_i-connected in G for every $i \in [\ell]$, then RSND is exactly the same as SND. If in G there exists some $i \in [\ell]$ such that s_i and t_i are *not* k_i-connected, then although there is no solution for SND, there is a meaningful RSND solution. We note that the fault-tolerance we

achieve is really "one less" than the given number (there are strict inequalities in the definitions). This is "off-by-one" from the related relative fault-tolerance literature (see the definitions in [12]), but makes the connection to the traditional SND formulation cleaner.

There has been recent work on a related network design model introduced by Adjiashvili [1–3, 11,14]. In this model, E is partitioned into "safe" and "unsafe" edges. Informally, in the Flex-SNDP problem we are given a graph $G = (V, E)$ with edge costs and with functions $p, q : V \times V \to \mathbb{Z}^+$. We must return a min cost subgraph such that for each vertex pair u, v, they are $p(u, v)$-connected after deleting any subset of up to $q(u, v)$ unsafe edges. Like RSND, Flex-SNDP is a natural generalization of SND. However, it is an absolute fault-tolerance problem since it does not consider the underlying connectivity of the input. No polynomial-time approximation algorithms are known for general Flex-SNDP, though there has been recent work on several special cases [3, 11, 13, 14].

The results of [17]. Although relative fault-tolerance is a natural and promising generalization of fault-tolerance, the results given in [17] for the RSND problem are quite limited. Outside of a 2-approximation algorithm for the special case in which all demands are identical, [17] is only able to give algorithms for some of the simplest RSND special cases. First, they give an extremely simple 2-approximation for the RSND special case where all demands are in $\{0, 1, 2\}$ (also known as the 2-RSND problem). The algorithm falls out of the observation that there is only a difference between a relative demand of 2 and an absolute demand of 2 when there is a cut of size one separating the vertex demand pair. Cuts of size one are very easy to handle, allowing for a simple and straightforward reduction to SND.

Cuts of size two or larger are significantly more difficult to reason about, and so the 2-RSND algorithm does not extend to larger demands. As a result of this more complex cut structure, [17] is only able to handle demands of value 3 (and reason about the size two cuts between them) when there is only *a single demand*, with value 3 (also known as the SD-3-RSND problem). Despite this being an extremely restricted special case of RSND, the algorithm and analysis given by [17] are quite complex, depending on a careful graph decomposition involving "important separators" (a concept from fixed-parameter tractability [29]). Moreover, this algorithm only achieved a 27/4-approximation for the problem, far from the 2-approximation (or even exact algorithm) that one might hope for.

The limited results of [17] show that while relative fault-tolerance is an attractive notion, applying it to the Survivable Network Design problem significantly changes the structure of the problem and makes it difficult to reason about and develop algorithms for. For example, while [17] only gives a 27/4-approximation for SD-3-RSND, there is an exact polynomial-time algorithm for the SND equivalent (by a simple reduction to the Min-Cost Flow problem). So one might worry that relative fault-tolerance is simply too difficult of a definition, and the results of [17] are limited precisely because nothing is possible for even slightly more general settings.

1.1 Our Results and Techniques

In this paper, we seek to alleviate this worry by providing improved bounds for generalizations of the settings considered in [17]. In particular, we study two natural generalizations of the SD-3-RSND problem (which [17] provided a 27/4-approximation for). First, rather than only a single demand with value at most 3, can we handle an *arbitrary* number of demands that are at most 3? Secondly, in the single demand setting, instead of only handling a demand of at most 3, can we generalize to *arbitrary* values?

1.1.1 3-RSND

We begin with the setting where all demands are at most 3, but there can be an arbitrary number of such demands. We call this the 3-RSND problem. Note that, as discussed, there are no previous results for this setting, and the most related result is a 27/4-approximation if there is only a single such demand [17]. We prove the following theorem.

Theorem 1.1. *There is a polynomial-time 2-approximation for the 3-RSND problem.*

To obtain this theorem, we use entirely different techniques from those used in [17]. Most notably, we use the *cactus representation* of the global minimum cuts (which in this case are 2-cuts) of the input graph. The cactus representation of global minimum cuts is well studied and has been leveraged in a number of settings (see [20, 21, 25, 28, 30] for a sample). While it can be defined and constructed for more general connectivity values, for our setting we can construct the cactus representation by contracting components with certain connectivity properties. This results in a so-called *cactus graph*, which at a high level is a "tree of cycles": every pair of cycles intersects on at most one component in the construction. This cactus graph now has a simple enough structure that it allows us to reduce the original problem to a simpler problem in each of the contracted components. That is, we are able to show that certain parts of the cactus are essentially "forced", while other parts are not necessary, so the only question that remains is what to do "inside" of each cactus vertex, i.e., each component. This reduction makes the connectivity demands inside each component more complicated, but fortunately we are guaranteed 3-connectivity between special vertices inside the component. Hence we can use Jain's 2-approximation for SND [26] without worrying about the relative nature of the demands.

1.1.2 SD-k-RSND

Our second improvement is orthogonal: rather than allowing for more demands of at most 3, we still restrict ourselves to a single demand but allow it to be a general constant k rather than 3. We call this the SD-k-RSND problem. As with the 3-RSND problem discussed earlier, there are no known results for this problem. We prove the following theorem:

Theorem 1.2. *There is a polynomial-time $2^{O(k^2)}$-approximation for the SD-k-RSND problem.*

To prove this, we extend the technique used by [17] for the $k = 3$ case. They construct a "chain" of 2-separators (cuts of size 2 that are also important separators) so that in each component in the chain, there are no 2-cuts between the incoming separator and the outgoing separator. They are then able to use this structure to characterize the connectivity requirement of any feasible solution restricted to that component. To extend this technique, we use important separators of size up to $k - 1$ to carefully construct a *hierarchy* of chains. The hierarchy has $k - 1$ levels of nested components, so that for each component in the ith level of the hierarchy, there are no cuts of size at most i between the incoming and outgoing separators. There are multiple ways of constructing such a hierarchy, but we prove that a particular construction yields a hierarchy with a number of useful but delicate properties within a single level and between different levels of the hierarchy. With these properties, we can characterize the complex connectivity requirement of any feasible solution when restricted to a component in the hierarchy. Once we have this structure theorem, we approximate the optimal solution in each component of the hierarchy via a recursive algorithm.

1.1.3 Simplification of k-EFTS

The k-Edge Fault Tolerant Subgraph problem (k-EFTS) is the special case of RSND where all demands are identical: every two nodes have a demand of exactly k. A 2-approximation for k-

EFTS was recently given in [17] via a somewhat complex proof; in particular, they defined a new property called *local weak supermodularity* and used it to show that Jain's iterative rounding still gave the same bounds in the relative setting. In Appendix D we give a simplification of this proof. It turns out that local weak supermodularity is not actually needed, and a more classical notion of $\mathcal{F}$-supermodularity suffices. This allows us to reduce to previous work in a more black-box manner.

2 Preliminaries

We will consider the following special cases of RSND (Definition 2):

- The k-RELATIVE SURVIVABLE NETWORK DESIGN problem (k-RSND) is the special case of RSND where $r(s,t) \leq k$ for all $s,t \in V$. In this paper we consider the case $k = 3$, namely, the 3-RSND problem.

- The SINGLE DEMAND k-RELATIVE SURVIVABLE NETWORK DESIGN problem (SD-k-RSND) is the special case of RSND where $r(s,t) = k$ for exactly one vertex pair $s,t \in V$ and there is no demand for any other vertex pairs (equivalently, all other demands are 0). We consider the full SD-k-RSND problem for arbitrary k.

- The k-EDGE-FAULT-TOLERANT-SUBGRAPH problem (k-EFTS) is the special case of RSND where $r(s,t) = k$ for all $s,t \in V$.

For each of the listed RSND problem variants, we will use the following notation and definitions throughout. Let $G = (V, E)$ be a (multi-)graph and H a spanning subgraph (or an edge subset) of G. For $A \subseteq V$, let $\delta_H(A)$ denote the set of edges in H with exactly one endpoint in A, and let $d_H(A) = |\delta_H(A)|$ be their number. Additionally, let $G[A]$ denote the subgraph of G induced by the vertex set A. Let $s, t \in V$. We say that A is an st-**set** if $s \in A$ and $t \notin A$, and that $\delta_G(A)$ (or $\delta_E(A)$) is an st-**cut** of G (induced by A). An st-cut $\delta_G(A)$ (or an st-set A) is G-**minimal** if $\delta_G(A)$ contains no other st-cut of G. Assuming G is connected, it is easy to see that $\delta_G(A)$ is G-minimal if and only if both $G[A]$ and $G[V \setminus A]$ are connected. One can also see that if an st-cut $X \subseteq E$ is G-minimal, then $X = \delta_G(A)$ for some $A \subseteq V$. Finally, let $\lambda_G(s,t)$ denote the size of a min st-cut in G.

By Theorem 17 of [17], we may assume without loss of generality that the input graph G of any RSND instance is 2-edge-connected (which we will call "2-connected").

3 2-Approximation for 3-RSND (and SD-3-RSND)

Given an RSND instance, we say that a vertex pair s, t is a k-**demand** if $r(s,t) = k$. We call a k-demand **relative** if the minimum st-cut has size less than k; that is, if $\lambda_G(s,t) < k$. A k-demand is then **ordinary** if $\lambda_G(s,t) \geq k$. Recall that SD-3-RSND has only one demand st, and that it is a 3-demand. The edges of any size 2 st-cut, or 2-st-cut, belong to any feasible solution so we call them **forced edges**. As a result, we can assume without loss of generality that they have cost 0.

3.1 Overview

We first give an overview of the theorems and proofs in this section. In order to prove Theorem 1.1, we will show that we can replace a single *relative* 3-demand by an equivalent set of *ordinary* 3-demands. More formally, we will prove the following.

Theorem 3.1. *Given an* SD-3-RSND *instance, there exists a polynomially computable set of ordinary 3-demands, D, such that for any $H \subseteq E$ that contains all forced edges, H is a feasible* SD-3-RSND *solution if and only if H satisfies all demands in D.*

This theorem reduces SD-3-RSND to the ordinary 3-SND problem (that is, the special case of SND where all demands are at most 3). In fact, this also gives us a lossless reduction from 3-RSND to 3-SND: Given a 3-RSND instance, we include the forced edges of all 3-demands into our solution, replace each relative 3-demand by an equivalent set of ordinary demands, and obtain an equivalent ordinary 3-SND instance. Since SND admits approximation ratio 2, this reduction from 3-RSND to 3-SND implies Theorem 1.1.

We will also show that SD-3-RSND is approximation equivalent to certain instances of a special case of 3-SND. Before we define this special case, we must give a definition. A vertex subset R is a **k-edge-connected subset** in a graph H if $\lambda_H(u, v) \geq k$ for all vertex pairs $u, v \in R$. Since the relation $\{(u, v) \in V \times V : \text{no } (k - 1)\text{-cut separates } u, v\}$ is transitive, this is equivalent to requiring that $\lambda_H(u, v) \geq k$ for pairs u, v that form a tree on R. We will prove that SD-3-RSND is approximation equivalent to special instances of the following problem:

4-SUBSET 3-EC
Input: A graph $J = (V', E')$ with edge costs, and a set $R \subseteq V'$ of at most 4 terminals.
Output: A min-cost subgraph H of J, such that R is 3-edge-connected in H.

More specifically, we will prove the following.

Theorem 3.2. *Let s and t be vertices in $J = (V', E')$, where J is the input graph to an instance of* 4-SUBSET 3-EC. SD-3-RSND *admits approximation ratio ρ if and only if* 4-SUBSET 3-EC *with the following properties* (A,B) *admits approximation ratio ρ:*

(A) $d_J(s) = d_J(t) = 2$ and R is the set of neighbors of s, t.

(B) If $d_J(A) = 2$ for some st-set A, then $A = \{s\}$ or $A = V' \setminus \{t\}$. Namely, if F is a set of 2 edges of J such that $J \setminus F$ has no st-path, then $F = \delta_J(s)$ or $F = \delta_J(t)$.

The general 4-SUBSET 3-EC problem admits approximation ratio 2, since it is a special case of SND. However, it is not actually known whether 4-SUBSET 3-EC is in P or is NP-hard. To the best of our knowledge, the complexity status of even 3-SUBSET 2-EC is open. Both the 4-SUBSET 3-EC and 3-SUBSET 2-EC problems are related to the special case of SND where the sum of the demands is bounded by a constant. This special case has been studied; for example, both [15] and [22] studied this and related questions. Nevertheless, the best known approximation algorithm for the 4-SUBSET 3-EC problem gives a 2-approximation. As such, our 2-approximation for 3-RSND is the best we can hope for. In the rest of this section, we prove Theorems 3.2, 3.1, and 1.1. All missing proofs can be found in Appendix A.

3.2   Cactus Representation and Definitions

We first give some definitions and describe the cactus representation. The relation $\{(u, v) \in V \times V : \text{no } (k - 1)\text{-cut separates } u, v\}$ is an equivalence, and we will call its equivalence classes **k-classes**. We construct a **cactus** $\mathcal{G}$ by shrinking every nontrivial 3-class (that is, every 3-class with at least 2 nodes) of the input graph G. Note that since G is 2-connected, $\mathcal{G}$ is a connected graph in which every two cycles have at most one node in common; see Fig. 1(a) for an example. Going forward, we will identify every 3-class with the corresponding node of $\mathcal{G}$. The edge pairs that belong to the same cycle of $\mathcal{G}$ are the 2-cuts of G.

5
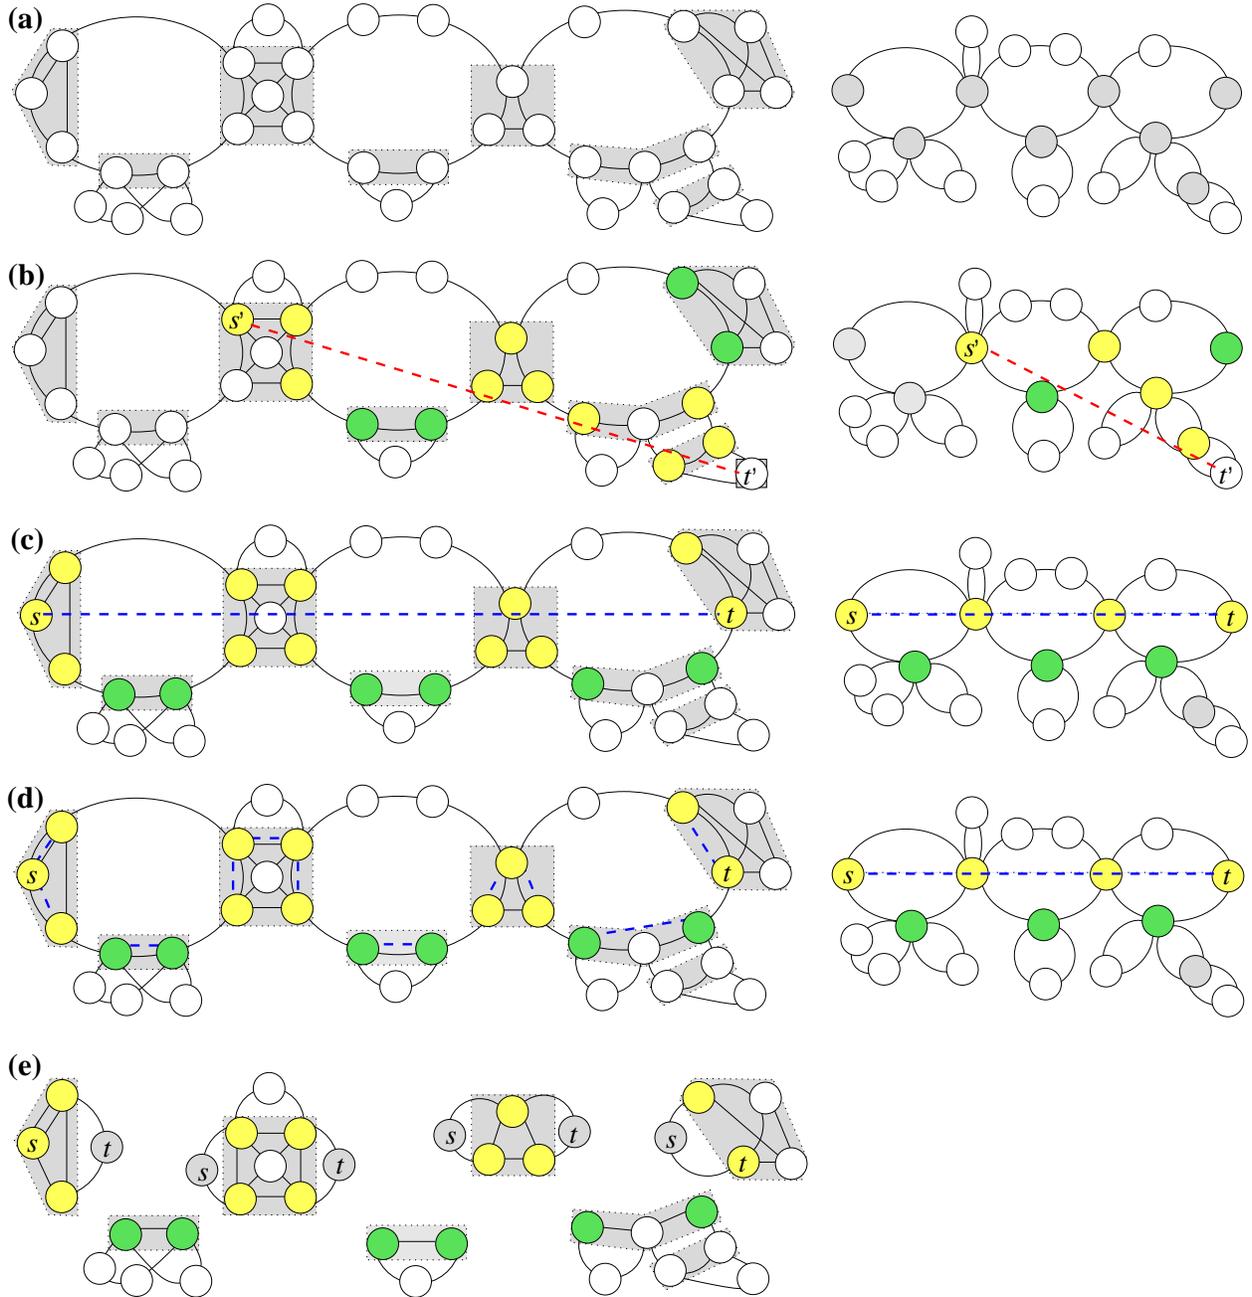
Figure 1: (a) 3-classes (shown by gray areas) and the cactus $\mathcal{G}$ of G; (b) 3-demands $s't'$ (red dashed line) and the attachment nodes of $s't'$-relevant 3-classes, where attachment nodes of central classes are yellow and of non-central classes are green; (c) 3-demands st (blue dashed line) and the attachment nodes of st-relevant 3-classes; (d) a set of ordinary 3-demands that is equivalent to the relative 3-demand st. (e) components w.r.t. the 3-demand st.

We will assume that vertex pair st is a relative 3-demand (otherwise we have an instance of SD-3-SND which can be solved exactly in polynomial time). We say that the *st*-**chain of cycles** of $\mathcal{G}$ consists of all the cycles (and their nodes) in $\mathcal{G}$ that contain a 2-st-cut. We say that the nodes, or 3-classes, on these cycles are *st*-**relevant**. Note that the set of edges in $\mathcal{G}$ that are in the st-chain of cycles are the forced edges. We also say that an st-relevant 3-class is **central** if it contains s or t, or belongs to two cycles of the st-chain. Additionally, the **attachment nodes** of an st-relevant 3-class are nodes in the 3-class that are either s or t, or are the ends of the edges (the **attachment edges**) that belong to some cycle in the st-chain of cycles. Since G is 2-connected, the number of attachment nodes in a non-central 3-class is exactly 2, while the number of attachment nodes in a central 3-class is between 2 and 4; see Fig. 1(b,c).

3.3   Proof of Theorems 3.2, 3.1, and 1.1

For the proof of Theorems 3.2 and 3.1, we associate with each st-relevant 3-class, C, a certain graph G_C which we call the **component of** C, obtained as follows (see Fig. 1(e)):

- If C is a non-central 3-class then, in the graph obtained from G by removing the two attachment edges of C, G_C is the connected component that contains C.

- If C is a central 3-class, then removing the attachment edges of C results in at least one and at most two connected components that do not contain C – one contains s and the other contains t, if any. We obtain G_C from G by contracting the connected component that contains s into node s, and contracting the connected component that contains t into node t.

We now modify the central components G_C to satisfy properties (A,B) from Theorem 3.2. Consider some central 3-class C, and consider its component $J = G_C$. If J does not contain one of the original nodes s or t, then it has properties (A,B) and no modification is needed. If J contains the original node s, then we rename s to s', add a new node s, and connect new s by two zero cost edges to s'. The obtained J now has properties (A,B). A similar transformation applies if J contains the original node t.

The following lemma is about both the non-central components and these *modified* central components; in the lemma, we show that for H to be a feasible SD-3-RSND solution, it is necessary and sufficient to satisfy certain connectivity properties within each component. Note that a subgraph H is a feasible solution to SD-3-RSND if for any $F \subseteq E$ with $|F| \leq 2$, the following holds: if there is an st-path in $G \setminus F$, then there is an st-path in $H \setminus F$.

Lemma 3.3. *Let H be a subgraph of G, and suppose that H contains all forced edges. Subgraph H is a feasible* SD-3-RSND *solution if and only if for every component J, the following holds (see Fig, 1(e)).*

(i) *If J is a non-central component, then $H[J]$ contains two edge-disjoint uv-paths, where u and v are the two attachment nodes of J.*

(ii) *If J is a central component, then $H[J]$ is a feasible solution to the* SD-3-RSND *instance in J (with demand $r(s,t) = 3$).*

Suppose that for the special SD-3-RSND instances that arise in the central components we can achieve approximation ratio α. Then, we can achieve ratio α for general SD-3-RSND by picking into our solution H three types of edge sets.

1. The forced edges.

7

2. A min-cost set of 2 edge-disjoint paths between the attachment nodes of each st-relevant non-central component.

3. An α-approximate solution in each st-relevant central component.

Note that edges picked in steps 1,2 do not invoke any cost in the approximation ratio, since by Lemma 3.3 we actually pick parts of an optimal solution. Thus we get that the approximability of SD-3-RSND is equivalent to the approximability of the very special instances that arise in the central components. We will now show that these special instances from the central components are in fact instances of 4-Subset 3-EC with properties (A,B) from Theorem 3.2, thus proving Theorem 3.2. We will consider only central components with 4 attachment nodes; other cases with 3 or 2 attachment nodes are similar.

In what follows, let $\mathcal{I}$ be an SD-3-RSND instance on input graph J with properties (A,B) (just as in our central components). Let $R = \{x, y, z, w\}$ where x, y are the neighbors of s and z, w are the neighbors of t and let H be a subgraph of J that includes the four forced edges sx, sy, zt, and wt. Now we prove that H is a feasible solution for $\mathcal{I}$ if and only if R is 3-connected in H. Lemma 3.4 gives the "only-if" direction, while Lemma 3.5 gives the "if" direction.

Lemma 3.4. *If subgraph H is a feasible solution for instance $\mathcal{I}$, then $R = \{x, y, z, w\}$ is a 3-edge-connected subset in H.*

Lemma 3.5. *If $R = \{x, y, z, w\}$ is a 3-edge-connected subset in subgraph H, then H is a feasible solution for instance $\mathcal{I}$.*

By Lemmas 3.4 and 3.5, H is a feasible solution for $\mathcal{I}$ if and only if H includes all forced edges and R is a 3-edge-connected subset—that is, R forms a feasible solution to 4-Subset 3-EC—in H. This, along with Lemma 3.3, implies that the approximability of SD-3-RSND is equivalent to that of 4-Subset 3-EC with properties (A,B), concluding the proof of Theorem 3.2.

Proof of Theorem 3.1 We will prove that a single relative 3-demand st can be replaced by an equivalent forest of ordinary 3-demands in polynomial time, where the trees in this forest span the sets of attachment nodes of the st-relevant 3-classes; see Fig. 1(d). Recall that by Lemma 3.3 and Lemmas 3.4,3.5, subgraph H is a feasible SD-3-RSND solution for 3-demand st if and only if the following holds for every st-relevant 3-class C:

(i) If C is central, then the set R_C of attachment nodes of C is a 3-connected subset in H.

(ii) If C is non-central, then $H[C]$ contains 2 edge-disjoint uv-paths, where u and v are the attachment nodes of C.

The first condition is equivalent to satisfying a clique of 3-demands on R_C.[1] For the second condition, consider a non-central st-relevant 3-class C with attachment nodes u, v. One can see that if H contains all forced edges and satisfies (i,ii) then the number of edge-disjoint uv-paths in H is larger by exactly 1 than their number in $H[C]$—the additional path (that exists in H but not in $H[C]$) goes along the cycle of the cactus that contains C, and there is exactly one such path. Thus, the demand $r(u, v) = 3$ is equivalent to requiring two edge-disjoint paths from u to v in R_C (in addition to including all forced edges).

---

[1]Recall that since the relation $\{(u, v) \in V \times V :$ no 2-cut separates $u, v\}$ is transitive, this is equivalent to having a tree of 3-demands on R_C.

Thus we obtain an equivalent 3-SND instance by replacing the single relative 3-demand st by a set D of 3-demands that form a clique (or, which is equivalent, a tree) on the set R_C of attachment nodes of every st-relevant 3-class C. These new demands can be computed in polynomial time, and they are ordinary 3-demands, since each R_C is a 3-edge-connected subset in G. This concludes the proof of Theorem 3.1.

Proof of Theorem 1.1   Finally, we give a 2-approximation for 3-RSND in the following theorem. We treat each demand in the 3-RSND instance as its own instance of SD-3-RSND, solve each SD-3-RSND instance, and include the edges of each solution in our output.

4   SD-k-RSND

We give a recursive $2^{O(k^2)}$-approximation algorithm for SD-k-RSND for arbitrary constant k (Theorem 1.2). The algorithm is a generalization of the SD-3-RSND algorithm from [17]. At a high level, the main idea is to partition the input graph using a hierarchy of important separators, prove a structure theorem that characterizes the required connectivity guarantees within each component of the hierarchy, and then achieve these guarantees using a variety of subroutines: a Weighted st Shortest-Path algorithm, the recursive SD-k-RSND approximation algorithm, and a Min-Cost Flow algorithm.

4.1   Hierarchical Chain Decomposition

In this section we define important separators and describe how to use them to construct a hierarchical k-chain decomposition of G.

Definition 3. Let X and Y be vertex sets of a graph G. An (X, Y)-*separator* of G is a set of edges S such that there is no path between any vertex $x \in X$ and any vertex $y \in Y$ in $G \setminus S$. An (X, Y)-separator S is *minimal* if no subset $S' \subset S$ is also an (X, Y)-separator. If $X = \{x\}$ and $Y = \{y\}$, we say that S is an (x, y)-separator.

In [17], the authors provide the following definition.

Definition 4 (Definition 20 of [17])**.** Let S be an (X, Y)-separator of graph G, and let R be the vertices reachable from X in $G \setminus S$. Then S is an *important* (X, Y)-separator if S is minimal and there is no (X, Y)-separator S' such that $|S'| \leq |S|$ and $R' \subset R$, where R' is the set of vertices reachable from X in $G \setminus S'$.

In Section 4.1 of [17], the authors describe how to construct the "$s - t$ 2-chain" of a graph G [2]. Here, we define the (X, Y) *h-chain* of G similarly, where X and Y are vertex sets and $h > 0$ is an integer. Instead of using size two important separators to partition the graph, we will use size h important separators. See Figure 2 for an example.

First, if there are no important (X, Y)-separators of size h in G, then the (X, Y) h-chain of G is simply G and we're done (the chain is a single component, G, with no separators). If such an important separator exists, then we first find an important (X, Y)-separator S_0^h of size h in G, and we let R_0^h be the set of vertices reachable from any vertex $x \in X$ in $G \setminus S_0^h$. We let $V_{(0,r)}^h$ be the vertices in R_0^h incident on S_0^h, and let $V_{(1,\ell)}^h$ be the nodes in $V \setminus R_0^h$ incident on S_0^h. We then proceed

---

[2]Note that all separator-based chain definitions given in this section are unrelated to the cactus-based chain definitions in Section 3.
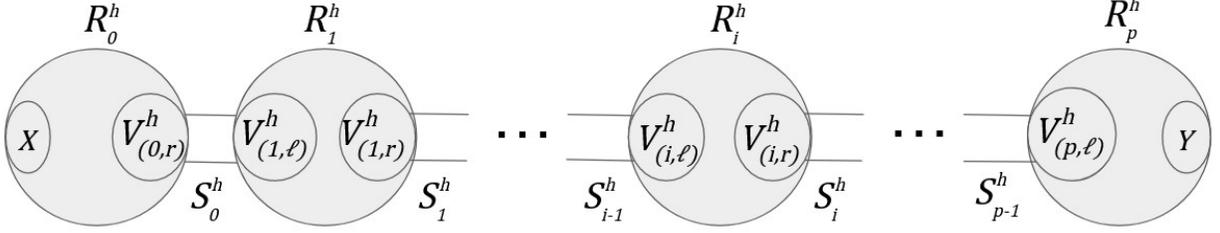
Figure 2: An (X, Y) h-chain of a graph G, where $h = 2$.

inductively. Given $V_{(i,\ell)}^h$, if there is no important $(V_{(i,\ell)}^h, Y)$-separator of size h in $G \setminus (\cup_{j=0}^{i-1} R_j^h)$ then the chain is finished. Otherwise, let S_i^h be such a separator, let R_i^h be the nodes reachable from $V_{(i,\ell)}^h$ in $(G \setminus (\cup_{j=0}^{i-1} R_j^h)) \setminus S_i^h$, let $V_{(i,r)}^h$ be the nodes in R_i^h incident on S_i^h, and let $V_{(i+1,\ell)}^h$ be the nodes in $V \setminus (\cup_{j=0}^i R_j^h)$ incident on S_i^h. After this process completes we have our (X, Y) h-chain, consisting of components $R_0^h, \ldots, R_p^h$ along with important separators $S_0^h, \ldots, S_{p-1}^h$ between the components. See Figure 2.

Next we note that by Lemma 17 of [17], we can find an important (X, Y)-separator of size h in polynomial time as long as h is a constant.

Lemma 4.1 (Lemma 17 of [17]). *Let $d \geq 0$. An important (X, Y)-separator of size d can be found in time $4^d \cdot n^{O(1)}$ (if one exists), where $n = |V|$.*

Constructing the Hierarchical k-chain Decomposition. Now we describe how to construct the *hierarchical k-chain decomposition* of G. We start by creating the (s, t) 2-chain of G. We say that each component of the (s, t) 2-chain is a *2-component* of G in the hierarchical chain decomposition, and that G is the *1-component* of the decomposition. We also say that G is the *parent component* of each 2-component.

We then proceed inductively. Let R_i^h be an h-component of the hierarchical k-chain decomposition. If $h = k - 1$, then the decomposition is finished. Otherwise, build the $(V_{(i,\ell)}^h, V_{(i,r)}^h)$ $(h+1)$-chain of R_i^h. The $(h + 1)$-chain consists of $(h + 1)$-components. Component R_i^h is the parent of these $(h + 1)$-components (and these $(h + 1)$-components are the children of R_i^h). Note that the vertex set $V_{(i,\ell)}^h$ is also in the first $(h+1)$-component in R_i^h and that $V_{(i,r)}^h$ is in the last $(h+1)$-component in R_i^h. After this process completes we have our hierarchical k-chain decomposition of G.

The set of all h-components can be ordered as follows: The h-component that contains s is the first component while the h-component that contains t is last. All other h-components are adjacent via a left important separator and a right important separator to a left neighbor h-component and a right neighbor h-component, respectively. Observe that in this hierarchical decomposition, an h-component R_i^h has no $(V_{(i,\ell)}^h, V_{(i,r)}^h)$-separator of size h or less.

4.2 Structure Theorem

Preliminaries. We say a subgraph H satisfies the RSND demand (X, Y, d) on input graph G if the following is true: If there is a path from at least one vertex in X to at least one vertex in Y in $G \setminus F$, where F is a set of at most $d - 1$ edges, then there is a path from at least one vertex in X to at least one vertex in Y in $H \setminus F$. Going forward, if $V_{(i,\ell)}^h = \{s\}$, then we consider S_{i-1}^h to be the empty set. Similarly, if $V_{(i,r)}^h = \{t\}$, then S_i^h is the empty set.

10

Fix an h-component R_i^h and let X be a vertex set such that $X \subseteq V_{(i,\ell)}^h$. We say that S_X is the set of edges in S_{i-1}^h incident on vertices in X. Similarly, if Y is a vertex set such that $Y \subseteq V_{(i,r)}^h$, we say that S_Y is the set of edges in S_i^h incident on vertices in Y. We will also use S to denote the set of all edges in an important separator in the hierarchical chain decomposition. Let H be a subgraph of G. We will also say that $G_i^h = G[R_i^h]$ is the subgraph of G induced by the vertex set R_i^h, and that $H_i^h = H[R_i^h]$ is the subgraph of H induced by R_i^h.

We can now use the hierarchical chain construction to give a structure lemma that characterizes feasible solutions. The lemma states that a subgraph H of G is a feasible solution to SD-k-RSND if and only if in the hierarchical k-chain decomposition of G, all edges in S are in H, and certain connectivity requirements between groups of vertices in $V_{(i,\ell)}^h$ and in $V_{(i,r)}^h$ are met in H_i^h for each component R_i^h in the decomposition.

Theorem 4.2 (Structure Theorem). *Subgraph H is a feasible solution to SD-k-RSND if and only if all edges in S are included in H, and for each h-component R_i^h in the hierarchical k-chain decomposition of input graph G, subgraph H_i^h satisfies the following:*

1. *H_i^h is a feasible solution to RSND on subgraph G_i^h with demands*

$$\Big\{ (X, Y, d) \ : \ X \subseteq V_{(i,\ell)}^h, \ Y \subseteq V_{(i,r)}^h, \ (X, Y) \neq \big(V_{(i,\ell)}^h, V_{(i,r)}^h\big),$$
$$d = \max(0, k + |S_X| + |S_Y| - |S_{i-1}^h| - |S_i^h|) \Big\}.$$

2. *H_i^h is a feasible solution to RSND on subgraph G_i^h with demand*

$$\Big(V_{(i,\ell)}^h, V_{(i,r)}^h, h + 1\Big).$$

3. *H_i^h is a feasible solution to RSND on subgraph G_i^h with demand*

$$\Big(V_{(i,\ell)}^h, V_{(i,r)}^h, k - 1\Big).$$

The proof of this structure theorem is long and involved; it can be found in Appendix B. We describe the proof here at a very high level. To prove the "only if" direction, we first assume that we are given some feasible solution H. Then for each of the properties in Theorem 4.2, we assume it is false and derive a contradiction by finding a fault set $F \subseteq E$ with $|F| < k$ where there is a path from s to t in $G \setminus F$, but not in $H \setminus F$. The exact construction of such an F depends on which of the properties of Theorem 4.2 we are analyzing.

To prove the 'if" direction, we assume that H satisfies the conditions of Theorem 4.2 and consider a fault set $F \subseteq E$ with $|F| < k$ where s and t are connected in $G \setminus F$. We want to show that s and t are connected in $H \setminus F$. We analyze a special "fault subchain" of the hierarchical chain decomposition that has specific properties, including that it contains all the edges in F. We prove that at least one vertex, v, at the end of that subchain is reachable from s in $H \setminus F$ (we also show that t is reachable from v in $H \setminus F$ to complete the proof). Specifically, we consider two cases. In Case 1, all fault edges in F are in the same $(k-1)$-level component in the hierarchical chain decomposition; this component is the fault subchain in this case. In Case 2 there is some level, h, such that F is not entirely contained in a single h-level component, but at level $h-1$, F is completely contained in a single $(h-1)$-level component, C. In this case, the fault subchain is the set of all h-level components in C.

For Case 1, we show that because all components in the hierarchy satisfy Property 2 of Theorem 4.2, there must be an st path in $H \setminus F$. The structure of F relative to the hierarchy is more complex in Case 2. As a result we must prove several key characteristics of the components in the fault subchain. With these characteristics, along with Properties 1, 3 of Theorem 4.2, we show via a highly technical inductive proof that at least one vertex at the end of the fault subchain is reachable from s in $H \setminus F$, and thus t is also reachable from s in $H \setminus F$.

4.3 Algorithm and Analysis

4.3.1 Algorithm

We can now use Theorem 4.2 (Structure Theorem) to give a $2^{O(k^2)}$-approximation algorithm for SD-k-RSND. Given a graph $G = (V, E)$ with edge weights $w : E \to \mathbb{R}_{\geq 0}$ and a single demand $\{(s, t, k)\}$, we first create the hierarchical k-chain decomposition of G in polynomial time, as described in Section 4.1. Within each component we run a set of algorithms to satisfy the RSND demands stated in Theorem 4.2. Our solution, H, includes the outputs of each of these algorithms along with S, the set of all edges in the separators of the hierarchical k-chain decomposition. We now describe the set of algorithms run on each component in the hierarchical k-chain decomposition. Fix a component R_i^h of the decomposition and let $X \subseteq V_{(i,\ell)}^h$, $Y \subseteq V_{(i,r)}^h$, and $d = \max(0, k + |S_X| + |S_Y| - |S_{i-1}| - |S_i|)$:

- **Base Case (Shortest st Path).** For each X, Y pair such that $d = 1$, contract the vertices in X and contract the vertices in Y to create super nodes x and y, respectively. We first check in polynomial time if x and y are connected in $G_i^h = G[R_i^h]$. If they are connected, then we create an instance of the Weighted st Shortest-Path problem on G_i^h (in polynomial time), using x and y as our source and destination nodes. For each edge $e \in E(R_i^h)$, set the weight of e to $w(e)$. Run a polynomial-time Weighted st Shortest-Path algorithm on this instance (e.g. Dijkstra's algorithm), and add to H all edges in the output of the algorithm.

- **Recursive Step.** For each X, Y pair such that $1 < d < k$, we create an instance of SD-d-RSND on G_i^h. Contract the vertices in X and contract the vertices in Y to create super nodes x and y, respectively. For each edge $e \in E(R_i^h)$, set the cost of e to $w(e)$. The set of RSND demands is just $\{(x, y, d)\}$. Run the recursive polynomial-time SD-d-RSND algorithm on this instance, where $d < k$. Add to H all edges in the output of the algorithm.

- **Final Recursive Step.** We now create an SD-$(k-1)$-RSND instance on G_i^h. Contract the vertices in $V_{(i,\ell)}^h$ and contract the vertices in $V_{(i,r)}^h$ to create super nodes v_ℓ and v_r, respectively. For each edge $e \in E(R_i^h)$, set the cost of e to $w(e)$. The set of RSND demands is just $\{(v_\ell, v_r, k-1)\}$. Run the recursive SD-$(k-1)$-RSND algorithm on this instance. Add to H all edges in the output of the algorithm.

- **Min-Cost Flow.** Finally, we create an instance of the Min-Cost Flow problem on G_i^h. Contract the vertices in $V_{(i,\ell)}^h$ and contract the vertices in $V_{(i,r)}^h$ to create super nodes v_ℓ and v_r, respectively. Let v_ℓ be the source node and v_r be the sink node. For each edge $e \in E(R_i^h)$, set the capacity of e to 1 and set the cost of e to $w(e)$. Require a minimum flow of $h + 1$, and run a polynomial-time Min-Cost Flow algorithm on this instance. Since all capacities are integer the algorithm will return an integral flow, so we add to H all edges with non-zero flow.

4.3.2 Analysis

The following lemma is essentially directly from Theorem 4.2 (Structure Theorem) and the description of the algorithm. The proof of this lemma can be found in Appendix C.

Lemma 4.3. *Let H be the output of the algorithm given in Section 4.3.1. Subgraph H is a feasible solution to the SD-k-RSND problem.*

In the next lemma, we give an approximation ratio for the SD-k-RSND algorithm. Let H^* denote the optimal solution, and for any set of edges $A \subseteq E$, let $w(A) = \sum_{e \in A} w(e)$. The next lemma follows from combining the approximation ratios of each of the subroutines used in the algorithm and solving the recurrence.

Lemma 4.4. $w(H) \leq 2^{O(k^2)} * w(H^*)$.

Proof. Fix level h of the hierarchical k-chain decomposition of G. Let $H_i^h = H[R_i^h]$, $G_i^h = G[R_i^h]$, and let $H_i^{h*} = H^*[R_i^h]$ be the subgraph of the optimal solution induced by R_i^h. Let $d = \max(0, k + |S_X| + |S_Y| - |S_{i-1}^h| - |S_i^h|)$. We also let $W_{i,X,Y}^h$ denote the subgraph of H_i^h returned by the Weighted st Shortest-Path algorithm run on G_i^h for contracted vertex subsets $X \subseteq V_{(i,\ell)}^h$ and $Y \subseteq V_{(i,r)}^h$ such that $d = 1$. Additionally, let $D_{i,X,Y}^h$ denote the subgraph of H_i^h returned by the SD-d-RSND algorithm run on G_i^h for contracted $X \subseteq V_{(i,\ell)}^h$ and contracted $Y \subseteq V_{(i,r)}^h$ such that $1 < d < k - 1$. Let $D_{i,\ell,r}^h$ denote the subgraph of H_i^h returned by the SD-$(k-1)$-RSND algorithm run on G_i^h for contracted $V_{(i,\ell)}^h$ and contracted $V_{(i,r)}^h$ with demand $k - 1$. We also let M_i^h denote the subgraph of H_i^h returned by the Min-Cost Flow algorithm run on G_i^h for contracted $V_{(i,\ell)}^h$ and contracted $V_{(i,r)}^h$.

Similarly, let $W_{i,X,Y}^{h*}$ be the optimal solution to the Weighted st Shortest-Path instance on G_i^h for contracted $X \subseteq V_{(i,\ell)}^h$ and contracted $Y \subseteq V_{(i,r)}^h$ such that $d = 1$; let $D_{i,X,Y}^{h*}$ be the optimal solution to the SD-d-RSND instance on G_i^h for contracted $X \subseteq V_{(i,\ell)}^h$ and contracted $Y \subseteq V_{(i,r)}^h$ such that $1 < d < k$; let $D_{i,\ell,r}^{h*}$ be the optimal solution to the SD-$(k-1)$-RSND instance on G_i^h for contracted $V_{(i,\ell)}^h$ and contracted $V_{(i,r)}^h$ with demand $k - 1$; let M_i^{h*} be the optimal solution to the Min-Cost Flow instance on G_i^h.

For each X, Y pair such that $d = 1$, subgraph $W_{i,X,Y}^h$ is given by an exact algorithm. Subgraph M_i^h is also given by an exact algorithm. Let $T(j)$ be the approximation ratio of our Single Demand RSND algorithm with demand (s, t, j). For each X, Y pair such that $1 < d < k$, subgraph $D_{i,X,Y}^h$ is given by our Single Demand RSND algorithm, a $T(d)$-approximation algorithm. Subgraph $D_{i,\ell,r}^h$ is also given by our algorithm, which in this case is a $T(k-1)$-approximation. Hence we have the following for each component H_i^h in the hierarchical k-chain decomposition:

$$w(W_{i,X,Y}^h) = w(W_{i,X,Y}^{h*}) \qquad\qquad \forall X, Y \text{ s.t. } d = 1 \qquad\qquad (1)$$

$$w(D_{i,X,Y}^h) \leq T(d) \cdot w(D_{i,X,Y}^{h*}) \qquad\qquad \forall X, Y \text{ s.t. } 1 < d < k \qquad\qquad (2)$$

$$w(D_{i,\ell,r}^h) \leq T(k-1) \cdot w(D_{i,\ell,r}^{h*}) \qquad\qquad\qquad\qquad\qquad (3)$$

$$w(M_i^h) = w(M_i^{h*}). \qquad\qquad\qquad\qquad\qquad (4)$$

Summing over all subsets of $V_{(i,\ell)}^h$ and $V_{(i,r)}^h$, we get the following from expressions (1) and (2):

$$\sum_{(X,Y):d=1} w(W_{i,X,Y}^h) = \sum_{(X,Y):d=1} w(W_{i,X,Y}^{h*})$$

$$\sum_{(X,Y):1<d<k} w(D_{i,X,Y}^h) \leq \sum_{(X,Y):1<d<k} T(d) \cdot w(D_{i,X,Y}^{h*}).$$

13

For each level h of the hierarchical k-chain decomposition, we denote the set of h-components as $R_0^h, R_1^h, \ldots R_{p_h}^h$. By summing over all components in the hierarchical k-chain decomposition, we get the following:

$$\sum_{h=2}^{k-1}\sum_{i=0}^{p_h}\sum_{(X,Y):d=1} w(W_{i,X,Y}^h) = \sum_{h=2}^{k-1}\sum_{i=0}^{p_h}\sum_{(X,Y):d=1} w(W_{i,X,Y}^{h*}) \tag{5}$$

$$\sum_{h=2}^{k-1}\sum_{i=0}^{p_h}\sum_{(X,Y):1<d<k} w(D_{i,X,Y}^h) \leq \sum_{h=2}^{k-1}\sum_{i=0}^{p_h}\sum_{(X,Y):1<d<k} T(d) \cdot w(D_{i,X,Y}^{h*}) \tag{6}$$

$$\sum_{h=2}^{k-1}\sum_{i=0}^{p_h} w(D_{i,\ell,r}^h) \leq T(k-1) \cdot \sum_{h=2}^{k-1}\sum_{i=0}^{p_h} w(D_{i,\ell,r}^{h*}) \tag{7}$$

$$\sum_{h=2}^{k-1}\sum_{i=0}^{p_h} w(M_i^h) = \sum_{h=2}^{k-1}\sum_{i=0}^{p_h} w(M_i^{h*}). \tag{8}$$

We also have that

$$w(H_i^h) \leq \sum_{(X,Y):d=1} w(W_{i,X,Y}^h) + \sum_{(X,Y):1<d<k} w(D_{i,X,Y}^h) + w(D_{i,\ell,r}^h) + w(M_i^h). \tag{9}$$

Summing inequality (9) over all components in the hierarchical k-chain decomposition, then substituting in expressions (5) through (8) gives the following:

$$\sum_{h=2}^{k-1}\sum_{i=0}^{p_h} w(H_i^h) \leq \sum_{h=2}^{k-1}\sum_{i=0}^{p_h}\sum_{(X,Y):d=1} w(W_{i,X,Y}^h) + \sum_{h=2}^{k-1}\sum_{i=0}^{p_h}\sum_{(X,Y):1<d<k} w(D_{i,X,Y}^h)$$

$$+ \sum_{h=2}^{k-1}\sum_{i=0}^{p_h} w(D_{i,\ell,r}^h) + \sum_{h=2}^{k-1}\sum_{i=0}^{p_h} w(M_i^h) \tag{10}$$

$$\leq \sum_{h=2}^{k-1}\sum_{i=0}^{p_h}\sum_{(X,Y):d=1} w(W_{i,X,Y}^{h*}) + \sum_{h=2}^{k-1}\sum_{i=0}^{p_h}\sum_{(X,Y):1<d<k} T(d) \cdot w(D_{i,X,Y}^{h*})$$

$$+ T(k-1) \cdot \sum_{h=2}^{k-1}\sum_{i=0}^{p_h} w(D_{i,\ell,r}^{h*}) + \sum_{h=2}^{k-1}\sum_{i=0}^{p_h} w(M_i^{h*}). \tag{11}$$

The optimal subgraph H^* is a feasible solution, so by Theorem 4.2, each RSND demand in the theorem statement must be satisfied in subgraph H_i^{h*}, for all h, i. For all RSND demands in the theorem statement to be satisfied in a subgraph H_i^{h*}, the set of edges $E(H_i^{h*})$ must be a feasible solution to each of the Weighted st Shortest Path, SD-k-RSND, and Min-Cost Flow instances we created on G_i^h. Therefore, the cost of H_i^{h*} must be at least the cost of the optimal solution to each of the instances. We have that

$$w(W_{i,X,Y}^{h*}) \leq w(H_i^{h*}) \qquad\qquad \forall X, Y \text{ s.t. } d = 1$$

$$w(D_{i,X,Y}^{h*}) \leq w(H_i^{h*}) \qquad\qquad \forall X, Y \text{ s.t. } 1 < d < k$$

$$w(D_{i,\ell,r}^{h*}) \leq w(H_i^{h*})$$

$$w(M_i^{h*}) \leq w(H_i^{h*}).$$

14

Let $C(j)$ be the number X, Y pairs such that $X \subseteq V_{(i,\ell)}$, $Y \subseteq V_{(i,r)}$, $(X,Y) \neq (V_{(i,\ell)}, V_{(i,r)})$, and $j = \max(0, k + |S_X| + |S_Y| - |S_{i-1}| - |S_i|)$. Summing over subsets of $V_{(i,\ell)}$ and $V_{(i,r)}$ and over all components in the hierarchical k-chain decomposition, we have the following:

$$\sum_{h=2}^{k-1} \sum_{i=0}^{p_h} \sum_{(X,Y):d=1} w(W_{i,X,Y}^{h*}) \leq \sum_{h=2}^{k-1} \sum_{i=0}^{p_h} C(1) \cdot w(H_i^{h*}) \tag{12}$$

$$\sum_{h=2}^{k-1} \sum_{i=0}^{p_h} \sum_{(X,Y):1<d<k} T(d) \cdot w(D_{i,X,Y}^{h*}) \leq \sum_{h=2}^{k-1} \sum_{i=0}^{p_h} \sum_{(X,Y):1<d<k} T(d) \cdot w(H_i^{h*})$$

$$\leq \sum_{h=2}^{k-1} \sum_{i=0}^{p_h} \sum_{j=2}^{k-1} C(j) \cdot T(j) \cdot w(H_i^{h*}) \tag{13}$$

$$\sum_{h=2}^{k-1} \sum_{i=0}^{p_h} w(D_{i,\ell,r}^{h*}) \leq \sum_{h=2}^{k-1} \sum_{i=0}^{p_h} w(H_i^{h*}) \tag{14}$$

$$\sum_{h=2}^{k-1} \sum_{i=0}^{p_h} w(M_i^{h*}) \leq \sum_{h=2}^{k-1} \sum_{i=0}^{p_h} w(H_i^{h*}). \tag{15}$$

Plugging inequalities (12) through (15) into inequality (11) gives the following:

$$\sum_{h=2}^{k-1} \sum_{i=0}^{p_h} w(H_i^h) \leq C(1) \sum_{h=2}^{k-1} \sum_{i=0}^{p_h} w(H_i^{h*}) + \sum_{j=2}^{k-1} C(j)T(j) \cdot \sum_{h=2}^{k-1} \sum_{i=0}^{p_h} w(H_i^{h*}) + T(k-1) \sum_{h=2}^{k-1} \sum_{i=0}^{p_h} w(H_i^{h*})$$

$$+ \sum_{h=2}^{k-1} \sum_{i=0}^{p_h} w(H_i^{h*})$$

$$= \left[C(1) + \sum_{j=2}^{k-1} C(j)T(j) + T(k-1) + 1 \right] \sum_{h=2}^{k-1} \sum_{i=0}^{p_h} w(H_i^{h*})$$

Now we account for the important separators, or edges between components, in the hierarchical k-chain decomposition. Let S be the separators of the decomposition and let S^* be the set of edges in these separators included in the optimal solution. Edge set S is included in H. By Theorem 4.2, any feasible solution must include all edges between components in the decomposition. We therefore have that $S = S^*$ and get the following:

$$w(H) = \sum_{h=2}^{k-1} \sum_{i=0}^{p_h} w(H_i^h) + w(S)$$

$$\leq \left[C(1) + \sum_{j=2}^{k-1} C(j)T(j) + T(k-1) + 1 \right] \sum_{h=2}^{k-1} \sum_{i=0}^{p_h} w(H_i^{h*}) + w(S^*)$$

$$= \left[C(1) + \sum_{j=2}^{k-1} C(j)T(j) + T(k-1) + 1 \right] w(H^*).$$

Therefore, a bound on the approximation ratio of the SD-k-RSND algorithm for demand k is given

by the following recurrence relation:

$$T(k) \leq C(1) + \sum_{j=2}^{k-1} C(j)T(j) + T(k-1) + 1$$

$$= \sum_{j=1}^{k-1} C(j)T(j) + T(k-1) + 1$$

$$\leq \sum_{j=1}^{k-1} C(j)T(k-1) + T(k-1) + 1$$

with base case $T(1) = 1$ (recall that the SD-k-RSND problem with $k = 1$ is solved exactly using a Weighted st Shortest Paths algorithm).

First, note that because there are at most $2^{2(k-1)}$ pairs of vertex sets X, Y such that $X \subseteq V_{(i,\ell)}^h$ and $Y \subseteq V_{(i,r)}^h$, we have that $\sum_{j=1}^{k-1} C(j) \leq 2^{2(k-1)}$. This gives the following:

$$T(k) \leq T(k-1) \sum_{j=1}^{k-1} C(j) + T(k-1) + 1$$

$$\leq T(k-1) \cdot 2^{2(k-1)} + T(k-1) + 1$$

$$= (2^{2(k-1)} + 1)T(k-1) + 1.$$

By solving the recurrence relation, we get the following

$$T(k) = O\left(\prod_{x=1}^{k-1} 2^{2(k-x)}\right)$$

$$= O\left(2^{\sum_{j=1}^{k} 2j}\right)$$

$$= O\left(2^{k(k+1)}\right)$$

$$= 2^{O(k^2)}.$$

Therefore, we have that $w(H) \leq 2^{O(k^2)} \cdot w(H^*)$. □

Theorem 1.2 is directly implied by Lemmas 4.3 and 4.4 together with the observation that the algorithm runs in polynomial time.

References

[1] David Adjiashvili, Felix Hommelsheim, and Moritz Mühlenthaler. Flexible graph connectivity: Approximating network design problems between 1- and 2-connectivity, 2020. `arXiv:1910.13297`.

[2] David Adjiashvili, Felix Hommelsheim, Moritz Mühlenthaler, and Oliver Schaudt. Fault-tolerant edge-disjoint paths – beyond uniform faults, 2020. `arXiv:2009.05382`.

[3] Ishan Bansal, Joseph Cheriyan, Logan Grout, and Sharat Ibrahimpur. Improved approximation algorithms by generalizing the primal-dual method beyond uncrossable functions, 2022. `arXiv:2209.11209`.

[4] Davide Bilò, Luciano Gualà, Stefano Leucci, and Guido Proietti. Multiple-edge-fault-tolerant approximate shortest-path trees, 2016. URL: `https://arxiv.org/abs/1601.04169`, `doi:10.48550/ARXIV.1601.04169`.

[5] Greg Bodwin, Michael Dinitz, and Yasamin Nazari. Vertex fault-tolerant emulators. In Mark Braverman, editor, *13th Innovations in Theoretical Computer Science Conference, ITCS 2022*, volume 215 of *LIPIcs*, pages 25:1–25:22. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2022. `doi:10.4230/LIPIcs.ITCS.2022.25`.

[6] Greg Bodwin, Michael Dinitz, and Yasamin Nazari. Epic Fail: Emulators can tolerate polynomially many edge faults for free. In *14th Innovations in Theoretical Computer Science Conference, ITCS 2023*, 2023.

[7] Greg Bodwin, Michael Dinitz, Merav Parter, and Virginia Vassilevska Williams. Optimal vertex fault tolerant spanners (for fixed stretch). In Artur Czumaj, editor, *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2018, New Orleans, LA, USA, January 7-10, 2018*, pages 1884–1900. SIAM, 2018.

[8] Greg Bodwin, Michael Dinitz, and Caleb Robelle. Optimal vertex fault-tolerant spanners in polynomial time. In *Proceedings of the Thirty-Second Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2021*, 2021.

[9] Greg Bodwin, Michael Dinitz, and Caleb Robelle. Optimal vertex fault-tolerant spanners in polynomial time. In Joseph (Seffi) Naor and Niv Buchbinder, editors, *Proceedings of the 2022 ACM-SIAM Symposium on Discrete Algorithms, SODA 2022*, pages 2924–2938. SIAM, 2022. `doi:10.1137/1.9781611976465.174`.

[10] Greg Bodwin and Shyamal Patel. A trivial yet optimal solution to vertex fault tolerant spanners. In *Proceedings of the 2019 ACM Symposium on Principles of Distributed Computing*, PODC '19, page 541–543, New York, NY, USA, 2019. Association for Computing Machinery. `doi:10.1145/3293611.3331588`.

[11] Sylvia Boyd, Joseph Cheriyan, Arash Haddadan, and Sharat Ibrahimpur. Approximation algorithms for flexible graph connectivity, 2022. `arXiv:2202.13298`.

[12] Shiri Chechik, Michael Langberg, David Peleg, and Liam Roditty. Fault tolerant spanners for general graphs. *SIAM J. Comput.*, 39(7):3403–3423, 2010.

[13] Chandra Chekuri and Rhea Jain. Approximating flexible graph connectivity via räcke tree based rounding, 2022. `arXiv:2211.08324`.

[14] Chandra Chekuri and Rhea Jain. Augmentation based approximation algorithms for flexible network design, 2022. `arXiv:2209.12273`.

[15] J. Cheriyan, B. Laekhanukit, G. Naves, and A. Vetta. Approximating rooted steiner networks. *ACM Trans. Algorithms*, 11(2):8:1–8:22, 2014.

[16] Joseph Cheriyan and Ramakrishna Thurimella. Approximating minimum-size k-connected spanning subgraphs via matching. *SIAM Journal on Computing*, 30(2):528–560, 2000. `doi: 10.1137/S009753979833920X`.

[17] M. Dinitz, A. Koranteng, and G. Kortsarz. Relative survivable network design. In *APPROX-RANDOM*, volume 245, pages 41:1–41:19, 2022.

[18] Michael Dinitz and Robert Krauthgamer. Fault-tolerant spanners: better and simpler. In *Proceedings of the 30th Annual ACM Symposium on Principles of Distributed Computing, PODC 2011, San Jose, CA, USA, June 6-8, 2011*, pages 169–178, 2011.

[19] Michael Dinitz and Caleb Robelle. Efficient and simple algorithms for fault-tolerant spanners. In Yuval Emek and Christian Cachin, editors, *PODC '20: ACM Symposium on Principles of Distributed Computing*, pages 493–500. ACM, 2020. `doi:10.1145/3382734.3405735`.

[20] Ye Dinitz and Jeffery Westbrook. Maintaining the classes of 4-edge-connectivity in a graph on-line. *Algorithmica*, 20:242–276, 1998.

[21] Yefim Dinitz and Zeev Nutov. A 2-level cactus model for the system of minimum and minimum+ 1 edge-cuts in a graph and its incremental maintenance. In *Proceedings of the twenty-seventh annual ACM symposium on Theory of computing*, pages 509–518, 1995.

[22] A. E. Feldmann, A. Mukherjee, and E. J. van Leeuwen. The parameterized complexity of the survivable network design problem. In *SOSA*, pages 37–56, 2022.

[23] Harold N. Gabow and Suzanne R. Gallagher. Iterated rounding algorithms for the smallest k-edge connected spanning subgraph. *SIAM Journal on Computing*, 41(1):61–103, 2012. `doi: 10.1137/080732572`.

[24] Harold N. Gabow, Michel X. Goemans, Éva Tardos, and David P. Williamson. Approximating the smallest k-edge connected spanning subgraph by lp-rounding. *Networks*, 53(4):345–357, 2009.

[25] Monika Rauch Henzinger. A static 2-approximation algorithm for vertex connectivity and incremental approximation algorithms for edge and vertex connectivity. *Journal of Algorithms*, 24(1):194–220, 1997.

[26] Kamal Jain. A factor 2 approximation algorithm for the generalized steiner network problem. *Combinatorica*, 21(1):39–60, 2001. `doi:10.1007/s004930170004`.

[27] Rohit Khandekar, Guy Kortsarz, and Zeev Nutov. Approximating fault-tolerant group-steiner problems. *Theoretical Computer Science*, 416:55–64, 2012.

[28] On-Hei Solomon Lo, Jens M. Schmidt, and Mikkel Thorup. Compact cactus representations of all non-trivial min-cuts. *Discret. Appl. Math.*, 303:296–304, 2021.

[29] Dániel Marx. Important separators and parameterized algorithms. In *International Workshop on Graph-Theoretic Concepts in Computer Science*, pages 5–10. Springer, 2011.

[30] Johannes A. La Poutre. Maintenance of 2- and 3-edge-connected components of graphs ii. *SIAM J. Comput.*, 29(5):1521–1549, 2000.

[31] 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`.

A Proofs from Section 3 (3-RSND)

First we must show that any RSND demand can be expressed as a set of cut-covering constraints.

Lemma A.1. *An RSND demand $r(s,t) = k$ is satisfied by a subgraph H if and only if for any G-minimal st-set A, $d_H(A) \geq \min\{k, d_G(A)\}$*

Proof. First we prove the "only if" direction. Suppose that demand $r(s,t) = k$ is satisfied by subgraph H, but suppose for the sake of contradiction that there is some G-minimal st-set, A, such that $d_H(A) < \min\{k, d_G(A)\}$. Then, $F = \delta_H(A)$ is a fault set of size at most $k - 1$ such that there is an st-path in $G \setminus F$ (since $G[A]$ and $G[V \setminus A]$ are connected and $|\delta_G(A) \setminus F| \geq 1$), but no st-path in $H \setminus F$. This contradicts the feasibility of H.

Now we show the "if" direction. Suppose that for any G-minimal st-set A, $d_H(A) \geq \min\{k, d_G(A)\}$. Suppose however for the sake of contradiction that H does not satisfy requirement $r(s,t) = k$. Then there is an H-minimal st-cut $X \subseteq E$ with $|X| < k$ such that s, t are not connected in $H \setminus X$ but are connected in $G \setminus X$. Edge set X is H-minimal, so we have that $X = \delta_H(A)$ for some vertex set A. Note that $\delta_G(A)$ is a G-minimal st-set with $d_G(A) > d_H(A)$. By our assumption, we have that $d_H(A) \geq \min\{k, d_G(A)\}$. This is a contradiction, since we have $d_H(A) < k$ and $d_H(A) < d_G(A)$. $\square$

Proof of Lemma 3.3. First suppose that H is a feasible SD-3-RSND solution. We will show that both (i) and (ii) are satisfied. Suppose to the contrary that (i) does not hold for some non-central component J. Let $H' = H[J]$. Let u and v be the attachment nodes of J, and assume without loss of generality that in G, any (simple) st-path has u before v, as in Fig, 1(a). Then, there exists a J-minimal uv-cut, $\delta_{H'}(A_J)$, such that $d_{H'}(A_J) \leq 1$ (recall that H' does not include any of the attachment edges of its associated 3-class). Let A be the union of A_J and all components that lie on some (simple) us-path of the cycle chain that do not contain v. It can be seen that $d_H(A) = d_{H'}(A_J) + 1 \leq 2$. Since u and v are 3-connected in G (they are in the same 3-class) and A is a uv-set, $d_G(A) \geq 3$. It is also not hard to see that $G[A]$ and $G[V \setminus A]$ are both connected, and thus $\delta_G(A)$ is a G-minimal st-cut. This all means that $F = \delta_H(A)$ is a fault set of size at most 2 such that there is an st path in $G \setminus F$, but no st path in $H \setminus F$. This contradicts the feasibility of H.

Now suppose that (ii) does not hold for some central component J. Then, there must be a J-minimal st-cut, $\delta_{H'}(A_J)$, such that $d_{H'}(A_J) < \min\{3, d_J(A_J)\}$ (by Lemma A.1). Let A be the union of A_J and all components of cycles that precede J in the cactus chain. Note that $d_H(A) = d_{H'}(A_J)$ and $d_G(A) = d_J(A_J)$. This gives $d_H(A) < \min\{3, d_G(A)\}$, contradicting the feasibility of H.

Now suppose that H satisfies (i,ii). We will show that H is a feasible SD-3-RSND solution. Let $F \subseteq E$ such that $G \setminus F$ has an st-path and $H \setminus F$ has no st-path. One can see that since F cannot contain two edges of the same cactus chain cycle, we have that F must have size at least 3; thus, if there is an st path in $G \setminus F$ for $|F| \leq 2$, then there is an st path in $H \setminus F$ as well. $\square$

Lemma A.2. *Subgraph H is a feasible solution for instance $\mathcal{I}$ if and only if for any st-set A such that $d_H(A) = 2$, $\delta_H(A) = \delta_J(s)$ or $\delta_H(A) = \delta_J(t)$.*
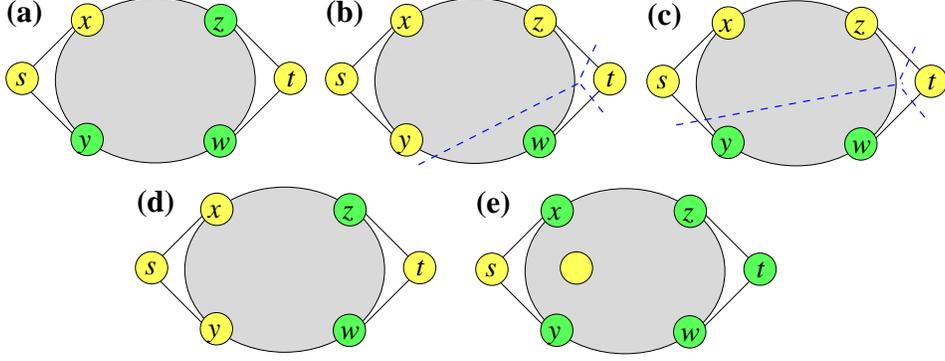
Figure 3: Illustration to the proof of Lemmas 3.4, 3.5. Nodes in A are yellow and in $V \setminus A$ are green.

Proof. Let H be a feasible solution for $\mathcal{I}$ and let A be an st-set with $d_H(A) = 2$. We will show that either $\delta_H(A) = \delta_J(s)$ or $\delta_H(A) = \delta_J(t)$. Let $F = \delta_H(A)$. Then, $H \setminus F$ has no st-path. Since H is a feasible solution for $\mathcal{I}$, $J \setminus F$ must also have no st-path. Component J is 2-connected, so $\delta_H(A) = \delta_J(A)$. By property (B) of Theorem 3.2, we then have that $\delta_J(A) = \delta_J(s)$ or $\delta_J(A) = \delta_J(t)$.

Now suppose that for any st-set A with $d_H(A) = 2$, we have that $\delta_H(A) = \delta_J(s)$ or $\delta_H(A) = \delta_J(t)$. We will show that H is a feasible solution for $\mathcal{I}$. For any st-set A such that $\delta_H(A) \neq \delta_J(s)$ and $\delta_H(A) \neq \delta_J(t)$, we have that $d_H(A) \geq 3$. Since H includes the forced edges sx, sy, zt, and wt, we also know that $\delta_H(s) = \delta_J(s)$ and $\delta_H(t) = \delta_J(t)$. Thus, we have that $d_H(A) \geq \min\{3, d_J(A)\}$ for all A such that A is an st-set. By Lemma A.1, H is a feasible solution for $\mathcal{I}$. $\square$

Proof of Lemma 3.4. Suppose that H is a feasible solution for $\mathcal{I}$, but suppose to the contrary that R is not a 3-connected subset in H. Then there is a vertex set A such that $d_H(A) = 2$ and $R \cap A, R \setminus A \neq \emptyset$. By Lemma A.2, A cannot be an st-set. Hence, without loss of generality, assume that $s, t \in A$. Now we consider several cases, see Fig. 3(a,b,c,d).

(a) $|R \cap A| = 1$, say $R \cap A = \{x\}$; see Fig. 3(a). Then $\{sy, tz, tw\} \subseteq \delta_H(A)$, contradicting that $d_H(A) = 2$.

(b) $|R \cap A| = 3$, say $R \setminus A = \{w\}$; see Fig. 3(b). Vertex set $A \setminus \{t\}$ is an st-set with $d_H(A \setminus \{t\}) = d_H(A) = 2$, contradicting (by Lemma A.2) that H is a feasible solution.

(c) $|R \cap A| = 2$ and $|R \cap \{x, y\}| = |R \cap \{z, w\}| = 1$, say $R \cap A = \{x, z\}$; see Fig. 3(c). Then $d_H(A \setminus \{t\}) = d_H(A) = 2$, contradicting (by Lemma A.2) that J is a feasible solution.

(d) $R \cap A = \{x, y\}$ or $R \cap A = \{z, w\}$, say $R \cap A = \{x, y\}$; see Fig. 3(d). Then $A \setminus \{t\}$ is an st-set with $d_H(A \setminus \{t\}) = d_H(A) - 2 = 0$, contradicting that H is a feasible solution.

In all cases we have a contradiction, thus the lemma holds. $\square$

Proof of Lemma 3.5. Suppose that R is a 3-connected subset in H, but suppose to the contrary that H is not a feasible solution. Then by Lemma A.2 there is an st-set A with $d_H(A) = 2$ such that $\delta_H(A) \neq \delta_J(s)$ and $\delta_H(A) \neq \delta_J(t)$. Since R is a 3-edge-connected subset in H, we must have $R \cap A = \emptyset$ or $R \subset A$. Without loss of generality, assume that $R \cap A = \emptyset$ (see Fig. 3(e)); otherwise we interchange the roles of s and t. But then, $sx, sy \in \delta_H(A)$, so we get the contradiction that $\delta_H(A) = \{sx, sy\} = \delta_J(s)$. $\square$

20

Proof of Theorem 1.1. We reduce the 3-RSND instance to an instance of SND. Let $D = \{d_1, d_2, \ldots, d_\ell\}$ be the set of all ℓ vertex demands in the input 3-RSND instance (where each $d_i \in D$ corresponds to a demand of the form (x, y, r), where x, y are vertices in G and $r \leq 3$). We will create a new set of ordinary demands, D', that characterizes the problem.

First note that G is 2-connected, so all 1- and 2-demands in the input are ordinary demands. We add them to D'. We also add all ordinary 3-demands from the input to D'. By Theorem 3.1, each *relative* 3-demand $d_i \in D$ is characterized by a set of ordinary 3-demands, D_i. Thus, for each relative 3-demand $d_i \in D$, we add all demands in D_i to D'. Demand set D' characterizes the 3-RSND instance, meaning that any optimal solution to the 3-RSND instance must also satisfy D'. All demands in D' are ordinary so we can run Jain's 2-approximation SND algorithm ([26]) on G under demands D'. By Theorem 3.1, we only add a polynomial number of demands to D', so the reduction (and the algorithm) runs in polynomial time. This gives us a 2-approximation for 3-RSND. $\qquad\square$

B Proof of Theorem 4.2 (Structure Theorem)

B.1 Connectivity Lemmas for Input Graph G

To prove Theorem 4.2, we will first need lemmas that describe the connectivity in G between s and any h-component, and between any h-component and t. We will say that the *level h chain* is the ordered sequence of all h-components, labelled $R_0^h, \ldots, R_p^h$ (where $s \in R_0^h$ and $t \in R_p^h$). Note that although all components in the level h chain are h-components, each h-component may have a left or right separator with size less than h. Let $E(R_i^h)$ denote the set of edges with both endpoints in component R_i^h. We define the subchain of a level h chain that starts at G_i^h and ends at G_j^h (where $i \leq j$) to be the subgraph of G induced by $\cup_{k=i}^j R_k^h$. An edge e is in a subchain that starts at G_i^h and ends at G_j^h if $e \in \cup_{k=i}^j E(R_k^h)$ or $e \in \cup_{k=i}^{j-1} S_k^h$. We define a subchain that starts at H_i^h and ends at H_j^h similarly. We may also use vertex sets $V_{(j,r)}^h$ to define the start and end of a subchain.

Lemma B.1. *Fix level h of the hierarchical k-chain decomposition, and consider the subchain that starts at s and ends at an h-component. Let R_i^h be this h-component. Then, in G, there is a path from s to each vertex in $V_{(i,r)}^h$; these paths only use edges within the level h subchain that begins at G_0^h and ends at G_i^h.*

Proof. Graph G is 2-edge connected, so the lemma statement is true when R_i^h contains t (that is, when $V_{(i,r)}^h = \{t\}$). Thus we will assume $t \notin R_i^h$. First we will show that for each h-component R_j^h such that $t \notin R_j^h$, there is a path in G from $V_{(j,\ell)}^h$ to each vertex in $V_{(j,r)}^h$. Fix such an h-component R_j^h, and suppose that component R_m^{h-1} is the parent of R_j^h in the hierarchical k-chain decomposition. Suppose for the sake of contradiction that there is no path in G from $V_{(j,\ell)}^h$ to a vertex $v \in V_{(j,r)}^h$. Let S_v be the set of edges in S_j^h that are incident on v. Then, $S_j^h \setminus S_v$ is a $(V_{(j,\ell)}^h, V_{(m,r)}^{h-1})$-separator with size strictly less than $|S_j^h|$. This means that S_j^h is not minimal and therefore not an important separator, giving a contradiction.

We have shown that in each h-component R_j^h such that $t \notin R_j^h$, there is a path in G from $V_{(j,\ell)}^h$ to each vertex in $V_{(j,r)}^h$. Therefore, for every component R_i^h, there is a path in G from s to each vertex in $V_{(i,r)}^h$, using only edges within the subchain. This can be seen via a proof by induction on the number of components into the level h subchain that begins at R_0^h and ends at R_i^h. $\qquad\square$

We now give a similar lemma describing the connectivity between any h-component and t.

Lemma B.2. *Fix level h of the hierarchical k-chain decomposition, and consider the subchain that starts at an h-component and ends at t. Let R_i^h be this h-component. Then, in G, there is a path from each vertex in $V_{(i,\ell)}^h$ to t; these paths only use edges within the level h subchain that begins at G_i^h and ends at G_p^h.*

Proof. Graph G is 2-edge connected, so the lemma statement is true when R_i^h contains s (that is, when $V_{(i,\ell)}^h = \{s\}$). Thus we will assume $s \notin R_i^h$. First we will show that for each h-component R_j^h such that $s \notin R_j^h$, there is a path in G from each vertex in $V_{(j,\ell)}^h$ to $V_{(j,r)}^h$. Fix such an h-component R_j^h, and suppose that component R_m^{h-1} is the parent of R_j^h in the hierarchical k-chain decomposition. Suppose for the sake of contradiction that there is no path in G from a vertex $v \in V_{(j,\ell)}^h$ to $V_{(j,r)}^h$. Let S_v be the set of edges in S_{j-1}^h that are incident on v. Then, $S_{j-1}^h \setminus S_v$ is a $(V_{(j-1,\ell)}^h, V_{(m,r)}^{h-1})$-separator with size strictly less than $|S_{j-1}^h|$. This means that S_{j-1}^h is not minimal and therefore not an important separator, giving a contradiction.

We have shown that in each h-component R_j^h such that $s \notin R_j^h$, there is a path in G from each vertex in $V_{(j,\ell)}^h$ to $V_{(j,r)}^h$. Therefore, for every component R_i^h, there is a path in G from $V_{(i,\ell)}^h$ to t, using only edges within the subchain. This can be seen via a proof by induction on the number of components into the level h chain that begins at R_i^h and ends at R_p^h. $\qquad\square$

B.2 Only If

We are now ready to prove that the properties stated in Theorem 4.2 are necessary in any feasible solution. Let Property 1 denote the first demand set given in Theorem 4.2—that is, the set $\left\{(X, Y, d) : X \subseteq V_{(i,\ell)}^h, Y \subseteq V_{(i,r)}^h, (X, Y) \neq (V_{(i,\ell)}^h, V_{(i,r)}^h), d = \max(0, k + |S_X| + |S_Y| - |S_{i-1}^h| - |S_i^h|)\right\}$. Let Property 2 denote the demand $(V_{(i,\ell)}^h, V_{(i,r)}^h, h+1)$ and Property 3 denote the demand $(V_{(i,\ell)}^h, V_{(i,r)}^h, k-1)$ in the union.

Important Separators. Suppose that subgraph H is a feasible solution, but suppose for the sake of contradiction that at level h in the hierarchical k-chain decomposition, there is an important separator S_i^h such that an edge $e \in S_i^h$ is not in H. Then, the edge set $E' = S_i^h \setminus \{e\}$, which has size at most $k-2$, separates s and t in H but not in G. By Lemmas B.1 and B.2, in $G \setminus E'$, there is a path from s to each vertex in $V_{(i,r)}^h$ and a path from each vertex in $V_{(i+1,\ell)}^h$ to t. The edge e is incident on a vertex in $V_{(i,r)}^h$ and a vertex in $V_{(i+1,\ell)}^h$. Therefore, putting everything together, there is a path from s to t in $G \setminus E'$. In H however, there is no path from $V_{(i,r)}^h$ to $V_{(i+1,\ell)}^h$, and therefore there is no path from s to t in $H \setminus E'$. This contradicts the assumption that H is a feasible solution.

Property 1. Now suppose H is feasible, but suppose for the sake of contradiction that at least one of the RSND demands in Property 1 is not satisfied in subgraph H_i^h. Specifically, fix subgraph H_i^h, let X and Y be vertex sets such that $X \subseteq V_{(i,\ell)}^h$, $Y \subseteq V_{(i,r)}^h$, and $(X, Y) \neq (V_{(i,\ell)}^h, V_{(i,r)}^h)$, and let $d = \max(0, k + |S_X| + |S_Y| - |S_{i-1}^h| - |S_i^h|)$. Suppose that the RSND demand (X, Y, d) is not satisfied in H_i^h. That is, there exists a fault set F_i with $|F_i| < d$ such that there is a path from X to Y in $G_i^h \setminus F_i$ but there is no path from X to Y in $H_i^h \setminus F_i$. Since we are assuming that the RSND demand (X, Y, d) is not satisfied in H_i^h, we can assume that $d > 0$ (since a demand of 0 is always satisfied). Therefore we assume $d = k + |S_X| + |S_Y| - |S_{i-1}^h| - |S_i^h| > 0$.

We now show that there exists a fault set F with $|F| < k$ such that s and t are connected in $G \setminus F$ but not in $H \setminus F$. Let $F = (S_{i-1}^h \setminus S_X) \cup (S_i^h \setminus S_Y) \cup F_i$. By Lemma B.1, and by the fact that all edges in S_X are still in $G \setminus F$, there is a path in $G \setminus F$ from s to each vertex in X, using

only edges before G_i^h in the level h chain. There is a path from X to Y in $G_i^h \setminus F_i$ (and in $G_i^h \setminus F$). By Lemma B.2, and by the fact that all edges in S_Y are still in $G \setminus F$, we also have that there is a path in $G \setminus F$ from each vertex in Y to t that only uses edges after G_i^h in the level h chain. Putting it all together, there is a path from s to t in $G \setminus F$. Note that because X and Y are not connected in $H \setminus F$, there is no st path in $H \setminus F$. Now we give an upper bound on the size of F:

$$
\begin{aligned}
|F| &\leq (|S_{i-1}^h| - |S_X|) + (|S_i^h| - |S_Y|) + (d-1) \\
&= |S_{i-1}^h| - |S_X| + |S_i^h| - |S_Y| + k + |S_X| + |S_Y| - |S_{i-1}^h| - |S_i^h| - 1 \\
&= k - 1.
\end{aligned}
$$

We have shown that $|F| \leq k - 1$. Thus, our construction of F contradicts the assumption that H is feasible, so all RSND demands in the set $\{(X, Y, d) : X \subseteq V_{(i,\ell)}^h, Y \subseteq V_{(i,r)}^h, (X, Y) \neq (V_{(i,\ell)}^h, V_{(i,r)}^h), d = \max(0, k + |S_X| + |S_Y| - |S_{i-1}^h| - |S_i^h|)\}$ on G_i^h are satisfied in H_i^h in all feasible solutions.

Property 2. Now, suppose H is feasible, but suppose for the sake of contradiction that the RSND demand in Property 2, $(V_{(i,\ell)}^h, V_{(i,r)}^h, h+1)$, is not satisfied in a fixed subgraph H_i^h. That is, there exists a fault set F_i with $|F_i| < h + 1 \leq k$ such that there is a path from $V_{(i,\ell)}^h$ to $V_{(i,r)}^h$ in $G_i^h \setminus F_i$ but there is no path from $V_{(i,\ell)}^h$ to $V_{(i,r)}^h$ in $H_i^h \setminus F_i$. By Lemma B.1, there is a path from s to each vertex in $V_{(i,\ell)}^h$ in $G \setminus F_i$, using only edges before G_i^h in the level h chain. There is a path from at least one vertex in $V_{(i,\ell)}^h$ to a vertex in $V_{(i,r)}^h$ in $G_i^h \setminus F_i$. By Lemma B.2, there is a path from each vertex in $V_{(i,r)}^h$ to t in $G \setminus F_i$, using only edges after G_i^h in the level h chain. Putting it all together, there is a path from s to t in $G \setminus F$. However, there is no path from $V_{(i,\ell)}^h$ to $V_{(i,r)}^h$ in $H_i^h \setminus F_i$, so there is no st path in $H \setminus F_i$. This contradicts the assumption that H is feasible, so the RSND demand $(V_{(i,\ell)}^h, V_{(i,r)}^h, h+1)$ on G_i^h must be satisfied in H_i^h in all feasible solutions.

Property 3. Finally, suppose H is feasible, but suppose for the sake of contradiction that the RSND demand in Property 3, $(V_{(i,\ell)}^h, V_{(i,r)}^h, k-1)$, is not satisfied in a fixed subgraph H_i^h. Then there exists a fault set F_i with $|F_i| < k - 1$ such that there is a path from $V_{(i,\ell)}^h$ to $V_{(i,r)}^h$ in $G_i^h \setminus F_i$ but there is no path from $V_{(i,\ell)}^h$ to $V_{(i,r)}^h$ in $H_i^h \setminus F_i$. By Lemmas B.1 and B.2, there is a path from s to t in $G \setminus F_i$. However, there is no path from $V_{(i,\ell)}^h$ to $V_{(i,r)}^h$ in $H_i^h \setminus F_i$, so there is no st path in $H \setminus F_i$. This contradicts the assumption that H is feasible, so the RSND demand $(V_{(i,\ell)}^h, V_{(i,r)}^h, k-1)$ on G_i^h must be satisfied in H_i^h in all feasible solutions.

B.3 If

We now prove that the properties stated in Theorem 4.2 are sufficient. In this section, we will let H be a subgraph of G such that all edges in S are in H and all RSND demands in the statement of Theorem 4.2 are satisfied forall subgraphs H_i^h. For all possible fault sets $F \subseteq E$ such that $|F| < k$, we will show that if s and t are connected in $G \setminus F$, they must also be connected in $H \setminus F$. Going forward, fix $F \subseteq E$ to be a fault set such that $|F| < k$.

B.3.1 Connectivity Lemmas for Subgraph H

We first give the equivalent of Lemmas B.1 and B.2 for connectivity in H. The following lemmas describe the connectivity in H between s and any h-component, and between any h-component

and t.

Lemma B.3. *Let H be a subgraph of G, and suppose that all properties in the statement of Theorem 4.2 are satisfied by each subgraph H_i^h. Fix level h of the hierarchical k-chain decomposition, and fix a subgraph H_i^h in the level h chain. Then, in H, there is a path from s to each vertex in $V_{(i,r)}^h$; these paths only use edges within the level h subchain that begins at H_0^h and ends at H_i^h.*

Proof. First we show that for any subgraph H_j^h, there is a nonzero RSND demand from the vertex set $V_{(j,\ell)}^h$ to each vertex $v \in V_{(j,r)}^h$ that is satisfied in H_j^h. That is, for all $v \in V_{(j,r)}^h$, the RSND demand $(V_{(j,\ell)}^h, v, d)$ with $d > 0$ on G_j^h is satisfied in H_j^h. This would mean that there is a path in H_j^h from $V_{(j,\ell)}^h$ to v since there is a path in G_j^h from $V_{(j,\ell)}^h$ to v (from Lemma B.1).

We have two cases: If $\{v\} = V_{(j,r)}^h$, then the RSND demand $(V_{(j,\ell)}^h, v, h+1)$ on G_j^h is a nonzero demand satisfied in H_j^h. Now suppose that $\{v\} \neq V_{(j,r)}^h$. To show that a nonzero RSND demand from $X = V_{(j,\ell)}^h$ to $Y = \{v\}$ is satisfied in H_j^h, we plug $|S_X| = |S_{j-1}^h|$ and $|S_Y| \geq 1$ into the formula $d = \max(0, k + |S_X| + |S_Y| - |S_{j-1}^h| - |S_j^h|)$. We plug in $|S_Y| \geq 1$ because v is adjacent to at least one edge in S_j^h:

$$
\begin{aligned}
d &= \max(0, k + |S_X| + |S_Y| - |S_{j-1}^h| - |S_j^h|) \\
&\geq k + |S_X| + |S_Y| - |S_{j-1}^h| - |S_j^h| \\
&\geq k + |S_{j-1}^h| + 1 - |S_{j-1}^h| - |S_j^h| \\
&= k + 1 - |S_j^h| \\
&> 0. \qquad\qquad\qquad\qquad\qquad\qquad\qquad (|S_j^h| \leq k - 1)
\end{aligned}
$$

We have shown that there is a positive RSND demand from $V_{(j,\ell)}^h$ to each vertex in $V_{(j,r)}^h$. By Lemma B.1, we have that in subgraph G_j^h, there is a path from $V_{(j,\ell)}^h$ to each vertex in $V_{(j,r)}^h$. Therefore, in H_j^h, there must be a path from $V_{(j,\ell)}^h$ to each vertex in $V_{(j,r)}^h$. This also implies that in H, each vertex in $V_{(i,r)}^h$ is reachable from s, using only edges within the subchain from H_0^h to H_i^h. This can be seen via a proof by induction on the number of components into the level h subchain that starts at H_0^h and ends at H_i^h. $\qquad\qquad\square$

We now give a similar lemma describing the connectivity between any h-component and t.

Lemma B.4. *Let H be a subgraph of G, and suppose that all properties in the statement of Theorem 4.2 are satisfied by each subgraph H_i^h. Fix level h of the hierarchical k-chain decomposition, and fix a subgraph H_i^h in the level h chain. Then, in H, there is a path from each vertex in $V_{(i,\ell)}^h$ to t; these paths only use edges within the level h subchain that begins at H_i^h and ends at H_p^h.*

Proof. First we show that for all $v \in V_{(j,\ell)}^h$, the RSND demand $(v, V_{(j,r)}^h d)$ with $d > 0$ on G_j^h is satisfied in H_j^h. This would mean that there is a path in H_j^h from v to $V_{(j,r)}^h$ since there is a path in G_j^h from v to $V_{(j,r)}^h$ (from Lemma B.2).

We have two cases: If $\{v\} = V_{(j,\ell)}^h$, then the RSND demand $(v, V_{(j,r)}^h, h+1)$ on G_j^h is a nonzero RSND demand satisfied in H_j^h. To show that a nonzero RSND demand from $X = \{v\}$ to $Y = V_{(j,r)}^h$ on G_j^h is satisfied in H_j^h, we plug $|S_X| \geq 1$ and $|S_Y| = |S_j^h|$ into the formula $d = \max(0, k + |S_X| + $

$|S_Y| - |S_{j-1}^h| - |S_j^h|)$:

$$
\begin{aligned}
d &= \max(0, k + |S_X| + |S_Y| - |S_{j-1}^h| - |S_j^h|) \\
&\geq k + |S_X| + |S_Y| - |S_{j-1}^h| - |S_j^h| \\
&\geq k + 1 + |S_j^h| - |S_{j-1}^h| - |S_j^h| \\
&= k + 1 - |S_{j-1}^h| \\
&> 0. \hspace{5cm} (|S_{j-1}^h| \leq k - 1)
\end{aligned}
$$

We have shown that there is a positive RSND demand from each vertex in $V_{(j,\ell)}^h$ to $V_{(j,r)}^h$. By Lemma B.2, we have that in subgraph G_j^h, there is a path from each vertex in $V_{(j,\ell)}^h$ to $V_{(j,r)}^h$. Therefore, in H_j^h, there must be a path from each vertex in $V_{(j,\ell)}^h$ to $V_{(j,r)}^h$. This also implies that in H, there is a path from each vertex in $V_{(i,\ell)}^h$ to t, using only edges within the subchain from H_i^h to H_p^h. This can be seen via a proof by induction on the number of components into the level h subchain that starts at H_i^h and ends at H_p^h. $\qquad\square$

B.3.2 Analyzing Fault Sets via the Hierarchical Chain Decomposition

We say that an edge e is inside a subgraph G_i^h in the hierarchical k-chain decomposition if $e \in E(R_i^h)$. That is, an edge is in a subgraph G_i^h if the edge has both endpoints in R_i^h. Edges in the left or right separator of G_i^h are not considered to be inside G_i^h. We now consider two cases for the locations of the edge faults in the hierarchical k-chain decomposition. In **Case 1**, all fault edges in F are inside the same $(k-1)$-component in the decomposition. In **Case 2**, there is some level h such that the edges in F are not all inside the same h-component, but at level $h-1$, the edges in F are inside a single $(h-1)$-component.

Case 1. We first consider Case 1, which is much simpler, and prove that there is an st path in $H \setminus F$ under this case. Suppose we are in Case 1, and let R_i^{k-1} be the $(k-1)$-component that contains all of F.

Lemma B.5. *Suppose that Case 1 applies in the hierarchical k-chain decomposition; that is, each edge in F is inside the same $(k-1)$-component R_i^{k-1}. Suppose also that there is an st path in $G \setminus F$. Then, there is an st path in $H \setminus F$.*

Proof. In G_i^{k-1}, there are at least k edge-disjoint paths from $V_{(i,\ell)}^{k-1}$ to $V_{(i,r)}^{k-1}$ (otherwise, G_i^{k-1} would have a $(V_{(i,\ell)}^{k-1}, V_{(i,r)}^{k-1})$-separator with size at most $k-1$, contradicting the structure of the hierarchical decomposition). In addition, the subgraph H_i^{k-1} satisfies the RSND demand $(V_{(i,\ell)}^{k-1}, V_{(i,r)}^{k-1}, h+1)$ on G_i^{k-1}, where $h+1 = k$. Therefore, in H_i^{k-1}, there are also at least k edge-disjoint paths from $V_{(i,\ell)}^{k-1}$ to $V_{(i,r)}^{k-1}$. Since F has size at most $k-1$, there must be a path in $H_i^{k-1} \setminus F$ from $V_{(i,\ell)}^{k-1}$ to $V_{(i,r)}^{k-1}$. By Lemma B.3, there is a path from s to each vertex in $V_{(i,\ell)}^{k-1}$, and by Lemma B.4 there is a path from each vertex in $V_{(i,r)}^{k-1}$ to t. Combining all aforementioned paths gives a path from s to t in $H \setminus F$. $\qquad\square$

Case 2. We now consider Case 2, which is more complex. In Case 2 there is some level h such that the edges in F are not all inside the same h-component, but at level $h-1$, the edges in F are inside a single $h-1$-component. Let R_j^{h-1} be this $(h-1)$-component. Going forward, we will focus on what we will call the fault subchain of G.

Definition 5. The *fault subchain* of G is the $(V_{(j,\ell)}^{h-1}, V_{(j,r)}^{h-1})$ h-chain of R_j^{h-1}. Let $R_0, \ldots, R_p$ be the h-components of the fault subchain, and let $S_0, \ldots, S_{p-1}$ be the important separators adjacent to components in the fault subchain. Note that the left and right separators of R_j^{h-1} are not included in the fault subchain.

We will show that there is a path from $V_{(0,\ell)}$ to $V_{(p,r)}$ in the fault subchain after all faults are removed. With Lemmas B.3 and B.4, this will be enough to prove that there is an st path in $H \setminus F$ given an st path in $G \setminus F$.

We first give some useful notation and definitions. We denote the prefix subchains of the fault subchain as follows. Let L_i be all vertices in $\cup_{k=0}^i R_k$. Let L_i^G be the subgraph of G induced by L_i, and let L_i^H be the subgraph of H induced by L_i. Let f_i be the number of edge faults in L_i^G.

Let $i \leq m$. In general, we will say that an edge $e \in S_m$ in the fault subchain is "reachable" from a vertex set $X \subseteq V_{(i,\ell)}$ in $L_m^G \setminus F$ (in $L_m^H \setminus F$) if there is a path in $L_m^G \setminus F$ (in $L_m^H \setminus F$) from X to the vertex in $V_{(m,r)}$ that is incident on e, and the path only uses edges in the subchain from G_i^h to G_m^h. Note that e can be considered reachable even if $e \in F$. We will also use S_i^G (S_i^H) to denote the set of edges in S_i that are reachable from $V_{(0,\ell)}$ in $L_i^G \setminus F$ (in $L_i^H \setminus F$). That is, an edge e is in S_i^G (in S_i^H) if in $L_i^G \setminus F$ (in $L_i^H \setminus F$), there is a path from $V_{(0,\ell)}$ to the vertex in $V_{(i,r)}$ incident on e. In the next lemma, we prove a lower bound for the size of S_i^G.

Lemma B.6 (Lower bound on $|S_i^G|$). *Suppose there is an st path in $G \setminus F$. Let L_i^G be a prefix subchain of the fault subchain such that $i \neq p$, and let f_i be the number of edge faults in L_i^G. The edge set S_i^G has size at least $h - f_i$.*

Proof. Let $F_i \subseteq F$ be the set of fault edges in the subchain L_i^G, where $i \neq p$. Suppose for the sake of contradiction that $|S_i^G| < h - f_i$. Observe that the edge set $F_i \cup S_i^G$ separates $V_{(0,\ell)} = V_{(j,\ell)}^{h-1}$ and $V_{(p,r)} = V_{(j,r)}^{h-1}$. However, $F_i \cup S_i^G$ has size at most $f_i + (h - f_i - 1) = h - 1$. This means that $F_i \cup S_i^G$ is a $(V_{(j,\ell)}^{h-1}, V_{(j,r)}^{h-1})$-separator with size at most $h - 1$. This is a contradiction, since R_j^{h-1}, which is an $(h-1)$-component, cannot have a $(V_{(j,\ell)}^{h-1}, V_{(j,r)}^{h-1})$-separator with size $h - 1$ or less. We therefore have that S_i^G has size at least $h - f_i$. $\square$

For any subgraph G_i^h in the fault subchain such that $0 < i < p$, we give a lower bound on the number of edge disjoint paths between any two vertex sets $X \subseteq V_{(i,\ell)}$ and $Y \subseteq V_{(i,r)}$.

Lemma B.7 (Edge-disjoint paths lower bound). *Let R_i be an h-component in the fault subchain such that $i \neq 0$ and $i \neq p$. Let $X \subseteq V_{(i,\ell)}$ and $Y \subseteq V_{(i,r)}$, and let S_X and S_Y be the edges in S_{i-1} and in S_i that are incident on vertices in X and in Y, respectively. Let $d = \max(0, k + |S_X| + |S_Y| - |S_{i-1}| - |S_i|)$. In subgraph G_i^h, there are at least $d - k + h$ edge disjoint paths from X to Y.*

Proof. Suppose for the sake of contradiction that there are at most $d - k + h - 1$ edge disjoint paths from X to Y in subgraph G_i^h, where $i \neq 0$ and $i \neq p$. Then there exists an edge set E' with size at most $d - k + h - 1$ such that X and Y are disconnected in $G_i^h \setminus E'$. Let $E_X = S_{i-1} \setminus S_X$ and let $E_Y = S_i \setminus S_Y$. Observe that $E' \cup E_X \cup E_Y$ separates $V_{(0,\ell)} = V_{(j,\ell)}^{h-1}$ and $V_{(p,r)} = V_{(j,r)}^{h-1}$. We give an upper bound on the number of edges in $E' \cup E_X \cup E_Y$:

$$
\begin{aligned}
|E' \cup E_X \cup E_Y| &= |E'| + |E_X| + |E_Y| \\
&\leq (d - k + h - 1) + (|S_{i-1}| - |S_X|) + (|S_i| - |S_Y|) \\
&= k + |S_X| + |S_Y| - |S_{i-1}| - |S_i| - k + h - 1 + |S_{i-1}| - |S_X| + |S_i| - |S_Y| \\
&= h - 1.
\end{aligned}
$$

26

We have shown that $E' \cup E_X \cup E_Y$ is a $(V_{(j,\ell)}^{h-1}, V_{(j,r)}^{h-1})$-separator with size at most $h - 1$. This is a contradiction, since R_j^{h-1} cannot have a $(V_{(j,\ell)}^{h-1}, V_{(j,r)}^{h-1})$-separator with size $h - 1$ or less. We therefore have that there are at least $d - k + h$ edge disjoint paths from X to Y. $\qquad \square$

The following lemma is the crux of the argument for the sufficient condition. The lemma gives guarantees on the size of S_i^H when $i \neq p$. Later, we will use this lemma to argue that at least one vertex in $V_{(p,r)}$ is reachable from $V_{(0,\ell)}$ in the fault subchain, after the removal of faults. Going forward, we will let $r_i = k - 1 - f_i$ be the maximum number of edge faults that are in $(\cup_{k=i}^{p-1} S_k) \cup (\cup_{k=i+1}^{p} E(R_k))$ (that is, the maximum number of edge faults that are not in L_i^G).

Lemma B.8 (Properties of S_i^H when $i \neq p$). *Suppose there is an st path in $G \setminus F$. Let L_i^G be a prefix subchain such that $i \neq p$, and let $r_i = k - 1 - f_i$. Then, at least one of the following is true:*

1. $S_i^H = S_i^G$

2. $|S_i^H| \geq r_i + 1$.

Proof. We give a proof by induction on i, where R_i is a component of the fault subchain.

Base Case. Consider R_0 of the fault subchain. We will show that either $S_0^H = S_0^G$ or $|S_0^H| \geq r_0 + 1$. Let $F_R = F \cap E(R_0)$ be the set of fault edges with both endpoints in R_0, and let $f_R = |F_R|$. We will use S_{-1} to denote the left separator of $V_{(0,\ell)}$.

Using the RSND demands given in the Theorem 4.2 statement, we will show that for certain vertex subsets $Y \subseteq V_{(0,r)}$, there exists an RSND demand $(V_{(0,\ell)}, Y, d)$ on G_0^h such that $d > f_R$. Using the RSND demands from $V_{(0,\ell)}$ to these subsets $Y \subseteq V_{(0,r)}$, we will show that $S_0^H = S_0^G$ or $|S_0^H| = r_0 + 1$. Recall that for a vertex set $Y \subseteq V_{(0,r)}$, S_Y is the set of edges in S_0 that are incident on vertices in Y. We will first show that for all subsets $Y \subseteq V_{(0,r)}$ such that $|S_Y| \geq h - r_0$, there exists an RSND demand, $(V_{(0,\ell)}, Y, d)$, such that $d > f_R$. If $Y = V_{(0,r)}$, then the RSND demand $(X, Y, k - 1)$ is satisfied in H_0^h, regardless of the value of $|S_Y|$. Since $f_R \leq k - 2$ (recall that the edge faults are not all inside the same h-component in the fault subchain), we have that the RSND demand $(X, Y, f_R + 1)$ is satisfied in H_0^h. Now we consider Y such that $Y \neq V_{(0,r)}$. In this case, let $Y \subseteq V_{(0,r)}$ such that $|S_Y| \geq h - r_0$, and let $X = V_{(0,\ell)}$. Since the fault subchain is made up of at least two h-components, $|S_0| = h$. Plugging in $|S_X| = |S_{-1}|$, $|S_Y| \geq h - r_0$, and $|S_0| = h$ into $d = \max(0, k + |S_X| + |S_Y| - |S_{-1}| - |S_0|)$, we have the following:

$$
\begin{aligned}
d &= \max(0, k + |S_X| + |S_Y| - |S_{-1}| - |S_0|) \\
&\geq k + |S_X| + |S_Y| - |S_{-1}| - |S_0| \\
&\geq k + (h - r_0) - (h) && (|S_X| = |S_{-1}|, \ |S_Y| \geq h - r_0, \ |S_0| = h) \\
&= k - r_0 \\
&= k - k + 1 + f_0 && (r_0 = k - 1 - f_0) \\
&= f_0 + 1 \\
&\geq f_R + 1. && (f_0 \geq f_R)
\end{aligned}
$$

We have that if $|S_Y| \geq h - r_0$, the RSND demand $(V_{(0,\ell)}, Y, f_R + 1)$ on G_0^h is satisfied by H_0^h. Now we consider two cases and show that in the first case (1), $|S_0^H| \geq r_0 + 1$ and in the second case (2), $S_0^H = S_0^G$.

- **(1) Suppose first that $|S_0^G| > r_0$.** This means that $S_0 \setminus S_0^G$, the set of edges in S_0 not reachable from $V_{(0,\ell)}$ in $G_0^h \setminus F_R$, has size at most $|S_0| - r_0 - 1 = h - r_0 - 1$. Let $Y \subseteq V_{(0,r)}$ such that $|S_Y| \geq h - r_0$. Then, S_Y must contain at least one edge in S_0^G, so at least one vertex in Y is incident on an edge in S_0^G. Thus, there is a path from $V_{(0,\ell)}$ to Y in $G_0^h \setminus F_R$. We have shown that for all $Y \subseteq V_{(0,r)}$ such that $|S_Y| \geq h - r_0$, there is a path in $G_0^h \setminus F$ from $V_{(0,\ell)}$ to Y. We also have that the RSND demands $\{(V_{(0,\ell)}, Y, f_R + 1) : Y \subseteq V_{(0,r)}, |S_Y| \geq h - r_0\}$ on G_0^h are satisfied in H_0^h. Therefore, in $H_0^h \setminus F_R$, there is also a path from $V_{(0,\ell)}$ to a set $Y \subseteq V_{(0,r)}$ if $|S_Y| \geq h - r_0$.

  Suppose for the sake of contradiction that $|S_0^H| < r_0 + 1$. This means that $S_0 \setminus S_0^H$ has size at least $h - r_0$. The set $S_0 \setminus S_0^H$ is not reachable from $V_{(0,\ell)}$ in $H_0^h \setminus F_R$. However, since $|S_0 \setminus S_0^H| \geq h - r_0$, there must be a path in $H_0^h \setminus F_R$ from X to U, where U is set of vertices incident on at least one edge in $S_0 \setminus S_0^H$. This is a contradiction, so we have that S_0^H must have size at least $r_0 + 1$.

- **(2) Now suppose that $|S_0^G| \leq r_0$.** This means that $S_0 \setminus S_0^G$ has size at least $h - r_0$. Let e be an edge in S_0^G. Additionally, let $E_U \subseteq S_0 \setminus S_0^G$ such that E_U has size at least $h - r_0 - 1$. Let $S_e = \{e\} \cup E_U$. Let Y_e be the set of vertices in $V_{(0,r)}$ that are incident on at least one edge in S_e, and let v_e be the vertex in Y_e incident on e. Let S_{Y_e} be the set of edges in S_0 that are incident on a vertex in Y_e (note that $S_e \subseteq S_{Y_e}$). Since S_e has size at least $h - r_0$ and $|S_{Y_e}| \geq |S_e|$, the RSND demand $(V_{(0,\ell)}, Y_e, f_R + 1)$ on G_0^h is satisfied in H_0^h. Edge e is the only edge in S_e that is reachable from $V_{(0,\ell)}$ in $G_0^h \setminus F_R$, which means that v_e is the only vertex in Y_e that is reachable from $V_{(0,\ell)}$ in $G_0^h \setminus F_R$. The RSND demand $(V_{(0,\ell)}, Y_e, f_R + 1)$ is satisfied in H_0^h. The only way for this demand to be satisfied is if there is a path from $V_{(0,\ell)}$ to v_e (and thus to e) in $H_0^h \setminus F_R$. Applying this argument to all $e \in S_0^G$, in H_0^h there must be a path from $V_{(0,\ell)}$ to each edge in S_0^G, so $S_0^H = S_0^G$.

We have shown that if $|S_0^G| \geq r_0 + 1$ then $|S_0^H| \geq r_0 + 1$, and that if $|S_0^G| \leq r_0$ then $S_0^H = S_0^G$.

Inductive Step. We will now show that when $i \neq p$, if the lemma is true for R_{i-1}, then it also holds for R_i. Let $F_S = F \cap S_{i-1}$ be the set of fault edges in the left separator of R_i. Let $F_R = F \cap E(R_i)$, $f_S = |F_S|$ and $f_R = |F_R|$. We also let X be the set of vertices in $V_{(i,\ell)}$ that are incident on at least one edge in $S_{i-1}^H \setminus F_S$. A vertex $v \in V_{(i,\ell)}$ is in X if there is a path from $V_{(0,\ell)}$ to v in $G \setminus F$ that only uses edges in $L_{i-1}^G \cup S_{i-1}$. We also let S_X be the set of edges in S_{i-1} incident on a vertex in X. Note that $S_{i-1} \setminus F_S \subseteq S_X$. Recall that f_{i-1} is the number of faults in L_{i-1}^G, and that $r_{i-1} = k - 1 - f_{i-1}$. We have two cases by the inductive hypothesis: (1) $S_{i-1}^H = S_{i-1}^G$ or (2) $|S_{i-1}^H| \geq r_{i-1} + 1$.

- **(1) Suppose first that $S_{i-1}^H = S_{i-1}^G$.** As in the base case, we first show that for all $Y \subseteq V_{(i,r)}$ such that $|S_Y| \geq h - r_i$, there exists an RSND demand (X, Y, d), with $d > f_R$, that is satisfied in H_i^h. Again, if $(X, Y) = (V_{(i,\ell)}, V_{(i,r)})$, then the RSND demand $(X, Y, f_R + 1)$ is satisfied in H_i^h. Now we consider X and Y such that $(X, Y) \neq (V_{(i,\ell)}, V_{(i,r)})$. In this case, let $Y \subseteq V_{(i,r)}$ such that $|S_Y| \geq h - r_i$. Since $i - 1 \neq p$, by Lemma B.6 we have that $|S_{i-1}^G| \geq h - f_{i-1}$. Therefore, $|S_X| \geq |S_{i-1}^H \setminus F_S| = |S_{i-1}^G \setminus F_S| \geq |S_{i-1}^G| - f_S \geq h - f_{i-1} - f_S$. Also note that since $0 < i < p$, we have that $|S_{i-1}| = |S_i| = h$. Plugging in $|S_X| \geq h - f_{i-1} - f_S$, $|S_Y| \geq h - r_i$,

and $|S_{i-1}| = |S_i| = h$ into $d = \max(0, k + |S_X| + |S_Y| - |S_{i-1}| - |S_i|)$, we get the following:

$$
\begin{aligned}
d &= \max(0, k + |S_X| + |S_Y| - |S_{i-1}| - |S_i|) \\
&\geq k + |S_X| + |S_Y| - |S_{i-1}| - |S_i| \\
&\geq k + (h - f_{i-1} - f_S) + (h - r_i) - |S_{i-1}| - |S_i| \qquad (|S_X| \geq h - f_{i-1} - f_S, |S_Y| \geq h - r_i) \\
&= k + h - f_{i-1} - f_S + h - r_i - (h) - (h) \qquad\qquad\qquad (|S_{i-1}| = |S_i| = h) \\
&= k - f_{i-1} - f_S - r_i \\
&= k - f_{i-1} - f_S - k + 1 + f_{i-1} + f_S + f_R \qquad\qquad (r_i = k - 1 - f_{i-1} - f_S - f_R) \\
&= f_R + 1.
\end{aligned}
$$

We have shown that if $Y \subseteq V_{(i,r)}$ and $|S_Y| \geq h - r_i$, the RSND demand $(X, Y, f_R + 1)$ on G_i^h is satisfied by H_i^h. The rest of the argument is identical to that of the Base Case.

- **(2) Now suppose that $|S_{i-1}^H| \geq r_{i-1} + 1$.** We will first show that for all $Y \subseteq V_{(i,r)}$ such that $|S_Y| \geq h - r_i$, there are at least $f_R + 1$ edge-disjoint paths in G_i^h from X to Y. Then, as in the previous case, we will show that for all $Y \subseteq V_{(i,r)}$ such that $|S_Y| \geq h - r_i$, the RSND demand $(X, Y, f_R + 1)$ on G_i^h is satisfied in H_i^h. Since $0 < i < p$, by Lemma B.7 there are at least $d - k + h$ edge-disjoint paths from $X \subseteq V_{(i,\ell)}$ to $Y \subseteq V_{(i,r)}$ in G_i^h, where $d = \max(0, k + |S_X| + |S_Y| - |S_{i-1}| - |S_i|)$. Additionally, $|S_X| \geq |S_{i-1}^H| - f_S \geq r_{i-1} + 1 - f_S$. Plugging in $|S_X| \geq r_{i-1} + 1 - f_S$, $|S_Y| \geq h - r_i$, and $|S_{i-1}| = |S_i| = h$ into $d = \max(0, k + |S_X| + |S_Y| - |S_{i-1}| - |S_i|)$, we get the following:

$$
\begin{aligned}
d - k + h &= \max(0, k + |S_X| + |S_Y| - |S_{i-1}| - |S_i|) - k + h \\
&\geq k + |S_X| + |S_Y| - |S_{i-1}| - |S_i| - k + h \\
&\geq (r_{i-1} + 1 - f_S) + (h - r_i) - |S_{i-1}| - |S_i| + h \\
&\qquad\qquad\qquad (|S_X| \geq r_{i-1} + 1 - f_S, |S_Y| \geq h - r_i) \\
&= r_{i-1} + 1 - f_S + h - r_i - (h) - (h) + h \qquad\qquad (|S_{i-1}| = |S_i| = h) \\
&= r_{i-1} + 1 - f_S - r_i \\
&= (r_i + f_R + f_S) + 1 - f_S - r_i \qquad\qquad\qquad (r_{i-1} = r_i + f_R + f_S) \\
&= f_R + 1.
\end{aligned}
$$

We have shown that if $Y \subseteq V_{(i,r)}$ and $|S_Y| \geq h - r_i$, then $d - k + h \geq f_R + 1$. This means that in G_i^h there are at least $f_R + 1$ edge disjoint paths from X to Y if $Y \subseteq V_{(i,r)}$ and $|S_Y| \geq h - r_i$. Therefore, in $G_i^h \setminus F_R$, there is a path from X to all $Y \subseteq V_{(i,r)}$ such that $|S_Y| \geq h - r_i$. As before, if $(X, Y) = (V_{(i,\ell)}, V_{(i,r)})$, then the RSND demand $(X, Y, f_R + 1)$ is satisfied in H_i^h. We now show that this demand is also satisfied when $(X, Y) \neq (V_{(i,\ell)}, V_{(i,r)})$:

$$
\begin{aligned}
d - k + h \geq f_R + 1 \implies d &\geq f_R + 1 + k - h \\
&\geq f_R + 1. \qquad\qquad\qquad (k - h > 0)
\end{aligned}
$$

Thus, if $Y \subseteq V_{(i,r)}$ with $|S_Y| \geq h - r_i$, then the RSND demand $(X, Y, f_R + 1)$ on G_i^h is satisfied by H_i^h. This, combined with the fact that in $G_i^h \setminus F_R$ there is a path from X to all $Y \subseteq V_{(i,r)}$ such that $|S_Y| \geq h - r_i$, implies that there is a path in $H_i^h \setminus F_R$ from X to all $Y \subseteq V_{(i,r)}$ such that $|S_Y| \geq h - r_i$. Now, suppose for the sake of contradiction that $|S_i^H| \leq r_i$. This means that $S_i \setminus S_i^H$ has size at least $h - r_i$. The set $S_i \setminus S_i^H$ is not reachable from X in $H_i^h \setminus F_R$. However, since $|S_i \setminus S_i^H| \geq h - r_i$, there must be a path in $H_i^h \setminus F$ from X to U, where U

is the set of vertices incident on at least one edge in $S_i \setminus S_i^H$. This is a contradiction. We therefore have that $|S_i^H| \geq r_i + 1$.

We have shown that for all i such that $i \neq p$, either $S_i^H = S_i^G$ or $|S_i^H| \geq r_i + 1$. $\qquad \square$

Now we use Lemma B.8 to show that there is a path in the fault subchain from $V_{(0,r)}$ to at least one vertex in $V_{(p,r)}$. The proof uses arguments similar to those in Lemma B.8.

Lemma B.9. *Suppose there is an st path in $G \setminus F$. There is a path from $V_{(0,r)}$ to $V_{(p,r)}$ in $L_p^H \setminus F$.*

Proof. Note that since F is not contained in a single h-component, the fault subchain has at least two components, so component R_{p-1} must exist. Let $F_S = F \cap S_{p-1}$, $F_R = F \cap E(R_p)$, $f_S = |F_S|$, and $f_R = |F_R|$. Let X be the set of vertices in $V_{(p,\ell)}$ that are incident on at least one edge in $S_{p-1}^H \setminus F_S$. In addition, let S_X be the set of edges in S_{p-1} incident on a vertex in X, and let S_p denote the right separator of R_p. By Lemma B.8, one of the following holds: (1) $|S_{p-1}^H| = |S_{p-1}^G|$ or (2) $|S_{p-1}^H| \geq r_{p-1} + 1$.

- **(1) Suppose first that $|S_{p-1}^H| = |S_{p-1}^G|$.** There is an st path in $G \setminus F$; therefore, there is a path in $G_p^h \setminus F_R$ from X to $V_{(p,r)}$. We now show that the RSND demand $(X, V_{(p,r)}, f_R + 1)$ on G_p^h is satisfied in H_p^h. If $X = V_{(p,\ell)}$, the RSND demand $(X, V_{(p,r)}, f_R + 1)$ is satisfied. Now suppose $X \neq V_{(p,\ell)}$. By Lemma B.6, we have that $|S_{p-1}^G| \geq h - f_{p-1}$. We therefore have that $|S_X| \geq |S_{p-1}^H \setminus F_S| = |S_{p-1}^G \setminus F_S| \geq h - f_{p-1} - f_S$. Let $Y = V_{(p,r)}$. We plug in $|S_X| \geq h - f_{p-1} - f_S$, $|S_Y| = |S_p|$, and $|S_{p-1}| = h$ into $d = \max(0, k + |S_X| + |S_Y| - |S_{p-1}| - |S_p|)$:

$$
\begin{aligned}
d &= \max(0, k + |S_X| + |S_Y| - |S_{p-1}| - |S_p|) \\
&\geq k + |S_X| + |S_Y| - |S_{p-1}| - |S_p| \\
&\geq k + (h - f_{p-1} - f_S) - (h) && (|S_X| \geq h - f_{p-1} - f_S, |S_Y| = |S_p|, |S_{p-1}| = h) \\
&= k - f_{p-1} - f_S \\
&\geq k - k + 1 + f_S + f_R - f_S && (f_{p-1} \leq k - 1 - f_S - f_R) \\
&= f_R + 1.
\end{aligned}
$$

  We have shown that the RSND demand $(X, V_{(p,r)}, f_R + 1)$ on G_p^h is satisfied in H_p^h. Since there is a path in $G_p^h \setminus F_R$ from X to $V_{(p,r)}$, this means there must also be a path in $H_p^h \setminus F_R$ from X to $V_{(p,r)}$. This, with Lemma B.8, implies a path from $V_{(0,\ell)}$ to $V_{(p,r)}$.

- **(2) Now suppose that $|S_{p-1}^H| \geq r_{p-1} + 1$.** This means that $|S_X| \geq r_{p-1} + 1 - f_S$. We will first show that in G_p^h, there are at least $f_R + 1$ edge-disjoint paths from X to $V_{(p,r)}$. Suppose for the sake of contradiction that there are at most f_R edge-disjoint paths X to $V_{(p,r)}$. Then there exists a set E' size at most f_R such that X and $V_{(p,r)}$ and not connected in $G_p^h \setminus E'$. Let $E_X = S_{p-1} \setminus S_X$. Observe that $E' \cup E_X$ separates $V_{(0,\ell)} = V_{(j,\ell)}^{h-1}$ and $V_{(p,r)} = V_{(j,r)}^{h-1}$. We give an upper bound on the number of edges in $E' \cup E_X$:

$$
\begin{aligned}
|E' \cup E_X| &= |E'| + |E_X| \\
&\leq f_R + (|S_{p-1}| - |S_X|) \\
&\leq f_R + h - r_{p-1} - 1 + f_S && (|S_{p-1}| = h, |S_X| \geq r_{p-1} + 1 - f_S) \\
&\leq f_R + h - f_R - f_S - 1 + f_S && (r_{p-1} \geq f_R + f_S) \\
&= h - 1.
\end{aligned}
$$

We have shown that $E' \cup E_X$ is a $(V^{h-1}_{(j,\ell)}, V^{h-1}_{(j,r)})$-separator with size at most $h - 1$. This is a contradiction, since R^{h-1}_j cannot have a $(V^{h-1}_{(j,\ell)}, V^{h-1}_{(j,r)})$-separator with size $h - 1$ or less. We therefore have that there are at least $f_R + 1$ edge disjoint paths from X to $V_{(p,r)}$ in G^h_p. This means that in $G^h_p \setminus F_R$, there is a path from X to $V_{(p,r)}$. We now show that the RSND demand $(X, V_{(p,r)}, f_R + 1)$ is satisfied in H^h_p. As before, if $X = V_{(p,\ell)}$, then the demand is satisfied. Suppose $X \neq V_{(p,\ell)}$. Plugging in $|S_X| \geq r_{p-1} + 1 - f_S$, $|S_Y| = |S_p|$, and $|S_{p-1}| = h$ into $d = \max(0, k + |S_X| + |S_Y| - |S_{p-1}| - |S_p|)$, we get the following:

$$
\begin{aligned}
d &= \max(0, k + |S_X| + |S_Y| - |S_{p-1}| - |S_p|) \\
&\geq k + |S_X| + |S_Y| - |S_{p-1}| - |S_p| \\
&\geq k + (r_{p-1} + 1 - f_S) - h && (|S_X| \geq r_{p-1} + 1 - f_S, |S_Y| = |S_p|, |S_{p-1}| = h) \\
&\geq k + (f_R + f_S) + 1 - f_S - h && (r_{p-1} \geq f_R + f_S) \\
&= f_R + 1 + k - h \\
&\geq f_R + 1. && (k - h > 0)
\end{aligned}
$$

We have shown that the RSND demand $(X, V_{(p,r)}, f_R + 1)$ on G^h_p is satisfied by H^h_p. Since there is a path in $G^h_p \setminus F_R$ from X to $V_{(p,r)}$, this means that there must also be a path in $H^h_p \setminus F_R$ from X to $V_{(p,r)}$. This, with Lemma B.8, implies there is a path from $V_{(0,\ell)}$ to $V_{(p,r)}$. $\qquad \square$

Corollary B.10. *Suppose that Case 2 applies in the hierarchical k-chain decomposition; that is, there is some level h such that the edges in F are not all inside the same h-component, but at level $h - 1$, the edges in F are inside a single $(h - 1)$-component. Suppose there is an st path in $G \setminus F$. Then, there is an st path in $H \setminus F$.*

Proof. Suppose neither s nor t is in the fault subchain. By Lemma B.3, there is a path in $H \setminus F$ from s to each vertex in $V_{(0,\ell)}$, using only edges in the level h subchain starting at s and ending at $V_{(0,\ell)}$. Since there must be a path from $V_{(0,\ell)}$ to $V_{(p,r)}$ in $L^G_p \setminus F$, by Lemma B.9, there is also a path from $V_{(0,\ell)}$ to $V_{(p,r)}$ in $L^H_p \setminus F$. Finally, by Lemma B.4, there is a path in $H \setminus F$ from each vertex in $V_{(p,r)}$ to t, using only edges in the level h subchain starting at $V_{(p,r)}$ and ending at t. Combining all aforementioned paths gives a path in $H \setminus F$ from s to t. If s or t is in the fault subchain, then a similar argument shows that there is an st path in $H \setminus F$. $\qquad \square$

With both Lemma B.5 (Case 1) and Corollary B.10 (Case 2), we have shown that if H satisfies the set of demands specified in Theorem 4.2, then H is a feasible solution to SD-k-RSND.

C Proof of Lemma 4.3

Proof of Lemma 4.3. Fix level h of the hierarchical k-chain decomposition. For each i, let H^h_i denote the subgraph of H induced by R^h_i, and let G^h_i denote the subgraph of G induced by R^h_i. We will show that H satisfies the conditions stated in Theorem 4.2, and hence is feasible. By construction, H contains S, or all edges in the important separators. First we show that the set of RSND demands $\{(X, Y, d) : X \subseteq V^h_{(i,\ell)}, Y \subseteq V^h_{(i,r)}, (X, Y) \neq (V^h_{(i,\ell)}, V^h_{(i,r)}), d = \max(0, k + |S_X| + |S_Y| - |S^h_{i-1}| - |S^h_i|)\}$ on G^h_i are satisfied in H^h_i.

Let $X \subseteq V^h_{(i,\ell)}$, $Y \subseteq V^h_{(i,r)}$, and $d = \max(0, k + |S_X| + |S_Y| - |S_{i-1}| - |S_i|)$, as in Theorem 4.2. First we show that for all X, Y pairs such that $d = 1$, the RSND demand (X, Y, d) on G^h_i is satisfied

in H_i^h. If there is a path from X to Y in G_i^h, we include in H_i^h a path from X to Y. Therefore, the demand $(X, Y, 1)$ is satisfied.

Now we show that for all X, Y pairs such that $1 < d < k$, the RSND demand (X, Y, d) on G_i^h is satisfied in H_i^h. We include in H_i^h all edges selected via the SD-d-RSND algorithm on contracted X and contracted Y with demand d. Therefore, the demand (X, Y, d) is satisfied. We have shown that if $d < k$, then the RSND demand (X, Y, d) is satisfied. Since the demand for all X, Y pairs such that $(X, Y) \neq (V_{(i,\ell)}^h, V_{(i,r)}^h)$ is at most $k-1$, we have shown that the set of demands $\left\{ (X, Y, d) : X \subseteq V_{(i,\ell)}^h, Y \subseteq V_{(i,r)}^h, (X, Y) \neq (V_{(i,\ell)}^h, V_{(i,r)}^h), d = \max(0, k + |S_X| + |S_Y| - |S_{i-1}^h| - |S_i^h|) \right\}$ on G_i^h is satisfied in H_i^h.

Now we show that the RSND demand $\left(V_{(i,\ell)}^h, V_{(i,r)}^h, k-1 \right)$ on G_i^h is satisfied in H_i^h. We include in H_i^h all edges selected via the SD-$(k-1)$-RSND algorithm on contracted $V_{(i,\ell)}^h$ and contracted $V_{(i,r)}^h$ with demand $k-1$. Therefore, the demand $(V_{(i,\ell)}^h, V_{(i,r)}^h, k-1)$ is satisfied.

Finally we show that the RSND demand $\left(V_{(i,\ell)}^h, V_{(i,r)}^h, h+1 \right)$ on G_i^h is satisfied in H_i^h. Due to the construction of the hierarchical k-chain decomposition, there are at least $h+1$ edge-disjoint paths from $V_{(i,\ell)}^h$ to $V_{(i,r)}^h$ in G_i^h. Thus, to show that the RSND demand is satisfied in H_i^h, it suffices to show that there are $h+1$ edge-disjoint paths from $V_{(i,\ell)}^h$ to $V_{(i,r)}^h$ in H_i^h. In H_i^h, we include the edges selected via a Min-Cost Flow algorithm from contracted $V_{(i,\ell)}^h$ to contracted $V_{(i,r)}^h$ with flow $h+1$. Since there are $h+1$ edge-disjoint paths from $V_{(i,\ell)}^h$ to $V_{(i,r)}^h$ in G_i^h, the Min-Cost Flow instance must have a feasible solution, and the algorithm will return $h+1$ edge-disjoint paths from $V_{(i,\ell)}^h$ to $V_{(i,r)}^h$ (i.e. the set of edges with non-zero flow). $\qquad\square$

D   Simplification of k-EFTS 2-approximation

As discussed, a 2-approximation for k-EFTS was given by [17]. They showed that while the "cut-requirement" function defined by k-EFTS was not weakly supermodular (so the algorithm of [26] could not be applied in a black-box way), it did have a property they dubbed "local weak supermodularity" which was enough to be able to use [26]. However, this was relatively involved, and required a careful case analysis of the cut requirement function. We show here that more standard techniques can be used to achieve the same results in a far more simple manner.

D.1   Weakly $\mathcal{F}$-Supermodular Definition and Properties

First we give some preliminary definitions. Two sets A, B **cross** if $A \cap B \neq \emptyset$ and $A \cup B \neq V$, and **co-cross** if $A \setminus B, B \setminus A \neq \emptyset$. Given $A, B \in \mathcal{F}$, we say that f satisfies the:

- **supermodular inequality** if $f(A) + f(B) \geq f(A \cap B) + f(A \cup B)$.

- **co-supermodular inequality** if $f(A) + f(B) \geq f(A \setminus B) + f(B \setminus A)$.

A set function f is **symmetric** if $f(A) = f(V \setminus A)$ for all $A \subseteq V$, and f is:

- **crossing supermodular** if f satisfies the supermodular inequality whenever A, B cross.

- **weakly supermodular** if f satisfies the supermodular or the co-supermodular inequality for all $A, B \subseteq V$.

Note that if f is symmetric crossing supermodular and A, B co-cross, then (by symmetry) f satisfies the co-supermodular inequality.
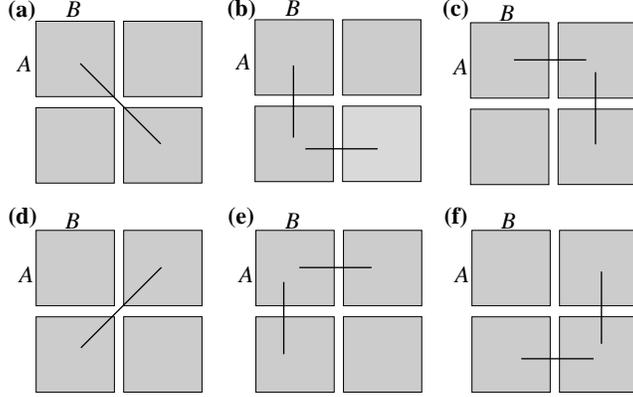
Figure 4: Illustration to the proof of Lemma D.1.

Given a set family $\mathcal{F}$ with $\emptyset, V \notin \mathcal{F}$, we may want to consider the restriction $f_{\mathcal{F}}$ of f to sets in $\mathcal{F}$. We say that f is **weakly $\mathcal{F}$-supermodular** if $f_{\mathcal{F}}$ is weakly supermodular, namely, if for all $A, B \in \mathcal{F}$ at least one of the following holds:

- $A \cap B, A \cup B \in \mathcal{F}$ and f satisfies the supermodular inequality.

- $A \setminus B, B \setminus A \in \mathcal{F}$ and f satisfies the co-supermodular inequality.

From Jain's result it follows that the problem of covering a weakly $\mathcal{F}$-supermodular function admits approximation ratio 2 (assuming we can solve certain LPs in polynomial time).

Lemma D.1. *Let $G = (V, E)$ be a graph, let $I \subseteq E$, let f be a symmetric crossing supermodular set function on $2^V \setminus \{\emptyset, V\}$. Then f is weakly $\mathcal{F}$-supermodular for $\mathcal{F} = \{A \subseteq V : d_I(A) \geq 1\}$.*

Proof. Let $A, B \in \mathcal{F}$. Then there is an edge in $\delta_I(A)$ and an edge in $\delta_I(B)$. Fig. 4 depicts all possible locations of (at least one of) such edges. In cases (a,b,c) we have $A \cap B, A \cup B \in \mathcal{F}$ and A, B cross – hence f satisfies the supermodular inequality. In cases (d,e,f) we have $A \setminus B, B \setminus A \in \mathcal{F}$ and A, B co-cross – hence f satisfies the co-supermodular inequality. $\square$

Remark: One can show that if f is symmetric crossing supermodular and $\mathcal{F}$ is **uncrossable** ($A \cap B, A \cup B \in \mathcal{F}$ or $A \setminus B, B \setminus A \in \mathcal{F}$ whenever $A, B \in \mathcal{F}$) and $\emptyset, V \notin \mathcal{F}$, then f is weakly $\mathcal{F}$-supermodular. The proof essentially shows that $\mathcal{F} = \{A \subseteq V : d_I(A) \geq 1\}$ is uncrossable. Alternatively, we may (roughly) say that if f satisfies *both* the supermodular and the co-supermodular inequality (whenever appropriate sets are not $\emptyset, V$), then f is weakly $\mathcal{F}$-supermodular.

D.2 k-EFTS 2-Approximation

Now we apply the weakly $\mathcal{F}$-supermodular concept to the k-EFTS problem. Note that k-EFTS is a particular case of the following meta-problem.

SET FUNCTION COVER
Input: A graph $G = (V, E)$ with edge costs, set function f defined on a family $\mathcal{F} \subseteq 2^V$.
Output: A min-cost subgraph H of G such that $d_H(A) \geq f(A)$.

Let $\mathcal{F} = 2^V \setminus \{\emptyset, V\}$. It is easy to see via Lemma A.1 that for k-EFTS in particular, we must cover the set function $f(A) = \min\{k, d_G(A)\}$ for $A \in \mathcal{F}$. If we already picked a set F of edges into our solution H, then we need to cover the residual function $\min\{d_G(A), k\} - d_F(A)$. We pick all the **forced edges** – edges that belong to cuts of G with at most k edges. During the iterative rounding algorithm we might pick additional edges. So let $F \subseteq E$ be a set of edges that includes all forced edges. Cuts $\delta_E(A)$ with $\delta_E(A) \subseteq F$ cannot and need not be covered, hence the function we need to cover is the restriction of $f(A) = k - d_F(A)$ (for $A \in 2^V \setminus \{\emptyset, V\}$) to the family $\mathcal{F} = \{A \subseteq V : d_{E \setminus F}(A) \geq 1\}$. In the setting of Lemma D.1, we have here $I = E \setminus F$. The function $k - d_F(A)$ is symmetric crossing supermodular, since $d_F(A)$ is symmetric and submodular. Hence by Lemma D.1 the function we need to cover is weakly $\mathcal{F}$-supermodular.

Theorem D.2. *The k-EFTS problem admits approximation ratio 2 for arbitrary costs and $1 + 4/k$ for unit costs.*

Proof. The set function we need to cover is the restriction of the function $f(A) = k - d_F(A)$ to the set family $\mathcal{F} = \{A \subseteq V : d_{E \setminus F}(A) \geq 1\}$. By Lemma D.1, we know that f is weakly $\mathcal{F}$-supermodular. The rest of the proof is as in [17]: one can solve LPs in polynomial time and thus apply iterative rounding. In the case of arbitrary costs we can use Jain's technique [26], while in the case of unit costs Gabow-Gallagher [23]. □