



מכון ויצמן למדע

WEIZMANN INSTITUTE OF SCIENCE

Thesis for the degree
Master of Science

עבודת גמר (תזה) לתואר
מוסמך למדעים

Submitted to the Scientific Council of the
Weizmann Institute of Science
Rehovot, Israel

מוגשת למועצה המדעית של
מכון ויצמן למדע
רחובות, ישראל

By
Ron Mosenzon

מאת
רון מוסנזון

מדללי זרימה באופן מדויק דורשים גודל לא חסום
Exact flow sparsification requires unbounded size

Advisor:
Robert Krauthgamer

מנחה:
רוברט קראוטגמר

July 2022

אב תשפ"ב

Exact Flow Sparsification Requires Unbounded Size

Ron Mosenzon

July 30, 2022

Abstract

Given a large edge-capacitated network G and a subset of k vertices called *terminals*, an (*exact*) *flow sparsifier* is a small network G' that preserves (exactly) all multicommodity flows that can be routed between the terminals. Flow sparsifiers were introduced by Leighton and Moitra [STOC 2010], and have been studied and used in many algorithmic contexts.

A fundamental question that remained open for over a decade, asks whether every k -terminal network admits an exact flow sparsifier whose size is bounded by some function $f(k)$ (regardless of the size of G or its capacities). We resolve this question in the negative by proving that there exist 6-terminal networks G whose flow sparsifiers G' must have arbitrarily large size. This unboundedness is perhaps surprising, since the analogous sparsification that preserves all terminal cuts (called *exact cut sparsifier* or *mimicking network*) admits sparsifiers of size $f_0(k) \leq 2^{2^k}$ [Hagerup, Katajainen, Nishimura, and Ragde, JCSS 1998].

We prove our results by analyzing the set of all feasible demands in the network, known as the *demand polytope*. We identify an invariant of this polytope, essentially the slope of certain facets, that can be made arbitrarily large even for $k = 6$, and implies an explicit lower bound on the size of the network. We further use this technique to answer, again in the negative, an open question of Seymour [JCTB 2015] regarding flow-sparsification that uses only contractions and preserves the infeasibility of one demand vector.

1 Introduction

Graph compression is a powerful paradigm in the design of graph algorithms. Where one reduces the size of the graph before performing heavy computation on it, while preserving properties of the graph that are needed for downstream computation. This paradigm is known to be extremely useful for faster computation, smaller memory requirement, as well as better accuracy when using approximation algorithms.

We study a form of graph compression known as *vertex sparsification*, where one is given an undirected network $G = (V, E, c)$ with edge capacities $c : E \rightarrow \mathbb{R}_{\geq 0}$, and a small set $T \subseteq V$ of vertices from G called *terminals*. The goal is to create a small network $G' = (V', E', c')$ called a vertex sparsifier, which may have different vertices, edges, and capacities than G , but includes the same set of terminals $T \subseteq V'$, and preserves some desired relationship between the terminals. We study flow sparsifiers, where the relationship of interest is the set of feasible multicommodity flows, as introduced by Leighton and Moitra [LM10].

Let us recall the relevant definitions. Given G and T as above, a pair of terminals $i \in \binom{T}{2}$ is referred to as a *commodity*, and the two terminals of each commodity are arbitrarily assigned to be the source and sink of that commodity. A *demand vector* $\mathbf{d} \in \mathbb{R}_{\geq 0}^{\binom{T}{2}}$ assigns a non-negative value d_i to each commodity $i \in \binom{T}{2}$. A multicommodity flow f that realizes \mathbf{d} is a collection of $\binom{|T|}{2}$ flow functions, namely, $f_i : E \rightarrow \mathbb{R}$ for each commodity $i \in \binom{T}{2}$, such that f_i ships d_i units of flow from the source to the sink of commodity i (satisfying flow conservation at every other vertex), and together these flows satisfy the the following capacity constraint

$$\forall e \in E, \quad \sum_{i \in \binom{T}{2}} |f_i(e)| \leq c(e).$$

Throughout, one fixes an arbitrary orientation of the edges of G , so that the sign of the flow $f_i(e)$ determines its direction. A demand vector \mathbf{d} is called *feasible* in G if there exists a multicommodity flow in G that realizes it. The set of all feasible demands in G , called the *demand polytope* (or *throughput polytope*, see e.g. [RBC07]) is defined as

$$\mathcal{D}(G) \stackrel{\text{def}}{=} \{\mathbf{d} \in \mathbb{R}_{\geq 0}^{\binom{T}{2}} : \mathbf{d} \text{ is feasible in } G\}.$$

It is not hard to see that $\mathcal{D}(G)$ is indeed a polytope.

Definition 1.1 (Exact Flow Sparsifier). A network $G' = (V', E', c')$ is called an *exact flow sparsifier* (or a flow sparsifier of quality 1) of a network $G = (V, E, c)$ if they share the same terminal set T and

$$\mathcal{D}(G') = \mathcal{D}(G).$$

1.1 Results

We answer the following problem, that is open since [LM10] and is well-known to experts: Does every k -terminal network G have an exact flow sparsifier of size bounded by some function $f(k)$? Interestingly, this question is open even if the sparsifier is allowed $O(1)$ -approximation, stated explicitly in [Moi11, Section 7] and [Chu12, Section 1]. Our main

result resolves this question in the negative (for exact sparsifiers) by showing that even 6-terminal networks can require arbitrarily large sparsifiers.

Theorem 1.2. *For every $k \geq 6$, there exists an infinite family $\{G_l\}_{l \in \mathbb{N}}$ of k -terminal networks, such that every exact flow sparsifier of every G_l must have at least l vertices.*

This theorem reveals a stark contrast with the analogous sparsification that preserves all terminal cuts, called *exact cut sparsifiers* (or *mimicking networks*), which admit constructions (exact cut sparsifiers) of size $f_0(k) \leq 2^{2^k}$ [HKNR98, KR14]. Furthermore, it is known that for networks G with at most 4 terminals, an exact cut sparsifier of G is also an exact flow sparsifier of G (because there is no flow-cut gap, see e.g. [AGK14]), and therefore such networks admit exact flow sparsifiers of size bounded by $f_0(4)$. This means that the condition $k \geq 6$ in the theorem is almost tight. We further note that the family of networks used in Theorem 1.2 excludes a fixed minor, namely, K_7 . We discuss these issues further in Section 5.

We further use our techniques to answer a question of Seymour [Sey15] regarding a related notion of flow-sparsification, that uses only contractions and preserves the infeasibility of one demand vector. This is discussed in Section 1.3.

1.2 Techniques and Proof Outline

Our proof of Theorem 1.2 is based on identifying a numerical feature of the demand polytope that forces a large sparsifier network. Some natural candidates are the maximum size of feasible demands, or the bit complexity required to represent the polytope. These candidates fail because even a small sparsifier can have arbitrarily large capacities, and thus can also represent an unbounded number of bits. To circumvent this problem, we identify a feature that can be upper bounded in a manner independent of the capacities, yet can grow arbitrarily large as the network size increases (even as the number of terminals remains fixed).

In essence, $\mathcal{D}(G)$ represents all possible tradeoffs between the values of the commodities. Since this is too complicated, we focus on the tradeoff between just two commodities, at the neighborhood of a single point $\mathbf{d} \in \mathbb{R}_{\geq 0}^{\binom{T}{2}}$. We then use this tradeoff as our numerical feature. We remark that our proof can be adapted to rely on the number of facets of the polytope, a numerical feature that is a widely used in polyhedral theory.¹

Definition 1.3 (Tradeoff). Let $\mathbf{d} \in \mathbb{R}_{\geq 0}^{\binom{T}{2}}$ be a point in the demand polytope $\mathcal{D}(G)$ of a network G with terminal set T . We say that $\tau \geq 0$ is a *feasible tradeoff* between commodity $i \in \binom{T}{2}$ and commodity $j \in \binom{T}{2}$ at the point \mathbf{d} if there exists a vector $\mathbf{d}^\Delta \in \mathbb{R}^{\binom{T}{2}}$ with entries $d_j^\Delta < 0$ and $d_i^\Delta = -\tau d_j^\Delta$ while all other entries are zero, such that $\mathbf{d} + \mathbf{d}^\Delta \in \mathcal{D}(G)$.² The

¹However, it seems more difficult to prove that the facet count can be arbitrarily large when the number of terminals is constant, as our proof method is essentially a reduction to tradeoff argument. Furthermore, we believe that our tradeoff feature is more versatile for analyzing demand polytopes, as exemplified by our refutation of Seymour’s conjecture [Sey15] in Section 1.3, which does not seem to have an analogous proof using facet counting.

²Observe that if $\tau > 0$ is a feasible tradeoff then also every $\tau' \in [0, \tau]$ is a feasible tradeoff, because the polytope $\mathcal{D}(G)$ is down-monotone.

maximum tradeoff between commodities i and j at the point \mathbf{d} is defined as

$$\tau_{\mathbf{d},i,j} \stackrel{\text{def}}{=} \sup\{\text{feasible tradeoff } \tau \geq 0 \text{ between commodities } i \text{ and } j \text{ at point } \mathbf{d}\}.$$

Observe that $\tau_{\mathbf{d},i,j}$ can be infinite. However, we are interested in finite $\tau_{\mathbf{d},i,j}$, and indeed the aforementioned numerical feature of the demand polytope is the largest finite $\tau_{\mathbf{d},i,j}$ over all \mathbf{d}, i, j . The maximum tradeoff $\tau_{\mathbf{d},i,j}$ has a geometric interpretation in terms of the demand polytope $\mathcal{D}(G)$. Consider the two-dimensional plane of points that differ from \mathbf{d} only in coordinates i and j , then the intersection of this plane with $\mathcal{D}(G)$ is a two-dimensional polytope. When \mathbf{d} is strictly inside this polytope, $\tau_{\mathbf{d},i,j}$ is infinite. However, when \mathbf{d} lies on a one-dimensional face of this two-dimensional polytope, $\tau_{\mathbf{d},i,j}$ is exactly the slope of this face, see Figure 1 for illustration. Furthermore, if the point \mathbf{d} is on exactly one facet of $\mathcal{D}(G)$, say a facet $\sum_{i' \in \binom{T}{2}} a_{i'} d_{i'} \leq b$, then the maximum tradeoff is $\frac{a_j}{a_i}$, which is essentially a “directional slope” of this facet.




Figure 1: The two-dimensional plane of points that differ from \mathbf{d} only in coordinates i and j . The region inside $\mathcal{D}(G)$ is shown in green stripes. The vector \mathbf{d}^Δ attains the maximum tradeoff $\tau_{\mathbf{d},i,j}$, which is the slope of the face on which \mathbf{d} lies.

Theorem 1.2 follows from the next two propositions, which we prove in Section 2 and Section 3, respectively.

Proposition 1.4 (Tradeoff Bound). *Let G be an n -vertex network with terminal set T , let $\mathbf{d} \in \mathcal{D}(G)$ be a point, and let $i^* \in \binom{T}{2}$ and $j^* \in \binom{T}{2}$ be two distinct commodities. If the maximum tradeoff $\tau_{\mathbf{d},i^*,j^*}$ is finite, then it is bounded by $\tau_{\mathbf{d},i^*,j^*} \leq (n^4)!$.*

Proposition 1.5 (Graph Family with Unbounded Tradeoff). *For every $M > 0$, there exists a 6-terminal network G_M , a point $\mathbf{d} \in \mathcal{D}(G_M)$, and two distinct commodities $i^* \in \binom{T}{2}$ and $j^* \in \binom{T}{2}$, for which the maximum tradeoff $\tau_{\mathbf{d},i^*,j^*}$ is finite and $\tau_{\mathbf{d},i^*,j^*} \geq M$.*

Proof of Theorem 1.2. For every $l \in \mathbb{N}$, choose $M > ((l - 1)^4)!$ and consider the network G_M from Proposition 1.5. It is then guaranteed by Proposition 1.4 that this network has no exact flow sparsifier with less than l vertices. \square

1.3 Seymour's Conjectures

Seymour [Sey15] investigated the following flow-sparsification problem.

Problem 1.6. Given a network $G = (V, E, c)$ with integral capacities together with an infeasible integral demand vector $\mathbf{d}^* \notin \mathcal{D}(G)$ on the set of commodities $\binom{V}{2}$, the goal is to apply on G edge contractions and obtain a small network G' in which the demand \mathbf{d}^* is still infeasible.³ Such a network G' is said to *compatible* with the problem instance (G, \mathbf{d}^*) .

Seymour [Sey15] proved the existence of compatible networks G' for (G, \mathbf{d}^*) with size bounded with respect to the *total demand*, i.e., the sum of the entries in the demand vector \mathbf{d}^* , and asked whether the size of a compatible network can be bounded independently of this parameter. Seymour actually posed two questions (see Section 1 and Table 1 in [Sey15]), which we rephrase and present as two conjectures:

- There exists $N > 0$, such that every instance of Problem 1.6 has a compatible network G' with at most N vertices.
- There exists $N_3 > 0$, such that every instance of Problem 1.6 with $\|\mathbf{d}^*\|_0 \leq 3$ has a compatible network G' with at most N_3 vertices. (As usual, $\|\cdot\|_0$ is the number of non-zero entries in a vector, and in our case it is the number of non-trivial commodities.)

The first conjecture clearly implies the second one, and is thus stronger.

We refute both conjectures. We present below the formal statement and its proof sketch, and provide the full proof in Section 4.

Theorem 1.7. *For every $N > 0$, there exists an instance of Problem 1.6 with $\|\mathbf{d}^*\|_0 \leq 3$, for which every compatible network must have more than N vertices.*

Proof sketch. In this sketch, we ignore the issues of bounded $\|\mathbf{d}^*\|_0$ and of integrality. Our goal in the proof is to construct an instance for which every compatible network must have a large tradeoff (i.e. have a point with a large finite maximum tradeoff), and thus, by Proposition 1.4, every compatible network must have a large number of vertices. To do this, we rely on two observations. The first one is that, because contraction operations can only eliminate/relax constraints on the flows, every network G' that is compatible with an instance (G, \mathbf{d}^*) must have $\mathcal{D}(G') \supseteq \mathcal{D}(G)$. The second observation is that the maximum tradeoff has an alternative characterization as follows:

Claim 1.8. *Let G be a network, and let i^* and j^* be two distinct commodities in $\binom{V}{2}$. Then, for every $\tau > 0$, the following two statements are equivalent.*

1. *There exists a point $\mathbf{d} \in \mathcal{D}(G)$ with a finite $\tau_{\mathbf{d}, i^*, j^*} > \tau$.*
2. *There exists a point $\mathbf{d}^* \notin \mathcal{D}(G)$, a point $\widehat{\mathbf{d}} \in \mathcal{D}(G)$ that results from \mathbf{d}^* by zeroing coordinate i^* , and a vector $\mathbf{d}^{\Delta*} \in \mathbb{R}^{\binom{V}{2}}$ with $d_{j^*}^{\Delta*} < 0$ and $d_{i^*}^{\Delta*} = -\tau d_{j^*}^{\Delta*}$ while all other entries are zero, such that $\mathbf{d}^* + \mathbf{d}^{\Delta*} \in \mathcal{D}(G)$. (See Figure 2 for illustration.)*

³Here, the contractions affect demands in the natural way: when an edge (u, v) is contracted, demands that involve its endpoints u or v carry over to contracted vertices. Notice that all vertices are viewed as terminals, but this is inconsequential.

Claim 1.8 tells us that in order to preserve a large tradeoff, it is enough to keep one specific point outside of the demand polytope and keep two specific points inside the demand polytope. Together with our first observation, this implies that by creating an instance from a network with a large tradeoff and using the right \mathbf{d}^* , we can force all compatible networks to have large tradeoff. The theorem follows from Proposition 1.5 and Proposition 1.4. \square




Figure 2: The condition for a large finite maximum tradeoff from statement 2 in Claim 1.8. The two-dimensional plane corresponds to points that differ from \mathbf{d}^* only in coordinates i^* and j^* , and the region inside $\mathcal{D}(G)$ is shown in green stripes.

1.4 Related Work

There is a large body of work on flow sparsification. It has begun with Leighton and Moitra [LM10], who studied approximate sparsifiers (in contrast to our focus on exact sparsifiers) that are supported *only* on the terminals. This setting was studied further by Charikar et al. [CLLM10], Makarychev and Makarychev [MM16], and Englert et al. [EGK⁺14], who tightened the approximation factor, called quality. Chuzhoy [Chu12, Chu16] introduced a more general setting of approximate flow sparsification, where the sparsifier can have non-terminal vertices, and she constructed sparsifiers whose size bound depends on the total capacity of all edges incident to the terminals, as well as a version of these sparsifiers that preserves both fractional and integral flows. Several improved bounds are known for important graph families, like quasi-bipartite graphs [AGK14], trees [GR16], and planar [EGK⁺14].

Another concept that is closely related to flow sparsification is *cut sparsification*. Exact cut sparsifiers were introduced by Hagerup, Katajainen, Nishimura, and Ragde [HKNR98], under the name *mimicking networks*, and tightened bounds for the general case and special graph families were proved in [CSWZ00, KR13, KR14, KPZ19, GHP20, KR20]. Approximate cut sparsifiers were introduced by Moitra [Moi09], and were often studied in conjunction with approximate flow sparsifier, because of the flow-cut gap. Another closely related notion is that of (exact) vertex-cut sparsifiers [KW20, HLW21, CDK⁺21, Liu20, BK22]. There are many other notions of sparsifiers, e.g., preserving other graph properties like spectrum or

distances, or sparsifiers that have few edges (but the same vertex set), and listing all of them here would be prohibitive.

1.5 Preliminaries

Throughout this paper, we use tilde to denote formal variables, such as \mathbf{x} .

Given a matrix $A \in \mathbb{R}^{m_1 \times m_2}$ with rows $\mathbf{a}_1^\top, \dots, \mathbf{a}_{m_1}^\top$, as well as a column vector $\mathbf{b} = (b_1, \dots, b_{m_1}) \in \mathbb{R}^{m_1}$, the *feasibility linear program* $A\mathbf{x} \geq \mathbf{b}$ is the problem of assigning values $\mathbf{x} = (x_1, \dots, x_{m_2})$ to the formal variables $\mathbf{x} = (\tilde{x}_1, \dots, \tilde{x}_{m_2})$, such that each of the inequalities $\{\mathbf{a}_i^\top \mathbf{x} \geq b_i\}_{i \in [m_1]}$ holds. Such an assignment \mathbf{x} is called a *feasible solution* of the linear program, and the set of all feasible solutions is called the *feasible region* of the linear program.

Definition 1.9. Let $A\mathbf{x} \geq \mathbf{b}$ be a linear program with m_2 variables, and let $S \subseteq [m_2]$ be a subset of the variables. Then, the *projection of the linear program* onto the variables of S is

$$\{\mathbf{x}_{|S} : \mathbf{x} \in \mathbb{R}^{m_2}, A\mathbf{x} \geq \mathbf{b}\},$$

where $\mathbf{x}_{|S}$ denotes the restriction of \mathbf{x} to the coordinates in S .

The following theorem is folklore.

Theorem 1.10. Let $A\mathbf{x} \geq \mathbf{b}$ be a linear program with m_1 constraints and m_2 variables, and suppose that $\max\{\mathbf{c}^\top \mathbf{x} : A\mathbf{x} \geq \mathbf{b}\}$ is finite. Then, there exists a subset $S_1 \subseteq [m_1]$ of constraints and a subset $S_2 \subseteq [m_2]$ of variables, such that the system of linear equalities $A_{|S_1 \times S_2} \mathbf{x}_{|S_2} = \mathbf{b}_{|S_1}$ has a unique solution, and such that extending this solution to the variables outside S_2 by setting them to 0 results in an optimal solution for the objective $\mathbf{c}^\top \mathbf{x}$ under constraint $A\mathbf{x} \geq \mathbf{b}$.

2 Tradeoff Bound (Proof of Proposition 1.4)

This section proves Proposition 1.4. For convenience, we first recall its statement.

Proposition 1.4 (Tradeoff Bound). Let G be an n -vertex network with terminal set T , let $\mathbf{d} \in \mathcal{D}(G)$ be a point, and let $i^* \in \binom{T}{2}$ and $j^* \in \binom{T}{2}$ be two distinct commodities. If the maximum tradeoff $\tau_{\mathbf{d}, i^*, j^*}$ is finite, then it is bounded by $\tau_{\mathbf{d}, i^*, j^*} \leq (n^4)!$.

Our proof of Proposition 1.4 relies on the following well-known lemma. For completeness, a proof of this lemma is presented in the appendix.

Lemma 2.1. $\mathcal{D}(G)$ can be described as the projection of a linear program $A\mathbf{x} \geq \mathbf{b}$ with at most n^4 variables onto $\binom{T}{2}$ of its variables that represent the commodities, such that all entries of A are from $\{-1, 0, 1\}$.

Let G, T, \mathbf{d}, i^* and j^* be as in Proposition 1.4, and let $A\mathbf{x} \geq \mathbf{b}$ be the linear program promised by Lemma 2.1. According to the lemma, $\mathcal{D}(G)$ is the projection of the program onto $\binom{T}{2}$ of its variables. Let $L : \mathbb{R}^{m_2} \rightarrow \mathbb{R}^{\binom{T}{2}}$ be the projection onto these variables. Let m_1 and m_2 be the number of constraints and the number of variables in the aforementioned

linear program. Then, $m_2 \leq n^4$. Since $\mathbf{d} \in \mathcal{D}(G)$, there exists a feasible solution \mathbf{x} of $A\mathbf{x} \geq \mathbf{b}$ such that $L(\mathbf{x}) = \mathbf{d}$. Let $EQ(\mathbf{x}) \subseteq [m_1]$ denote the set of inequalities in $A\mathbf{x} \geq \mathbf{b}$ that are tight at this solution \mathbf{x} , and let $A_{EQ(\mathbf{x})}\mathbf{x} \geq \mathbf{b}_{EQ(\mathbf{x})}$ be the linear program defined by this subset of the constraints. We use the following claim to complete the proof of Proposition 1.4.

Claim 2.2. *A value $\tau \geq 0$ is a feasible tradeoff at the point \mathbf{d} between commodities i^* and j^* iff there exists a solution $\mathbf{x}^\Delta \in \mathbb{R}^{m_2}$ of the linear program $A_{EQ(\mathbf{x})}\mathbf{x} \geq \mathbf{0}$, such that the vector $L\mathbf{x}^\Delta$ has $(L\mathbf{x}^\Delta)_{j^*} = -1$ and $(L\mathbf{x}^\Delta)_{i^*} = \tau$ while all other entries are zero.*

Proof. For the first direction, assume $\tau \geq 0$ is a feasible tradeoff. Then, by the definition of a feasible tradeoff, there exists a vector $\mathbf{d}^\Delta \in \mathbb{R}^{\binom{T}{2}}$ with $d_{j^*}^\Delta < 0$ and $d_{i^*}^\Delta = -\tau d_{j^*}^\Delta$ and all other entries being zero, such that $\mathbf{d} + \mathbf{d}^\Delta \in \mathcal{D}(G)$. Thus, there exists a solution \mathbf{x}' of $A\mathbf{x} \geq \mathbf{b}$ such that $L(\mathbf{x}') = \mathbf{d} + \mathbf{d}^\Delta$. Let $\mathbf{x}^\Delta \stackrel{\text{def}}{=} \mathbf{x}' - \mathbf{x}$. Then, by the linearity of L , we get that $L(\mathbf{x}^\Delta) = L(\mathbf{x}') - L(\mathbf{x}) = \mathbf{d}^\Delta$. Furthermore, since \mathbf{x}' is a solution of $A\mathbf{x} \geq \mathbf{b}$, it is also a solution of $A_{EQ(\mathbf{x})}\mathbf{x} \geq \mathbf{b}_{EQ(\mathbf{x})}$, and thus

$$A_{EQ(\mathbf{x})}\mathbf{x}^\Delta = A_{EQ(\mathbf{x})}(\mathbf{x}' - \mathbf{x}) \geq \mathbf{b}_{EQ(\mathbf{x})} - A_{EQ(\mathbf{x})}\mathbf{x} = \mathbf{0}$$

Since $(L\mathbf{x}^\Delta)_{j^*} = -|d_{j^*}^\Delta|$ and $(L\mathbf{x}^\Delta)_{i^*} = \tau|d_{j^*}^\Delta|$, the claim follows by scaling \mathbf{x}^Δ .

For the other direction, assume there exists \mathbf{x}^Δ as in Claim 2.2. Then, for every $\epsilon > 0$, it holds that $(L(\epsilon\mathbf{x}^\Delta))_{j^*} < 0$ and $(L(\epsilon\mathbf{x}^\Delta))_{i^*} = -\tau(L(\epsilon\mathbf{x}^\Delta))_{j^*}$ and that all other entries of $L(\epsilon\mathbf{x}^\Delta)$ are zero. Furthermore, for each $\epsilon > 0$ it also holds that $A_{EQ(\mathbf{x})}(\epsilon\mathbf{x}^\Delta) \geq \mathbf{0}$, and thus

$$A_{EQ(\mathbf{x})}(\mathbf{x} + \epsilon\mathbf{x}^\Delta) \geq \mathbf{b}_{EQ(\mathbf{x})}$$

We wish to choose ϵ such that $\mathbf{x} + \epsilon\mathbf{x}^\Delta$ is also a solution of the original program $A\mathbf{x} \geq \mathbf{b}$. To see that this is possible, notice that the constraints that appear in $A\mathbf{x} \geq \mathbf{b}$ but not in $A_{EQ(\mathbf{x})}\mathbf{x} \geq \mathbf{b}_{EQ(\mathbf{x})}$ are not tight at \mathbf{x} . Thus, by choosing a small enough $\epsilon > 0$, we get that $\mathbf{x} + \epsilon\mathbf{x}^\Delta$ is feasible for the original linear program, which means that $L(\mathbf{x} + \epsilon\mathbf{x}^\Delta) \in \mathcal{D}(G)$. By the linearity of L we get that $\mathbf{d} + L(\epsilon\mathbf{x}^\Delta) \in \mathcal{D}(G)$, and thus, τ is a feasible tradeoff between i^* and j^* at the point \mathbf{d} . \square

Using Claim 2.2, we now construct a linear program for finding the maximum tradeoff $\tau_{\mathbf{d},i^*,j^*}$ at the point \mathbf{d} . For each $i \in \binom{T}{2}$, let \tilde{x}_i denote the variable that L projects onto commodity i . We start with the linear program $A_{EQ(\mathbf{x})}\mathbf{x} \geq \mathbf{0}$ and add the constraint $\tilde{x}_{j^*} = -1$ as well as a constraint $\tilde{x}_i = 0$ for every $i \in \binom{T}{2} \setminus \{i^*, j^*\}$. Lastly, we add the optimization objective to maximize \tilde{x}_{i^*} . Let $A'\mathbf{x} \geq \mathbf{b}'$ denote the new linear program. Notice that the width of the matrix A' is m_2 , and that all entries of A' and \mathbf{b}' are from $\{-1, 0, 1\}$.

Remember that one of the conditions of Proposition 1.4 is that the maximum tradeoff $\tau_{\mathbf{d},i^*,j^*}$ is finite. Thus, by Theorem 1.10, there exists a submatrix A'' of A' and a vector \mathbf{b}'' whose entries are a subset of the entries of \mathbf{b}' , such that the maximum tradeoff $\tau_{\mathbf{d},i^*,j^*}$ is the i^* th entry of the unique solution to the system $A''\mathbf{x} = \mathbf{b}''$ of linear equalities. So, by Cramer's rule, $\tau_{\mathbf{d},i^*,j^*}$ is equal to the ratio of the determinants of two matrices whose entries are all from $\{-1, 0, 1\}$ and whose size is at most m_2 . Such matrices must have integer determinants whose absolute value is at most $m_2!$, and thus $\tau_{\mathbf{d},i^*,j^*} \leq m_2!$. By Lemma 2.1, $m_2 \leq n^4$, and this concludes the proof of Proposition 1.4.

3 Graph Family with Unbounded Tradeoff (Proof of Proposition 1.5)

This section proves Proposition 1.5. For convenience, we first recall its statement.

Proposition 1.5 (Graph Family with Unbounded Tradeoff). *For every $M > 0$, there exists a 6-terminal network G_M , a point $\mathbf{d} \in \mathcal{D}(G_M)$, and two distinct commodities $i^* \in \binom{T}{2}$ and $j^* \in \binom{T}{2}$, for which the maximum tradeoff $\tau_{\mathbf{d},i^*,j^*}$ is finite and $\tau_{\mathbf{d},i^*,j^*} \geq M$.*

At a high level, we construct a network G_M where both commodities i^* and j^* have to cross a cut $(V_{\text{left}}, V_{\text{right}})$. The idea is that the flow of commodity j^* would have to use a path that crosses this cut M times, while the flow of commodity i^* only has to cross it once. To force commodity j^* to use such a path, we introduce additional commodities that saturate all the edges that connect this path to the rest of the network.

We now present the network G_M , assuming for simplicity that M is an odd integer. The vertices of G_M are naturally divided into two sets, which we denote by V_{left} and V_{right} . Each side includes three terminals, which are denoted by $t_{\text{top_left}}, t_{\text{middle_left}}, t_{\text{bottom_left}}$ and $t_{\text{top_right}}, t_{\text{middle_right}}, t_{\text{bottom_right}}$. The network includes a path of length M from $t_{\text{middle_left}}$ to $t_{\text{middle_right}}$, whose internal vertices are non-terminals, and whose vertices (including the endpoints) alternate between the two sides V_{left} and V_{right} . Let e_1, \dots, e_M denote the edges of this path. Thus, each of e_1, \dots, e_M has one endpoint in V_{left} and one endpoint in V_{right} , which will be denoted by $\text{left}(e_i)$ and $\text{right}(e_i)$. For example, $\text{left}(e_1) = t_{\text{middle_left}}$ and $\text{right}(e_1) = \text{right}(e_2)$. The edges e_1, \dots, e_M have capacity 2. Furthermore, for each edge e_i in this path, the network includes edges of capacity 1 connecting $\text{left}(e_i)$ to $t_{\text{top_left}}$ and $t_{\text{bottom_left}}$, and similarly connecting $\text{right}(e_i)$ to $t_{\text{top_right}}, t_{\text{bottom_right}}$. Denote these edges by $e_{i,\text{top_left}}, e_{i,\text{bottom_left}}, e_{i,\text{top_right}}$ and $e_{i,\text{bottom_right}}$, respectively. For convenience, we allow edges to appear with multiplicity, e.g., $e_{1,\text{top_right}}$ and $e_{2,\text{top_right}}$ are two parallel edges of capacity 1 because $\text{right}(e_1) = \text{right}(e_2)$, although it could effectively be modeled by a single edge of capacity 2. Lastly, the network G_M includes an edge $(t_{\text{top_left}}, t_{\text{top_right}})$ of capacity $2M$. This concludes the description of the network G_M , see Figure 3 for illustration.

Let i^* and j^* be the commodities $i^* \stackrel{\text{def}}{=} (t_{\text{top_left}}, t_{\text{top_right}})$ and $j^* \stackrel{\text{def}}{=} (t_{\text{middle_left}}, t_{\text{middle_right}})$. The demand vector \mathbf{d} which we will be analyzing has a demand of M for the commodity $(t_{\text{bottom_left}}, t_{\text{top_right}})$, a demand of M for the commodity $(t_{\text{top_left}}, t_{\text{bottom_right}})$, a demand of 2 for commodity $j^* = (t_{\text{middle_left}}, t_{\text{middle_right}})$, and a demand of zero for every other commodity. It is easy to see that this vector is in the demand polytope; as depicted in Figure 4a, we can realize \mathbf{d} by sending 2 units of flow along the path (e_1, \dots, e_M) , sending M units of flow along the M paths of the form $(t_{\text{bottom_left}}, \text{left}(e_i), t_{\text{top_left}}, t_{\text{top_right}})$, and sending M units of flow along the M paths of the form $(t_{\text{bottom_right}}, \text{right}(e_i), t_{\text{top_right}}, t_{\text{top_left}})$. These $2M$ paths all overlap on the edge $(t_{\text{top_left}}, t_{\text{top_right}})$, but this is fine because this edge has capacity $2M$.

In order to prove Proposition 1.5, we need to show that M is a feasible tradeoff between commodities i^* and j^* at the point \mathbf{d} , as well as show that the maximum tradeoff $\tau_{\mathbf{d},i^*,j^*}$ is finite.

Let $\mathbf{d}^\Delta \in \mathbb{R}^{\binom{T}{2}}$ be the vector such that $d_{j^*}^\Delta = -2$ and $d_{i^*}^\Delta = -M d_{j^*}^\Delta = 2M$, and every other entry of \mathbf{d}^Δ is zero. To show that M is a feasible tradeoff between commodities i^* and




Figure 3: The network G_5 .

j^* at the point \mathbf{d} , we need to show that $\mathbf{d}' \stackrel{\text{def}}{=} \mathbf{d} + \mathbf{d}^\Delta$ is in the demand polytope $\mathcal{D}(G_M)$. To get some intuition, let us start by explicitly stating the demand vector \mathbf{d}' : It has a demand of $2M$ for commodity $i^* = (t_{\text{top_left}}, t_{\text{top_right}})$, a demand of M for the commodity $(t_{\text{bottom_left}}, t_{\text{top_right}})$, a demand of M for the commodity $(t_{\text{top_left}}, t_{\text{bottom_right}})$, and zero for all other commodities. Now, it is not hard to see that \mathbf{d}' is indeed in the demand polytope; as depicted in Figure 4b, it can be realized by sending the $2M$ units of flow of commodity i^* along the edge $(t_{\text{top_left}}, t_{\text{top_right}})$ of capacity $2M$, while sending M units of flow along the M paths of the form $(t_{\text{bottom_left}}, \text{left}(e_i), \text{right}(e_i), t_{\text{top_right}})$ and sending M units of flow along the M paths of the form $(t_{\text{bottom_right}}, \text{right}(e_i), \text{left}(e_i), t_{\text{top_left}})$.

In order to complete the proof of Proposition 1.5, we just need to show that the maximum tradeoff $\tau_{\mathbf{d}, i^*, j^*}$ is finite. To do this, we will use the following claim to show that every feasible tradeoff τ must satisfy $\tau \leq M$.

Claim 3.1. *Let \mathbf{d}'' be a demand vector in the demand polytope $\mathcal{D}(G_M)$ that has a demand of M for the commodity $(t_{\text{bottom_left}}, t_{\text{top_right}})$, and a demand of M for the commodity $(t_{\text{top_left}}, t_{\text{bottom_right}})$; the other entries in \mathbf{d}'' can be arbitrary. Then, it must hold that $d''_{i^*} + M \cdot d''_{j^*} \leq 2M$, which means that*

$$(-1) \cdot \frac{d''_{i^*} - d_{i^*}}{d''_{j^*} - d_{j^*}} = \frac{d''_{i^*}}{2 - d''_{j^*}} \leq M.$$

Proof. Fix some multicommodity flow that realizes the demand vector \mathbf{d}'' . Let $S_{\text{top}} \stackrel{\text{def}}{=} \{t_{\text{top_left}}, t_{\text{top_right}}\}$. Since the cut $(S_{\text{top}}, \overline{S_{\text{top}}})$ separates the endpoints of each of the two commodities $(t_{\text{bottom_left}}, t_{\text{top_right}})$ and $(t_{\text{top_left}}, t_{\text{bottom_right}})$, and since the total capacity of this cut is equal to the sum of the demands on these two commodities, we get that all edges

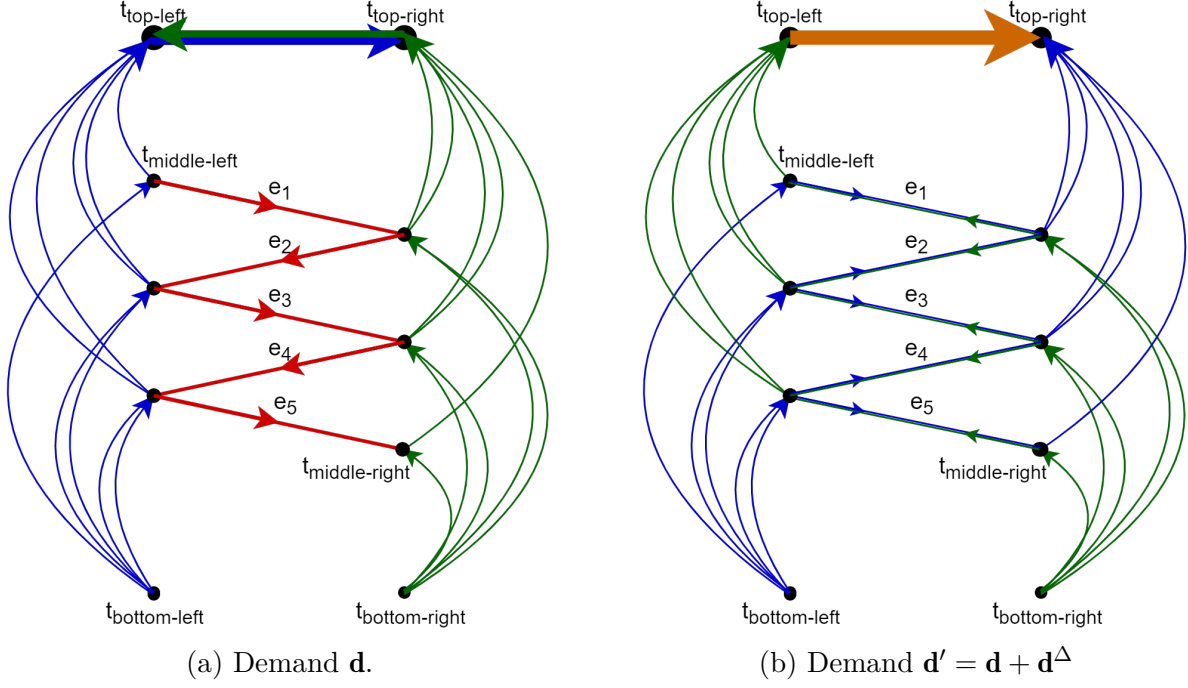


Figure 4: Multicommodity flows that realize the demand vectors \mathbf{d} and \mathbf{d}' , using colors to distinguish between the commodities.

crossing this cut must be saturated by the flows of the two commodities $(t_{\text{bottom_left}}, t_{\text{top_right}})$ and $(t_{\text{top_left}}, t_{\text{bottom_right}})$. Thus, the flow of commodity j^* cannot use any of the edges that cross this cut. By letting $S_{\text{bottom}} \stackrel{\text{def}}{=} \{t_{\text{bottom_left}}, t_{\text{bottom_right}}\}$ and making the exact same argument, we get that the flow of commodity j^* cannot use any of the edges that cross the cut $(S_{\text{bottom}}, \overline{S_{\text{bottom}}})$. However, the only path between the endpoints of commodity $j^* = (t_{\text{middle_left}}, t_{\text{middle_right}})$ that does not use any edges that cross either of the two aforementioned cuts is the path (e_1, \dots, e_M) . This means that every unit of flow of commodity j^* must pass M times through the cut $(V_{\text{left}}, V_{\text{right}})$, and therefore the flow of commodity j^* must use up a total of at least $M \cdot d_{j^*}''$ capacity from edges crossing the cut $(V_{\text{left}}, V_{\text{right}})$. Furthermore, because the cut $(V_{\text{left}}, V_{\text{right}})$ separates the endpoints of each of the other three commodities, the flows of commodities $(t_{\text{bottom_left}}, t_{\text{top_right}})$ and $(t_{\text{top_left}}, t_{\text{bottom_right}})$ must each use at least M capacity from edges crossing this cut, and the flow of commodity i^* must use at least d_{i^*}'' capacity from edges crossing this cut. But the total capacity of the cut $(V_{\text{left}}, V_{\text{right}})$ is $4M$, and thus

$$M \cdot d_{j^*}'' + M + M + d_{i^*}'' \leq 4M,$$

which gives us the desired inequality $d_{i^*}'' + M \cdot d_{j^*}'' \leq 2M$. □

This completes the proof of Proposition 1.5.

4 Refutation of Seymour's conjectures (Proof of Theorem 1.7)

We now prove one direction of Claim 1.8 (namely, that statement 2 implies statement 1), and then use it to prove Theorem 1.7. We omit the other direction of the claim because our proof of Theorem 1.7 does not use it, at least not in a blackbox manner; instead, our proof follows its steps implicitly and picks a vector \mathbf{d}^* with additional properties not promised by Claim 1.8 (namely, having integral entries and a small number of non-zero entries).

Claim 1.8. *Let G be a network, and let i^* and j^* be two distinct commodities in (V_2) . Then, for every $\tau > 0$, the following two statements are equivalent.*

1. *There exists a point $\mathbf{d} \in \mathcal{D}(G)$ with a finite $\tau_{\mathbf{d}, i^*, j^*} > \tau$.*
2. *There exists a point $\mathbf{d}^* \notin \mathcal{D}(G)$, a point $\widehat{\mathbf{d}} \in \mathcal{D}(G)$ that results from \mathbf{d}^* by zeroing coordinate i^* , and a vector $\mathbf{d}^{\Delta*} \in \mathbb{R}^{\binom{T}{2}}$ with $d_{j^*}^{\Delta*} < 0$ and $d_{i^*}^{\Delta*} = -\tau d_{j^*}^{\Delta*}$ while all other entries are zero, such that $\mathbf{d}^* + \mathbf{d}^{\Delta*} \in \mathcal{D}(G)$. (See Figure 2 for illustration.)*

Proof of Claim 1.8 (the direction (2) \rightarrow (1)). Since the points $\mathbf{d}^* \notin \mathcal{D}(G)$ and $\widehat{\mathbf{d}} \in \mathcal{D}(G)$ differ only in the i^* th coordinate, the line between them must intersect some facet $\sum_{i \in (V_2)} a_i \tilde{d}_i \leq b$ of the polytope $\mathcal{D}(G)$ for which $a_{i^*} > 0$. Let \mathbf{d} be the point at which this intersection occurs. Then, \mathbf{d} is in the polytope, and the maximum tradeoff $\tau_{\mathbf{d}, i^*, j^*}$ is at most $\frac{a_{j^*}}{a_{i^*}}$, which in particular means it is finite. Furthermore, since \mathbf{d} is on the line between \mathbf{d}^* and $\widehat{\mathbf{d}}$, it differs from these two points only in the i^* th coordinate, and by the down monotonicity of $\mathcal{D}(G)$ this implies that $d_{i^*} < d_{i^*}^*$. Thus, \mathbf{d} must have a feasible tradeoff larger than τ ; by defining $\mathbf{d}^{\Delta} \stackrel{\text{def}}{=} \mathbf{d}^* + \mathbf{d}^{\Delta*} - \mathbf{d}$, we get that $\mathbf{d} + \mathbf{d}^{\Delta} = \mathbf{d}^* + \mathbf{d}^{\Delta*} \in \mathcal{D}(G)$ and that $d_{i^*}^{\Delta} > d_{i^*}^*$ while every other coordinate of \mathbf{d}^{Δ} is the same as in $\mathbf{d}^{\Delta*}$. This means that $\tau_{\mathbf{d}, i^*, j^*} > \tau$, which concludes the proof. \square

Theorem 1.7. *For every $N > 0$, there exists an instance of Problem 1.6 with $\|\mathbf{d}^*\|_0 \leq 3$, for which every compatible network must have more than N vertices.*

To keep things simple, we first prove the theorem while dropping the condition $\|\mathbf{d}^*\|_0 \leq 3$, which is equivalent to refuting the first conjecture presented in Section 1.3, and then show how to adapt the proof to satisfy the condition $\|\mathbf{d}^*\|_0 \leq 3$, thereby refuting the second conjecture in Section 1.3.

Proof of Theorem 1.7 (without the condition $\|\mathbf{d}^\|_0 \leq 3$).* Fix any integer $N > 0$. Our goal is to provide an instance of Problem 1.6 in which the demand polytope of every compatible network must include a point with a finite maximum tradeoff greater than $(N^4)!$. Let $M \stackrel{\text{def}}{=} (N^4)! + 1$, and let G_M be the network from our proof of Proposition 1.5. Furthermore let $i^*, j^*, \mathbf{d}, \mathbf{d}^{\Delta}$ be as in that proof. Let the vector \mathbf{d}^* be the result of increasing coordinate i^* of \mathbf{d} by 1, and let $\mathbf{d}^{\Delta*} \stackrel{\text{def}}{=} \mathbf{d} + \mathbf{d}^{\Delta} - \mathbf{d}^*$. Thus, $d_{j^*}^{\Delta*} = -2$ and $d_{i^*}^{\Delta*} = 2M - 1 > 2 \cdot (N^4)!$, and every other entry of $\mathbf{d}^{\Delta*}$ is zero. Furthermore, $\mathbf{d}^* \notin \mathcal{D}(G_M)$. Notice that according to the construction of G_M and \mathbf{d}^* , the capacities of G_M and the entries of \mathbf{d}^* are all integral, and $\|\mathbf{d}^*\|_0 = 4$.

Now consider the instance of Problem 1.6 comprising of these network G_M and demand \mathbf{d}^* . Every compatible network G' of this instance must satisfy $\mathbf{d}^* \notin \mathcal{D}(G')$. Furthermore, as G' results from G_M by edge-contraction operations, it must satisfy $\mathcal{D}(G_M) \subseteq \mathcal{D}(G')$ and thus $\mathbf{d} \in \mathcal{D}(G')$ and $\mathbf{d}^* + \mathbf{d}^{\Delta*} = \mathbf{d} + \mathbf{d}^{\Delta} \in \mathcal{D}(G')$. Using Claim 1.8 with $\widehat{\mathbf{d}} = \mathbf{d}$, every compatible network G' must include a point that has a finite maximum tradeoff larger than $\frac{d^{\Delta*}_{i^*}}{|d^{\Delta*}_{j^*}|} > \frac{2 \cdot (N^4)!}{2}$, and the theorem follows using Proposition 1.4. \square

Proof of Theorem 1.7 (with the condition $\|\mathbf{d}^\|_0 \leq 3$).* Fix N , and let (G_M, \mathbf{d}^*) be the instance of Problem 1.6 from the previous proof. Thus, every compatible network G' of (G_M, \mathbf{d}^*) must have at least N vertices. However, the condition $\|\mathbf{d}^*\|_0 \leq 3$ does not hold, in fact $\|\mathbf{d}^*\|_0 = 4$.

Remove one unit from the capacity of the edge $(t_{\text{top_left}}, t_{\text{top_right}})$ in G_M as well as one unit from the demand of commodity $i^* = (t_{\text{top_left}}, t_{\text{top_right}})$ in \mathbf{d}^* , and let (G'_M, \mathbf{d}'_*) be the resulting instance. The fact that \mathbf{d}^* is not feasible in G_M implies that \mathbf{d}'_* is not feasible in G'_M , so this is indeed a valid instance of Problem 1.6. Furthermore, $\|\mathbf{d}'_*\|_0 \leq 3$ holds. It thus remains to show that every compatible network G'_- of the instance (G'_M, \mathbf{d}'_*) must have at least N vertices.

Let G'_- be any compatible network of the instance (G'_M, \mathbf{d}'_*) . This implies that G'_- is formed from G'_M using contractions, and that \mathbf{d}'_* is infeasible in G'_- . Let $t'_{\text{top_left}}$ and $t'_{\text{top_right}}$ be the vertices in G'_- that correspond to the vertices $t_{\text{top_left}}$ and $t_{\text{top_right}}$ in G'_M . The vertices $t'_{\text{top_left}}$ and $t'_{\text{top_right}}$ may have been formed from multiple vertices of G'_M through contractions, and may even be the same vertex (i.e. if $t_{\text{top_left}}$ and $t_{\text{top_right}}$ were merged together by contractions). By adding 1 to the capacity of the edge $(t'_{\text{top_left}}, t'_{\text{top_right}})$ in G'_- , we get a network G' that can be created from G_M using contractions, and satisfies that \mathbf{d}^* is infeasible in G' . Thus, G' is a compatible network of (G_M, \mathbf{d}^*) , which means it has at least N vertices. The theorem follows since G'_- has the same number of vertices as G' . \square

5 Future Directions

Our main result shows that there is no function $f : \mathbb{N} \rightarrow \mathbb{N}$ such that every k -terminal network G admits an exact flow sparsifier of size bounded by $f(k)$; in fact, this is not possible even for $k = 6$. We present here three interesting directions for future work.

Direction 1. *Extend our main result to planar networks G (and other graph families).*

Our framework offers a potential method to prove such a result; all one has to do is show that planar networks can have arbitrarily large finite tradeoffs (i.e., extend Proposition 1.5). A weaker version of such a result could restrict the sparsifier graph to be planar as well. We point out that Theorem 1.2 holds for input graphs excluding K_7 as a minor, and one can try to determine the smallest excluded minor for such results.

Direction 2. *Extend from exact sparsifiers to quality $1 + \epsilon$. That is, provide bounds on $f(k, \epsilon)$, the smallest value for which every k -terminal network G admits a $(1 + \epsilon)$ -quality flow sparsifier G' of size at most $f(k, \epsilon)$. The quality bound here means that G' has the same terminals as G and satisfies $\mathcal{D}(G) \subseteq \mathcal{D}(G') \subseteq (1 + \epsilon)\mathcal{D}(G)$.*

Our results imply a lower bound on $f(k, \epsilon)$, although it is somewhat weak, namely, $f(k, \epsilon) = \Omega\left(\left(\frac{\log(1/\epsilon)}{\log \log(1/\epsilon)}\right)^{1/4}\right)$ for every $k \geq 6$ and $\epsilon \in (0, \frac{1}{4})$. This can be strengthened by an $\Omega(k)$ factor, by simply taking $k/6$ copies of a network with 6 terminals. Using different techniques, we can prove stronger lower bounds $f(k, \epsilon) = \tilde{\Omega}(\frac{k}{\epsilon^{1/6}})$ and $f(k, \epsilon) = \tilde{\Omega}(\min\{2^{k/8}, \frac{k}{\epsilon^{1/2}}\})$ for every $k \geq 10$ and $\epsilon \in (0, \frac{1}{2})$, where the notation $\tilde{\Omega}(\cdot)$ hides $\text{polylog}(1/\epsilon)$ factors. In fact, it seems quite challenging to get any upper bound on $f(k, \epsilon)$, and this gap remains a fascinating open question.

Direction 3. *Extend the result to $k = 5$.*

It is known that for every $k \leq 4$, there exists $N'_k > 0$ such that every k -terminal network G has an exact flow sparsifier of size at most N'_k , see e.g. [AGK14]. Our main result shows that this is not possible for any $k \geq 6$, but it remains open for $k = 5$.

References

- [AGK14] A. Andoni, A. Gupta, and R. Krauthgamer. Towards $(1+\epsilon)$ -approximate flow sparsifiers. In *25th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 279–293, 2014. doi:10.1137/1.9781611973402.20.
- [BK22] Itai Boneh and Robert Krauthgamer. Optimal vertex-cut sparsification of quasi-bipartite graphs. *CoRR*, abs/2207.01459, 2022. arXiv:2207.01459.
- [CDK⁺21] Parinya Chalermsook, Syamantak Das, Yunbum Kook, Bundit Laekhanukit, Yang P. Liu, Richard Peng, Mark Sellke, and Daniel Vaz. Vertex sparsification for edge connectivity. In *32nd Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA’21, page 1206–1225. SIAM, 2021. doi:10.1137/1.9781611976465.74.
- [Chu12] Julia Chuzhoy. On vertex sparsifiers with Steiner nodes. In *44th symposium on Theory of Computing*, pages 673–688. ACM, 2012. doi:10.1145/2213977.2214039.
- [Chu16] Julia Chuzhoy. Routing in undirected graphs with constant congestion. *SIAM Journal on Computing*, 45(4):1490–1532, 2016. doi:10.1137/130910464.
- [CLLM10] Moses Charikar, Tom Leighton, Shi Li, and Ankur Moitra. Vertex sparsifiers and abstract rounding algorithms. In *51st Annual Symposium on Foundations of Computer Science*, pages 265–274. IEEE Computer Society, 2010. doi:10.1109/FOCS.2010.32.
- [CSWZ00] S. Chaudhuri, K. V. Subrahmanyam, F. Wagner, and C. D. Zaroliagis. Computing mimicking networks. *Algorithmica*, 26:31–49, 2000. doi:10.1007/s004539910003.
- [EGK⁺14] M. Englert, A. Gupta, R. Krauthgamer, H. Räcke, I. Talgam-Cohen, and K. Talwar. Vertex sparsifiers: New results from old techniques. *SIAM Journal on Computing*, 43(4):1239–1262, 2014. arXiv:1006.4586, doi:10.1137/130908440.
- [GHP20] Gramoz Goranci, Monika Henzinger, and Pan Peng. Improved guarantees for vertex sparsification in planar graphs. *SIAM Journal on Discrete Mathematics*, 34(1):130–162, 2020. doi:10.1137/17M1163153.

- [GR16] Gramoz Goranci and Harald Räcke. Vertex sparsification in trees. In *International Workshop on Approximation and Online Algorithms*, pages 103–115. Springer, 2016.
- [HKNR98] Torben Hagerup, Jyrki Katajainen, Naomi Nishimura, and Prabhakar Ragde. Characterizing multiterminal flow networks and computing flows in networks of small treewidth. *J. Comput. Syst. Sci.*, 57:366–375, 1998. doi:[10.1006/jcss.1998.1592](https://doi.org/10.1006/jcss.1998.1592).
- [HLW21] Zhiyang He, Jason Li, and Magnus Wahlström. Near-linear-time, optimal vertex cut sparsifiers in directed acyclic graphs. In *29th Annual European Symposium on Algorithms, ESA 2021*, volume 204 of *LIPICs*, pages 52:1–52:14. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2021. doi:[10.4230/LIPICs.ESA.2021.52](https://doi.org/10.4230/LIPICs.ESA.2021.52).
- [KPZ19] Nikolai Karpov, Marcin Pilipczuk, and Anna Zych-Pawlewicz. An exponential lower bound for cut sparsifiers in planar graphs. *Algorithmica*, 81(10):4029–4042, 2019. doi:[10.1007/s00453-018-0504-8](https://doi.org/10.1007/s00453-018-0504-8).
- [KR13] Robert Krauthgamer and Inbal Rika. Mimicking networks and succinct representations of terminal cuts. In *24th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1789–1799. SIAM, 2013. doi:[10.1137/1.9781611973105.128](https://doi.org/10.1137/1.9781611973105.128).
- [KR14] Arindam Khan and Prasad Raghavendra. On mimicking networks representing minimum terminal cuts. *Information Processing Letters*, 114(7):365–371, 2014. doi:[10.1016/j.ipl.2014.02.011](https://doi.org/10.1016/j.ipl.2014.02.011).
- [KR20] Robert Krauthgamer and Havana Inbal Rika. Refined vertex sparsifiers of planar graphs. *SIAM Journal on Discrete Mathematics*, 34(1):101–129, 2020. doi:[10.1137/17M1151225](https://doi.org/10.1137/17M1151225).
- [KW20] Stefan Kratsch and Magnus Wahlström. Representative sets and irrelevant vertices: New tools for kernelization. *J. ACM*, 67(3), 2020. doi:[10.1145/3390887](https://doi.org/10.1145/3390887).
- [Liu20] Yang P. Liu. Vertex sparsification for edge connectivity in polynomial time. *CoRR*, abs/2011.15101, 2020. arXiv:[2011.15101](https://arxiv.org/abs/2011.15101).
- [LM10] F. Thomson Leighton and Ankur Moitra. Extensions and limits to vertex sparsification. In *42nd ACM symposium on Theory of computing, STOC*, pages 47–56. ACM, 2010. doi:[10.1145/1806689.1806698](https://doi.org/10.1145/1806689.1806698).
- [MM16] Konstantin Makarychev and Yury Makarychev. Metric extension operators, vertex sparsifiers and Lipschitz extendability. *Israel Journal of Mathematics*, 212(2):913–959, 2016. doi:[10.1007/s11856-016-1315-8](https://doi.org/10.1007/s11856-016-1315-8).
- [Moi09] Ankur Moitra. Approximation algorithms for multicommodity-type problems with guarantees independent of the graph size. In *50th Annual Symposium on Foundations of Computer Science, FOCS*, pages 3–12. IEEE, 2009. doi:[10.1109/FOCS.2009.28](https://doi.org/10.1109/FOCS.2009.28).
- [Moi11] Ankur Moitra. *Vertex sparsification and universal rounding algorithms*. PhD thesis, Massachusetts Institute of Technology, 2011. URL: <https://dspace.mit.edu/handle/1721.1/66019>.
- [RBC07] Gábor Rétvári, József J. Biró, and Tibor Cinkler. Fairness in capacitated networks: A polyhedral approach. In *IEEE INFOCOM 2007-26th IEEE International Conference*

on *Computer Communications*, pages 1604–1612. IEEE, 2007. doi:10.1109/INFCOM.2007.188.

[Sey15] Paul Seymour. Criticality for multicommodity flows. *Journal of Combinatorial Theory, Series B*, 110:136–179, 2015. doi:10.1016/j.jctb.2014.08.001.

A Proof of Lemma 2.1

In this section we present a linear program with the properties promised by Lemma 2.1. The linear program that we use is the most standard linear program for the multicommodity flow problem. For each commodity $i \in \binom{T}{2}$ we have a variable \tilde{d}_i representing the number of units of flow shipped by the flow of the i -th commodity, as well as variables $\tilde{f}_{i,e}$ representing the value of f_i at each edge $e \in E$. Then, we have a constraint

$$\tilde{d}_i \geq 0,$$

and for each vertex $v \in V$, we have a flow-conservation constraint for the flow f_i at the vertex v . If $v \notin \{source(i), sink(i)\}$ then the flow-conservation constraint is

$$\sum_{e \text{ entering } v} \tilde{f}_{i,e} + \sum_{e \text{ leaving } v} (-1) \cdot \tilde{f}_{i,e} = 0.$$

However, if $v = source(i)$ then the flow-conservation constraint is

$$\tilde{d}_i + \sum_{e \text{ entering } v} \tilde{f}_{i,e} + \sum_{e \text{ leaving } v} (-1) \cdot \tilde{f}_{i,e} = 0,$$

and if $v = sink(i)$ then the constraint is

$$(-1) \cdot \tilde{d}_i + \sum_{e \text{ entering } v} \tilde{f}_{i,e} + \sum_{e \text{ leaving } v} (-1) \cdot \tilde{f}_{i,e} = 0$$

Lastly, aside from including each of the aforementioned variables and constraints for each commodity $i \in \binom{T}{2}$, we need to include constraints that force the multicommodity flow to respect the capacities of the edges. To do this, for each edge $e \in E$, and each possible choice of signs for the commodities $\{s_i \in \{1, -1\}\}_{i \in \binom{T}{2}}$, we add a constraint

$$\sum_{i \in \binom{T}{2}} s_i \cdot \tilde{f}_{i,e} \leq c(e).$$

Observe that these constraints are together equivalent to the non-linear constraint $\sum_{i \in \binom{T}{2}} |\tilde{f}_{i,e}| \leq c(e)$ that we need to enforce.

This concludes the presentation of the linear program. It is easy to see that the solutions of the program are in one-to-one correspondence to the multicommodity flows in the network, and thus that projecting the feasible region of the program onto the variables $\{\tilde{d}_i\}_{i \in \binom{T}{2}}$ gives the demand polytope $\mathcal{D}(G)$. Furthermore, the coefficients of the variables in the program

are all from $\{-1, 0, 1\}$. Lastly, since we have exactly $(1 + |E|)$ variables for each commodity $i \in \binom{T}{2}$, we get that the total number of variables in the program is

$$(1 + |E|) \cdot \binom{T}{2} \leq n^2 \cdot |T|^2 \leq n^4.$$

This concludes the proof of Lemma 2.1.