On Extremal Combinatorics and Its Applications in Matroid Theory

Zhiyang He

April 2021

Department of Mathematical Sciences Carnegie Mellon University Pittsburgh, Pennsylvania, United States, 15213

> Thesis Committee: Boris Bukh, Advisor Thomas Bohman Po-Shen Loh

Submitted in partial fulfillment of the requirements for the degree of Master of Science in Mathematical Sciences

Abstract

In this thesis, we study two problems that stem from the research of Tur´an number of even cycles, and one problem on the kernelization of minimum vertex cuts. In Chapter 2, we study the maximum number of hyperedges which may be in an r-uniform hypergraph under the restriction that no pair of vertices has more than t Berge paths of length k between them. When r = t = 2, this is the even-cycle problem asking for $ex(n, C_{2k})$. We extend results of Füredi and Simonovits and of Conlon, who studied the problem when r = 2.

In Chapter 3, we prove $ex(n, C_{2k}) \leq (16\sqrt{5}\sqrt{k \log k} + o(1)) \cdot n^{1+1/k}$. We improve on Bukh and Jiang's method used in their 2017 publication, thereby reducing the best known upper bound by a factor of $\sqrt{5 \log k}$.

In Chapter 4, we present a linear algorithm to construct minimum vertex cut sparsifiers of size $\Theta(k^2)$ in directed acyclic graphs with k terminals. Previously, Kratsch and Wahlström constructed a vertex cut sparsifier with $O(k^3)$ vertices via the theory of representative families on matroids. We draw inspiration from the renowned Bollobás's Two-Families Theorem in extremal combinatorics and introduce the use of skew-symmetry into Kratsch and Wahlström's methods, which leads to the stated improvement.

Acknowledgements

The following acknowledgements are ordered (mostly) chronologically.

I would like to thank Professor Michael Tait for being my first advisor, and thereby introducing me to the world of theoretical research. Chapter 2 is joint work with him, and we would like to thank two anonymous referees for a very helpful and detailed reading of the first version of the manuscript.

I would also like to thank Professor Boris Bukh for being a long time advisor of mine as well as advising this thesis. I would not be where I am today without his advice and teaching. Chapter 3 is joint work with him, and we would like to thank two anonymous referees for very helpful comments on an earlier version of the manuscript.

I further want to thank Professor Anupam Gupta and Dr. Jason Li for leading me into the study of theoretical computer science. Their patient guidance has a tremendous impact on my current pursuit in the field. Chapter 4 is joint work with them and Professor Magnus Wahlström.

Additionally, I would like to thank Professor Thomas Bohman and Professor Po-Shen Loh for serving on my committee, and more importantly, for being amazing professors throughout my years at Carnegie Mellon University.

Finally, I would like to thank my parents for their continuous support throughout my academic endeavor.

Contents

1	Introduction		2
	1.1	The Even-Cycle Problem	2
	1.2	An Extremal Result in Matroid Theory	3
2	The Even-Cycle Problem in Hypergraphs		
	2.1	Introduction	4
	2.2	Proof of Theorem 2.1.2	5
	2.3	Preliminaries for Theorem 2.1.4	8
	2.4	Proof of Theorem 4	10
	2.5	Conclusion	14
3	Improving the Constant		
	3.1	Introduction	16
	3.2	Graph Reduction and Exploration	17
	3.3	Results on Θ -Graphs	19
	3.4	Search for Θ -graphs	20
	3.5	Proof of Theorem 3.2.1	30
	3.6	Potential Improvements	31
4	Representative Sets Method in Matroid Theory		
	4.1	Introduction	35
	4.2	Preliminaries	37
	4.3	Vertex Cut Sparsifier for DAGs	40
	4.4	Lower Bound Constructions	43
	4.5	Conclusions	45

Chapter 1

Introduction

1.1 The Even-Cycle Problem

The study of extremal combinatorics has a heavy focus on unavoidable structures. A common theme that lies at the heart of many important problems in the field is the following question: when the size of the system grows to infinity, what are the structures that must, or must not, appear? A natural follow-up to this question is, can we derive upper and lower bounds on the size of our systems (for these structural characteristics to emerge)? These types of inquiries are best embodied by one of the most well-studied groups of problems in extremal graph theory – the Tuŕan type problems, which study the following notion in various settings.

Definition 1.1.1 (Tur'an Number). Given a graph F, we denote by ex(n, F) the maximum number of edges that a graph on n vertices can have while not containing F as a subgraph. Similarly, for a family of graphs \mathcal{F} , $ex(n, \mathcal{F})$ requires that no element of \mathcal{F} is present.

The first result, known as Mantel's Theorem, was proven by Mantel [41] in 1907. Since then, extensive amount of works have been established, among which is the celebrated Erdős–Stone–Simonovits Theorem [19]: $ex(n, F) = (1 - \frac{1}{\chi(F)-1} + o(1))\binom{n}{2}$, where $\chi(F)$ is the chromatic number of F. This result, proven in 1946, essentially solved Turán's Problem for all graphs F with $\chi(F) > 2$. However, for bipartite F it only gives $ex(n, F) = o(n^2)$, and determining the order of magnitude for the Turán number of a bipartite graph in general is a difficult problem (see [23] for a survey).

When $F = C_{2k}$, the study of $ex(n, C_{2k})$ is known as the even-cycle problem. This problem was first studied in 1938 by Erdős [16] and has since become a central problem in extremal graph theory. A general upper bound of $ex(n, C_{2k}) = O_k(n^{1+1/k})$ was first

published by Bondy and Simonovits [4]. Since then, improvements have been made to the multiplicative constant [8,47,49], and constructions have been found showing that the order of magnitude is correct for $k \in \{2,3,5\}$ [1,5,17,50]. However, besides C_4 , C_6 , and C_{10} , the precise order of magnitude is unknown.

In the next two chapters, we present new results on the even-cycle problem. In Chapter 2, we study the problem in hypergraphs and extend well-known results in simple graphs to *r*-uniform hypergraphs [31]. The upper bound we derive employs a reduction lemma and generalizes the results of Bondy and Simonovits [4]. The lower bound construction utilizes a random polynomial method developed by Blagojević, Bukh and Karasev [2] as well as results from algebraic geometry. In Chapter 3, we return to simple graphs and present an improvement on the multiplicative constant in the bound $ex(n, C_{2k}) = O_k(n^{1+1/k})$ [29]. Our proof is an improved and simplified version of Bukh and Jiang's methods [8], with a different delivery.

1.2 An Extremal Result in Matroid Theory

In Chapter 4, we deviate from our previous discussion on Tuŕan problems and explore another topic of similar flavour – linear dependency structures in extremal set systems. To start with, we present the famous Bollobás's Two Families Theorem.

Theorem 1.2.1. Let $\mathcal{A} = \{A_1, \dots, A_m\}$ and $\mathcal{B} = \{B_1, \dots, B_m\}$ be families of sets such that for all $i \in [m]$, $|A_i| \leq r$ and $|B_i| \leq s$. Suppose for all i, we have $A_i \cap B_i = \emptyset$ and for all $j \neq i$, $A_i \cap B_j \neq \emptyset$. Then $m \leq \binom{r+s}{s}$.

This theorem was first proved by Bollobás [3] in 1965, and was later generalized into matroid settings by Lovász (see Theorem 4.7 in [38]). An algorithmic version of this theorem in matroid settings, which we refer to as the representative set families method, was presented by Marx [42], and applied by Kratsch and Wahlström in their landmark paper [35] to derive new results on many problems in kernelization.

It is worth noting, however, that Theorem 1.2.1 has a skew-symmetric version, where the same result holds with the conditions relaxed from "for all $j \neq i$, $A_i \cap B_j \neq \emptyset$ " to "for all j > i, $A_i \cap B_j \neq \emptyset$ ". More importantly, this skew-symmetry can be propagated through the above line of work to derive a skew-symmetric representative set families method. In Chapter 4 of this thesis, we apply this method to the minimum vertex cut sparsifier problem, which is studied by Kratsch and Wahlström in [35], and derive a stronger result [30].

Chapter 2

The Even-Cycle Problem in Hypergraphs

2.1 Introduction

In this chapter, we focus on a generalization of the even-cycle problem to hypergraphs. Note that a C_{2k} is a pair of internally disjoint paths of length k between a fixed pair of vertices. Analogously, we will study hypergraphs where we forbid a certain number of paths of length k between vertices. In a hypergraph H, a *Berge path* of length k is a set of distinct vertices v_0, v_1, \dots, v_k and a set of distinct hyperedges h_1, h_2, \dots, h_k such that $\{v_{i-1}, v_i\} \in h_i$ for $1 \le i \le k$. Note that the hyperedges $\{h_i\}$ could intersect in many different ways and in general two Berge paths of length k need not be isomorphic. We call the vertices v_0, \dots, v_k the *core vertices* of the Berge path. Note that given a Berge path, the core vertices need not be unique.

The definition of Berge paths and cycles was extended to arbitrary graphs by Gerbner and Palmer [24]. We say a hypergraph H is a *Berge*-F if there is a bijection $\phi : E(F) \rightarrow E(H)$ such that $e \subseteq \phi(e)$ for all $e \in E(F)$. Again note that two hypergraphs H_1 and H_2 may each be a Berge-F and could be non-isomorphic. We denote by $F^{B,r}$ the set of all r-uniform hypergraphs which are a Berge-F and we will write F^B when the uniformity is fixed. Extending the extremal notation, given a family of r-uniform hypergraphs \mathcal{F} , we denote the maximum number of hyperedges in an n vertex r-uniform hypergraph which does not contain any $F \in \mathcal{F}$ as a subhypergraph by $ex_r(n, \mathcal{F})$.

In this chapter, we are interested in studying $ex_r(n, \mathcal{F})$ when \mathcal{F} is a family of Berge theta graphs. A *theta graph*, denoted by $\Theta_{k,t}$, is the (2-uniform) graph given by a set of t internally disjoint paths of length k between a fixed pair of vertices. Note that $\Theta_{k,2} = C_{2k}$

and so the study of the Turán number of $\Theta_{k,t}$ generalizes the even-cycle problem. We will study *r*-uniform hypergraphs which do not contain a Berge- $\Theta_{k,t}$, i.e. we will study the Turán number for the family $\Theta_{k,t}^B$. An alternative definition of a hypergraph in $\Theta_{k,t}^B$ is a set of distinct vertices $x, y, v_1^1, \dots, v_{k-1}^1, \dots, v_1^t, \dots, v_{k-1}^t$ and a set of distinct *r*-edges $h_1^1, \dots, h_k^1, \dots, h_1^t, \dots, h_k^t$ such that $\{x, v_1^i\} \subset h_1^i, \{v_{j-1}^i, v_j^i\} \subset h_j^i$, and $\{v_{k-1}^i, y\} \subset h_k^i$ for $1 \leq i \leq t$ and $2 \leq j \leq k - 1$. That is, we are forbidding that a pair of vertices have *t* Berge paths of length *k* between them with disjoint internal core vertices. For (2-uniform) graphs, the study of $ex(n, \Theta_{k,t})$ has been investigated in [9, 14, 20]. In particular, Faudree and Simonovits gave a more general version of the upper bound in the even-cycle problem.

Theorem 2.1.1 (Faudree and Simonovits [20]). *Given two integers* k *and* t*, there exists a constant* $c_{k,t} > 0$ *, such that*

$$\exp(n, \Theta_{k,t}) \le c_{k,t} n^{1+\frac{1}{k}}.$$

Our first main result is an extension of this bound to *r*-uniform hypergraphs.

Theorem 2.1.2. For fixed r, given two integers k and t, there exists a constant $c_{r,k,t}$, such that

$$\exp_r(n, \Theta^B_{k,t}) \le c_{r,k,t} n^{1+\frac{1}{k}}.$$

Recently, Conlon [14] complemented the upper bound of Faudree and Simonovits and showed that the order of magnitude is correct when t is large enough relative to k.

Theorem 2.1.3 (Conlon [14]). For any natural number $k \ge 2$, there exists a natural number t such that

$$\operatorname{ex}(n,\Theta_{k,t}) = \Omega_k(n^{1+\frac{1}{k}}).$$

Our second theorem shows that this is also the case in higher uniformities.

Theorem 2.1.4. For fixed r and any natural number $k \ge 2$, there exists a natural number t such that

$$\operatorname{ex}_{r}(n, \Theta_{k,t}^{B}) = \Omega_{k,r}(n^{1+\frac{1}{k}}).$$

In Section 2.2 we prove Theorem 2.1.2 and in Section 2.4 we prove Theorem 2.1.4

2.2 Proof of Theorem 2.1.2

After submiting the work in this chapter, we learned that Theorem 2.1.2 also follows from a more general theorem of Gerbner, Methuku, and Vizer (remark 13 in [26]). Their

theorem states that if F is graph with a vertex whose deletion makes the graph acyclic, then $ex_r(n, Berge - F) = O(ex(n, F))$. Since $\Theta_{k,t}$ has this property, the upper bound follows. We include our proof as well, as it gives an explicit multiplicative constant. We further comment that this multiplicative constant was later improved by Gerbner, Methuku and Palmer in [25].

The proof of Theorem 2.1.2 is inspired by a reduction lemma of Győri and Lemons [28] (see also Lemma 2.16 of [37]). To prove Theorem 2.1.2 from Theorem 2.1.1, we prove the following lemma, which says that given any *r*-uniform $\Theta_{k,t}^B$ -free hypergraph, we can reduce it to a $\Theta_{k,t}$ -free (2-uniform) graph which has a constant proportion of the edges of the original *r*-uniform hypergraph.

Lemma 2.2.1 (Reduction Lemma). Let $2 \le m < r$, and H be a r-uniform hypergraph without $\Theta_{k,t}^B$. Define the m-uniform hypergraph G in the following way: Order the edges of H arbitrarily and, going through the edges one by one, pick an m-set from each hyperedge of H to be in G where the m set chosen is the one which has been chosen the fewest number of times previously (break ties arbitrarily). Let $M_{k,i,r,m} = \sum_{j=1}^{k+1} {mk(i-1)+jm-m \choose r-m} + k + 1$. Then the hypergraph G will have no $\Theta_{k,t}^{B,m}$ and each of its edges will have multiplicity no more than $M_{k,t,r,m}$.

Proof. Let's first fix the notation we'll use in the rest of the proof. We will call the edges of H as hyperedges and the edges of G just as edges. We'll use h_i to denote the hyperedges, and e_i to denote edges. The edge set of H and G will be E(H) and E(G), respectively, and their (common) vertex set will be denoted as V. For every $h \in E(H)$, the m-set chosen from h will be denoted as e(h).

First note that *G* is $\Theta_{k,t}^{B,m}$ -free. Indeed, for any Berge path in *G*, say $v_0e_1v_1\cdots e_kv_k$, there is a corresponding Berge path in $H v_0h_1...h_kv_k$ where h_i is the hyperedge such that $e_i = e(h_i)$. Therefore, a $\Theta_{k,t}^{B,m}$ in *G* would imply that there is a $\Theta_{k,t}^{B,r}$ in *H*.

Now assume the lemma is not true and that there is an edge in G which was chosen more than $M := M_{k,t,r,m}$ times. It suffices for us to construct a $\Theta_{k,t}^{B,m}$ in G, which will give us a $\Theta_{k,t}^{B,r}$ in H and results in a contradiction. Let e be an edge in G with some multiplicity at least M+1, and $x, y \in e$. We will construct paths from x to y starting with a path of length 2, with the goal to construct t paths of length k, yielding a contradiction. Consider the last hyperedge which contributed to the multiplicity of e, call it h. Since every time we select an m-set from a hyperedge in H, we select the one with the least multiplicity, we know that every other m-set in h must have multiplicity at least M. Since r > 2, we know there exists $v_1 \in h \setminus \{x, y\}$. Therefore there exists a hyperedge $h_1 \in E(H)$ such that $\{x, v_1\} \subseteq e(h_1)$. On the other hand, note that $v_1, y \in h \setminus \{x\}$, thus there exists $h'_1 \in E(H) \setminus \{h_1\}$ such that $\{v_1, y\} \subseteq e(h'_1)$. This gives us a length 2 path from x to y, namely x, v_1, y connected by h_1 and h'_1 . Note that $e(h_1)$ and $e(h'_1)$ also have multiplicity at least M, and e(h) is not part of the path.

To extend this result, we prove the following claim, from which the construction of a $\Theta_{k,t}^B$ in *G* follows easily follows easily.

Claim 2.2.1. Let $S \subset V$ be a "forbidden set", and $M = \sum_{j=1}^{k-1} {\binom{|S|+jm-m}{r-m}} + k + 1$. For $i \leq k-2$, suppose we have vertices $x, v_1, v_2, \cdots, v_i, y$ and edges e_1, \cdots, e_{i+1} that form a path in G in the order given, and the last edge $e_{i+1} \in G$ with $v_i, y \in e_{i+1}$ has multiplicity at least $M - \sum_{j=1}^{i} {\binom{|S|+jm-m}{r-m}} - (i-1) + 1$. Let $S' = S \cup e_1 \cup \cdots \cup e_{i+1}$. Then we can find $v_{i+1} \notin S'$ such that $x, v_1, v_2, \cdots, v_i, v_{i+1}, y$ forms a path, and the last edge in this new path containing v_{i+1}, y has multiplicity at least $M - \sum_{j=1}^{i+1} {\binom{|S|+jm-m}{r-m}} - i + 1$.

Proof of Claim 2.2.1. To find the vertex v_{i+1} , first we will find a new hyperedge $h' \in E(H)$ such that $e_{i+1} \subset h'$ and $h' \not\subseteq S'$. Note that the set S' has cardinality less than |S| + (i+1)m as we included all i+1 edges in the already existing path. We would like to make sure the next edge we choose for the new path is different from all previous edges. The number of hyperedges h such that $e(h) = e_{i+1}$ and $h \subset S'$ is at most $\binom{|S|+(i+1)m-m}{r-m}$. Now we let $E' = \{h \in E(H) \mid e(h) = e_{i+1}, h \not\subseteq S'\}$, then we have

$$|E'| \ge |\{h \in E(H) \mid e(h) = e_{i+1}\}| - \binom{|S| + (i+1)m - m}{r - m}$$
$$\ge M - \sum_{j=1}^{i+1} \binom{|S| + jm - m}{r - m} - (i-1) + 1 \ge 5.$$

Now we pick h' to be the last edge (in the original ordering) of E', which then implies that every *m*-set in h', besides e_{i+1} , has multiplicity at least $M - \sum_{j=1}^{i+1} {\binom{|S|+jm-m}{r-m}} - i + 1$. Since $h' \not\subseteq S'$, we can find $v_{i+1} \in h'$ such that $v_{i+1} \notin S'$. Now we choose two *m*-sets from h', namely e'_{i+1} and e_{i+2} where $v_i, v_{i+1} \in e'_{i+1}$ and $v_{i+1}, y \in e_{i+2}$. This gives us a length i+2 path from *x* to *y*, namely the path $x, v_1, \dots, v_i, v_{i+1}, y$ where the last two edges are e'_{i+1} and e_{i+2} , and the previous edges are the same as in the old length i+1 path. The last edge, e_{i+2} , has multiplicity at least $M - \sum_{j=1}^{i+1} {\binom{|S|+jm-m}{r-m}} - i + 1$, as desired. Note that the edge e(h') is not part of the path, and as we discard the edge e_{i+1} in the original path, we add in two new edges, namely e'_{i+1} and e_{i+2} .

With this claim, we can now build the first path from x to y by induction on k. The base case when k = 2 is already constructed before the statement of the claim. In the induction step, when we already have a length k - 1 path, apply the claim with $S = \emptyset$

and we obtain a length k path, as desired. Note that if all we need is just one path, then we just need the multiplicity of the edge e to be at least $M_{k,1,r,m}$.

To build t paths, we will construct each path separately. Assume that we have built i-1 paths. To build the *i*th path, the forbidden set S would be the union of all edges in previous paths. This ensures that the paths we are building are vertex-independent and edge-distinct. It then follows that to build t vertex-independent paths from x to y, it suffices for us to have that the multiplicity of the edge e to be at least $M_{k,t,r,m}$. This gives us a $\Theta_{k,t}^{B,m}$ in the graph G. If we then choose hyperedges in H that gives the edges in this $\Theta_{k,t}^{B,m}$ in G have multiplicity at least 1, and these hyperedges are guaranteed to be distinct since for any two hyperedges chosen, say h_1 and h_2 , $e(h_1) \neq e(h_2)$, and therefore $h_1 \neq h_2$. This leads to a contradiction. Therefore all edges in G must have multiplicity at most $M_{k,t,r,m}$.

By Theorem 2.1.1, there exists a constant $c_{k,t}$ such that $ex(n, \Theta_{k,t}) \leq c_{k,t}n^{1+\frac{1}{k}}$. Now assume that a *r*-uniform hypergraph has more than $M_{k,t,r,2}c_{k,t}n^{1+\frac{1}{k}}$ edges. Then if we reduce this hypergraph into a (2-uniform) graph with the scheme described above, we will either have a graph with more than $c_{k,t}n^{1+\frac{1}{k}}$ edges, or it will have at least one edge with more than $M_{k,t,r,2}$ multiplicity. Both cases imply that there is a $\Theta_{k,t}$ in the reduced graph, which leads to a $\Theta_{k,t}^B$ in the original hypergraph. This completes the proof of Theorem 2.1.2.

2.3 Preliminaries for Theorem 2.1.4

To prove Theorem 2.1.4, given any natural number $k \ge 2$, we need to construct a $\Theta_{k,t}^B$ -free *r*-uniform hypergraph with $\Omega(n^{1+\frac{1}{k}})$ edges where *t* is a large enough constant depending only on *k* and *r*. Bukh [6], building from work of Blagojević, Bukh and Karasev [2], found an elegant random algebraic construction for Turán type problems, and we will use this method to construct our hypergraphs. This method has recently been used with success in both the graph [7,9,14] and hypergraph setting [39].

Let k and r be fixed, and for q a prime power let \mathbb{F}_q be the finite field of order q. We will work with polynomials in kr variables over \mathbb{F}_q . Let P_d be the set of such polynomials of degree at most d. That is, P_d consists of linear combinations of monomials $\prod_{i=1}^{kr} x_i^{\alpha_i}$ where $\sum \alpha_i \leq d$. For the remainder of this chapter we will use the term *random polynomial* to denote a polynomial chosen uniformly at random from P_d . Note that the distribution of random polynomials is equivalent to choosing the coefficient of each monomial $\prod_{i=1}^{kr} x_i^{\alpha_i}$ independently and uniformly from \mathbb{F}_q . We will need to know the probability that a random polynomial vanishes on a fixed set of points. In particular, if we fix one point, since the constant term of a random polynomial is chosen uniformly from \mathbb{F}_q , we have the following lemma.

Lemma 2.3.1. Let t be a natural number. If f is a random polynomial of degree d in t variables, then, for any fixed $x \in \mathbb{F}_a^t$,

$$\mathbb{P}[f(x) = 0] = \frac{1}{q}.$$

The next lemma, proved as Lemma 2.3 in [7] and Lemma 2 in [14] extends the conclusion of Lemma 2.3.1.

Lemma 2.3.2. Let t be a natural number. Assume x_1, \dots, x_z are z distinct points in \mathbb{F}_q^t . Suppose $q > {z \choose 2}$ and $d \ge z - 1$. Then if f is a random polynomial in t variables of degree d,

$$\mathbb{P}[f(x_i) = 0 \text{ for all } i = 1, \cdots, z] = \frac{1}{q^z}.$$

We now define the graphs that we will be interested in. Let $N = q^k$. We will construct an *r*-partite *r*-uniform hypergraph on rN vertices as follows. Let V_1, \ldots, V_r be the partite sets, each a distinct copy of \mathbb{F}_q^k . Choose $f_1, f_2, \cdots, f_{k(r-1)-1} : \mathbb{F}_q^{kr} \to \mathbb{F}_q$ to be random polynomials of degree at most d := k(2k + 1), chosen independently. For v_1, \ldots, v_r with $v_i \in V_i$, we declare (v_1, \cdots, v_r) to be an edge if and only if

$$f_1(v_1, v_2, \cdots, v_r) = f_2(v_1, v_2, \cdots, v_r) = \cdots = f_{k(r-1)-1}(v_1, v_2, \cdots, v_r) = 0.$$

We use the term *random polynomial graph* to describe the distribution of hypergraphs obtained this way. Since these polynomials are chosen independently, we know from Lemma 2.3.1 that the probability of a given hyperedge is in a random polynomial graph is $q^{1-k(r-1)}$. The total number of possible hyperedges is $N^r = q^{kr}$. Therefore the expected number of hyperedges is $q^{kr}q^{1-k(r-1)} = q^{k+1} = N^{1+1/k}$.

We will be interested in subgraphs that appear in a random polynomial graph. Since hyperedges appear when a system of polynomials vanishes, we will describe subgraphs as varieties. Let $\overline{\mathbb{F}}_q$ be the algebraic closure of \mathbb{F}_q and let *t* be a natural number. A *variety* over $\overline{\mathbb{F}}_q$ is a set of the form:

$$W = \{x \in \overline{\mathbb{F}}_{q}^{t} : f_{1}(x) = f_{2}(x) = \dots = f_{s}(x) = 0\}$$

for a collection of polynomials $f_1, \dots, f_s : \overline{\mathbb{F}}_q^t \to \overline{\mathbb{F}}_q$. In other words, a variety is the set of common roots of a set of polynomials. We say W is defined over \mathbb{F}_q if the coefficients

of these polynomials are from \mathbb{F}_q and write $W(\mathbb{F}_q) = W \cap \mathbb{F}_q^t$. We say W has complexity at most M if s, t and the maximum degree of the polynomials are all bounded by M. We will be very interested in how many points can be on a variety, and we will use the following theorem, proved by Bukh and Conlon ([7] Lemma 2.7) using tools from algebraic geometry.

Theorem 2.3.3. Let q be sufficiently large and suppose W and D are varieties over $\overline{\mathbb{F}}_q$ of complexity at most M which are defined over \mathbb{F}_q . Then one of the following holds:

- $|W(\mathbb{F}_q) \setminus D(\mathbb{F}_q)| \leq c_M$, where c_M depends only on M, or
- $|W(\mathbb{F}_q) \setminus D(\mathbb{F}_q)| \ge q/2.$

Using these tools, we will show in the following section that with high probability we may modify a random polynomial graph to obtain a $\Theta_{k,t}^B$ -free hypergraph with $\Omega(n^{1+1/k})$ edges.

2.4 Proof of Theorem 4

In this section we show that there exists an n vertex $\Theta_{k,t}^B$ -free hypergraph with $\Omega_{k,r}(n^{1+1/k})$ hyperedges. Our proof is an adaption of [14] to the hypergraph setting. Let q be a sufficiently large prime power and let G be a random polynomial graph on $n = rN = rq^k$ vertices defined in the previous section. As noted before the expected number of hyperedges in G is

$$N^{1+1/k} = \Omega(n^{1+1/k}).$$
(2.1)

We are interested in $\Theta_{k,t}^B$ as a subgraph of G, and so we will be interested in Berge paths between vertices. Suppose now that x and y are two fixed vertices in G and let S be the set of Berge-paths with length k between them. We will be interested in the moments of the random variable |S|. Let m be fixed and note that $|S|^m$ is the number of ordered collections of m Berge paths of length k between x and y. These paths can be overlapping or identical, and the total number of hyperedges in any collection of m paths is at most km. The probability that any hyperedge is in G is $q^{1-k(r-1)}$ because the polynomials are chosen independently, and so since q is sufficiently large, Lemma 2.3.2 implies that for $z - 1 \le d$ the probability of any particular collection with z hyperedges is in G is $q^{z(1-k(r-1))}$. Now if we denote $P_{m,z}$ as the number of ordered collections of m paths between x and y such that their union has z hyperedges in total, we have:

$$\mathbb{E}[|S|^{m}] = \sum_{z=1}^{km} P_{m,z} q^{z-zk(r-1)},$$

as long as $km \leq d$. Now we shall estimate $P_{m,z}$ by estimating the maximum number of vertices there can be in any particular collection with z hyperedges. Namely, if a collection has hyperedges h_1, \dots, h_z , then we want to estimate $\max |\bigcup_{i \in [z]} h_i|$.

Claim 2.4.1. If the union of m Berge-paths from x to y, each of length k, has z edges in their union $\{h_i\}_{i=1}^{z}$, then

$$\left| \bigcup_{i=1}^{z} h_i \setminus \{x, y\} \right| \le \frac{zk(r-1) - z}{k}.$$

Proof. Given such a collection, assume $| \cup h_i \setminus \{x, y\}| = n_0$. Let P_1, \ldots, P_m be the set of paths. We will think about P_1, \ldots, P_m as both a set of vertices and as a set of edges, and we let n_i and z_i be the number of vertices and edges respectively in $P_i \setminus (P_1 \cup \cdots \cup P_{i-1})$. Let $z'_i = z_i$ if $n_i > 0$ and $z'_i = 0$ if $n_i = 0$.

If $n_i > 0$, we first give a lower bound on z_i . If we consider the edges in $P_i \setminus (P_1 \cup \cdots \cup P_{i-1})$, this set of edges will form some number of disjoint Berge paths. Since consecutive edges in a Berge path must overlap, then in each of these smaller Berge paths, at least 2 vertices will have already been seen in $P_1 \cup \cdots \cup P_{i-1}$. Then if one of the smaller Berge paths has a edges, there will be at most a(r-1) - 1 vertices contained in these edges that were not in $P_1 \cup \cdots \cup P_{i-1}$. Since z_i is the sum of the number of edges in each of the smaller Berge paths, we have $z_i \geq \frac{n_i+1}{r-1}$. Let r' be the number of n_i which are greater than 0. Let

$$z' = \sum_{i=1}^{m} z'_i \le r'k,$$

and so $r' \geq \frac{z'}{k}$. On the other hand,

$$z' \ge \sum_{i:n_i>0} \frac{n_i+1}{r-1} = \frac{n_0+r'}{r-1} \ge \frac{n_0}{r-1} + \frac{z'}{k(r-1)}$$

This implies

$$n_0 \le \frac{z'k(r-1) - z'}{k}.$$

Since $z' \leq z$ the result follows.

We can now bound $P_{m,z}$.

Claim 2.4.2. Assume that $km \leq d$. Then $P_{m,z} = O_{k,r}(q^{zk(r-1)-z})$

Proof. We want to count the number of all possible ordered collections of m Berge-paths between x and y with z edges in their union. Let $V = \frac{zk(r-1)-z}{k}$ be the upper bound on

the number of vertices, then we first choose the vertices that will be in the collection. There are less than $(rN)^V$ number of ways to do this. Then we choose z hyperedges, each having r vertices. The number such choices is bounded by $\binom{V}{r}^z$. Last we choose 2 vertices from each hyperedge to make up the ordered set of core vertices. This is bounded by r^{2z} , which eventually gives us

$$r^{V}\binom{V}{r}^{z}r^{2z}N^{V} = O_{k,r}(N^{V}) = O_{k,r}(q^{zk(r-1)-z}),$$

where the implied constant uses that m is bounded above by a constant depending on k. Note that every collection of m paths is counted by this method in at least one way, giving the upper bound on $P_{m,z}$.

Thus, when $m \le 2k + 1$ we have $z \le km \le d$ (since d = k(2k + 1)) and we may apply Lemma 2.3.2 to find

$$\mathbb{E}[|S|^{m}] = \sum_{z=1}^{km} P_{m,z} q^{z-zk(r-1)} \le km C_{k,r} := C.$$
(2.2)

Where $C_{k,r}$ is a constant dependent on k and r, and C is used to simplify our notation.

Now if we want to apply tools from algebraic geometry, we need to write S as a variety. However, this cannot be done directly since there is no fixed set of polynomials whose set of common roots is exactly S. Therefore, we use the following analysis to bound |S|. Any path in S is a sequence of core vertices and edges $(x, h_1, v_1, h_2, \dots, v_{k-1}, h_k, y)$. We may partition the set of paths into which partite set each v_i is in. That is, S can be partitioned into disjoint sets depending on which partite sets each core vertex belongs to. Namely, we can let $S_{t_1,\dots,t_{k-1}}$ denote the set of paths from x to y such that the *i*th core vertex v_i belongs to V_{t_i} .

Now if we let σ denote any length k-1 tuple from $[r]^{k-1}$, then we have

$$S = \bigcup_{\sigma \in [r]^{k-1}} S_{\sigma},$$

and this is a disjoint union.

Fix any arbitrary S_{σ} . For notation, we denote the core vertices in an arbitrary path as v_1, \dots, v_{k-1} and the non core vertices in hyperedge h_i as w_1^i, \dots, w_{r-2}^i . We also need to make sure that the non core vertices are ordered based on their partite sets. In other words, if $w_{j_1}^i \in V_{t_1}$ and $w_{j_2}^i \in V_{t_2}$ where $j_1 < j_2$, then $t_1 \le t_2$.

Given such a sequence $p = (v_1, \dots, v_{k-1}, w_1^1, \dots, w_{r-2}^1, \dots, w_1^k, \dots, w_{r-2}^k)$ which denotes a path (that is, p is a vector ordered with the core vertices first and the non

core vertices after), let the polynomials $f_{i,1}, \dots, f_{i,k}$ be extensions to the polynomial f_i . Namely,

$$f_{i,1}(p) = f_i(x, v_1, w_1^1, \cdots, w_{r-2}^1),$$

$$f_{i,2}(p) = f_i(v_1, v_2, w_1^2, \cdots, w_{r-2}^2),$$

$$\cdots$$

$$f_{i,k}(p) = f_i(v_{k-1}, y, w_1^k, \cdots, w_{r-2}^k),$$

where the inputs to each f_i are reordered according to which partite sets the vertices are in. For example, if $x \in V_2$, $v_1 \in V_3$, then

$$f_{i,1}(p) = f_i(w_1^1, x, v_1, w_2^1, \cdots, w_{r-2}^1)$$

Now we may define the variety T_{σ} as

$$\{p \in \mathbb{F}_q^{k(r-2)+k-1} : f_{i,1}(p) = \dots = f_{i,k}(p) = 0 \text{ for all } i \text{ in } [k(r-1)-1]\},\$$

where $p \in \mathbb{F}_q^{k(r-2)+k-1}$ runs over sequences

$$(v_1, \cdots, v_{k-1}, w_1^1, \cdots, w_{r-2}^1, \cdots, w_1^k, \cdots, w_{r-2}^k).$$

With this restriction on ordering, we see that $S_{\sigma} \subseteq T_{\sigma}(\mathbb{F}_q)$. Note that T_{σ} contains all of the paths in S_{σ} , but may also contain walks that are not paths.

If $T_{\sigma}(\mathbb{F}_q)$ contains a degenerate walk $x, v_1, v_2, \dots, v_{k-1}, y$, then one of the following three conditions must be true: $x = v_b$ for some $b \in [k-1]$, $v_a = y$ for some $a \in [k-1]$, or $v_a = v_b$ for some $a \neq b \in [k-1]$. Therefore we can consider the collections of sets:

$$W_{0,b} = T_{\sigma} \cap \{v_1, \cdots, v_k, w_1^1, \cdots, w_{r-2}^1, \cdots, w_1^k, \cdots, w_{r-2}^k : x = v_b\},\$$

$$W_{a,b} = T_{\sigma} \cap \{v_1, \cdots, v_k, w_1^1, \cdots, w_{r-2}^1, \cdots, w_1^k, \cdots, w_{r-2}^k : v_a = v_b\},\$$

$$W_{a,0} = T_{\sigma} \cap \{v_1, \cdots, v_k, w_1^1, \cdots, w_{r-2}^1, \cdots, w_1^k, \cdots, w_{r-2}^k : v_a = y\}.$$

Each of these sets is also a variety with complexity bounded in terms of k and r. Let W be the union of all of the $W_{0,b}$, $W_{a,b}$, and $W_{a,0}$ where the union is over all a and b with $1 \le a < b \le k - 1$. If X and Y are varieties with complexity bounded in terms of k and r, then we claim that $X \cup Y$ is also a variety with complexity bounded in terms of k and r. To see this, if X is the set of points where the set of polynomials $\{f_i\}$ vanish and Y is the set of points where the polynomials $\{g_j\}$ vanish, then $X \cup Y$ is exactly the set of points where the polynomials $\{f_ig_j\}$ vanish. Since W is the union of $O(k^2)$ varieties, we have that W is also a variety with complexity bounded in terms of k and r. Since

 $S_{\sigma} = T_{\sigma} \setminus W$, we may apply Theorem 2.3.3 to S_{σ} (for algebraic geometry background see [48] and for a similar discussion of applying Theorem 2.3.3 see Section 3 of [14]).

Now if we put everything together, we see that there exists a constant c_{σ} , dependent on k and r, such that either $|S_{\sigma}| \leq c_{\sigma}$ or $|S_{\sigma}| \geq \frac{q}{2}$. This conclusion holds true for any arbitrary $\sigma \in [r]^{k-1}$. Looking at S, we see that either $|S| = \sum_{\sigma \in [r]^{k-1}} |S_{\sigma}| \leq C_{k,r}$ for some constant $C_{k,r}$ dependent on k and r, or $|S| > C_{k,r}$, which implies that there exists σ such that $|S_{\sigma}| > c_{\sigma}$ and therefore $|S| \geq |S_{\sigma}| > \frac{q}{2}$. Now by (2.2) and Markov's inequality, for $m \leq 2k + 1$ we have

$$\mathbb{P}[|S| > C_{k,r}] = \mathbb{P}[|S| > \frac{q}{2}] = \mathbb{P}[|S|^m > (q/2)^m] \le \frac{C}{(q/2)^m}.$$

Call a pair of vertices (x, y) bad if it has more than $C_{k,r}$ length k paths between them. If B is the random variable denoting the number of bad pairs, then we have

$$\mathbb{E}(B) \le n^2 \times \frac{C}{(q/2)^m} = O_{k,r}(q^{2k-m}) = O_{k,r}\left(\frac{1}{q}\right),$$

when we take m = 2k + 1. Therefore by Markov's Inequality, $\mathbb{P}[B \ge 1] \to 0$ as $n \to \infty$. Now let X be the number of edges, then by (2.1) the expected number of edges in G is $N^{1+1/k}$. The variance is

$$\mathbb{E}[X^2] - \mathbb{E}[X]^2 = \sum_{i,j} \mathbb{E}[H_i H_j] - \mathbb{E}[H_i] \mathbb{E}[H_j] = N^r (q^{1-k(r-1)} - q^{2-2k(r-1)}) \le \mathbb{E}[X]$$

where H_i are indicator random variables for hyperedges and for $i \neq j$, $\mathbb{E}[H_iH_j] = \mathbb{E}[H_i]\mathbb{E}[H_j]$ because a pair of events that 2 edges appear or not are independent due to Lemma 2.3.2. By Chebyshev's Inequality,

$$\mathbb{P}[|X - N^{1+1/k}| \ge \frac{1}{2}N^{1+1/k}] \le \frac{4}{N^{1+1/k}}$$

which goes to zero. Therefore with high probability this hypergraph on rN vertices has $\Omega_{k,r}(N^{1+1/k})$ edges, and it contains no $\Theta^B_{k,C_{k,r}+1}$. This completes the proof of Theorem 2.1.4.

2.5 Conclusion

In this chapter we showed that for fixed k, r, t there is a constant $c_{k,r,t}$ such that

$$\operatorname{ex}_{r}(n, \Theta_{k,t}^{B}) \leq c_{r,k,t} n^{1+\frac{1}{k}},$$

and that this order of magnitude is correct when t is large enough relative to r and k. That is, for fixed k and r there is a constant $c_{k,r}$ such that

$$\operatorname{ex}_{r}(n, \Theta^{B}_{k, c_{k,r}}) = \Omega_{k,r}(n^{1+\frac{1}{k}}).$$

We end with some open questions. First, it would be interesting to determine the dependence on r. Even when t = 2 and $k \in \{2,3,5\}$ this dependence is unknown. For example, it is known that $ex_r(n, C_4^B) = \Theta(n^{3/2})$ when $2 \le r \le 6$, but the order of magnitude is unknown for $r \ge 7$ (c.f. [27]). It would also be interesting to determine how the multiplicative constant of $ex_r(n, \Theta_{k,t}^B)$ depends on t when $n \gg t \gg 1$ and r and k are fixed. Finally, in this chapter we worked with the least restrictive definition of paths between vertices in hypergraphs. One could forbid only certain types of paths and demand that no pair of vertices have more than t of these paths between them.

Chapter 3

Improving the Constant

3.1 Introduction

In this chapter, we present an improvement to the constant factor in the bound $ex(n, C_{2k}) = O_k(n^{1+1/k})$. To discuss methods that lead to upper bounds on $ex(n, C_{2k})$, we first show a simple derivation of $ex(n, \{C_3, C_4, \dots, C_{2k}\}) \leq cn^{1+1/k}$ for some constant c. Consider a graph containing $\Theta(n^{1+1/k})$ edges with girth at least 2k + 1, and reduce it to a graph with minimum degree $\Theta(n^{1/k})$. Fix arbitrary vertex v, we start a Breadth-First Search (BFS) at v and observe that for the first k levels of the breadth-fist search tree, every level must expand by a factor of $\Theta(n^{1/k})$ compared to the previous level. In particular, no two vertices with depth less than k can have common neighbors with greater depth. Since the kth level cannot have more than n vertices, the bound follows. We present this derivation since the best upper bounds on $ex(n, C_{2k})$ are, in essence, all established using this same approach. As we will see shortly, employing this method imposes fundamental limitations to the results derivable.

The first important upper bound on $ex(n, C_{2k})$ was proved by Bondy and Simonovits [4] in 1974, where they showed $ex(n, C_{2k}) \leq 20kn^{1+1/k}$. This result is subsequently improved through a line of researches, most recently by Pikhurko [47] to $ex(n, C_{2k}) \leq (k - 1)n^{1+1/k} + O_k(n)$ in 2010 and by Bukh and Jiang [8] to $ex(n, C_{2k}) \leq 80\sqrt{k}\log_e kn^{1+1/k} + O_k(n)$ in 2017. Our main contribution in this chapter is the following theorem.

Theorem 3.1.1. Fix k, let G be a n-vertex graph where $n \ge (20k)^{4k^3+2k^2}$. If

$$|E(G)| > 16\sqrt{5}\sqrt{k\log_e k} \cdot n^{1+1/k} + 8000k^4n^{1+(2k-1)/(2k^2)},$$

then G contains a copy of C_{2k} .

For the rest of this chapter, we will abbreviate \log_e as \log . Our approach is an improved version of Bukh–Jiang's approach, and therefore suffers the same limitation as all BFS arguments. More specifically, consider a bipartite graph G with bipartition V_1, V_2 such that $|V_1| = n, |V_2| = n/(k-1)$. The BFS argument for the girth problem can be exploited to show that $e(G) \leq c(k-1)^{-1/2}n^{1+1/k}$ if G has girth at least 2k+1 (For a more detailed argument, see [46]). Now if we duplicate each vertex in V_2 into k-1 copies, we obtain a graph on 2n vertices and $c\sqrt{k-1}n^{1+1/k}$ edges with no C_{2k} . Therefore, the best upper bound on $ex(n, C_{2k})$ derivable from the BFS argument is $c\sqrt{kn^{1+1/k}}$ for fixed constant c. To break the $O(\sqrt{kn^{1+1/k}})$ threshold would require a different approach.

Our result improves the best known bound for $ex(n, C_{2k})$ by a factor of $\sqrt{5 \log k}$, taking us one step closer to the limitation of the method. Before discussing the proof, we would like to point out the following facts about this chapter. This chapter is modified from Bukh–Jiang's manuscript. While we had made global modifications to Bukh–Jiang's methods, most improvements made are local. In particular, it is highly similar to their paper mathematically, with a few statements and minor proofs largely unmodified. The author's intentions in writing this chapter this way are to give a more intuitive delivery of Bukh–Jiang's methods, to present simplifications and improvements that lead to a better result, and to avoid confusing readers with different notations and proof structures that demonstrate the same ideas. Therefore, this chapter adopted the same notation with some unmodified definitions from Bukh–Jiang's paper, and reshaped the delivery structure and language to uncover the underlying ideas and intuition.

To begin our proof, in Section 3.2 we describe the graph structures used in this chapter. For a detailed discussion of how our methods relate and differ from Pikhurko's and Bukh–Jiang's work, please see Section 3.3.

3.2 Graph Reduction and Exploration

To employ the breadth-first search approach, we first process our graph to gain control over the degrees of vertices. Classically, the graph is reduced to have minimum degree $O_k(n^{1/k})$ at the expense of half of the edges. However, in Bukh–Jiang's approach and in our approach, control over maximum degree is also required. Bukh and Jiang modified the BFS process to avoid vertices of high degrees, while we make use of the following reduction lemma.

Reduction Lemma. Fix $\alpha \in (0,1)$, let $\gamma = (20/\alpha)^{-2/\alpha}$. Let $d_{\min}(G), d_{\max}(G)$ denote the minimum and maximum degree of a graph G, respectively. If a graph G on n vertices has at least $cn^{1+\alpha}$ edges, then it contains a subgraph G' such that $|V(G')| \geq c\gamma n^{\alpha/2}$, $|E(G)| \geq (c/4)v(G')^{1+\alpha}$, $d_{\min}(G') \geq (c/2)v(G')^{\alpha}$, and $d_{\max}(G')/d_{\min}(G') \leq 1/\gamma$.

An initial version of this lemma was first proved by Erdös and Simonovits [18], and various forms of this lemma occur in other works. Bukh and Jiang proved a slightly different version of this lemma in their addendum. By slightly modifying their proof, we obtain the above lemma. This proof is included in the Appendix for completeness.

With this structure in mind, our real result in this chapter is the following theorem.

Theorem 3.2.1. Fix $k \ge 4$, let $\Delta = \sqrt{k}(20k)^{2k}$, and let $d \ge \max(2\sqrt{5}\sqrt{k\log k}n^{1/k}, (20k)^{4k^2+2k})$. If G is a graph on n vertices such that $d_{\min}(G) \ge 2d + 5k^2$ and $d_{\max}(G) \le \Delta d$, then G contains a copy of C_{2k} .

Theorem 3.1.1 then follows from Reduction Lemma and Theorem 3.2.1.

Proof of Theorem 3.1.1. Assume a graph H on m vertices has more than $16\sqrt{10}\sqrt{k\log k} \cdot n^{1+1/k} + 8000k^4n^{1+(2k-1)/(2k^2)}$ edges, then we can find a bipartite subgraph H' with at least half of its edges. Using the Reduction Lemma, we find a subgraph G on $n \ge 4\sqrt{10}\sqrt{k\log k}\gamma m^{1/(2k)}$ vertices and at least $2\sqrt{10}\sqrt{k\log k}n^{1+1/k}$ edges, where $\gamma = (20k)^{-2k}$. Now we compute the minimum degree in G.

Let $c = e(H)/m^{1+1/k}$, we have that $c \ge 8\sqrt{10}\sqrt{k\log k} + 4000k^4/m^{1/(2k^2)}$, which implies that

$$d_{\min}(G) \ge \frac{c}{2} n^{1/k} \ge 4\sqrt{10}\sqrt{k\log k} n^{1/k} + \frac{2000k^4(\gamma m^{1/(2k)})^{1/k}}{m^{1/(2k^2)}} \ge 4\sqrt{10}\sqrt{k\log k} n^{1/k} + 5k^2.$$

Now from Pikhurko's Result [47], we know that if $d_{\min}(G) \ge kn^{1/k}$, then G contains a C_{2k} . Therefore, Reduction Lemma implies $d_{\max}(G) \le (20k)^{2k}kn^{1/k}$. Let $d = 2\sqrt{10}\sqrt{k\log kn^{1/k}}$, $\Delta = \sqrt{k}(20k)^{2k}$. Theorem 3.2.1 completes the proof.

To prove Theorem 3.2.1, we elaborate further on the graph structures. Let G be a graph as in the statement of Theorem 3.2.1. Fix arbitrary vertex v of G and start a breadth-first search process at v. Let V_i be the set of vertices at minimum distance ifrom v for $i \in [k]$. We recall the following definition of a tri-layered graph, which is the basis of our discussions in Section 3.4, from Bukh–Jiang.

Definition 3.2.1 (Bukh–Jiang [8]). A graph G is called trilayered if its vertex set can be partitioned into V_1, V_2, V_3 such that all edges in G are between V_1, V_2 or between V_2, V_3 . For arbitrary G, we use $G[V_1, V_2, V_3]$ to denote the induced trilayered graph of G on V_1, V_2 and V_3 . For $A, B, C, D \in \mathbb{R}$, we say that a trilayered graph has minimum degree [A : B, C : D] if the minimum degree from V_1 to V_2 , V_2 to V_1 , V_2 to V_3 and V_3 to V_2 are at least A, B, C, D, respectively.

The last ingredient we need is the following definition of a Θ -graph, which is at the core of all our future discussions.

Definition 3.2.2. A Θ -graph is a cycle of length at least 2k with a chord. That is, an edge outside of the cycle connecting two vertices of the cycle.

The rest of the chapter is organized as follows: In Section 3.3, we recall several important results from Pikhurko and Bukh–Jiang, which prove the non-existence of Θ -graphs in the trilayered subgraphs formed by our breadth-first search exploration. In Section 3.4, which contains our main improvements in this chapter, we argue that if certain conditions hold, then a trilayered graph satisfies certain minimum degree condition must be present. We then embed a Θ -graph in such subgraphs, contradicting our result from Section 3.3. In Section 3.5, we show that either the aforementioned conditions hold, or the levels from exploration expand exponentially. Final computations then prove Theorem 3.2.1.

3.3 Results on Θ -Graphs

To argue for non-existence of Θ -graphs in our exploration, we recall results of Pikhurko.

Lemma 3.3.1 (Lemma 2.2 in [47]). Let $k \ge 3$. Any bipartite graph H of minimum degree at least k contains a Θ -graph.

Corollary 3.3.2. Let $k \ge 3$. Any bipartite graph H of average degree at least 2k contains a Θ -graph.

Lemma 3.3.3 (Claim 3.1 in [47]). Suppose G contains no C_{2k} . For $1 \le i \le k-1$, neither of $G[V_i]$ and $G[V_i, V_{i+1}]$ contains a bipartite Θ -graph.

Using these results, Pikhurko showed that every level must expand by a factor of roughly d/k compared to the previous level. The bound $ex(n, C_{2k}) \leq O(kn^{1+1/k})$ then followed. Bukh and Jiang improved on his method by analyzing three consecutive levels, proving a better expansion ratio among them. They employed the following technical definition, which generalized Θ -graphs to three levels, and proved the next lemma in conjunction.

Definition 3.3.1. Let G be a trilayered graph with layers V_1, V_2, V_3 . A Θ -graph T in G is well-placed if every vertex of T in V_2 is adjacent to some vertex of V_1 not in T.

Lemma 3.3.4 (Lemma 10 in [8]). Suppose G contains no C_{2k} . For $1 \le i \le k - 1$, the graph $G[V_{i-1}, V_i, V_{i+1}]$ contains no well-placed Θ -graphs.

Note that Lemma 3.3.4 is analogous to Lemma 3.3.3. To prove statements equivalent to Lemma 3.3.1 in trilayered graphs, Bukh and Jiang analyzed trilayered subgraphs with specific minimum degree structures. They first determined sufficient conditions for the existence of such trilayered graphs, then showed that if such subgraphs exist, a (well-placed) Θ -graph could be embedded inside. Finally, they argued that either the preceding conditions hold, or the levels must expand by an average factor of $O(\frac{d}{\sqrt{k} \log k})$. Their result followed.

In this chapter, we follow the same proof structure. We improve on Bukh–Jiang's result by weakening the conditions required for minimum degree trilayered subgraphs to be present, and presenting a better method to embed well-placed Θ -graphs in such subgraphs. These changes, presented in the following sections, lead to our $O(\sqrt{\log k})$ improvement on the best-known upper bound for $ex(n, C_{2k})$.

3.4 Search for Θ -graphs

In this section, we present the central arguments of this chapter. Our results are summarized in the following lemma, which states sufficient conditions for the existence of (well-placed) Θ -graphs.

Lemma 3.4.1. Let G be a trilayered graph with layers V_1 , V_2 , V_3 , such that $d_{\min}(G) \ge 2d + 5k^2$ and $d_{\max}(G) \le \Delta d$. If the following conditions hold:

$$d \cdot e(V_1, V_2) \ge 40k \log k |V_3|, \tag{3.1}$$

$$e(V_1, V_2) \ge 6k(\log k + 1)^2 (2\Delta k)^{2k-1} |V_1|,$$
(3.2)

$$e(V_1, V_2) \ge 20(\log k + 1)|V_2| \tag{3.3}$$

then either there is a Θ -graph in $G[V_1, V_2]$, or there is a well-placed Θ -graph in $G[V_1, V_2, V_3]$.

This lemma is an improvement over Lemma 6 in Bukh–Jiang. We removed two of the conditions and improved the last condition by a factor of $(\log k + 1)$.

To prove Lemma 3.4.1, the rest of this section is organized as follows: In Lemma 3.4.2, we show that given a trilayered graph formed by three consecutive levels in our BFS process, either we can find a trilayered subgraph with desired minimum degree structure, or we can find a trilayered subgraph with stronger constraints on its edges. This process can then be iterated — In Lemma 3.4.3, we prove that under the conditions stated in Lemma 3.4.1, Lemma 3.4.2 can be iterated to show the existence of a desired trilayered subgraphs. Finally, in Lemma 3.4.4, we show that a (well-placed) Θ -graph can be embedded in such subgraphs, which completes the proof.

Without further delay, we now quote the following result, which is Lemma 7 in Bukh–Jiang.

Lemma 3.4.2. Let a, A, B, C, D be positive real numbers. Suppose G is a trilayered graph with layers V_1 , V_2 , V_3 and the degree of every vertex in V_2 is at least $d + 4k^2 + C$. Assume also that

$$a \cdot e(V_1, V_2) \ge (A + k + 1)|V_1| + B|V_2|. \tag{3.4}$$

Then one of the following holds:

- I) There is a Θ -graph in $G[V_1, V_2]$.
- II) There exist non-empty subsets $V'_1 \subset V_1$, $V'_2 \subset V_2$, $V'_3 \subset V_3$ such that the induced trilayered subgraph $G[V'_1, V'_2, V'_3]$ has minimum degree at least [A : B, C : D].
- III) There is a subset $\widetilde{V}_2 \subset V_2$ such that $e(V_1, \widetilde{V}_2) \ge (1-a)e(V_1, V_2)$, and $|\widetilde{V}_2| \le D|V_3|/d$.

A proof of this lemma, as presented in Bukh–Jiang, is included in the Appendix for completeness. Here the parameter *a* can be interpreted as the edge loss ratio. More specifically, in case (I) the proof of Lemma 3.4.1 is complete, and similarly in case (II) we are done by Lemma 3.4.4. In case (III), we found \tilde{V}_2 that shrinks proportionally compared to V_2 , while being adjacent to most edges between V_1 and V_2 . We can then apply Lemma 3.4.2 in $G[V_1, \tilde{V}_2, V_3]$, thereby iterating this process. In the end, we will either obtain a subset of vertices of V_2 with an overly high average degree, or lands in case (I) or (II). This procedure is done precisely in the following lemma, which is Lemma 8 in Bukh–Jiang.

Lemma 3.4.3. Let G be a trilayered graph with layers V_1, V_2, V_3 satisfying conditions (3.1), (3.2) and (3.3). Let C be a positive real number, such that the minimum degree from V_2 to V_3 is at least $d + 4k^2 + C$. Then one of the following holds:

- I) There is a Θ -graph in $G[V_1, V_2]$.
- II) There are non-empty subsets $V'_1 \subset V_1$, $V'_2 \subset V_2$, and $V'_3 \subset V_3$ such that the induced trilayered subgraph $G[V'_1, V'_2, V'_3]$ has minimum degree at least [A : B, C : D], where A, B are real numbers and C, D are integers. Moreover, $B \ge 5$, and

$$A \ge 2k(\Delta D)^{D-1},\tag{3.5}$$

$$(B-4)D \ge 2k. \tag{3.6}$$

Proof. This proof is an improved version of Bukh–Jiang's proof. Assume for the sake of contradiction that neither of the conclusions are true. We will first show that the conditions of Lemma 3.4.2 hold for a tuple of well defined A, B and D. Due to our assumptions, the only probable conclusion of Lemma 3.4.2 would be (III), which gives us $\tilde{V}_2 \subseteq V_2$. We then iterate this procedure for $t = \log k$ steps on \tilde{V}_2 and subsequent

subsets of V_2 . This process will generate a chain of sets $V_2^{(t)} \subseteq V_2^{(t-1)} \subseteq \cdots \subseteq V_2^{(1)} \subseteq V_2^{(0)} = V_2$. Finally, we will show a contradiction in $V_2^{(t)}$ to conclude the proof.

Let $a_i = \frac{1}{t-i+1}$, where *i* ranges from 0 to t-1. Let $V_2^{(0)} = V_2$. For $V_2^{(i)}$ that is well defined, set

$$d_{i} = e(V_{1}, V_{2}^{(i)}) / |V_{2}^{(i)}|,$$

$$A_{i} = a_{i}e(V_{1}, V_{2}^{(i)}) / 2|V_{1}| - k - 1,$$

$$B_{i} = a_{i}d_{i} / 4 + 5,$$

$$D_{i} = |\min(2k, 10k / a_{i}d_{i})|.$$

Here d_i is the average degree from $V_2^{(i)}$ to V_1 . We prove the following simple claim.

Claim 3.4.1. For all $i = 0, \dots, t - 1$, we have $d_i < 2k$.

Proof of Claim 3.4.1. If we have $|V_2^{(i)}| \leq |V_1|$, then

$$\frac{2e(V_1, V_2^{(i)})}{|V_1| + |V_2^{(i)}|} \ge \frac{e(V_1, V_2^{(i)})}{|V_1|} \ge \frac{1}{t+1} \frac{e(V_1, V_2)}{|V_1|} \ge 2k,$$

which then implies outcome (I) by Corollary 3.3.2. On the other hand, if $|V_2^{(i)}| \ge |V_1|$ and $d_i \ge 2k$, then

$$\frac{2e(V_1, V_2^{(i)})}{|V_1| + |V_2^{(i)}|} \ge \frac{e(V_1, V_2^{(i)})}{|V_2^{(i)}|} = d_i \ge 2k,$$

which again leads to outcome (I). Therefore $d_i < 2k$.

Now note that A_i , B_i , D_i satisfy constraints (3.5) and (3.6). Indeed, (3.6) follows as long as $d_i < 2k$, and for (3.5), we have by (3.2)

$$\begin{split} A_i &= a_i e(V_1, V_2^{(i)})/2 |V_1| - k - 1 \geq \frac{1}{2(t+1)^2} \frac{e(V_1, V_2)}{|V_2|} - k - 1 \\ &\stackrel{\textbf{(3.2)}}{\geq} 3k (2\Delta k)^{2k-1} - k - 1 \geq 2k (\Delta D_i)^{D_i - 1}. \end{split}$$

Therefore, if we apply Lemma 3.4.2, the only possible outcome is (III). The following claim is the key to our iteration process.

Claim 3.4.2. For $V_2^{(i)}$ that is well-defined, condition (3.4) of Lemma 3.4.2 hold with respect to the above defined a_i, A_i, B_i, C, D_i . Moreover, let $V_2^{(i+1)} \subseteq V_2^{(i)}$ be the set derived

from Lemma 3.4.2. We have the following invariants:

$$e(V_1, V_2^{(i+1)}) \ge (1 - a_i)e(V_1, V_2^{(i)}),$$
(3.7)

$$d_{i+1} \ge a_i d_i \frac{t-i}{t+1} \frac{d \cdot e(V_1, V_2)}{10k|V_3|}.$$
(3.8)

Proof of Claim 3.4.2. This proof will proceed by induction. We first show that condition (3.4) holds for i = 0.

$$(A_0 + k + 1)|V_1| + B_i|V_2| = \frac{3}{4}a_0e(V_1, V_2) + 5|V_2|$$

$$\stackrel{(3.3)}{\leq} \frac{3}{4}a_0e(V_1, V_2) + \frac{1}{4(t+1)}e(V_1, V_2) = a_0e(V_1, V_2).$$

Therefore, we can apply Lemma 3.4.2 to obtain $\widetilde{V}_2 \subseteq V_2$ as in outcome (III). Set $V_2^{(1)} = \widetilde{V}_2$. Invariant (3.7) then follows directly from the conclusions of Lemma 3.4.2. For (3.8), since $|V_2^{(1)}| \leq D_0|V_3|/d$, we have

$$d_1 = \frac{e(V_1, V_2^{(1)})}{|V_2^{(1)}|} \ge \frac{(1 - a_0)e(V_1, V_2)}{D_0|V_3|/d} \ge (1 - a_0)a_0d_0\frac{d \cdot e(V_1, V_2)}{10k|V_3|}.$$

This completes the proof for the base case. For induction, note that iterative application of (3.7) gives

$$e(V_1, V_2^{(i)}) \ge e(V_1, V_2) \prod_{j=0}^{i-1} (1 - a_j) = \frac{t - i + 1}{t + 1} e(V_1, V_2).$$
 (3.9)

This inequality helps us show condition (3.4) again. Indeed,

$$\begin{split} (A_i + k + 1)|V_1| + B_i|V_2^{(i)}| &= \frac{3}{4}a_i e(V_1, V_2^{(i)}) + 5|V_2^{(i)}| \le \frac{3}{4}a_i e(V_1, V_2^{(i)}) + 5|V_2| \\ &\stackrel{(3.3)}{\le} \frac{3}{4}a_i e(V_1, V_2^{(i)}) + \frac{1}{4(t+1)}e(V_1, V_2) \\ &\stackrel{(3.9)}{\le} \frac{3}{4}a_i e(V_1, V_2^{(i)}) + \frac{t+1}{4(t+1)(t-i+1)}e(V_1, V_2^{(i)}) \\ &= \frac{3}{4}a_i e(V_1, V_2^{(i)}) + \frac{1}{4}a_i e(V_1, V_2^{(i)}) = a_i e(V_1, V_2^{(i)}). \end{split}$$

Therefore, by Lemma 3.4.2 again, there is a subset $V_2^{(i+1)} \subset V_2^{(i)}$ satisfying (3.7), and

$$|V_2^{(i+1)}| \le D_i |V_3| / d$$

This implies

$$d_{i+1} = \frac{e(V_1, V_2^{(i+1)})}{|V_2^{(i+1)}|} \ge \frac{(1-a_i)e(V_1, V_2^{(i)})}{D_i|V_3|/d} \ge (1-a_i)a_i d_i \frac{d}{10k|V_3|}e(V_1, V_2^{(i)})$$
$$\ge (1-a_i)a_i d_i \frac{de(V_1, V_2)}{10k|V_3|} \prod_{j=0}^{i-1} (1-a_j) = a_i d_i \frac{t-i}{t+1} \frac{de(V_1, V_2)}{10k|V_3|}.$$

Therefore invariant (3.8) holds. This complete the proof of this claim.

Through iterative application of this claim, we obtain our desired chain of subsets $V_2^{(t)} \subseteq \cdots \subseteq V_2^{(1)} \subseteq V_2^{(0)}$. For simplicity of notation, let $F = \frac{d \cdot e(V_1, V_2)}{10k|V_3|}$. By (3.8), we have

$$d_{i} \geq d_{0} \cdot F^{i} \prod_{j=0}^{i-1} a_{j} \frac{t-j}{t+1} = d_{0} \cdot F^{i} \prod_{j=0}^{i-1} \frac{t-j}{t-j+1} \frac{1}{t+1}$$
$$= d_{0} \cdot \left(\frac{F}{t+1}\right)^{i} \frac{t-i+1}{t+1} \stackrel{(3.1)}{\geq} d_{0} \cdot 4^{i} \left(\frac{t}{t+1}\right)^{t} \frac{t-i+1}{t+1}$$
$$\geq d_{0} \cdot 4^{i} e^{-1} \frac{t-i+1}{t+1}.$$

Therefore we have

$$\frac{d_0}{d_i} \le \frac{e \cdot (t+1)}{4^i (t-i+1)}.$$
(3.10)

We now analyze the end results of our iteration, $V_2^{(t)}$ and d_t . Observe that $V_2^{(t)}$ preserves a good portion of the edges from V_2 to V_1 (invariant (3.7)), while having exponentially large average degree (equation (3.10)). Similar to Claim 3.4.1, we have $d_t < 2k$, which then implies

$$d_0 \le d_t e \cdot (t+1)/4^t < 2ke \cdot (t+1)/4^t < 20(t+1).$$

This contradicts condition (3.3). Therefore we conclude that the iteration must stop before *t* steps, resulting in either outcome (I) or outcome (II). \Box

Remark. The contraction rate of $V_2^{(i)}$ could be shown explicitly in the above proof. Specifically,

$$|V_2^{(i+1)}| \le D_i |V_3|/d \le \frac{1}{a_i d_i} \frac{10k|V_3|}{d} = \frac{d_0}{Fa_i d_i} |V_2| \stackrel{(3.1)}{\le} \frac{e \cdot (t+1)}{4^{i+1}t} |V_2|.$$

This bound confirms our intuition that $V_2^{(i)}$ shrinks exponentially.

We now come to the last piece of the puzzle: proving the existence of a Θ -graph. The following lemma, while following the same scheme as in Lemma 9 of Bukh–Jiang,

presents a different method to embed an arbitrarily long path under the assumption that no (well-placed) Θ -graphs exist. More details on such distinctions are discussed after the proof.

Lemma 3.4.4. Let G be a trilayered graph with layers V_1, V_2, V_3 and minimum degree at least [A : B, d + k : D], where A, B are real numbers and D is an integer. Suppose $B \ge 5$, and

$$A \ge 2k(\Delta D)^{D-1}, (B-4)D \ge 2k.$$
(3.11)

Assume that every vertex in V_2 has at most Δd neighbors in V_3 . Then there is a Θ -graph in $G[V_2, V_3]$, or there is a well-placed Θ -graph in G.

Proof. Assume that neither of the conclusions are true. In this proof, we will utilize this assumption to embed an arbitrarily long path P in G, contradicting the finiteness of the graph. P will have the form $v_0 \leftrightarrow v_1 \leftrightarrow \cdots \leftrightarrow v_l$, where $v_1, \cdots, v_l \in V_1$ and each pair v_i, v_{i+1} is connected by a path of length 2D alternating between V_2 and V_3 .

To utilize the assumption of no well-placed Θ -graph, we strengthen the statement by maintaining the following property while building the path:

Definition 3.4.1. A path P is called good if every vertex in $V_2 \cap P$ has at least one neighbor in $V_1 \setminus P$.

This property enables us to make arguments of the form "either the path could be extended, or we can find a well-placed Θ -graph", as we will see later in the proof.

We start our construction with a random vertex v_0 from V_1 . Inductively, assume that a good path $P = v_0 \iff v_1 \iff \cdots \iff v_{l-1}$ has been constructed, we wish to extend it to $v_0 \iff \cdots \iff v_l$. We make the following observations.

Claim 3.4.3. For all $i = 0, \dots, l-1$, v_i cannot have k or more neighbors in $V_2 \cap P$.

Proof of Claim 3.4.3. If v_i has at least k neighbors in $V_2 \cap P$, then we can follow the path and build a Θ -graph with a chord through v_i . This Θ -graph is well-placed since P is a good path.

Claim 3.4.4. Given a good path Q, let $u \in V_2 \cap Q$ be a vertex adjacent to the last vertex of Q (note that this last vertex can belong to either V_1 or V_3). Then u has less than $t = \lceil B/2 \rceil$ neighbors in $V_1 \cap Q$.

Proof of Claim 3.4.4. If u has neighbors v_{k_1}, \dots, v_{k_t} , where $k_1 < k_2 < \dots < k_t$, then the path $v_{k_2} \leftrightarrow u$ and the edge uv_{k_2} form a cycle of length at least

$$2D(t-2) + 2 \ge 2D(B/2 - 2) + 2 = D(B - 4) + 2 \ge 2k.$$

This cycle, together with the chord uv_{k_3} , forms a Θ -graph spanning over V_1, V_2, V_3 . Moreover, this Θ -graph is well-placed since Q is a good path, and u is adjacent to v_{k_1} which is not part of the Θ -graph. This contradicts our assumption.

Note that by Claim 3.4.3, there are at least A - k ways to extend v_{l-1} to another vertex in $V_2 \setminus P$, and Claim 3.4.4 ensures that all of these extensions are good. Denote $U_0 = N(v_{l-1}) \setminus P$, where $N(\cdot)$ is the usual notation for neighborhood. The following claim, which is the heart of our embedding scheme, states that a large portion of these good extensions in U_0 can be extended further inductively in a vertex-disjoint manner.

Claim 3.4.5. For $i = 0, 1, \dots, D-1$, there exist sets $U_i \subset V_2$ such that for each $u \in U_i$, there exists a path Q(u) from U_0 to u of length 2i that alternates between V_2 and V_3 . Moreover, Q(u) is a good extension of P, and for every pair $u, v \in U_i$, Q(u) and Q(v) are vertex disjoint. Furthermore,

$$|U_i| \ge -3k + A\left(\frac{1}{8(2k+1)\Delta}\right)^i \prod_{i=1}^{D-1} \frac{D-i}{i+1}.$$

Proof of Claim 3.4.5. We prove this claim by induction, where the base case with i = 0 is true as stated. Assume the claim is true for *i*, we want to find U_{i+1} by extending paths from U_i .

For arbitrary $u \in U_i$, let P_u denote the concatenation of paths P and Q(u). By similar argument as in Claim 3.4.3, we see that u cannot have more than k neighbors in $P_u \cap V_3$. Therefore, u has at least d neighbors in V_3 that does not land on P_u . These neighbors are our candidates for extending P_u , and we filter these candidates with the following procedure. Define three sets S_1 , F_1 and T, where S_1 , $F_1 \subset U_i$ and $T \subset V_3$. Intuitively, we want S_1 to be the set of vertices with successful extensions to V_3 , F_1 to be $U_i \setminus S_1$, and T to be the set of potential extensions from S_1 to V_3 . Set them to be empty initually, consider the following procedure.

We claim that when this procedure terminates, $|S_1| \ge |U_i|/2$. Indeed, if $|F_1| > |U_i|/2$, then $|T| < |S_1| \frac{d}{2k+1} < \frac{|U_i|}{2} \frac{d}{2k+1}$. Moreover, every vertex u in F_1 has at least d neighbors in $V_3 \setminus P_u$, which means at least $\frac{2kd}{2k+1}$ edges adjacent to u land in T. Therefore,

$$e(F_1, T) \ge |F_1| \frac{2kd}{2k+1} > \frac{2kd}{2k+1} \frac{|U_i|}{2},$$

Algorithm 1

- 1: Pick a vertex *u* randomly from $U_i \setminus (S_1 \cup F_1)$.
- 2: Let $M_u = (N(u) \cap V_3) \setminus (T \cup P_u)$. If $|M_u| \ge \frac{d}{2k+1}$, then randomly select $\frac{d}{2k+1}$ vertices in M_u to put into T, and denote these vertices as T_u . Put u into S_1 .
- 3: Otherwise, put $u \in F_1$ and move on to the next itration. Terminate this procedure if $S_1 \cup F_1 = U_i$.

which implies $e(F_1, T)/|T| > 2k$. By Lemma 3.3.1, there exists a Θ -graph in $G[V_2, V_3]$, which is a contradiction. Thus $|S_1| \ge |U_i|/2$.

We extend the previous notations to vertices in T. For $v \in T_u$ (as defined in Procedure 1), let Q(v) be the path Q(u)v, and $P_v = P_u v$. Note that the paths $\{Q(v)\}_{v \in T}$ are not necessarily pairwise vertex disjoint, since v could be on the path Q(w) for some $w \in U_i, w \neq u$. This issue will be resolved later. For now, we make the following observation concerning extending vertices in T back to V_2 .

Observation 3.4.6. For an arbitrary vertex $v \in T$, it has at least D-i neighbors in $V_2 \setminus P_v$ or in the last 2k vertices of P. To see that, suppose we call an edge vw where $w \in P_v \cap V_2$ long if the distance between v, w is at least 2k through the path P_v and short otherwise. If v has a long edge vw, then v cannot have any other neighbors in $P_v \cap V_2$, for otherwise there would be a well-placed Θ -graph. Moreover, since $|Q(v) \cap V_2| = i$, we see that v has at most i neighbors on Q(v). Our claim then follows.

Utilizing this observation, we will extend every vertex in S_1 greedily, while maintaining that all extensions land in different vertices in V_2 . As in procedure 1, we define sets $S_2, F_2 \subset S_1, D \subset V_2$, where S_2 denotes the set of vertices with successful 2-step extensions, and $F_2 = S_1 \setminus S_2$. D denotes the set of endpoints of successful extensions. We set them to be empty initially, and consider the following procedure.

Algorithm 2

- 1: Pick a vertex u arbitrarily from $S_1 \setminus (S_2 \cup F_2)$.
- 2: If there exists $v \in N(u) \cap T$ and $w \in N(v) \setminus (P_v \cup D)$, then we can successfully extend P_u to P_uvw . Put u into S_2 and put w into D.
- 3: If such vertices do not exist, put u into F_2 and move on to the next itration. Terminate this procedure if $S_2 \cup F_2 = S_1$.

Let $\epsilon = \frac{D-i}{4(2k+1)\Delta}$. We claim that when this procedure terminates, $|S_2| \ge \epsilon |S_1| - 2k$. To see that, we know every vertex $u \in F_2$ cannot be extended, which means all of its possible extensions land in D or the last 2k vertices of P. If $|F_2| > (1 - \epsilon)S_1 + 2k > \epsilon$ $(1 - \epsilon)S_1$, by Procedure 1 and Observation 3.4.6, the number of failed extension must be at least

$$|F_2| \cdot \frac{d}{2k+1} \cdot (D-i) \ge \frac{(D-i)(1-\epsilon)d}{2(2k+1)} |U_i|.$$

Since $|D| = |S_2|$, all these failed extensions must land in a set of size less than $\epsilon |S_1|$. The average degree on this set would then be at least

$$\frac{(D-i)(1-\epsilon)d}{2(2k+1)}|U_i|\cdot\frac{1}{\epsilon|U_i|} > 2(1-\epsilon)\Delta d > \Delta d,$$

which is a contradiction to the assumption that no vertices in V_2 has more than Δd neighbors in V_3 . Therefore we have at least $|S_2| \ge \epsilon |U_i| - 2k$ successful extensions.

The next step is to filter these extensions such that they are pairwise vertex disjoint. What we have constructed so far is a set Q of length 2i + 2 paths from U_0 to D such that if we choose any two paths p_1, p_2 from Q, their first 2i vertices would be disjoint, and their last two vertices would also be disjoint. Therefore every path could only overlap with at most 2i + 2 other paths in Q, which implies there exists a set of pairwise disjoint paths Q' such that $|Q'| \ge |Q|/(2i + 2)$. Let $U_{i+1} \subset D$ be the set of endpoints of these paths, we have

$$|U_{i+1}| = |Q'| \ge |Q|/(2i+2) = |S_2|/(2i+2) \ge \frac{\epsilon}{2(i+1)}|U_i| - 2k \ge \frac{D-i}{i+1}\frac{1}{8(2k+1)\Delta}|U_i| - 2k \ge \frac{D-i}{2(i+1)}|U_i| - 2k \ge \frac{D-i}{2(i+1)}|U_$$

which satisfies the stated bound. All of these extensions are good by Claim 3.4.4.

Now from condition (3.11), we see that U_{D-1} is non-empty. Let $Q = v_0 \iff \cdots \iff v_{l-1} \iff u$ be an arbitrary extension with $u \in U_{D-1}$. By Claim 3.4.4, $(N(u) \cap V_1) \setminus Q$ is non-empty. Let v_l be chosen arbitrarily from this set, and let the new path be Qv_l . We prove one last claim to finish the proof.

Claim 3.4.7. The path Qv_l is good.

Proof of Claim 3.4.7. We show that for any $w \in V_2 \cap Q$, w has at most 2t - 2 neighbors in $V_1 \cap Qv_l$. By Claim 3.4.4, w has fewer than t neighbors in $Q \cap V_1$ that precede w in Q. We want to apply the same argument to the reverseal of Qv_l .

Consider the sub-path $Q' = v_l \iff w$ of Q. Since Q is a good path, w can't have t or more neighbors in $V_1 \cap Q$. Therefore, assume w has neighbors $v_{k_1}, \cdots, v_{k_t} \in V_1 \cap Q'$, where $v_{k_t} = v_l$ and $k_1 < k_2 < \cdots < k_t$. Then the path $v_{k_{t-1}} \iff w$, together with the edges $wv_{k_{t-1}}$ forms a cycle of length at least 2k, with chords through v. This Θ -graph is well-placed since the path Q is good, and $v_{k_{t-1}} \iff w$ does not go through $v_{k_t} = v_l$, which

means vertices of this Θ -graph in V_2 can use v_l to satisfy the well-placed condition. We conclude that w must have less than t neighbors in $Q' \cap V_1$. Since 2t - 2 < B, the path Qv_l is good.

Therefore, we can construct an arbitrarily long path in G, which is a contradiction. We conclude that a (well-placed) Θ -graph must exist.

Remark. This result is stronger than Lemma 9 in Bukh–Jiang, in the sense that Bukh and Jiang showed how to embed one extension inductively, while we presented a method to embed multiple vertex-disjoint extensions simultaneously. We also note that $|U_i|$ can be made arbitrarily large by increasing A, which only affects the magnitude of n. Therefore, our methods embed many "parallel" paths concurrently.

Utilizing Lemma 3.4.2, 3.4.3 and 3.4.4, we now prove Lemma 3.4.1.

Proof of Lemma 3.4.1. Given a graph *G* satisfying the conditions in Lemma 3.4.1, we first apply Lemm 3.4.3 with C = d + k. If the lemma results in outcome (I), then our claim holds. If the lemma results in outcome (II), then we apply Lemma 3.4.4 on the resulting trilayered subgraph and our claim holds. This proves Lemma 3.4.1.

We now proceed to prove Theorem 3.2.1.

3.5 Proof of Theorem 3.2.1

In this section, we prove that under the conditions of Theorem 3.2.1, we have for all i

$$|V_{i+1}| \ge (d^2/20k\log k)|V_{i-1}|. \tag{3.12}$$

We introduce the following auxiliary conditions, which will be proved by induction on *i*.

$$e(V_i, V_{i+1}) \ge 2d|V_i|,$$
(3.13)

$$e(V_i, V_{i+1}) \le 2k|V_{i+1}|, \tag{3.14}$$

$$|V_{i+1}| \ge k^{-1} d|V_i|, \tag{3.15}$$

These inequalities hold for i = 0. Assuming the inductive hypothesis, we know that the minimum degree in the graph is at least $2d + 5k^2$. Therefore

$$e(V_i, V_{i+1}) \ge (2d + 5k^2)|V_i| - e(V_{i-1}, V_i) \stackrel{(3.14)}{\ge} (2d + 5k^2 - 2k)|V_i| \ge 2d|V_i|$$

This inequality implies that V_i has average degree at least 2d in $G[V_i, V_{i+1}]$. Moreover, if (3.14) is false, then V_{i+1} has average degree at least 2k in $G[V_i, V_{i+1}]$. By Corollary 3.3.2, this leads to a contradiction. Therefore (3.14) is true, and (3.15) is a consequence of (3.13) and (3.14). This completes the proof for the auxiliary claims.

We now move on to prove (3.12). Assume for the sake of contradiction that (3.12) is false, we will show that the conditions of Lemma 3.4.1 hold, which then leads to a contradiction.

Assume (3.1) is false. We have

$$2d^{2}|V_{i-1}| \stackrel{(3.13)}{\leq} de(V_{i-1}, V_{i}) \leq 40k \log k |V_{i+1}|, \\ |V_{i+1}| \geq \frac{d^{2}}{20k \log k} |V_{i-1}|.$$

This contradicts with the assumption that (3.12) is false.

(3.2) follows from the fact that $d \ge (20k)^{4k^2+2k}$. We have

$$6k(\log k + 1)^2 (2\Delta k)^{2k-1}, \le 6k^3 (2k^{3/2}(20k)^{2k}))^{2k-1}$$

$$\le (20k)^{4k^2 - 2k} \cdot 6k^3 \cdot (2k)^{3k} \le 2d \le e(V_i, V_{i+1})/|V_i|.$$

Finally, if (3.3) is false, we have

$$2d|V_{i-1}| \stackrel{(3.13)}{\leq} e(V_{i-1}, V_i) \leq 20(\log k + 1)|V_i| \stackrel{(3.15)}{\leq} 40\log k \frac{k}{d}|V_{i+1}|,$$
$$|V_{i+1}| \geq \frac{d^2}{20k\log k}|V_{i-1}|.$$

This again implies (3.12). Therefore (3.12) hold for all *i*.

We now conclude the proof of Theorem 3.2.1. If k is even, applying (3.12) k/2 times results in

$$|V_k| \ge \frac{d^k}{(20k\log k)^{k/2}}$$

If k is odd, applying (3.12) (k-1)/2 times results in

$$|V_k| \ge \frac{d^{k-1}}{(20k\log k)^{(k-1)/2}} |V_1| \ge \frac{d^k}{(20k\log k)^{k/2}}.$$

Since $|V_k| < n$, we must have $d < \sqrt{20k \log k} n^{1/k}$.

3.6 Potential Improvements

This section is dedicated to devoted readers who intend to improve Theorem 3.2.1 using our methods.

The idea of using Θ -graphs in the BFS approach originated from Pikhurko's work [47]. The most critical component of this combination is the embedding scheme of a Θ -graph in specific graph structures. For reference, Pikhurko utilized Lemma 3.3.1, while Bukh and Jiang and the author utilized different versions of Lemma 3.4.4. In essence, all three proofs are driven by their respective embedding methods. Therefore, if one intends to improve the upper bound on $ex(n, C_{2k})$ following this approach, one shall investigate potential structures and schemes to embed Θ -graphs.

We investigated the following structure in particular.

Definition 3.6.1. For $A, B, C, D \in \mathbb{R}$, we say that a trilayered graph G on vertex sets V_1, V_2, V_3 has degree [A : B, C : (2 : D)] if there exists a partition of V_2 into V_2^B and V_2^C , such that the minimum degree from V_1 to V_2, V_2^B to V_1, V_2 to V_3, V_3 to V_2^C , and V_3 to V_2^B are at least A, B, C, 2, D, respectively.

This definition is inspired by two observations. First of all, the proof of Lemma 3.4.2 found that if an [A : B, C : D] structure cannot be found, then V_2 can be partitioned into two sets \tilde{V}_2 and $V_2 \setminus \tilde{V}_2$, such that the former has high density with V_3 and the latter has high density with V_1 (see Appendix). Let $V_C = \tilde{V}_2$ and $V_B = V_2 \setminus \tilde{V}_2$, the existence of a [A : B, C : (2, D)] structure is likely with respect to such graph partitions. Second, using the ideas in our proof of Lemma 3.4.4 and Bukh and Jiang's proof of their Lemma 9, we can prove the following result.

Lemma 3.6.1. Under the same constraints on A, B, D as in Lemma 3.4.4, if G is a trylayered graph on $V_1, V_2 = V_B \cup V_C, V_3$ with minimum degree at least [A : B, C : (2, D)], then there is a Θ -graph in $G[V_2, V_3]$ or there is a well-placed Θ -graph in G.

Therefore, if one is able to show, under weaker conditions in comparison to Lemma 3.4.1, that either an [A : B, C : D] structure exists or an [A : B, C : (2 : D)] structure exists, then one could improve our bound. We were able to prove an analog of Lemma 3.4.2 for the [A : B, C : (2 : D)] structure, but was unable to derive an analog of Lemma 3.4.3.

It is also worth pointing out that the constant factor of the upper bound proved by this chapter is not fully optimized. In particular, we believe that the bound can be further improved by constant factors if instead of using the Reduction Lemma, we employ a modified breadth-first search algorithm (see Bukh–Jiang Section 1) to bound the maximum degree in our graph. In this chapter, we decided not to present the proof with a modified BFS since we believe the current proof with Reduction Lemma is cleaner and more applicable to further problems.

Appendix

Bukh–Jiang's Proof of Reduction Lemma (slightly modified) Let H be a subgraph of G that maximizes the ratio $e(H)/v(H)^{1+\alpha/2}$. By the assumption on e(G), this ratio is at least $cn^{\alpha/2}$. Since $e(H) \leq v(H)^2/2$, it then follows that $v(H)^{1-\alpha/2} \geq 2cn^{\alpha/2}$. Let S be subset of V(H) consisting of $\gamma v(H)$ vertices of largest degrees. We consider two cases.

Suppose at least e(H)/4 edges of H are incident to vertices in S. Set $\eta = 2\gamma/\alpha$. By averaging, we can find a set $T \subset V(H) \setminus S$ of $\eta v(H)$ elements that is incident to at least fraction $\eta/(1-\gamma)$ of edges leaving S. Hence, $e(S \cup T) \ge (\frac{\eta}{1-\gamma})e(H)/4 \ge \eta e(H)/4$. Let H' be the subgraph of H induced by $S \cup T$. Since

$$(\gamma + \eta)^{1+\alpha/2} = \gamma^{1+\alpha/2} (1 + 2/\alpha)^{1+\alpha/2} \le (3/\alpha)^{1+\alpha/2} \gamma^{1+\alpha/2}$$
$$\le (3^{3/2}/\alpha^{1+\alpha/2}) \gamma^{1+\alpha/2} \le (10/\alpha) \gamma^{1+\alpha/2} \le \gamma/2,$$

we have

$$\frac{e(H')}{v(H')^{1+\alpha/2}} \ge \frac{\eta e(H)}{2\gamma v(H)^{1+\alpha/2}} = \frac{e(H)}{\alpha v(H)^{1+\alpha/2}} > \frac{e(H)}{v(H)^{1+\alpha/2}},$$

contradictory to the choice of H.

Therefore, we may assume that S is incident to fewer than e(H)/4 edges of H. Thus the minimum degree of a vertex in S is at most $\frac{e(H)}{2|S|} = \frac{e(H)}{2\gamma v(H)}$. Removing edges incident to S from H then leaves a graph H' with maximum degree at most $\frac{e(H)}{2\gamma v(H)}$ (since S consists of vertices of highest degrees in H) and at least 3e(H)/4 edges. In particular, average degree of H' is at least 3e(H)/(2v(H)).

Now we remove vertices of degree less than e(H)/(2v(H)) repeatedly to obtain G'. Since the number of edges removed is less than e(H)/2, G' would have at least $\gamma v(H)$ vertices and e(H)/4 edges. Each vertex in this graph has degree between e(H)/2v(H) and $e(H)/2\gamma v(H)$, and we have $e(G') \ge e(H)/4 \ge (c/4)n^{\alpha/2}v(H)^{1+\alpha/2} \ge (c/4)v(G)^{1+\alpha/2}$. Finally, since $e(H)/v(H) \ge cn^{\alpha/2}v(H)^{\alpha} \ge cv(G')^{\alpha}$, we are done.

Bukh–Jiang's Proof of Lemma 3.4.2 We suppose that alternative (I) does not hold. Then, by Corollary 3.3.2, the average degree of every subgraph of $G[V_1, V_2]$ is at most 2k.

Consider the process that aims to construct a subgraph satisfying (II). The process starts with $V'_1 = V_1$, $V'_2 = V_2$ and $V'_3 = V_3$, and at each step removes one of the vertices that violate the minimum degree condition on $G[V'_1, V'_2, V'_3]$. The process stops when either no vertices are left, or the minimum degree of $G[V'_1, V'_2, V'_3]$ is at least [A : B, C : D]. Since in the latter case we are done, we assume that this process eventually removes every vertex of G.

Let *R* be the vertices of V_2 that were removed because at the time of removal they had fewer than *C* neighbors in V'_3 . Put

 $E' \stackrel{\text{def}}{=} \{uv \in E(G) : u \in V_2, v \in V_3, \text{ and } v \text{ was removed before } u\},\$ $S \stackrel{\text{def}}{=} \{v \in V_2 : v \text{ has at least } 4k^2 \text{ neighbors in } V_1\}.$

Note that $|E'| \leq D|V_3|$. We cannot have $|S| \geq |V_1|/k$, for otherwise the average degree of the bipartite graph $G[V_1, S]$ would be at least $\frac{4k}{1+1/k} \geq 2k$. So $|S| \leq |V_1|/k$.

The average degree condition on $G[V_1, S]$ implies that

$$e(V_1, S) \le k(|V_1| + |S|) \le (k+1)|V_1|.$$

Let *u* be any vertex in $R \setminus S$. Since it is connected to at least $(d+4k^2+C)-4k^2 = d+C$ vertices of V_3 , it must be adjacent to at least *d* edges of *E'*. Thus,

$$|R \setminus S| \le |E'|/d \le D|V_3|/d.$$

Assume that the conclusion (III) does not hold with $\tilde{V}_2 = R \setminus S$. Then $e(V_1, R \setminus S) < (1-a)e(V_1, V_2)$. Since the total number of edges between V_1 and V_2 that were

removed due to the minimal degree conditions on V_1 and V_2 is at most $A|V_1|$ and $B|V_2|$ respectively, we conclude that

$$e(V_1, V_2) \le e(V_1, S) + e(V_1, R \setminus S) + A|V_1| + B|V_2|$$

$$< (k+1)|V_1| + (1-a)e(V_1, V_2) + A|V_1| + B|V_2|,$$
(3.16)

implying that

$$a \cdot e(V_1, V_2) < (A + k + 1)|V_1| + B|V_2|.$$

The contradiction with (3.4) completes the proof.

Chapter 4

Representative Sets Method in Matroid Theory

4.1 Introduction

Let G = (V, E) be an unweighted, directed graph, and let $S, T \subset V$ be sets of terminals. In the vertex sparsifier problem, our goal is to construct a smaller graph H, called the *vertex sparsifier*, that preserves the cut structure of S, T in G. More precisely, H should include all vertices in S, T, and for all subsets $A \subseteq S$ and $B \subseteq T$, the size of the min-cut separating A and B is the same in G and H. Here, we allow the min-cut to contain vertices from A and B.

A landmark result of Kratsch and Wahlström proved the first bound on the size of a vertex sparsifier that is polynomial in the number of terminals. When S, T have size k, the vertex sparsifier has $O(k^3)$ vertices. Kratsch and Wahlström's main insight is to phrase the problem in terms of constructing representative families on a certain matroid, after which they can appeal to the rich theory on representative families [42, 45]. Their result, also known as the *cut-covering lemma* in the areas of fixed-parameter tractability and kernelization, has led to many new algorithmic developments [32–35]. Nevertheless, despite the recent surge in applications of the cut-covering lemma, the original bound of $O(k^3)$ has yet to be improved.

In this chapter, we observe the ordered version of the representative family method, and use it to give a sparsifier on $O(k^2)$ vertices in directed acyclic graphs. This matches known lower bounds of $\Omega(k^2)$. Furthermore, unlike previous versions, our new algorithm runs in linear time in the size of the graph, and computes a cover for all *furthest* min-cuts between subsets of the terminals. We expect this may lead to further applications in the theory of kernelization. The central method we use is the following theorem.

Theorem 4.1.1. Suppose $\mathcal{F} \subseteq \mathcal{P}(\mathbb{F}^d)$ for some field \mathbb{F} and for all $B \in \mathcal{F}$, |B| = s. Let

 $\mathcal{A} = \{ A \subseteq \mathbb{F}^d \mid |A| \leq r \text{ and } \exists B \in \mathcal{F} \text{ s.t. } A \uplus B \text{ is linearly independent} \}.$

Fix any ordering σ of \mathcal{F} , namely $\mathcal{F} = \{B_1, B_2, \cdots, B_n\}$, and suppose $d \ge r+s$. Then there exists $\mathcal{B} \subseteq \mathcal{F}$, $\mathcal{B} = \{B_{i_1}, B_{i_2}, \cdots, B_{i_m}\}$ where $i_1 < i_2 < \ldots < i_m$, such that

- (a) For all $A \in A$, there exists $B_{i_k} \in B$ where $A \uplus B_{i_k}$ is independent and for all $j \in [n], j > i_k$, $A \uplus B_j$ is dependent. Note that B_j is not necessarily in B.
- (b) $m \leq {\binom{d}{s}}$, and we can find \mathcal{B} algorithmically using $O({\binom{d}{s}}ns^{\omega} + {\binom{d}{s}}^{\omega-1}n)$ field operations over \mathbb{F} (in particular, in time linear in $|\mathcal{F}|$).

Technically, this version is equivalent to the weighted version of the representative family method (see, e.g., Fomin et al. [21]), as any input with all weights distinct enforces a corresponding total ordering on the elements¹. However, the difference in focus is significant, as a total element ordering can carry semantic meaning that is obscured when implemented using weights.

Applying Theorem 4.1.1 to vertex cut sparsifiers, we obtain the main result of this chapter.

Theorem 4.1.2. Given a directed acyclic graph G = (V, E) with terminal sets S, T of size k, we can find a vertex cut sparsifier of G of size $\Theta(k^2)$ algorithmically in time $O((m+n)k^{O(1)})$.

In order to get Theorem 4.1.2 to run in linear time, two further obstacles need to be overcome. The first is to compute a representation for the matroids underlying the result, known as *gammoids* in time linear in m + n. The usual method for representing gammoids goes via the class of *transversal matroids*, however, this requires taking the inverse of an $n \times n$ matrix. Luckily, we observe that an older construction of Mason [43] can be used to represent gammoids more efficiently over DAGs; see Lemma 4.2.1.

The other obstacle is that the method of Kratsch and Wahlström [35] is inherently iterative. They use a representative family computation to find *essential vertices* – i.e., vertices that have to be included in the sparsifier if the result is to be correct – then eliminate one non-essential vertex at a time until all vertices are deemed potentially essential. Using Theorem 4.1.2 and the topological ordering of a DAG, we show that a

¹In the first version of this manuscript, we presented a proof of Theorem 4.1.1 that runs in polynomial time. It was later pointed out to us that this theorem is equivalent to Theorem 3.7 in [21], which runs in linear time. We thereby refer the readers to the proof in [21], as our proof shares similar underlying ideas with theirs.

single representative family computation lets us find all vertices contained in furthest min-cuts between subsets of the terminals, thereby removing the need for iteration. In fact, a cut-covering set for furthest (or, symmetrically, closest) min-cuts was not previously known.

4.2 Preliminaries

Throughout this chapter, all graphs are directed and unweighted. We begin with standard terminology on cuts and cut sparsifiers.

Definition 4.2.1 (Vertex Cut). Given a directed unweighted graph G = (V, E) with sets $X, Y \subseteq V$, a set $C \subseteq V$ is a vertex cut of (X, Y) if after removing C from G, there does not exist a path from a vertex in X to a vertex in Y. We denote the size of a minimum vertex cut between X, Y in G as mincut_G(X, Y).

Definition 4.2.2 (Vertex Cut Sparsifier). Consider a directed unweighted graph G = (V, E) with sets $S, T \subseteq V$. A directed unweighted graph H = (V', F) is a vertex cut sparsifier of G if

(a) $S,T \subseteq V'$.

(b) For all $X \subseteq S, Y \subseteq T$, mincut_G(X, Y) = mincut_H<math>(X, Y).

The problem we consider in this paper is the minimum size of a vertex sparsifier.

Problem (Minimum Vertex Cut Sparsifier). Given a directed unweighted graph G = (V, E) with terminal sets $S, T \subseteq V$, what is the minimum number of vertices in a vertex cut sparsifier of G?

Vertex cut sparsifiers were first introduced by Moitra [44] in the approximation algorithms setting; see also [11, 13, 40]. Recently, they have found applications in fast graph algorithms, especially in the dynamic setting [10, 12, 15]. For the specific problem above, Kratsch and Wahlström [35] obtained the bound $O(k^3)$, where |S| = |T| = k. Their application was in the fixed-parameter tractability setting, specifically in constructing kernels for cut-based problems. Our proof utilizes similar matroid-theoretic techniques as theirs, which we introduce next.

Definition 4.2.3 (Matroid). Given a finite ground set U, a set system M = (E, I) where $E \subseteq U, I \subseteq \mathcal{P}(E)$ is called a matroid if

- (a) $\emptyset \in I$.
- (b) For $X, Y \subseteq E$, if $Y \in I$ and $X \subseteq Y$, then $X \in I$.
- (c) If $X, Y \in I$ and |X| < |Y|, then there exists $y \in Y \setminus X$ such that $X \cup \{y\} \in I$.

Central to our proof is the use of gammoids and their representations.

Definition 4.2.4 (Gammoid). Given a graph G = (V, E) and a vertex subset S, the gammoid on S is the matroid M = (V, I) where I contains all subsets $T \subseteq V$ such that there exist |T| vertex-disjoint paths from S to T in G.

Definition 4.2.5 (Matroid disjoint union). Given two matroids on disjoint ground sets, their matroid disjoint union is the matroid whose ground set is the union of their ground sets, and a set is independent if the corresponding part in each matroid is independent.

Definition 4.2.6 (Representable matroid). Given a field F, a matroid M = (E, I) is representable over F if there exists a matrix A over the field F such that there exists a bijective mapping from E to the columns of A, where a set $S \subseteq E$ is independent if and only if its corresponding set of columns of A are linearly independent.

In particular, it is well known in matroid theory that gammoids are representable in randomized polynomial time; see Marx [42]. However, to control the running time, we revisit an older representation by Mason [43], and note that it leads to a representation in near-linear time in |V| + |E| on DAGs. We also note that the disjoint union of two representable matroids is representable.

Lemma 4.2.1 (Construction of Gammoid Representation on DAGs). Given a directed acyclic graph G = (V, E), a set $S \subseteq V$, and $\varepsilon > 0$, with |V| = n, |E| = m and |S| = k, a representation of the gammoid on S of dimension k over a finite field with entries of bitlength $\ell = O(k \log n + \log(1/\varepsilon))$ can be constructed in randomized time $\tilde{O}((n + m)k\ell)$ with one-sided error at most ε , where \tilde{O} hides factors logarithmic in ℓ .

Proof. We review the construction of Mason [43]. Associate a variable x_{uv} to every edge $(u, v) \in E$, where all variables x_{uv} are formally independent. For two vertices $u, v \in V$, define the path polynomial

$$P(u,v) = \sum_{P: u \rightsquigarrow v} \prod_{(u,v) \in E(P)} x_{uv}$$

where *P* ranges over all directed paths from *u* to *v* in *G*. Define the matrix *M* with rows indexed by *S* and columns by *V*, where for $u \in S$, $v \in V$ we have M(u, v) = P(u, v). Then *M* is a representation of the gammoid on *S*.

Indeed, on the one hand let $T \subseteq V$ be a basis of the gammoid. By definition, there is a vertex-disjoint flow linking S to T. Instantiate the variables x_{uv} by letting $x_{uv} = 1$ for every edge used in one of these paths, and $x_{uv} = 0$ otherwise. Under this evaluation M[S,T] is a permutation matrix, hence independent, which shows that det $M[S,T] \neq 0$. On the other hand, let $T \subseteq V$ and let $C \subseteq V$ be an (S,T)-min cut, |C| < |T|. Then M[S,T] factors as $M[S,T] = M[S,C] \cdot M'[C,T]$ for a matrix M'[C,T], hence the rank of M[S,T] is at most |C|. Here, M'[C,T] is defined as M, except that vertices of C have been turned into sources.

To get a representation over a finite field \mathbb{F} , we pick a sufficiently large field \mathbb{F} and replace every variable x_{uv} by a random value from \mathbb{F} . For the success probability, note that any dependent set in M remains dependent after such a replacement. Therefore, it is enough to consider the probability that det $M[S, B] \neq 0$ for every basis B of the gammoid. For this, we observe that the entries M(u, v) are polynomials of degree at most n, hence det M[S, B] has degree at most nk. Furthermore, the number of bases is at most $\binom{n}{k} \leq n^k$. Let $\mathbb{F} = GF(2^\ell)$ where $2^\ell > (1/\varepsilon)(nk)n^k$, i.e., $\ell = \Theta(k \log n + \log(1/\varepsilon))$. By Schwartz-Zippel, the probability that det M[S, B] = 0 for a given basis B is at most $nk/|\mathbb{F}| \leq \varepsilon n^{-k}$, hence by the union bound the probability that this occurs for at least one basis B is at most ε . We note that field arithmetic over \mathbb{F} can be performed in time $\tilde{O}(\ell)$.

It remains to evaluate the vectors $R_v = (P(s, v))_{s \in S}$ for $v \in V$ quickly. For simplicity, we assume without loss of generality that the vertices $s \in S$ are sources in G; this can be achieved by introducing a new vertex s' for every $s \in S$, with a single edge (s', s), and replacing s by s' in S. Note that this does not change the resulting gammoid. Let $V = \{v_1, \ldots, v_n\}$ where (v_1, \ldots, v_n) is a topological ordering of G, starting with the vertices of S; this can be computed in time O(m + n) by standard methods. Note that P(s, s) = 1 for $s \in S$, hence the vectors R_s , $s \in S$ are unit vectors making up the standard basis for \mathbb{F}^k . For all other vertices $v \in V$, note

$$P(s,v) = \sum_{u \in N^-(v)} P(s,u) x_{uv},$$

where $P(s, u) = R_u(s)$ has already been computed due to the topological ordering. Hence R_v can be computed using $O(kd^-(v))$ field operations from the previously computed vectors. Performing this across all variables v_1, \ldots, v_n uses O((n + m)k) field operations, hence the total running time is bounded by $\tilde{O}((n + m)k\ell)$, as stated. \Box

Given these definitions, we now apply Theorem 4.1.1 to representable matroids and obtain the following Theorem, which we will use in our proof.

Theorem 4.2.2 (Algorithmic Skew-Symmetric Bollobás's Theorem on Representable Matroids). Consider a representable matroid M = (E, I) of rank r + s. Let \mathcal{A} be a family of sets of size at most r, and let $\mathcal{F} = \{B_1, B_2, \dots, B_n\}$ be a family of sets of size s such that for all $A \in \mathcal{A}$, exists $B \in \mathcal{F}$, $A \uplus B \in I$. Then there exists $\mathcal{B} = \{B_{i_1}, \dots, B_{i_m}\} \subseteq F$, $m \leq {r+s \choose s}$, such that for all $A \in \mathcal{A}$, exists j where $A \uplus B_{i_j} \in I$ and for all $k > i_j, A \uplus B_k \notin I$. \mathcal{B} can be found in time $O((m + n)k^{O(1)})$.

4.3 Vertex Cut Sparsifier for DAGs

In this section, we prove our main result, Theorem 4.1.2. We first borrow the following key concepts from Kratsch and Wahlström [35].

Definition 4.3.1 (Essential Vertex). A vertex $v \in V \setminus (S \cup T)$ is called essential if there exists $X \subseteq S, Y \subseteq T$ such that v belongs to every minimum vertex cut between X, Y.

Definition 4.3.2 (Neighborhood Closure). For a digraph G = (V, E) and a vertex $v \in G$, the neighborhood closure operation is defined by removing v from G and adding an edge from every in-neighbor of v to every out-neighbor of v. The new graph is denoted by $cl_v(G)$.

Definition 4.3.3 (Closest Set). For sets of vertices $X, A \subseteq V$, A is closest to X if A is the unique min-cut between X and A.

We remark that neighborhood closure is exactly what we call vertex elimination, but we prefer to keep the same terminology from [35]. We introduce the following definitions to simplify our discussions.

Definition 4.3.4. For sets $X \subseteq S, Y \subseteq T$, let C be a vertex cut for X, Y. Let G' be the subgraph formed by the union of all paths from X to Y. The left-hand side of C, denotes L(C), is the set of vertices that are still reachable from X in G' after C is removed. Similarly, the right-hand side of C, denoted R(C), is the set of vertices that can still reach Y in G' after C is removed.

Definition 4.3.5 (Saturation). Let C be a vertex cut for $X \subseteq S, Y \subseteq T$. For a vertex $v \in C$, we say that v and C are saturated by X if there exists |C| + 1 paths from X to C that are vertex disjoint except for two paths that both ends at v. Similarly, v and C are saturated by Y if there exists |C| + 1 paths from C to Y that are vertex disjoint except for two paths that both starts at v.

The following three lemmas are stated and proved as Reduction Rule 1, Proposition 1 and Lemma 5 in [35], respectively. Their proofs, as presented in [35] and Chapter 11.6 of [22], are included in the Appendix for completeness.

Lemma 4.3.1 (Closure Lemma). If $v \in V \setminus (S \cup T)$ is not an essential vertex, then $cl_v(G)$ is a vertex cut sparsifier of G.

Lemma 4.3.2 (Closest Cut Lemma). Let C be a vertex cut for $X \subseteq S, Y \subseteq T$ that is closest to X or Y, then for all vertices $v \in C$, v and C are saturated by X or Y, respectively.

Lemma 4.3.3 (Essential Vertex Lemma). Let v be an essential vertex with respect to $X \subseteq S, Y \subseteq T$. Let C be any minimum vertex cut between X, Y. Then v and C are

saturated by both X and Y.

Using these results, Kratsch and Wahlström presented an inspired construction of matroids which, combined with the theory of representative families, resulted in their $O(k^3)$ bound. We now present a modified construction that utilizes skew-symmetry in Theorem 4.2.2, resulting in Theorem 4.1.2.

Theorem 4.3.4. For a directed acyclic graph G = (V, E) with terminal sets S and T, let $k = |S \cup T|$. Then there exists a set of vertices P of size $O(k^2)$ such that for each pair of $X \subseteq S, Y \subseteq T$, the min-(X, Y) cut that is closest to Y is contained in P. This set can be found in time $\tilde{O}(nk^{2\omega-1} + mk^2)$, and a sparsifier on P can then be constructed in the same asymptotic running time.

Proof. Let $G_R = (V, E_R)$ be the graph G with the direction of all edges reversed. We make the following modification to our graphs G, G_R . For each vertex $v \in V \setminus (S \cup T)$, add a vertex v' into V and for each directed edge $(u, v) \in E$, add (u, v') into E. Denote this new directed graph G' = (V', E'), and similarly construct $G'_R = (V', E'_R)$. Note that G', G'_R are both acyclic. Enumerate V in a reverse topological ordering, namely $V = (v_1, v_2, \cdots, v_n)$ where v_i cannot reach any v_j for j > i.

Let $M_1 = (E_1, I_1)$ be the gammoid constructed on the graph G and the set of terminals S in G, and let $M_2 = (E_2, I_2)$ be the gammoid constructed on the graph G'_R and the set of terminals T in G'_R . To distinguish between elements of E_1 and E_2 , we label vertices in E_1 as $v_1, \ldots, v_n, v'_1, \ldots, v'_n$, and elements in E_2 as $\overline{v_1}, \ldots, \overline{v_n}, \overline{v_1}', \ldots, \overline{v_n}'$. For any set of vertices $U \subseteq V$, denote the respective sets in E_1 and E_2 as U_1 and U_2 . Let M be the disjoint union of matroids M_1 and M_2 , so M is representable and it has rank O(k).

Observe that for any $X \subseteq S$, $Y \subseteq T$, the min-cuts between X and Y in G are the same as in G' because the vertex copies v' we added to G have no outgoing edges. Therefore, G and G' have the same set of closest cuts. We now construct families of sets \mathcal{A}, \mathcal{F} that satisfy the conditions of Theorem 4.2.2. For a min-cut C between X, Y that is closest to Y, and a vertex $v \in C$, define

$$A_{(C,v)} = (S_1 \setminus X_1 \cup C_1 \setminus \{v\}) \cup (T_2 \setminus Y_2 \cup C_2).$$

Let \mathcal{A} consists of $A_{(C,v)}$ for all such cut-vertex pairs. Define

$$\mathcal{F} = \{ B_v = \{ v, \overline{v}' \} \mid v \in V \}.$$

We prove the following claim.

Claim 4.3.1. For each $A = A_{(C,v)}$, B_v is the unique set in \mathcal{F} such that $A \uplus B_v$ is independent in M, and for all u > v in the reverse topological order of V, $A \uplus B_u$ is dependent in M.

Proof of Claim 4.3.1. We first show that $A \uplus B_v$ is independent. Since M is a disjoint union matroid, we need to show that $(A \cap E_1) \uplus \{v\} = S_1 \setminus X_1 \cup C_1$ is independent in M_1 and $(A \cap E_2) \uplus \{\overline{v'}\} = T_2 \setminus Y_2 \cup C_2 \cup \{\overline{v'}\}$ is independent in M_2 .

Since both M_1 and M_2 are gammoids, we need to show existence of vertex disjoint paths from S_1 to $S_1 \setminus X_1 \cup C_1$. First note that singleton paths can cover all vertices in $S_1 \setminus X_1$. Since C_1 is a min-cut between X_1 and Y_1 , by duality there exists vertex disjoint paths from X_1 to C_1 . Therefore $S_1 \setminus X_1 \cup C_1$ is independent in M_1 . Similarly, singleton paths can cover all vertices in $T_2 \setminus Y_2$. It suffices for us to show the existence of vertex disjoint paths from Y_2 to $C_2 \cup \{\overline{v'}\}$.

By Lemma 4.3.2, there exists $|C_2| + 1$ paths from Y_2 to C_2 that are vertex disjoint except for two paths that both ends at \overline{v} . Therefore, we can redirect one of these two paths to end at \overline{v}' , and we obtain $|C_2| + 1$ vertex disjoint paths from Y_2 to $C_2 \cup {\overline{v}'}$. This proves independence.

Now fix u > v in the reverse topological ordering, so that there does not exist a path from v to u. We want to show that either $(A \cap E_1) \uplus \{u\}$ is dependent in M_1 , or $(A \cap E_2) \uplus \{\overline{u'}\}$ is dependent in M_2 . Consider four possible cases:

- *u* is not on any path from *X* to *Y*. Assume for the sake of contradiction that both (*A* ∩ *E*₁) ⊎ {*u*} and (*A* ∩ *E*₂) ⊎ {*ū*'} are independent. Then there must exist a path from *X* to *u* and a path from *u* to *Y*, which forms a path from *X* to *Y* through *u* (since *G* is acyclic), contradiction.
- u ∈ L(v) (see Definition 4.3.4). Then any path from u to Y (or from Y to u in G_R) must intersect with C, which means there does not exist vertex disjoint paths from Y to C ∪ {u} in G_R. Therefore (A ∩ E₂) ⊎ {ū'} is dependent.
- *u* ∈ *R*(*v*). Then any path from *X* to *u* must intersect *C*. Assume for the sake of contradiction that (*A* ∩ *E*₁) ⊎ {*u*} is independent, then there exists vertex disjoint paths from *X* to *C* \ {*v*} ∪ {*u*}, which means there is a path from *X* to *u* that goes through *v*. However, since *u* > *v* in the topological ordering, there does not exist paths from *v* to *u*. This is a contradiction, so (*A* ∩ *E*₁) ⊎ {*u*} is dependent.

• $u \in C$. Then $u \in (A \cap E_2)$, which implies $(A \cap E_2) \uplus \{u\}$ is dependent.

This completes the proof.

Let *P* be the collection of vertices that Theorem 4.2.2 finds. Then by the above claim, for each pair of $X \subseteq S, Y \subseteq T$ and their min-cut *C* closest to *Y*, all vertices in *C* must be

in *P*. Note that this also implies that all essential vertices are in *P*. For the final running time, note that computing the gammoids takes time $\tilde{O}((m+n)k^2)$ by Lemma 4.2.1 with $\varepsilon = 1/n^{O(1)}$, and computing the representative sets takes time $\tilde{O}(nk^{2\omega-1})$ by taking d = k and s = 2 in Theorem 4.1.1.

To construct the final sparsifier H on P, for each vertex $u \in P$, we run a depth-first search starting at u on the graph G_u , defined to be G minus the out-edges of vertices in $P \setminus \{u\}$. For each vertex $v \in P \setminus \{u\}$ that is reachable from u in G_u , we add an edge (u, v)to H. Note that this procedure returns the same graph H as the one that sequentially applies Lemma 4.3.1 on all vertices not in P, but achieves a shorter runtime of $O(k^2m)$. To see the equivalence, observe that in both cases, there is an edge (u, v) in the final sparsifier if and only if there is a path from u to v in G whose internal vertices are disjoint from P. We conclude that the output graph H is a valid sparsifier.

We make two remarks about this proof. First of all, we never explicitly compute the sets $A_{(C,v)}$. They are used only to prove the existence of a vertex cut sparsifier. More importantly, the critical difference between our proof and Kratsch and Wahlström's proof is that our proof gives an asymmetrical construction. Kratsch and Wahlström utilized the property that essential vertices can be saturated from both sides (see Lemma 4.3.3) to give a symmetrical construction using three matroids, while we use the topological ordering on DAGs to give a construction that only requires one side of saturation.

We now present tight lower bound constructions in the following section.

4.4 Lower Bound Constructions

In this section we present two constructions that have $\Omega(k^2)$ essential vertices, which implies that their vertex cut sparsifiers must have size at least $\Omega(k^2)$. The first construction is presented by Kratsch and Wahlström in [35].

Construction 4.4.1. Let S and T be two vertex sets of size 2n. Enumerate them as

$$S = \{v_1, v'_1, v_2, v'_2, \cdots, v_n, v'_n\}, T = \{u_1, u'_1, u_2, u'_2, \cdots, u_n, u'_n\}.$$

For each $i, j \in [n]$, create a vertex $w_{i,j}$ and add edges from v_i, v'_i to $w_{i,j}$, and from $w_{i,j}$ to u_j, u'_j . Then $w_{i,j}$ is an essential vertex with respect to $X = \{v_i, v'_i\}$ and $Y = \{u_j, u'_j\}$.

The second construction is a variant of a grid. Similar constructions are present in many related works; see, for example, [36].

Construction 4.4.2. Consider the following grid of vertices in Figure 4.1.



Figure 4.1: Grid construction

Let S and T be the boundary vertices marked as red, namely

 $S = T = \{t_{0,1}, \cdots, t_{0,n}, t_{1,0}, \cdots, t_{n,0}, t_{n+1,1}, \cdots, t_{n+1,n}, t_{1,n+1}, \cdots, t_{n,n+1}, t_{n+1,n+1}\}.$

Formally, the grid is constructed by adding the following edges.

- (a) For each row $i \in [n]$. add horizontal edges from left to right between consecutive vertices from $t_{i,0}$ to $t_{i,n+1}$. Note that we exclude the top and bottom rows of terminals.
- (b) For each column $i \in [n]$, add vertical edges from top to bottom between consecutive vertices from $t_{0,i}$ to $t_{n+1,i}$. Note that we exclude the left and right columns of terminals.
- (c) For each vertex $v_{i,j}$ where $i, j \le n$, add a diagonal edge going to the immediate vertex in the bottom right direction of $v_{i,j}$.

We show that the vertices $v_{1,1}, \dots, v_{n,n}$ are all essential. For each $i = 1, \dots, n$, consider

$$X = \{t_{1,0}, \cdots, t_{n,0}, t_{0,1}, \cdots, t_{0,i}, t_{n+1,1}, \cdots, t_{n+1,i}\}, Y = T \setminus X.$$

In other words, we partition the sets of terminals into two sets, namely the left and right sides of the *i*th column. Observe that the column $C = \{v_{1,i}, \dots, v_{n,i}\}$ is a min-cut for X, Y. Moreover, every vertex $v \in C$ is essential with respect to X, Y. Applying this argument to all columns, we see that all vertices $v_{1,1}, \dots, v_{n,n}$ are essential.

These two constructions exhibit a dichotomy. Call a min-cut C between X, Y essential if exists $v \in C$ that is essential with respect to X, Y. Construction 4.4.1 demonstrate that a directed acyclic graph could have $O(k^2)$ essential cuts, each containing O(1) essential vertices. On the other hand, Construction 4.4.2 gives a directed acyclic graph with O(k) essential cuts, each of size O(k). We do not observe any construction of directed graphs that falls outside of these two cases.

4.5 Conclusions

After the presented results were derived, the authors made subsequent attempts to improve the bound over general directed graphs. While we have found evidence suggesting that $O(k^2)$ is the correct bound in general directed graphs, we were not able to find a complete proof.

Appendix

Proof of Lemma 4.3.1

Lemma. If $v \in V \setminus (S \cup T)$ is not an essential vertex, then $cl_v(G)$ is a vertex cut sparsifier of G.

Proof. Let $H = cl_v(G)$, it suffices for us to show that for all $X \subseteq S, Y \subseteq T$,

 $\operatorname{mincut}_{G}(X, Y) = \operatorname{mincut}_{H}(X, Y).$

We first show $\operatorname{mincut}_G(X, Y) \leq \operatorname{mincut}_H(X, Y)$. Let *C* be a min-cut between *X*, *Y* in *H*, we show that *C* is a valid vertex cut in *G*. Suppose not, then there exists a path *P* in *G* from *X* to *Y* that does not intersect with *C*. If $v \notin P$, then *P* is also a path in *H*, which contradicts the fact that *C* is a cut. If $v \in P$, let *u* and *w* be *v*'s predecessor and successor in *P*. According to the closure operation on *v*, there is an edge *uw* in *H*. Therefore we again have a path from *X* to *Y* in *H* not intersecting *C*, contradiction.

Next we show $\operatorname{mincut}_G(X, Y) \ge \operatorname{mincut}_H(X, Y)$. Since v is not essential, there exists a min-cut C between X, Y such that $v \notin C$. We show that C is a cut in H. Suppose not, then there exists a path P in H from X to Y not crossing C. For every edge uw on P that is created by the closure operation, replace it by the path u, v, w and denote the new walk P'. After truncating all the cycles of P', we obtain a path from X to Y in Gnot crossing C, contradiction. We conclude that $\operatorname{mincut}_G(X, Y) = \operatorname{mincut}_H(X, Y)$. \Box

Proof of Lemma 4.3.2

Lemma. Let C be a vertex cut for $X \subseteq S, Y \subseteq T$ that is closest to X or Y, then for all vertices $v \in C$, v and C are saturated by X or Y, respectively.

Proof. We prove this lemma for cuts closest to X, as the other case is symmetrical. Assume for the sake of contradiction there exists $v \in C$ such that v and C is not saturated by X. Add a sink-only copy of v into our digraph G and call the new graph G', then by duality there must exists a cut C' of size |C| between X and $C \cup \{v'\}$ in G'. We now consider a few cases.

First note that we can't have v, v' both in C', as otherwise $C' \setminus \{v'\}$ would be a smaller (X, Y) cut in G. Now if $v \in C'$ and $v' \notin C'$, then by duality there is a path from X to v that does not intersect $c' \setminus \{v\}$, which can be redirected to v', contradicting the fact that C' is a cut. Similarly, if $v' \in C'$ and $v \notin C$ we have a contradiction. Therefore $v, v' \notin C'$, which implies C is not the unique (X, C) min-cut, contradiction. \Box

Proof of Lemma 4.3.3

Lemma. Let v be an essential vertex with respect to $X \subseteq S, Y \subseteq T$. Let C be any minimum vertex cut between X, Y. Then v and C are saturated by both X and Y.

Proof. We slightly modify our graph. Add a vertex v' to G, and add edges such that v' has the same in-neighbors and out-neighbors as v. Denote the new graph as G', we show that $C' = C \cup \{v'\}$ is a min-cut between X, Y in G'. This lemma then follows from Menger's Theorem.

Assume for the sake of contradiction that C' is not a min-cut between X, Y, and let D be a min-cut between X, Y where |D| < |C'| = |C| + 1. We consider a few cases.

- (a) v ∈ D and v' ∉ D. Since D is a min-cut, by Menger's theorem there exists a path P from X to Y such P ∩ D = {v}. If we replace v by v' in P, we obtain a valid path P' from X to Y that does not cross D, which is a contradiction to the fact that D is a cut.
- (b) $v' \in D$ and $v \notin D$. The same argument from the previous case works.
- (c) $v, v' \notin D$. Then *D* is a valid cut in *G*, and $|D| \leq |C|$. Therefore *D* is a min-cut in *G* that does not contain *v*, which is a contradiction.
- (d) $v, v' \in D$. Then consider $D' = D \setminus \{v'\}$. D' is a valid cut in G, and $|D'| \le |C| 1$, which contradicts the fact that C is a min-cut in G.

Therefore we conclude that *D* is a min-cut in *G'*, which completes the proof. \Box

Bibliography

- [1] C. T. Benson. Minimal regular graphs of girths eight and twelve. *Canad. J. Math.*, 18:1091–1094, 1966.
- [2] P. V. M. Blagojević, B. Bukh, and R. Karasev. Turán numbers for $K_{s,t}$ -free graphs: Topological obstructions and algebraic constructions. *Israel Journal of Mathematics*, 197(1):199–214, 2013.
- [3] Béla Bollobás. On generalized graphs. Acta Mathematica Hungarica, 16(3-4):447– 452, 1965.
- [4] J. A. Bondy and M. Simonovits. Cycles of even length in graphs. *J. Combinatorial Theory Ser. B*, 16:97–105, 1974.
- [5] William G Brown. On graphs that do not contain a thomsen graph. *Canadian Mathematical Bulletin*, 9(3):281–285, 1966.
- [6] B. Bukh. Random algebraic construction of extremal graphs. *Bulletin of the London Mathematical Society*, 47(6):939–945, 2015.
- [7] Boris Bukh and David Conlon. Rational exponents in extremal graph theory. *Journal of the European Mathematical Society*, 20(7):1747–1757, 2018.
- [8] Boris Bukh and Zilin Jiang. A bound on the number of edges in graphs without an even cycle. *Combinatorics, Probability and Computing*, 26(1):1–15, 2017.
- [9] Boris Bukh and Michael Tait. Turán numbers of theta graphs. *Combinatorics, Probability and Computing*, 29(4):495–507, 2020.
- [10] Parinya Chalermsook, Syamantak Das, Bundit Laekhanukit, Yunbum Kook, Yang P Liu, Richard Peng, Mark Sellke, and Daniel Vaz. Vertex sparsification for edge connectivity. arXiv preprint arXiv:2007.07862, 2020.
- [11] Moses Charikar, Tom Leighton, Shi Li, and Ankur Moitra. Vertex sparsifiers and ab-

stract rounding algorithms. In 2010 IEEE 51st Annual Symposium on Foundations of Computer Science, pages 265–274. IEEE, 2010.

- [12] Li Chen, Gramoz Goranci, Monika Henzinger, Richard Peng, and Thatchaphol Saranurak. Fast dynamic cuts, distances and effective resistances via vertex sparsifiers. *arXiv preprint arXiv:2005.02368*, 2020.
- [13] Julia Chuzhoy. On vertex sparsifiers with steiner nodes. In *Proceedings of the fortyfourth annual ACM symposium on Theory of computing*, pages 673–688, 2012.
- [14] David Conlon. Graphs with few paths of prescribed length between any two vertices. *Bulletin of the London Mathematical Society*, 51(6):1015–1021, 2019.
- [15] David Durfee, Yu Gao, Gramoz Goranci, and Richard Peng. Fully dynamic spectral vertex sparsifiers and applications. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*, pages 914–925, 2019.
- [16] P. Erdős. On sequences of integers no one of which divides the product of two others and on some related problems. *Isvestia Nauchno-Issl. Inst. Mat. i Meh. Tomsk*, 2:74–82, 1938.
- [17] P. Erdős, A. Rényi, and V. T. Sós. On a problem of graph theory. *Studia Sci. Math. Hungar.*, 1:215–235, 1966.
- [18] P. Erdős and M. Simonovits. Some extremal problems in graph theory, combinatorial theory and its applications (proc. colloq. balatonf ured, 1969). 1970.
- [19] Paul Erdös and Arthur H Stone. On the structure of linear graphs. *Bulletin of the American Mathematical Society*, 52(12):1087–1091, 1946.
- [20] R. J. Faudree and M. Simonovits. On a class of degenerate extremal graph problems. *Combinatorica*, 3(1):83–93, 1983.
- [21] Fedor V. Fomin, Daniel Lokshtanov, Fahad Panolan, and Saket Saurabh. Efficient computation of representative families with applications in parameterized and exact algorithms. J. ACM, 63(4):29:1–29:60, 2016.
- [22] Fedor V Fomin, Daniel Lokshtanov, Saket Saurabh, and Meirav Zehavi. *Kernelization: theory of parameterized preprocessing*. Cambridge University Press, 2019.
- [23] Z. Füredi and M. Simonovits. The history of degenerate (bipartite) extremal graph problems. In *Erdős Centennial*, pages 169–264. Springer, 2013.
- [24] D. Gerbner and C. Palmer. Extremal results for Berge hypergraphs. SIAM J. Discrete

Math., 31(4):2314–2327, 2017.

- [25] Dániel Gerbner, Abhishek Methuku, and Cory Palmer. General lemmas for berge– turán hypergraph problems. *European Journal of Combinatorics*, 86:103082, 2020.
- [26] Dániel Gerbner, Abhishek Methuku, and Máté Vizer. Asymptotics for the Turán number of Berge-*K*_{2,t}. *Journal of Combinatorial Theory, Series B*, 2019.
- [27] D. Grósz, A. Methuku, and C. Tompkins. Uniformity thresholds for the asymptotic size of extremal Berge-*F*-free hypergraphs. *Electronic Notes in Discrete Mathematics*, 61:527–533, 2017.
- [28] Ervin Győri and Nathan Lemons. 3-uniform hypergraphs avoiding a given odd cycle. *Combinatorica*, 32(2):187–203, 2012.
- [29] Zhiyang He. A new upper bound on extremal number of even cycles. *The Electronic Journal of Combinatorics*, pages P2–41, 2021.
- [30] 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), pages 52–1. Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 2021.
- [31] Zhiyang He and Michael Tait. Hypergraphs with few berge paths of fixed length between vertices. *SIAM Journal on Discrete Mathematics*, 33(3):1472–1481, 2019.
- [32] Eva-Maria C Hols and Stefan Kratsch. A randomized polynomial kernel for subset feedback vertex set. *Theory of Computing Systems*, 62(1):63–92, 2018.
- [33] Stefan Kratsch. A randomized polynomial kernelization for vertex cover with a smaller parameter. *SIAM Journal on Discrete Mathematics*, 32(3):1806–1839, 2018.
- [34] Stefan Kratsch and Magnus Wahlström. Compression via matroids: a randomized polynomial kernel for odd cycle transversal. *ACM Transactions on Algorithms* (*TALG*), 10(4):1–15, 2014.
- [35] Stefan Kratsch and Magnus Wahlström. Representative sets and irrelevant vertices: New tools for kernelization. *J. ACM*, 67(3):16:1–16:50, 2020.
- [36] Robert Krauthgamer and Inbal Rika. Mimicking networks and succinct representations of terminal cuts. In *Proceedings of the twenty-fourth annual ACM-SIAM symposium on Discrete algorithms*, pages 1789–1799. SIAM, 2013.

- [37] Nathan W Lemons. *Turán problems for hypergraphs*. PhD thesis, dissertation, Central European University, Budapest, Hungary, 2008.
- [38] László Lovász. Flats in matroids and geometric graphs. In *Combinatorial Surveys* (*Proc. 6th British Combinatorial Conference*, pages 45–86, 1977.
- [39] J. Ma, X. Yuan, and M. Zhang. Some extremal results on complete degenerate hypergraphs. *J. Combin. Theory Ser. A*, 154:598–609, 2018.
- [40] Konstantin Makarychev and Yury Makarychev. Metric extension operators, vertex sparsifiers and lipschitz extendability. In *2010 IEEE 51st Annual Symposium on Foundations of Computer Science*, pages 255–264. IEEE, 2010.
- [41] W. Mantel. Solution to problem 28. *H. Gouwentak, W. Mantel, J. Teixeira de Mattes, F. Schuh, and WA Wythoff, Wiskundige Opgaven,* 10:60–61, 1907.
- [42] Dániel Marx. A parameterized view on matroid optimization problems. *Theoretical Computer Science*, 410(44):4471–4479, 2009.
- [43] John H Mason. On a class of matroids arising from paths in graphs. *Proceedings of the London Mathematical Society*, 3(1):55–74, 1972.
- [44] Ankur Moitra. Approximation algorithms for multicommodity-type problems with guarantees independent of the graph size. In *2009 50th Annual IEEE Symposium on Foundations of Computer Science*, pages 3–12. IEEE, 2009.
- [45] Burkhard Monien. How to find long paths efficiently. In *North-Holland Mathematics Studies*, volume 109, pages 239–254. Elsevier, 1985.
- [46] Assaf Naor and Jacques Verstraëte. A note on bipartite graphs without 2k-cycles. *Combinatorics, Probability and Computing*, 14(5-6):845–849, 2005.
- [47] O. Pikhurko. A note on the Turán function of even cycles. *Proc. Amer. Math. Soc.*, 140(11):3687–3692, 2012.
- [48] I. R. Shafarevich and M. Reid. Basic algebraic geometry, volume 2. Springer, 1994.
- [49] J. Verstraëte. On arithmetic progressions of cycle lengths in graphs. *Combin. Probab. Comput.*, 9(4):369–373, 2000.
- [50] R. Wenger. Extremal graphs with no C_4 's, C_6 's, or C_{10} 's. Journal of Combinatorial *Theory, Series B*, 52(1):113–116, 1991.