A New Upper Bound on the
Turán Number of Even Cycles

Zhiyang He[*]

Department of Mathematical Sciences
Carnegie Mellon University
Pittsburgh, U.S.A.

`szh@andrew.cmu.edu`

Abstract

In this paper, we prove $\mathrm{ex}(n, C_{2k}) \leqslant (16\sqrt{5}\sqrt{k \log k}+o(1))\cdot n^{1+1/k}$. This improves on a result of Bukh and Jiang from 2017, thereby reducing the best known upper bound by a factor of $\sqrt{5 \log k}$.

Mathematics Subject Classifications: 05C35, 05C38, 05D99

Introduction

In the field of extremal graph theory, Turán's Problem, introduced by Turán [8] in 1941, asks the following question: Given a graph F, what is the maximum number of edges that a graph on n vertices can have while not containing F as a subgraph? This number, denoted $\mathrm{ex}(n, F)$, is now referred to as the Turán number or the extremal number of F. Similarly, for a family of graphs $\mathcal{F}$, $\mathrm{ex}(n, F)$ requires that no element of $\mathcal{F}$ is present.

The first result, known as Mantel's Theorem, was proven by Mantel [5] in 1907. Since then, extensive amount of works have been established, among which is the celebrated Erdős–Stone–Simonovits Theorem [3]: $\mathrm{ex}(n, F) = (1-\frac{1}{\chi(F)-1}+o(1))\binom{n}{2}$, 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, the case for bipartite graphs is left open. For the two most studied families of bipartite graphs, complete bipartite graphs $K_{s,t}$ and even cycles C_{2k}, the magnitude of $\mathrm{ex}(n, F)$ is not known for general s, t, and k. In particular, it has been known for decades that $\mathrm{ex}(n, C_{2k}) = O_k(n^{1+1/k})$, while a matching lower bound has not been established.

To discuss methods that lead to upper bounds on $\mathrm{ex}(n, C_{2k})$, we first show a simple derivation of $\mathrm{ex}(n, \{C_3, C_4, \cdots, C_{2k}\}) \leqslant cn^{1+1/k}$ 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 $\mathrm{ex}(n, C_{2k})$ 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 $\mathrm{ex}(n, C_{2k})$ was proved by Bondy and Simonovits [1] in 1974, where they showed $\mathrm{ex}(n, C_{2k}) \leqslant 20kn^{1+1/k}$. This result is subsequently improved through a line of researches, most recently by Pikhurko [7] to $\mathrm{ex}(n, C_{2k}) \leqslant (k-1)n^{1+1/k} + O_k(n)$ in 2010 and by Bukh and Jiang [2] to $\mathrm{ex}(n, C_{2k}) \leqslant 80\sqrt{k} \log_e kn^{1+1/k} + O_k(n)$ in 2017. Our main contribution in this paper is the following theorem.

Theorem 1. *Fix k, let G be a n-vertex graph where $n \geqslant (20k)^{4k^3+2k^2}$. If*

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

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

For the rest of this paper, 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 $e(G) \leqslant c(k-1)^{-1/2}n^{1+1/k}$ if G has girth at least $2k+1$ (For a more detailed argument, see [6]). 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 $\mathrm{ex}(n, C_{2k})$ derivable from the BFS argument is $c\sqrt{k}n^{1+1/k}$ for fixed constant c. To break the $O(\sqrt{k}n^{1+1/k})$ threshold would require a different approach.

Our result improves the best known bound for $\mathrm{ex}(n, C_{2k})$ 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 paper. This paper 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 paper 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 paper 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 1 we describe the graph structures used in this paper. For a detailed discussion of how our methods relate and differ from Pikhurko's and Bukh–Jiang's work, please see Section 2.

1 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 $|V(G')| \geqslant c\gamma n^{\alpha/2}$, $|E(G)| \geqslant (c/4)v(G')^{1+\alpha}$, $d_{\min}(G') \geqslant (c/2)v(G')^\alpha$, and $d_{\max}(G')/d_{\min}(G') \leqslant 1/\gamma$.

An initial version of this lemma was first proved by Erdős and Simonovits [4], 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 paper is the following theorem.

Theorem 2. *Fix $k \geqslant 4$, let $\Delta = \sqrt{k}(20k)^{2k}$, and let*

$$d \geqslant \max(2\sqrt{5}\sqrt{k \log k}n^{1/k}, (20k)^{4k^2+2k}).$$

If G is a graph on n vertices such that $d_{min}(G) \geqslant 2d + 5k^2$ and $d_{max}(G) \leqslant \Delta d$, then G contains a copy of C_{2k}.

Theorem 1 then follows from Reduction Lemma and Theorem 2.

Proof of Theorem 1. Assume a graph H on m vertices has more than $16\sqrt{10}\sqrt{k \log k} \cdot n^{1+1/k} + 8000k^4 n^{1+(2k-1)/(2k^2)}$ 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 $n \geqslant 4\sqrt{10}\sqrt{k \log k}\gamma m^{1/(2k)}$ 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 $c \geqslant 8\sqrt{10}\sqrt{k \log k}+4000k^4/m^{1/(2k^2)}$, which implies that

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

Now from Pikhurko's Result [7], we know that if $d_{\min}(G) \geqslant kn^{1/k}$, then G contains a C_{2k}. Therefore, Reduction Lemma implies $d_{\max}(G) \leqslant (20k)^{2k}kn^{1/k}$. Let $d = 2\sqrt{10}\sqrt{k \log k}n^{1/k}$, $\Delta = \sqrt{k}(20k)^{2k}$. Theorem 2 completes the proof. $\qquad\square$

To prove Theorem 2, we elaborate further on the graph structures. Let G be a graph as in the statement of Theorem 2. 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 i from 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, from Bukh–Jiang.

Definition 3 (Bukh–Jiang [2])**.** 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 4. 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 paper is organized as follows: In Section 2, 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, which contains our main improvements in this paper, 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 2. In Section 4, we show that either the aforementioned conditions hold, or the levels from exploration expand exponentially. Final computations then prove Theorem 2.

2 Results on Θ-Graphs

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

Lemma 5 (Lemma 2.2 in [7])**.** *Let $k \geqslant 3$. Any bipartite graph H of minimum degree at least k contains a Θ-graph.*

Corollary 6. *Let $k \geqslant 3$. Any bipartite graph H of average degree at least $2k$ contains a Θ-graph.*

Lemma 7 (Claim 3.1 in [7])**.** *Suppose G contains no C_{2k}. For $1 \leqslant i \leqslant k-1$, 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 $\mathrm{ex}(n, C_{2k}) \leqslant O(kn^{1+1/k})$ 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 8. 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 9 (Lemma 10 in [2]). *Suppose G contains no C_{2k}. For $1 \leqslant i \leqslant k-1$, the graph $G[V_{i-1}, V_i, V_{i+1}]$ contains no well-placed Θ-graphs.*

Note that Lemma 9 is analogous to Lemma 7. To prove statements equivalent to Lemma 5 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 paper, 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 $\mathrm{ex}(n, C_{2k})$.

3   Search for Θ-graphs

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

Lemma 10. *Let G be a trilayered graph with layers V_1, V_2, V_3, such that $d_{min}(G) \geqslant 2d + 5k^2$ and $d_{max}(G) \leqslant \Delta d$. If the following conditions hold:*

$$d \cdot e(V_1, V_2) \geqslant 40k \log k |V_3|, \tag{1}$$

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

$$e(V_1, V_2) \geqslant 20(\log k + 1)|V_2| \tag{3}$$

then there is either a Θ-graph in $G[V_1, V_2]$, or 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 10, the rest of this section is organized as follows: In Lemma 11, 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 12, we prove that under the conditions stated in Lemma 10, Lemma 11 can be iterated to show the existence of a desired trilayered subgraphs. Finally, in Lemma 16, 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 11. *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*

$$a \cdot e(V_1, V_2) \geqslant (A + k + 1)|V_1| + B|V_2|. \tag{4}$$

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 $e(V_1, \widetilde{V_2}) \geqslant (1-a)e(V_1, V_2)$, and $|\widetilde{V_2}| \leqslant D|V_3|/d$.*

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 10 is complete, and similarly in case (II) we are done by Lemma 16. In case (III), we found $\widetilde{V_2}$ that shrinks proportionally compared to V_2, while being adjacent to most edges between V_1 and V_2. We can then apply Lemma 11 in $G[V_1, \widetilde{V_2}, V_3]$, 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 12. *Let G be a trilayered graph with layers V_1, V_2, V_3 satisfying conditions (1), (2) and (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 $B \geqslant 5$, and*

$$A \geqslant 2k(\Delta D)^{D-1}, \tag{5}$$
$$(B - 4)D \geqslant 2k. \tag{6}$$

The following proof is an improved version of Bukh–Jiang's proof.

Proof. Assume for the sake of contradiction that neither of the conclusions are true. We will first show that the conditions of Lemma 11 hold for a tuple of well defined A, B and D. Due to our assumptions, the only probable conclusion of Lemma 11 would be (III), which gives us $\widetilde{V_2} \subseteq V_2$. We then iterate this procedure for $t = \log k$ steps on $\widetilde{V_2}$ 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

$$d_i = e(V_1, V_2^{(i)})/|V_2^{(i)}|,$$
$$A_i = a_i e(V_1, V_2^{(i)})/2|V_1| - k - 1,$$
$$B_i = a_i d_i/4 + 5,$$
$$D_i = \lfloor \min(2k, 10k/a_i d_i) \rfloor.$$

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

Claim 13. *For all $i = 0, \cdots, t-1$, we have $d_i < 2k$.*

Proof of Claim 13. If we have $|V_2^{(i)}| \leqslant |V_1|$, then

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

which then implies outcome (I) by Corollary 6. On the other hand, if $|V_2^{(i)}| \geqslant |V_1|$ and $d_i \geqslant 2k$, then

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

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

Now note that A_i, B_i, D_i satisfy constraints (5) and (6). Indeed, (6) follows as long as $d_i < 2k$, and for (5), we have by (2)

$$A_i = a_i e(V_1, V_2^{(i)})/2|V_1| - k - 1 \geqslant \frac{1}{2(t+1)^2} \frac{e(V_1, V_2)}{|V_2|} - k - 1$$

$$\overset{(2)}{\geqslant} 3k(2\Delta k)^{2k-1} - k - 1 \geqslant 2k(\Delta D_i)^{D_i - 1}.$$

Therefore, if we apply Lemma 11, the only possible outcome is (III). The following claim is the key to our iteration process.

Claim 14. *For $V_2^{(i)}$ that is well-defined, condition (4) of Lemma 11 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 11. We have the following invariants:*

$$e(V_1, V_2^{(i+1)}) \geqslant (1 - a_i)e(V_1, V_2^{(i)}), \tag{7}$$

$$d_{i+1} \geqslant a_i d_i \frac{t-i}{t+1} \frac{d \cdot e(V_1, V_2)}{10k|V_3|}. \tag{8}$$

Proof of Claim 14. This proof will proceed by induction. We first show that condition (4) holds for $i = 0$.

$$(A_0 + k + 1)|V_1| + B_i|V_2| = \frac{3}{4}a_0 e(V_1, V_2) + 5|V_2|$$

$$\overset{(3)}{\leqslant} \frac{3}{4}a_0 e(V_1, V_2) + \frac{1}{4(t+1)}e(V_1, V_2) = a_0 e(V_1, V_2).$$

Therefore, we can apply Lemma 11 to obtain $\widetilde{V}_2 \subseteq V_2$ as in outcome (III). Set $V_2^{(1)} = \widetilde{V}_2$. Invariant (7) then follows directly from the conclusions of Lemma 11. For (8), since $|V_2^{(1)}| \leqslant D_0|V_3|/d$, we have

$$d_1 = \frac{e(V_1, V_2^{(1)})}{|V_2^{(1)}|} \geqslant \frac{(1-a_0)e(V_1, V_2)}{D_0|V_3|/d} \geqslant (1-a_0)a_0 d_0 \frac{d \cdot e(V_1, V_2)}{10k|V_3|}.$$

This completes the proof for the base case. For induction, note that iterative application of (7) gives

$$e(V_1, V_2^{(i)}) \geqslant e(V_1, V_2) \prod_{j=0}^{i-1}(1 - a_j) = \frac{t - i + 1}{t + 1}e(V_1, V_2). \tag{9}$$

This inequality helps us show condition (4) again. Indeed,

$$(A_i + k + 1)|V_1| + B_i|V_2^{(i)}| = \frac{3}{4}a_i e(V_1, V_2^{(i)}) + 5|V_2^{(i)}| \leqslant \frac{3}{4}a_i e(V_1, V_2^{(i)}) + 5|V_2|$$

$$\overset{(3)}{\leqslant} \frac{3}{4}a_i e(V_1, V_2^{(i)}) + \frac{1}{4(t+1)}e(V_1, V_2)$$

$$\overset{(9)}{\leqslant} \frac{3}{4}a_i e(V_1, V_2^{(i)}) + \frac{t+1}{4(t+1)(t-i+1)}e(V_1, V_2^{(i)})$$

$$= \frac{3}{4}a_i e(V_1, V_2^{(i)}) + \frac{1}{4}a_i e(V_1, V_2^{(i)}) = a_i e(V_1, V_2^{(i)}).$$

Therefore, by Lemma 11 again, there is a subset $V_2^{(i+1)} \subset V_2^{(i)}$ satisfying (7), and

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

This implies

$$d_{i+1} = \frac{e(V_1, V_2^{(i+1)})}{|V_2^{(i+1)}|} \geqslant \frac{(1-a_i)e(V_1, V_2^{(i)})}{D_i|V_3|/d} \geqslant (1-a_i)a_i d_i \frac{d}{10k|V_3|}e(V_1, V_2^{(i)})$$

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

Therefore invariant (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 (8), we have

$$
\begin{aligned}
d_i &\geqslant d_0 \cdot F^i \prod_{j=0}^{i-1} a_j \frac{t-j}{t+1} = d_0 \cdot F^i \prod_{j=0}^{i-1} \frac{t-j}{t-j+1} \frac{1}{t+1} \\
&= d_0 \cdot \left(\frac{F}{t+1}\right)^i \frac{t-i+1}{t+1} \overset{(1)}{\geqslant} d_0 \cdot 4^i \left(\frac{t}{t+1}\right)^t \frac{t-i+1}{t+1} \\
&\geqslant d_0 \cdot 4^i e^{-1} \frac{t-i+1}{t+1}.
\end{aligned}
$$

Therefore we have

$$
\frac{d_0}{d_i} \leqslant \frac{e \cdot (t+1)}{4^i(t-i+1)}. \tag{10}
$$

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 (7)), while having exponentially large average degree (equation (10)). Similar to Claim 13, we have $d_t < 2k$, which then implies

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

This contradicts condition (3). Therefore we conclude that the iteration must stop before t steps, resulting in either outcome (I) or outcome (II). $\qquad\square$

Remark 15. The contraction rate of $V_2^{(i)}$ could be shown explicitly in the above proof. Specifically,

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

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 16. *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 $B \geqslant 5$, and*

$$
A \geqslant 2k(\Delta D)^{D-1}, (B-4)D \geqslant 2k. \tag{11}
$$

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 $v_0 \longleftrightarrow v_1 \longleftrightarrow \cdots \longleftrightarrow v_l$, 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 17. 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 $P = v_0 \longleftrightarrow v_1 \longleftrightarrow \cdots \longleftrightarrow v_{l-1}$ has been constructed, we wish to extend it to $v_0 \longleftrightarrow \cdots \longleftrightarrow v_l$. We make the following observations.

Claim 18. *For all $i = 0, \cdots, l-1$, v_i cannot have k or more neighbors in $V_2 \cap P$.*

Proof of Claim 18. 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 19. *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 19. If u has neighbors $v_{k_1}, \cdots, v_{k_t}$, where $k_1 < k_2 < \cdots < k_t$, then the path $v_{k_2} \longleftrightarrow u$ and the edge uv_{k_2} form a cycle of length at least

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

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 18, there are at least $A - k$ ways to extend v_{l-1} to another vertex in $V_2 \setminus P$, and Claim 19 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 20. *For $i = 0, 1, \cdots, D-1$, 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,*

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

Proof of Claim 20. 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 18, 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.

Procedure 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 $|M_u| \geqslant \frac{d}{2k+1}$, 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 iteration. Terminate this procedure if $S_1 \cup F_1 = U_i$.

We claim that when this procedure terminates, $|S_1| \geqslant |U_i|/2$. 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,

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

which implies $e(F_1, T)/|T| > 2k$. By Lemma 5, there exists a Θ-graph in $G[V_2, V_3]$, which is a contradiction. Thus $|S_1| \geqslant |U_i|/2$.

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 21. *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,

---

Procedure 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 $P_u vw$. 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 $S_2 \cup F_2 = S_1$.

---

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.

Let $\epsilon = \frac{D-i}{4(2k+1)\Delta}$. We claim that when this procedure terminates, $|S_2| \geqslant \epsilon |S_1| - 2k$. 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 $|F_2| > (1-\epsilon)S_1 + 2k > (1-\epsilon)S_1$, by Procedure 1 and Observation 21, the number of failed extension must be at least

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

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 $|S_2| \geqslant \epsilon |U_i| - 2k$ 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 $|Q'| \geqslant |Q|/(2i+2)$. Let $U_{i+1} \subset D$ be the set of endpoints of these paths, we have

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

which satisfies the stated bound. All of these extensions are good by Claim 19. ∎

Now from condition (11), we see that U_{D-1} is non-empty. Let $Q = v_0 \leftrightsquigarrow \cdots \leftrightsquigarrow v_{l-1} \leftrightsquigarrow u$ be an arbitrary extension with $u \in U_{D-1}$. By Claim 19, $(N(u) \cap V_1) \setminus Q$ 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 22. *The path Qv_l is good.*

Proof of Claim 22. 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 19, 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 $Q' = v_l \leftrightsquigarrow w$ 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 $v_{k_{t-1}} \leftrightsquigarrow w$, 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 $v_{k_{t-1}} \leftrightsquigarrow w$ 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. $\blacksquare$

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

Remark 23. 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 11, 12 and 16, we now prove Lemma 10.

Proof of Lemma 10. Given a graph G satisfying the conditions in Lemma 10, we first apply Lemm 12 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 16 on the resulting trilayered subgraph and our claim holds. This proves Lemma 10. $\square$

We now proceed to prove Theorem 2.

4 Proof of Theorem 2

In this section, we prove that under the conditions of Theorem 2, we have for all i

$$|V_{i+1}| \geqslant (d^2/20k \log k)|V_{i-1}|. \tag{12}$$

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

$$e(V_i, V_{i+1}) \geqslant 2d|V_i|, \tag{13}$$

$$e(V_i, V_{i+1}) \leqslant 2k|V_{i+1}|, \tag{14}$$

$$|V_{i+1}| \geqslant k^{-1}d|V_i|, \tag{15}$$

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

$$e(V_i, V_{i+1}) \geqslant (2d + 5k^2)|V_i| - e(V_{i-1}, V_i) \overset{(14)}{\geqslant} (2d + 5k^2 - 2k)|V_i| \geqslant 2d|V_i|$$

This inequality implies that V_i has average degree at least $2d$ in $G[V_i, V_{i+1}]$. Moreover, if (14) is false, then V_{i+1} has average degree at least $2k$ in $G[V_i, V_{i+1}]$. By Corollary 6, this leads to a contradiction. Therefore (14) is true, and (15) is a consequence of (13) and (14). This completes the proof for the auxiliary claims.

We now move on to prove (12). Assume for the sake of contradiction that (12) is false, we will show that the conditions of Lemma 10 hold, which then leads to a contradiction.

Assume (1) is false. We have

$$2d^2|V_{i-1}| \overset{(13)}{\leqslant} de(V_{i-1}, V_i) \leqslant 40k \log k|V_{i+1}|,$$
$$|V_{i+1}| \geqslant \frac{d^2}{20k \log k}|V_{i-1}|.$$

This contradicts with the assumption that (12) is false.

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

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

Finally, if (3) is false, we have

$$2d|V_{i-1}| \overset{(13)}{\leqslant} e(V_{i-1}, V_i) \leqslant 20(\log k + 1)|V_i| \overset{(15)}{\leqslant} 40 \log k \frac{k}{d}|V_{i+1}|,$$
$$|V_{i+1}| \geqslant \frac{d^2}{20k \log k}|V_{i-1}|.$$

This again implies (12). Therefore (12) hold for all i.

We now conclude the proof of Theorem 2. If k is even, applying (12) $k/2$ times results in

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

If k is odd, applying (12) $(k-1)/2$ times results in

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

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

5 Potential Improvements

This section is dedicated to devoted readers who intend to improve Theorem 2 using our methods.

The idea of using Θ-graphs in the BFS approach originated from Pikhurko's work [7]. The most critical component of this combination is the embedding scheme of a Θ-graph in specific graph structures. For reference, Pikhurko utilized Lemma 5, while Bukh and Jiang and the author utilized different versions of Lemma 16. In essence, all three proofs are driven by their respective embedding methods. Therefore, if one intends to improve the upper bound on $\mathrm{ex}(n, C_{2k})$ following this approach, one shall investigate potential structures and schemes to embed Θ-graphs.

We investigated the following structure in particular.

Definition 24. 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 11 found that if an $[A : B, C : D]$ structure cannot be found, then V_2 can be partitioned into two sets $\widetilde{V_2}$ and $V_2 \setminus \widetilde{V_2}$, such that the former has high density with V_3 and the latter has high density with V_1 (see Appendix). Let $V_C = \widetilde{V_2}$ and $V_B = V_2 \setminus \widetilde{V_2}$, 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 16 and Bukh and Jiang's proof of their Lemma 9, we can prove the following result.

Lemma 25. *Under the same constraints on A, B, D as in Lemma 16, 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 10, 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 11 for the $[A : B, C : (2 : D)]$ structure, but was unable to derive an analog of Lemma 12.

It is also worth pointing out that the constant factor of the upper bound proved by this paper 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 paper, we decided not to present the proof with a modified BFS since we believe the current proof using the Reduction Lemma is cleaner and more applicable to further problems.

Acknowledgement

thor would also like to thank the anonymous referees for many helpful comments and suggestions.

References

[1] J. A. Bondy and M. Simonovits. Cycles of even length in graphs. *J. Combinatorial Theory Ser. B*, 16:97–105, 1974.

[2] 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.

[3] P. Erdős and A. H. Stone. On the structure of linear graphs. *Bull. Amer. Math. Soc.*, 52:1087–1091, 1946.

[4] P. Erdős and M. Simonovits. Some extremal problems in graph theory, combinatorial theory and its applications 1 (proc. colloq. balatonf ured, 1969). 1970.

[5] 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.

[6] Assaf Naor and Jacques Verstraëte. A note on bipartite graphs without 2k-cycles. *Combinatorics, Probability and Computing*, 14(5-6):845–849, 2005.

[7] O. Pikhurko. A note on the Turán function of even cycles. *Proc. Amer. Math. Soc.*, 140(11):3687–3692, 2012. https://homepages.warwick.ac.uk/~maskat/Papers/EvenCycle.pdf.

[8] P. Turán. On an extremal problem in graph theory (in Hungarian). *Mat. Fiz. Lapok*, 48:436–452, 1941.

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 $e(H) \leqslant v(H)^2/2$, it then follows that $v(H)^{1-\alpha/2} \geqslant 2cn^{\alpha/2}$. 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, $e(S \cup T) \geqslant (\frac{\eta}{1-\gamma})e(H)/4 \geqslant \eta e(H)/4$. Let H' be the subgraph of H induced by $S \cup T$. Since

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

we have

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

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 $e(G') \geqslant e(H)/4 \geqslant (c/4)n^{\alpha/2}v(H)^{1+\alpha/2} \geqslant (c/4)v(G)^{1+\alpha/2}$. Finally, since $e(H)/v(H) \geqslant cn^{\alpha/2}v(H)^{\alpha} \geqslant cv(G')^{\alpha}$, we are done. $\qquad\square$

Bukh–Jiang's Proof of Lemma 11 We suppose that alternative (I) does not hold. Then, by Corollary 6, 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 $|E'| \leqslant D|V_3|$. We cannot have $|S| \geqslant |V_1|/k$, for otherwise the average degree of the bipartite graph $G[V_1, S]$ would be at least $\frac{4k}{1+1/k} \geqslant 2k$. So $|S| \leqslant |V_1|/k$.

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

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

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,

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

Assume that the conclusion (III) does not hold with $\widetilde{V}_2 = R \setminus S$. 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

$$\begin{aligned}
e(V_1, V_2) &\leqslant e(V_1, S) + e(V_1, R \setminus S) + A|V_1| + B|V_2| \\
&< (k+1)|V_1| + (1-a)e(V_1, V_2) + A|V_1| + B|V_2|,
\end{aligned}$$

implying that

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

The contradiction with (4) completes the proof. □