

THE MAXIMUM NUMBER OF COMPLETE SUBGRAPHS OF FIXED SIZE IN A GRAPH WITH GIVEN MAXIMUM DEGREE

JONATHAN CUTLER AND A. J. RADCLIFFE

ABSTRACT. In this paper, we make progress on a question related to one of Galvin that has attracted substantial attention recently. The question is that of determining among all graphs G with n vertices and $\Delta(G) \leq r$, which has the most complete subgraphs of size t , for $t \geq 3$. The conjectured extremal graph is $aK_{r+1} \cup K_b$, where $n = a(r+1) + b$ with $0 \leq b \leq r$. Gan, Loh, and Sudakov proved the conjecture when $a \leq 1$, and also reduced the general conjecture to the case $t = 3$. We prove the conjecture for $r \leq 6$ and also establish a weaker form of the conjecture for all r .

1. INTRODUCTION

Galvin [6] conjectured that the graph on n vertices with minimum degree at least d having the most independent sets is $K_{d,n-d}$, provided $n \geq 2d$. He proved the conjecture for $d = 1$, and also for fixed d and large n . Alexander, Mink, and the first author [1] proved this conjecture for bipartite graphs. Further, Galvin conjectured that the same graph has the largest number of independent sets of each size $t \geq 3$, even though it has fewer independent sets of size two than a d -regular graph. Engbers and Galvin [5] proved this stronger conjecture for $d = 2, 3$. Other progress on the stronger conjecture was made by Law and McDiarmid [8] and Alexander and Mink [2]. Quite recently, Gan, Loh, and Sudakov [7] proved this conjecture. We write $i_t(G)$ for the number of independent sets in G of size t .

Theorem 1.1 (Gan, Loh, Sudakov [7]). *If G is a graph with n vertices having $\delta(G) \geq d$ where $n \geq 2d$ then*

$$i_t(G) \leq i_t(K_{d,n-d}),$$

provided $t \geq 3$.

This left unanswered the question of what is the optimal graph when $n < 2d$. In this range it is cleaner to state the problem in its complementary version: counting complete subgraphs in a graph of given maximum degree. In fact, this is the version of the problem on which Gan, Loh, and Sudakov worked. The current authors [4] answered this question for the total number of complete subgraphs. For a graph G , we let $k(G)$ be the number of complete subgraphs (which we also call cliques) in G and $k_t(G)$ be the number of cliques of size t in G .

The first author was sponsored by the National Security Agency under Grant H98230-15-1-0016. The United States Government is authorized to reproduce and distribute reprints notwithstanding any copyright notation herein.

Theorem 1.2 (Cutler, Radcliffe [4]). *For all $n, r \in \mathbb{N}$, write $n = a(r+1) + b$ with $0 \leq b \leq r$. If G is a graph on n vertices with $\Delta(G) \leq r$, then*

$$k(G) \leq k(aK_{r+1} \cup K_b),$$

with equality if and only if $G = aK_{r+1} \cup K_b$, or $r = 2$ and $G = (a-1)K_3 \cup C_4$ or $(a-1)K_3 \cup C_5$.

The case $a = 1$ of this theorem corresponds (in the complement) to Galvin's original conjecture (i.e., $n \geq 2d$). The case $b = 0$ was proved independently by Engbers and Galvin [5] and Wood [10]. What now remains is the level set version to Theorem 1.2, conjectured by Gan, Loh, and Sudakov.

Main Conjecture (Gan, Loh, Sudakov [7]). *For all $n, r \in \mathbb{N}$, write $n = a(r+1) + b$ with $0 \leq b \leq r$. If G is a graph on n vertices with $\Delta(G) \leq r$ and $t \geq 3$ is an integer, then*

$$k_t(G) \leq k_t(aK_{r+1} \cup K_b).$$

Gan, Loh, and Sudakov [7] resolved this conjecture in the case $a \leq 1$. One striking consequence of their approach is a proof that the case $t = 3$, i.e., the case of triangles, implies the general result.

Theorem 1.3 (Gan, Loh, Sudakov [7]). *If the Gan-Loh-Sudakov Conjecture is true for $t = 3$, then it is true for $t > 3$.*

In this paper we make partial progress on the Gan-Loh-Sudakov Conjecture.

Main Theorem. *Let G be a graph with n vertices and $\Delta(G) \leq r$. Write $n = a(r+1) + b$ where $0 \leq b \leq r$.*

(1) *There exists $a' \geq 0$ and some graph H with $n(H) \leq \lfloor 2r/(3\sqrt{3}) \rfloor (r+1) + b$ such that*

$$n = a'(r+1) + n(H) \quad \text{and} \quad k_t(G) \leq k_t(a'K_{r+1} \cup H).$$

(2) *If $r \leq 6$, then $k_3(G) \leq k_3(aK_{r+1} \cup K_b)$, i.e., the Gan-Loh-Sudakov Conjecture is true for $r \leq 6$.*

Following a similar approach to that in [4], we analyze structures in G that we call *clusters*, maximal cliques with maximum-size common neighborhoods. We aim to prove that clusters are either *foldable* or *dischargeable*. In Section 2, we introduce our techniques. In Section 2.1, we discuss a folding operation which performs a local modification, turning a cluster and its neighborhood into a K_{r+1} . A cluster is foldable if this increases the number of triangles. In Section 2.2, we discuss a parameter μ associated with a cluster which determines its foldability. In Section 2.3, we explain a discharging technique for proving that the average number of common neighbors of adjacent vertices is small. Finally, in Section 3, we prove the Main Theorem.

Our notation is standard. We let $[n] = \{1, 2, \dots, n\}$ and, for a set S , we let $\binom{S}{k} = \{X \subseteq S : |X| = k\}$. Also, for sets of vertices U and V , we let $[U, V] = \{uv : u \in U, v \in V\}$.

2. FOLDING AND DISCHARGING

2.1. Folding. Our approach to the proof of the Main Theorem is as follows. We consider a graph G on n vertices with $\Delta(G) \leq r$. We assign weights to the edges of G : if $xy \in E(G)$ then we set

$$w(xy) = |N(x) \cap N(y)|,$$

the number of common neighbors of x and y . Equivalently, $w(xy)$ is the number of triangles of G containing the edge xy . Clearly,

$$k_3(G) = \frac{1}{3} \sum_{xy \in E(G)} w(xy).$$

Thus if the average weight of edges is small then there will not be many triangles in total. Note that, since $\Delta(G) \leq r$, we have $w(xy) \leq r - 1$ for all $xy \in E(G)$.

Definition. An edge $xy \in E(G)$ is called *tight* if $w(xy) = r - 1$. A clique in G , all of whose edges are tight, is called a *tight clique*. A maximal tight clique is a *cluster*.

Note that the edge xy is tight exactly if $N[x] = N[y]$ (where $N[x]$ is the closed neighborhood of x) and $d(x) = d(y) = r$. Hence, the tight edge graph is the disjoint union of clusters of G .

If T is a cluster, we will want to think about its set of common neighbors, $S = \bigcap_{x \in T} N(x)$. Each vertex of T has closed neighborhood $T \cup S$. Vertices in S can have neighbors in T , S , and the rest of the graph. All vertices in T have degree r , but this is not necessarily true for vertices in S .

Definition. Suppose that G is a graph with $\Delta(G) \leq r$ and $T \subseteq V(G)$ is a cluster of size t . We let $S_T = \bigcap_{x \in T} N(x)$ and define a new graph by converting $T \cup S_T$ into a clique (of size $r + 1$) and deleting all the edges $[S_T, V(G) \setminus (T \cup S_T)]$. In other words we define the *folding of G at T* by

$$G_T = G + \binom{S_T}{2} - [S_T, V(G) \setminus (T \cup S_T)].$$

Also, we define

$$R_T = \overline{G[S_T]},$$

i.e., the graph on S_T whose edges are those not in G . Note that since T is a maximal tight clique, $\delta(R) \geq 1$.

Throughout the remainder of the paper, if T is a cluster, we consistently write t for $|T|$, S for S_T , s for $|S_T|$, and R for R_T , when there is no potential for confusion.

Lemma 2.1. *With the notation of the previous definition, for all $x \in S$ we have*

$$|N_G(x) \setminus (T \cup S)| \leq d_R(x),$$

and, in particular, $\Delta(G_T) \leq r$.

Proof. The first inequality is immediate and the second follows from the fact that, after folding, vertices in $T \cup S$ have degree exactly r and no other vertex has its degree increased. \square

We will show that in a wide range of circumstances, we have $k_3(G_T) \geq k_3(G)$. In the rest of this section, we bound the difference $k_3(G_T) - k_3(G)$. We will consider the induced subgraphs of R having three vertices, and their configurations.

Definition. If H is a graph on three vertices, then we let $\#(G, H)$ be the number of 3-subsets of $V(G)$ inducing a graph isomorphic to H .

Lemma 2.2. *If T is a cluster, then*

$$k_3(G_T) - k_3(G) \geq t \cdot e(R) - 2\#(R, \Delta) - \#(R, \Lambda).$$

Proof. We start by proving that

$$(1) \quad k_3(G_T) - k_3(G) \geq t \cdot e(R) + \#(R, \Delta) + \#(R, \Lambda) + \#(R, \text{I.}) - \sum_{v \in S} \binom{d(v)}{2} - \sum_{vw \notin E(R)} \min(d(v), d(w)),$$

where we temporarily write $d(v)$ for $d_R(v)$. To prove (1), we will count exactly the new triangles in G_T and bound the number of triangles lost from G . The triangles gained are exactly triples in $S \cup T$ containing an edge of R . This number is $t \cdot e(R) + \#(R, \Delta) + \#(R, \Lambda) + \#(R, \text{I.})$. Triangles lost from G contain either one vertex of S or two vertices of S not adjacent in R . There are at most $\sum_v \binom{d(v)}{2}$ of the first type since such a triangle contains two G -neighbors of v outside $S \cup T$ that are adjacent in G by Lemma 2.1. Triangles of the second type contain two vertices $v, w \in S$ which are adjacent in G (and so nonadjacent in R) and a common G -neighbor of v and w outside $S \cup T$, of which there are at most $\min(d(v), d(w))$ again by Lemma 2.1.

We will now bound the summation terms by counting configurations of triples. Note first that

$$(2) \quad \sum_{v \in S} \binom{d(v)}{2} = \#(R, \Lambda) + 3\#(R, \Delta).$$

The left hand side counts the number of pairs of incident edges of R . Each such pair of edges determines a unique triple of vertices. Triples of the form Δ contribute three incident pairs; triples of the form Λ contribute one. The other summation term we bound as follows.

$$(3) \quad \sum_{vw \notin E(R)} \min(d(v), d(w)) \leq \frac{1}{2} \sum_{vw \notin E(R)} (d(v) + d(w)) = \#(R, \Lambda) + \#(R, \text{I.}).$$

The first bound is clear. The second equality follows from the fact that the sum counts edge/non-edge incidences, and these incidences each determine a unique triple. Triangles and empty triples have no such incidences. Two-edge and one-edge triples each have two edge/non-edge incidences. Combining (1), (2), and (3), the lemma is proved. \square

Definition. Given a graph G , we let $\mu(G) = 2\#(G, \Delta) + \#(G, \Lambda)$. We call triples inducing Δ , of course, *triangles*, and those inducing Λ , *cherries*.

Corollary 2.3. *The net gain of triangles when folding at T is at least $t \cdot e(R) - \mu(R)$. Hence, if $t \cdot e(R) \geq \mu(R)$, then $k_3(G_T) \geq k_3(G)$.*

Definition. We say that T is *foldable* if $t \cdot e(R) \geq \mu(R)$.

2.2. Extremal results for $\mu(G)$. In this section, we answer two natural extremal questions about the parameter μ which will be useful in the proof of the Main Theorem. Firstly, we prove that given n and m , the maximum value of μ for graphs with n vertices and m edges is attained by the lex graph. Secondly, we prove that the maximum value of μ among graphs with n vertices, $m < n - 1$ edges, and minimum degree at least one is attained by the union of a star of maximum size and a collection of disjoint edges.

Definition. The *lex order* on $\mathcal{P}([n])$ is defined by $A < B$ if $\min(A \triangle B) \in A$. The *lex graph* on n vertices and m edges, denoted $L(n, m)$, is the graph with vertex set $[n]$ and edge set consisting of the initial segment of size m according to the lex order on $\binom{[n]}{2}$.

We will adopt a standard approach to show that the lex graph is extremal. We first prove that μ is maximized by some threshold graph, and then show that among threshold graphs, lex is best. There are many equivalent definitions of threshold graphs (see, e.g., [9]) but in this paper, we find the following most convenient.

Definition. A *threshold graph* is a graph that is constructed inductively from K_1 by successively adding either an isolated vertex or a dominating vertex. We think of the vertex set of a threshold graph on n vertices as $[n]$. We label the vertices “backwards”, so that vertex n is the initial K_1 , vertex $n - 1$ is the first vertex added, and so on. The *code* of a threshold graph with n vertices is the binary sequence σ of length $n - 1$ such that $\sigma_i = 1$ if the vertex i was a dominating vertex when added and $\sigma_i = 0$ if it was an isolate. For a binary sequence $\sigma \in \{0, 1\}^{n-1}$, we let $T(\sigma)$ denote the threshold graph on $[n]$ with code σ . Note that (vertices that correspond to) 1s are adjacent everything to their right and 1s to their left, while 0s are adjacent only to 1s to their left.

We will write an expression such as $1^a 0^b 1^c \dots$ for the sequence consisting of a 1s, b 0s, c 1s, etc. Also, given sequences σ and τ , we write $\sigma.\tau$ for their concatenation. Thus, for instance, $T(1^4 0^5)$ is the join¹ of K_4 and E_6 , whereas $T(0^6 1^3) = K_4 \cup E_6$. For a binary sequence σ , we let $|\sigma|$ be its length. Also, we let $|\sigma|_0$ and $|\sigma|_1$ be the number of 0s and 1s, respectively, in σ .

The lex graph $L(n, m)$ is a threshold graph. Its code is of the form $1^a 0^b 1^x 0^c$ for some $a, b, c \geq 0$ and $x \in \{0, 1\}$. As m increases, 1s march across from the right, accumulating at the left.

Our approach to proving that μ is maximized by some threshold graph is to use a compression that makes a graph “more threshold”. We will show that this compression increases the value of μ . This implies, as in [3], that some threshold graph maximizes μ . Let G be any graph, and let x and y be distinct vertices in G . The choice of x and y defines a natural partition of $V(G \setminus \{x, y\})$ into four parts: vertices that are adjacent only to x , vertices

¹We define the *join* of graphs G and H with disjoint vertex sets, denoted $G \vee H$, to be the graph with vertex set $V(G) \cup V(H)$ and edge set $E(G) \cup E(H) \cup \{xy : x \in V(G), y \in V(H)\}$.

Type of edge	Configurations in G	Configurations in $G_{x \rightarrow y}$
1	$x\text{-}\Delta$	$y\text{-}\Delta$
2	$x\text{-}\wedge$	$y\text{-}\wedge$
3	$x\text{-}\Delta$ and $y\text{-}\wedge$	$y\text{-}\Delta$ and $x\text{-}\wedge$
4	$x\text{-}\wedge$	$y\text{-}\wedge$
5	$x\text{-}\wedge$	$y\text{-}\wedge$
6	—	—
7	$x\text{-}\wedge$ and $y\text{-}\wedge$	$y\text{-}\Delta$
8	—	$y\text{-}\wedge$

TABLE 1. Contributions to μ in G and $G_{x \rightarrow y}$.

adjacent only to y , vertices adjacent to both, and vertices adjacent to neither. We write

$$\begin{aligned} A_{x\bar{y}} &= \{v \in V(G \setminus \{x, y\}) : v \sim x, v \not\sim y\}, \\ A_{xy} &= \{v \in V(G \setminus \{x, y\}) : v \sim x, v \sim y\}, \\ A_{\bar{x}y} &= \{v \in V(G \setminus \{x, y\}) : v \not\sim x, v \sim y\}, \text{ and} \\ A_{\bar{x}\bar{y}} &= \{v \in V(G \setminus \{x, y\}) : v \not\sim x, v \not\sim y\}. \end{aligned}$$

Definition. The *compression of G from x to y* , denoted $G_{x \rightarrow y}$, is the graph obtained from G by deleting all edges between x and $A_{x\bar{y}}$ and adding all edges from y to $A_{x\bar{y}}$.

The following lemma from [3] summarizes a useful property of the compression operator. When we use the lemma, \mathcal{G} will be the set of graphs with n vertices and m edges having $\mu(G)$ maximum.

Lemma 2.4 (Cutler, Radcliffe [3]). *Suppose that \mathcal{G} is a family of graphs on a fixed vertex set V . If \mathcal{G} is closed under all compressions, i.e., for any $G' \in \mathcal{G}$ and any $x, y \in V$ we also have $G'_{x \rightarrow y} \in \mathcal{G}$, then \mathcal{G} contains a threshold graph.*

We now show that compression does not decrease μ .

Lemma 2.5. *If G is a graph with distinct vertices x and y , then $\mu(G) \leq \mu(G_{x \rightarrow y})$.*

Proof. We will compute $\mu(G_{x \rightarrow y}) - \mu(G)$ by counting triangles and cherries that appear in only one of the two graphs. Such triples must consist of a pair with at least one vertex from $A_{x\bar{y}}$ together with x or y . Figure 1 shows almost all of the types of pairs we need to consider, and Table 1 shows the contribution of each kind of pair to both $\mu(G)$ and $\mu(G_{x \rightarrow y})$. In Table 1, we use $x\text{-}\Delta$ to denote a pair that forms a triangle with x , and so on. Note that edges of type 7 contribute to two cherries in G and one triangle in $G_{x \rightarrow y}$, but these combinations make the same contribution to μ . The only type of triple not covered in Table 1 is one of the form axy with $a \in A_{x\bar{y}}$. If $x \not\sim y$, then this triple contributes neither to $\mu(G)$ nor to $\mu(G_{x \rightarrow y})$. If $x \sim y$, the triple axy is a cherry in both G and $G_{x \rightarrow y}$, though with different edges. \square

Theorem 2.6. *If G is a graph with n vertices and m edges, then $\mu(G) \leq \mu(L(n, m))$.*




FIGURE 1. Types of pairs involved in changing triples under compression. The edges drawn are in both G and $G_{x \rightarrow y}$. In G , x is joined to all of $A_{x\bar{y}}$, whereas in $G_{x \rightarrow y}$, those edges are transferred to y .

Proof. Let

$$\mathcal{G} = \{G' : n(G') = n, e(G') = m, \mu(G') \text{ is maximum}\}.$$

By Lemma 2.4, \mathcal{G} contains a threshold graph and so we may assume that $G \in \mathcal{G}$ and is threshold. Say $G = T(\sigma)$. We will show that σ cannot be of the form $\phi.01.\psi.10.\rho$ (even if one or all of ϕ , ψ , or ρ are empty) nor $\phi.01.\psi.1$. If so, we will compare G to $G' = T(\sigma')$ where, in the first case, $\sigma' = \phi.10.\psi.01.\rho$, and in the second, $\sigma' = \phi.10.\psi.0$. The second case is that where the second 01 extends past the end of the code to involve vertex n . Our proof works the same in both cases with $|\rho| = 0$ in the second case. It is easy to check that $e(G') = e(G)$, and we will show that $\mu(G') > \mu(G)$. We will then show that the only threshold graphs whose codes do not contain these patterns are lex graphs.

Our first job is to show $\mu(G') \geq \mu(G)$. Many of the triples that contribute to $\mu(G)$ and $\mu(G')$ effectively stay the same between the two graphs. In σ , the 01 is in positions a and $a + 1$ where $a = |\phi| + 1$ and the 10 is in positions b and $b + 1$ where $b = |\phi| + |\psi| + 3$. If we define the involution $f = (a \ a + 1)(b \ b + 1)$, the product of two transpositions, then typically a triple Q is a triangle or cherry in G exactly if $f(Q)$ is a triangle or cherry in G' . We compute the difference $\mu(G') - \mu(G)$ by enumerating the exceptions. In fact, we have $G' = f(G) + \{a, a + 1\} - \{b, b + 1\}$. Thus all triangles or cherries gained or lost must contain either $\{a, a + 1\}$ or $\{b, b + 1\}$. Considering triangles first, there are triangles in G of the form $Q = \{x, b, b + 1\}$ where $\sigma_x = 1$ and $x < b$ for which $f(Q)$ is not a triangle in G' . In compensation, there are triangles $Q' = \{y, a, a + 1\}$ where $\sigma'_y = 1$ and $y < a$ for which $f(Q')$

is not a triangle in G . Hence, the contribution of triangles to $\mu(G') - \mu(G)$ is

$$-2(|\phi|_1 + 1 + |\psi|_1 - |\phi|_1) = -2(1 + |\psi|_1).$$

As for cherries, G has cherries of the form $R = \{b, b + 1, x\}$ where $x > b + 1$ and also cherries of the form $R = \{y, a, a + 1\}$ where $\sigma_y = 1$ and $y < a$, for which $f(R)$ is not a cherry in G' . Again, in compensation, there are cherries in G' of the form $\{y, b, b + 1\}$ with $\sigma'_y = 1$ and $y < b$ and $\{a, a + 1, x\}$ with $x > a + 1$ whose images under f are not cherries in G . Thus, the contribution to $\mu(G') - \mu(G)$ is

$$(|\psi| + 2 + |\rho|) - |\rho| + (|\phi|_1 + 1 + |\psi|_1) - |\phi|_1 = |\psi| + |\psi|_1 + 3.$$

Combining the two contributions, we see that

$$\mu(G') - \mu(G) = |\psi| + |\psi|_1 + 3 - 2(1 + |\psi|_1) = |\psi|_0 + 1 > 0.$$

It only remains to show that if σ is not of either of the forms above, then $T(\sigma)$ is the lex graph. We will call a 01 substring a *rise*, and a 10 substring a *fall*. We have shown that σ does not have a rise followed by a (disjoint) fall, nor does it have a rise followed by a final 1. We wish to show that σ is the code of a lex graph. Suppose first that σ has no rises; in that case σ is of the form $1^a 0^b$ for some non-negative a, b and hence G is a lex graph. Otherwise we look for the first rise, and split σ there: $\sigma = \phi'.01.\psi' = \phi.\psi$, where $\phi = \phi'.0$ and $\psi = 1.\psi'$. Since ϕ does not contain a rise it is of the form $1^a 0^b$ with $a \geq 0, b \geq 1$. Also, by the condition on σ we know that ψ' does not contain a fall, nor does it end with a 1. This implies that $\psi = 0^c$ for some $c \geq 0$. Hence $\sigma = 1^a 0^b 10^c$ is the code for a lex graph. \square

In our application to our proof of the Main Theorem, the graphs for which we want to bound μ are required to have minimum degree at least one.

Theorem 2.7. *If G is a graph with n vertices, $m < n - 1$ edges, and minimum degree at least one (so that $m \geq n/2$), then $\mu(G) \leq \mu(K_{1,p} \cup qK_2)$ where $p = 2m - n + 1$ and $q = n - m - 1$.*

Proof. We may suppose that G is a graph satisfying these conditions maximizing μ and having the least number of components. If G has components C_1, C_2, \dots, C_k , then $\mu(G) = \sum_{i=1}^k \mu(C_i)$. Thus, by Theorem 2.6, we may assume that each component is a lex graph. Our first aim is to prove that we may suppose that every component is a star. It cannot be that each component is a non-tree, for then G would have at least n edges. Suppose, then, that C_1 is a non-tree component and C_2 is a tree component (and so therefore a star). We have $e(C_1) + e(C_2) \geq n(C_1) + n(C_2) - 1$, so we have enough edges to build a connected lex graph L on $V(C_1) \cup V(C_2)$. Thus, $\mu(L) \geq \mu(C_1) + \mu(C_2)$ and we have decreased the number of components, a contradiction.

Lastly, we will show that all but at most one component of G is an edge. Suppose $k, \ell \geq 2$ and G contains components $K_{1,k}$ and $K_{1,\ell}$. We can replace these with $K_{1,k+\ell-1} \cup K_2$ and strictly increase μ , since

$$\binom{k}{2} + \binom{\ell}{2} < \binom{k+\ell-1}{2}.$$

This contradiction finishes the proof. \square

2.3. Discharging. In the proof of the Main Theorem, one case occurs when the average weight of all edges of G is at most $r - 2$. In this section, we give a technique to bound the weight of edges around a cluster. The only edges having weight $r - 1$ are, of course, in clusters. It is convenient to let the *benefit* of an edge be defined by $\bar{w}(xy) = r - 2 - w(xy)$. An edge uv with u in a cluster T and $v \in S = S_T$ will have benefit $d_R(v) - 1$. Note that if e is a tight edge, then $\bar{w}(e) = -1$ and all other edges have $\bar{w}(e)$ nonnegative.

Such an edge may be incident with two clusters, so we will use the benefit of these edges by transferring $\frac{1}{2}(d_R(v) - 1)$ from such an edge to the edges in T . This “discharging” will sometimes be enough to show that average weight of an edge in G is at most $r - 2$, i.e., that the average benefit is nonnegative. This is sufficient to show that $k_3(G) \leq k_3(aK_{r+1} \cup K_b)$, as we prove in the final section.

Definition. A cluster T is *dischargeable* if the associated graph R satisfies

$$\sum_{v \in S} (d_R(v) - 1) \geq t - 1,$$

i.e., if $2e(R) \geq s + t - 1$.

Lemma 2.8. *If every cluster of G is dischargeable, then the average weight of all edges is at most $r - 2$, i.e., the total benefit of all edges is nonnegative.*

Proof. We partition the edges of G into those in clusters, those associated with clusters (meaning they are incident to a vertex in a cluster, but not in a cluster), and others. Edges not incident to clusters all have nonnegative benefit. Given a cluster T , writing $e \sim T$ if e is associated with T , we compute the benefit associated with that cluster as the benefit of the edges in the cluster plus half of the benefit of the edges associated with it. We have

$$\begin{aligned} \sum_{e \in E(G[T])} \bar{w}(e) + \frac{1}{2} \sum_{e \sim T} \bar{w}(e) &= -\binom{t}{2} + \frac{1}{2} \sum_{v \in S} t(d_R(v) - 1) \\ &= \frac{t}{2} \left(-(t - 1) + \sum_{v \in S} (d_R(v) - 1) \right). \end{aligned}$$

Thus, if T is dischargeable, the benefit associated with T is nonnegative. Since any edge is associated with at most two clusters, we have, summing over clusters T ,

$$\begin{aligned} \sum_{e \in E(G)} \bar{w}(e) &\geq \sum_T \left(\sum_{e \in E(G[T])} \bar{w}(e) + \frac{1}{2} \sum_{e \sim T} \bar{w}(e) \right) \\ &\geq \sum_T \frac{|T|}{2} \left(-(|T| - 1) + \sum_{v \in S_T} (d_{R_T}(v) - 1) \right), \end{aligned}$$

and if every cluster is dischargeable then every term in this final sum is non-negative. \square

3. PROOF OF MAIN THEOREM

In this section, we first prove that every cluster is either dischargeable or foldable.

Lemma 3.1. *If T is a cluster in G , then T is either dischargeable or foldable.*

In order to prove this, we need a couple of preliminary lemmas.

Lemma 3.2. *If T is a cluster with $R = K_{1,p} \cup qK_2$, then T is either dischargeable or foldable.*

Proof. Note that we have $s = p + 1 + 2q$ and $e(R) = p + q$. Also, $\mu(R) = \binom{p}{2}$. Then T is dischargeable if

$$2(p + q) = 2e(R) \geq s + t - 1 = p + 2q + t.$$

This holds when $t \leq p$. If not, i.e., if $t > p$, then

$$t \cdot e(R) = t(p + q) > \binom{p}{2} = \mu(R),$$

and so T is foldable. \square

Lemma 3.3. *If G is any graph, then $\mu(G) \leq \frac{2}{3}(n - 2)e(G)$. In particular, if T is a cluster satisfying $t \geq \frac{2}{3}(s - 2)$, then T is foldable.*

Proof. Let G be a graph with vertex set V and edge set E . For $Q \in \binom{V}{3}$, we define $\mu(Q)$ by $\mu(\Delta) = 2$, $\mu(\wedge) = 1$, and $\mu(\mathcal{I}) = \mu(\cdot) = 0$. That is, $\mu(Q)$ is the contribution of Q to $\mu(G)$. Also, we let $e(Q) = e(G[Q])$ and $E(Q) = E(G[Q])$. We compute as follows.

$$\mu(G) = \sum_{Q \in \binom{V}{3}} \mu(Q) = \sum_{\substack{Q \in \binom{V}{3} \\ e(Q) \neq 0}} \frac{1}{e(Q)} \sum_{e \in E(Q)} \mu(Q) = \sum_{e \in E} \sum_{Q \supseteq e} \frac{\mu(Q)}{e(Q)} \leq \frac{2}{3}(n - 2)e(G).$$

The inequality comes from the fact that $\mu(Q)/e(Q) \leq 2/3$ for every triple Q containing edges.

Now, for the second claim, we have $\mu(R) \leq \frac{2}{3}(s - 2)e(R)$. Therefore, by hypothesis,

$$t \cdot e(R) \geq \frac{2}{3}(s - 2)e(R) \geq \mu(R),$$

and so T is foldable. \square

We are now ready to prove our main lemma.

Proof of Lemma 3.1. By Lemma 3.3, we are done if $t \geq \frac{2}{3}(s - 2)$. Also, T is dischargeable if $2e(R) \geq s + t - 1$. Thus, in the remaining cases, we have

$$e(R) < \frac{s + t - 1}{2} < \frac{5}{6}s - \frac{7}{6}.$$

Let $m = e(R) < s - 1$. We know that whether T is foldable depends only on t , m and $\mu(R)$, whereas whether T is dischargeable depends only on t and m . By Theorem 2.7, since $\delta(R) \geq 1$, we have that $\mu(R) \leq \mu(K_{1,p} \cup qK_2)$ where $q = s - m - 1$ and $p = 2m - s + 1$. Thus, if T would be foldable with $R = K_{1,p} \cup qK_2$, then T is foldable. Therefore, T is foldable (dischargeable) if T is foldable (dischargeable) with $R = K_{1,p} \cup qK_2$. By Lemma 3.2, we are done. \square

It only remains to show that, provided $a \geq 2r/(3\sqrt{3})$, G must contain a cluster, i.e., G must contain an edge of weight $r - 1$.

Lemma 3.4. *If $a \geq 2r/(3\sqrt{3})$, then*

$$\sum_{e \in E(aK_{r+1} \cup K_b)} w(e) > (r-2) \frac{r(a(r+1)+b)}{2},$$

i.e., any graph of maximum degree at most r on $a(r+1)+b$ vertices either has fewer triangles than $aK_{r+1} \cup K_b$ or has an edge of weight $r-1$.

Proof. The proof is a straightforward calculation. □

We are now ready to prove the main result of the paper.

Proof of Main Theorem. It suffices to prove that, if G is a graph with $\Delta(G) \leq r$ having the greatest number of triangles, then, provided $a \geq 2r/(3\sqrt{3})$, G contains a copy of K_{r+1} . By Lemma 3.4, the average weight of edges in G is such that some edge must have weight $r-1$. Thus, G must have a cluster, and by Theorem 3.1, all clusters are foldable or dischargeable. If all clusters were dischargeable, the average weight would be at most $r-2$, contradicting Lemma 3.4. If G has a foldable cluster T , then $k_3(G) \leq k_3(G_T)$ and G_T contains a K_{r+1} so we are done by induction.

We now consider the case when $r \leq 6$. We apply our techniques to the graph H guaranteed by the above, so let $n = n(H)$ and redefine a and b as usual. It suffices to show that if $r \leq 6$ and H is a graph on at most $2r(r+1)/(3\sqrt{3}) + r$ vertices having average weight of edges at most $r-2$, then $k_3(H) \leq k_3(aK_{r+1} \cup K_b)$. I.e., we need to verify

$$(r-2) \frac{r(a(r+1)+b)}{6} \leq a \binom{r+1}{3} + \binom{b}{3},$$

when $1 < a < 2r/(3\sqrt{3})$. The lower bound follows from the result of Gan, Loh, and Sudakov [7]. For $r \leq 5$, there are no such a . When $r = 6$, we need to verify the inequality for $a = 2$ and $0 \leq b \leq 6$, which is straightforward. Note that the inequality fails $r = 7$, $a = 2$, and $b = 4$, where our proof technique breaks down. □

REFERENCES

1. James Alexander, Jonathan Cutler, and Tim Mink, *Independent sets in graphs with given minimum degree*, Electron. J. Combin. **19** (2012), no. 3, Paper 37, 11.
2. James Alexander and Tim Mink, *A new method for enumerating independent sets of a fixed size in general graphs*, to appear in J. Graph Theory (2015).
3. Jonathan Cutler and A. J. Radcliffe, *Extremal graphs for homomorphisms II*, J. Graph Theory **76** (2014), 42–59.
4. ———, *The maximum number of complete subgraphs in a graph with given maximum degree*, J. Combin. Theory Ser. B **104** (2014), 60–71.
5. John Engbers and David Galvin, *Counting independent sets of a fixed size in graphs with a given minimum degree*, J. Graph Theory **76** (2014), no. 2, 149–168.
6. David Galvin, *Two problems on independent sets in graphs*, Discrete Math. **311** (2011), no. 20, 2105–2112.
7. Wenyang Gan, Po-Shen Loh, and Benny Sudakov, *Maximizing the number of independent sets of a fixed size*, Combin. Probab. Comput. **24** (2015), no. 3, 521–527.
8. Hiu-Fai Law and Colin McDiarmid, *On independent sets in graphs with given minimum degree*, Combin. Probab. Comput. **22** (2013), no. 6, 874–884.

9. N. V. R. Mahadev and U. N. Peled, *Threshold graphs and related topics*, Annals of Discrete Mathematics, vol. 56, North-Holland Publishing Co., Amsterdam, 1995.
10. David R. Wood, *On the maximum number of cliques in a graph*, Graphs Combin. **23** (2007), no. 3, 337–352.

DEPARTMENT OF MATHEMATICAL SCIENCES, MONTCLAIR STATE UNIVERSITY, MONTCLAIR, NJ 07043
E-mail address: jonathan.cutler@montclair.edu

DEPARTMENT OF MATHEMATICS, UNIVERSITY OF NEBRASKA-LINCOLN, LINCOLN, NE
E-mail address: aradcliffe1@math.unl.edu