

Counting independent sets of a fixed size in graphs with a given minimum degree

John Engbers* David Galvin†

February 19, 2014

Abstract

Galvin showed that for all fixed δ and sufficiently large n , the n -vertex graph with minimum degree δ that admits the most independent sets is the complete bipartite graph $K_{\delta, n-\delta}$. He conjectured that except perhaps for some small values of t , the same graph yields the maximum count of independent sets of size t for each possible t . Evidence for this conjecture was recently provided by Alexander, Cutler, and Mink, who showed that for all triples (n, δ, t) with $t \geq 3$, no n -vertex *bipartite* graph with minimum degree δ admits more independent sets of size t than $K_{\delta, n-\delta}$.

Here we make further progress. We show that for all triples (n, δ, t) with $\delta \leq 3$ and $t \geq 3$, no n -vertex graph with minimum degree δ admits more independent sets of size t than $K_{\delta, n-\delta}$, and we obtain the same conclusion for $\delta > 3$ and $t \geq 2\delta + 1$. Our proofs lead us naturally to the study of an interesting family of critical graphs, namely those of minimum degree δ whose minimum degree drops on deletion of an edge or a vertex.

1 Introduction and statement of results

An independent set (a.k.a. stable set) in a graph is a set of vertices spanning no edges. For a graph $G = (V, E)$ (always assumed to be simple and finite in this paper), denote by $i(G)$ the number of independent sets in G . In [15] this quantity is referred to as the *Fibonacci number* of G , motivated by the fact that for the path graph P_n its value is a Fibonacci number. It has also been studied in the field of molecular chemistry, where it is referred to as the *Merrifield-Simmons index* of G [13].

*john.engbers@marquette.edu; Department of Mathematics, Statistics and Computer Science, Marquette University, Milwaukee, WI 53201.

†dgalvin1@nd.edu; Department of Mathematics, University of Notre Dame, Notre Dame IN 46556. Supported by National Security Agency grants H98230-10-1-0364 and H98230-13-1-0248 and by the Simons Foundation.

A natural extremal enumerative question is the following: as G ranges over some family \mathcal{G} , what is the maximum value attained by $i(G)$, and which graphs achieve this maximum? This question has been addressed for numerous families. Proding and Tichy [15] considered the family of n -vertex trees, and showed that the maximum is uniquely attained by the star $K_{1,n-1}$. Granville, motivated by a question in combinatorial group theory, raised the question for the family of n -vertex, d -regular graphs (see [2] for more details). An approximate answer – $i(G) \leq 2^{n/2(1+o(1))}$ for all such G , where $o(1) \rightarrow 0$ as $d \rightarrow \infty$ – was given by Alon in [2], and he speculated a more exact result, that the maximizing graph, at least in the case $2d|n$, is the disjoint union of $n/2d$ copies of $K_{d,d}$. This speculation was confirmed for *bipartite* G by Kahn [10], and for general regular G by Zhao [18]. The family of n -vertex, m -edge graphs was considered by Cutler and Radcliffe in [5], and they observed that it is a corollary of the Kruskal-Katona theorem that the lex graph $L(n, m)$ (on vertex set $\{1, \dots, n\}$, with edges being the first m pairs in lexicographic order) maximizes $i(G)$ in this class. Zykov [19] considered the family of graphs with a fixed number of vertices and fixed independence number, and showed that the maximum is attained by the complement of a certain Turán graph. (Zykov was actually considering cliques in a graph with given clique number, but by complementation this is equivalent to considering independent sets in a graph with given independence number.) Other papers addressing questions of this kind include [9], [11], [14] and [16].

Having resolved the question of maximizing $i(G)$ for G in a particular family, it is natural to ask which graph maximizes $i_t(G)$, the number of independent sets of size t in G , for each possible t . For many families, it turns out that the graph which maximizes $i(G)$ also maximizes $i_t(G)$ for all t . Wingard [17] showed this for trees, Zykov [19] showed this for graphs with a given independence number (see [4] for a short proof), and Cutler and Radcliffe [4] showed this for graphs on a fixed number of edges (again, as a corollary of Kruskal-Katona). In [10], Kahn conjectured that for all $2d|n$ and all t , no n -vertex, d -regular graph admits more independent sets of size t than the disjoint union of $n/2d$ copies of $K_{d,d}$; this conjecture remains open, although asymptotic evidence appears in [3].

The focus of this paper is the family $\mathcal{G}(n, \delta)$ of n -vertex graphs with minimum degree δ . One might imagine that, since removing edges increases the count of independent sets, the graph in $\mathcal{G}(n, \delta)$ that maximizes the count of independent sets would be δ -regular (or close to), but this turns out not to be the case, even for $\delta = 1$. The following result is from [7].

Theorem 1.1. *For $n \geq 2$ and $G \in \mathcal{G}(n, 1)$, we have $i(G) \leq i(K_{1,n-1})$. For $\delta \geq 2$, $n \geq 4\delta^2$ and $G \in \mathcal{G}(n, \delta)$, we have $i(G) \leq i(K_{\delta, n-\delta})$.*

What about maximizing $i_t(G)$ for each t ? The family $\mathcal{G}(n, \delta)$ is an example of a family for which the maximizer of the total count is *not* the maximizer for each individual t . Indeed, consider the case $t = 2$. Maximizing the number of independent

sets of size two is the same as minimizing the number of edges, and it is easy to see that for all fixed δ and sufficiently large n , there are n -vertex graphs with minimum degree at least δ which have fewer edges than $K_{\delta, n-\delta}$ (consider for example a δ -regular graph, or one which has one vertex of degree $\delta + 1$ and the rest of degree δ). However, we expect that anomalies like this occur for very few values of t . Indeed, the following conjecture is made in [7].

Conjecture 1.2. *For each $\delta \geq 1$ there is a $C(\delta)$ such that for all $t \geq C(\delta)$, $n \geq 2\delta$ and $G \in \mathcal{G}(n, \delta)$, we have*

$$i_t(G) \leq i_t(K_{\delta, n-\delta}) = \binom{n-\delta}{t} + \binom{\delta}{t}.$$

The case $\delta = 1$ of Conjecture 1.2 is proved in [7], with $C(1)$ as small as it possibly can be, namely $C(1) = 3$. In [1], Alexander, Cutler and Mink looked at the subfamily $\mathcal{G}^{\text{bip}}(n, \delta)$ of *bipartite* graphs in $\mathcal{G}(n, \delta)$, and resolved the conjecture in the strongest possible way for this family.

Theorem 1.3. *For $\delta \geq 1$, $n \geq 2\delta$, $t \geq 3$ and $G \in \mathcal{G}^{\text{bip}}(n, \delta)$, we have $i_t(G) \leq i_t(K_{\delta, n-\delta})$.*

This provides good evidence for the truth of the strongest possible form of Conjecture 1.2, namely that we may take $C(\delta) = 3$.

The purpose of this paper is to make significant progress towards this strongest possible conjecture. We completely resolve the cases $\delta = 2$ and 3, and for larger δ we deal with all but a small fraction of cases. In what follows, E_k is the graph on k vertices with no edges, and $G_1 \vee G_2$ indicates the join of graphs G_1 and G_2 .

Theorem 1.4. *1. For $\delta = 2$, $t \geq 3$ and $G \in \mathcal{G}(n, 2)$, we have $i_t(G) \leq i_t(K_{2, n-2})$. For $n \geq 5$ and $3 \leq t \leq n - 2$ we have equality iff $G = H \vee E_{n-2}$, where H is any graph on two vertices.*

2. For $\delta = 3$, $t \geq 3$ and $G \in \mathcal{G}(n, 3)$, we have $i_t(G) \leq i_t(K_{3, n-3})$. For $n \geq 6$ and $t = 3$ we have equality iff $G = K_{3, n-3}$; for $n \geq 7$ and $4 \leq t \leq n - 3$ we have equality iff $G = H \vee E_{n-3}$, where H is any graph on three vertices.

3. For $\delta \geq 3$, $t \geq 2\delta + 1$ and $G \in \mathcal{G}(n, \delta)$, we have $i_t(G) \leq i_t(K_{\delta, n-\delta})$. For $n \geq 3\delta + 1$ and $2\delta + 1 \leq t \leq n - \delta$ we have equality iff $G = H \vee E_{n-\delta}$, where H is any graph on δ vertices.

(Note that there is some overlap between parts 2 and 3 above.) Part 1 above provides an alternate proof of the $\delta = 2$ case of the total count of independent sets, originally proved in [7].

Corollary 1.5. *For $n \geq 4$ and $G \in \mathcal{G}(n, 2)$, we have $i(G) \leq i(K_{2, n-2})$. For $n = 4$ and $n \geq 6$ there is equality iff $G = K_{2, n-2}$.*

Proof. The result is trivial for $n = 4$. For $n = 5$, it is easily verified by inspection, and we find that both C_5 and $K_{2,3}$ have the same total number of independent sets. So we may assume $n \geq 6$.

We clearly have $i(K'_{2,n-2}) < i(K_{2,n-2})$, where $K'_{2,n-2}$ is the graph obtained from $K_{2,n-2}$ by joining the two vertices in the partition class of size 2. For all $G \in \mathcal{G}(n, 2)$ different from both $K_{2,n-2}$ and $K'_{2,n-2}$, Theorem 1.4 part 1 tells us that $i_t(G) \leq i_t(K_{2,n-2}) - 1$ for $3 \leq t \leq n - 2$. For $t = 0, 1, n - 1$ and n we have $i_t(G) = i_t(K_{2,n-2})$ (with the values being 1, n , 0 and 0, respectively). We have $i_2(G) \leq \binom{n}{2} - n$ (this is the number of non-edges in a 2-regular graph), and so

$$i_2(G) \leq i_2(K_{2,n-2}) + \binom{n}{2} - n - \binom{n-2}{2} - 1 = i_2(K_{2,n-2}) + n - 4. \quad (1)$$

Putting all this together we get $i(G) \leq i(K_{2,n-2})$.

If G is not 2-regular then we have strict inequality in (1) and so $i(G) < i(K_{2,n-2})$. If G is 2-regular, then (as we will show presently) we have $i_3(G) < i_3(K_{2,n-2}) - 1$ and so again $i(G) < i(K_{2,n-2})$. To see the inequality concerning independent sets of size 3 note that in any 2-regular graph the number of independent sets of size 3 that include a fixed vertex v is the number of non-edges in the graph induced by the $n - 3$ vertices $V \setminus \{v, x, y\}$ (where x and y are the neighbors of v), which is at most $\binom{n-3}{2} - (n - 4)$. It follows that

$$i_3(G) \leq \frac{1}{3} \left(n \left(\binom{n-3}{2} - (n-4) \right) \right) < \binom{n-2}{3} - 1.$$

□

The paper is laid out as follows. In Section 2 we make some easy preliminary observations that will be used throughout the rest of the paper, and we introduce the ideas of *ordered independent sets* and *min-critical graphs*. In Section 3 we deal with the case $\delta = 2$ (part 1 of Theorem 1.4). We begin Section 4 with the proof of part 3 of Theorem 1.4, and then explain how the argument can be improved (within the class of min-critical graphs). This improvement will be an important ingredient in the $\delta = 3$ case (part 2 of Theorem 1.4) whose proof we present in Section 5. Finally we present some concluding remarks and conjectures in Section 6.

Notation: Throughout the paper we use $N(v)$ for the set of vertices adjacent to v , and $d(v)$ for $|N(v)|$. We write $u \sim v$ to indicate that u and v are adjacent (and $u \not\sim v$ to indicate that they are not). We use $G[Y]$ to denote the subgraph induced by a subset Y of the vertices, and $E(Y)$ for the edge set of this subgraph. Finally, for $t \in \mathbb{N}$ we use $x^{\underline{t}}$ to indicate the falling power $x(x-1)\dots(x-(t-1))$.

2 Preliminary remarks

For integers n , δ and t , let $P(n, \delta, t)$ denote the statement that for every $G \in \mathcal{G}(n, \delta)$, we have $i_t(G) \leq i_t(K_{\delta, n-\delta})$. A key observation (Lemma 2.1 and Corollary 2.2 below) is that if we prove $P(n, \delta, t)$ for some triple (n, δ, t) with $t \geq \delta + 1$, we automatically have $P(n, \delta, t + 1)$. The proof introduces the idea of considering *ordered* independent sets, that is, independent sets in which an order is placed on the vertices.

Lemma 2.1. *For $\delta \geq 2$ and $t \geq \delta + 1$, if $G \in \mathcal{G}(n, \delta)$ satisfies $i_t(G) \leq i_t(K_{\delta, n-\delta})$ then $i_{t+1}(G) \leq i_{t+1}(K_{\delta, n-\delta})$. Moreover, if $t < n - \delta$ and $i_t(G) < i_t(K_{\delta, n-\delta})$ then $i_{t+1}(G) < i_{t+1}(K_{\delta, n-\delta})$.*

Corollary 2.2. *For $\delta \geq 2$ and $t \geq \delta + 1$, $P(n, \delta, t) \Rightarrow P(n, \delta, t + 1)$.*

Proof. Fix $G \in \mathcal{G}(n, \delta)$. By hypothesis, the number of ordered independent sets in G of size t is at most $(n - \delta)^t$. For each ordered independent set of size t in G there are at most $n - (t + \delta)$ vertices that can be added to it to form an ordered independent set of size $t + 1$ (no vertex of the independent set can be chosen, nor can any neighbor of any particular vertex in the independent set). This leads to a bound on the number of ordered independent sets in G of size $t + 1$ of $(n - \delta)^t(n - (t + \delta)) = (n - \delta)^{t+1}$. Dividing by $(t + 1)!$, we find that $i_{t+1}(G) \leq \binom{n-\delta}{t+1} = i_{t+1}(K_{\delta, n-\delta})$.

If we have $i_t(G) < \binom{n-\delta}{t}$ then we have strict inequality in the count of ordered independent sets of size t , and so also as long as $n - (\delta + t) > 0$ we have strict inequality in the count for $t + 1$, and so $i_{t+1}(G) < \binom{n-\delta}{t+1}$. \square

Given Corollary 2.2, in order to prove $P(n, \delta, t)$ for $n \geq n(\delta)$ and $t \geq t(\delta)$ it will be enough to prove $P(n, \delta, t(\delta))$. Many of our proofs will be by induction on n , and will be considerably aided by the following simple observation.

Lemma 2.3. *Fix $t \geq 3$. Suppose we know $P(m, \delta, t)$ for all $m < n$. Let $G \in \mathcal{G}(n, \delta)$ be such that there is $v \in V(G)$ with $G - v \in \mathcal{G}(n - 1, \delta)$ (that is, $G - v$ has minimum degree δ). Then $i_t(G) \leq i_t(K_{\delta, n-\delta})$. Equality can only occur if all of 1) $i_t(G - v) = i_t(K_{\delta, n-1-\delta})$, 2) $G - v - N(v)$ is empty (has no edges), and 3) $d(v) = \delta$ hold.*

Proof. Counting first the independent sets of size t in G that do not include v and then those that do, and bounding the former by our hypothesis on $P(m, \delta, t)$ and

the latter by the number of subsets of size $t - 1$ in $G - v - N(v)$, we have

$$\begin{aligned}
i_t(G) &= i_t(G - v) + i_{t-1}(G - v - N(v)) \\
&\leq i_t(K_{\delta, n-1-\delta}) + i_{t-1}(E_{n-1-d(v)}) \\
&\leq \binom{n-1-\delta}{t} + \binom{\delta}{t} + \binom{n-1-\delta}{t-1} \\
&= \binom{n-\delta}{t} + \binom{\delta}{t} \\
&= i_t(K_{\delta, n-\delta}).
\end{aligned}$$

The statement concerning equality is evident. \square

Lemma 2.3 allows us to focus on graphs with the property that each vertex has a neighbor of degree δ . Another simple lemma further restricts the graphs that must be considered.

Lemma 2.4. *If G' is obtained from G by deleting edges, then for each t we have $i_t(G) \leq i_t(G')$.*

This leads to the following definition.

Definition 2.5. Fix $\delta \geq 1$. A graph G with minimum degree δ is *edge-min-critical* if for any edge e in G , the minimum degree of $G - e$ is $\delta - 1$. It is *vertex-min-critical* if for any vertex v in G , the minimum degree of $G - v$ is $\delta - 1$. If it is both edge-min-critical and vertex-min-critical, we say that G is *min-critical*.

Lemmas 2.3 and 2.4 allow us to concentrate mostly on min-critical graphs. In Section 3 (specifically Lemma 3.2) we obtain structural information about min-critical graphs in the case $\delta = 2$, while much of Section 5 is concerned with the same problem for $\delta = 3$.

An easy upper bound on the number of independent sets of size $t \geq 1$ in a graph with minimum degree δ is

$$i_t(G) \leq \frac{n(n - (\delta + 1))(n - (\delta + 2)) \cdots (n - (\delta + (t - 1)))}{t!}. \quad (2)$$

This bound is tight only when it simultaneously happens that all vertices have degree δ and every pair of non-adjacent vertices share the same δ neighbors. These two conditions do not tend to occur simultaneously, and we will improve the upper bound in (2) by considering carefully what happens when one or both of them fail. To begin this process, it will be helpful to distinguish between vertices with degree δ and those with degree larger than δ . Set

$$V_{=\delta} = \{v \in V(G) : d(v) = \delta\}$$

and

$$V_{>\delta} = \{v \in V(G) : d(v) > \delta\}.$$

The main thrust of most of our proofs is that a min-critical graph must have at least one of a small list of different structures in it, and we exploit the presence of a structure to significantly improve the easy upper bound.

3 Proof of Theorem 1.4, part 1 ($\delta = 2$)

Recall that we want to show that for $\delta = 2$, $t \geq 3$ and $G \in \mathcal{G}(n, 2)$, we have $i_t(G) \leq i_t(K_{2, n-2})$, and that for $n \geq 5$ and $3 \leq t \leq n - 2$ we have equality iff $G = K_{2, n-2}$ or $K'_{2, n-2}$ (obtained from G by joining the two vertices in the partition class of size 2). We concern ourselves initially with the inequality, and discuss the cases of equality at the end. By Corollary 2.2, it is enough to consider $t = 3$, and we will prove this case by induction on n , the base cases $n \leq 5$ being trivial. So from here on we assume that $n > 5$ and that $P(m, 2, 3)$ has been established for all $m < n$, and let $G \in \mathcal{G}(n, 2)$ be given. By Lemmas 2.3 and 2.4 we may assume that G is min-critical.

We begin with two lemmas, the first of which is well-known (see e.g. [8]), and the second of which gives structural information about min-critical graphs (in the case $\delta = 2$).

Lemma 3.1. *Let $k \geq 1$ and $0 \leq t \leq k + 1$. In the k -path P_k we have*

$$i_t(P_k) = \binom{k+1-t}{t}.$$

Let $k \geq 3$ and $0 \leq t \leq k - 1$. In the k -cycle C_k we have

$$i_t(C_k) = \binom{k-t}{t} + \binom{k-t-1}{t-1}.$$

Lemma 3.2. *Fix $\delta = 2$. Let G be a connected n -vertex min-critical graph. Either*

1. G is a cycle or
2. $V(G)$ can be partitioned as $Y_1 \cup Y_2$ with $2 \leq |Y_1| \leq n - 3$ in such a way that Y_1 induces a path, Y_2 induces a graph with minimum degree 2, each endvertex of the path induced by Y_1 has exactly one edge to Y_2 , the endpoints of these two edges in Y_2 are either the same or non-adjacent, and there are no other edges from Y_1 to Y_2 (see Figure 1).

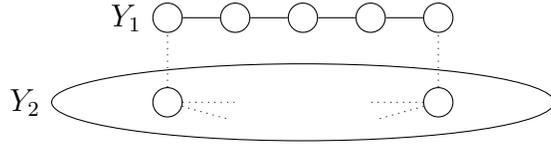


Figure 1: An example of a partition of $V(G)$ from Lemma 3.2.

Proof. If G is not a cycle, then it has some vertices of degree greater than 2. If there is exactly one such vertex, say v , then by parity considerations $d(v)$ is even and at least 4. Since all degrees are even, the edge set may be partitioned into cycles. Take any cycle through v and remove v from it to get a path whose vertex set can be taken to be Y_1 .

There remains the case when G has at least two vertices with degree larger than 2. Since G is edge-min-critical, $V_{>\delta}$ forms an independent set and so there is a path on at least 3 vertices joining distinct vertices $v_1, v_2 \in V_{>\delta}$, all of whose internal vertices u_1, \dots, u_k have degree 2 (the shortest path joining two vertices in $V_{>\delta}$ would work). Since G is vertex-min-critical we must in fact have $k \geq 2$, since otherwise u_1 would be a vertex whose deletion leaves a graph with minimum degree 2. We may now take $Y_1 = \{u_1, \dots, u_k\}$. Note that the Y_2 endpoints (v_1 and v_2) of the edges from u_1 and u_k to Y_2 are both in $V_{>\delta}$ and so are non-adjacent. \square

Armed with Lemmas 3.1 and 3.2 we now show that for min-critical G we have

$$i_3(G) < i_3(K_{2,n-2}) = \binom{n-2}{3}.$$

If G is the n -cycle, then we are done by Lemma 3.1. If G is a disjoint union of cycles, then choose one such, of length k , and call its vertex set Y_1 , and set $Y_2 = V \setminus Y_1$. We will count the number of independent sets of size 3 in G by considering how the independent set splits across Y_1 and Y_2 .

For $k \geq 4$ Lemma 3.1 gives a count of $\binom{k-3}{3} + \binom{k-4}{2}$ for the number of independent sets of size 3 in Y_1 , and this is still a valid upper bound when $k = 3$. By induction there are at most $\binom{n-k-2}{3}$ independent sets of size 3 in Y_2 . There are $(\binom{k-1}{2} - 1)(n-k)$ independent sets with two vertices in Y_1 and one in Y_2 (the first factor here simply counting the number of non-edges in a k -cycle). Finally, there are $k(\binom{n-k-1}{2} - 1)$ independent sets with one vertex in Y_1 and two in Y_2 (the second factor counting the number of non-edges in a 2-regular graph on $n-k$ vertices). The sum of these bounds is easily seen to be $\binom{n-2}{3} - k$, so strictly smaller than $\binom{n-2}{3}$.

We may now assume that G has a component that is not 2-regular. Choose one such component. Let Y_1 be as constructed in Lemma 3.2 and let Y_2 be augmented by including the vertex sets of all other components. Denote by v_1, v_2 the neighbors in Y_2 of the endpoints of the path. Note that it is possible that $v_1 = v_2$, but if not then

by Lemma 3.2 we have $v_1 \approx v_2$. We will bound $i_3(G)$ from above by considering the possible splitting of independent sets across Y_1 and Y_2 .

By Lemma 3.1, there are $\binom{k-2}{3}$ independent sets of size 3 in Y_1 , and by induction there are at most $\binom{n-k-2}{3}$ independent sets of size 3 in Y_2 .

The number of independent sets of size 3 in G that have two vertices in Y_1 and one in Y_2 is at most

$$\binom{k-3}{2} (n-k) + \left(\binom{k-1}{2} - \binom{k-3}{2} \right) (n-k-1).$$

The first term above counts those independent sets in which neither endpoint of the k -path is among the two vertices from Y_1 , and uses Lemma 3.1. The second term bounds from above the number of independent sets in which at least one endpoint of the k -path is among the two vertices from Y_1 , and again uses Lemma 3.1. (Note that when $k=2$ the application of Lemma 3.1 is not valid, since when we remove the endvertices we are dealing with a path of length 0, outside the range of validity of the lemma; however, the above bound is valid for $k=2$, since it equals 1 in this case.) Finally, the number of independent sets of size 3 in G that have one vertex in Y_1 and two in Y_2 is at most

$$(k-2) \left(\binom{n-k}{2} - |E(Y_2)| \right) + \sum_{i=1}^2 \left(\binom{n-k-1}{2} - |E(Y_2)| + d_{Y_2}(v_i) \right).$$

The first term here counts the number of independent sets in which the one vertex from Y_1 is not an endvertex, the second factor being simply the number of non-edges in $G[Y_2]$. The second term counts those with the vertex from Y_1 being the neighbor of v_i , the second factor being the number of non-edges in $G[Y_2] - v_i$.

The sum of all of these bounds, when subtracted from $\binom{n-2}{3}$, simplifies to

$$-(k-1)n + k^2 + k - 3 + k|E(Y_2)| - d_{Y_2}(v_1) - d_{Y_2}(v_2), \quad (3)$$

a quantity which we wish to show is strictly positive.

Suppose first that Y_1 can be chosen so that $v_1 \neq v_2$. Recall that in this case $v_1 \approx v_2$, so $d_{Y_2}(v_1) + d_{Y_2}(v_2) \leq |E(Y_2)|$. Combining this with $|E(Y_2)| \geq n-k$ we get that (3) is at least $2k-3$, which is indeed strictly positive for $k \geq 2$.

If $v_1 = v_2 = v$, then we first note that

$$|E(Y_2)| = \frac{1}{2} \sum_{w \in Y_2} d_{Y_2}(w) \geq \frac{d_{Y_2}(v)}{2} + (n-k-1)$$

(since $G[Y_2]$ has minimum degree 2). Inserting into (3) we find that (3) is at least

$$n-3 + \left(\frac{k}{2} - 2 \right) d_{Y_2}(v). \quad (4)$$

This is clearly strictly positive for $k \geq 4$, and for $k = 3$ strict positivity follows from $d_{Y_2}(v) < 2(n - 3)$, which is true since in fact $d_{Y_2}(v) < n - 3$ in this case.

If $k = 2$, then (4) is strictly positive unless $d_{Y_2}(v) = n - 3$ (the largest possible value it can take in this case). There is just one min-critical graph G with the property that for all possible choices of Y_1 satisfying the conclusions of Lemma 3.2 we have $|Y_1| = 2$, $v_1 = v_2 = v$ and $d_{Y_2}(v) = n - 3$; this is the windmill graph (see Figure 2) consisting of $(n - 1)/2$ triangles with a single vertex in common to all the triangles, and otherwise no overlap between the vertex sets (note that the degree condition on v forces G to be connected). A direct count gives $(n - 1)(n - 3)(n - 5)/6 < \binom{n-2}{3}$ independent sets of size 3 in this particular graph.

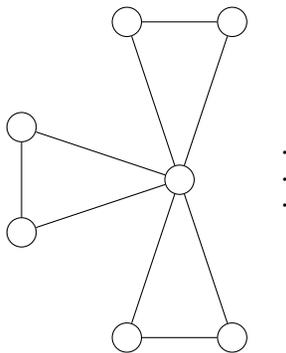


Figure 2: The windmill graph.

This completes the proof that $i_t(G) \leq i_t(K_{2,n-2})$ for all $t \geq 3$ and $G \in \mathcal{G}(n, 2)$. We now turn to considering the cases where equality holds in the range $n \geq 5$ and $3 \leq t \leq n - 2$. For $t = 3$ and $n = 5$, by inspection we see that we have equality iff $G = K_{2,3}$ or $K'_{2,3}$ (obtained from $K_{2,3}$ by adding an edge inside the partition class of size 2). For larger n , we prove by induction that equality can be achieved only for these two graphs. If a graph G is not edge-min-critical, we delete edges until we obtain a graph G' which is edge-min-critical, using Lemma 2.4 to get $i_t(G) \leq i_t(G')$. If G' is min-critical, then the discussion in this section shows that we cannot achieve equality. If G' is not vertex-min-critical, Lemma 2.3 and our induction hypothesis shows that we only achieve equality for G' if there is $v \in V(G')$ with $G' - v = K_{2,n-3}$ or $K'_{2,n-3}$, $G' - v - N(v)$ empty, and $d(v) = 2$. First, notice that $G' - v = K'_{2,n-3}$ implies that G' is not edge-min-critical, so equality can only occur when $G' - v = K_{2,n-3}$. If $G' - v = K_{2,n-3}$, the second and third conditions tell us that $N(v)$ is exactly the partition class of size 2 in $K_{2,n-3}$, that is, that $G' = K_{2,n-2}$. From here it is evident that equality can only occur for $G = K_{2,n-2}$ or $K'_{2,n-2}$.

Now for each fixed $n \geq 5$, we conclude from Lemma 2.1 that for $3 \leq t \leq n - 2$ we *cannot* have equality unless $G = K_{2,n-2}$ or $K'_{2,n-2}$; and since the equality is trivial for these two cases, the proof is complete.

4 Proof of Theorem 1.4, part 3 ($\delta \geq 3$)

Throughout this section we set $h = |V_{>\delta}|$ and $\ell = |V_{=\delta}|$; note that $h + \ell = n$. We begin this section with the proof of Theorem 1.4 part 3; we then show how the method used may be improved to obtain a stronger result within the class of min-critical graphs (Lemma 4.1 below), a result which will play a role in the proof of Theorem 1.4, part 2 ($\delta = 3$) that will be given in Section 5.

Recall that we are trying to show that for $\delta \geq 3$, $t \geq 2\delta + 1$ and $G \in \mathcal{G}(n, \delta)$, we have $i_t(G) \leq i_t(K_{\delta, n-\delta})$, and that for $n \geq 3\delta + 1$ and $2\delta + 1 \leq t \leq n - \delta$ there is equality iff G is obtained from $K_{\delta, n-\delta}$ by adding some edges inside the partition class of size δ . As with Theorem 1.4 part 1 we begin with the inequality and discuss cases of equality at the end.

By Corollary 2.2 it is enough to consider $t = 2\delta + 1$. We prove $P(n, \delta, 2\delta + 1)$ by induction on n . For $n < 3\delta + 1$ the result is trivial, since in this range all $G \in \mathcal{G}(n, \delta)$ have $i_t(G) = 0$. It is also trivial for $n = 3\delta + 1$, since the only graphs G in $\mathcal{G}(n, \delta)$ with $i_t(G) > 0$ in this case are those that are obtained from $K_{\delta, n-\delta}$ by the addition of some edges inside the partition class of size δ , and all such G have $i_t(G) = 1 = i_t(K_{\delta, n-\delta})$. So from now on we assume $n \geq 3\delta + 2$ and that $P(m, \delta, 2\delta + 1)$ is true for all $m < n$, and we seek to establish $P(n, \delta, 2\delta + 1)$.

By Lemmas 2.3 and 2.4 we may restrict attention to G which are min-critical. To allow the induction to proceed, we need to show that the number of ordered independent sets of size $2\delta + 1$ in G is at most $(n - \delta)^{2\delta+1}$.

We do so by constructing an ordered independent set $(v_1, \dots, v_{2\delta+1})$ sequentially, at each step keeping track of how many choices there are for the next vertex. We break the count into cases, according to whether v_1 comes from $V_{>\delta}$ or $V_{=\delta}$. In the first case ($v_1 \in V_{>\delta}$) there are at most

$$\begin{aligned} h(n - (\delta + 2))(n - (\delta + 3)) \cdots (n - (3\delta + 1)) &= \frac{h}{n} (n(n - (\delta + 2))^{2\delta}) \\ &< \frac{h}{n} (n - \delta)^{2\delta+1} \end{aligned} \quad (5)$$

ordered independent sets of size $2\delta + 1$ (recall $|V_{>\delta}| = h$), since once v_1 has been chosen there are at most $n - (\delta + 2)$ choices for v_2 , then at most $n - (\delta + 3)$ choices for v_3 , and so on.

In the second case ($v_1 \in V_{=\delta}$) there are, by the same reasoning, at most

$$\ell(n - (\delta + 1))(n - (\delta + 2)) \cdots (n - 2\delta)$$

choices for the ordered independent set $(v_1, \dots, v_{2\delta+1})$ (recall $|V_{=\delta}| = \ell$). The key observation now is that since G is vertex-min-critical there can be at most $\delta - 1$ vertices distinct from v_1 with the same neighborhood as v_1 . It follows that one of v_2

through $v_{\delta+1}$ has a neighbor w outside of $N(v_1)$. Since w cannot be included in the independent set, there are at most

$$(n - (2\delta + 2))(n - (2\delta + 3)) \cdots (n - (3\delta + 1))$$

choices for $(v_{\delta+2}, \dots, v_{2\delta+1})$. Combining these bounds, there are at most

$$\frac{\ell}{n} \left(n \frac{(n - (\delta + 1))^{2\delta+1}}{n - (2\delta + 1)} \right) < \frac{\ell}{n} (n - \delta)^{2\delta+1}$$

ordered independent sets of size $2\delta + 1$ that begin with a vertex from $V_{=\delta}$. Combining with (5) we get $i_{2\delta+1}(G) < (n - \delta)^{2\delta+1}/(2\delta + 1)!$, as required.

This completes the proof that $i_t(G) \leq i_t(K_{\delta, n-\delta})$ for all $t \geq 2\delta + 1$ and $G \in \mathcal{G}(n, \delta)$. We now turn to considering the cases where equality holds in the range $n \geq 3\delta + 1$ and $2\delta + 1 \leq t \leq n - \delta$. For $t = 2\delta + 1$ and $n = 3\delta + 1$, we clearly have equality iff G is obtained from $K_{\delta, 2\delta+1}$ by adding some edges inside the partition class of size δ . For larger n , we prove by induction that equality can be achieved only for a graph of this form. If a graph G is not edge-min-critical, we delete edges until we obtain a graph G' which is edge-min-critical, using Lemma 2.4 to get $i_t(G) \leq i_t(G')$. If G' is min-critical, then the discussion in this section shows that we cannot achieve equality. If G' is not vertex-min-critical, Lemma 2.3 and our induction hypothesis shows that we only achieve equality for G' if there is $v \in V(G')$ with $G' - v$ obtained from $K_{\delta, n-\delta-1}$ by adding some edges inside the partition class of size δ , $G' - v - N(v)$ empty, and $d(v) = \delta$. First, notice that the cases where $G' - v \neq K_{\delta, n-\delta-1}$ imply that G' is not edge-min-critical, so in fact equality can only occur when $G' - v = K_{\delta, n-\delta-1}$. Since $d(v) = \delta$ the neighborhood of v cannot include all of the partition class of size $n - 1 - \delta$. If it fails to include a vertex of the partition class of size δ , there must be an edge in $G - v - N(v)$; so in fact, $N(v)$ is exactly the partition class of size δ and $G' = K_{\delta, n-\delta}$. From here it is evident that equality can only occur for G obtained from $K_{\delta, n-\delta}$ by adding some edges inside the partition class of size δ .

Now for each fixed $n \geq 3\delta + 1$, we conclude from Lemma 2.1 that for $2\delta + 1 \leq t \leq n - \delta$ we *cannot* have equality unless G is obtained from $K_{\delta, n-\delta}$ by adding some edges inside the partition class of size δ ; and since the equality is trivial in these cases, the proof is complete.

The ideas introduced here to bound the number of ordered independent sets in a min-critical graph can be modified to give a result that covers a slightly larger range of t , at the expense of requiring n to be a little larger. Specifically we have the following:

Lemma 4.1. *For all $\delta \geq 3$, $t \geq \delta + 1$, $n \geq 3.2\delta$ and vertex-min-critical $G \in \mathcal{G}(n, \delta)$, we have $i_t(G) < i_t(K_{\delta, n-\delta})$. For $\delta = 3$ and $t = 4$ we get the same conclusion for vertex-min-critical $G \in \mathcal{G}(n, 3)$ with $n \geq 8$.*

Remark. The constant 3.2 has not been optimized here, but rather chosen for convenience.

Proof. By Lemma 2.1 it is enough to consider $t = \delta + 1$. The argument breaks into two cases, depending on whether G has at most $\delta - 2$ vertices with degree larger than m (a parameter to be specified later), or at least $\delta - 1$. The intuition is that in the former case, after an initial vertex v has been chosen for an ordered independent set, many choices for the second vertex should have at least two neighbors outside of $N(v)$, which reduces subsequent options, whereas in the latter case, an initial choice of one of the at least $\delta - 1$ vertices with large degree should lead to few ordered independent sets.

First suppose that G has at most $\delta - 2$ vertices with degree larger than m . Just as in (5), a simple upper bound on the number of ordered independent sets of size t whose first vertex is in $V_{>\delta}$ is

$$\frac{h}{n} (n(n - (\delta + 2))(n - (\delta + 3)) \cdots (n - (2\delta + 1))) < \frac{h}{n} (n - \delta)^{\delta+1}. \quad (6)$$

There are ℓ choices for the first vertex v of an ordered independent set that begins with a vertex from $V_{=\delta}$. For each such v , we consider the number of extensions to an ordered independent set of size $\delta + 1$. This is at most

$$x(n - (\delta + 2))^{\delta-1} + y(n - (\delta + 3))^{\delta-1} + z(n - (\delta + 4))^{\delta-1} \quad (7)$$

where x is the number of vertices in $V(G) \setminus (\{v\} \cup N(v))$ that have no neighbors outside $N(v)$, y is the number with one neighbor outside $N(v)$, and z is the number with at least 2 neighbors outside $N(v)$. Note that $x + y + z = n - \delta - 1$, and that by vertex-min-criticality $x \leq \delta - 1$.

Let u_1 and u_2 be the two lowest degree neighbors of v . By vertex-min-criticality and our assumption on the number of vertices with degree greater than m , the sum of the degrees of u_1 and u_2 is at most $\delta + m$. Each vertex counted by y is adjacent to either u_1 or u_2 , so counting edges incident to u_1 and u_2 there are at most $m + \delta - 2x - 2$ such vertices.

For fixed x we obtain an upper bound on (7) by taking y as large as possible, so we should take $y = m + \delta - 2x - 2$ and $z = n - m - 2\delta + x + 1$. With these choices of y and z , a little calculus shows that we obtain an upper bound by taking x as large as possible, that is, $x = \delta - 1$. This leads to an upper bound on the number of ordered independent sets of size t whose first vertex is in $V_{=\delta}$ of

$$\ell \left((\delta - 1)(n - (\delta + 2))^{\delta-1} + (m - \delta)(n - (\delta + 3))^{\delta-1} + (n - m - \delta)(n - (\delta + 4))^{\delta-1} \right). \quad (8)$$

Combining with (6) we see that we are done (for the case G has at most $\delta - 2$ vertices with degree larger than m) as long as we can show that (8) is strictly less

than $\ell(n - \delta)^{\delta+1}/n$, or equivalently that

$$\begin{aligned} & n(\delta - 1)(n - (\delta + 2))(n - (\delta + 3)) + \\ & n(m - \delta)(n - (\delta + 3))(n - (2\delta + 1)) + < (n - \delta)^4. \end{aligned} \quad (9)$$

$$n(n - m - \delta)(n - (2\delta + 1))(n - (2\delta + 2))$$

We will return to this presently; but first we consider the case where G has at least $\delta - 1$ vertices with degree larger than m . An ordered independent set of size $\delta + 1$ in this case either begins with one of $\delta - 1$ vertices of largest degree, in which case there are strictly fewer than $(n - m - 1)^\delta$ extensions, or it begins with one of the remaining $n - \delta + 1$ vertices. For each such vertex v in this second case, the second vertex chosen is either one of the $k = k(v) \leq \delta - 1$ vertices that have the same neighborhood as v , in which case there are at most $(n - (\delta + 2))^{\delta-1}$ extensions, or it is one of the $n - d(v) - 1 - k$ vertices that have a neighbor that is not a neighbor of v , in which case there are at most $(n - (\delta + 3))^{\delta-1}$ extensions. We get an upper bound on the total number of extensions in this second case (starting with a vertex not among the $\delta - 1$ of largest degree) by taking k as large as possible and $d(v)$ as small as possible; this leads to a strict upper bound on the number of ordered independent sets of size $\delta + 1$ in the case G has at least $\delta - 1$ vertices with degree larger than m of

$$(\delta - 1)(n - m - 1)^\delta + (n - \delta + 1) \left((\delta - 1)(n - (\delta + 2))^{\delta-1} + (n - 2\delta)(n - (\delta + 3))^{\delta-1} \right).$$

We wish to show that this is at most $(n - \delta)^{\delta+1}$. As long as $m \geq \delta$ we have $n - m - i \leq n - \delta - i$, and so what we want is implied by

$$\begin{aligned} & (\delta - 1)(n - m - 1)(n - m - 2) + \\ & (n - \delta + 1)(\delta - 1)(n - (\delta + 2)) + \leq (n - \delta)^3. \end{aligned} \quad (10)$$

$$(n - \delta + 1)(n - 2\delta)(n - (2\delta + 1))$$

Setting $m = n/2$, we find that for $\delta \geq 3$, both (9) and (10) hold for all $n \geq 3.2\delta$. Indeed, in both cases at $n = 3.2\delta$ the right-hand side minus the left-hand side is a polynomial in δ (a quartic in the first case and a cubic in the second) that is easily seen to be positive for all $\delta \geq 3$; and in both cases we can check that for each fixed $\delta \geq 3$, when viewed as a function of n the right-hand side minus the left-hand side has positive derivative for all $n \geq 3.2\delta$. This completes the proof of the first statement. It is an easy check that both (9) and (10) hold for all $n \geq 8$ in the case $\delta = 3$, completing the proof of the lemma. \square

5 Proof of Theorem 1.4, part 2 ($\delta = 3$)

Recall that we are trying to show that for $\delta = 3$, $t \geq 3$ and $G \in \mathcal{G}(n, 3)$, we have $i_t(G) \leq i_t(K_{3,n-3})$, and that for $n \geq 6$ and $t = 3$ we have equality iff $G = K_{3,n-3}$,

while for $n \geq 7$ and $4 \leq t \leq n - 3$ we have equality iff G is obtained from $K_{3,n-3}$ by adding some edges inside the partition class of size 3.

For $t = 4$ and $n \geq 7$ we prove the result (including the characterization of uniqueness) by induction on n , with the base case $n = 7$ trivial. For $n \geq 8$, Lemma 4.1 gives strict inequality for all vertex-min-critical G , so we may assume that we are working with a G which is not vertex-min-critical. Lemma 2.3 now gives the inequality $i_4(G) \leq i_4(K_{3,n-3})$, and the characterization of cases of equality goes through exactly as it did for Theorem 1.4 parts 1 and 3. The result for larger t (including the characterization of uniqueness) now follows from Lemma 2.1.

For $t = 3$, we also argue by induction on n , with the base case $n = 6$ trivial. For $n \geq 7$, if G is not vertex-min-critical then the inequality $i_3(G) \leq i_3(K_{3,n-3})$ follows from Lemma 2.3, and the fact that there is equality in this case only for $G = K_{3,n-3}$ follows exactly as it did in the proofs of Theorem 1.4 parts 1 and 3. So we may assume that G is vertex-min-critical. We will also assume that G is edge-min-critical (this assumption is justified because in what follows we will show $i_3(G) < i_3(K_{3,n-3})$, and restoring the edges removed to achieve edge-min-criticality maintains the strictness of the inequality). Our study of min-critical 3-regular graphs will be based on a case analysis that adds ever more structure to the G under consideration. A useful preliminary observation is the following.

Lemma 5.1. *Fix $\delta = 3$. If a min-critical graph G has a vertex w of degree $n - 3$ or greater, then $i_3(G) < i_3(K_{3,n-3})$.*

Proof. If $d(w) > n - 3$ then there are no independent sets of size 3 containing w , and by Theorem 1.4 part 1 the number of independent sets of size 3 in $G - w$ (a graph of minimum degree 2) is at most $\binom{n-3}{3} < i_3(K_{3,n-3})$. If $d(w) = n - 3$ and the two non-neighbors of w are adjacent, then we get the same bound. If they are not adjacent (so there is one independent set of size 3 containing w) and $G - w$ is not extremal among minimum degree 2 graphs for the count of independent sets of size 3, then we also get the same bound, since now $i_3(G - w) \leq \binom{n-3}{3} - 1$. If $G - w$ is extremal it is either $K_{2,n-3}$ or $K'_{2,n-3}$, and in either case w must be adjacent to everything in the partition class of size $n - 3$ (to ensure that G has minimum degree 3), and then, since the non-neighbors of w are non-adjacent, it must be that $G = K_{3,n-3}$, a contradiction since we are assuming that G is min-critical (recall that $n \geq 7$ so $n - 3 \geq 4$). \square

5.1 Regular G

If G is 3-regular then we have $i_3(G) < \binom{n-3}{3} + 1$. We see this by considering ordered independent sets of size 3. Given an initial vertex v , we extend to an ordered independent set of size 3 by adding ordered non-edges from $V \setminus (N(v) \cup \{v\})$. Since G is 3-regular there are $3n$ ordered edges in total, with at most 18 of them

adjacent either to v or to something in $N(v)$. This means that the number of ordered independent sets of size 3 in G is at most

$$n((n-4)(n-5) - (3n-18)) < (n-3)(n-4)(n-5) + 6$$

with the inequality valid as long as $n \geq 7$. So from here on we may assume that G is not 3-regular, or equivalently that $V_{>3} \neq \emptyset$.

Remark. The argument above generalizes to show that δ -regular graphs have at most $\binom{n-\delta}{3} + \binom{\delta}{3}$ independent sets of size 3, with equality only possible when $n = 2\delta$.

Let $v \in V(G)$ have a neighbor in $V_{>3}$. By min-criticality $d(v) = 3$. Let w_1, w_2 , and w_3 be the neighbors of v , listed in decreasing order of degree, so $d(w_1) = d$, $d(w_2) = x$ and $d(w_3) = 3$ satisfy $3 \leq x \leq d \leq n-4$, the last inequality by Lemma 5.1 as well as $d > 3$ (see Figure 3).

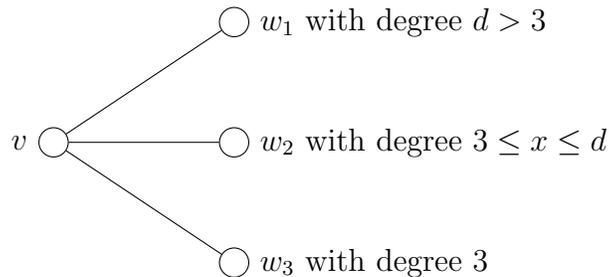


Figure 3: The generic situation from the end of Section 5.1 on.

5.2 No edge between w_3 and w_2

We now proceed by a case analysis that depends on the value of x as well as on the set of edges present among the w_i 's. The first case we consider is $w_3 \not\sim w_2$. In this case we give upper bounds on the number of independent sets of size 3 which contain v and the number which do not. There are $\binom{n-4}{2} - |E(Y)|$ independent sets of size 3 which include v , where $Y = V \setminus (N(v) \cup \{v\})$. We bound $|E(Y)|$ from below by putting a lower bound on the sum of the degrees in Y and then subtracting off the number of edges from Y to $\{v\} \cup N(v)$. This gives

$$|E(Y)| \geq \frac{3(n-4) - 2 - (d-1) - (x-1)}{2} = \frac{3(n-4) - x - d}{2}. \quad (11)$$

To bound the number of independent sets of size 3 which don't include v , we begin by forming G' from G by deleting v and (to restore minimum degree 3) adding an edge between w_3 and w_2 (we will later account for independent sets in G that contain both

w_2 and w_3). The number of independent sets of size 3 in G' is, by induction, at most $i_3(K_{3,n-4})$. But in fact, we may assume that the count is strictly smaller than this. To see this, note that if we get exactly $i_3(K_{3,n-4})$, then by induction $G' = K_{3,n-4}$. For $n = 7$ this forces G to have a vertex of degree 4 and so $i_3(G) < i_3(K_{3,4})$ by Lemma 5.1. For $n > 7$, w_3 must be in the partition class of size $n - 4$ in G' (to have degree 3) so since $w_2 \sim w_3$ (in G'), w_2 must be in the partition class of size 3. To avoid creating a vertex of degree $n - 3$ in G , w_1 must be in the partition class of size $n - 4$. But then all other vertices in the partition class of size $n - 4$ only have neighbors of degree $n - 4$ (in G), contradicting min-criticality.

So we may now assume that the number of independent sets of size 3 in G which do not include v is at most

$$\binom{n-4}{3} + (n-x-2), \quad (12)$$

the extra $n - x - 2$ being an upper bound on the number of independent sets of size 3 in G that include both w_3 and w_2 . Combining (11) and (12) we find that in this case

$$i_3(G) \leq \binom{n-4}{2} - \frac{3(n-4) - x - d}{2} + \binom{n-4}{3} + (n-x-2). \quad (13)$$

As long as $d < n + x - 6$ this is strictly smaller than $i_3(K_{3,n-3})$. Since $x \geq 3$ and $d < n - 3$, this completes the case $w_3 \approx w_2$.

5.3 Edge between w_3 and w_2 , no edge between w_3 and w_1 , degree of w_2 large

The next case we consider is $w_3 \sim w_2$, $w_3 \approx w_1$, and $x > 3$. In this case we can run an almost identical argument to that of Section 5.2, this time adding the edge from w_1 to w_3 when counting the number of independent sets of size 3 that don't include v . We add 1 to the right-hand side of (11) (to account for the fact that there is now only one edge from w_3 to Y instead of 2, and only $x - 2$ edges from w_2 to Y instead of $x - 1$) and replace (12) with $\binom{n-4}{3} + 1 + (n - d - 2)$ (the 1 since in this case we do not need strict inequality in the induction step). Using $x \leq d$ in this latter expression, we get the same inequality as (13).

5.4 Edge between w_3 and w_2 , edge between w_3 and w_1 , degree of w_2 large

Next we consider the case $w_3 \sim w_2$, $w_3 \sim w_1$, and $x > 3$. Here we must have $w_1 \approx w_2$, since otherwise G would not be edge-min-critical. The situation is illustrated in Figure 4. To bound $i_3(G)$, we consider v and w_3 . Arguing as in Section 5.2 (around

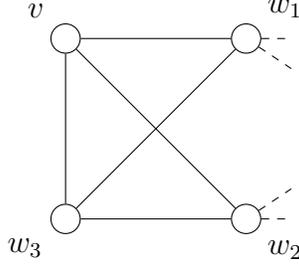


Figure 4: The situation in Section 5.4.

(11)), the number of independent sets including one of v , w_3 is at most

$$2 \left(\binom{n-4}{2} - \frac{3(n-4) - (d-2) - (x-2)}{2} \right).$$

To obtain an upper bound on the number of independent sets including neither v nor w_3 , we delete both vertices, add an edge from w_1 to w_2 (to restore minimum degree 3) and use induction to get a bound of

$$\binom{n-5}{3} + 1 + (n-d-2)$$

(where $n-d-2$ bounds from above the number of independent sets containing both w_1 and w_2). Since $x \leq n-2$ the sum of these two bounds is strictly smaller than $i_3(K_{3,n-3})$.

5.5 None of the above

If there is no v of degree 3 that puts us into one of the previous cases, then every v of degree 3 that has a neighbor w_1 of degree strictly greater than 3 may be assumed to have two others of degree 3, w_2 and w_3 say, with vw_2w_3 a triangle (see Figure 5).

Since every neighbor of a vertex of degree greater than 3 has degree exactly 3 (by min-criticality) it follows that for every w_1 of degree greater than 3, every neighbor of w_1 is a vertex of a triangle all of whose vertices have degree 3. We claim that two of these triangles must be vertex disjoint. Indeed, if w_1 has two neighbors a and b with $a \sim b$ then the triangles associated with a and b must be the same, and by considering degrees we see that the triangle associated with any other neighbor of w_1 must be vertex disjoint from it. If a and b are not adjacent and their associated triangles have no vertex in common, then we are done; but if they have a vertex in common then (again by considering degrees) they must have two vertices in common, and the triangle associated with any other neighbor of w_1 must be vertex disjoint from both.

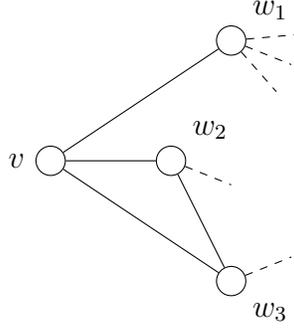


Figure 5: The situation in Section 5.5.

By suitable relabeling, we may therefore assume that G has distinct vertices w_1 (of degree greater than 3) and x, y_2, y_3, v, w_2 and w_3 (all of degree 3), with x and v adjacent to w_1 , and with xy_2y_3 and vw_2w_3 forming triangles (see Figure 6). By considering degrees, we may also assume that the w_i 's and y_i 's are ordered so that $w_i \approx y_i$ for $i = 2, 3$.

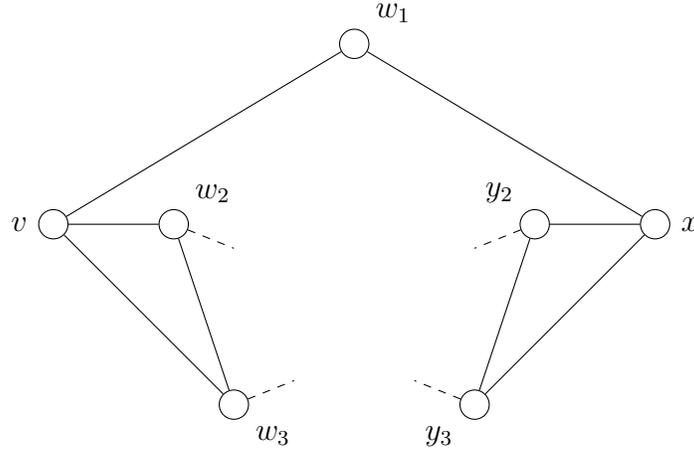


Figure 6: The forced structure in Section 5.5, before modification.

From G we create G' by removing the edges w_2w_3 and y_2y_3 , and adding the edges w_2y_2 and w_3y_3 (see Figure 7). We will argue that $i_3(G) \leq i_3(G')$; but then by the argument of Section 5.2 we have $i_3(G') < i_3(K_{3,n-3})$, and the proof will be complete.

Independent sets of size 3 in G partition into $I_{w_2y_2}$ (those containing both w_2 and y_2 , and so neither of y_3, w_3), $I_{w_3y_3}$ (containing both w_3 and y_3), and I_{rest} , the rest. Independent sets of size 3 in G' partition into $I'_{w_2w_3}$, $I'_{y_2y_3}$, and I'_{rest} . We have

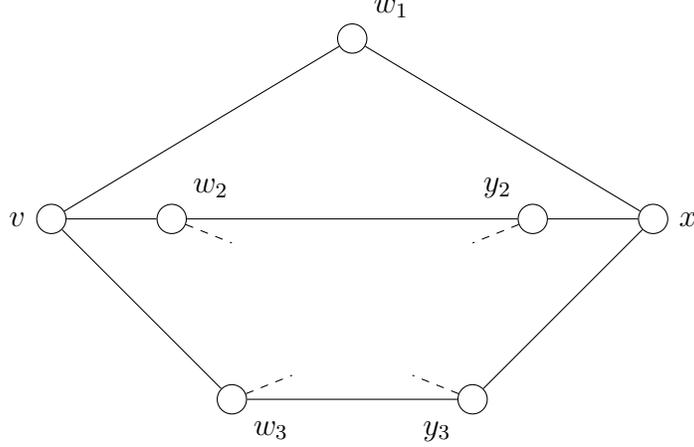


Figure 7: The forced structure in Section 5.5, after modification (i.e. in G').

$|I_{\text{rest}}| = |I'_{\text{rest}}|$ (in fact $I_{\text{rest}} = I'_{\text{rest}}$). We will show $i_3(G) \leq i_3(G')$ by exhibiting an injection from $I_{w_2y_2}$ into $I'_{w_2w_3}$ and one from $I_{w_3y_3}$ into $I'_{y_2y_3}$.

If it happens that for every independent set $\{w_2, y_2, a\}$ in G , the set $\{w_2, w_3, a\}$ is also an independent set in G' , then we have a simple injection from $I_{w_2y_2}$ into $I'_{w_2w_3}$. There is only one way it can happen that $\{w_2, y_2, a'\}$ is an independent set in G but $\{w_2, w_3, a'\}$ is not one in G' ; this is when a' is the neighbor of w_3 that is not v or w_2 . If $\{w_2, y_2, a'\}$ is indeed an independent set in G in this case, then letting b' be the neighbor of y_2 that is not x or y_3 , we find that $\{w_2, w_3, b'\}$ is an independent set in G' , but $\{w_2, y_2, b'\}$ is not one in G . So in this case we get an injection from $I_{w_2y_2}$ into $I'_{w_2w_3}$ by sending $\{w_2, y_2, a\}$ to $\{w_2, w_3, a\}$ for all $a \neq a'$, and sending $\{w_2, y_2, a'\}$ to $\{w_2, w_3, b'\}$. The injection from $I_{w_3y_3}$ into $I'_{y_2y_3}$ is almost identical and we omit the details.

6 Concluding remarks

There now seems to be ample evidence to extend Conjecture 1.2 as follows.

Conjecture 6.1. *For each $\delta \geq 1$, $n \geq 2\delta$, $t \geq 3$ and $G \in \mathcal{G}(n, \delta)$, we have $i_t(G) \leq i_t(K_{\delta, n-\delta})$.*

Throughout we have considered $n \geq 2\delta$, that is, δ small compared to n . It is natural to ask what happens in the complementary range $\delta > n/2$. In the range $n \geq 2\delta$ we (conjecturally) maximize the count of independent sets by extracting as large an independent set as possible (one of size $n - \delta$). In the range $\delta > n/2$ this is still the largest independent set size, but now it is possible to have many disjoint independent sets of this size. The following conjecture seems quite reasonable.

Conjecture 6.2. For $\delta \geq 1$, $n \geq \delta + 1$ and $G \in \mathcal{G}(n, \delta)$, we have $i(G) \leq i(K_{n-\delta, n-\delta, \dots, n-\delta, x})$, where $0 \leq x < n - \delta$ satisfies $n \equiv x \pmod{n - \delta}$.

Question 6.3. For $\delta \geq 1$, $n \geq \delta + 1$ and $t \geq 3$, which $G \in \mathcal{G}(n, \delta)$ maximizes $i_t(G)$?

When $n - \delta$ divides n (that is, $x = 0$), both Conjecture 6.2 and Question 6.3 turn out to be easy; in this case (2) gives that for all $1 \leq t \leq n - \delta$ and all $G \in \mathcal{G}(n, \delta)$ we have $i_t(G) \leq i_t(K_{n-\delta, n-\delta, \dots, n-\delta})$ and so also $i(G) \leq i(K_{n-\delta, n-\delta, \dots, n-\delta})$ (the case $n = 2\delta$ was observed in [1]). The problem seems considerably more delicate when $x \neq 0$.

Lemmas 2.3 and 2.4 allow us in the present paper to focus attention on the class of edge-min-critical and vertex-min-critical graphs. Lemma 3.2 gives us a good understanding of min-critical graphs in the case $\delta = 2$, and the bulk of Section 5 concerns structural properties of min-critical graphs for $\delta = 3$. It is clear that approaching even the case $\delta = 4$ by similar arguments would be considerable work. Any answer to the following question would help significantly.

Question 6.4. For $\delta \geq 4$, what can be said about the structure of edge-min-critical and vertex-min-critical graphs?

Notes added in proof: While this paper was under review, we learned that Cutler and Radcliffe [6] have proved Conjecture 6.2 (although their methods do not seem adaptable to Question 6.3), and we also learned that McDiarmid and Law [12] have made progress on Conjecture 6.1 that overlaps our Theorem 1.4; specifically, they prove Conjecture 6.1 for all $\delta = o(n^{1/3})$, $t \geq 3$ and n sufficiently large.

Acknowledgement: We thank the referees for very careful readings of the paper.

References

- [1] J. Alexander, J. Cutler, and T. Mink, *Independent sets in graphs with given minimum degree*, Electronic Journal of Combinatorics **19** (2012), #P37.
- [2] N. Alon, *Independent sets in regular graphs and sum-free subsets of finite groups*, Israel Journal of Mathematics **73** (1991), 247-256.
- [3] T. Carroll, D. Galvin, and P. Tetali, *Matchings and independent sets of a fixed size in regular graphs*, Journal of Combinatorial Theory Series A **116** (2009), 1219-1227.
- [4] J. Cutler and A. J. Radcliffe, *Extremal problems for independent set enumeration*, Electronic Journal of Combinatorics **18(1)** (2011), #P169.
- [5] ———, *Extremal graphs for homomorphisms*, Journal of Graph Theory **67** (2011), 261-284.
- [6] ———, *The maximum number of complete subgraphs in a graph with given maximum degree*, arXiv:1306.1803.
- [7] D. Galvin, *Two problems on independent sets in graphs*, Discrete Mathematics **311** (2011), 2105-2112.

- [8] G. Hopkins and W. Staton, *Some identities arising from the Fibonacci numbers of certain graphs*, Fibonacci Quarterly **22** (1984), 255-258.
- [9] H. Hua, *A sharp upper bound for the number of stable sets in graphs with given number of cut edges*, Applied Mathematics Letters **22** (2009), 1380-1385.
- [10] J. Kahn, *An entropy approach to the hard-core model on bipartite graphs*, Combinatorics, Probability & Computing **10** (2001), 219-237.
- [11] C. Lin and S. Lin, *Trees and forests with large and small independent indices*, Chinese Journal of Mathematics **23** (1995), 199-210.
- [12] C. McDiarmid and H.-F. Law, *Independent sets in graphs with given minimum degree*, arXiv:1210.1497.
- [13] R. Merrifield and H. Simmons, *Topological Methods in Chemistry*, Wiley, New York, 1989.
- [14] A. Pedersen and P. Vestergaard, *Bounds on the number of vertex independent sets in a graph*, Taiwanese Journal of Mathematics **10** (2006), 1575-1587.
- [15] H. Prodinger and R. Tichy, *Fibonacci numbers of graphs*, The Fibonacci Quarterly **20** (1982), 16-21.
- [16] A.A. Sapozhenko, *On the number of independent sets in bipartite graphs with large minimum degree*, DIMACS Technical Report (2000), no. 2000-25.
- [17] G. Wingard, *Properties and applications of the Fibonacci polynomial of a graph*, Ph.D. thesis, University of Mississippi, May 1995.
- [18] Y. Zhao, *The Number of Independent Sets in a Regular Graph*, Combinatorics, Probability & Computing **19** (2010), 315-320.
- [19] A.A. Zykov, *On some properties of linear complexes*, Mat. Sbornik N.S. **24(66)** (1949), 163-188.