THE PSEUDOFOREST ANALOGUE FOR THE STRONG NINE DRAGON TREE CONJECTURE IS TRUE

LOGAN GROUT AND BENJAMIN MOORE

ABSTRACT. We prove that for any positive integers k and d, if a graph G has maximum average degree at most $2k + \frac{2d}{d+k+1}$, then G decomposes into $k+1$ pseudoforests $C_1, \ldots, C_{k+1}$ so that for at least one of the pseudoforests, each connected component has at most d edges.

1. INTRODUCTION

For any graph G, a *decomposition* of G is a set of edge disjoint subgraphs of G such that the union of their edges sets is the edge set of the graph. Graph decompositions are a particularly heavily studied area of graph theory, and one of the most beautiful results about graph decompositions is the Nash-William characterization of when a graph decomposes into k forests.

Theorem 1.1 ([8], Nash-Williams Theorem). *A graph G decomposes into k forests if and only if*

$$\max_{H \subseteq G, v(H) \geq 2} \frac{e(H)}{v(H) - 1} \leq k.$$

Here, $H \subseteq G$ is taken to mean H is a subgraph of G. We will refer to $\max_{H \subseteq G, v(H) \geq 2} \frac{e(H)}{v(H)-1}$ as the *fractional arboricity* of G. It is not hard to see that the fractional arboricity of a graph might not be integral. For instance, the fractional arboricity of a cycle on n vertices is $\frac{n}{n-1}$. By the Nash-William characterization, this implies that cycles decompose into two forests, and you cannot decompose a cycle into a single forest. Additionally, it is easy to see that cycles decompose into a forest and a matching, which is quite a bit stronger than just saying that cycles decompose into two forests. As the fractional arboricity of cycles is much closer to 1 than to 2, you might speculate that this is the reason you get this extra structure in the decomposition. In general, one might speculate that if the fractional arboricity of a graph is much closer to $k-1$ than to k, then not only does G decompose into k forests, but that one of the forests can be assumed to be a matching. Intuitively, this should be believable. If the fractional arboricity is very close to $k-1$ but still greater than $k-1$, then the Nash-William characterization says that you need k forests, but just barely. Such a graph would roughly look like the union of $k-1$ forests, and then a few edges left over. It is reasonable to believe that you can force these left over edges to form a matching. Even more generally, you could conjecture that there exists an $\varepsilon \in (0,1)$ such that if the fractional arboricity of G is at most $k + \varepsilon$, then G decomposes into $k+1$ forests, such that one of the forests has maximum degree at most d. The Nine Dragon Tree Conjecture (now proven by Hongbi Jiang and Daqing Yang [4]), proposed by Mickeal Montassier, Patrice Ossona de Mendez, André Raspaud, and Xuding Zhu in [6] does just this.

Theorem 1.2 ([4], Nine Dragon Tree Theorem)**.** *Let k and d be positive integers. If the fractional arboricity of G is at most $k + \frac{d}{k+d+1}$, then G decomposes into $k+1$ forests such that one of the forests has maximum degree at most d.*

In [6], it was shown that the fractional arboricity bound is sharp for arbitrarily large graphs. Despite this, Montassier et al. proposed a significant strengthening of the Nine Dragon Tree Theorem, aptly named the Strong Nine Dragon Tree Conjecture.

Conjecture 1.3 ([6], Strong Nine Dragon Tree Conjecture)**.** *Let k and d be positive integers. If the fractional arboricity of G is at most $k + \frac{d}{k+d+1}$, then G decomposes into $k+1$ forests such that one of the forests has each connected component containing at most d edges.*

This conjecture is wide open. The $d = 1$ case was proven by Yang in [9]. The $k = 1$ and $d = 2$ case was proven by Kim, Kostochka, West, Wu, and Zhu [5]. Recently, the second author resolved the $d \le k + 1$ case [7]. Moreover, it was shown that the conjecture is true if you replace "d edges" in the conclusion with a function $f(d, k)$ edges.

This paper will focus on pseudoforest decompositions. Recall, a *pseudoforest* is a graph where each connected component contains at most one cycle. All of the above results or conjectures have pseudoforest analogues. The pseudoforest analogue of Nash-Williams Theorem is Hakimi's Theorem.

Theorem 1.4 ([3], Hakimi's Theorem)**.** *A graph G decomposes into k pseudoforests if and only if the maximum average degree of G is at most $2k$.*

Here, the maximum average degree of a graph G is

$$\mathrm{mad}(G) := \max_{H \subseteq G} \frac{2e(H)}{v(H)}.$$

Genghua Fan, Yan Li, Nine Song, and Daqing Yang [2] proved the pseudoforest analogue of the Nine Dragon Tree Theorem.

Theorem 1.5 ([2])**.** *Let k and d be positive integers. If $mad(G) \le 2k + \frac{2d}{d+k+1}$, then G decomposes into $k+1$ pseudoforests, such that one of the pseudoforests has maximum degree at most d.*

Fan et al. also showed the maximum average degree bound is best possible.

Theorem 1.6 ([2])**.** *For any positive integers k and d, there are arbitrarily large graphs G and an edge e where $mad(G - e) = 2k + \frac{2d}{k+d+1}$, but G does not decompose into $k+1$ pseudoforests where one of the pseudoforests has maximum degree d.*

The main result of this paper is that the pseudoforest analogue of the Strong Nine Dragon Tree Conjecture is true.

Theorem 1.7. *Let k and d be positive integers. If $mad(G) \le 2k + \frac{2d}{d+k+1}$, then G decomposes into $k+1$ pseudoforests $T_1, \ldots, T_k, F$, such that each connected component of F contains at most d edges.*

By Theorem 1.6, the maximum average degree bound given in Theorem 1.7 is sharp. Some of the cases of Theorem 1.7 were known. As having maximum degree one is the same as having one of the pseudoforests be a matching, Theorem 1.5 implies Theorem 1.7 when $d = 1$. Interestingly, the proof given in [5] of the Strong Nine Dragon Tree Conjecture when

$k = 1$ and $d = 2$ implies Theorem 1.7 when $k = 1$ and $d = 2$. Prior to our result, all other cases were open.

As a template for how the proof of Theorem 1.7 will proceed, we will give a proof of the non-trivial direction of Hakimi's Theorem. The proof we give appears in [2] (and is perhaps the first time it appeared in print, we are not aware of any earlier proofs), and the proof of the Nine Dragon Tree Theorem for pseudoforests in some sense follows this proof of Hakimi's Theorem. However, we feel that our proof more faithfully follows this proof of Hakimi's Theorem, and thus leads to the stronger result while having a shorter proof.

Before proceeding, we need some definitions. Given a graph G, an *orientation* of G is obtained from $E(G)$ by taking each edge xy, and replacing xy with exactly one of the arcs (x, y) or (y, x). To reverse the direction of an arc (x, y) is to replace (x, y) with the arc (y, x). For any vertex v, let $d(v)$ denote the degree of v, and $d^+(v)$ denote the outdegree of v. A directed path P from u to v is a path P oriented so that v is the only vertex with no outgoing edge. The next observation is easy and well known.

Observation 1.8. A graph G is a pseudoforest if and only if G admits an orientation where every vertex has outdegree at most one.

From this observation, we get an important corollary.

Corollary 1.9. *A graph admits a decomposition into k pseudoforests if and only if it admits an orientation such that every vertex has outdegree at most k.*

For a proof of Corollary 1.9 we refer the reader to Corollary 1.2 and Theorem 1.1 of [2]. Alternatively, here is a short proof due to a referee. Given an orientation where each vertex has out degree at most k, colour the tails incident to each vertex with distinct colours from $1, \ldots, k$. Now each colour class of edges induces a subgraph with maximum outdegree at most one, hence it is a pseudoforest. For the converse, given k pseudoforests, we orient each vertex to have maximum outdegree at most 1. The union of these oriented pseudoforests is an orientation of G with maximum outdegree at most k. We will use Corollary 1.9 repeatedly and implicitly throughout our proofs. With that, we can give a proof of Hakimi's Theorem.

Proof of Theorem 1.4. We only prove that a graph with maximum average degree $2k$ decomposes into k pseudoforests, as the other direction is trivial.

Suppose towards a contradiction that G has maximum average degree at most $2k$, but G does not decompose into k pseudoforests. Then G does not admit an orientation such that each vertex has outdegree at most k.

Consider an orientation $\vec{G}$ of G that minimizes the sum

$$\rho := \sum_{v \in V(G)} \max\{0, d^+(v) - k\}.$$

If this sum is zero, then we have a desired decomposition, a contradiction. Thus there is a vertex $v \in V(G)$ such that v has outdegree at least $k + 1$. If there is a directed path P from v to x such that x has outdegree at most $k - 1$, then we can reverse the directions on all of the arcs on P and obtain a decomposition with smaller ρ value, a contradiction. Consider the subgraph H induced by vertices which are reachable from directed paths of v. That is, $x \in V(H)$ if there is a directed path P which starts at v and ends at x. Then all vertices in H have outdegree at least k, and v has outdegree at least $k + 1$. But this implies that the average degree of H is strictly larger than $2k$, a contradiction. $\qquad\square$

Now we will give a high level overview of how our proof will proceed. We will take a pseudoforest decomposition $C_1, \ldots, C_t, F$ where we will try and bound the size of each connected component in F. In the above proof of Hakimi's Theorem, the bad situation was a vertex which had too large outdegree. Now, the bad situation is that there is a component which is too large. In the proof of Hakimi's Theorem, we searched for special paths to augment on from a vertex which had too large outdegree, and in our proof, we will search for paths to augment on from a component which is too large. In the proof of Hakimi's theorem, we identified a situation where we could augment our decomposition and obtain a better decomposition, namely, directed paths from a vertex with too large outdegree to a vertex with small outdegree. In our proof, we will identify similar situations, namely, finding two components which are small enough to augment our decomposition, or finding a large component which has at least two small components nearby to perform augmentations. Then we will show that when these configurations are removed, either we have a decomposition satisfying Theorem 1.7 or our graph actually had too large maximum average degree to begin with.

The paper is structured as follows. In Section 2 we give the necessary definitions on how we will pick our counterexample, and prove basic properties about the counterexample. In Section 3, we describe how we will augment our decomposition in certain situations. In Section 4, we show how to use this augmentation strategy to either find an optimal decomposition or show that our graph has too large maximum average degree.

Finally, to fix some notation we will let $v(G) = |V(G)|$ and $e(G) = |E(G)|$. All other undefined graph theory terminology can be found in [1], or any other standard graph theory textbook.

2. Picking the counterexample

In this section we describe how we will pick our counterexample. Fix k and d as positive integers, and suppose that G is a vertex minimal counterexample to Theorem 1.7 for the fixed values of k and d.

Our first step will be to obtain desirable orientations of G. In particular, the orientations we will demand will imply that G decomposes into k pseudoforests each with $v(G)$ edges, and one left over pseudoforest. For this, we use a lemma proved in [2] (Lemma 2.1). Technically, we need a stronger lemma, however the same proof as Lemma 2.1 will suffice. We give a proof for completeness sake only, there is no new idea needed.

Lemma 2.1 ([2]). *If G is a vertex minimal counterexample to Theorem 1.7, then there exists an orientation of G such that for all $v \in V(G)$, we have $k \leq d^+(v) \leq k + 1$.*

Proof. Suppose no such orientation exists. As G has maximum average degree at most $2k+2$, by Hakimi's Theorem, G admits an orientation so that every vertex has outdegree at most $k + 1$. Orient G so that every vertex has outdegree at most $k + 1$, and that the sum

$$\rho := \sum_{v \in V(G)} \max\{0, k - d^+(v)\}$$

is minimized. Observe that if ρ is zero, then we have a desirable orientation.

First we claim there is a vertex v with outdegree $k+1$. If not, then all vertices have outdegree at most k, and by Hakimi's Theorem G decomposes into k pseudoforests, contradicting that G is a counterexample to Theorem 1.7.

Now we claim there is no directed path P from a vertex v with outdegree $k + 1$ to a vertex u with outdegree at most $k - 1$. Suppose towards a contradiction that P is such a path.

Then reversing the orientation on all of the arcs in P gives a new orientation, where v has outdegree k, all internal vertices have the same outdegree, and the outdegree of u increases by one. But this contradicts that we picked our orientation to minimize ρ.

Let S be the set of vertices in G with out degree at most $k - 1$, and let S' be the set of vertices which have a directed path to a vertex in S. Observe that every vertex in S' has outdegree at most k. Let $\bar{S}' = V(G) - S'$. Then every edge with one endpoint lying in S' and one endpoint in $\bar{S}'$ is directed from S' to $\bar{S}'$. Observe that $|\bar{S}'| < v(G)$.

As G is a vertex minimal counterexample we can decompose $G[\bar{S}']$ into $k+1$ pseudoforests such that one of the pseudoforests has each connected component containing at most d edges. Additionally, as every vertex in S' has outdegree at most k, by Hakimi's Theorem we can decompose $G[S']$ into k pseudoforests $C_1, \ldots, C_k$. Thus we only need to deal with the edges between $\bar{S}'$ and S'. Observe that if v has t arcs $(v, u_1), \ldots, (v, u_t)$ where $u_i \in \bar{S}'$ for all $i \in \{1, \ldots, t\}$, then v has outdegree at most $k - t$ in $G[S']$. Thus v has outdegree zero in at least t of the pseudoforests $C_1, \ldots, C_k$. Therefore we can add the arcs $(v, u_1), \ldots, (v, u_t)$ to t of the pseudoforests so that the result is a pseudoforest. As all arcs between S' and $\bar{S}'$ are oriented from S' to $\bar{S}'$, we now have a decomposition of G which satisfies Theorem 1.7. But this contradicts that G is a counterexample to Theorem 1.7. $\qquad\square$

Let $\mathcal{F}$ be the set of orientations of $E(G)$ with $k \leq d^+(v) \leq k+1$ for each vertex $v \in V(G)$. A useful way of keeping track of our pseudoforest decomposition will be to colour the edges blue and red, where the edges coloured red will induce a pseudoforest. This will be the pseudoforest where we will want to bound the size of each connected component.

Definition 2.2. Suppose G is oriented such that $k \leq d^+(v) \leq k+1$ for each vertex $v \in V(G)$. Then a *red-blue colouring of G* is a (non-proper)colouring of the edges where for any vertex $v \in V(G)$, we colour k outgoing arcs of v blue; if after this there is an uncoloured outgoing arc, colour this arc red.

Note that given an orientation in $\mathcal{F}$, one can generate many different red-blue colourings. As a graph decomposes into k pseudoforests if and only if it admits an orientation where each vertex has outdegree at most k, we obtain the following observation.

Observation 2.3. Given a red-blue colouring of G, we can decompose our graph G into $k+1$ pseudoforests such that k of the pseudoforests have all of their edges coloured blue, and the other pseudoforest has all of its edges coloured red.

Observe that one red-blue colouring can give rise to many different pseudoforest decompositions. Given a pseudoforest decomposition obtained from Observation 2.3 we will say a pseudoforest which has all arcs coloured blue is a *blue pseudoforest*, and the pseudoforest with all arcs coloured red is the *red pseudoforest*.

Definition 2.4. Let f be a red-blue colouring of G, and let $C_1, \ldots, C_k, F$ be a pseudoforest decomposition obtained from f by Observation 2.3. Then we say that $C_1, \ldots, C_k, F$ *is a pseudoforest decomposition generated from f*. We will always use the convention that F is the red pseudoforest, and each C_i is a blue pseudoforest.

As G is a counterexample, in every pseudoforest decomposition generated from a red-blue colouring, there is a component of the red pseudoforest which has more than d edges. We define a *residue function* which simply measures how close a decomposition is to satisfying Theorem 1.7.
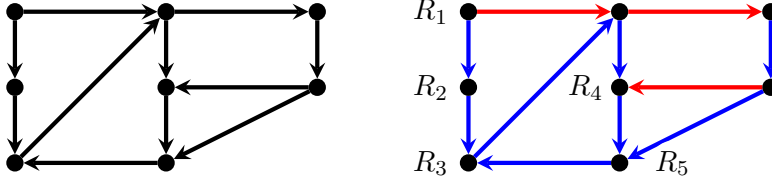
FIGURE 1. In this example we assume $k = 1$ and $d = 1$. On the left, the orientation is in $\mathcal{F}$. On the right we have one possible red-blue colouring generated by this orientation. Here, the entire graph would be the exploration subgraph, and assuming R_1 is the root, $(R_1, R_2, R_3, R_4, R_5)$ is the smallest legal order. Lastly, the isolated vertices are the small components (and are in fact the only possible small components when $k = 1$ and $d = 1$)

Definition 2.5. Let f be a red-blue colouring and $C_1, \ldots, C_k, F$ be a pseudoforest decomposition generated by f. Let $\mathcal{T}$ be the set of components of F. Then the *residue function*, denoted ρ, is

$$\rho(F) = \sum_{K \in \mathcal{T}} \max\{e(K) - d, 0\}.$$

Using a red-blue colouring, and the resulting pseudoforest decomposition, we define an induced subgraph of G which we will focus our attention on. Intuitively, this subgraph should be thought of as an "exploration" subgraph similar to how in the proof of Hakimi's theorem we "explored" from a vertex which had too large outdegree. Here we will "explore" from a component which is too large.

Definition 2.6. Suppose that f is a red-blue colouring of G, and suppose $D = (C_1, \ldots, C_k, F)$ is a pseudoforest decomposition generated from f. Let R be a component of F such that $e(R) > d$. We define the *exploration subgraph* $H_{f,D,R}$ in the following manner. Let $S \subseteq V(G)$ where $v \in S$ if and only if there exists a path $P = v_1, \ldots, v_m$ such that $v_m = v$, $v_1 \in V(R)$, and either $v_i v_{i+1}$ is an arc (v_i, v_{i+1}) coloured blue, or $v_i v_{i+1}$ is an arbitrarily directed arc coloured red. Then we let $H_{f,D,R}$ be the graph induced by S.

Given a particular exploration subgraph $H_{f,D,R}$, we say R is the *root component*. We say the *red components* of $H_{f,D,R}$ are the components of F contained in $H_{f,D,R}$.

It might not be clear why we made this particular definition for $H_{f,D,R}$, however the next observation shows that for any exploration subgraph $H_{f,D,R}$, the red edge density must be low. Before stating the observation, we fix some notation. Given a subgraph K of G, we will let $E_b(K)$ and $E_r(K)$ denote the sets of edges of K coloured blue and red, respectively. We let $e_b(K) = |E_b(K)|$ and $e_r(K) = |E_r(K)|$.

Observation 2.7. For any red-blue colouring f, any pseudoforest decomposition D generated from f, and any choice of root component R, the exploration subgraph $H_{f,D,R}$ satisfies

$$\frac{e_r(H_{f,D,R})}{v(H_{f,D,R})} \leq \frac{d}{d + k + 1}.$$

Proof. Suppose towards a contradiction that

$$\frac{e_r(H_{f,D,R})}{v(H_{f,D,R})} > \frac{d}{d + k + 1}.$$

As $H_{f,D,R}$ is an induced subgraph defined by directed paths, and every vertex $v \in V(G)$ has k outgoing blue edges, each vertex in $H_{f,D,R}$ has k outgoing blue edges. Thus,

$$\frac{e_b(H_{f,D,R})}{v(H_{f,D,R})} = k.$$

Then we have

$$\frac{\mathrm{mad}(G)}{2} \geq \frac{e(H_{f,D,R})}{v(H_{f,D,R})} = \frac{e_r(H_{f,D,R})}{v(H_{f,D,R})} + \frac{e_b(H_{f,D,R})}{v(H_{f,D,R})} > k + \frac{d}{d+k+1}.$$

But this contradicts that G has $\mathrm{mad}(G) \leq 2k + \frac{2d}{d+k+1}$.

$\square$

For the entire proof, we will be attempting to show that we can augment a given decomposition in such a way that either we obtain a decomposition satisfying Theorem 1.7 or we can find a exploration subgraph $H_{f,D,R}$ which contradicts Observation 2.7.

As Observation 2.7 allows us to focus only on red edges, it is natural to focus on red components which have small average degree. With this in mind, we define the notion of a *small red component*.

Definition 2.8. Let $C_1, \ldots, C_t, F$ be a pseudoforest decomposition generated by a red-blue colouring. Let K be a subgraph of F. Then K is a *small red subgraph* if

$$\frac{e_r(K)}{v(K)} < \frac{d}{d+k+1}.$$

If K is connected, we say K *is a small red component*

In particular, we will be interested in the case when K is connected and a small red component. When K is connected the red subgraph is actually isomorphic to a tree, and we can rewrite the density bound in the definition in a more convenient manner.

Observation 2.9. Let K be a connected small red subgraph. Then K is acyclic, and further

$$e_r(K) < \frac{d}{k+1}.$$

Proof. First, suppose that K is not acyclic. Then K contains exactly one cycle. As K is connected, it follows that

$$\frac{e_r(K)}{v(K)} = 1.$$

But $\frac{d}{k+d+1} < 1$ as d and k are positive integers, and hence K would not be a small red subgraph. Thus we can assume that K is acyclic, so $e(K) = v(K) - 1$. Thus

$$\frac{e(K)}{v(K)} = \frac{e(K)}{e(K)+1} < \frac{d}{d+k+1}.$$

Therefore

$$e(K)(d+k+1) < d(e(K)+1).$$

Simplifying, we see that this is equivalent to

$$e(K) < \frac{d}{k+1}.$$

$\square$

We will want to augment our decomposition, and we will want a measure of progress that our decomposition is improving. Of course, if we reduce the residue function that clearly improves the decomposition. However, this might not always be possible, so we will introduce a notion of a "legal order" of the red components. This order keeps track of the number of edges in components which are "close" to the root component, with the idea being that if we can continually perform augmentations to make components "closer" to the root component have fewer edges without creating any large components, then we eventually reduce the number of edges in the root component, which improves the residue function. We formalize this in the following manner.

Definition 2.10. We call an ordering $(R_1, \ldots, R_t)$ of the red components of $H_{f,D,R}$ *legal* if all red components are in the ordering, R_1 is the root component, and for all $j \in \{2, \ldots, t\}$ there exists an integer i with $1 \leq i < j$ such that there is a blue arc (u, v) such that $u \in V(R_i)$ and $v \in V(R_j)$.

Let $(R_1, \ldots, R_t)$ be a legal ordering. We will say that R_i is a *parent* of R_j if $i < j$ and there is a blue arc (v_i, v_j) where $v_i \in R_i$ and $v_j \in R_j$. In this definition a red component may have many parents. To remedy this, if a red component has multiple parents, we arbitrarily pick one such red component and designate it as the only parent. If R_i is the parent of R_j, then we say that R_j is a *child* of R_i. We say a red component R_i is an *ancestor* of R_j if we can find a sequence of red components $R_{i_1}, \ldots, R_{i_m}$ such that $R_{i_1} = R_i$, $R_{j_m} = R_j$, and R_{i_q} is the parent of $R_{i_{q+1}}$ for all $q \in \{1, \ldots, m-1\}$. An important definition is that of vertices witnessing a legal order.

Definition 2.11. Given a legal order $(R_1, \ldots, R_t)$, we say a vertex v *witnesses the legal order for* R_j if there is a blue arc (u, v) such that $u \in R_i$ and $v \in R_j$ and $i < j$.

Observe that there may be many vertices which witness the legal order for a given red component. More importantly, for every component which is not the root, there exists a vertex which witnesses the legal order. We also want to compare two different legal orders.

Definition 2.12. Let $(R_1, \ldots, R_t)$ and $(R'_1, \ldots, R'_{t'})$ be two legal orders. We will say $(R_1, \ldots, R_t)$ is *smaller* than $(R'_1, \ldots, R'_{t'})$ if the sequence $(e(R_1), \ldots, e(R_t))$ is smaller lexicographically than $(e(R'_1), \ldots, e(R'_{t'}))$.

With Definition 2.12, we will pick our minimal counterexample G in the following manner. First, $v(G)$ is minimized.

After this we pick an orientation in $\mathcal{F}$, a red-blue colouring f of this orientation, a pseudoforest decomposition $D = (C_1, \ldots, C_k, F)$ generated from f, such that the number of cycles in F is minimized. Subject to this, we minimize the residue function ρ. Finally, we pick a smallest possible legal order $(R_1, \ldots, R_t)$.

From here on out, we will assume we are working with a counterexample picked in the manner described. The point of minimizing the number of cycles in F is slightly unintuitive compared to minimizing the residue function and minimizing the legal order. However, we minimize the number of cycles because when we augment we will need to ensure our decomposition is in fact a pseudoforest decomposition, and our augmentations will never create more cycles in F. Hence by minimizing the number of cycles in F first, we can easily take care of the cases where cycles occur, which allows us to focus on the more important cases where the components are acyclic.

It was pointed out to the second author via personal communication by Daqing Yang that one can remove the condition to minimize the number of cycles in F by instead modifying the

definition of legal order to take into account the maximum average degree instead of just the number of edges. This is a nice approach and an argument could be made that this approach is more intuitive. However the proof ends up being fundamentally the same, and the authors prefer this approach.

3. Augmenting the decomposition

In this section we describe a very simple operation which will mostly be how we augment our decomposition. Let f be the red-blue colouring of our counterexample, and let $C_1, \ldots, C_k$ be the blue pseudoforests, and F the red pseudoforest. Let $(R_1, \ldots, R_t)$ be the legal ordering picked for our counterexample. As some notation, given a vertex $x \in V(H_{f,D,R})$, we let R^x denote the red component of $H_{f,D,R}$ containing x.

Definition 3.1. Let (x, y) be a blue arc. Further suppose that R^y is acyclic, and suppose that $e = xv$ is an arbitrarily oriented red arc incident to x. To *exchange e and (x, y)* is to perform the following procedure. First, take the maximal directed red path in R^y starting at y, say $Q = y, v_1, \ldots, v_l$ where (v_i, v_{i+1}) is a red arc for $i \in \{1, \ldots, l\}$ and (y, v_1) is a red arc, and reverse the direction of all arcs of Q. Second, change the colour of (x, y) to red and reorient (x, y) to (y, x). Finally, change the colour of e to blue, and if e is oriented (v, x), reorient to (x, v).

See Figure 2 for an illustration. We note that exchanging on an edge e and (x, y) is well-defined. This is because R^y is acyclic (and hence a tree), and thus there is a unique maximal directed path in R^y which starts at y.

Observation 3.2. Suppose we exchange the edge $e = xv$ and (x, y). Then the resulting orientation is in $\mathcal{F}$, and the resulting colouring is a red-blue colouring of this orientation.

Proof. Let us first check the outdegrees of vertices after the exchange. Let $Q = y, v_1, \ldots, v_l$ be the maximal directed red path in R^y before exchanging the edges. First suppose that Q is not just y. Then all of the internal vertices on this path have the same outdegree after reversing as before. On the other hand, the outdegree of y decreases by 1, and the outdegree of v_l increases by one. As Q is a maximal red path, this implies that the outdegree of v_l before reversing the arcs on Q was k, and hence after reversing the arcs this outdegree is $k+1$. The outdegree of y drops by one after reversing the arcs on Q, but we reverse the arc (x, y) to (y, x), and hence the outdegree of y remains the same as before.

If Q is just y, then the outdegree of y before exchanging was k, and then as we reorient (x, y) to (y, x), the outdegree of y is now $k+1$.

Focusing on x now, if e is oriented from x to v, then the outdegree of x is initially $k+1$, and after the exchange it ends up being k; otherwise, we reorient (v, x) to (x, v) and so the outdegree of x remains the same as before the exchange. Lastly, the outdegree of v remains the same if e was oriented (x, v), and otherwise the outdegree of v initially was $k+1$, and after reorienting becomes k. Thus the resulting orientation is in $\mathcal{F}$.

Now we will see that this new colouring is a red-blue colouring of the orientation. Note after exchanging e and (x, y), y has exactly one outgoing edge coloured red. If x had no outgoing red edge before, it still has no outgoing red edge, and if it did have an outgoing red edge, then the outdegree of x dropped by one, and now x has no outgoing red edge. Finally, if Q was not just y, then v_l now has no outgoing red edge. It follows that the resulting colouring is in fact a red-blue colouring. $\qquad\square$
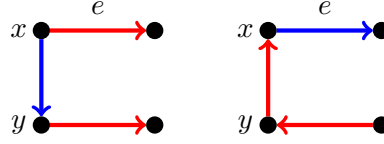
FIGURE 2.   An example of an exchange on an edge e and (x, y)

To avoid repetitively mentioning it, we will implicitly make use of Observation 3.2. Now we begin to impose some structure on our decomposition. First we make an observation which allows us to effectively ignore parent components with cycles (such components will still exist, but for the purposes of our argument we will not need to worry about them).

Observation 3.3. Let (x, y) be a blue arc such that R^x is distinct from R^y and further R^y is a tree. Then x does not lie in a cycle of F.

Proof. Suppose towards a contradiction that x lies in a cycle of F. Let e be an edge incident to x which lies in the cycle coloured red. Now exchange (x, y) and e. As (x, y) was an arc between two distinct red components, and e was in the cycle coloured red, after performing the exchange, we reduce the number of cycles in F by one. However, this contradicts that we picked our counterexample to have the fewest number of red cycles.   □

With this we can show that given two components K and C, where K is the parent of C, and C is acyclic, that $e_r(K) + e_r(C) \geq d$.

Lemma 3.4. *Let R^x and R^y be red components such that R^y is the child of R^x, R^y does not contain a cycle, and (x, y) is a blue arc from x to y. Then $e_r(R^x) + e_r(R^y) \geq d$.*

Proof. Suppose towards a contradiction that $e_r(R^x) + e_r(R^y) < d$. Hence $e_r(R^x) < d$. Thus R^x is not the root. Let w be a vertex which witnesses the legal order for R^x (w exists as R^x is not the root). By Observation 3.3 we know that x does not belong to a cycle of R^x.
    Case 1: $w \neq x$.
    Let e be the edge incident to x in R^x such that e lies on the path from x to w in R^x. Then exchange (x, y) and e. As $e_r(R^x) + e_r(R^y) < d$, all resulting red components have fewer than d edges, and hence we do not increase the residue function. Furthermore, we claim we can find a smaller legal order. Let R_i be the component in the legal order corresponding to R^x. Then consider the new legal order where the components $R_1, \ldots, R_{i-1}$ remain in the same position, we replace R_i with the new red component containing w, and then complete the order arbitrarily. By how we picked e, $e(R^w)$ is strictly smaller than $e(R_i)$, and hence we have found a smaller legal order, a contradiction.
    Case 2: $w = x$.
    We refer the reader to Figure 3 for an illustration. As R^x is not the root component, let R^{x_1} be the closest ancestor of R^x such that $e(R^{x_1}) \geq 1$ (there is an ancestor with this property, as the root has at least one edge). Let $R^{x_n}, R^{x_{n-1}}, \ldots, R^{x_1}$ be a sequence of red components such that for $i \in \{2, \ldots, n\}$, R^{x_i} is the child of $R^{x_{i-1}}$ and R^x is the child of R^{x_n}. Up to relabelling the vertices, there is a path $P = x_1, \ldots, x_n, x, y$ such that (x_i, x_{i+1}) is an arc coloured blue, and (x_n, x) is an arc coloured blue. Let e be a red edge incident to x_n. Now do the following. Colour (x, y) red, and reverse the direction of all arcs in P. Colour e blue, and orient e away from x_1. By the argument in our proof of Observation 3.2, the resulting orientation is in $\mathcal{F}$ and the colouring described is a red-blue colouring. Furthermore as $e_r(R^x) + e_r(R^y) < d$,
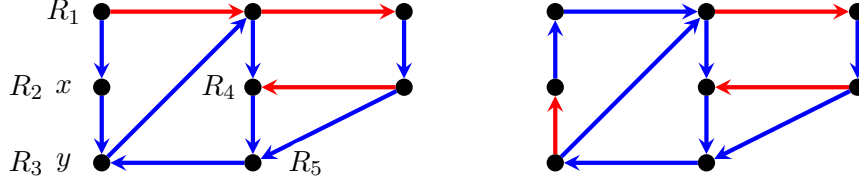
FIGURE 3. An illustration of Case 2 in Lemma 3.4, where $k = 1$ and $d = 1$ and $x = w$. A vertex in the root component was the ancestor, so in this case we reduce the residue function.

all resulting red components have at most d edges, and hence the residue function did not increase (in the event that R^{x_1} is the root, the residue function strictly decreases, so we assume that R^{x_1} is not the root). Finally, we can find a smaller legal order in this orientation, as we simply take the same legal order up to the component containing x_1, and then complete the remaining order arbitrarily. As the component containing x_1 has at least one fewer edge now, this order is a smaller legal order, a contradiction. $\square$

We note the following important special case of Lemma 3.4, that small red components do not have small red children.

Corollary 3.5. *If K is a small red component, then K does not have any small red children.*

Proof. Suppose towards a contradiction that K has a small red child C. As K is small, then $e_r(K) < \frac{d}{k+1}$. Similarly, $e_r(C) < \frac{d}{k+1}$. But then $e_r(K) + e_r(C) < \frac{2d}{k+1} \leq d$, contradicting Lemma 3.4. $\square$

Now we will show that every red component has at most k small red children.

Lemma 3.6. *If K is a red component, then K has at most k small children.*

Proof. Suppose towards a contradiction that K has at least $k + 1$ distinct small children. Then by the pigeon-hole principle, there are two distinct small children C_1 and C_2 such that there are blue arcs (x, x'), (y, y') so that $x \neq y$, $x, y \in V(K)$, $x' \in V(C_1)$ and $y' \in V(C_2)$. By Observation 3.3 we can assume that neither x nor y lies in a red cycle in K. Consider a path $P_{x,y}$ in K from x to y (not directed). Let e_x be the edge incident to x in $P_{x,y}$ and e_y be the edge incident to y in $P_{x,y}$. Let K_x denote the component of $K - e_y$ which contains x, and let K_y denote the component of $K - e_x$ which contains y.

Claim 3.7. $e_r(K_x) \leq e_r(T_2)$ *and* $e_r(K_y) \leq e_r(T_1)$.

Proof. By symmetry, we will only show that $e_r(K_x) \leq e_r(T_2)$. So suppose towards a contradiction that $e_r(K_x) > e_r(T_2)$. Then exchange on (y, y') and e_y. As $e_r(K_x) > e(T_2)$, the residue function does not increase. Observe that if K is the root, then the residue function will decrease, and that will give a contradiction. Thus we may assume that K is not the root. We claim we can find a smaller legal order. If there is a vertex which witnesses the legal order for K in K_x, then taking the same legal order up to K and then replacing K with K_x gives a smaller legal order. Similarly, if there is no vertex which witnesses the legal order in K_x, then because $e_r(K_x) > e_r(T_2)$, taking the same legal order up to K and replacing K with the component containing y of $K - e_y$ after the exchange, and filling in the rest of the order arbitrarily gives a smaller legal order. In both cases, this is a contradiction. $\square$

Note that either $e_r(K) \leq e_r(K_x) + e_r(K_y)$ or $e_r(K) \leq e_r(K_x) + e_r(K_y) - 1$ (the first case occurs if $e_y \neq e_x$, and the second occurs if $e_y = e_x$). Since each T_i is a small child, Claim 3.7 (together with Observation 2.9) implies that

$$e_r(K_x) + e_r(K_y) < \frac{d}{k+1} + \frac{d}{k+1} \leq d.$$

Hence, $e_r(K) \leq d$. Thus we can assume that K is not the root. Let w be a vertex which witnesses the legal order. Without loss of generality, we can assume that $w \in V(K_x)$. Then exchange on (y, y') and e_y. We do not increase the residue function as $e_r(K_y) \leq e_r(T_1) < \frac{d}{k+1}$. However, we can find a smaller legal order by taking the same legal order up to K, and replacing K with K_x, and completing this order arbitrarily. But this contradicts our choice of legal order, a contradiction. $\qquad \square$

We are now in position to prove the theorem.

4. BOUNDING THE MAXIMUM AVERAGE DEGREE

In this section, we give a counting argument to show that our chosen exploration subgraph has too large average degree. We make the following definition for ease of notation.

Definition 4.1. Let K be a red component, and let $K_1, \ldots, K_q$ be the small red children of K. We will let K_C denote the subgraph with vertex set $V(K_C) = V(K) \cup V(K_1) \cup \cdots \cup V(K_q)$, that contains all red edges from $K, K_1, \ldots, K_q$.

Lemma 4.2. *Let K be a red component which is not small. Then K_C is not small. Further, if $e_r(K) > d$, then*

$$\frac{e_r(K_C)}{v(K_C)} > \frac{d}{d+k+1}.$$

Proof. First, observe that if K has no small children then $K_C = K$ and hence is not small. If $e_r(K) > d$, then as K is connected, $v(K) \leq e_r(K) + 1$ and hence $e_r(K)/v(K) \geq (d+1)/(d+2) > d/(d+k+1)$. Thus we can suppose that K has small children $K_1, \ldots, K_q$. By Lemma 3.6, $q \leq k$. By Lemma 3.4, we know for every $i \in \{1, \ldots, q\}$ the inequality $e_r(K) + e_r(K_i) \geq d$ holds. As $e_r(K_i) \geq 0$ for all $i \in \{1, \ldots, q\}$, it follows that $e_r(K_i) \geq \max\{0, d - e(K)\}$ for all $i \in \{1, \ldots, q\}$.

Then a quick calculation shows

$$\begin{aligned}
\frac{e_r(K_C)}{v(K_C)} &= \frac{e_r(K) + \sum_{i=1}^{q} e_r(K_i)}{v(K) + \sum_{i=1}^{q} v(K_i)} \\
&\geq \frac{e_r(K) + \sum_{i=1}^{q} \max\{0, d - e_r(K)\}}{e_r(K) + 1 + q + \sum_{i=1}^{q} \max\{0, d - e_r(K)\}} \\
&\geq \frac{e_r(K) + \sum_{i=1}^{q} \max\{0, d - e_r(K)\}}{e_r(K) + 1 + k + \sum_{i=1}^{q} \max\{0, d - e_r(K)\}}.
\end{aligned}$$

The first equality is simply applying the definition of K_C. The first inequality uses that $e_r(K_i) \geq \max\{0, d - e_r(K)\}$, and that as K_i is small, K_i is a tree, so $v(K_i) = e_r(K_i) + 1$. Finally, the second inequality is using that $q \leq k$. Now we split this into two cases based on whether or not $\max\{0, d - e_r(K)\}$ is 0 or $d - e_r(K)$.

Case 1: $\max\{0, d - e_r(K)\} = 0$.

If $\max\{0, d - e_r(K)\} = 0$, then $e_r(K) \geq d$. Thus it follows that,

$$\frac{e_r(K) + \sum_{i=1}^{q} \max\{0, d - e_r(K)\}}{e_r(K) + 1 + k + \sum_{i=1}^{q} \max\{0, d - e_r(K)\}} = \frac{e_r(K)}{e_r(K) + k + 1}$$

$$\geq \frac{d}{d + k + 1}.$$

Further, if $e_r(K) > d$, the above inequality is strict.

Case 2: $\max\{0, d - e_r(K)\} = d - e_r(K)$. As $e_r(K) \leq d$, we only need to show that K is not small. Calculating we obtain,

$$\frac{e_r(K) + \sum_{i=1}^{q} \max\{0, d - e_r(K)\}}{e_r(K) + 1 + k + \sum_{i=1}^{q} \max\{0, d - e_r(K)\}} = \frac{e_r(K) + q(d - e_r(K))}{e_r(K) + q(d - e_r(K)) + k + 1}$$

$$\geq \frac{d - e_r(K) + e_r(K)}{e_r(K) + d - e_r(K) + k + 1}$$

$$= \frac{d}{d + k + 1}.$$

$\square$

Now we finish the proof. Let $\mathcal{R}$ denote the set of red components of $H_{f,D,R}$ which are not small. By Corollary 3.5 it follows that,

$$V(H_{f,D,R}) = \bigcup_{K \in \mathcal{R}} V(K_C).$$

This follows since a small component cannot have a small child. Therefore it follows that:

$$E_r(H_{f,D,R}) = \bigcup_{K \in \mathcal{R}} E(K_C).$$

Now we bound the maximum average degree of $H_{f,D,R}$. By Lemma 4.2, we have

$$\frac{e_r(H_{f,D,R})}{v(H_{f,D,R})} = \frac{\sum_{K \in \mathcal{R}} e_r(K_C)}{\sum_{K \in \mathcal{R}} v(K_C)}$$

$$> \frac{d}{d + k + 1}.$$

Here, equality holds in the first line because we chose unique parents for components. The strict inequality follows as K_C is not small for any $K \in \mathcal{R}$ by Lemma 4.2, and further the root component satisfies $e(R) > d$. However, this contradicts Observation 2.7. Theorem 1.7 follows.

REFERENCES

[1] J.A Bondy and U.S.R Murty. *Graph Theory*. Springer, 2008.
[2] Genghua Fan, Yan Li, Ning Song, and Daqing Yang. Decomposing a graph into pseudoforests with one having bounded degree. *Journal of Combinatorial Theory, Series B*, 115:72 − 95, 2015.
[3] S.L. Hakimi. On the degrees of the vertices of a directed graph. *Journal of the Franklin Institute*, 279(4):290 − 308, 1965.

[4] Hongbi Jiang and Daqing Yang. Decomposing a graph into forests: The nine dragon tree conjecture is true. *Combinatorica*, 37(6):1125–1137, Dec 2017.

[5] Seog-Jin Kim, Alexandr V. Kostochka, Douglas B. West, Hehui Wu, and Xuding Zhu. Decomposition of sparse graphs into forests and a graph with bounded degree. *Journal of Graph Theory*, 74(4):369–391, 2013.

[6] Mickael Montassier, Patrice Ossona de Mendez, André Raspaud, and Xuding Zhu. Decomposing a graph into forests. *Journal of Combinatorial Theory, Series B*, 102(1):38 – 52, 2012.

[7] Benjamin Moore. An approximate version of the strong nine dragon tree conjecture, 2019. arXiv:1909.07946.

[8] C. St.J. A. Nash-Williams. Decomposition of finite graphs into forests. *Journal of the London Mathematical Society*, s1-39(1):12–12, 1964.

[9] Daqing Yang. Decomposing a graph into forests and a matching. *Journal of Combinatorial Theory, Series B*, 131:40 – 54, 2018.

(Logan Grout) DEPARTEMENT OF COMBINATORICS AND OPTIMIZATION, UNIVERSITY OF WATERLOO, WATERLOO, ON, CANADA

  E-mail address: `lcgrout@uwaterloo.ca`

(Benjamin Moore) DEPARTMENT OF COMBINATORICS AND OPTIMIZATION, UNIVERSITY OF WATERLOO, WATERLOO, ON, CANADA

  E-mail address: `brmoore@uwaterloo.ca`