

Acyclic subgraphs of tournaments with high chromatic number

Jacob Fox, Matthew Kwan and Benny Sudakov

ABSTRACT

We prove that every n -vertex tournament G has an acyclic subgraph with chromatic number at least $n^{5/9-o(1)}$, while there exists an n -vertex tournament G whose every acyclic subgraph has chromatic number at most $n^{3/4+o(1)}$. This establishes in a strong form a conjecture of Nassar and Yuster and improves on another result of theirs. Our proof combines probabilistic and spectral techniques together with some additional ideas. In particular, we prove a lemma showing that every tournament with many transitive subtournaments has a large subtournament that is almost transitive. This may be of independent interest.

1. Introduction

An *orientation* of a graph G is an assignment of a direction to each edge. There is a long history of surprising connections between colourings and orientations of graphs (see [8] for a survey). Perhaps the most famous of these is the Gallai–Hasse–Roy–Vitaver theorem [5, 7 13, 14], which states that the chromatic number of any graph G can be equivalently defined to be the minimum, over all orientations of G , of the length of the longest directed path in that orientation, plus one. This implies that for any oriented graph G with chromatic number n , there is a directed path with n vertices.

It is very natural to ask whether there are oriented graphs H other than directed paths that must necessarily appear in any oriented graph G with sufficiently large chromatic number. Since Erdős famously proved that there are graphs with arbitrarily large girth and chromatic number, we can only hope to prove results of this type when H is an oriented forest. As a far-reaching extension of the Gallai–Hasse–Roy–Vitaver theorem, Burr [3] conjectured in 1980 that any $(2k-2)$ -chromatic oriented graph contains a copy of every oriented tree on k vertices. The first general result in this direction is also due to Burr [3], who proved that any oriented graph G with chromatic number $n = (k-1)^2$ contains a copy of every oriented k -vertex tree.

Burr’s conjecture has remained widely open in the 40 years since it was proposed, even for relatively simple trees such as paths with arbitrary orientation. A special case that has attracted a lot of attention is the case where G is a *tournament*: an orientation of the complete n -vertex graph. This special case of Burr’s conjecture is known as *Sumner’s conjecture*, and following a sequence of partial results it was resolved for large n in a tour de force by Kühn, Mycroft and Osthus [10]. We remark that both Burr’s and Sumner’s conjectures, if true, are best possible: there is a tournament on $(2k-3)$ vertices which does not contain a certain oriented tree on k vertices.

Received 15 December 2019; revised 11 November 2020; published online 5 January 2021.

2020 *Mathematics Subject Classification* 05C15, 05C20, 05D40 (primary).

The first author is supported by a Packard Fellowship and by NSF Award DMS-1855635. The second author is supported in part by SNSF project 178493 and NSF Award DMS-1953990. The third author is supported in part by SNSF grant 200021-196965.

© 2021 The Authors. The publishing rights in this article are licensed to the London Mathematical Society under an exclusive licence.

Recently, Addario-Berry, Havet, Sales, Reed and Thomassé [1] managed to improve Burr's original bound, showing that every oriented graph with chromatic number $n = k^2/2 - k/2 + 1$ contains every oriented tree on k vertices. In the same paper, they also made the interesting observation that if a k -chromatic oriented graph has no directed cycle, then it contains every oriented tree of order k , and suggested that one approach to an improved bound towards Burr's conjecture could be to prove that every k -chromatic oriented digraph has an acyclic subgraph with large chromatic number.

PROBLEM 1.1. What is the minimum integer $f(k)$ such that every $f(k)$ -chromatic oriented graph has an acyclic k -chromatic subgraph?

It is easy to see that the edges of any oriented graph can be partitioned into two acyclic subgraphs (fix an ordering of the vertex set, and consider the graph G_1 of 'forwards' edges and the graph G_2 of 'backwards' edges). Using the well-known inequality $\chi(G_1 \cup G_2) \leq \chi(G_1)\chi(G_2)$, it follows that $f(k) \leq k^2$. By taking a bit more care with this argument, the aforementioned authors were able to prove that $f(k) \leq k^2 - 2k + 2$, but it seems that new ideas will be required to substantially improve this bound.

In the spirit of Sumner's conjecture, it is natural to consider the restriction of Problem 1.1 to the setting of tournaments. This was done by Nassar and Yuster [11], who asked the following.

PROBLEM 1.2. What is the minimum integer $g(k)$ such that every $g(k)$ -vertex tournament has an acyclic k -chromatic subgraph?

Nassar and Yuster were able to prove that $k^{8/7}/4 \leq g(k) \leq k^2 - (2 - 1/\sqrt{2})k + 2$, and further conjectured that $g(k) = o(k^2)$. Our first result proves their conjecture in a strong form.

THEOREM 1.3. *Every n -vertex tournament G has an acyclic subgraph with chromatic number at least $n^{5/9-o(1)}$. That is, $g(k) \leq k^{9/5+o(1)}$.*

The rough idea for the proof of Theorem 1.3 is to consider a *random* acyclic subgraph of G , and observe that this is likely to have high chromatic number unless G contains many transitive subtournaments. We then prove an approximate structural result (Lemma 2.2) showing that the existence of many transitive subtournaments implies that G has a large almost-transitive subtournament, from which it is possible to deduce in a different way that G has an acyclic subgraph with high chromatic number. The details of the proof are presented in Section 2.

In addition, we are also able to improve Nassar and Yuster's lower bound on $g(k)$.

THEOREM 1.4. *There is an n -vertex tournament G such that every acyclic subgraph has chromatic number at most $O(n^{3/4} \log n)$. That is, $g(k) = \Omega((k/\log k)^{4/3})$.*

Since $f(k) \geq g(k)$, this theorem gives a corresponding lower bound on $f(k)$, so an immediate consequence is that one cannot hope to prove bounds stronger than $k^{4/3-o(1)}$ for Burr's conjecture via Problem 1.1. We prove Theorem 1.4 in Section 3 using a construction due to Rödl and Winkler [12], in which a random orientation is defined in terms of a projective plane. This involves some non-trivial analysis of increasing subsequences in random permutations (see Section 3.2).

For the sake of clarity of presentation, we omit floor and ceiling signs where they are not crucial. All logarithms are base e unless otherwise stated.

2. Finding an acyclic subgraph with high chromatic number

In this section, we prove Theorem 1.3. Note that the maximal acyclic subgraphs of an oriented graph G are all obtained by taking some ordering π on the vertices and considering the graph G_π obtained by including the edges that are oriented ‘forwards’ according to π . Given an n -vertex tournament G , we need to show that there is an ordering π for which G_π has high chromatic number. First, this easily follows from consideration of a random ordering if G has few small transitive subtournaments. Let TT_k be the k -vertex transitive tournament.

LEMMA 2.1. *If an n -vertex tournament G has fewer than $k!$ copies of TT_k , then there is an ordering π with $\alpha(G_\pi) \leq k$, so $\chi(G_\pi) \geq n/k$.*

Proof of Lemma 2.1. Consider uniformly random π . For a particular set S of k vertices, the probability that S is an independent set in G_π is either $1/k!$ if $G[S]$ is transitive, or zero if it is not. So, by the union bound we have $\alpha(G_\pi) < k$ with positive probability. \square

Given Lemma 2.1, it suffices to consider the case where G has many copies of TT_k . The most important ingredient is the following lemma, showing that in this case G has a large subtournament that is almost transitive in a certain sense. Say that an n -vertex tournament is q -almost-transitive if there is an ordering such that for every vertex v , at most q edges incident to v are oriented ‘backwards’ according to π (so a 0-almost transitive tournament is transitive).

LEMMA 2.2. *Let $k = n^{\Omega(1)}$. If an n -vertex tournament G has at least $k!$ copies of TT_k then it has a set of $n' \geq k^{2-o(1)}$ vertices which induces an $(n'/k^{1-o(1)})$ -almost-transitive subtournament G' .*

The other ingredient we need is the following lemma, showing that almost-transitive tournaments have acyclic subgraphs with high chromatic number.

LEMMA 2.3. *Suppose $n^{1/3} \leq q \leq n^{1-\Omega(1)}$. An n -vertex q -almost-transitive tournament G has an ordering π with $\alpha(G_\pi) \leq n^{1/4+o(1)}q^{1/4}$, so $\chi(G_\pi) \geq n^{3/4-o(1)}q^{-1/4}$.*

Before proving these lemmas, we give the short deduction of Theorem 1.3.

Proof of Theorem 1.3 given Lemmas 2.1 to 2.3. Let $k = n^{4/9}$. If there are fewer than $k!$ copies of TT_k , then Lemma 2.1 shows that there is some π such that $\chi(G_\pi) \geq n/k = n^{5/9}$ and we are done. Otherwise, by Lemma 2.2, G has an $(n'/k^{1-o(1)})$ -almost-transitive subtournament G' on $n' \geq k^{2-o(1)}$ vertices, which by Lemma 2.3 has an acyclic subgraph G'_ρ such that

$$\chi(G'_\rho) \geq \left(k^{2-o(1)}\right)^{3/4-o(1)} \left(k^{1+o(1)}\right)^{-1/4} = n^{5/9-o(1)},$$

as desired. \square

2.1. Iterative alignment and refinement

In this subsection, we prove Lemma 2.2. Most of the subsection will be spent proving the following lemma, which should be viewed as a more technical variant of Lemma 2.2.

LEMMA 2.4. *Let n be sufficiently large and let G be an n -vertex tournament which has at least M copies of TT_k . If $q = 2n(\log n)^2/k$, then there is $k - 3k/\log n \leq k' \leq k$ such that G has an induced q -almost-transitive subtournament G' on at least $e^{-5}kM^{1/k}$ vertices containing at least $e^{-3k}M$ copies of $\text{TT}_{k'}$.*

The key aspect in which Lemma 2.4 is weaker than Lemma 2.2 is that the almost-transitivity parameter $q = 2n(\log n)^2/k$ depends on the number of vertices n of G , whereas Lemma 2.2 demands an almost-transitivity parameter with a similar dependence on the number of vertices n' of G' . In order to overcome this issue, we simply iterate Lemma 2.4 as follows.

Proof of Lemma 2.2, given Lemma 2.4. Let $t = \sqrt{\log n}$. We obtain a sequence of subtournaments $G = G_0 \supseteq G_1 \supseteq \dots \supseteq G_t$ by iteratively applying Lemma 2.4 (where in each step, we use the guarantees from the previous step on the number of vertices and the number of small transitive subtournaments). For each $0 \leq i \leq t$, let n_i be the number of vertices in G_i . We let k_i be the value for which we obtain a bound on the number of copies of TT_{k_i} in G_i . Thus, $k_0 = k$ and for $1 \leq i \leq t$ we have $k_i \geq k_{i-1}(1 - 3/\log n_{i-1})$. Let $q_i = 2n_{i-1}(\log n_{i-1})^2/k_{i-1}$. Then, each G_i is q_i -almost-transitive, has at least $M_i := e^{-3(k_0 + \dots + k_{i-1})}k! \geq e^{-3ik}k!$ copies of TT_{k_i} , and for $i > 0$, has $n_i \geq e^{-5}k_{i-1}(M_{i-1})^{1/k_{i-1}}$ vertices.

Combining our lower bounds on n_i and M_{i-1} , and using Stirling's approximation, we obtain $n_i \geq \Omega(e^{-3i}k_{i-1}k)$. Recalling that $k_i \geq k_{i-1}(1 - 3/\log n_{i-1})$ and that $k = n^{\Omega(1)}$, it is straightforward to inductively prove that $n_i \geq k^{2-o(1)}$ and $k_i \geq ke^{-O(i/\log n)} = (1 - o(1))k$ for every $i \leq t = \sqrt{\log n}$. It follows that each $q_i = n_{i-1}k^{o(1)-1}$.

There is some i such that $\log n_i - \log n_{i-1} = \log(n_i/n_{i-1}) \leq (1/t)\log n \leq \sqrt{\log n}$, meaning that $n_i/n_{i-1} = k^{o(1)}$, so $q_i = n_i/k^{1-o(1)}$. Then G_i satisfies the required properties. \square

To prove Lemma 2.2, it therefore suffices to prove Lemma 2.4. We need some very simple auxiliary lemmas.

LEMMA 2.5. *Let H be a hypergraph with m non-empty edges and n vertices. Then it has an induced subhypergraph with minimum degree at least m/n .*

Proof. Iteratively delete a vertex of degree less than m/n together with the edges touching it, until no such vertex remains. In total we deleted fewer than $(m/n) \cdot n = m$ edges, so we are left with a non-empty induced subhypergraph that satisfies the assertion of the lemma. \square

LEMMA 2.6. *Every n -vertex tournament G has a vertex whose indegree and outdegree are both at least $(n - 2)/4$.*

Proof. Let A (respectively, B) be the set of all vertices whose indegree (respectively, outdegree) is at least $(n - 2)/4$. For every vertex, the sum of its indegree and outdegree is exactly $n - 1 \geq 2(n - 2)/4$, so every vertex is in A or B . Without loss of generality, suppose $|A| \geq |B|$. Then observe that the induced subtournament $G[A]$ has average outdegree $(|A| - 1)/2 \geq (n/2 - 1)/2 = (n - 2)/4$, so it has a vertex whose outdegree is at least this average. This vertex has the desired property. \square

Now we prove Lemma 2.4.

Proof of Lemma 2.4. We iteratively build sequences of vertex sets

$$\emptyset = W_0 \subset W_1 \subset \dots \subset W_t \subset V(G), \quad V(G) = V'_0 \supseteq V_0 \supset V'_1 \supseteq V_1 \supset \dots \supset V'_t \supseteq V_t,$$

for some $t < 3k/\log n$, in such a way that each $|W_i| = i$ and each $V'_i \cap W_i, V_i \cap W_i = \emptyset$. The idea is that the W_i will be sets of vertices each inducing a copy of TT_i , and V'_i, V_i will be sets of vertices which can be used to extend W_i into a copy of TT_k in many different ways. We will add vertices one-by-one to form the W_i , in such a way that $G[V'_i], G[V_i]$ gradually get 'closer to being transitive', and then we will take $G' = G[V_t]$. In order to explain how to iteratively build our desired sets, we need to make some definitions.

- Let H'_i be the $(k - i)$ -uniform hypergraph with vertex set V'_i whose hyperedges are those $(k - i)$ -vertex subsets S for which $G[S \cup W_i]$ forms a copy of TT_k .
- Let N'_i be the number of edges in H'_i . Equivalently, N'_i is the number of copies of TT_k in $G[V'_i \cup W_i]$ which include all vertices of W_i . In particular, $N'_0 \geq M$.
- Define V_i from V'_i by letting $H_i = H'_i[V_i]$ be a non-empty-induced subgraph of H'_i with minimum degree at least $N'_i/|V'_i|$, which exists by Lemma 2.5. Let $N_i \geq N'_i/|V'_i| \geq N'_i/n$ be the number of edges in H_i .
- For each i and each $v \in V_i$, note that $W_i \cup \{v\}$ induces a transitive subtournament on $i + 1$ vertices (because v is contained in at least one edge of H_i). In this transitive tournament, there are $i + 1$ possibilities for the position of v relative to the vertices in W_i . Partition V_i into subsets $V_{i,0} \cup \dots \cup V_{i,i}$ according to these $i + 1$ possibilities: let $V_{i,j}$ be the set of vertices $v \in V_i$ which have indegree exactly j in $G[W_i \cup \{v\}]$.
- For vertices $x \in V_{i,j}$ and $y \in V_{i,q}$ with $j < q$ we say that x and y are *incompatible* if the edge between x and y is oriented towards x (so then there can be no transitive tournament containing $W_i \cup \{x, y\}$).

Now, we will choose $w \in V_i$ (which we will then add to W_i to form W_{i+1}) according to one of the following two procedures.

- *Refinement.* Consider a largest $V_{i,j}$. By Lemma 2.6, $G[V_{i,j}]$ has a vertex with indegree and outdegree at least $(|V_{i,j}| - 2)/4 \geq |V_{i,j}|/4 - 1$, which we take as w .
- *Alignment.* Let $\varepsilon = (\log n)^2/k$. If there is a vertex which is incompatible with $\varepsilon|V_i|$ other vertices of V_i , then take such a vertex as w .

Note that alignment may not always be possible, but refinement is always possible. After choosing $w \in V'_i$ according to one of the above two procedures, we then set $W_{i+1} = W_i \cup \{w\}$ and let V'_{i+1} be obtained from V_i by deleting w and all vertices incompatible with it.

The purpose of the alignment steps is to ensure that most edges between the $V_{i,j}$ are oriented in a single direction, and the purpose of the refinement steps is to ensure that the $V_{i,j}$ are not too large (so there are not very many edges inside the $V_{i,j}$, whose orientations we do not control). We make some simple observations about both procedures.

CLAIM. *No matter how we decide whether to do alignments or refinements in our $t < 3k/\log n$ steps, the following always hold.*

- (a) *If step i is an alignment step, then $|V'_{i+1}| \leq (1 - \varepsilon)|V'_i|$.*
- (b) $N_t \geq e^{-3k}M$.
- (c) $|V_t| \geq e^{-5}kM^{1/k}$. □

Proof. First, (a) follows from the definition of an alignment step. Next, observe that each H'_{i+1} is a *link hypergraph* of H_i , obtained by taking all edges incident to a vertex $w \in H_i$ and deleting w from each of them. Since every vertex in H_i has degree at least N'_i/n , it follows that $N'_{i+1} \geq N'_i/n$. Iterating this $t < 3k/\log n$ times, then using the fact that $N_t \geq N'_t/n$, yields (b). Finally, since H_t is a $(k - t)$ -uniform hypergraph, note that

$$N_t \leq \binom{|V_t|}{k - t} \leq \left(\frac{e|V_t|}{k - t}\right)^{k-t},$$

so

$$|V_t| \geq \left(\frac{k - t}{e}\right)(N_t)^{1/(k-t)}.$$

Combining this with (b) yields (c), for large n . □

We will end up taking $G[V_t]$ as our subtournament G' , so properties (b) and (c) in the above claim ensure that it has enough vertices and enough copies of $\text{TT}_{k'}$, for some $k' \geq k - 3k/\log n$. We just need to decide how to choose whether to do alignments or refinements, in such a way that $G[V_t]$ ends up being q -almost-transitive. These choices are very simple: we first do $s := k/\log n$ refinement steps, then do alignment steps as long as is possible. By (a), each alignment step decreases V_i' by a factor of $(1 - \varepsilon)$, so it is not possible to perform more than $\log_{1/(1-\varepsilon)} n < (2/\varepsilon) \log n = 2k/\log n$ alignment steps (using the inequality $-\log(1 - \varepsilon) > \varepsilon/2$ for small $\varepsilon > 0$). Thus, the total number of steps will be $t < 3k/\log n$, as we have been assuming.

CLAIM. For $i \geq s = k/\log n$, we have

$$\max_j |V_{i,j}| \leq \frac{4n}{s+1}. \quad \square$$

Proof. For $i \leq s$, the way the parts $V_{i,j}$ evolve is as follows. At each step, we ‘split’ a largest part $V_{i,j}$ into two sets, each of size at most $3|V_{i,j}|/4$ (removing a single element w), and then we delete some further elements in each part. For steps $i > s$, the parts $V_{i,j}$ only get smaller, so it suffices to consider $\max_j |V_{s,j}|$.

Consider the following alternative procedure to iteratively build a partition of $\{1, \dots, n\}$. Start with the trivial partition \mathcal{P}_0 into one part, and for each step i consider a largest part $A \in \mathcal{P}_i$, split it into two parts of sizes $\lceil |A|/4 \rceil$ and $\lfloor 3|A|/4 \rfloor$ and call the resulting partition \mathcal{P}_{i+1} . This procedure ‘dominates’ the procedure used to build the $V_{i,j}$ in the sense that for each $i \leq s$, we have

$$\max_j |V_{i,j}| \leq \max_{A \in \mathcal{P}_i} |A|.$$

But note that in this second procedure we always maintain the property

$$\max_{A \in \mathcal{P}_i} |A| \leq 4 \min_{A \in \mathcal{P}_i} |A|,$$

so $\max_{A \in \mathcal{P}_s} |A|$ is at most four times the average size of the $A \in \mathcal{P}_s$, which is $n/(s+1)$. \square

Now, after all our alignment steps have finished, in V_t there is no vertex which is incompatible with more than $\varepsilon|V_t|$ other vertices of V_t . Consider an arbitrary ordering of V_t such that for $j < q$, all the elements of $V_{t,j}$ appear before all the elements of $V_{t,q}$. In $G[V_t]$, for any vertex v there are a total of at most $\varepsilon|V_t| + 4n/(s+1) \leq 2n(\log n)^2/k = q$ edges adjacent to v which are oriented ‘against’ our chosen ordering (for large n). So, $G[V_t]$ is q -almost-transitive, as desired. \square

2.2. Biased random orderings

In this subsection, we prove Lemma 2.3. If G is q -almost-transitive, then we can choose ρ to be the ordering defining almost-transitivity, yielding a very dense graph G_ρ whose complement has maximum degree at most q . One can give a lower bound on $\chi(G_\rho)$ simply using the number of edges of G_ρ , but to prove Lemma 2.3 we need to do better than this. We will consider a random perturbation π of ρ , which has the same ‘large-scale’ structure as ρ , but on a ‘small scale’ it resembles a uniformly random permutation.

We first need a simple auxiliary lemma about uniform random permutations.

LEMMA 2.7. Let G be an n -vertex tournament and let π be a random permutation of its vertex set. Then with probability $1 - n^{-\omega(1)}$, we have $\alpha(G_\pi) = O(\sqrt{n})$.

Proof. For any k , the expected number of size- k independent sets in G_π is $N_G/k! \leq \binom{n}{k}/k!$, where N_G is the number of copies of TT_k in G . If k is a large multiple of \sqrt{n} , this number is $n^{-\omega(1)}$. □

Now we prove Lemma 2.3.

Proof of Lemma 2.3. Let $s = \sqrt{n/q}$. Consider the ordering ρ defining almost-transitivity (so for each vertex v , only q edges incident to v are oriented ‘against’ ρ , to some set of vertices which we call $B(v)$). Starting from the ordering ρ , we divide the vertex set into s contiguous blocks A_1, \dots, A_s of size n/s , and then randomly permute the indices within each block, to obtain a random ordering π (retaining the large-scale structure of ρ , but introducing some ‘local’ randomness).

By Lemma 2.7 and the union bound, with positive probability π has the property that for each $1 \leq i \leq s$ and each vertex v , we have

$$\alpha(G_\pi[A_i]) = O\left(\sqrt{n/s}\right), \tag{2.1}$$

$$\alpha(G[A_i \cap B(v)]) \leq n^{o(1)} + O\left(\sqrt{|A_i \cap B(v)|}\right). \tag{2.2}$$

Consider such an outcome of π .

Now consider an independent set S in G_π . Let v be the first vertex in S (according to the ordering π), and let A_i be the block that contains v . By the definition of π , we must have $S \subseteq A_i \cup B(v)$. Next, equation (2.1) implies $\alpha(G_\pi[S \cap A_i]) = O(\sqrt{n/s})$. Hence by equation (2.2) and convexity of the function $f(x) = \sqrt{x}$, we have

$$\alpha(G_\pi[S \setminus A_i]) \leq \sum_j \left(n^{o(1)} + O\left(\sqrt{|A_j \cap B(v)|}\right) \right) \leq sn^{o(1)} + O(\sqrt{qs}),$$

where the sum is over all $j \neq i$ with $A_j \cap B(v) \neq \emptyset$. It follows that $\alpha(G_\pi[S]) \leq \alpha(G_\pi[A_i]) + \alpha(G_\pi[S \setminus A_i]) \leq O(\sqrt{n/s} + \sqrt{qs}) + sn^{o(1)} \leq n^{1/4+o(1)}q^{1/4}$, recalling that $q \geq n^{1/3}$ (therefore $s \leq n^{1/4}q^{1/4}$). □

3. A tournament with easily colourable acyclic subgraphs

In this section, we prove Theorem 1.4. We will study the following construction, which was previously introduced by Rödl and Winkler [12] (see also [2]).

DEFINITION 3.1. Let G be a tournament obtained as follows. First, consider a projective plane P of order $t - 1$, with $t^2 - t + 1$ lines and $t^2 - t + 1$ points, t points on every line and t lines going through every point. Then, G will be a tournament on $(t^2 - t + 1)k$ vertices, where for each point of the projective plane there are k associated vertices of G (so, our vertices are divided into $t^2 - t + 1$ ‘buckets’). For the edges inside the buckets, choose their orientation arbitrarily. For the edges between the buckets, we do the following. For each line of P , consider the kt corresponding vertices of G . Choose a random ordering of these vertices, and orient all remaining edges between these vertices according to the ordering. Since every pair of points in P are contained in exactly one line, we can do this independently for each line.

Now the following lemma will imply Theorem 1.4 almost immediately. Recall from the previous section that for an oriented graph G , we define G_π to be the subgraph obtained by including the edges that are oriented ‘forwards’ according to π .

LEMMA 3.2. *There is an absolute constant C such that the following holds. Let G be as in Definition 3.1, with $k = t^2$ and t sufficiently large. With positive probability, G has the following property. For any ordering π of the vertices of G and any subset X of at least $100t^3 \log t$ vertices of G , we have $\alpha(G_\pi[X]) \geq |X|/(16t^3)$.*

Before turning to the proof of Lemma 3.2, we give the short deduction of Theorem 1.4.

Proof of Theorem 1.4. For any n , by Bertrand’s postulate we can find a prime $t - 1 = \Theta(n^{1/4})$, with $(t^2 - t + 1)t^2 \geq n$. So, there is a projective plane of order $t - 1$ and we can consider G as in Definition 3.1, with $k = t^2$. Consider a particular outcome of G satisfying the condition in Lemma 3.2. Then, consider any acyclic subgraph G' of G , which is a subgraph of some G_π (simply take π to be an ordering corresponding to a transitive closure of G').

Now we construct a proper colouring of G_π as follows. First greedily take maximum independent sets as colour classes, until there are fewer than $100t^3 \log t$ vertices remaining. Then give each remaining vertex its own unique colour. Let ℓ be the number of colours coming from the greedy phase of this procedure, so $k(t^2 - t + 1)(1 - 1/(16t^3))^{\ell-1} \geq 100t^3 \log t$ by the property in Lemma 3.2, and therefore $\ell = O(t^3 \log t)$. We deduce that $\chi(G_\pi) = \ell + 100t^3 \log t = O(t^3 \log t)$.

Finally, we can simply delete a few vertices from G to obtain a tournament on exactly n vertices with the desired properties, using the fact that the chromatic number of a subgraph is at most the chromatic number of the graph. □

The first ingredient in the proof of Lemma 3.2 is the fact that, in G , all large sets of vertices are well distributed between lines of the underlying projective plane.

LEMMA 3.3. *Let G be as in Definition 3.1, and let X be a set of at least $9kt$ vertices of G . Then at most $t^2/2$ lines of the projective plane underlying G contain fewer than $|X|/(2t)$ vertices of X .*

We also need the following lemma giving strong bounds on the probability that a random permutation fails to contain a long increasing subsequence, where each index in the subsequence comes from a different ‘bucket’.

LEMMA 3.4. *Divide the interval $\{1, \dots, m\}$ into ‘buckets’ of size at most k , where $m \leq k^2$. Consider a random permutation σ of $\{1, \dots, m\}$. Let $\ell = m/(8k)$. With probability at least $1 - e^{-m/24}$, there is a sequence $1 \leq i_1 < \dots < i_\ell \leq m$ with each i_j coming from a different bucket, such that $\sigma(i_1) < \dots < \sigma(i_\ell)$.*

We prove the above two lemmas at the end of the section, but first we deduce Lemma 3.2.

Proof of Lemma 3.2. Consider any ordering π and vertex set X , with $|X| \geq 100t^3 \log t$ as in the lemma statement. By Lemma 3.3, there are at least $t^2 - t + 1 - t^2/2 = t^2/2 - t + 1$ lines which contain at least $|X|/(2t)$ vertices of X . For any such line $\Lambda \subseteq V(G)$, let \mathcal{E}_Λ be the event that there is an independent set of size $|X|/(16t^3)$ in $G_\pi[\Lambda \cap X]$ with each vertex coming from a different bucket. Note that each of the \mathcal{E}_Λ are mutually independent, and by Lemma 3.4 each $\Pr(\mathcal{E}_\Lambda) \geq 1 - e^{-|X|/(48t)}$. So, the probability no \mathcal{E}_Λ holds is at most $e^{-(t^2/2-t+1)|X|/(48t)} \leq e^{-t|X|/97}$. This probability is small enough to take the union bound over all $|V(G)|! \leq e^{t^4 \log t}$ orderings π and at most $2^{|V(G)|} = e^{O(t^4)}$ subsets X . □

3.1. Quasi-randomness

In this section, we prove Lemma 3.3. We will want a bipartite version of the expander mixing lemma, due to Haemers [6, Theorem 5.1] (see also [4, Lemma 8] for a version with notation closer to what we use here). Note that, for a bipartite graph H , the non-zero singular values of the bipartite adjacency matrix of H are in correspondence with the positive eigenvalues of the (non-bipartite) adjacency matrix of H .

LEMMA 3.5. *Let H be a biregular bipartite graph with parts A and B , where every vertex in A has degree a and every vertex in B has degree b . Let M be the $|A| \times |B|$ bipartite adjacency matrix of H , and let $\sigma_1 \geq \sigma_2 \geq \dots$ be the singular values of M . Then for any sets $X \subseteq A$ and $Y \subseteq B$, we have*

$$\left| e(X, Y) - \frac{\sqrt{ab}}{\sqrt{|A||B|}} |X||Y| \right| \leq \sigma_2 \sqrt{|X||Y|}.$$

To apply Lemma 3.5, we will want to study the singular values of a bipartite graph related to Definition 3.1.

LEMMA 3.6. *Let G be as in Definition 3.1, and let H be the bipartite graph with parts A and B defined as follows. Let A be the set of vertices of G , let B be the set of lines in the projective plane underlying G , and put an edge between a vertex and a line if the vertex is in a bucket corresponding to a point on that line. Let M be the bipartite adjacency matrix of H . Then the non-zero singular values of M are \sqrt{kt} with multiplicity 1, and $\sqrt{k(t-1)}$ with multiplicity $t^2 - t$.*

Proof. Let L be the $(t^2 - t + 1) \times (t^2 - t + 1)$ point-line incidence matrix associated with the projective plane underlying G . Then M can be represented as the Kronecker (tensor) product $K \otimes L$, where K is the $k \times 1$ all-ones matrix. Note that $L^T L$ has t for all diagonal entries, and 1 for all off-diagonal entries, because in an order- $(t + 1)$ projective plane, every point lies in t lines, and every pair of points lie in exactly one line. $L^T L$ has eigenvalues $t + (t^2 - t + 1) - 1 = t^2$ (with eigenvector $(1, \dots, 1)$) and $t - 1$ (with eigenvectors of the form $(-1, 0, \dots, 0, 1, 0, \dots)$). Also, $K^T K$ is the 1×1 matrix whose entry is k . So, $M^T M = (K^T \otimes L^T)(K \otimes L) = (K^T K \otimes L^T L) = kL^T L$, and the desired claim about singular values follows. □

Finally, combining the above two lemmas yields Lemma 3.3, as follows.

Proof of Lemma 3.3. Consider the bipartite graph H in Lemma 3.6. Note that this graph has parts of size $t^2 - t + 1$ and $k(t^2 - t + 1)$. All the degrees in the first part are kt and all the degrees in the second part are t . Suppose more than $t^2/2$ lines contain fewer than $|X|/(2t)$ vertices. Then in H we have $X \subseteq A$ and $Y \subseteq B$ with $|X| \geq 9kt$, $|Y| > t^2/2$ and $e(X, Y) < |Y||X|/(2t)$. Then by Lemma 3.5 (with $\sigma_2 = \sqrt{k(t-1)}$, coming from Lemma 3.6),

$$(1 + o(1)) \frac{|X||Y|}{2t} \leq \left| e(X, Y) - \frac{1 + o(1)}{t} |X||Y| \right| \leq \sqrt{k(t-1)} \sqrt{|X||Y|},$$

a contradiction. □

3.2. Long increasing subsequences

Proof of Lemma 3.4. First, consider independent random variables $\alpha_1, \dots, \alpha_m$ each uniform in the interval $(0, 1]$ (these are distinct with probability 1), and note that we can define a

uniformly random permutation $\sigma : \{1, \dots, m\} \rightarrow \{1, \dots, m\}$ by taking $\alpha_{\sigma(1)} < \dots < \alpha_{\sigma(m)}$. Then, divide the interval $(0, 1]$ into $m/(2k)$ equal-sized sub-intervals $I_1, \dots, I_{m/(2k)}$ (so $I_y = (2(y-1)k/m, 2yk/m]$ for each y). We then define a random zero-one matrix $M \in \{0, 1\}^{(m/(2k)) \times m}$, where $M(y, x) = 1$ if $\alpha_x \in I_y$, and $M(y, x) = 0$ otherwise. Note that the columns of this matrix are independent, each uniform among all zero-one vectors with exactly one ‘1’.

Next, in addition to the existing division into buckets, we also divide $\{1, \dots, m\}$ into m/k (discrete) intervals we call ‘slices’: for $1 \leq r \leq m/k$, let $S_r = \{(r-1)k + 1, \dots, rk\}$. For $0 \leq q < m/(2k)$, define the ‘ q th diagonal strip’ $D_q = \{(y, x) : 1 \leq y \leq m/(2k), x \in S_{y+q}\}$ (which we interpret as a set of positions in our matrix M), and note that these strips each have size $m/2$ and are disjoint for different q .

Now, for each $0 \leq q < m/(2k)$ consider the following procedure to find a suitable sequence $x_1 < \dots < x_\ell$ (which we call ‘phase q ’). We order the positions in D_q from left to right, and scan through the corresponding entries of M until we see a ‘1’. We take the corresponding column as x_1 , then we continue scanning, skipping all entries that share a bucket or a slice with x_1 (without exposing their values) until we find another ‘1’, whose column we take as x_2 . We continue in this fashion, continually scanning along D_q and skipping all elements that share a bucket or a slice with a previously selected element, until we have constructed a sequence $x_1 < \dots < x_\ell$ of length ℓ or we have reached the end of the entries indexed by D_q . In the latter case, we say that the procedure fails, and denote the corresponding event by \mathcal{E}_q . Note that during phase q we exposed only some subset of the entries indexed by D_q (though the entries are dependent, we indirectly revealed some partial information about entries in other diagonal strips). Among the entries we exposed, at most ℓ of these are ‘1’s. See Figure 3.1 below.




FIGURE 3.1. An example outcome of the first two phases. In this case, $k = 3$, $m = 24$ and we are looking for an increasing sequence of length $\ell = 3$ (these values of k, m, ℓ were chosen to yield a simple picture, but note that they do not actually satisfy the conditions in the lemma statement). First, in phase 0 we searched from left to right through the light grey cells in D_0 . We first saw a ‘1’ in column 4, meaning that we take $x_1 = 4$. Then we continued searching, skipping columns that share a bucket or slice with x_1 (the skipped cells are labelled with ‘*’; here 4 and 8 share a bucket). We eventually found a second ‘1’ in column 11, so we take $x_2 = 11$, but we were not able to find a third ‘1’ so the first phase failed (\mathcal{E}_0 occurred). In phase 1, we scanned through the dark grey cells in D_1 in the same way. This phase succeeded, finding $(x_1, x_2, x_3) = (5, 12, 14)$.

Observe that if any of the \mathcal{E}_q do not occur (that is, some phase q succeeds), then the sequence $x_1 < \dots < x_\ell$ produced in phase q satisfies the conditions of the lemma. It therefore suffices to show that $\Pr(\bigcap_q \mathcal{E}_q) \leq e^{-m/24}$. We will accomplish this by showing that $\Pr(\mathcal{E}_q | \mathcal{E}_0 \cup \dots \cup \mathcal{E}_{q-1}) \leq e^{-k/12}$ for each of the $m/(2k)$ choices of q . So, fix some q , and condition on any outcome of phases 0 to $q-1$.

Now, suppose that phase q failed (meaning \mathcal{E}_q occurred). Since each bucket and slice has size at most k , during the scanning procedure we skipped fewer than $2\ell k$ entries (out of the $m/2$ total entries indexed by D_q) for sharing a bucket or slice with a previously chosen element. Also, since each phase exposes at most ℓ ‘1’s, there are at most ℓq columns in which we saw a ‘1’ in previous phases.

Now, for each $(y, x) \in D_q$ for which we have not already seen a ‘1’ in column x in a previous phase, the probability that $M_{y,x} = 1$ is at least $2k/m$ (it could be greater than this if we have already seen a ‘0’ in column x in a previous phase). So, $\Pr(\mathcal{E}_q | \mathcal{E}_0 \cup \dots \cup \mathcal{E}_{q-1})$ is upper

bounded by $\Pr(X < \ell)$, where X has the binomial distribution $\text{Bin}(m/2 - 2\ell k - \ell q, 2k/m)$. Recalling that $\ell = m/(8k)$, $m \leq k^2$ and $q < m/(2k)$ (so $\mathbb{E}X \geq 3k/8$ and $3k/8 - \ell \geq k/4$), this probability is at most

$$\exp\left(-\frac{(k/4)^2}{2(3k/8)}\right) = e^{-k/12}$$

by a Chernoff bound (see, for example, [9, Theorem 2.1]). □

4. Concluding remarks

Recall that $g(k)$ is the minimum integer such that every $g(k)$ -vertex tournament has an acyclic k -chromatic subgraph. Here we proved that $k^{4/3-o(1)} \leq g(k) \leq k^{9/5+o(1)}$. We further suspect that the lower bound may be closer to the truth and propose the following conjecture.

CONJECTURE 4.1. Every n -vertex tournament has an acyclic subgraph with chromatic number at least $n^{3/4-o(1)}$. That is, $g(k) \leq k^{4/3+o(1)}$.

Our proof of the lower bound $g(k) \geq k^{4/3-o(1)}$ is based on a certain randomized construction involving a projective plane, but we wonder if the following simpler construction may also attain the same bound. Suppose that $n = q^2$ is a perfect square, and consider the tournament on the vertex set $\{1, \dots, n\}$ where, for $i < j$, the edge ij is oriented from i to j unless $i - j$ is divisible by q , in which case it is oriented from j to i . This particular tournament was used by Nassar and Yuster to prove their lower bound $g(k) \geq n^{8/7}/4$.

Acknowledgements. The second author would like to thank Asaf Ferber for helpful discussions about multiple-exposure arguments for random permutations (which turned out to be helpful for the proof of Lemma 3.4). The third author would like to thank Noga Alon for bringing the problem of Nassar and Yuster to his attention.

References

1. L. ADDARIO-BERRY, F. HAVET, C. L. SALES, B. REED and S. THOMASSÉ, ‘Oriented trees in digraphs’, *Discrete Math.* 313 (2013) 967–974.
2. N. ALON, J. PACH and J. SOLYMOSSI, ‘Ramsey-type theorems with forbidden subgraphs’, *Combinatorica* 21 (2001) 155–170. (Paul Erdős and his mathematics (Budapest, 1999).)
3. S. A. BURR, ‘Subtrees of directed graphs and hypergraphs’, *Congr. Numer.* 28 (1980) 227–239.
4. S. DE WINTER, J. SCHILLEWAERT and J. VERSTRAETE, ‘Large incidence-free sets in geometries’, *Electron. J. Combin.* 19 (2012), 24.
5. T. GALLAI, ‘On directed paths and circuits’, *Theory of graphs* (eds P. Erdős and G. Katona; Academic Press, New York, 1968) 115–118.
6. W. H. HAEMERS, ‘Interlacing eigenvalues and graphs’, *Linear Algebra Appl.* 226/228 (1995) 593–616.
7. M. HASSE, ‘Zur algebraischen Begründung der Graphentheorie. I’, *Math. Nachr.* 28 (1964/65) 275–290.
8. F. HAVET, ‘Orientations and colouring of graphs’, Lecture notes of SGT Oleron, France, 2013.
9. S. JANSON, T. ŁUCZAK and A. RUCINSKI, *Random graphs*, Wiley-Interscience Series in Discrete Mathematics and Optimization (Wiley-Interscience, New York, NY, 2000).
10. D. KÜHN, R. MYCROFT and D. OSTHUS, ‘A proof of Sumner’s universal tournament conjecture for large tournaments’, *Proc. Lond. Math. Soc.* (3) 102 (2011) 731–766.
11. S. NASSAR and R. YUSTER, ‘Acyclic subgraphs with high chromatic number’, *European J. Combin.* 75 (2019) 11–18.
12. V. RÖDL and P. WINKLER, ‘A Ramsey-type theorem for orderings of a graph’, *SIAM J. Discrete Math.* 2 (1989) 402–406.
13. B. ROY, ‘Nombre chromatique et plus longs chemins d’un graphe’, *Rev. Française Informat. Recherche Opérationnelle* 1 (1967) 129–132.
14. L. M. VITAYER, ‘Determination of minimal coloring of vertices of a graph by means of Boolean powers of the incidence matrix’, *Dokl. Akad. Nauk* 147 (1962) 758–759.

Jacob Fox and Matthew Kwan
Department of Mathematics
Stanford University
450 Jane Stanford Way
Stanford, CA 94305
USA

jacobfox@stanford.edu
mattkwan@stanford.edu

Benny Sudakov
Department of Mathematics
ETH
Ramistrasse 101
Zurich 8092
Switzerland

benny.sudakov@gmail.com