In this talk, I will discuss dependent random choice, which is a surprisingly powerful technique in extremal graph theory and related areas. Almost all of the material is drawn from the excellent survey of Fox and Sudakov.

1 Background: Kővári–Sós–Turán

Kővári, Sós, and Turán introduced an extremely simple, but surprisignly versatile averaging technique in their study of the problem of Zarankiewicz, aka the extremal numbers of bipartite graphs. In my opinion, the "right" way to think of dependent random choice is as a generalization/strengthening of the Kővári–Sós–Turán technique, so I want to start with their result. The following is a convenient way of stating (a special case) of their result.

Theorem 1.1. *Fix an integer s. Let G be an n-vertex graph with average degree $d = \varepsilon n$. Then G contains an s-tuple of vertices which has at least $\varepsilon^s n - \binom{s}{2}$ common neighbors.*

In every application, the additive error $\binom{s}{2}$ is going to be much smaller than the main term $\varepsilon^s n$. So, roughly speaking, we are able to find an s-tuple of vertices with $\gtrsim \varepsilon^s n$ common neighbors. This is essentially best possible, since in a random graph of edge density ε, w.h.p. every s-set of vertices has $(\varepsilon^s \pm o(1))n$ common neighbors.

Proof. Let $u_1, \ldots, u_s$ be iid uniformly random vertices of G, and let $S = \{u_1, \ldots, u_s\}$. For a fixed vertex v, the probability that v is a common neighbor of S is $(\deg(v)/n)^s$. Therefore, if we let Y denote the set of common neighbors of S, we see that

$$\mathbb{E}[|Y|] = \sum_{v \in V(G)} \Pr(v \in Y) = \sum_{v \in V(G)} \left(\frac{\deg(v)}{n}\right)^s \geq n \cdot \left(\frac{d}{n}\right)^s = \varepsilon^s n, \tag{1}$$

where the inequality follows from convexity of the function $x \mapsto x^s$. Hence, there exists such a set S with at least $\varepsilon^s n$ common neighbors.

Unfortunately, we are not yet done, since it is possible that the vertices $u_1, \ldots, u_s$ are not all distinct, and thus that $|S| < s$. Note that $\Pr(u_i = u_j) = 1/n$, and hence $\Pr(|S| < s) \leq \binom{s}{2}/n$ by the union bound. Note too that any set of vertices in G has at most n common neighbors. Therefore,

$$\begin{aligned} \varepsilon^s n &\leq \mathbb{E}[|Y|] \\ &= \mathbb{E}[|Y| \mid |S| = s] \Pr(|S| = s) + \mathbb{E}[|Y| \mid |S| < s] \Pr(|S| < s) \\ &\leq \mathbb{E}[|Y| \mid |S| = s] + n \cdot \frac{\binom{s}{2}}{n}. \end{aligned}$$

Rearranging gives the desired result. □

As an immediate corollary, we see that for any fixed integers $s \leq t$, every n-vertex graph G with $\Omega_{s,t}(n^{2-1/s})$ edges contains a copy of $K_{s,t}$. Indeed, in such a graph, the average degree is $\Omega(n^{1-1/s}) = \Omega(n^{-1/s}) \cdot n$. Letting $\varepsilon = \Omega(n^{-1/s})$ and choosing the implicit constant large enough, we see that G contains an s-set with at least $\varepsilon^s n - \binom{s}{2} \geq t$ common neighbors, i.e. a copy of $K_{s,t}$.

2 Basic dependent random choice

In Theorem 1.1, we found a *single* s-set with many common neighbors, where "many" means at least roughly $\varepsilon^s n$. However, in our example of a random graph, we actually see that *every* s-set has this property. Is it possible to strengthen Theorem 1.1 to ensure that all s-sets have many common neighbors?

Trivially, the answer is no. For example, let G be the disjoint union of two copies of $K_{n/2}$. Then G has average degree around $n/2$. But most s-sets, namely those with vertices in both components, have no common neighbors at all. So something as strong as what was suggested in the previous paragraph clearly cannot happen.

Nonetheless, it is clear what the "problem" with this graph G is: it's disconnected! All the s-sets which lie in only one component do indeed have many common neighbors. This suggests that perhaps we can get a result for all s-sets, as long as we allow ourselves to restrict this "all" to all s-sets coming from a fixed subset of the vertices. This is indeed true, and is the content of the next theorem.

Theorem 2.1. *Fix an integer s. Let G be an n-vertex graph with average degree $d = \varepsilon n$. Let m be any integer. Then there exists $W \subseteq V(G)$ with the following properties:*

- *Every s-tuple of vertices in W has at least m common neighbors, and*
- *$|W| \geq \varepsilon^s n - m^s$.*

In the proof of Theorem 1.1, we picked a *uniformly* random s-tuple of vertices and showed that in expectation, it had many common neighbors. In this proof, we will use a non-independent random choice (hence the name of the technique). Namely, we will first pick a random s-tuple of vertices, and we will then define Y to be the set of common neighbors of this s-tuple. Then Y is a random, but certainly not uniformly random, set of vertices. We then delete some vertices from Y to obtain W, and show that this choice of W satisfies the desired properties with positive probability. Here are the details of the proof.

Proof. Let $u_1, \ldots, u_s$ be iid uniformly random vertices of G, and let $S = \{u_1, \ldots, u_s\}$. Let Y denote the set of common neighbors of S. By (1), we see that $\mathbb{E}[|Y|] \geq \varepsilon^s n$. Let Z denote the number of s-tuples in Y which have at most m common neighbors. We now estimate $\mathbb{E}[Z]$.

Fix an s-tuple T of vertices with at most m common neighbors. The event that $T \subseteq Y$ is the same as the event that $u_1, \ldots, u_s$ are all in the common neighborhood of T. As this common neighborhood has size at most m, we see that $\Pr(T \subseteq Y) \leq (m/n)^s$. As there are at most $\binom{n}{s}$ choices for T, we conclude that

$$\mathbb{E}[Z] \leq \binom{n}{s} \cdot \left(\frac{m}{n}\right)^s \leq m^s.$$

Therefore, $\mathbb{E}[|Y| - Z] \geq \varepsilon^s n - m^s$. Thus, for some fixed realization of S, we have that the corresponding realizations of Y, Z satisfy $|Y| - Z \geq \varepsilon^s n - m^s$. Let W be a subset of Y obtained by deleting one vertex from each s-tuple counted from Z. Then W satisfies the two desired properties. $\square$

As an immediate corollary, we obtain the following generalization of the Kővári–Sós–Turán bound on the extremal number of complete bipartite graphs. The result is originally due to Füredi, and its simplified proof using dependent random choice is due to Alon, Krivelevich, and Sudakov.

Theorem 2.2. *Let H be a bipartite graph with bipartition $A \cup B$, and suppose that every vertex in B has degree at most s. Then $\mathrm{ex}(n, H) = O_H(n^{2-1/s})$. In other words, if G is an n-vertex graph with average degree $\Omega_H(n^{1-1/s})$, then G contains a copy of H.*

Plugging in $H = K_{s,t}$ recovers our earlier result, but this theorem is of course much more general.

Proof. Let G have average degree $d = \varepsilon n$, where $\varepsilon = \Omega_H(n^{-1/s})$. We apply Theorem 2.1 with $m = |B|$, to conclude that there is $W \subseteq V(G)$ satisfying two properties. First, every s-tuple of vertices in W has at least m common neighbors, and second,

$$|W| \geq \varepsilon^s n - m^s \geq \Omega_H(1) - O_H(1) \geq |A|,$$

by choosing the implicit constants appropriately. Pick an arbitrary embedding of A into W. Now proceed one by one along the vertices of B. Each one has at most s neighbors in A, and by the choice of W, each such s-tuple has at least $|B|$ common neighbors in G. So we may arbitrarily embed the vertices of B one at a time, and find a copy of H in G. □

In fact, if we examine the proof of Theorem 2.1, we see that we prove Theorem 2.2 with a weaker assumption on H. Namely, note that by the construction of W in the proof of Theorem 2.1, we know that all vertices in W have s common neighbors in G, namely[1] $u_1, \ldots, u_s$. We can also use these vertices to embed vertices of B. By doing this, we can weaken the assumption in Theorem 2.2 to say that all vertices of B, with at most s exceptions, have degree at most s. This implies, for example, an upper bound of $O(n^{3/2})$ on the extremal number of $K_{3,3} \setminus e$.

3 Slightly less basic dependent random choice

If we examine the proof of Theorem 2.1, we notice that the parameter s plays two different roles. On the one hand, it is a parameter in the statement, controlling the sizes of tuples we're interested in. On the other hand, it's also a parameter in the proof, controlling how many random vertices are chosen in defining the common neighborhood from which we build W. Other than simplifying some computations (and sufficing in some applications), there is no reason to make these two numbers equal; making them distinct yields the following generalization of Theorem 2.1.

Theorem 3.1. *Fix an integer s. Let G be an n-vertex graph with average degree $d = \varepsilon n$. Let m, t be any integers. Then there exists $W \subseteq V(G)$ with the following properties:*

[1] This is not 100% true, as these vertices may not be distinct. But it is easy to modify the proof to ensure their distinctness (and obtain a slightly worse bound on $|W|$).

- *Every s-tuple of vertices in W has at least m common neighbors, and*
- $|W| \geq \varepsilon^t n - m^t n^{s-t}$.

Proof. The proof is identical to that of Theorem 2.1. We now pick $u_1, \ldots, u_t$ at random, and let Y be their common neighborhood. Then by (1), we know that $\mathbb{E}[|Y|] \geq \varepsilon^t n$. On the other hand, letting Z denote the number of s-tuples in Y with fewer than m common neighbors, we see that $\mathbb{E}[Z] \leq \binom{n}{s}(m/n)^t \leq m^t n^{s-t}$. Applying linearity of expectation as before yields the desired result. □

In most applications of Theorem 3.1, the added flexibility over Theorem 2.1 is used as follows. We want to take m to be fairly large, e.g. $m \approx \sqrt{n}$. In Theorem 2.1, this is not possible, as then m^s will be much larger than $\varepsilon^s n$, so the set W we get is empty (and thus useless). But Theorem 3.1 lets us get around this problem by choosing $t \gg s$, so that the term $m^t n^{s-t} = n^s(m/n)^t$ becomes very small. Of course, we pay for this in that the term $\varepsilon^t n$ is also smaller, but in many applications there is plenty of wiggle room on this main term, so shrinking it is OK. I will show two applications that demonstrate this idea.

The first application can be viewed as a refinement of Theorem 2.2. Let $\widehat{K_r}$ denote the 1-subdivision of K_r; this is a bipartite graph with parts of size r and $\binom{r}{2}$, where every pair of vertices in the first part has a unique common neighbor in the second part, and where every vertex in the second part has degree two. By Theorem 2.2, we know that $\mathrm{ex}(n, \widehat{K_r}) = O_r(n^{3/2})$. However, Erdős asked about what can be said for graphs of constant edge density: if G is an n-vertex graph with average degree εn, what is the largest r for which we can guarantee $\widehat{K_r} \subseteq G$? As $\widehat{K_r}$ has $r + \binom{r}{2} = \Theta(r^2)$ vertices, it is clear that we must have $r = O(\sqrt{n})$ if $\widehat{K_r} \subseteq G$. The next result, due to Alon, Krivelevich, and Sudakov, shows that this is tight up to a constant factor.

Theorem 3.2. *Let G be an n-vertex graph with average degree εn, and let $r = \varepsilon^{3/2}\sqrt{n}$. Then $\widehat{K_r} \subseteq G$.*

Proof. Let $s = 2, m = r^2/2 = \varepsilon^3 n/2$, and $t = \log n/(2 \log \frac{1}{\varepsilon})$. Note that $\varepsilon^t = n^{-1/2}$, and hence

$$\varepsilon^t n = \sqrt{n} \qquad \text{and} \qquad m^t n^{s-t} = n^2 \left(\frac{m}{n}\right)^t = n^2 \cdot \frac{\varepsilon^{3t}}{2^t} \leq \frac{\sqrt{n}}{2}.$$

Therefore, by Theorem 3.1, there is $W \subseteq V(G)$ with $|W| \geq \sqrt{n}/2 \geq r$ such that every pair of vertices in W has at least $m \geq \binom{r}{2}$ common neighbors. By the same argument as in the proof of Theorem 2.2, we can embed the first part of $\widehat{K_r}$ in W and the second half among the common neighbors. □

The next result is a combination of Ramsey's theorem and the Kővári–Sós–Turán theorem. Recall that Ramsey's theorem implies that every n-vertex graph contains a clique or coclique[2] of order $\frac{1}{2} \log n$. Additionally, Theorem 1.1 implies that if G is an n-vertex graph with average degree εn, then G contains a complete bipartite subgraph with parts of size

[2]aka an independent set

$\log n/(2\log\frac{1}{\varepsilon})$. It would be nice if we could guarantee that this complete bipartite subgraph is an induced subgraph, but this is clearly impossible, since G might be a disjoint union of copies of $K_{\varepsilon n}$. However, the following result, essentially implicit in work of Fox and Sudakov, shows that we can have it be "almost" induced.

Theorem 3.3. *Let G be an n-vertex graph with average degree εn, and let $s = \log n/(4\log\frac{1}{\varepsilon})$. Then G contains a copy of $K_{s,s}$ such that each part is either a clique or a coclique.*

Note that a naive way of proving such a result is to first find a big complete bipartite subgraph, and then apply Ramsey's theorem in each part. But doing this would yield a bound of order $\log\log n$; the point of this theorem is that one can avoid doing the two steps one after another, and thus only lose one logarithm. This result is tight up to the constant factor, as shown by a random graph of edge density ε.

Proof. Let $s = \log n/(4\log\frac{1}{\varepsilon}), m = \sqrt{n}$, and $t = 2s$. Note that $\varepsilon^t = n^{-1/2}$, so

$$\varepsilon^t n = \sqrt{n} \qquad \text{and} \qquad m^t n^{s-t} = (\sqrt{n})^{2s} \cdot n^{-s} = 1.$$

By Theorem 3.1, we find $W \subseteq V(G)$ with $|W| \geq \sqrt{n} - 1$ such that every s-set of vertices from W has at least m common neighbors. Applying Ramsey's theorem to the induced subgraph on W, we find a clique or coclique of order $\frac{1}{2}\log(\sqrt{n}-1) \geq s$. This yields a clique or coclique of order s, which has at least $m = \sqrt{n}$ common neighbors. Applying Ramsey's theorem again in this common neighborhood gives the desired result. □

4 More advanced dependent random choice

Recall that in Theorem 1.1, we were able to guarantee the existence of an s-tuple whose common neighborhood is roughly as big as it would be in a random graph: if the edge density is ε, then we can find an s-tuple with roughly $\varepsilon^s n$ common neighbors. However, in the versions of dependent random choice we saw above, we were not able to get such a strong bound: we instead guarantee m common neighbors to each s-tuple, where m is necessarily much smaller than $\varepsilon^s n$.

It turns out that one can get a version of the dependent random choice lemma where $m \sim \varepsilon^s n$, but it requires weakening our condition slightly. Rather than getting that *every* s-tuple has many common neighbors, we instead get that *almost every* s-tuple has many common neighbors. The following is a representative example of this type of result.

Theorem 4.1. *Fix an integer s and a real number $\delta \in (0,1)$. Let G be an n-vertex graph with average degree $d = \varepsilon n$, and suppose that $d \geq 2s^2$. Let $m = \frac{\delta}{2}\cdot\varepsilon^s n$. There exists $W \subseteq V(G)$ with the following properties:*

- *The number of s-tuples of vertices in W with at least m common neighbors is at least $(1-\delta)\binom{|W|}{s}$, and*
- *$|W| \geq \varepsilon n/2$.*

Proof. For this proof, we pick a single random vertex u, and let Y be its common neighborhood. Then $\mathbb{E}[|Y|] = \varepsilon n$. Let Z denote the number of subsets of Y with at most m common neighbors; as before, we know that $\mathbb{E}[Z] \leq \binom{n}{s} \cdot (m/n) \leq mn^{s-1}/s!$.

The trick now is to consider the sth moment of the random variable $|Y|$. By Jensen's inequality, we know that $\mathbb{E}[|Y|^s] \geq \mathbb{E}[|Y|]^s$, and therefore

$$\mathbb{E}\left[|Y|^s - \frac{\mathbb{E}[|Y|]^s}{2\mathbb{E}[Z]} Z - \frac{\mathbb{E}[|Y|]^s}{2}\right] = \mathbb{E}[|Y|^s] - \mathbb{E}[|Y|]^s \geq 0.$$

Therefore, there exists a realization such that

$$|Y|^s \geq \frac{\mathbb{E}[|Y|]^s}{2\mathbb{E}[Z]} Z + \frac{\mathbb{E}[|Y|]^s}{2} \geq \frac{\varepsilon^s n^s}{mn^{s-1}/s!} Z + \frac{\varepsilon^s n^s}{2} = \frac{s!\varepsilon^s n}{m} Z + \frac{\varepsilon^s n^s}{2} = \frac{2s!}{\delta} Z + \frac{\varepsilon^s n^s}{2}.$$

In particular, we find that $|Y| \geq \varepsilon n/2^{1/s} \geq \varepsilon n/2$. Additionally, we see that

$$Z \leq \frac{\delta}{2s!}|Y|^s \leq \delta \binom{|Y|}{s},$$

where the final step uses the fact that $|Y| \geq \varepsilon n/2 \geq s^2 \geq 2\binom{s}{2}$ to conclude that

$$\binom{|Y|}{s} = \frac{|Y|^s}{s!} \prod_{i=0}^{s-1} \left(1 - \frac{i}{|Y|}\right) \geq \frac{|Y|^s}{s!} \left(1 - \sum_{i=1}^{s-1} \frac{i}{|Y|}\right) = \frac{|Y|^s}{s!} \left(1 - \frac{\binom{s}{2}}{|Y|}\right) \geq \frac{|Y|^s}{2s!}.$$

This shows that setting $W = Y$ gives the desired result. □

An important application of this version of dependent random choice is the Balog–Szemerédi–Gowers theorem in additive combinatorics[3]. We begin with a graph theory result, which is a quick consequence of Theorem 4.1.

Lemma 4.2. *There exist absolute constants $c, C > 0$ such that the following holds. Let $G = (A, B, E)$ be a bipartite graph with at least $\varepsilon|A||B|$ edges. There exist $A' \subseteq A, B' \subseteq B$ with $|A'| \geq c\varepsilon^C|A|, |B'| \geq c\varepsilon^C|B|$ with the property that for every $(a, b) \in A' \times B'$, the number of three-edge paths (a, x, y, b) joining them is at least $c\varepsilon^C|A||B|$.*

Note that this result is best possible up to the values of c, C; indeed, if $|A| = |B|$ and G is d-regular with $d = \varepsilon|A|$, then the number of three-edge paths between any two vertices is at most $d^2 = \varepsilon^2|A||B|$. Note too that it's crucial that we count paths of length three (specifically, length greater than one), for in a random bipartite graph of edge density ε, the largest complete bipartite subgraph has only logarithmically many vertices in one part.

[3] A good reference for the following material is Chapter 7.13 of Zhao's book *Graph Theory and Additive Combinatorics.*

Proof. Throughout this proof, c and C will represent arbitrary absolute constants, whose value can change from one line to the next.

By discarding low-degree vertices from A, we may assume that every vertex in A has at least $\varepsilon|B|/2$ neighbors in B. We now apply Theorem 4.1 to the resulting graph, with $s = 2$ and $\delta = \varepsilon/10$. By doing this[4], we find $W \subseteq A$ with $|W| \geq c\varepsilon|A|$, and with the property that all but a δ-fraction of pairs of vertices in W have at least $c\varepsilon^C|B|$ common neighbors in B.

Recall that we discarded the low-degree vertices from A, so in particular every vertex of W has at least $\varepsilon|B|/2$ neighbors in B. Thus, the induced subgraph on (W, B) has edge density at least $\varepsilon/2$. Let B' denote the set of vertices in B with at least $\varepsilon|W|/4$ neighbors in W, so that $|B'| \geq c\varepsilon|B|$.

Recall that there are at most $\delta\binom{|W|}{2}$ pairs of vertices in W which are "bad", meaning that they have fewer than $c\varepsilon^C|B|$ common neighbors in B. Let $A' \subseteq W$ denote the set of vertices of W that participate in at most $2\delta|W| = \varepsilon|W|/5$ bad pairs. Then $|A'| \geq c\varepsilon^C|A|$, and we claim that (A', B') satisfy the desired property.

To see this, fix $(a, b) \in A' \times B'$. Recall that b has at least $\varepsilon|W|/4$ neighbors in W, and that a is in a bad pair with at most $\varepsilon|W|/5$ other vertices of W. Thus, there are at least $\varepsilon|W|/20$ vertices $y \in W$ that are neighbors of b and don't form a bad pair with a. For each such y, by definition, there are at least $c\varepsilon^C|B|$ common neighbors x of a and y. This yields, in total, $c\varepsilon^C|A||B|$ paths of length three connecting a and b. □

The reason we care about these paths of length three is that they have an important consequence in additive combinatorics. Let X be a finite subset of some abelian group. We define the *sumset*

$$X + X \coloneqq \{x + y : x, y \in X\}$$

and the *additive energy*

$$E(X) \coloneqq |\{(w, x, y, z) \in X : w + x = y + z\}|.$$

It is easy to see that $|X| \leq |X + X| \leq |X|^2$ and that $|X|^2 \leq E(X) \leq |X|^3$. In general, we say that X "has a lot of additive structure" if $|X + X| \approx |X|$, or if $E(X) \approx |X|^3$. These are both natural notions of additive structure, and they both arise naturally in many contexts It would be nice to say that these two notions are equivalent. In one direction, a simple application of Cauchy–Schwarz shows that if $|X + X| \leq K|X|$, then $E(X) \geq |X|^3/K$, so the first notion of additive structure implies the second. However, the reverse direction is simply false, as shown by letting X be the union of an arithmetic progression of length n and n random elements. Then $|X + X| = \Theta(|X|^2)$, but $E(X) = \Theta(|X|^3)$. Thus, this set is additively structured in the second sense, but not in the first. Of course, it's clear what the "problem" is in this example: X contains a large subset which is additively structured in both senses. The Balog–Szemerédi–Gowers theorem says that this is in fact the only obstruction to equivalence.

[4]Strictly speaking, this doesn't quite follow from Theorem 4.1, since that was stated for a general graph, and here we want to apply it in a way that respects the bipartition, so that $W \subseteq A$. But it is easy to see that if we rerun the proof of Theorem 4.1, but now select $u \in B$ uniformly at random, we necessarily find that $W \subseteq A$, and everything else goes through as before.

Theorem 4.3 (Balog–Szemerédi–Gowers)**.** *There exists an absolute constant $C > 0$ such that the following holds. Suppose X is a finite subset of some abelian group satisfying $E(X) \geq |X|^3/K$. Then there exists $Y \subseteq X$ with $|Y| \geq |X|/K^C$ and $|Y + Y| \leq K^C|Y|$.*

This theorem is extremely useful in many applications, as it allows one to convert between the two notions of additive structure while invoking only a polynomial loss in the parameters. The original proof of Balog and Szemerédi had much worse quantitative dependence; the proof with the polynomial dependence is due to Gowers, and is one of the earliest applications of dependent random choice.

I won't prove Theorem 4.3 in full. Instead, I will prove the following result; there is a short and elementary argument showing that it implies Theorem 4.3. To state it, we need the following terminology: given a bipartite graph $G = (A, B, E)$, where the vertex sets A, B are subsets of some abelian group, we define the *restricted sumset*

$$A \underset{G}{+} B \coloneqq \{a + b : (a, b) \in E(G)\}.$$

Lemma 4.4. *Let $G = (A, B, E)$ be a bipartite graph, where A, B are subsets of some abelian group with $|A| = |B| = n$. Suppose that G has at least n^2/K edges and that $|A \underset{G}{+} B| \leq Kn$. Then there exist $A' \subseteq A, B' \subseteq B$ with $|A'|, |B'| \geq n/K^C$ and $|A' + B'| \leq K^C n$.*

Proof. We apply Lemma 4.2 with $\varepsilon = 1/K$, and claim that the resulting sets A', B' satisfy the desired result. Certainly the size lower bound follows immediately from the conclusion of Lemma 4.2, so it remains to prove the bound on the sumset. By definition, for every $(a, b) \in A' \times B'$, there exist at least n^2/K^C choices of $x \in B, y \in A$ such that (a, x, y, b) is a path in G. Note that

$$a + b = (a + x) - (x + y) + (y + b)$$

and that $(a+x), (x+y), (y+b) \in A \underset{G}{+} B$. Therefore, every element of $A' + B'$ can be written as a signed sum of three elements from $A \underset{G}{+} B$ in at least n^2/K^C different ways. Therefore,

$$|A' + B'| \leq \frac{\left|A \underset{G}{+} B\right|^3}{n^2/K^C} \leq \frac{K^3 n^3}{n^2/K^C} = K^{C+3} n. \qquad \square$$