

RESEARCH ARTICLE

The Turán number of the grid

Domagoj Bradač | Oliver Janzer | Benny Sudakov | István Tomon

Department of Mathematics, ETH Zürich,
Switzerland**Correspondence**Oliver Janzer, Department of
Mathematics, ETH, Zürich, Switzerland.
Email: oliver.janzer@math.ethz.ch**Funding information**SNSF, Grant/Award Number:
200021_196965; ETH Zürich,
Grant/Award Number: 20-1 FEL-35**Abstract**

For a positive integer t , let F_t denote the graph of the $t \times t$ grid. Motivated by a 50-year-old conjecture of Erdős about Turán numbers of r -degenerate graphs, we prove that there exists a constant $C = C(t)$ such that $\text{ex}(n, F_t) \leq Cn^{3/2}$. This bound is tight up to the value of C . One of the interesting ingredients of our proof is a novel way of using the tensor power trick.

MSC 2020

05C35 (primary)

1 | INTRODUCTION

For a graph H and a positive integer n , the Turán number (or extremal number), denoted $\text{ex}(n, H)$, is the maximum number of edges in an H -free graph on n vertices. Estimating this function for various choices of H is one of the most important problems in extremal graph theory. The celebrated Erdős–Stone–Simonovits theorem [8, 10] states that $\text{ex}(n, H) = (1 - \frac{1}{\chi(H)-1} + o(1)) \binom{n}{2}$, which determines the asymptotics of $\text{ex}(n, H)$ whenever $\chi(H) \geq 3$. For bipartite graphs the formula only gives $\text{ex}(n, H) = o(n^2)$, but it is known [18] that in this case there is some $\varepsilon = \varepsilon(H) > 0$ such that $\text{ex}(n, H) = O(n^{2-\varepsilon})$. It is conjectured that for every graph there is a (rational) number α such that $\text{ex}(n, H) = \Theta(n^\alpha)$. However, there are relatively few bipartite graphs H for which this is known. The existence of a suitable α has not been established even for some very simple graphs such as the cycle C_8 , the complete bipartite graph $K_{4,4}$ and the cube Q_3 . For a survey on the topic, we refer the reader to [14].

In general, it is not even clear which graph parameters determine the asymptotics of the Turán number of a bipartite graph. Nevertheless, in 1966, Erdős conjectured that $\text{ex}(n, H) = O(n^{2-1/r})$ holds for every r -degenerate bipartite graph H . If true, this estimate would be best possible. Füredi [12] proved a special case of the conjecture when H has maximum degree r on one side of the bipartition. Alon, Krivelevich and Sudakov [1] gave a new proof of this result, using a

powerful probabilistic technique, called dependent random choice (see, for example, [11] for a description of this method and a brief history). They also proved a bound of $O(n^{2-\frac{1}{4r}})$ for general r -degenerate H . However, Erdős’ conjecture in full generality remains widely open even in the case $r = 2$.

In this paper we study Turán numbers of grids. For an integer $t \geq 2$, the grid graph F_t has vertex set $[t] \times [t]$ with two vertices joined by an edge if they differ in exactly one coordinate and in that coordinate they differ by exactly one. It is not difficult to see that F_t is a 2-degenerate graph, and therefore by Erdős’ conjecture, one can expect its Turán number to be $O(n^{3/2})$. Recently, extremal problems involving grids have attracted considerable attention. Let us briefly mention a few of these papers. Clemens, Miralaei, Reding, Schacht and Taraz [5] gave an upper bound for the size Ramsey number of the $t \times t$ grid. Their bound was improved very recently by Conlon, Nenadov and Trujić [6]. Kim, Lee and Lee [17] proved Sidorenko’s conjecture for grids (in arbitrary dimension). Füredi and Ruszinkó [13] studied the maximum number of hyperedges that a t -uniform linear hypergraph can have without containing a certain $t \times t$ hypergraph grid; see also [15].

Our main result determines $\text{ex}(n, F_t)$ up to a multiplicative constant, adding grids to the rather short list of families of bipartite graphs whose Turán number is known.

Theorem 1.1. *For any integer $t \geq 2$, there exist positive real numbers $c = c(t)$ and $C = C(t)$ such that*

$$cn^{3/2} \leq \text{ex}(n, F_t) \leq Cn^{3/2}.$$

In fact, our proof works for a slightly more general class of graphs. For graphs G and H , the Cartesian product $G \square H$ is the graph whose vertex set is $V(G) \times V(H)$ and in which (u, v) is adjacent to (u', v') if and only if $u = u'$ and $vv' \in E(H)$ or $uu' \in E(G)$ and $v = v'$. Writing P_t for the path with t vertices, our Theorem 1.1 states that $\text{ex}(n, P_t \square P_t) = \Theta_t(n^{3/2})$. We can prove the following more general version.

Theorem 1.2. *For any tree T and any path P (each with at least one edge), there exist positive real numbers c and C such that*

$$cn^{3/2} \leq \text{ex}(n, T \square P) \leq Cn^{3/2}.$$

Since $T \square P$ contains the 4-cycle as a subgraph, the lower bound in Theorem 1.2 follows from $\text{ex}(n, C_4) = \Theta(n^{3/2})$ (see, for example, [18]).

In the next section, we outline the proof of the weaker bound $\text{ex}(n, F_t) = O_t(n^{3/2}(\log n)^t)$. We get rid of the $(\log n)^t$ factor using the tensor power trick. This technique was used by Alon and Ruzsa [2] to give an elementary proof of the celebrated Sidorenko’s conjecture for trees, which implies the Blakley–Roy matrix inequality [3]. The technique has also been used in many other areas, see, for example, Tao [20], who has collected a number of these applications. Our way of using this tool somewhat differs from previous applications. Notably, we cannot deduce the correct bound on $\text{ex}(n, F_t)$ by proving the aforementioned weaker bound and applying it as a black box to the k th tensor power G^k of the given graph G ; instead we work directly with G^k while keeping in mind the original graph G . To the best of our knowledge, this is the first application of the tensor power trick to problems about Turán numbers.

1.1 | An outline of the proof

In this section we give a sketch of the proof of $\text{ex}(n, F_t) = O_t(n^{3/2}(\log n)^t)$. Let G be an n -vertex graph with $\alpha n^{3/2}$ edges, where $\alpha = (\log n)^t$. We want to show that G contains the $t \times t$ grid as a subgraph. We may assume by standard reduction results that G is almost regular, that is, $\Delta(G) \leq K\delta(G)$ for some absolute constant K .

An important element of the proof will be a careful count of the $2 \times t$ grids (that is, subgraphs isomorphic to $P_2 \square P_t$) in G . This graph will be the building block for our $t \times t$ grids. More precisely, note that we can obtain an $(s+1) \times t$ grid by gluing a $2 \times t$ grid to an $s \times t$ grid along some side of length $t-1$. Thus, by successively attaching $2 \times t$ grids, we end up with a $t \times t$ grid as long as the attached $2 \times t$ grids do not intersect the already existing grid in any vertex other than the ones on the boundary that are to be glued.

It is not too hard to see that G contains roughly at least as many $2 \times t$ grids as a random graph with the same edge density, namely, $\alpha^{3t-2} n^{t/2+1}$. By the maximum degree condition, the number of paths of length $t-1$ in G is at most about $\alpha^{t-1} n^{t/2+1/2}$ (ignoring absolute constant factors). This is promising as it shows that on average a path of length $t-1$ can be extended to many (at least $\alpha^{2t-1} n^{1/2}$) $2 \times t$ grids. As long as we can make sure that not too many of these extensions contain any fixed vertex, we can conclude that G contains a $t \times t$ grid. However, this is a significant obstacle since in general we do not have any good upper bound on the number of extensions of a path to a $2 \times t$ grid which contain a given vertex. For this reason, we will look at a certain subfamily \mathcal{F} of all $2 \times t$ grids such that $|\mathcal{F}|$ is still large enough, but for which we will be able to bound efficiently the number of members of \mathcal{F} that extend a fixed path of length $t-1$ and contain a fixed vertex.

In order to describe \mathcal{F} , it is helpful to label the vertices of the $2 \times t$ grid. Let us call them $x_1, x_2, \dots, x_t, y_1, y_2, \dots, y_t$ where $x_1 x_2 \dots x_t$ and $y_1 y_2 \dots y_t$ are paths of length $t-1$, and $x_i y_i$ is an edge for every i . Now \mathcal{F} will consist of those $2 \times t$ grids in G for which the codegree of x_{i+1} and y_i (in G) is at most s_i for each $1 \leq i \leq t-1$, where s_1, s_2, \dots, s_{t-1} depend on the graph G . We claim that for suitable choices of $s_1, \dots, s_{t-1} \geq \alpha$, $|\mathcal{F}|$ is at least about $n(\alpha n^{1/2})^t s_1 s_2 \dots s_{t-1} / (\log n)^{t-1}$. Note that there is a trade-off between the control on the codegrees and the lower bound on $|\mathcal{F}|$: the weaker control we have on the codegrees, the bigger \mathcal{F} is guaranteed to be.

Before we explain how we obtain the lower bound on $|\mathcal{F}|$, let us see how this guarantees that a $t \times t$ grid can be found. The number of paths of length $t-1$ in G is about $n(\alpha n^{1/2})^{t-1}$, so on average a path of length $t-1$ extends to at least $\alpha n^{1/2} s_1 s_2 \dots s_{t-1} / (\log n)^{t-1}$ members of \mathcal{F} . What we need to show is that only a small proportion of those can contain a fixed vertex. For this, fix the (images of the) vertices x_1, \dots, x_t and let us count how many members of \mathcal{F} extend these vertices which map y_j to a fixed vertex $u \in V(G)$ where $2 \leq j \leq t$ is given (the case $j=1$ is slightly different but arguably easier). Note that y_1 has to be a neighbour of x_1 , so there are at most $\Delta(G) \approx \alpha n^{1/2}$ choices for it. Then y_2 has to be a common neighbour of x_2 and y_1 , both of which are already fixed and need to have codegree at most s_1 by the definition of \mathcal{F} , so there are at most s_1 choices for y_2 . Similarly, since y_i needs to be a common neighbour of x_i and y_{i-1} , once y_1, y_2, \dots, y_{i-1} have been chosen, there are at most s_{i-1} choices for y_i . Moreover, there is a unique choice for y_j as we are counting only those extensions which map y_j to u . Altogether, we get at most $\alpha n^{1/2} s_1 s_2 \dots s_{j-2} s_j \dots s_{t-1}$ extensions mapping y_j to u . Since $s_{j-1} \geq \alpha = (\log n)^t$, we have that $\alpha n^{1/2} s_1 s_2 \dots s_{j-2} s_j \dots s_{t-1} = o(\alpha n^{1/2} s_1 s_2 \dots s_{t-1} / (\log n)^{t-1})$, so the number of extensions of any x_1, \dots, x_t to a member of \mathcal{F} mapping y_j to u is indeed negligible compared to the typical number of extensions of a fixed x_1, \dots, x_t to a member of \mathcal{F} . Moreover, since the conditions on \mathcal{F} are symmetrical with respect to x and y , only a small proportion of all extensions of a given

embedding of the y_i s to a member of \mathcal{F} use any given vertex. It is not hard to see that this implies that a $t \times t$ grid can be found in G .

We conclude this outline by sketching how to find enough $2 \times t$ grids satisfying the codegree conditions. We first choose x_1 and y_1 for which there are $n \cdot \alpha n^{1/2}$ possibilities. Then we choose x_2 to be a neighbour of x_1 which can be done in $\alpha n^{1/2}$ many ways. By dyadic pigeonholing, there is some s_1 such that for at least $1/\log n$ proportion of all the partial $2 \times t$ grids that were constructed so far, the codegree of x_2 and y_1 is around s_1 (up to a constant factor). For each of these partial grids, we can choose y_2 in around s_1 many ways. Then we choose x_3 to be an arbitrary neighbour of x_2 and again apply dyadic pigeonholing to keep at least $1/\log n$ proportion of our partial grids in all of which x_3 and y_2 have codegree around s_2 . Continuing this process, we end up with the desired number of $2 \times t$ grids. In turning this into a rigorous proof, the difficulty is making sure that s_1, s_2, \dots, s_{t-1} are at least α (or even that there is always at least one choice for y_i which is different from $x_1, y_1, x_2, y_2, \dots, x_{i-1}, y_{i-1}$).

Notation. For a graph G , we use $\delta(G), \bar{d}(G)$ and $\Delta(G)$ to denote the minimum, average and maximum degree of G , respectively. For vertices $u, v \in V(G)$, we write $d_G(u, v)$ for the codegree of u and v in G . We omit floor and ceiling signs whenever this does not affect the argument. All logarithms are to the base 2.

2 | THE PROOF OF THEOREM 1.2

A graph G is called K -almost-regular if $\Delta(G) \leq K\delta(G)$. A well-known regularization lemma of Erdős and Simonovits [9] allows us to restrict our attention to almost-regular graphs. We use the following version of the lemma proved by Jiang and Seiver.

Lemma 2.1 (Jiang–Seiver [16]). *Let ε, c be positive reals, where $\varepsilon < 1$ and $c \geq 1$. Let n be a positive integer that is sufficiently large as a function of ε . Let G be a graph on n vertices with $e(G) \geq cn^{1+\varepsilon}$. Then G contains a K -almost-regular subgraph G' on $m \geq n^{\frac{\varepsilon-\varepsilon^2}{2+2\varepsilon}}$ vertices such that $e(G') \geq \frac{2c}{5}m^{1+\varepsilon}$ and $K = 20 \cdot 2^{\frac{1}{\varepsilon^2}+1}$.*

As discussed in the proof outline, the first step of our argument is to find many $2 \times t$ grids with bounds on the codegree of certain pairs. In order to do this, we would like to assume that for each edge uv , there are $\Omega(\bar{d}(G))$ neighbours w of u such that $d(v, w) \geq C$, where C is some large constant. If C is much larger than t , then this allows us to extend a $2 \times s$ grid to a $2 \times (s + 1)$ grid in $\Omega(\bar{d}(G) \cdot C)$ many ways. The following lemma allows us to find a large subgraph with the property mentioned above.

Lemma 2.2. *Let G be an n -vertex graph with $\alpha n^{3/2}$ edges. Then G has a subgraph H on the same vertex set such that $e(H) \geq \frac{1}{2}\alpha n^{3/2}$ and for any $uv \in E(H)$, u has at least $\frac{1}{8}\alpha n^{1/2}$ neighbours w in H such that $d_H(v, w) \geq \alpha^2/32$.*

Proof. Let us define a sequence $G_0 \supset G_1 \supset G_2 \supset \dots$ of graphs as follows. We set $G_0 = G$. Having defined G_i , if there is a vertex $u \in V(G_i)$ with $1 \leq d_{G_i}(u) \leq \frac{1}{4}\alpha n^{1/2}$, then choose such a vertex and let G_{i+1} be the graph obtained from G_i by deleting all edges incident to u . Call this deletion type 1. If no such vertex exists but there is an edge $uv \in E(G_i)$ for which u has less than $\frac{1}{8}\alpha n^{1/2}$ neighbours

w in G_i satisfying $d_{G_i}(v, w) \geq \alpha^2/32$, then let $G_{i+1} = G_i - uv$. Call this deletion type 2. If no such vertex or edge exists, then set $H = G_i$ and terminate the process.

Now for any $uv \in E(H)$, it follows immediately from the definition that u has at least $\frac{1}{8}\alpha n^{1/2}$ neighbours w in H such that $d_H(v, w) \geq \alpha^2/32$. It remains to check that $e(H) \geq \frac{1}{2}\alpha n^{3/2}$, or equivalently that $e(G) - e(H) \leq \frac{1}{2}\alpha n^{3/2}$.

We shall prove that at most $\frac{1}{4}\alpha n^{3/2}$ edges are deleted in type 1 deletions and at most $\frac{1}{4}\alpha n^{3/2}$ edges are deleted in type 2 deletions. Indeed, there are at most n type 1 deletion steps and each of them removes at most $\frac{1}{4}\alpha n^{1/2}$ edges, so it is clear that at most $\frac{1}{4}\alpha n^{3/2}$ edges are removed during type 1 deletions.

Since in each type 2 deletion, we remove precisely one edge, it suffices to prove that there are at most $\frac{1}{4}\alpha n^{3/2}$ such deletions throughout the process. Assume that the edge uv gets deleted from G_i because u has less than $\frac{1}{8}\alpha n^{1/2}$ neighbours w in G_i such that $d_{G_i}(v, w) \geq \alpha^2/32$. Since no deletion of type 1 was applied to G_i , but there is an edge in G_i incident to u , we have $d_{G_i}(u) > \frac{1}{4}\alpha n^{1/2}$. It follows that u has more than $\frac{1}{8}\alpha n^{1/2}$ neighbours w in G_i such that $d_{G_i}(v, w) < \alpha^2/32$. For each such v , u is a common neighbour of v and w , so $d_{G_{i+1}}(v, w) = d_{G_i}(v, w) - 1$. However, the condition $d_{G_i}(v, w) < \alpha^2/32$ shows that any pair (v, w) of vertices can ‘play this role’ in at most $\alpha^2/32$ type 2 deletions. Thus, there are at most $\frac{n^2 \cdot \alpha^2/32}{\frac{1}{8}\alpha n^{1/2}} = \frac{1}{4}\alpha n^{3/2}$ type 2 deletions in the process. \square

Combining the previous two lemmas, we have the following.

Lemma 2.3. *Let G be an n -vertex graph with $\alpha n^{3/2}$ edges, where $\alpha \geq 10$ and n is sufficiently large. Then G has a subgraph H on $m \geq n^{1/12}$ vertices such that for some $\alpha' \geq \alpha/20$, $e(H) \geq \alpha' m^{3/2}$, $\Delta(H) \leq 12000\alpha' m^{1/2}$ and for any $uv \in E(H)$, u has at least $\alpha' m^{1/2}$ neighbours w in H satisfying $d_H(v, w) \geq \alpha'$.*

Proof. We apply Lemma 2.1 to G to obtain a K -almost-regular subgraph H_0 on $m \geq n^{1/12}$ vertices with $e(H_0) = \alpha_0 m^{3/2}$, where $K = 640$ and $\alpha_0 \geq \frac{2}{5}\alpha \geq 4$. Set $\alpha' = \alpha_0/8 \geq \alpha/20$. Applying Lemma 2.2 to H_0 gives us a subgraph $H \subseteq H_0$ such that $e(H) \geq \frac{1}{2}\alpha_0 m^{3/2} \geq \alpha' m^{3/2}$ and for any $uv \in E(H)$, u has at least $\frac{1}{8}\alpha_0 m^{1/2} = \alpha' m^{1/2}$ neighbours w in H such that $d_H(v, w) \geq \alpha_0^2/32 \geq \alpha'$. Finally, $\Delta(H) \leq \Delta(H_0) \leq K\delta(H_0) \leq K\bar{d}(H_0) = 2K\alpha_0 m^{1/2} = 16K\alpha' m^{1/2} < 12000\alpha' m^{1/2}$, as claimed. \square

By Lemma 2.3, in order to prove Theorem 1.2 it suffices to consider host graphs G on n vertices, where n is sufficiently large, which satisfy the following, for a sufficiently large constant α :

- a) $e(G) \geq \alpha n^{3/2}$,
- b) $\Delta(G) \leq K\alpha n^{1/2}$, where $K = 12000$,
- c) for any $uv \in E(G)$, u has at least $\alpha n^{1/2}$ neighbours w in G such that $d_G(v, w) \geq \alpha$.

From now on, we fix a large constant α to be chosen later and let G be a graph on n vertices, which satisfies a)–c).

Let G^k be the graph whose vertex set is $V(G)^k$ and in which (u_1, \dots, u_k) and (v_1, \dots, v_k) are adjacent if $u_j v_j \in E(G)$ for all $1 \leq j \leq k$. This graph is called the k th tensor power of G . We use bold characters to denote vertices of the graph G^k and for $\mathbf{x} \in G^k$, $j \in [k]$, we denote by $\mathbf{x}(j)$ the j th coordinate of \mathbf{x} , which is a vertex of G .

Definition 2.4. In a graph H , a t -ladder is a $2t$ -tuple $(x_1, y_1, \dots, x_t, y_t) \in V(H)^{2t}$ such that for each $1 \leq i \leq t-1$, $x_i x_{i+1}, y_i y_{i+1}, x_i y_i \in E(H)$ and $x_t y_t \in E(H)$.

For real numbers s_1, \dots, s_{t-1} , we call a t -ladder $(\mathbf{x}_1, \mathbf{y}_1, \dots, \mathbf{x}_t, \mathbf{y}_t)$ in $G^k(s_1, s_2, \dots, s_{t-1})$ -good if

- for each $1 \leq i \leq t-1$, $d_{G^k}(\mathbf{x}_{i+1}, \mathbf{y}_i) \leq s_i$,
- for each $1 \leq i \leq t-1$, $1 \leq j \leq k$, $d_G(\mathbf{x}_{i+1}(j), \mathbf{y}_i(j)) \geq \alpha$ and
- for each $1 \leq j \leq k$, the vertices $\mathbf{x}_1(j), \dots, \mathbf{x}_t(j), \mathbf{y}_1(j), \dots, \mathbf{y}_t(j)$ are distinct.

Lemma 2.5. *There exist real numbers $s_1, s_2, \dots, s_{t-1} \geq 1$ for which the number of (s_1, \dots, s_{t-1}) -good t -ladders in G^k is at least*

$$\frac{\alpha^{tk} n^{(t/2+1)k} \prod_{i=1}^{t-1} s_i}{(4^{k+1} \log n^k)^{t-1}}.$$

Proof. We prove the statement by induction on t . Since $e(G) \geq \alpha n^{3/2}$, we have $e(G^k) \geq \alpha^k n^{3k/2}$, so the statement holds for $t = 1$. Assume we have already found real numbers $s_1, \dots, s_{t-2} \geq 1$ and at least

$$\frac{\alpha^{(t-1)k} n^{((t-1)/2+1)k} \prod_{i=1}^{t-2} s_i}{(4^{k+1} \log n^k)^{t-2}}$$

(s_1, \dots, s_{t-2}) -good $(t-1)$ -ladders in G^k . Note that for any such $(t-1)$ -ladder $(\mathbf{x}_1, \mathbf{y}_1, \dots, \mathbf{x}_{t-1}, \mathbf{y}_{t-1})$ and any $1 \leq j \leq k$, we have $\mathbf{x}_{t-1}(j)\mathbf{y}_{t-1}(j) \in E(G)$. Hence, by (c), $\mathbf{x}_{t-1}(j)$ has at least $\alpha n^{1/2}$ neighbours w in G such that $d_G(w, \mathbf{y}_{t-1}(j)) \geq \alpha$. So \mathbf{x}_{t-1} has at least $(\frac{1}{2}\alpha n^{1/2})^k$ neighbours \mathbf{x}_t in G^k such that for all $1 \leq j \leq k$, we have that $d_G(\mathbf{x}_t(j), \mathbf{y}_{t-1}(j)) \geq \alpha$ and $\mathbf{x}_t(j)$ is distinct from $\mathbf{x}_1(j), \dots, \mathbf{x}_{t-1}(j), \mathbf{y}_1(j), \dots, \mathbf{y}_{t-1}(j)$. So there are at least

$$\frac{\alpha^{(t-1)k} n^{((t-1)/2+1)k} \prod_{i=1}^{t-2} s_i}{(4^{k+1} \log n^k)^{t-2}} \cdot \left(\frac{1}{2}\alpha n^{1/2}\right)^k$$

$(2t-1)$ -tuples $(\mathbf{x}_1, \mathbf{y}_1, \mathbf{x}_2, \dots, \mathbf{y}_{t-1}, \mathbf{x}_t)$ in G^k such that $(\mathbf{x}_1, \mathbf{y}_1, \mathbf{x}_2, \dots, \mathbf{x}_{t-1}, \mathbf{y}_{t-1})$ is an (s_1, \dots, s_{t-2}) -good $(t-1)$ -ladder, \mathbf{x}_t is a neighbour of \mathbf{x}_{t-1} in G^k and for all $1 \leq j \leq k$, we have that $\mathbf{x}_t(j)$ is distinct from $\mathbf{x}_1(j), \dots, \mathbf{x}_{t-1}(j), \mathbf{y}_1(j), \dots, \mathbf{y}_{t-1}(j)$ and $d_G(\mathbf{x}_t(j), \mathbf{y}_{t-1}(j)) \geq \alpha$. Since $v(G^k) = n^k$, by dyadic pigeonholing, there is some $s_{t-1} \geq 1$ (in fact, $s_{t-1} \geq \alpha^k$) such that in at least $1/\log n^k$ proportion of these tuples, we have $s_{t-1}/2 < d_{G^k}(\mathbf{x}_t, \mathbf{y}_{t-1}) \leq s_{t-1}$. For each such tuple, the number of ways to choose \mathbf{y}_t to be a common neighbour of \mathbf{x}_t and \mathbf{y}_{t-1} in G^k is precisely $\prod_{j=1}^k d_G(\mathbf{x}_t(j), \mathbf{y}_{t-1}(j))$, and (since $d_G(\mathbf{x}_t(j), \mathbf{y}_{t-1}(j)) \geq \alpha \geq 4t$), at least $\prod_{j=1}^k \frac{1}{2}d_G(\mathbf{x}_t(j), \mathbf{y}_{t-1}(j))$ of these choices satisfy that $\mathbf{y}_t(j)$ is distinct from $\mathbf{x}_1(j), \mathbf{y}_1(j), \mathbf{x}_2(j), \mathbf{y}_2(j), \dots, \mathbf{x}_t(j)$. We have $\prod_{j=1}^k \frac{1}{2}d_G(\mathbf{x}_t(j), \mathbf{y}_{t-1}(j)) = 2^{-k}d_{G^k}(\mathbf{x}_t, \mathbf{y}_{t-1}) \geq 2^{-(k+1)}s_{t-1}$, so there are at least

$$\frac{\alpha^{(t-1)k} n^{((t-1)/2+1)k} \prod_{i=1}^{t-2} s_i}{(4^{k+1} \log n^k)^{t-2}} \cdot \left(\frac{1}{2}\alpha n^{1/2}\right)^k \cdot \frac{s_{t-1}}{2^{k+1} \log n^k} > \frac{\alpha^{tk} n^{(t/2+1)k} \prod_{i=1}^{t-1} s_i}{(4^{k+1} \log n^k)^{t-1}}$$

(s_1, \dots, s_{t-1}) -good t -ladders in G^k . This completes the induction step. □

Recall that in the proof outline we showed that for a fixed path (x_1, \dots, x_t) and a fixed additional vertex u , there are few ways to extend (x_1, \dots, x_t) to (an analogue of) a good ladder containing u . Directly translating this argument to G^k would give us that for any $(t - 1)$ -path $(\mathbf{x}_1, \dots, \mathbf{x}_t)$ in G^k and any $\mathbf{u} \in V(G^k)$, there are few ways to extend $(\mathbf{x}_1, \dots, \mathbf{x}_t)$ to an (s_1, \dots, s_{t-1}) -good t -ladder containing \mathbf{u} . However, and this is where the advantage of working with G^k is revealed, we can draw a similar conclusion even if we count the extensions that contain some fixed vertices of G in some subset $J \subset [k]$ (rather than all) of the coordinates. Naturally, the bound on the number of extensions gets stronger the larger the set J .

Lemma 2.6. *Let $s_1, \dots, s_{t-1} \geq 1, 1 \leq \ell \leq t$, let $J \subset [k]$ and consider fixed vertices $u_j \in V(G)$ for every $j \in J$.*

- (a) *For any fixed $\mathbf{x}_1, \dots, \mathbf{x}_t \in V(G^k)$, the number of (s_1, \dots, s_{t-1}) -good t -ladders in G^k of the form $(\mathbf{x}_1, \mathbf{y}_1, \mathbf{x}_2, \mathbf{y}_2, \dots, \mathbf{x}_t, \mathbf{y}_t)$ with $\mathbf{y}_\ell(j) = u_j$, for all $j \in J$, is at most $\Delta(G)^k \prod_{i=1}^{t-1} s_i \cdot \alpha^{-|J|}$.*
- (b) *For any fixed $\mathbf{y}_1, \dots, \mathbf{y}_t \in V(G^k)$, the number of (s_1, \dots, s_{t-1}) -good t -ladders in G^k of the form $(\mathbf{x}_1, \mathbf{y}_1, \mathbf{x}_2, \mathbf{y}_2, \dots, \mathbf{x}_t, \mathbf{y}_t)$ with $\mathbf{x}_\ell(j) = u_j$, for all $j \in J$, is at most $\Delta(G)^k \prod_{i=1}^{t-1} s_i \cdot \alpha^{-|J|}$.*

Proof. We first prove (a). Let us first assume that $\ell = 1$. Note that for every $j \in [k]$, $\mathbf{y}_1(j)$ is a neighbour of $\mathbf{x}_1(j)$ in G and for each $j \in J$, $\mathbf{y}_1(j) = u_j$. Hence, there are at most $\Delta(G)^{k-|J|}$ possibilities for \mathbf{y}_1 . Furthermore, for each $2 \leq i \leq t$, \mathbf{y}_i is a common neighbour of \mathbf{x}_i and \mathbf{y}_{i-1} in G^k . Since in every (s_1, \dots, s_{t-1}) -good t -ladder, $d_{G^k}(\mathbf{x}_i, \mathbf{y}_{i-1}) \leq s_{i-1}$, there are at most s_{i-1} possibilities for \mathbf{y}_i given $\mathbf{y}_1, \dots, \mathbf{y}_{i-1}$. Altogether we get at most $\Delta(G)^{k-|J|} \prod_{i=1}^{t-1} s_i \leq \Delta(G)^k \prod_{i=1}^{t-1} s_i \cdot \alpha^{-|J|}$ possibilities, where we used $\Delta(G) \geq \alpha$. This completes the proof of (a) in the case $\ell = 1$.

Assume now $\ell \geq 2$. Since \mathbf{y}_1 is a neighbour of \mathbf{x}_1 , there are at most $\Delta(G^k) = \Delta(G)^k$ choices for it. For any $2 \leq i \leq \ell - 1$, \mathbf{y}_i is a common neighbour of \mathbf{x}_i and \mathbf{y}_{i-1} in G^k , so given $\mathbf{y}_1, \dots, \mathbf{y}_{i-1}$, there are at most s_{i-1} choices for it because we are only counting (s_1, \dots, s_{t-1}) -good t -ladders in G^k . Given $\mathbf{y}_1, \dots, \mathbf{y}_{\ell-1}$, there are at most $\prod_{j \in [k] \setminus J} d_G(\mathbf{x}_\ell(j), \mathbf{y}_{\ell-1}(j))$ choices for \mathbf{y}_ℓ since $\mathbf{y}_\ell(j)$ is given for each $j \in J$. Note that

$$\prod_{j \in [k] \setminus J} d_G(\mathbf{x}_\ell(j), \mathbf{y}_{\ell-1}(j)) = \frac{\prod_{j \in [k]} d_G(\mathbf{x}_\ell(j), \mathbf{y}_{\ell-1}(j))}{\prod_{j \in J} d_G(\mathbf{x}_\ell(j), \mathbf{y}_{\ell-1}(j))} = \frac{d_{G^k}(\mathbf{x}_\ell, \mathbf{y}_{\ell-1})}{\prod_{j \in J} d_G(\mathbf{x}_\ell(j), \mathbf{y}_{\ell-1}(j))} \leq \frac{s_{\ell-1}}{\alpha^{|J|}},$$

where the last inequality holds because we are counting (s_1, \dots, s_{t-1}) -good t -ladders in G^k .

Finally, for any $\ell + 1 \leq i \leq t$, there are at most s_{i-1} possibilities for \mathbf{y}_i , given $\mathbf{y}_1, \dots, \mathbf{y}_{i-1}$, so (a) is proved.

To prove part (b), observe the following symmetry. If $(\mathbf{x}_1, \mathbf{y}_1, \mathbf{x}_2, \mathbf{y}_2, \dots, \mathbf{x}_t, \mathbf{y}_t)$ is an (s_1, \dots, s_{t-1}) -good t -ladder in G^k , then $(\mathbf{y}_t, \mathbf{x}_t, \mathbf{y}_{t-1}, \mathbf{x}_{t-1}, \dots, \mathbf{y}_1, \mathbf{x}_1)$ is an $(s_{t-1}, s_{t-2}, \dots, s_1)$ -good t -ladder in G^k . Now, (b) follows from (a). □

We are now ready to prove Theorem 1.2. We build an auxiliary graph whose vertices are the copies of P_t in G^k and the edges correspond to (s_1, \dots, s_{t-1}) -good t -ladders, for an appropriate choice of s_1, \dots, s_{t-1} . An embedding of T into this graph gives k homomorphic copies of $T \square P_t$ in G , one for each coordinate of G^k . Using the previous lemma allows us to greedily find an embedding such that at least one of these copies has distinct vertices, that is, the vertices of this coordinate form a genuine copy of $T \square P_t$ in G . We make these arguments precise below.

Proof of Theorem 1.2. Let T be a tree on $r \geq 2$ vertices and let $t \geq 2$. Recall that it is enough to prove that $T \square P_t$ is a subgraph of a given graph G satisfying (a)–(c). Let $\alpha = (16K)^{r^2 t^3}$ and let $k = k(t, r, n)$ be large enough so that $2^k > k \cdot 4t^2 r \log n$. Let s_1, \dots, s_{t-1} be real numbers given by Lemma 2.5 and let \mathcal{F} be the family of (s_1, \dots, s_{t-1}) -good t -ladders in G^k . We define a graph \mathcal{G} as follows. The vertex-set of \mathcal{G} is the set of t -tuples $(\mathbf{x}_1, \dots, \mathbf{x}_t) \in V(G^k)^t$ forming a path in G^k such that for all $j \in [k]$, the vertices $\mathbf{x}_1(j), \dots, \mathbf{x}_t(j)$ are distinct. We let $(\mathbf{x}_1, \dots, \mathbf{x}_t)$ and $(\mathbf{y}_1, \dots, \mathbf{y}_t)$ form an edge in \mathcal{G} if $(\mathbf{x}_1, \mathbf{y}_1, \mathbf{x}_2, \mathbf{y}_2, \dots, \mathbf{x}_t, \mathbf{y}_t)$ or $(\mathbf{y}_1, \mathbf{x}_1, \mathbf{y}_2, \mathbf{x}_2, \dots, \mathbf{y}_t, \mathbf{x}_t)$ is in \mathcal{F} .

The vertices of \mathcal{G} correspond to paths of length $t - 1$ in G^k , so $|V(\mathcal{G})| \leq |V(G^k)| \Delta(G^k)^{t-1} = |V(G)|^k \Delta(G)^{k(t-1)} \leq n^k (K\alpha n^{1/2})^{k(t-1)}$. Also, any element of \mathcal{F} gives an edge in \mathcal{G} and at most two elements of \mathcal{F} give the same edge, so

$$|E(\mathcal{G})| \geq |\mathcal{F}|/2 \geq \frac{\alpha^{tk} n^{(t/2+1)k} \prod_{i=1}^{t-1} s_i}{2(4^{k+1} \log n^k)^{t-1}}.$$

By successively removing vertices of degree less than $\bar{d}(\mathcal{G})/2$ from \mathcal{G} , we obtain a non-empty subgraph \mathcal{H} of \mathcal{G} with

$$\delta(\mathcal{H}) \geq \frac{1}{2} \bar{d}(\mathcal{G}) = \frac{|E(\mathcal{G})|}{|V(\mathcal{G})|} \geq \frac{(\alpha n^{1/2})^k \prod_{i=1}^{t-1} s_i}{2(4^{k+1} K^k \log n^k)^{t-1}}. \tag{1}$$

Let w_1, w_2, \dots, w_r be a 1-degenerate ordering of the vertices of T , that is, an ordering such that for all $2 \leq p \leq r$, there is a unique $z < p$ with $w_p w_z \in E(T)$.

Claim. There exists an embedding of T into \mathcal{H} which satisfies the following. Writing $(\mathbf{x}_1^p, \mathbf{x}_2^p, \dots, \mathbf{x}_t^p)$ for the vertex of \mathcal{H} into which w_p is embedded, for every $1 \leq q < p \leq r$ and $1 \leq i, \ell \leq t$, there are at most $\frac{k}{\binom{r}{t}^2}$ values of $j \in [k]$ such that $\mathbf{x}_\ell^p(j) = \mathbf{x}_i^q(j)$.

Assume we are given the embedding as in the claim and let us show how it implies the theorem. For every $1 \leq p \leq r$ and every $j \in [k]$, since $(\mathbf{x}_1^p, \mathbf{x}_2^p, \dots, \mathbf{x}_t^p) \in V(\mathcal{H})$, it follows that the vertices $\mathbf{x}_1^p(j), \mathbf{x}_2^p(j), \dots, \mathbf{x}_t^p(j)$ form a path in G . Similarly, whenever $w_p w_z \in E(T)$, the vertices $(\mathbf{x}_1^p, \mathbf{x}_2^p, \dots, \mathbf{x}_t^p)$ and $(\mathbf{x}_1^z, \mathbf{x}_2^z, \dots, \mathbf{x}_t^z)$ form an edge in \mathcal{H} , implying that $(\mathbf{x}_i^p(j), \mathbf{x}_i^z(j)) \in E(G)$ for all $i \in [t], j \in [k]$. Hence, if for some $j \in [k]$, the vertices $\mathbf{x}_i^p(j), 1 \leq p \leq r, 1 \leq i \leq t$, are all distinct, then they form a copy of $T \square P_t$ in G . The existence of such an index j follows by simple counting. Indeed, summing over all $1 \leq p < q \leq r, 1 \leq i, \ell \leq t$, there are at most $\binom{r}{2} t^2 \frac{k}{\binom{r}{t}^2} < k$ indices j for which the vertices $\mathbf{x}_i^p(j), 1 \leq p \leq r, 1 \leq i \leq t$ are not all distinct.

Proof of Claim. We iteratively find the desired embedding. We map w_1 to an arbitrary vertex $(\mathbf{x}_1^1, \dots, \mathbf{x}_t^1) \in \mathcal{H}$. Now, fix $2 \leq p \leq r$ and suppose that for each $1 \leq q \leq p - 1$, $(\mathbf{x}_1^q, \dots, \mathbf{x}_t^q)$ has already been found. Let $1 \leq z < p$ be the unique index such that $w_p w_z \in E(T)$. We choose $(\mathbf{x}_1^p, \dots, \mathbf{x}_t^p)$ to be a suitable neighbour of $(\mathbf{x}_1^z, \dots, \mathbf{x}_t^z)$ in \mathcal{H} . For any $1 \leq q \leq p - 1, 1 \leq i, \ell \leq t$ and $J \subset [k]$ of size $\lceil k/(rt)^2 \rceil$, Lemma 2.6 shows that $(\mathbf{x}_1^z, \dots, \mathbf{x}_t^z)$ has at most $\frac{2\Delta(G)^k \prod_{i=1}^{t-1} s_i}{\alpha^{\lceil k/(rt)^2 \rceil}}$ neighbours $(\mathbf{x}_1^p, \dots, \mathbf{x}_t^p)$ in \mathcal{H} with $\mathbf{x}_\ell^p(j) = \mathbf{x}_i^q(j)$ for all $j \in J$. Summing over all choices for i, ℓ, q and J , we

find that the number of neighbours of $(\mathbf{x}_1^z, \dots, \mathbf{x}_t^z)$ that are not suitable is at most

$$\begin{aligned} t^2 r 2^k \cdot \frac{2\Delta(G)^k \prod_{i=1}^{t-1} s_i}{\alpha^{\lceil k/(rt)^2 \rceil}} &\leq 2t^2 r \frac{(2K\alpha n^{1/2})^k \prod_{i=1}^{t-1} s_i}{\alpha^{\lceil k/(rt)^2 \rceil}} \leq 2t^2 r \frac{(\alpha n^{1/2})^k \prod_{i=1}^{t-1} s_i}{(16K)^{k(t-1)}} \\ &\leq \frac{4t^2 r (\log n^k)^{t-1}}{2^{k(t-1)}} \cdot \frac{(\alpha n^{1/2})^k \prod_{i=1}^{t-1} s_i}{2(4^{k+1} K^k \log n^k)^{t-1}} < \delta(H), \end{aligned}$$

where in the second inequality we used that $\alpha^{\lceil k/(rt)^2 \rceil} \geq (16K)^{kt}$ and in the last inequality we used $2^k > 4t^2 r k \log n$ as well as (1). It follows that $(\mathbf{x}_1^z, \dots, \mathbf{x}_t^z)$ has a suitable neighbour in \mathcal{H} , proving the claim. \square

3 | CONCLUDING REMARKS

We conjecture the following generalization of Theorem 1.2.

Conjecture 3.1. *For any two trees T and S (each with at least one edge), there exist positive real numbers c and C such that*

$$cn^{3/2} \leq \text{ex}(n, T \square S) \leq Cn^{3/2}.$$

Note that if T and S are trees, then $T \square S$ is 2-degenerate. This observation provides some evidence towards Conjecture 3.1 since the conjecture of Erdős [7], mentioned in the introduction, asserts that if H is an r -degenerate bipartite graph, then $\text{ex}(n, H) = O(n^{2-1/r})$. Let us briefly discuss why our methods fall short of proving Conjecture 3.1. We can define $K_2 \square S$ to be an S -ladder and prove an analogue of Lemma 2.5. That is, we can show that there are many S -ladders in which for every vertex of S apart from the root, the codegree of the vertex in the first copy of S and its parent in the other copy of S is small. While the analogue of Lemma 2.6 (a) holds, part (b) does not. This is because, informally speaking, we cannot reverse the direction in which a tree is built unless it is a path. However, as pointed out by a referee, our proof can easily be adapted to prove Conjecture 3.1 in the special case when one of the trees is a star since in that case the analogue of part (b) of Lemma 2.6 is not needed.

It would also be very interesting to study grids of higher dimension. For positive integers d and t , the d -dimensional grid $F_t^{(d)}$ has vertex set $[t]^d$ with two vertices joined by an edge if they differ in exactly one coordinate and in that coordinate they differ by exactly one. Since $F_t^{(d)}$ is d -degenerate, the following conjecture is natural.

Conjecture 3.2. *For any positive integers d and t , there is a constant C such that*

$$\text{ex}(n, F_t^{(d)}) \leq Cn^{2-1/d}.$$

It is likely that this bound is not tight and can be improved. On the other hand, there is the probabilistic lower bound $\text{ex}(n, F_t^{(d)}) = \Omega(n^{2 - \frac{t^d - 2}{(t^d - t^{d-1})^{d-1}}})$, which is close to $n^{2-1/d}$ when t is large.

For the smallest non-trivial case of the cube, it is known that $\text{ex}(n, F_2^{(3)}) = O(n^{8/5})$, a stronger bound than the one asserted by Conjecture 3.2.

It would also be interesting to determine the correct dependence of $\text{ex}(n, F_t)$ on t . Our proof implies that the Turán number of F_t is at most $e^{O(t^5)}n^{3/2}$. On the other hand, the following construction shows that $\text{ex}(n, F_t) \geq ct^{1/2}n^{3/2}$. Let G_0 be a C_4 -free graph with $\frac{n}{t-1}$ vertices and about $(\frac{n}{t-1})^{3/2}$ edges. Let G be the $(t-1)$ -blowup of G_0 . Then G is an n -vertex graph and $e(G) \approx (t-1)^2(\frac{n}{t-1})^{3/2} \geq ct^{1/2}n^{3/2}$. We claim that G does not contain a $t \times t$ grid. Indeed, since G_0 is C_4 -free, any 4-cycle in G must have an opposite pair of vertices which come from the same vertex of G_0 . It is known (see, for example, [19]) that if a diagonal is placed in each unit square of a $t \times t$ grid, then there is a path along these diagonals from one side of the grid to the opposite side (that is, either from top to bottom or from left to right). Note that this path contains at least t vertices of the grid. Hence, a $t \times t$ grid in G would have to contain at least t vertices which come from the same vertex in G_0 , showing that such a grid cannot be found in G .

Finally, we remark that we can apply the tensor power trick similarly as it is done in this paper to another extremal problem (that is, working directly with the k -th power rather than with the original graph/hypergraph). This problem asks to determine the asymptotics of the Turán number of the r -uniform hypergraph $K_{2,t}^{(r)}$ whose vertex set consists of disjoint sets X, Y_1, \dots, Y_t , where $|X| = 2$ and $|Y_1| = \dots = |Y_t| = r-1$, and whose edge set is $\{\{x\} \cup Y_i : x \in X, 1 \leq i \leq t\}$. However, for that problem, there is a different proof which does not require tensorization, see [4] for more details. It would be interesting to find other instances where the tensor power technique can be applied in a similar way.

ACKNOWLEDGEMENT

We are grateful to the referees for their careful review and valuable suggestions. Research of Domagoj Bradač, Benny Sudakov and István Tomon is supported in part by SNSF grant 200021_196965. Research of Oliver Janzer is supported by an ETH Zürich Postdoctoral Fellowship 20-1 FEL-35.

Open access funding provided by Eidgenössische Technische Hochschule Zurich.

JOURNAL INFORMATION

The *Bulletin of the London Mathematical Society* is wholly owned and managed by the London Mathematical Society, a not-for-profit Charity registered with the UK Charity Commission. All surplus income from its publishing programme is used to support mathematicians and mathematics research in the form of research grants, conference grants, prizes, initiatives for early career researchers and the promotion of mathematics.

REFERENCES

1. N. Alon, M. Krivelevich, and B. Sudakov, *Turán numbers of bipartite graphs and related Ramsey-type questions*, *Comb. Probab. Comput.* **12** (2003), 477–494.
2. N. Alon and I. Z. Ruzsa, *Non-averaging subsets and non-vanishing transversals*, *J. Combin. Theory Ser. A* **86** (1999), no. 1, 1–13.
3. G. R. Blakley and P. Roy, *A Hölder type inequality for symmetric matrices with nonnegative entries*, *Proc. Am. Math. Soc.* **16** (1965), no. 6, 1244–1245.
4. D. Bradač, L. Gishboliner, O. Janzer, and B. Sudakov, *Asymptotics of the hypergraph bipartite Turán problem*, preprint, 2022.
5. D. Clemens, M. Miralaei, D. Reding, M. Schacht, and A. Taraz, *On the size-Ramsey number of grid graphs*, *Comb. Probab. Comput.* **30** (2021), no. 5, 670–685.

6. D. Conlon, R. Nenadov, and M. Trujić, *On the size-Ramsey number of grids*, arXiv preprint arXiv:2202.01654, 2022.
7. P. Erdős, *Some recent results on extremal problems in graph theory (Results)*, Theory of Graphs (Internat. Sympos., Rome, 1966), 1967, pp. 117–123.
8. P. Erdős and M. Simonovits, *A limit theorem in graph theory*, Studia Sci. Math. Hungar. **1** (1966), 51–57.
9. P. Erdős and M. Simonovits, *Some extremal problems in graph theory*, Combinatorial theory and its applications, I (Proc. Colloq., Balatonfüred, 1969), 1970, pp. 377–390.
10. P. Erdős and A. H. Stone, *On the structure of linear graphs*, Bull. Amer. Math. Soc. **52** (1946), no. 12, 1087–1091.
11. J. Fox and B. Sudakov, *Dependent random choice*, Random Structures Algorithms **38** (2011), 1–32.
12. Z. Füredi, *On a Turán type problem of Erdős*, Combinatorica **11** (1991), no. 1, 75–79.
13. Z. Füredi and M. Ruszinkó, *Uniform hypergraphs containing no grids*, Adv. Mathematics **240** (2013), 302–324.
14. Z. Füredi and M. Simonovits, *The history of degenerate (bipartite) extremal graph problems*, Erdős centennial, volume 25 of Bolyai Soc. Math. Stud., János Bolyai Math. Soc., Budapest, 2013, pp. 169–264.
15. L. Gishboliner and A. Shapira, *Constructing dense grid-free linear 3-graphs*, Proc. Amer. Math. Soc. **150** (2022), no. 01, 69–74.
16. T. Jiang and R. Seiver, *Turán numbers of subdivided graphs*, SIAM J. Discret. Math. **26** (2012), 1238–1255.
17. J. H. Kim, C. Lee, and J. Lee, *Two approaches to Sidorenko’s conjecture*, Trans. Amer. Math. Soc. **368** (2016), no. 7, 5057–5074.
18. T. Kővári, V. Sós, and P. Turán, *On a problem of K. Zarankiewicz*, Colloquium Math. **3** (1954), 50–57.
19. *Mathoverflow thread: Sperner’s lemma and paths from one side to the opposite one in a grid*, <https://mathoverflow.net/questions/112067/sperners-lemma-and-paths-from-one-side-to-the-opposite-one-in-a-grid/359066#359066>.
20. T. Tao, *Tricks wiki article: The tensor power trick*, <https://terrytao.wordpress.com/2008/08/25/tricks-wiki-article-the-tensor-product-trick/>.