Sidorenko's conjecture for blow-ups

David Conlon^{*} Joonkyung Lee[†]

Abstract

A celebrated conjecture of Sidorenko and Erdős–Simonovits states that, for all bipartite graphs H, quasirandom graphs contain asymptotically the minimum number of copies of H taken over all graphs with the same order and edge density. This conjecture has attracted considerable interest over the last decade and is now known to hold for a broad range of bipartite graphs, with the overall trend saying that a graph satisfies the conjecture if it can be built from simple building blocks such as trees in a certain recursive fashion.

Our contribution here, which goes beyond this paradigm, is to show that the conjecture holds for any bipartite graph H with bipartition $A \cup B$ where the number of vertices in B of degree ksatisfies a certain divisibility condition for each k. As a corollary, we have that for every bipartite graph H with bipartition $A \cup B$, there is a positive integer p such that the blow-up H_A^p formed by taking p vertex-disjoint copies of H and gluing all copies of A along corresponding vertices satisfies the conjecture. Another way of viewing this latter result is that for every bipartite Hthere is a positive integer p such that an L^p -version of Sidorenko's conjecture holds for H.

1 Introduction

One of the central problems in extremal graph theory is to estimate the minimum number of copies of a graph H which can be contained in another graph G of given order and edge density. Even when H is a triangle, this problem is highly non-trivial and was only solved fully by Razborov [12, 13] in 2008, who used it as the first test case for his influential flag algebra technique. His result was then extended to K_4 by Nikiforov [11] and to all K_r by Reiher [14] using further ideas.

Part of the difficulty in proving these results is that the behaviour of the minimum number of copies of K_r as a function of the edge density is surprisingly complicated. On the other hand, when H is a bipartite graph, conjectures of Erdős and Simonovits [5] and Sidorenko [15] suggest that the minimum should be extremely simple, being asymptotically equal to the number of copies of H in a quasirandom graph of the same density.

This attractive conjecture, usually known as Sidorenko's conjecture, is best stated in terms of homomorphisms. A homomorphism from a graph H to a graph G is a mapping $f: V(H) \to V(G)$ such that (f(u), f(v)) is an edge of G whenever (u, v) is an edge of H. If $h_H(G)$ is the number of homomorphisms from H to G, we write $t_H(G) = h_H(G)/|G|^{|H|}$ for the homomorphism density, the probability that a uniform random mapping from V(H) to V(G) is a homomorphism. The conjecture is then as follows.

^{*}Department of Mathematics, California Institute of Technology, Pasadena, CA 91125, USA. E-mail: dconlon@caltech.edu. Research supported by a Royal Society University Research Fellowship and by ERC Starting Grant 676632.

[†]Department of Mathematics, University College London, Gower Street, London WC1E 6BT, UK. E-mail: joonkyung.lee@ucl.ac.uk. Research supported by ERC Consolidator Grant PEPCo 724903 and ERC Starting Grant 676632.

Sidorenko's conjecture. For any bipartite graph H and any graph G,

$$t_H(G) \ge t_{K_2}(G)^{e(H)}.$$

Sidorenko [15] himself showed that the conjecture holds for some simple classes of bipartite graph, namely, complete bipartite graphs, even cycles and trees, and for bipartite graphs with at most four vertices on one side. There the matter stood for some time until work of Hatami [7], connecting it with a question of Lovász [10] about which graphs define norms, revived interest in the conjecture. In particular, he showed that hypercubes have a certain weak norming property and, hence, that they satisfy Sidorenko's conjecture.

The first significant breakthrough on the conjecture was made by Conlon, Fox and Sudakov [1], who used the dependent random choice technique [6] to show that if H is a bipartite graph with a vertex which is complete to the other side, then Sidorenko's conjecture holds for H. As a corollary, they showed that this implies an approximate version of the conjecture. An important further advance was then made by Li and Szegedy [9], who initiated the application of entropy methods to the conjecture, at first in the guise of logarithmic convexity inequalities. In particular, they found a remarkably concise proof of the result of Conlon, Fox and Sudakov and extended this result to a more general class, which they referred to as reflection trees.

These ideas were developed further by Kim, Lee and Lee [8], who proved the conjecture for what they called tree-arrangeable graphs, and then pushed to their (seemingly) natural conclusion by Conlon, Kim, Lee and Lee [2, 3] and, independently, by Szegedy [16]. These works give broad classes of graphs for which Sidorenko's conjecture holds, though it is somewhat hard to do justice to these classes in this limited space. However, the overall trend is that a graph may be shown to satisfy the conjecture if it can be built from simple building blocks such as trees (or weakly norming graphs [4]) in a certain recursive fashion. The main result of this paper is the following, which we believe moves beyond the confines of this paradigm.

Theorem 1.1. Let H be a bipartite graph with bipartition $A \cup B$, $\max_{b \in B} \deg(b) = r$ and, for each $1 \leq k \leq r$, let d_k be the number of vertices with degree k in B. Then, if $\binom{|A|}{r}\binom{r}{k}$ divides d_k for each $1 \leq k \leq r$, Sidorenko's conjecture holds for H.

In the proof, critical use is made of a simple mechanism that we call the Hölder trick, which allows us to convert a graph of the type described in Theorem 1.1 into a simpler graph to which we can apply the existing techniques. This trick was first observed in [2, Section 3], but was not exploited to its full potential. To illustrate this key idea, we will discuss the notorious example $K_{5,5} \setminus C_{10}$ in the next section, showing that its 'square' satisfies the conjecture.

When the bipartite graph H is regular on one side of the bipartition, a stronger statement, without any divisibility condition, can be obtained which already implies the aforementioned result about $K_{5,5} \setminus C_{10}$.

Theorem 1.2. Let H be a bipartite graph with partition $A \cup B$ and $\deg(b) = r$ for all $b \in B$. Then, provided $|B| \ge \binom{|A|}{r}$, Sidorenko's conjecture holds for H.

Given a bipartite graph H with bipartition $A \cup B$ and a positive integer p, its blow-up H_A^p , or 'p-th power', relative to A is defined to be the graph formed by taking p vertex-disjoint copies of H and gluing all copies of A along corresponding vertices. That is, we replace each vertex in B with an independent set of order p and connect every vertex in A that was joined to $b \in B$ to every vertex in the corresponding independent set. Since $\binom{|A|}{r}\binom{r}{k}$ divides |A|!/(|A|-r)! for each k, a simple corollary of Theorem 1.1, again generalising the $K_{5,5} \setminus C_{10}$ case, is then as follows.

Corollary 1.3. For every bipartite graph H with bipartition $A \cup B$, there is a positive integer p such that H_A^p satisfies Sidorenko's conjecture. In particular, p = |A|! always suffices.

This can be viewed as saying that for any bipartite graph H there is an integer p such that an L^p -version of the conjecture holds for H. To see this, suppose that |A| = m and identify the set A with $[m] = \{1, 2, \ldots, m\}$. Now, writing $x_A = (x_1, \ldots, x_m)$, where $x_i \in V(G)$ for all $i = 1, 2, \ldots, m$, consider the function $t_H(G; x_A)$ which counts the proportion of mappings f from V(H) to V(G) with $f(i) = x_i$ for all $i = 1, 2, \ldots, m$ which are homomorphisms. Sidorenko's conjecture for H is clearly equivalent to the statement that $\mathbb{E}_{x_A} t_H(G; x_A) \geq t_{K_2}(G)^{e(H)}$, whereas Corollary 1.3 says that for any H there is a positive integer p such that $\mathbb{E}_{x_A} t_H(G; x_A)^p \geq t_{K_2}(G)^{p \cdot e(H)}$.

Another interesting corollary of Theorem 1.1 is as follows.

Corollary 1.4. For any bipartite graph H, there is another bipartite graph H' such that Sidorenko's conjecture holds for the disjoint union of H and H'.

To see this, suppose that H has bipartition $A \cup B$ with |A| = m, $\max_{b \in B} \deg(b) = r$ and, for each $1 \leq k \leq r$, d_k is the number of vertices of degree k in B. Let H' be a bipartite graph between A' and B', where |A'| = r and B' has exactly $\lceil d_k/(m+r)! \rceil (m+r)! - d_k$ vertices with degree k for each $1 \leq k \leq r$. It is then easy to check that the conditions of Theorem 1.1 are satisfied for $H \cup H'$.

For convenience of notation, we will use the language of graphons throughout the paper. A graphon is a symmetric measurable function W from $[0,1]^2$ to [0,1], where symmetric in this context means that W(x,y) = W(y,x) for all $(x,y) \in [0,1]^2$. Very roughly, this may be seen as a continuous analogue of the adjacency matrix of a graph. The homomorphism density $t_H(W)$ of a graph H in a graphon W is then given by

$$t_H(W) = \mathbb{E}\left[\prod_{ij\in E(H)} W(x_i, x_j)\right] = \int_{[0,1]^{v(H)}} \prod_{ij\in E(H)} W(x_i, x_j) \ d\mu^{v(H)},$$

where μ is the Lebesgue measure on [0, 1]. Note that we will typically abbreviate integrals with expectations, as above. In this language, Sidorenko's conjecture for a given H is equivalent to saying that

$$t_H(W) \ge t_{K_2}(W)^{e(H)}$$

for every graphon W. It is this statement that we will prove in the cases of interest.

2 A motivating example

We now take a closer look at the graph $M := K_{5,5} \setminus C_{10}$, the simplest graph for which Sidorenko's conjecture is not known, showing that its 'square' does satisfy the conjecture. Since M is vertex-transitive, we don't need to distinguish which part of the bipartition is glued and we can simply write M^2 rather than M_A^2 .

The main result of this short section relates the homomorphism density of M^2 to the homomorphism density of another graph. To define this graph, let \mathcal{H} be a family of subsets of a finite set A. The (A, \mathcal{H}) -incidence graph is the bipartite graph on $A \cup \mathcal{H}$ such that $a \in A$ and $F \in \mathcal{H}$ are adjacent if and only if $a \in F$. For $r \leq m$, the (m, r)-incidence graph is then the (A, \mathcal{H}) -incidence graph where A = [m] and $\mathcal{H} = {[m] \choose r}$.

Theorem 2.1. Let $F_{5,3}$ be the (5,3)-incidence graph. Then, for every graphon W,

$$t_{M^2}(W) \ge t_{F_{5,3}}(W).$$

By a result of the authors [4] (discussed in Appendix A), $F_{5,3}$ is a weakly norming graph, which in turn implies that it satisfies Sidorenko's conjecture, i.e., for any graphon W, $t_{F_{5,3}}(W) \ge t_{K_2}(W)^{30}$. Together with Theorem 2.1, this implies that M^2 also satisfies the conjecture.

For the proof of Theorem 2.1, we introduce some notation that we will use throughout the paper. Given a vector (x_1, \ldots, x_m) and $I = \{i_1, \ldots, i_s\} \subseteq [m]$ with $i_1 < \cdots < i_s$, let x_I be the vector $(x_{i_1}, \ldots, x_{i_s})$. Let S_m be the set of all permutations of [m] and let \mathcal{F} be a family of functions $\{f_I : I \subseteq [m]\}$, where each f_I is a measurable function from $[0, 1]^I$ to [0, 1]. Then, for each $f_I \in \mathcal{F}$, let

$$\tilde{f}_I(x_I) := \left(\prod_{\phi \in S_m} f_{\phi(I)}(x_I)\right)^{1/m}$$

We will need Hölder's inequality in the following form:

$$\mathbb{E}\left[\prod_{I\subseteq[m]}\tilde{f}_I(x_I)\right] \le \prod_{\phi\in S_m} \mathbb{E}\left[\prod_{I\subseteq[m]}f_{\phi(I)}(x_I)\right]^{1/m!}.$$
(1)

Finally, suppose H is a bipartite graph with bipartition $A \cup B$. For a subset F of A and a graphon W, let $\rho(x_F) := \mathbb{E}_y \prod_{i \in F} W(x_i, y)$. Then $t_H(W)$ can be rephrased as

$$t_H(W) = \mathbb{E}\left[\prod_{v \in B} \rho(x_{N(v)})\right],$$

where the expectation, both here and in the proof below, is over x_A .

Proof of Theorem 2.1. Note that M is isomorphic to the bipartite graph on $Z \cup Z'$ where Z and Z' are two disjoint copies of the group \mathbb{Z}_5 and $i \in Z'$ is adjacent to i - 1, i and i + 1 in Z. Let Z_i be the subset $\{i - 1, i, i + 1\}$ of Z. Then

$$t_{M^2}(W) = \mathbb{E}\left[\prod_{i=1}^5 \rho(x_{Z_i})^2\right].$$

Now let \mathcal{F} be the set $\{f_I : I \subseteq [5]\}$, where $f_I(x_I) = \rho(x_I)^2$ if $I = Z_i$ and $f_I = 1$ otherwise. Then, since Z_1, \ldots, Z_5 constitute half the triples in $\binom{A}{3}$, $\tilde{f}_I(x_I) = \rho(x_I)$ if |I| = 3 and $\tilde{f}_I = 1$ otherwise. Moreover,

$$t_{M^2}(W) = \mathbb{E}\left[\prod_{i=1}^5 \rho(x_{Z_i})^2\right] = \mathbb{E}\left[\prod_{I \subseteq [5]} f_{\phi(I)}(x_I)\right]$$

for each $\phi \in S_5$. Hence, Hölder's inequality (1) implies that

$$t_{F_{5,3}}(W) = \mathbb{E}\left[\prod_{T \in \binom{5}{3}} \rho(x_T)\right] = \mathbb{E}\left[\prod_{I \subseteq [5]} \tilde{f}_I(x_I)\right]$$
$$\leq \prod_{\phi \in S_5} \mathbb{E}\left[\prod_{I \subseteq [5]} f_{\phi(I)}(x_I)\right]^{1/5!} = \prod_{\phi \in S_5} \mathbb{E}\left[\prod_{i=1}^5 \rho(x_{Z_i})^2\right]^{1/5!} = t_{M^2}(W),$$

as required.

3 Weighted homomorphism densities

Let $\alpha = (\alpha_v)_{v \in B}$ be a vector indexed by $v \in B$ with non-negative coordinates. Define the α -weighted homomorphism density of H by

$$t_H^{\alpha}(W) = \mathbb{E}\left[\prod_{v \in B} \rho(x_{N(v)})^{\alpha_v}\right],$$

where the expectation here and below, unless otherwise indicated, is over x_A . In particular, if $\alpha_v = p$ for each $v \in B$, then $t^{\alpha}_H(W) = t_{H^p_A}(W)$. We say that the weight vector $\alpha = (\alpha_v)_{v \in B}$ is symmetric if $\alpha_u = \alpha_v$ whenever deg $(u) = \deg(v)$, i.e., it assigns the same weight to vertices with the same degree. When α is symmetric, we simply write $\alpha = (\alpha_k)_{k=1}^r$, where r is the maximum degree of B and $\alpha_k := \alpha_v$ for all v with deg(v) = k.

For $r \leq m$, define the (m, r)-downset graph to be the (A, \mathcal{H}) -incidence graph with A = [m] and $\mathcal{H} = \binom{A}{\leq r}$. The main result of this section states that, for certain symmetric integer weight vectors α , a 'weighted' version of Sidorenko's conjecture holds for the (m, r)-downset graph. This can also be interpreted as saying that Sidorenko's conjecture holds for certain blow-ups of the (m, r)-downset graph, where vertices in $\binom{[m]}{\leq r}$ of different degrees may be blown up by different amounts.

Theorem 3.1. Suppose $r \leq m$ and let H be the (m,r)-downset graph and $\alpha = (\alpha_k)_{k=1}^r$ be a symmetric integer weight vector such that $\binom{m-k}{r-k}$ divides α_k for each $1 \leq k \leq r$ and $\alpha_r > 0$. Then, for every graphon W,

$$t_H^{\alpha}(W) \ge t_{K_2}(W)^{e_{\alpha}(H)},\tag{2}$$

where $e_{\alpha}(H) := \sum_{v \in B} \alpha_v \deg(v)$.

The proof has three steps. First, we construct a weakly norming (r + 1)-graph \mathcal{H}_{α} and a measurable function $W_{\alpha} : [0, 1]^{r+1} \to [0, 1]$ such that

$$t_H^{\alpha}(W) = t_{\mathcal{H}_{\alpha}}(W_{\alpha}).$$

Then, by using the fact that \mathcal{H}_{α} is weakly norming, we obtain

$$t_{\mathcal{H}_{\alpha}}(W_{\alpha}) \ge t_{\mathcal{G}_{\alpha}}(W_{\alpha})^{e(\mathcal{H}_{\alpha})/e(\mathcal{G}_{\alpha})},$$

where \mathcal{G}_{α} is a subgraph of \mathcal{H}_{α} with 'simpler' structure. We then conclude by showing that

$$t_{\mathcal{G}_{\alpha}}(W_{\alpha}) \ge t_{K_2}(W)^{q_{\alpha,H}},$$

where $q_{\alpha,H}$ is the 'correct' exponent to yield (2).

Throughout this section, let $\beta_k := \alpha_k / {\binom{m-k}{r-k}}$ for $k = 1, 2, \dots, r$ and $\beta := \beta_1 \beta_2 \dots \beta_r$. The divisibility condition in Theorem 3.1 ensures that each β_k is an integer. For an *r*-set *F*, let $U_0(F) := F$ and, for $1 \le k \le r$, let $U_k(F)$ be the union of β_k disjoint copies of ${\binom{F}{k}}$. We say that $(u_0, u_1, u_2, \dots, u_r), u_i \in U_i(F)$, is an *F*-chain if the corresponding subsets of *F* form a chain under containment.

Let $V_0 = [m]$ and, for $1 \le k \le r$, let V_k be the disjoint union of all $U_k(F)$, $F \in {\binom{[m]}{r}}$. In other words, each $v \in V_k$ corresponds to the *j*-th copy of the *k*-set F' in $U_k(F)$ for some $1 \le j \le \beta_k$, $F' \in {\binom{F}{k}}$, and $F \in {\binom{[m]}{r}}$. Since V_k consists of β_k copies of each $F' \in {\binom{[m]}{k}}$ for each $F \in {\binom{[m]}{r}}$ with $F' \subseteq F$, we have $|V_k| = {\binom{m}{k}} {\binom{m-k}{r-k}} \beta_k = {\binom{m}{r}} {\binom{r}{k}} \beta_k$. For each $v \in V_k$, let $c_k(v)$ denote the set $u \in U_k(F)$ that corresponds to v. We then define \mathcal{H}_{α} to be the (r+1)-partite (r+1)-graph on $V = V_0 \cup V_1 \cup \cdots \cup V_r$ where (v_0, v_1, \cdots, v_r) , $v_i \in V_i$, is an edge if and only if the corresponding subsets $c_0(v_0), c_1(v_1), \cdots, c_r(v_r)$ form an F-chain for some $F \in {[m] \choose r}$.

To help understand this definition, let us also define the subgraph \mathcal{G}_{α} which will appear in Lemmas 3.3 and 3.4. Write $U'_i = U_i([r])$ and let \mathcal{G}_{α} be the (r + 1)-partite (r + 1)-graph on $U'_0 \cup U'_1 \cup \cdots \cup U'_r$ where (u_0, u_1, \cdots, u_r) , $u_i \in U'_i$, is an edge if and only if it is an [r]-chain. Clearly, \mathcal{G}_{α} is isomorphic to each subgraph of \mathcal{H}_{α} induced on $U_0(F) \cup U_1(F) \cup \cdots \cup U_r(F)$ for an rset F. Moreover, these induced subgraphs isomorphic to \mathcal{G}_{α} are edge disjoint, so $e(\mathcal{H}_{\alpha}) = {m \choose r} e(\mathcal{G}_{\alpha})$.

We make two further remarks about the definition of \mathcal{H}_{α} . First, note that V_0 has a different status to the rest of the V_k in that, for each $F \in \binom{[m]}{r}$, $U_0(F)$ is identified with the subset F of $V_0 = [m]$, while, for all other k, the $U_k(F)$ are all disjoint subsets of V_k . Thus, the $\binom{m}{r}$ edge-disjoint copies of \mathcal{G}_{α} that decompose \mathcal{H}_{α} are close to being vertex disjoint in the sense that they can only intersect in V_0 . Second, if any $\beta_k = 0$ for k < r, we simply ignore the k-th coordinate. This will reduce the uniformity of the hypergraph, but all of our arguments still go through in this case. For convenience of notation, we assume in what follows that $\beta_k \neq 0$ for all k.

Now let $W_{\alpha}: [0,1]^{r+1} \to [0,1]$ be the function

$$W_{\alpha}(x, z_1, z_2, \cdots, z_r) = \prod_{k=1}^r W(x, z_k)^{q_k},$$

where

$$q_k = \frac{\beta_k}{\beta(k-1)!(r-k)!}$$

Our first step is contained in the next lemma.

Lemma 3.2. Suppose $r \leq m$ and let H be the (m, r)-downset graph and $\alpha = (\alpha_k)_{k=1}^r$ be a symmetric integer weight vector such that $\binom{m-k}{r-k}$ divides α_k for each $1 \leq k \leq r$. Then, for every graphon W,

$$t_H^{\alpha}(W) = t_{\mathcal{H}_{\alpha}}(W_{\alpha}).$$

Proof. For $F \in {\binom{[m]}{r}}$, let $g(x_F) := \prod_{F' \subseteq F} \rho(x_{F'})^{\beta_{|F'|}}$. Then, since any k-set is contained in exactly ${\binom{m-k}{r-k}}$ r-sets, we may rewrite $t_H^{\alpha}(W)$ as

$$t_{H}^{\alpha}(W) = \mathbb{E}\left[\prod_{k=1}^{r} \prod_{F \in \binom{[m]}{k}} \rho(x_{F})^{\alpha_{k}}\right] = \mathbb{E}\left[\prod_{F \in \binom{[m]}{r}} g(x_{F})\right].$$
(3)

For each k-set F', note that

$$\rho(x_{F'})^{\beta_k} = \mathbb{E}_{y_1, \cdots, y_{\beta_k}} \prod_{i \in F'} \prod_{j=1}^{\beta_k} W(x_i, y_j).$$

Now relabel all the y_j , $j = 1, 2, \dots, \beta_k$, for each $F' \in {F \choose k}$, by mutually independent uniform random variables z_v , where v ranges over those $v \in U_k(F)$. Let $U(F) := U_1(F) \cup U_2(F) \cup \dots \cup U_r(F)$ and write $i \sim v$ if i is contained in the k-subset of F that corresponds to v. Then, for each $F \in {[m] \choose r}$,

$$g(x_F) = \prod_{F' \subseteq F} \rho(x_{F'})^{\beta_{|F'|}} = \mathbb{E}_{z_v: v \in U(F)} \left[\prod_{v \in U(F)} \prod_{i \sim v} W(x_i, z_v) \right],$$

since the $z_v, v \in U(F)$, are mutually independent. We may repeat this step for each *r*-set *F* while assigning mutually independent random variables z_v for all $v \in \bigcup_{F \in \binom{[m]}{U(F)}} U(F)$. Then, by (3),

$$t_{H}^{\alpha}(W) = \mathbb{E}\left[\prod_{F \in \binom{[m]}{r}} g(x_{F})\right] = \mathbb{E}_{x_{[m]}}\left[\prod_{F \in \binom{[m]}{r}} \mathbb{E}_{z_{v}:v \in U(F)}\left[\prod_{v \in U(F)} \prod_{i \sim v} W(x_{i}, z_{v})\right]\right]$$
$$= \mathbb{E}\left[\prod_{F \in \binom{[m]}{r}} \prod_{v \in U(F)} \prod_{i \sim v} W(x_{i}, z_{v})\right],$$
(4)

where the last equality follows from mutual independence. We shall verify that the right-hand side equals $t_{\mathcal{H}_{\alpha}}(W_{\alpha})$ by comparing the exponents of the $W(x_i, z_v)$ in (4) with those in

$$t_{\mathcal{H}_{\alpha}}(W_{\alpha}) = \mathbb{E}\left[\prod_{(i,v_1,\cdots,v_r)\in E(\mathcal{H}_{\alpha})}\prod_{k=1}^r W(x_i, z_{v_k})^{q_k}\right].$$
(5)

For each pair consisting of $i \in [m]$ and $v \in U_k(F)$ with $i \sim v$, the term $W(x_i, z_v)$ appears exactly once in the product in (4). On the other hand, there exist $\beta(k-1)!(r-k)!/\beta_k$ F-chains that contain $\{i\}$ and v. Thus, there are $\beta(k-1)!(r-k)!/\beta_k$ hyperedges in \mathcal{H}_{α} containing i and v. Hence, by expanding the product in (5), each $W(x_i, z_v)$ receives the exponent

$$\frac{\beta(k-1)!(r-k)!q_k}{\beta_k} = 1,$$

where we used the definition of q_k .

The next lemma gives a lower bound for the homomorphism density $t_{\mathcal{H}_{\alpha}}(W_{\alpha})$ in terms of $t_{\mathcal{G}_{\alpha}}(W_{\alpha})$. This follows in a straightforward manner from a result of the authors [4], so we have consigned the proof, and a broader discussion of weakly norming hypergraphs, to an appendix.

Lemma 3.3. The (r+1)-graph \mathcal{H}_{α} is weakly norming. In particular, for every (r+1)-graphon W_{α} ,

$$t_{\mathcal{H}_{\alpha}}(W_{\alpha}) \ge t_{\mathcal{G}_{\alpha}}(W_{\alpha})^{\binom{m}{r}}$$

The following lemma is the final ingredient we need to prove Theorem 3.1.

Lemma 3.4. If $\binom{m-k}{r-k}$ divides α_k for each $1 \leq k \leq r$ and $\alpha_r > 0$, then, for every graphon W,

$$t_{\mathcal{G}_{\alpha}}(W_{\alpha}) \ge t_{K_2}(W)^{q_{\alpha,H}},$$

where $q_{\alpha,H} = e_{\alpha}(H) / {m \choose r}$.

Proof. Let $U' = U'_1 \cup U'_2 \cup \cdots \cup U'_r$, where the U'_i are as in the definition of \mathcal{G}_{α} . Then, following the proof of Lemma 3.2, we have

$$t_{\mathcal{G}_{\alpha}}(W_{\alpha}) = \mathbb{E}\left[\prod_{(i,u_1,\cdots,u_r)\in E(\mathcal{G}_{\alpha})}\prod_{k=1}^r W(x_i,z_{u_k})^{q_k}\right] = \mathbb{E}\left[\prod_{i\sim u,i\in[r],u\in U'}W(x_i,z_u)\right],$$

where $i \sim u$ means that i is contained in the subset of [r] corresponding to u. Hence, we may write

$$t_{\mathcal{G}_{\alpha}}(W_{\alpha}) = \mathbb{E}\left[\mathbb{E}_{z_{u}:u\in U'}\left[\prod_{i\sim u,i\in[r],u\in U'}W(x_{i},z_{u})\right]\right] = \mathbb{E}\left[\rho(x_{[r]})^{\beta_{r}}\prod_{F\subsetneq[r]}\rho(x_{F})^{\beta_{|F|}}\right] = t_{J}^{\beta}(W),\quad(6)$$

where J is the (r, r)-downset graph and $\beta = (\beta_i)_{i=1}^r$ is a symmetric integer weight vector with $\beta_r = \alpha_r \ge 1$. Since $t_J^{\beta}(W)$ can be interpreted as the homomorphism density in W of a bipartite graph with at least one vertex complete to the other side, the result of Conlon, Fox and Sudakov [1] implies that

 $t_J^{\beta}(W) \ge t_{K_2}(W)^{e_{\beta}(J)},$

where $e_{\beta}(J) = \sum_{k=1}^{r} {\binom{r}{k}} k \beta_k$. By the elementary identity ${\binom{m}{r}} {\binom{r}{k}} = {\binom{m}{k}} {\binom{m-k}{r-k}}$,

$$e_{\beta}(J) = \frac{1}{\binom{m}{r}} \sum_{k=1}^{r} \binom{m}{k} \binom{m-k}{r-k} k\beta_{k} = \frac{1}{\binom{m}{r}} \sum_{k=1}^{r} \binom{m}{k} k\alpha_{k} = \frac{e_{\alpha}(H)}{\binom{m}{r}},$$

as desired.

Proof of Theorem 3.1. By Lemmas 3.2, 3.3, and 3.4, we obtain

$$t_{H}^{\alpha}(W) = t_{\mathcal{H}_{\alpha}}(W_{\alpha}) \ge t_{\mathcal{G}_{\alpha}}(W_{\alpha})^{\binom{m}{r}} \ge t_{K_{2}}(W)^{e_{\alpha}(H)},$$

as required.

4 The Hölder trick

The argument in Section 2 that M^2 satisfies Sidorenko's conjecture for $M = K_{5,5} \setminus C_{10}$ had two ingredients, the fact that the (5,3)-incidence graph $F_{5,3}$ is weakly norming and, therefore, satisfies the conjecture, and the inequality $t_{M^2}(W) \ge t_{F_{5,3}}(W)$ for all graphons W. The main result of the previous section may be seen as a generalisation of the fact that $F_{5,3}$ satisfies the conjecture. To complete the proof of Theorem 1.1, we now generalise the inequality. This will again be a simple consequence of Hölder's inequality.

Theorem 4.1. Let H be a bipartite graph with bipartition $A \cup B$ and $\max_{b \in B} \deg(b) = r$ and, for each $1 \leq k \leq r$, let $\alpha_k := d_k / \binom{|A|}{k}$, where d_k is the number of vertices in B of degree k. Then, for every graphon W,

$$t_H(W) \ge t_J^{\alpha}(W),$$

where J is the (|A|, r)-downset graph and $\alpha = (\alpha_k)_{k=1}^r$. Moreover, if deg(b) = r for all $b \in B$, then, for every graphon W,

$$t_H(W) \ge t_F^{\alpha}(W),$$

where F is the (|A|, r)-incidence graph and $\alpha = (\alpha_k)_{k=1}^r$ is the symmetric weight vector with $\alpha_r = |B|/{|A| \choose r}$ and $\alpha_k = 0$ for $k \neq r$.

Theorem 1.1 is an immediate consequence of this result and Theorem 3.1. Recall the statement, that Sidorenko's conjecture holds for any bipartite graph H with bipartition $A \cup B$ and $\max_{b \in B} \deg(b) = r$ such that the number of vertices of degree k is divisible by $\binom{|A|}{r}\binom{r}{k} = \binom{|A|}{k}\binom{|A|-k}{r-k}$ for each $1 \le k \le r$.

Proof of Theorem 1.1. As in Theorem 4.1, for each $1 \le k \le r$, let d_k be the number of vertices in B of degree k and $\alpha_k := \frac{d_k}{\binom{|A|}{k}}$. Then, by Theorem 4.1,

$$t_H(W) \ge t_J^{\alpha}(W),$$

where J is the (|A|, r)-downset graph. By the divisibility assumption, $\alpha = (\alpha_k)_{k=1}^r$ is a nonnegative integer vector such that $\binom{|A|-k}{r-k}$ divides α_k for each $1 \leq k \leq r$ and $\alpha_r > 0$. Hence, we may apply Theorem 3.1 to conclude that

$$t_J^{\alpha}(W) \ge t_{K_2}(W)^{e_{\alpha}(J)}.$$

Therefore, since

$$e_{\alpha}(J) = \sum_{k=1}^{r} {\binom{|A|}{k}} \alpha_k k = \sum_{k=1}^{r} d_k k = e(H),$$

Sidorenko's conjecture holds for H.

Theorem 1.2 is a similar consequence of the second part of Theorem 4.1 with one slight twist. *Proof of Theorem 1.2.* With the slight abuse of notation $\alpha = \alpha_r = |B|/{|A| \choose r}$, we have

 Γ Γ Γ Γ α

$$t_F^{\alpha}(W) = \mathbb{E}\left[\prod_{I \in \binom{A}{r}} \rho(x_I)^{\alpha}\right] \ge \mathbb{E}\left[\prod_{I \in \binom{A}{r}} \rho(x_I)\right] = t_F(W)^{\alpha},$$

where the inequality follows from the convexity of $f(x) = x^{\alpha}$ when $\alpha \ge 1$. Therefore,

$$t_H(W) \ge t_F^{\alpha}(W) \ge t_F(W)^{\alpha} \ge t_{K_2}(W)^{\alpha r \binom{|A|}{r}} = t_{K_2}(W)^{e(H)},$$

where the first inequality is the second part of Theorem 1.2, the third inequality follows from the fact that the (|A|, r)-incidence graph F is weakly norming [4] and, therefore, satisfies Sidorenko's conjecture, and the final equality comes from the straightforward computation $\alpha r \binom{|A|}{r} = r|B| = e(H)$.

We now return to the proof of Theorem 4.1.

Proof of Theorem 4.1. Let m := |A| and identify the set A with [m]. Let $\mathcal{F} = \{f_I : I \subseteq [m]\}$ be the family of functions given by $f_I(x_I) = \rho(x_I)^{c_I}$, where c_I is the number of vertices $b \in B$ such that N(b) = I. If I and I' are subsets of [m] of size k, then there are exactly k!(m-k)! permutations ϕ that map I onto I' and, for every such ϕ , $f_{\phi(I)}(x_I) = \rho(x_I)^{c_I'}$. Thus, for each I of size k,

$$\tilde{f}_{I}(x_{I}) = \left(\prod_{\phi \in S_{m}} f_{\phi(I)}(x_{I})\right)^{1/m!} = \rho(x_{I})^{\sum_{|K|=k} c_{K}/\binom{m}{k}} = \rho(x_{I})^{\alpha_{k}}.$$

On the other hand, since each $\phi \in S_m$ is just a relabeling of A = [m],

$$t_H(W) = \mathbb{E}\left[\prod_{I \subseteq [m]} f_{\phi(I)}(x_I)\right].$$

Therefore, Hölder's inequality in the form (1) gives

$$t_J^{\alpha}(W) = \mathbb{E}\left[\prod_{I \subseteq [m]} \rho(x_I)^{\alpha_{|I|}}\right] = \mathbb{E}\left[\prod_{I \subseteq [m]} \tilde{f}_I(x_I)\right] \le \prod_{\phi \in S_m} \mathbb{E}\left[\prod_{I \subseteq [m]} f_{\phi(I)}(x_I)\right]^{1/m!} = t_H(W),$$

as desired. Moreover, if deg(b) = r for all $b \in B$, then $\alpha_r = |B|/\binom{m}{r}$ and $\alpha_k = 0$ for $k \neq r$. Thus,

$$t_H(W) \ge \mathbb{E}\left[\prod_{I \in \binom{[m]}{r}} \rho(x_I)^{\alpha_r}\right] = t_F^{\alpha}(W),$$

which proves the second part of the statement.

5 Concluding remarks

As a closing remark, we note another corollary of Theorem 1.1 which we believe to be of independent interest. Put briefly, it says that a certain local version of Sidorenko's conjecture holds for any bipartite graph H.

Corollary 5.1. For every bipartite graph H with bipartition $A \cup B$ and every $\varepsilon, q > 0$, there is $n_0 = n_0(\varepsilon, H, q)$ such that every n-vertex graph G with $n \ge n_0$ and $q = t_{K_2}(G)$ has an |A|-tuple x_A of distinct vertices such that $t_H(G; x_A) \ge (1 - \varepsilon)q^{e(H)}$.

Proof. Let a = |A| and n = |V(G)|. Let K be the set of non-degenerate a-tuples in $V(G)^A$, noting that $|V(G)^A \setminus K| \leq an^{a-1}$. Then

$$\mathbb{E}_{x_A} t_H(G; x_A)^k = \frac{1}{n^a} \left(\sum_{x_A \in K} t_H(G; x_A)^k + \sum_{x_A \notin K} t_H(G; x_A)^k \right) \le \frac{1}{n^a} \sum_{x_A \in K} t_H(G; x_A)^k + \frac{a}{n}.$$

Therefore, by letting k = a!, Corollary 1.3 implies that

$$\mathbb{E}_{x_A \in K} t_H(G; x_A)^k \ge \frac{n^a}{n(n-1)\cdots(n-a+1)} \left(\mathbb{E}_{x_A} t_H(G; x_A)^k - \frac{a}{n} \right)$$
$$\ge (1-\varepsilon) t_{K_2}(G)^{k \cdot e(H)},$$

provided *n* is sufficiently large in terms of ϵ , *H* and *q*. There must then be $x_A \in K$ with $t_H(G; x_A) \ge (1-\varepsilon)t_{K_2}(G)^{e(H)}$, since otherwise we would have $\mathbb{E}_{x_A \in K} t_H(G; x_A)^k < (1-\varepsilon)^k t_{K_2}(G)^{k \cdot e(H)}$, contradicting the inequality above.

Acknowledgements. We are greatly indebted to Yufei Zhao and also to Leonardo Nagami Coregliano and Sasha Razborov for spotting a substantial error in an earlier version of this paper. The upshot of the resulting changes is the divisibility condition in Theorem 1.1, which was not present in the previous version. We are also grateful to the anonymous referees for their detailed reviews.

References

- D. Conlon, J. Fox, and B. Sudakov. An approximate version of Sidorenko's conjecture. *Geom. Funct. Anal.*, 20:1354–1366, 2010.
- [2] D. Conlon, J. H. Kim, C. Lee, and J. Lee. Some advances on Sidorenko's conjecture. J. London Math. Soc., 98:593–608, 2018.
- [3] D. Conlon, J. H. Kim, C. Lee, and J. Lee. Sidorenko's conjecture for higher tree decompositions. Unpublished note, available at arXiv:1805.02238 [math.CO], 2018.
- [4] D. Conlon and J. Lee. Finite reflection groups and graph norms. Adv. Math., 315:130–165, 2017.
- [5] P. Erdős and M. Simonovits. Supersaturated graphs and hypergraphs. *Combinatorica*, 3:181– 192, 1983.
- [6] J. Fox and B. Sudakov. Dependent random choice. Random Structures Algorithms, 38:68–99, 2011.
- [7] H. Hatami. Graph norms and Sidorenko's conjecture. Israel J. Math., 175:125–150, 2010.
- [8] J. H. Kim, C. Lee, and J. Lee. Two approaches to Sidorenko's conjecture. Trans. Amer. Math. Soc., 368:5057–5074, 2016.
- [9] J. Li and B. Szegedy. On the logarithmic calculus and Sidorenko's conjecture. To appear in Combinatorica.
- [10] L. Lovász. Graph homomorphisms: Open problems. Manuscript available at http://www.cs.elte.hu/~lovasz/problems.pdf, 2008.
- [11] V. Nikiforov. The number of cliques in graphs of given order and size. Trans. Amer. Math. Soc., 363:1599–1618, 2011.
- [12] A. A. Razborov. Flag algebras. J. Symbolic Logic, 72:1239–1282, 2007.
- [13] A. A. Razborov. On the minimal density of triangles in graphs. Combin. Probab. Comput., 17:603–618, 2008.
- [14] C. Reiher. The clique density theorem. Ann. of Math., 184:683–707, 2016.
- [15] A. Sidorenko. A correlation inequality for bipartite graphs. Graphs Combin., 9:201–204, 1993.
- [16] B. Szegedy. An information theoretic approach to Sidorenko's conjecture. Preprint available at arXiv:1406.6738 [math.CO].

A Weakly norming hypergraphs

Following Hatami [7], we say that a hypergraph H is weakly norming if the functional $\|\cdot\|_{r(H)}$ defined by

$$||W||_{r(H)} = t_H(|W|)^{1/e(H)}$$

is a norm on the space of bounded symmetric measurable functions. As shown by Hatami [7], any weakly norming hypergraph H has the property that for any $H' \subseteq H$ and any graphon W,

$$||W||_{r(H)} \ge ||W||_{r(H')}$$

or, in the language of homomorphism densities,

$$t_H(W) \ge t_{H'}(W)^{e(H)/e(H')}$$

In particular, this implies that any weakly norming hypergraph satisfies Sidorenko's conjecture.

The main result of interest to us here is a result of the authors [4, Theorem 5.1] saying that a certain class of hypergraphs, which we call reflection hypergraphs, are weakly norming. To define this class, suppose that W is a finite reflection group, T is a set of simple reflections in W, and T_1, \dots, T_r are subsets of T. Then the $(T_1, \dots, T_r; T, W)$ -reflection hypergraph is the r-partite r-graph whose parts are the cosets of the subgroup W_k generated by T_k for each $k = 1, \dots, r$, with an edge for every r-tuple of the form (wW_1, \dots, wW_r) with $w \in W$. An r-graph is then said to be a reflection hypergraph if it is isomorphic to the $(T_1, \dots, T_r; T, W)$ -reflection hypergraph for some choice of parameters.

Theorem A.1. Reflection hypergraphs are weakly norming.

In order to prove Lemma 3.3, it therefore suffices to show that \mathcal{H}_{α} is a reflection hypergraph. If $\beta_k = 1$ for each k, then it is not hard to construct a reflection hypergraph that is isomorphic to \mathcal{G}_{α} . Let s_{ij} be the permutation in S_r that swaps i and j. It is well-known that $T := \{s_{i(i+1)} : i = 1, 2, \dots, r-1\}$ is a set of simple reflections. We claim that the $(T_0, T_1, T_2, \dots, T_r; T, S_r)$ -reflection hypergraph with the choice $T_0 = T_1, T_i = T \setminus \{s_{i(i+1)}\}, 1 \leq i < r, T_r = T$, is isomorphic to \mathcal{G}_{α} when $\beta_1 = \beta_2 = \dots = \beta_r = 1$. To see this, observe first that each $T_k, 1 \leq k < r$, generates the subgroup

$$W_k := \{ \sigma \in S_r : \sigma(j) \le k \text{ for each } j \le k \},\$$

which is isomorphic to $S_k \times S_{r-k}$. Thus, each coset wW_k is the set of permutations in S_r that map [k] onto some k-subset $F \in {[r] \choose k}$, which allows us to identify each coset with the k-set F. For example, the chain $\{1\} = \{1\} \subset \{1,2\} \subset \cdots \subset [r-1] \subset [r]$, which is an edge in \mathcal{G}_{α} , corresponds to the (r+1)-tuple $(W_1, W_1, \cdots, W_{r-1}, S_r)$. Similarly, every edge in \mathcal{G}_{α} , identified by a chain $\{i\} =$ $F_1 \subset \cdots \subset F_{r-1} \subset [r]$, where $|F_k| = k$, corresponds to an (r+1)-tuple $(wW_1, wW_1, \cdots, wW_{r-1}, S_r)$ of cosets, where $w \in S_r$ is a permutation that maps 1 to *i*. This proves the claim.

Recall now that \mathcal{H}_{α} consists of $\binom{m}{r}$ edge-disjoint (r+1)-graphs $\mathcal{H}_{\alpha}[U_0(F) \cup U(F)]$, $F \in \binom{[m]}{r}$, each of which is isomorphic to \mathcal{G}_{α} . To realise \mathcal{H}_{α} with $\beta_i = 1$ for all $i = 1, 2, \cdots, r$ as a reflection hypergraph, one may add new reflections $s_{j(j+1)}$, for $j = r, \cdots, m-1$, to T, and amend the T_k , $0 \leq k \leq r$, to generate $\binom{m}{r}$ copies of each coset wW_k , $1 \leq k \leq r$, in the $(T_0, T_1, \cdots, T_r; T, S_r)$ reflection hypergraph. More explicitly, \mathcal{H}_{α} is isomorphic to the $(T'_0, T'_1, T'_2, \cdots, T'_r; T', S_m)$ -reflection hypergraph, where $T' = T \cup \{s_{j(j+1)} : r \leq j < m\}$, $T'_0 = T_0 \cup \{s_{j(j+1)} : r \leq j < m\}$ and $T'_k = T_k \cup \{s_{j(j+1)} : r+1 \leq j < m\}$ for $0 < k \leq r$. To see this, note first that the cosets of the subgroup generated by T'_0 and T'_r correspond to singletons and r-subsets in [m], respectively. For $1 \leq k < r, T'_k = T' \setminus \{s_{k(k+1)}, s_{r(r+1)}\}$ generates the subgroup

$$W'_k := \{ \sigma \in S_m : \sigma(i) \le k \text{ if } 1 \le i \le k, \ k < \sigma(j) \le r \text{ if } k < j \le r \},\$$

which is isomorphic to $S_k \times S_{r-k} \times S_{m-r}$. Thus, each coset of W'_k is identified with a disjoint pair (F', F'') consisting of a k-set F' and an (r-k)-set F'', which corresponds to an element in

 $U_k(F' \cup F'')$. For instance, W'_k corresponds to the copy of [k] in $U_k([r])$. Then it is easy to check that each edge in the $(T'_0, T'_1, T'_2, \cdots, T'_r; T', S_m)$ -reflection hypergraph corresponds to an *F*-chain for some *r*-set *F*.

We also remark that these constructions generalise to the case $\beta_k \in \{0, 1\}$ for $k = 1, 2, \dots, r$. If some $\beta_k = 0$, one may simply delete the corresponding T_k or T'_k when constructing \mathcal{G}_{α} and \mathcal{H}_{α} , respectively. By doing so, reflection *t*-graphs, where *t* is the number of positive β_k , are obtained.

Given the union of r disjoint sets $A_1 \cup A_2 \cup \cdots \cup A_r$, let $\overline{A}_k = \bigcup_{i=1}^r A_i \setminus A_k$. Moreover, if \mathcal{H} is an r-partite r-graph with r-partition $A_1 \cup A_2 \cup \cdots \cup A_r$, the blow-up $\mathcal{H}^p_{\overline{A}_k}$ is obtained by taking pvertex-disjoint copies of \mathcal{H} and gluing all copies of \overline{A}_k along corresponding vertices. The following lemma then says that blow-ups of reflection hypergraphs are also reflection hypergraphs.

Lemma A.2. Let \mathcal{H} be an r-partite r-graph on $A_1 \cup A_2 \cup \cdots \cup A_r$. If \mathcal{H} is a reflection hypergraph and k and p are positive integers with $1 \leq k \leq r$, then the blow-up $\mathcal{H}^p_{\overline{A}_k}$ is also a reflection hypergraph.

Proof. Suppose that \mathcal{H} is isomorphic to the $(T_1, T_2, \cdots, T_r; T, W)$ -reflection hypergraph, where W is a finite reflection group, T is a set of simple reflections in W and $T_i \subseteq T$, $i = 1, 2, \cdots, r$. Let $P = \{\sigma_{12}, \sigma_{23}, \cdots, \sigma_{(p-1)p}\}$ be a set of p-1 new reflections such that each element of P commutes with all the elements in W and P generates a group which is isomorphic to S_p , where σ_{ij} corresponds to the permutation swapping i and j in [p]. Let $T' := T \cup P$, $T'_i := T_i \cup P$ for $i \neq k$, $T'_k := T_k \cup P \setminus \{\sigma_{12}\}$, and let W' be the new reflection group generated by T', i.e., $W' \cong W \times S_p$. Let \mathcal{H}' be the $(T'_1, \cdots, T'_r; T', W')$ -reflection hypergraph. We claim that \mathcal{H}' is isomorphic to $\mathcal{H}^p_{\overline{A}_k}$. Let $A'_i := \{wW'_i : w \in W\}$. Then \mathcal{H}' is an r-graph on the new reflections in T, whereas $A'_k = \bigcup^p_{j=1}(A_k \times \sigma_{1j}S_{[2,p]})$. Let $\phi: V(\mathcal{H}') \to V(\mathcal{H})$ be the map defined by

 $\phi(wW_i \times S_p) := wW_i$ for each $i \neq k$ and $\phi(wW_k \times \sigma_{1j}S_{[2,p]}) := wW_k$ for each $1 \leq j \leq p$.

One may check that ϕ is a surjective hypergraph homomorphism from \mathcal{H}' to \mathcal{H} . Moreover, it is injective on $\bigcup_{i \neq k} A'_i$, *p*-to-1 on A'_k , and maps every non-edge in \mathcal{H}' to a non-edge in \mathcal{H} . Therefore, \mathcal{H}' is obtained by gluing *p* vertex-disjoint copies of the $(T_1, \cdots, T_r; T, W)$ -reflection hypergraph along the vertices that are not cosets of W_k and, hence, it is isomorphic to $\mathcal{H}^p_{\overline{A}_i}$.

Proof of Lemma 3.3. Let \mathcal{H} be the hypergraph \mathcal{H}_{α} with $\beta_k \in \{0, 1\}$ for all k. By applying the lemma above with this \mathcal{H} , one may blow-up each vertex in the vertex set V_k of \mathcal{H} to β_k copies. Thus, by repeating this blowing-up process, we conclude that \mathcal{H}_{α} is a reflection hypergraph for any given non-negative integers $\beta_1, \beta_2, \cdots, \beta_r$.