

Some advances on Sidorenko's conjecture

David Conlon* Jeong Han Kim† Choongbum Lee‡ Joonkyung Lee§

Abstract

A bipartite graph H is said to have Sidorenko's property if the probability that the uniform random mapping from $V(H)$ to the vertex set of any graph G is a homomorphism is at least the product over all edges in H of the probability that the edge is mapped to an edge of G . In this paper, we provide three distinct families of bipartite graphs that have Sidorenko's property. First, using branching random walks, we develop an embedding algorithm which allows us to prove that bipartite graphs admitting a certain type of tree decomposition have Sidorenko's property. Second, we use the concept of locally dense graphs to prove that subdivisions of certain graphs, including cliques, have Sidorenko's property. Third, we prove that if H has Sidorenko's property, then the Cartesian product of H with an even cycle also has Sidorenko's property.

1 Introduction

For two graphs H and G , a homomorphism from H to G is a mapping $g : V(H) \rightarrow V(G)$ such that $\{g(v), g(w)\}$ is an edge in G whenever $\{v, w\}$ is an edge in H . Let $\text{Hom}(H, G)$ denote the set of all homomorphisms from H to G . A beautiful conjecture of Sidorenko [13] (also made by Erdős and Simonovits [14] in a slightly different form) asserts the following.

Conjecture 1.1. *For all bipartite graphs H and all graphs G ,*

$$|\text{Hom}(H, G)| \geq |V(G)|^{|V(H)|} \left(\frac{2|E(G)|}{|V(G)|^2} \right)^{|E(H)|}. \quad (1)$$

That is, the probability that a uniform random mapping from $V(H)$ to $V(G)$ is a homomorphism is at least the product over all edges of the probability that the edge is mapped to an edge of G .

*Mathematical Institute, Oxford OX2 6GG, United Kingdom. Email: david.conlon@maths.ox.ac.uk. Research supported by a Royal Society University Research Fellowship.

†School of Computational Sciences, Korea Institute for Advanced Study (KIAS), Seoul, South Korea. Email: jhkim@kias.re.kr. The author was supported by a National Research Foundation of Korea (NRF) Grant funded by the Korean Government (MSIP) (NRF-2012R1A2A2A01018585) and KIAS Internal Research Fund CG046001. This work was partially carried out while the author was visiting Microsoft Research, Redmond and Microsoft Research, New England.

‡Department of Mathematics, MIT, Cambridge, MA 02139-4307. Email: cb_lee@math.mit.edu. Research supported by NSF Grant DMS-1362326.

§Mathematical Institute, Oxford OX2 6GG, United Kingdom. Email: joonkyung.lee@maths.ox.ac.uk. Supported by the ILJU Foundation of Education and Culture.




Figure 1: A strongly tree-decomposable graph (each colored tree represents $H[X]$ for some $X \in \mathcal{F}$).

We say that a graph H has Sidorenko's property if (1) holds for all graphs G . While Sidorenko [13] himself noted that the conjecture holds for some simple graphs such as trees, even cycles and complete bipartite graphs, a spate of recent work [1, 6, 7, 9], some of which we will describe below, has greatly expanded the class of graphs known to have Sidorenko's property. In this paper, we further this progress, providing three more families of graphs that satisfy the conjecture.

Our first family consists of those graphs which are amenable to a certain embedding process using branching random walks. Though the method can be applied more broadly, we will focus on using the technique to show that graphs admitting a particular type of tree decomposition have Sidorenko's property. A tree decomposition of a graph H , a concept introduced by Halin [5] and developed by Robertson and Seymour [12], is a pair $(\mathcal{F}, \mathcal{T})$ consisting of a family \mathcal{F} of vertex subsets of H and a tree \mathcal{T} on vertex set \mathcal{F} satisfying

1. $\bigcup_{X \in \mathcal{F}} X = V(H)$,
2. for each $\{v, w\} \in E(H)$, there exists a set $X \in \mathcal{F}$ such that $v, w \in X$, and
3. for $X, Y, Z \in \mathcal{F}$, $X \cap Y \subseteq Z$ whenever Z lies on the path from X to Y in \mathcal{T} .

A strong tree decomposition of a graph H is a tree decomposition $(\mathcal{F}, \mathcal{T})$ of H satisfying the following two extra conditions:

1. The induced subgraphs $H[X]$, $X \in \mathcal{F}$, are edge-disjoint trees.
2. For every pair $X, Y \in \mathcal{F}$ which are adjacent in \mathcal{T} , there is an isomorphism between the minimum subtrees of $H[X]$ and $H[Y]$ containing $X \cap Y$ that fixes $X \cap Y$.

A graph H is strongly tree-decomposable if it admits a strong tree decomposition. One can easily verify that every strongly tree-decomposable graph is bipartite. Our first theorem is as follows.

Theorem 1.2. *If H is strongly tree-decomposable, then H has Sidorenko's property.*

The family of strongly tree-decomposable graphs completely contains some other families for which the conjecture was previously known. For example, the family of reflection trees studied by

Li and Szegedy [9] consists exactly of those strongly tree-decomposable graphs for which \mathcal{T} is a star, while the family of tree-arrangeable graphs considered in [7] consists of those graphs H such that $\mathcal{F} = \{N_H(a) \cup \{a\} \mid a \in A\}$ for an independent set A in $V(H)$. However, our new family is broader than either of these, Figure 1 showing an example of a strongly tree-decomposable graph that is neither a reflection tree nor tree arrangeable.

The second family of graphs we consider consists of subdivisions of certain graphs. The subdivision of a graph H is the graph obtained from H by replacing its edges with paths of length two that are internally vertex disjoint. As a corollary of a more general theorem relating Sidorenko’s conjecture to a conjecture about subgraphs of locally dense graphs, we obtain the following results.

Theorem 1.3. (i) *For every positive integer k , the subdivision of K_k has Sidorenko’s property.*
(ii) *If H has Sidorenko’s property, then the subdivision of H also has Sidorenko’s property.*

The third family of graphs are those formed through taking certain Cartesian products. The Cartesian product $H_1 \square H_2$ of two graphs H_1 and H_2 is the graph on $V(H_1) \times V(H_2)$ where (x_1, x_2) and (y_1, y_2) are adjacent if and only if either (i) x_1 and y_1 are adjacent in H_1 and $x_2 = y_2$, or (ii) x_2 and y_2 are adjacent in H_2 and $x_1 = y_1$. In [7], it was proved that if H is a graph having Sidorenko’s property and T is a tree, then $H \square T$ also has Sidorenko’s property. In particular, this result implies that grid graphs of all dimensions have Sidorenko’s property. Here, we prove that a similar result holds when trees are replaced by even cycles.

Theorem 1.4. *Let $k \geq 2$ be a positive integer. If H has Sidorenko’s property, then $H \square C_{2k}$ also has Sidorenko’s property.*

In recent, independent work, Szegedy [15] developed a recursive procedure that generates a large family of graphs having Sidorenko’s property. Szegedy’s result is related to, but more general than, our Theorem 1.2. For example, it shows that 2-dimensional grids, which are not strongly tree-decomposable (but are amenable to other methods [7]), have Sidorenko’s property. Nevertheless, there are generalizations of Theorem 1.2, obtained through further iteration analogous to the way in which one builds strongly tree-decomposable graphs from trees, which bring our result in line with Szegedy’s. Since the statements of these more general results are quite technical, we have chosen to omit further discussion here. However, we refer the reader to the supplementary note [4] for further information.

It is less clear how Szegedy’s method relates to Theorems 1.3 and 1.4. Szegedy has himself shown that his class is closed both under taking Cartesian products with trees and under taking subdivisions and it is plausible that it is also closed under taking Cartesian products with even cycles. This would give a satisfactory analogue of Theorem 1.4, but, even if this doesn’t hold, our result shows that one can simply add this new operation to the definition of the class.

Parczyk [10] has also used Szegedy’s method to recover Theorem 1.3(i), that the subdivision of the complete graph has Sidorenko’s property, but it is unlikely that his method can be used to recover all the results of Section 3 below, since they relate Sidorenko’s conjecture to another conjecture whose status is undetermined. Things are muddled further by the fact that the results of Section 3 may also be generalized to hypergraphs. For example, if H is a k -uniform hypergraph for which the analogue of Sidorenko’s conjecture (or the Kohayakawa–Nagle–Rödl–Schacht conjecture)

holds, then we may subdivide each edge by adding a new vertex connected to each of the k vertices in that edge to produce another graph which satisfies Sidorenko's conjecture. These considerations suggest the following open problem, which lies outside the reach of current methods.

Problem 1.5. *Show that the subdivision of an arbitrary graph H has Sidorenko's property.*

2 Strongly tree-decomposable graphs

The proof of Theorem 1.2 is based on a randomized algorithm that produces a homomorphism in $\text{Hom}(H, G)$, with the desired bound for $|\text{Hom}(H, G)|$ following from an appropriate estimate on the entropy of the random homomorphism.

Before getting into the proof, we give a brief overview of entropy. Suppose that X is a random variable taking values in some finite set S . The entropy of X is defined by

$$H(X) = \sum_{x \in S} \mathbf{P}[X = x] \log \left(\frac{1}{\mathbf{P}[X = x]} \right).$$

Given two random variables X and Y over the same space, the relative entropy of Y with respect to X is

$$H(Y|X) = H(X, Y) - H(X).$$

For three random variables X, Y , and Z , we say that X and Y are conditionally independent given Z if the distribution of X conditioned on $(Y, Z) = (y, z)$ is identical to the distribution of X conditioned on $Z = z$ for all possible values y, z . Finally, we note some simple properties of entropy which will be crucial in what follows.

Lemma 2.1. *Let X, Y , and Z be random variables and suppose that X takes values in a set S . Then*

(i) $H(X) \leq \log |S|$,

(ii) $H(X|Y, Z) = H(X|Z)$ if X and Y are conditionally independent given Z .

2.1 Trees

We start with the case where $H = T$ is a tree. Let r be a fixed vertex of T and think of T as a tree rooted at r . To generate a (not necessarily uniform) random homomorphism \mathbf{w}_T from T to G , we consider the T -branching random walk, or simply T -BRW, on G . If we fix an order $r = u_0, u_1, \dots, u_t$ of all vertices in T in which no child precedes its parent, this algorithm proceeds as follows:

1. For the root r , $\mathbf{w}_T(r)$ is a random vertex in G chosen according to the stationary distribution, that is, with probability proportional to its degree,

$$\mathbf{P}[\mathbf{w}_T(r) = y] = \frac{\deg_G(y)}{2|E(G)|}.$$

2. Suppose that u_i , $i \geq 1$, is the first vertex that has not yet been embedded. Then, if v_i is the parent of u_i , $\mathbf{w}_T(u_i)$ is a vertex chosen uniformly at random from all vertices that are adjacent to $\mathbf{w}_T(v_i)$ in G . (Note that $v_i = u_j$ for some $j < i$, as v_i must precede u_i .)

In others words, when the child u_i of v_i is born, it resides on a vertex of G chosen uniformly at random from all vertices that are adjacent to $\mathbf{w}_T(v_i)$ in G .

3. Repeat Step 2 until all vertices have been embedded.

Unless G has no edge, \mathbf{w}_T is clearly a homomorphism from T to G . Moreover, the distribution of \mathbf{w}_T does not depend on either the choice of the root vertex r or the order of the remaining vertices. To see this, fix $h \in \text{Hom}(T, G)$ and note that

$$\begin{aligned} \mathbf{P}[\mathbf{w}_T = h] &= \frac{\deg_G(h(r))}{2|E(G)|} \left(\frac{1}{\deg_G(h(r))} \right)^{d_r} \prod_{v \in V(T) \setminus \{r\}} \left(\frac{1}{\deg_G(h(v))} \right)^{d_v-1} \\ &= \frac{1}{2|E(G)|} \prod_{v \in V(T)} \left(\frac{1}{\deg_G(h(v))} \right)^{d_v-1}, \end{aligned} \quad (2)$$

where $d_v = \deg_T(v)$. As the last term depends only upon T and G , the claim follows.

We will call the distribution of \mathbf{w}_T on $\text{Hom}(T, G)$ the T -BRW distribution and write $p_r : \text{Hom}(T, G) \rightarrow [0, 1]$ for the probability distribution function of the T -BRW, i.e.,

$$p_r(h) = \mathbf{P}[\mathbf{w}_T = h].$$

Notice that (2) implies

$$p_r(h) > 0 \text{ for all } h \in \text{Hom}(T, G). \quad (3)$$

The T -BRW also has an interesting projection property: if S is a fixed subtree of T , then, for an arbitrary homomorphism $h : V(T) \rightarrow V(G)$, the map h restricted to $V(S)$, which is denoted by $h|_{V(S)}$, or simply $h|_S$, induces a homomorphism from S to G . Hence, $\mathbf{w}_T|_S$ induces another probability distribution over $\text{Hom}(S, G)$. The following proposition asserts that $\mathbf{w}_T|_S$ is identically distributed to \mathbf{w}_S . For a subset X of $V(T)$ and $h : X \rightarrow V(G)$, let $\mathcal{G}_h(T)$ be the set of all homomorphisms from T to G that extend h , i.e., $\mathcal{G}_h(T) = \{g \in \text{Hom}(T, G) : g = h \text{ on } X\}$.

Proposition 2.2. *Let S be a subtree of a tree T and $h \in \text{Hom}(S, G)$. Then,*

$$\mathbf{P}[\mathbf{w}_T = h \text{ on } S] = p_S(h),$$

where \mathbf{w}_T is the T -BRW on G , or equivalently

$$\sum_{g \in \mathcal{G}_h(T)} p_T(g) = p_S(h).$$

Proof. As the distribution of \mathbf{w}_T is independent of the root as well as the order of the vertices, one may regard \mathbf{w}_T as a T -BRW that maps vertices in S before any other vertices in T . Then $\mathbf{w}_T|_S$ and \mathbf{w}_S are constructed in the same way and the desired equality follows. \square

For ordered edges (u, v) and (s, t) of T and G , respectively, one may apply Proposition 2.2 to the tree S consisting of the single edge (u, v) and the homomorphism h on S with $(h(u), h(v)) = (s, t)$ to obtain

$$\mathbf{P}\left[(\mathbf{w}_T(u), \mathbf{w}_T(v)) = (s, t)\right] = p_s(h) = \frac{\deg_G(s)}{2|E(G)|} \frac{1}{\deg_G(s)} = \frac{1}{2|E(G)|}.$$

In other words, the T -BRW \mathbf{w}_T restricted to any edge (u, v) of T has the uniform distribution on the set of all ordered edges in G . Therefore,

$$H(\mathbf{w}_T(u), \mathbf{w}_T(v)) = \log(2|E(G)|). \quad (4)$$

We conclude this subsection by giving a lower bound for the entropy of the T -BRW \mathbf{w}_T , which in turn implies that trees satisfy Sidorenko's property.

Proposition 2.3. *For a tree T and a graph G , the T -BRW \mathbf{w}_T on G satisfies*

$$H(\mathbf{w}_T) \geq \log\left(|V(G)|^{|V(T)|} \left(\frac{2|E(G)|}{|V(G)|^2}\right)^{|E(T)|}\right)$$

and, therefore,

$$|\text{Hom}(T, G)| \geq |V(G)|^{|V(T)|} \left(\frac{2|E(G)|}{|V(G)|^2}\right)^{|E(T)|}.$$

Proof. The second part easily follows from the first part and Lemma 2.1(i), which implies that

$$H(\mathbf{w}_T) \leq \log|\text{Hom}(T, G)|.$$

For the first part, observe, by the definition of conditional entropy and Lemma 2.1(ii), that

$$\begin{aligned} H(\mathbf{w}_T) &= H(\mathbf{w}_T(r)) + \sum_{i=1}^t H(\mathbf{w}_T(u_i) | \mathbf{w}_T(u_0, \dots, u_{i-1})) \\ &= H(\mathbf{w}_T(r)) + \sum_{i=1}^t H(\mathbf{w}_T(u_i) | \mathbf{w}_T(v_i)), \end{aligned}$$

where v_i is the parent of u_i . Denote by C_v the set of all children of v . Rearranging the last sum, we have that

$$\begin{aligned} H(\mathbf{w}_T) &= H(\mathbf{w}_T(r)) + \sum_{v \in V(T)} \sum_{u \in C_v} H(\mathbf{w}_T(u) | \mathbf{w}_T(v)) \\ &= H(\mathbf{w}_T(r)) + \sum_{v \in V(T)} \sum_{u \in C_v} (H(\mathbf{w}_T(u), \mathbf{w}_T(v)) - H(\mathbf{w}_T(v))). \end{aligned}$$

Since, by (4),

$$\sum_{v \in V(T)} \sum_{u \in C_v} H(\mathbf{w}_T(u), \mathbf{w}_T(v)) = \sum_{uv \in E(T)} H(\mathbf{w}_T(u), \mathbf{w}_T(v)) = |E(T)| \log(2|E(G)|)$$

and, by Lemma 2.1(i),

$$-H(\mathbf{w}_T(r)) + \sum_{v \in V(T)} \sum_{u \in C_v} H(\mathbf{w}_T(u)) = \sum_{v \in V(T)} (\deg_T(v) - 1) H(\mathbf{w}_T(v)) \leq (2|E(T)| - |V(T)|) \log |V(G)|,$$

we have that

$$H(\mathbf{w}_T) \geq \log \left(|V(G)|^{|V(T)|} \left(\frac{2|E(G)|}{|V(G)|^2} \right)^{|E(T)|} \right),$$

as desired. \square

2.2 Branching random walks starting from a partial embedding

In order to embed strongly tree-decomposable graphs, we will first need to understand T -BRWs starting from a partial embedding. Let T be a tree, G a graph with at least one edge, and X a subset of $V(T)$. Let S be the minimum subtree of T that contains X . If $X = \emptyset$, then S consists of a single fixed vertex of T . Take an ordering of the vertices of T such that the vertices of S precede all other vertices and, if the tree T is seen as being rooted at the first vertex, no child precedes its parent. The T -branching random walk on G , or simply T -BRW, starting from a partial embedding $h : X \rightarrow V(G)$ is then defined as follows:

1. If $\mathcal{G}_h(S) := \{g \in \text{Hom}(S, G) : g|_X = h\}$ is empty, then there is no T -BRW starting with h . Otherwise, $\mathbf{w}_T(\cdot | h) : V(T) \rightarrow V(G)$ restricted to S is defined according to the distribution

$$\mathbf{P} \left[\mathbf{w}_T(\cdot | h) = y \text{ on } S \right] = \begin{cases} \frac{p_S(y)}{\sum_{g \in \mathcal{G}_h(S)} p_S(g)} & \text{if } y \in \mathcal{G}_h(S) \\ 0 & \text{if } y \notin \mathcal{G}_h(S) . \end{cases}$$

(Recall that p_S is the probability distribution function of the S -BRW on $\text{Hom}(S, G)$ and $p_S(g) > 0$ for all $g \in \text{Hom}(S, G)$ as seen in (3) above).

2. Once $\mathbf{w}_T(\cdot | h)$ is defined on S , the embedding $\mathbf{w}_T(\cdot | h)$ may be defined in the same way that was used for the construction of the T -BRW. That is, for the first vertex u that has not yet been embedded and its parent v , $\mathbf{w}_T(u | h)$ is a vertex chosen uniformly at random from all vertices adjacent to $\mathbf{w}_T(v | h)$ in G .
3. Repeat Step 2 until all vertices have been embedded.

If $X = \emptyset$, the T -BRW starting from $h : X \rightarrow V(G)$ is the usual T -BRW. Moreover, it is not difficult to show that the T -BRW starting from h has the same distribution as the usual T -BRW conditioned on $\mathbf{w}_T|_X = h$. We conclude this subsection by proving this fact.

Proposition 2.4. *With the same notation as above and $x \in \text{Hom}(T, G)$,*

$$\mathbf{P} \left[\mathbf{w}_T(\cdot | h) = x \text{ on } S \right] = \mathbf{P} \left[\mathbf{w}_T = x \text{ on } S \mid \mathbf{w}_T = h \text{ on } X \right]$$

and

$$\mathbf{P}[\mathbf{w}_T(\cdot | h) = x] = \mathbf{P} \left[\mathbf{w}_T = x \mid \mathbf{w}_T = h \text{ on } X \right].$$

In particular, the distribution of $\mathbf{w}_T(\cdot | h)$ is independent of the order of vertices in T .

Proof. As Proposition 2.2 gives

$$p_S(x|S) = \mathbf{P}[\mathbf{w}_T = x \text{ on } S]$$

and

$$\sum_{g \in \mathcal{G}_h(S)} p_S(g) = \sum_{g \in \mathcal{G}_h(S)} \sum_{f \in \mathcal{G}_g(T)} p_T(f) = \sum_{f \in \mathcal{G}_h(T)} p_T(f) = \mathbf{P}[\mathbf{w}_T = h \text{ on } X],$$

we have that

$$\mathbf{P}[\mathbf{w}_T(\cdot|h) = x \text{ on } S] = \frac{p_S(x|S)}{\sum_{g \in \mathcal{G}_h(S)} p_S(g)} = \frac{\mathbf{P}[\mathbf{w}_T = x \text{ on } S]}{\mathbf{P}[\mathbf{w}_T = h \text{ on } X]} = \mathbf{P}[\mathbf{w}_T = x \text{ on } S \mid \mathbf{w}_T = h \text{ on } X].$$

Since \mathbf{w}_T and $\mathbf{w}_T(\cdot|h)$ are constructed in the same way once their values on S have been determined, we also have

$$\begin{aligned} \mathbf{P}[\mathbf{w}_T(\cdot|h) = x] &= \mathbf{P}[\mathbf{w}_T(\cdot|h) = x \text{ on } S] \mathbf{P}[\mathbf{w}_T(\cdot|h) = x \mid \mathbf{w}_T(\cdot|h) = x \text{ on } S] \\ &= \mathbf{P}[\mathbf{w}_T = x \text{ on } S \mid \mathbf{w}_T = h \text{ on } X] \mathbf{P}[\mathbf{w}_T = x \mid \mathbf{w}_T = x \text{ on } S] \\ &= \mathbf{P}[\mathbf{w}_T = x \mid \mathbf{w}_T = h \text{ on } X], \end{aligned}$$

as required. □

2.3 Embedding strongly tree-decomposable graphs

Recall that a graph H is strongly tree-decomposable if it has a tree decomposition $(\mathcal{F}, \mathcal{T})$ satisfying two extra conditions:

1. The induced subgraphs $H[X]$, $X \in \mathcal{F}$, are edge-disjoint trees.
2. For every pair $X, Y \in \mathcal{F}$ which are adjacent in \mathcal{T} , there is an isomorphism that fixes $X \cap Y$ between the minimum subtrees of $H[X]$ and $H[Y]$ containing $X \cap Y$.

Suppose that G has at least one edge. We will attempt to construct a homomorphism \mathbf{w} of H into G by regarding \mathcal{T} as a rooted tree with root R , for a fixed $R \in \mathcal{F}$, and taking a fixed order $R = X_0, X_1, X_2, \dots, X_t$ of all sets in \mathcal{F} in which no child precedes its parent. We then proceed as follows:

1. Take $\mathbf{w}_{H[R]}$, the $H[R]$ -BRW, and set $\mathbf{w} = \mathbf{w}_{H[R]}$ on R . Mark R as explored.
2. Suppose \mathbf{w} is defined on $Z_{i-1} := \bigcup_{j=0}^{i-1} X_j$, $i \geq 1$, and X_i is the first set that has not yet been explored. Take $\mathbf{w}_{H[X_i]}(\cdot|h)$, the $H[X_i]$ -BRW starting from $h = \mathbf{w}|_{X_i \cap Z_{i-1}}$, and set $\mathbf{w} = \mathbf{w}_{H[X_i]}(\cdot|h)$ on $X_i \setminus Z_{i-1}$, provided that there exists a $H[X_i]$ -BRW starting from h . Mark X_i as explored.

If there is no $H[X_i]$ -BRW starting from h , then the process stops and there is no output. (We will presently show that this case cannot occur.)

3. Repeat Step 2 until every set in \mathcal{F} has been explored.

We first show that there is always an output \mathbf{w} and it is a homomorphism from H to G .

Lemma 2.5. *The above process always outputs a homomorphism from H to G . Moreover, the distribution of the homomorphism does not depend on the choice of the order X_1, \dots, X_t of the sets in \mathcal{F} provided that X_0 is a fixed set R .*

Proof. We will prove inductively that \mathbf{w} is a homomorphism from $H[Z_i]$ to G for $i = 0, 1, \dots, t$. For $i = 0$, this follows since G has at least one edge and $H[R]$ is a tree. Hence, the $H[R]$ -BRW is well-defined and yields a homomorphism from $H[R]$ to G . For $i \geq 1$, suppose $\mathbf{w}|_{Z_{i-1}}$ is a homomorphism. Since X_i is a leaf of the subtree $\mathcal{T}_i := \mathcal{T}[\{X_0, X_1, X_2, \dots, X_i\}]$ of \mathcal{T} , every non-trivial path in \mathcal{T}_i containing X_i must contain its parent Y_i and the third property of tree decompositions gives

$$X_i \cap Z_{i-1} = X_i \cap Y_i.$$

Since H is strongly tree-decomposable, there is an isomorphism that fixes $X_i \cap Y_i$ between the minimum subtrees of $H[X_i]$ and $H[Y_i]$ containing $X_i \cap Y_i$. Since $Y_i \subset Z_{i-1}$, the relevant subtree of $H[Y_i]$ has already been embedded by \mathbf{w} . In particular, there is a homomorphism, say g , from the relevant subtree of $H[X_i]$ to G with $g = \mathbf{w}$ on $X_i \cap Y_i$. Hence, by the results of Section 2.2, there exists an $H[X_i]$ -BRW starting from $\mathbf{w}|_{X_i \cap Y_i}$. This shows that \mathbf{w} is a well-defined mapping from $Z_i = Z_{i-1} \cup X_i$ to $V(G)$.

To show that \mathbf{w} is a homomorphism on Z_i , suppose $\{u, v\}$ is an edge of H with $u, v \in Z_i$. Since both $\mathbf{w}|_{Z_{i-1}}$ and $\mathbf{w}|_{X_i}$ are homomorphisms, it is enough to show that either $u, v \in Z_{i-1}$ or $u, v \in X_i$, though not necessarily exclusively. If not, we may assume that $u \in X_i$ and $v \in X_j \subset Z_{i-1}$ for some $j < i$. We now take a set $U \in \mathcal{F}$ that contains both u and v . Such a set exists by the second property of tree decompositions. If U is a descendant of X_i in \mathcal{T} , then the path from U to X_j in \mathcal{T} must go through X_i . As $v \in U \cap X_j$, we have $v \in X_i$ and hence $u, v \in X_i$, contradicting our assumption. If U is not a descendant of X_i in \mathcal{T} , then the path from X_i to U in \mathcal{T} must contain Y_i , the parent of X_i . Since $u \in X_i \cap U$, we have $u \in Y_i \subset Z_{i-1}$ and hence $u, v \in Z_{i-1}$, again contradicting our assumption.

Therefore, the process constructing \mathbf{w} continues until \mathbf{w} has been extended to all sets in \mathcal{F} . Since $\bigcup_{X \in \mathcal{F}} X = V(H)$ (by the first property of tree decompositions), \mathbf{w} is a homomorphism from H to G . Moreover, using

$$\mathbf{P}[\mathbf{w} = h] = \mathbf{P}[\mathbf{w} = h \text{ on } X_0] \prod_{i=1}^t \mathbf{P}[\mathbf{w} = h \text{ on } X_i \mid \mathbf{w} = h \text{ on } Z_{i-1}]$$

and

$$\mathbf{P}[\mathbf{w} = h \text{ on } X_i \mid \mathbf{w} = h \text{ on } Z_{i-1}] = \mathbf{P}[\mathbf{w} = h \text{ on } X_i \mid \mathbf{w} = h \text{ on } X_i \cap Y_i],$$

we easily see that if C_Y is the set of children of Y in \mathcal{T} , then

$$\begin{aligned} \mathbf{P}[\mathbf{w} = h] &= \mathbf{P}[\mathbf{w} = h \text{ on } X_0] \prod_{i=1}^t \mathbf{P}[\mathbf{w} = h \text{ on } X_i \mid \mathbf{w} = h \text{ on } X_i \cap Y_i] \\ &= \mathbf{P}[\mathbf{w} = h \text{ on } X_0] \prod_{Y \in \mathcal{F}} \prod_{X \in C_Y} \mathbf{P}[\mathbf{w} = h \text{ on } X \mid \mathbf{w} = h \text{ on } X \cap Y]. \end{aligned}$$

Similarly, for any other order $R = X_0 = X'_0, X'_1, X'_2, \dots, X'_t$ in which no child precedes its parent,

$$\mathbf{P}[\mathbf{w}' = h] = \mathbf{P}[\mathbf{w}' = h \text{ on } X'_0] \prod_{Y \in \mathcal{F}} \prod_{X \in \mathcal{C}_Y} \mathbf{P}[\mathbf{w}' = h \text{ on } X \mid \mathbf{w}' = h \text{ on } X \cap Y].$$

On $X_0 = X'_0 = R$, both \mathbf{w} and \mathbf{w}' are $H[R]$ -BRWs and, hence,

$$\mathbf{P}[\mathbf{w}' = h \text{ on } X'_0] = \mathbf{P}[\mathbf{w} = h \text{ on } X_0].$$

Further, for all adjacent pairs $X, Y \in \mathcal{F}$, conditioned on $\mathbf{w}' = \mathbf{w} = h$ on $X \cap Y$, both $\mathbf{w}|_X$ and $\mathbf{w}'|_X$ are $H[X]$ -BRWs starting from h and

$$\mathbf{P}[\mathbf{w}' = h \text{ on } X \mid \mathbf{w}' = h \text{ on } X \cap Y] = \mathbf{P}[\mathbf{w} = h \text{ on } X \mid \mathbf{w} = h \text{ on } X \cap Y].$$

Thus, $\mathbf{P}[\mathbf{w} = h] = \mathbf{P}[\mathbf{w}' = h]$, which means that the distribution of \mathbf{w} is independent of the order. \square

In the following proposition, we show that the choice of root R is also irrelevant.

Proposition 2.6. *The distribution of the homomorphism \mathbf{w} described above does not depend on the choice of root R . In particular, $\mathbf{w}|_X$ is the $H[X]$ -BRW for each $X \in \mathcal{F}$, i.e.,*

$$\mathbf{P}[\mathbf{w} = h \text{ on } X] = p_{H[X]}(h|_X)$$

for $h \in \text{Hom}(H, G)$.

Proof. Since X in the second part may be chosen as the root of \mathcal{T} , the equality easily follows from the first part. To prove the first part, let \mathbf{w}_1 and \mathbf{w}_2 be the random homomorphisms constructed using the above procedure with roots $Y_1, Y_2 \in \mathcal{F}$, respectively. We will show that \mathbf{w}_1 and \mathbf{w}_2 have the same distribution whenever Y_1 and Y_2 are adjacent in \mathcal{T} . The general case follows easily from repeated application of this special case.

Since the order of sets in \mathcal{F} other than the root does not affect the distributions of \mathbf{w}_1 and \mathbf{w}_2 , we may embed Y_2 immediately after Y_1 when Y_1 is the root. If, on the other hand, Y_2 is the root, then Y_1 is to be embedded immediately after Y_2 . Order all remaining sets in \mathcal{F} in the exact same order for both cases. It therefore suffices to prove that, for each $h \in \text{Hom}(H, G)$,

$$\begin{aligned} & \mathbf{P}[\mathbf{w}_1 = h \text{ on } Y_1] \cdot \mathbf{P}[\mathbf{w}_1 = h \text{ on } Y_2 \mid \mathbf{w}_1 = h \text{ on } Y_1 \cap Y_2] \\ &= \mathbf{P}[\mathbf{w}_2 = h \text{ on } Y_2] \cdot \mathbf{P}[\mathbf{w}_2 = h \text{ on } Y_1 \mid \mathbf{w}_2 = h \text{ on } Y_1 \cap Y_2]. \end{aligned} \quad (5)$$

We will first show that if S_1 and S_2 are the minimum subtrees of $H[Y_1]$ and $H[Y_2]$ containing $Y_1 \cap Y_2$, respectively, then

$$\mathbf{P}[\mathbf{w}_1 = h \text{ on } Y_2 \mid \mathbf{w}_1 = h \text{ on } Y_1 \cap Y_2] = \frac{\mathbf{P}[\mathbf{w}_2 = h \text{ on } Y_2]}{\sum_{g \in \mathcal{G}_h(S_2)} p_{S_2}(g)},$$

and, similarly,

$$\mathbf{P}[\mathbf{w}_2 = h \text{ on } Y_1 \mid \mathbf{w}_2 = h \text{ on } Y_1 \cap Y_2] = \frac{\mathbf{P}[\mathbf{w}_1 = h \text{ on } Y_1]}{\sum_{g \in \mathcal{G}_h(S_1)} p_{S_1}(g)},$$

where $\mathcal{G}_h(S_i) = \{g \in \text{Hom}(S_i, G) : g = h \text{ on } Y_1 \cap Y_2\}$. The first equation follows from noting that

$$\begin{aligned} \mathbf{P}[\mathbf{w}_1 = h \text{ on } Y_2 \mid \mathbf{w}_1 = h \text{ on } Y_1 \cap Y_2] &= \mathbf{P}[\mathbf{w}_1 = h \text{ on } S_2 \mid \mathbf{w}_1 = h \text{ on } Y_1 \cap Y_2] \\ &\quad \times \mathbf{P}[\mathbf{w}_1 = h \text{ on } Y_2 \mid \mathbf{w}_1 = h \text{ on } S_2] \\ &= \frac{p_{S_2}(h|_{S_2})}{\sum_{g \in \mathcal{G}_h(S_2)} p_{S_2}(g)} \mathbf{P}[\mathbf{w}_1 = h \text{ on } Y_2 \mid \mathbf{w}_1 = h \text{ on } S_2], \end{aligned}$$

and

$$\begin{aligned} \mathbf{P}[\mathbf{w}_2 = h \text{ on } Y_2] &= \mathbf{P}[\mathbf{w}_2 = h \text{ on } S_2] \mathbf{P}[\mathbf{w}_2 = h \text{ on } Y_2 \mid \mathbf{w}_2 = h \text{ on } S_2] \\ &= p_{S_2}(h|_{S_2}) \mathbf{P}[\mathbf{w}_1 = h \text{ on } Y_2 \mid \mathbf{w}_1 = h \text{ on } S_2], \end{aligned}$$

where, in the last equality, we used the fact that \mathbf{w}_1 and \mathbf{w}_2 are constructed in the same way once their values on S_2 are determined to be h . The same argument also works for the second equation.

The desired equality (5) follows since

$$\sum_{g \in \mathcal{G}_h(S_1)} p_{S_1}(g) = \sum_{g \in \mathcal{G}_h(S_2)} p_{S_2}(g),$$

which, in turn, easily follows from the fact that there is an isomorphism from S_1 to S_2 that fixes $Y_1 \cap Y_2$. \square

We are now in a position to prove Theorem 1.2.

Proof of Theorem 1.2. Let H be a strongly tree-decomposable graph with tree decomposition $(\mathcal{F}, \mathcal{T})$ and let G be a graph with at least one edge. Using

$$H(\mathbf{w}) = H(\mathbf{w}(R)) + \sum_{i=1}^t H(\mathbf{w}(X_i) \mid \mathbf{w}(Z_{i-1}))$$

and, from Lemma 2.1(ii),

$$H(\mathbf{w}(X_i) \mid \mathbf{w}(Z_{i-1})) = H(\mathbf{w}(X_i) \mid \mathbf{w}(X_i \cap Z_{i-1})) = H(\mathbf{w}(X_i) \mid \mathbf{w}(X_i \cap Y_i)),$$

where Y_i is the parent of X_i in the tree \mathcal{T} rooted at R , we have that

$$H(\mathbf{w}) = H(\mathbf{w}(R)) + \sum_{i=1}^t H(\mathbf{w}(X_i) \mid \mathbf{w}(X_i \cap Y_i)) = H(\mathbf{w}(R)) + \sum_{Y \in \mathcal{F}} \sum_{X \in C_Y} H(\mathbf{w}(X) \mid \mathbf{w}(X \cap Y)),$$

where C_Y is the set of children of Y . As $H(\mathbf{w}(X)|\mathbf{w}(X \cap Y)) = H(\mathbf{w}(X)) - H(\mathbf{w}(X \cap Y))$, we also have that

$$\begin{aligned} H(\mathbf{w}) &= H(\mathbf{w}(R)) + \sum_{Y \in \mathcal{F}} \sum_{X \in C_Y} H(\mathbf{w}(X)) - H(\mathbf{w}(X \cap Y)) \\ &= \sum_{X \in \mathcal{F}} H(\mathbf{w}(X)) - \sum_{XY \in E(\mathcal{T})} H(\mathbf{w}(X \cap Y)) \\ &\geq \sum_{X \in \mathcal{F}} H(\mathbf{w}(X)) - \sum_{XY \in E(\mathcal{T})} |X \cap Y| \log |V(G)|, \end{aligned} \tag{6}$$

where the inequality follows from Lemma 2.1(i). Proposition 2.6 implies that, for each $X \in \mathcal{F}$, the distribution of $\mathbf{w}|_X$ is the $H[X]$ -BRW distribution and, therefore, by Proposition 2.3,

$$H(\mathbf{w}(X)) \geq \log \left(|V(G)|^{|X|} \left(\frac{2|E(G)|}{|V(G)|^2} \right)^{|E(H[X])|} \right) = |E(H[X])| \log \frac{2|E(G)|}{|V(G)|^2} + |X| \log |V(G)|.$$

Since strong tree decomposability implies that

$$\sum_{X \in \mathcal{F}} |E(H[X])| = |E(H)| \quad \text{and} \quad \sum_{X \in \mathcal{F}} |X| - \sum_{XY \in E(\mathcal{T})} |X \cap Y| = |V(H)|,$$

it follows that

$$\begin{aligned} H(\mathbf{w}) &\geq \sum_{X \in \mathcal{F}} |E(H[X])| \log \frac{2|E(G)|}{|V(G)|^2} + \left(\sum_{X \in \mathcal{F}} |X| - \sum_{XY \in E(\mathcal{T})} |X \cap Y| \right) \log |V(G)| \\ &= |E(H)| \log \frac{2|E(G)|}{|V(G)|^2} + |V(H)| \log |V(G)| \\ &= \log \left(|V(G)|^{|V(H)|} \left(\frac{2|E(G)|}{|V(G)|^2} \right)^{|E(H)|} \right). \end{aligned}$$

Since $H(\mathbf{w}) \leq \log |\text{Hom}(H, G)|$, we finally have that

$$\log |\text{Hom}(H, G)| \geq \log \left(|V(G)|^{|V(H)|} \left(\frac{2|E(G)|}{|V(G)|^2} \right)^{|E(H)|} \right),$$

or

$$|\text{Hom}(H, G)| \geq |V(G)|^{|V(H)|} \left(\frac{2|E(G)|}{|V(G)|^2} \right)^{|E(H)|},$$

as desired. □

3 Subdivisions

We say that a graph G is (ρ, d) -dense if G has density at least d on every $U \subseteq V(G)$ with $|U| \geq \rho|V(G)|$. A beautiful conjecture of Kohayakawa, Nagle, Rödl, and Schacht [8] states that for any graph H and any $\gamma, d > 0$, there should exist $\rho > 0$ such that any (ρ, d) -dense graph G on a sufficiently large number of vertices contains at least $(1 - \gamma)|V(G)|^{|V(H)|} d^{|E(H)|}$ labeled copies of H .

If H is a bipartite graph which satisfies Sidorenko's conjecture, it clearly satisfies this conjecture with $\rho = 1$. However, the conjecture is also known for some non-bipartite graphs, including complete graphs, complete multipartite graphs, the line graph of the cube [3], and odd cycles [11]. Here we show that any graph for which this conjecture holds may be subdivided to produce a graph that satisfies Sidorenko's conjecture.

Theorem 3.1. *Suppose that H is a graph for which there exists a function $f : (0, 1] \rightarrow (0, 1]$ and a constant $a > 0$ such that any $(f(d), d)$ -dense graph G contains at least*

$$a|V(G)|^{|V(H)|}d^{|E(H)|}$$

copies of H . Then the subdivision of H has Sidorenko's property.

We note that both parts of Theorem 1.3 follow easily from this theorem and the discussion above. It will therefore suffice to prove Theorem 3.1. For our proof, it will be convenient to have a version of the hypothesis that applies to weighted graphs as well as graphs.

Lemma 3.2. *Suppose that H is a graph for which there exists a function $f : (0, 1] \rightarrow (0, 1]$ and a constant $a > 0$ such that any $(f(d), d)$ -dense graph G contains at least*

$$a|V(G)|^{|V(H)|}d^{|E(H)|}$$

copies of H . Then, for any edge-weighted graph $W : V(G) \times V(G) \rightarrow [0, 1]$ such that $\sum_{u, u' \in U} W(u, u') \geq d \binom{|U|}{2}$ for all $U \subseteq V(G)$ with $|U| \geq f(d/2)|V(G)|$,

$$\sum_{x_1, \dots, x_h \in V(G)} \prod_{(i, j) \in E(H)} W(x_i, x_j) \geq \frac{a}{2} |V(G)|^{|V(H)|} d^{|E(H)|}.$$

We omit the proof of this lemma, since it is a standard application of Azuma's inequality (see, for example, Corollary 9.7 in [2]).

The following lemma, due to Fox (see [7]), says that in order to prove that a graph H has Sidorenko's property, it suffices to prove that all graphs whose maximum degree is bounded in terms of the average degree contain roughly the correct number of copies of H .

Lemma 3.3. *Let H be a bipartite graph. If there exists a constant c , depending only on H , such that*

$$|\text{Hom}(H, G)| \geq c|V(G)|^{|V(H)|} \left(\frac{2|E(G)|}{|V(G)|^2} \right)^{|E(H)|}$$

for all graphs G with maximum degree at most $\frac{4|E(G)|}{|V(G)|}$, then H has Sidorenko's property.

In proving Theorem 3.1, we may therefore assume that the maximum degree of G is at most twice the average degree d . We will now argue that any graph of this type contains a large subgraph where the minimum degree is at least $d/4$.

Lemma 3.4. *If G is a graph on n vertices with average degree d and maximum degree at most $2d$, there exists a subgraph G' with at least $2^{-6}n$ vertices such that the minimum degree of G' is at least $d/4$.*

Proof. We will prove the result by removing one vertex at a time to form graphs $G_0 = G, G_1, \dots, G_t$, eventually arriving at the required graph. Since the maximum degree of G is at most $2d$, we see that $d(G_i) \leq 2d$ for all i , where $d(G_i)$ is the average degree of G_i . Suppose now that we have formed the graph G_i , but it does not satisfy the required condition. That is, there is a vertex v_{i+1} of degree less than $d/4$. We remove this vertex to form G_{i+1} . Since we start from a graph of average degree at least d , it follows that $d(G_i) \geq d$ for all i . Note that

$$\begin{aligned} d(G_{i+1}) &= \frac{2|E(G_{i+1})|}{|V(G_{i+1})|} = \frac{2|E(G_i)| - 2d(v_{i+1})}{|V(G_i)| - 1} \geq \frac{2|E(G_i)| - d/2}{|V(G_i)| - 1} \\ &= d(G_i) + \frac{d(G_i) - d/2}{|V(G_i)| - 1} \geq d(G_i) + \frac{d}{2(|V(G_i)| - 1)}. \end{aligned}$$

Therefore,

$$d(G_t) \geq d + \frac{d}{2} \left(\frac{1}{n-1} + \frac{1}{n-2} + \dots + \frac{1}{n-t} \right),$$

which contradicts our observation that $d(G_t) \leq 2d$ when $t = (1 - 2^{-6})n$. Hence, for some $i < t$, G_i satisfied the required condition, completing the proof. \square

The key observation in our proof is the next lemma, which says that if a graph has large minimum degree, then the average codegree, taken over all pairs of vertices, in any large set U is itself large.

Lemma 3.5. *Suppose that G is a graph with n vertices and minimum degree δ . Then, provided $\delta|U| \geq 3n$,*

$$\sum_{\substack{\{u, u'\} \subset U \\ u \neq u'}} d(u, u') \geq \frac{\delta^2}{2n} \binom{|U|}{2},$$

where $d(u, u')$ is the codegree of u and u' .

Proof. If we write $d_U(v)$ for the number of neighbors of a vertex v in U , we see, by convexity, that

$$\begin{aligned} \sum_{\substack{\{u, u'\} \subset U \\ u \neq u'}} d(u, u') &= \sum_{v \in V(G)} \binom{d_U(v)}{2} \geq n \binom{\sum_v d_U(v)/n}{2} \\ &= n \binom{\sum_{u \in U} d(u)/n}{2} \geq n \binom{\delta|U|/n}{2} \geq \frac{\delta^2}{2n} \binom{|U|}{2}, \end{aligned}$$

as required. \square

We now have all the necessary ingredients for proving Theorem 3.1.

Proof of Theorem 3.1. Let H_1 be the subdivision of H . Suppose that G is a graph with maximum degree at most $\frac{4|E(G)|}{|V(G)|}$. By Lemma 3.3, it will be enough to prove that there exists a constant c , depending only on H , such that

$$|\text{Hom}(H_1, G)| \geq c|V(G)|^{|V(H)|+|E(H)|} \left(\frac{2|E(G)|}{|V(G)|^2} \right)^{2|E(H)|}. \quad (7)$$

Note that we may assume $|V(G)|$ is sufficiently large, since it is easy to see that verifying inequality (7) for a graph G is equivalent to verifying it for a balanced blow-up of G .

By Lemma 3.4, G has a subgraph G' with at least $2^{-6}|V(G)|$ vertices and minimum degree at least $\frac{|E(G)|}{2|V(G)|}$. Now define a weighted graph W on $V(G')$ by $W(u, v) = d(u, v)/|V(G')|$. By Lemma 3.5, the density of W on every set $U \subseteq V(G')$ with $|U| \geq \frac{6|V(G)|}{|E(G)|}|V(G')|$ is at least

$$\frac{1}{\binom{|U|}{2}} \sum_{\substack{\{u, u'\} \subseteq U \\ u \neq u'}} W(u, u') \geq \frac{1}{\binom{|U|}{2}} \sum_{\substack{\{u, u'\} \subseteq U \\ u \neq u'}} \frac{d(u, u')}{|V(G')|} \geq \frac{|E(G)|^2}{8|V(G)|^4}.$$

Taking $d = |E(G)|^2/8|V(G)|^4$, we have that $6|V(G)|/|E(G)| \leq f(d/2)$ for $|V(G)|$ sufficiently large. Therefore, by our assumption on H and Lemma 3.2,

$$\sum_{x_1, \dots, x_h \in V(G')} \prod_{(i, j) \in E(H)} W(x_i, x_j) \geq \frac{a}{2} |V(G')|^{|V(H)|} \left(\frac{|E(G)|^2}{8|V(G)|^4} \right)^{|E(H)|}.$$

Expanding out W , this is the same as

$$\begin{aligned} \sum_{x_1, \dots, x_h \in V(G')} \prod_{(i, j) \in E(H)} d(x_i, x_j) &\geq a |V(G')|^{|V(H)|+|E(H)|} \left(\frac{|E(G)|}{4|V(G)|^2} \right)^{2|E(H)|} \\ &\geq 2^{-6|V(H)|-12|E(H)|} a |V(G')|^{|V(H)|+|E(H)|} \left(\frac{2|E(G)|}{|V(G)|^2} \right)^{2|E(H)|}. \end{aligned}$$

But the first expression counts homomorphisms from H_1 to G' , while the last expression is of the required type (7). This completes the proof. \square

As an application of Theorem 3.1, we may prove the following more general theorem. We define the $K_{2,t}$ -replacement of a graph H to be the graph obtained from H by replacing each of its edges with a copy of $K_{2,t}$, identifying the edge's endpoints with the two vertices on one side of the corresponding copy of $K_{2,t}$. Note that a $K_{2,1}$ -replacement is equivalent to a subdivision.

Theorem 3.6. *Suppose that H is a graph for which there exists an increasing function $f : (0, 1] \rightarrow (0, 1]$ and a constant $a > 0$ such that any $(f(d), d)$ -dense graph G contains at least*

$$a |V(G)|^{|V(H)|} d^{|E(H)|}$$

copies of H . Then the $K_{2,t}$ -replacement of H has Sidorenko's property.

Proof. For a graph G and an ordered sequence of $h = |V(H)|$ vertices x_1, \dots, x_h , let $H_1(x_1, \dots, x_h)$ be the number of homomorphisms from the subdivision H_1 of H to G such that the i th vertex of H is mapped to x_i . By Theorem 3.1,

$$\sum_{x_1, \dots, x_h} H_1(x_1, \dots, x_h) \geq |V(G)|^{|V(H)|+|E(H)|} \left(\frac{2|E(G)|}{|V(G)|^2} \right)^{2|E(H)|}.$$

Therefore, by convexity,

$$\begin{aligned} \sum_{x_1, \dots, x_h} H_1(x_1, \dots, x_h)^t &\geq |V(G)|^{|V(H)|} \left(\frac{\sum_{x_1, \dots, x_h} H_1(x_1, \dots, x_h)}{|V(G)|^{|V(H)|}} \right)^t \\ &\geq |V(G)|^{|V(H)|+t|E(H)|} \left(\frac{2|E(G)|}{|V(G)|^2} \right)^{2t|E(H)|}. \end{aligned}$$

Since the first expression counts homomorphisms from the $K_{2,t}$ -replacement of H to G , this completes the proof. \square

It is also possible to mix replacements in certain ways. For example, suppose that H is a triangle and consider

$$X = \mathbb{E}_{x,y,z} d(x,y)^r d(y,z)^s d(z,x)^t.$$

This is a normalised count for the number of homomorphisms to G of the graph formed by replacing the edges of the triangle by a $K_{2,r}$, a $K_{2,s}$, and a $K_{2,t}$, respectively. Note, by symmetry, that X^3 can be expressed in the form

$$X^3 = \mathbb{E}[d(x,y)^r d(y,z)^s d(z,x)^t] \mathbb{E}[d(x,y)^s d(y,z)^t d(z,x)^r] \mathbb{E}[d(x,y)^t d(y,z)^r d(z,x)^s].$$

Hence, by Hölder's inequality with exponent 3,

$$X \geq \mathbb{E}[(d(x,y)d(y,z)d(z,x))^{(r+s+t)/3}] \geq \mathbb{E}[d(x,y)d(y,z)d(z,x)]^{(r+s+t)/3},$$

where the last inequality follows by convexity. Therefore, the graph obtained by replacing the edges of the triangle by a $K_{2,r}$, a $K_{2,s}$, and a $K_{2,t}$ has Sidorenko's property. This argument easily generalises to all complete graphs and some other edge-transitive graphs.

4 Cartesian products with even cycles

Let H, K , and G be graphs. In [7], Kim, Lee, and Lee studied homomorphisms from $H \square K$ to G by relating them to homomorphisms from H to the following auxiliary graph constructed from K and G .

Definition 4.1. *Given graphs K and G , let $\psi_K(G)$ be the graph with vertex set $\text{Hom}(K, G)$ in which two vertices $h_1, h_2 \in \text{Hom}(K, G)$ are adjacent if and only if $h_1(v)$ and $h_2(v)$ are adjacent in G for all $v \in V(K)$.*

The following lemma was proved in [7].

Lemma 4.2. *For all graphs H, K and G , there is a one-to-one mapping between $\text{Hom}(H \square K, G)$ and $\text{Hom}(H, \psi_K(G))$. In particular,*

$$|\text{Hom}(H \square K, G)| = |\text{Hom}(H, \psi_K(G))|.$$

We are now ready to prove Theorem 1.4.

Proof of Theorem 1.4. Let $k \geq 2$ be a given positive integer. It suffices to prove the theorem when H has no isolated vertex, since deleting isolated vertices in H only removes connected components of $H \square J$ each of which are isomorphic to J .

Let G be a graph with n vertices and let

$$p := \frac{2|E(G)|}{n^2} = \frac{|\text{Hom}(K_2, G)|}{n^2}, \quad \alpha := \frac{|\text{Hom}(C_4, G)|}{p^4 n^4}, \quad \beta_k := \frac{|\text{Hom}(C_{2k}, G)|}{p^{2k} n^{2k}},$$

for $k \geq 2$. By Lemma 3.3, we may assume that G has maximum degree at most $2pn$. Since C_4 and C_{2k} have Sidorenko's property, α and β_k are both at least 1. Using that $|\text{Hom}(C_{2k}, G)| = \sum_{i=1}^n \lambda_i^{2k}$ for the eigenvalues λ_i of the adjacency matrix of G , we also have

$$\beta_k p^{2k} n^{2k} \leq \alpha \lambda^{2k-4} p^4 n^4 \quad \text{for } \lambda = \max_i |\lambda_i| \text{ and } k \geq 2.$$

Since λ is at most the maximum degree of G , we know that $\lambda \leq 2pn$. Thus the above yields

$$\beta_k \leq 2^{2k-4} \alpha. \quad (8)$$

Writing $v = |V(H)|$ and $e = |E(H)|$, we apply Lemma 4.2 and then use the fact that H has Sidorenko's property to obtain

$$|\text{Hom}(H \square C_{2k}, G)| = |\text{Hom}(H, \psi_{C_{2k}}(G))| \geq |V(\psi_{C_{2k}}(G))|^v \left(\frac{2|E(\psi_{C_{2k}}(G))|}{|V(\psi_{C_{2k}}(G))|^2} \right)^e.$$

As $|V(\psi_{C_{2k}}(G))| = |\text{Hom}(C_{2k}, G)| = \beta_k p^{2k} n^{2k}$, it follows that

$$|\text{Hom}(H \square C_{2k}, G)| \geq (\beta_k p^{2k} n^{2k})^v \left(\frac{2|E(\psi_{C_{2k}}(G))|}{(\beta_k p^{2k} n^{2k})^2} \right)^e. \quad (9)$$

Lemma 4.2 also yields

$$2|E(\psi_{C_{2k}}(G))| = |\text{Hom}(K_2, \psi_{C_{2k}}(G))| = |\text{Hom}(K_2 \square C_{2k}, G)| = |\text{Hom}(C_{2k}, \psi_{K_2}(G))|.$$

Since C_{2k} has Sidorenko's property, we obtain

$$\begin{aligned} 2|E(\psi_{C_{2k}}(G))| &\geq |V(\psi_{K_2}(G))|^{2k} \left(\frac{2|E(\psi_{K_2}(G))|}{|V(\psi_{K_2}(G))|^2} \right)^{2k} \\ &= \left(\frac{2|E(\psi_{K_2}(G))|}{|\text{Hom}(K_2, G)|} \right)^{2k}. \end{aligned}$$

Applying Lemma 4.2 once more, we have

$$2|E(\psi_{K_2}(G))| = |\text{Hom}(K_2, \psi_{K_2}(G))| = |\text{Hom}(K_2 \square K_2, G)| = |\text{Hom}(C_4, G)| = \alpha p^4 n^4,$$

which, together with $|\text{Hom}(K_2, G)| = 2|E(G)| = pn^2$, yields

$$2|E(\psi_{C_{2k}}(G))| \geq \left(\frac{\alpha p^4 n^4}{pn^2} \right)^{2k} = \alpha^{2k} p^{6k} n^{4k}.$$

Together with (9), this inequality implies that

$$|\text{Hom}(H \square C_{2k}, G)| \geq (\beta_k p^{2k} n^{2k})^v \left(\frac{\alpha^{2k} p^{2k}}{\beta_k^2} \right)^e = \alpha^{2ek} \beta_k^{v-2e} p^{2kv+2ke} n^{2kv}.$$

Since H has no isolated vertex, we have $v - 2e \leq 0$ and hence, by (8), $\beta_k^{v-2e} \geq 2^{(2k-4)(v-2e)} \alpha^{v-2e}$. Therefore, since $\alpha \geq 1$,

$$|\text{Hom}(H \square C_{2k}, G)| \geq 2^{(2k-4)(v-2e)} \alpha^{2ek+(v-2e)} n^{2kv} p^{2kv+2ke} \geq 2^{(2k-4)(v-2e)} n^{2kv} p^{2kv+2ke}.$$

Since $H \square C_{2k}$ has $2kv$ vertices and $2kv + 2ke$ edges, the result now follows from Lemma 3.3. \square

Acknowledgements. We would like to thank Olaf Parczyk for sharing his Master's thesis with us. We would also like to thank Balázs Szegedy for a number of valuable discussions.

References

- [1] D. Conlon, J. Fox, and B. Sudakov, An approximate version of Sidorenko's conjecture, Geom. Funct. Anal. **20** (2010), 1354–1366.
- [2] D. Conlon and W. T. Gowers, Combinatorial theorems in sparse random sets, Ann. of Math. **184** (2016), 367–454.
- [3] D. Conlon, H. Hàn, Y. Person, and M. Schacht, Weak quasi-randomness for uniform hypergraphs, Random Structures Algorithms **40** (2012), 1–38.
- [4] D. Conlon, J. H. Kim, C. Lee, and J. Lee, Sidorenko's conjecture for higher tree decompositions, unpublished note.
- [5] R. Halin, S-functions for graphs, J. Geom. **8** (1976), 171–186.
- [6] H. Hatami, Graph norms and Sidorenko's conjecture, Israel J. Math. **175** (2010), 125–150.
- [7] J. H. Kim, C. Lee, and J. Lee, Two approaches to Sidorenko's conjecture, Trans. Amer. Math. Soc. **368** (2016), 5057–5074.
- [8] Y. Kohayakawa, B. Nagle, V. Rödl, and M. Schacht, Weak regularity and linear hypergraphs, J. Combin. Theory Ser. B **100** (2010), 151–160.
- [9] J. X. Li and B. Szegedy, On the logarithmic calculus and Sidorenko's conjecture, to appear in Combinatorica.
- [10] O. Parczyk, On Sidorenko's conjecture, Master's thesis, Freie Universität, Berlin.
- [11] C. Reiher, Counting odd cycles in locally dense graphs, J. Combin. Theory Ser. B **105** (2014), 1–5.

- [12] N. Robertson and P. D. Seymour, Graph minors III: Planar tree-width, J. Combin. Theory Ser. B **36** (1984), 49–64.
- [13] A. Sidorenko, Inequalities for functionals generated by bipartite graphs, Diskret. Mat. **3** (1991), 50–65 (in Russian), Discrete Math. Appl. **2** (1992), 489–504 (in English).
- [14] M. Simonovits, Extremal graph problems, degenerate extremal problems and super-saturated graphs, in Progress in graph theory (Waterloo, Ont., 1982), 419–437, Academic Press, Toronto, ON, 1984.
- [15] B. Szegedy, An information theoretic approach to Sidorenko’s conjecture, arXiv:1406.6738v3 [math.CO].