

Finite reflection groups and graph norms

David Conlon*

Joonkyung Lee†

Abstract

Given a graph H on vertex set $\{1, 2, \dots, n\}$ and a function $f : [0, 1]^2 \rightarrow \mathbb{R}$, define

$$\|f\|_H := \left| \int \prod_{ij \in E(H)} f(x_i, x_j) d\mu^{|V(H)|} \right|^{1/|E(H)|},$$

where μ is the Lebesgue measure on $[0, 1]$. We say that H is *norming* if $\|\cdot\|_H$ is a semi-norm. A similar notion $\|\cdot\|_{r(H)}$ is defined by $\|f\|_{r(H)} := \| \|f\|_H \|$ and H is said to be *weakly norming* if $\|\cdot\|_{r(H)}$ is a norm. Classical results show that weakly norming graphs are necessarily bipartite. In the other direction, Hatami showed that even cycles, complete bipartite graphs, and hypercubes are all weakly norming. We demonstrate that any graph whose edges percolate in an appropriate way under the action of a certain natural family of automorphisms is weakly norming. This result includes all previously known examples of weakly norming graphs, but also allows us to identify a much broader class arising from finite reflection groups. We include several applications of our results. In particular, we define and compare a number of generalisations of Gowers' octahedral norms and we prove some new instances of Sidorenko's conjecture.

1 Introduction

Let H be a graph on vertex set $\{1, 2, \dots, n\}$ and $f : [0, 1]^2 \rightarrow \mathbb{R}$ be a bounded Lebesgue measurable function. Consider the integral

$$\int \prod_{ij \in E(H)} f(x_i, x_j) d\mu^{|V(H)|}, \quad (1)$$

where μ is the Lebesgue measure on $[0, 1]$. If we choose f so as to model the adjacency matrix of a graph G , the integral above corresponds to the homomorphism density $t_H(G)$ of H in G , which plays a central role in extremal graph theory. Similar expressions also appear naturally in other areas, particularly in statistical physics.

Our concern in this paper will be with the natural question, proposed by Lovász [17], of determining when the integral (1) defines a (semi-)norm. Formally, we say that a graph H is *norming* if the functional defined by

$$\|f\|_H := \left| \int \prod_{ij \in E(H)} f(x_i, x_j) d\mu^{|V(H)|} \right|^{1/|E(H)|} \quad (2)$$

*Mathematical Institute, Oxford OX2 6GG, United Kingdom. E-mail: david.conlon@maths.ox.ac.uk. Research supported by a Royal Society University Research Fellowship and by ERC Starting Grant 676632.

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

is a semi-norm, and *weakly norming* if

$$\|f\|_{r(H)} := \left(\int \prod_{ij \in E(H)} |f(x_i, x_j)| d\mu^{|V(H)|} \right)^{1/|E(H)|} \quad (3)$$

is a norm. As one might expect from the name, it is easy to check that every norming graph is weakly norming. While our focus will usually be on weakly norming graphs, we will often discuss how analogous results can be derived for norming graphs under an extra technical condition.

The first in-depth study of (weakly) norming graphs was undertaken by Hatami [12]. A moment's thought shows that H is necessarily bipartite whenever it is weakly norming, because otherwise it can be zero for strictly positive functions. In [12], Hatami showed that the n -dimensional hypercube Q_n and the complete bipartite graph $K_{m,n}$ are weakly norming. He also observed that the functionals $\|\cdot\|_{C_{2k}}$ correspond to the classical Schatten–von Neumann norms and, hence, even cycles are norming. In addition, Lovász [17] showed that the complete bipartite graph $K_{n,n}$ minus a perfect matching is weakly norming.

We generalise these results, finding a much larger class of (weakly) norming graphs that includes all of the known examples. To give an indication of our results, suppose that k and r are integers with $k \leq r$ and \mathcal{P} is a polytope. Consider the bipartite graph between k -faces and r -faces of \mathcal{P} indicating their incidence. That is, we place an edge between a k -face and an r -face if one contains the other. We call this graph the (k, r) -incidence graph of the polytope \mathcal{P} . We then have the following theorem.

Theorem 1.1. *(k, r) -incidence graphs of regular polytopes are weakly norming.*

For example, in an n -dimensional simplex, the k -faces and r -faces naturally correspond to $(k+1)$ -element and $(r+1)$ -element subsets of $[n]$. Therefore, the (k, r) -incidence graph of an n -simplex is the inclusion graph between $(k+1)$ -sets and $(r+1)$ -sets. In particular, the $(0, 1)$ -incidence graph is the 1-subdivision of K_n , the $(0, n-2)$ -incidence graph is $K_{n,n}$ minus a perfect matching, and the $(0, n-1)$ -incidence graph is the star $K_{1,n}$, which by tensor powering shows that $K_{m,n}$ is also weakly norming. Even cycles C_{2k} are the $(0, 1)$ -incidence graphs of regular k -gons and thus are weakly norming. More generally, by considering the $(0, 1)$ -incidence graph of any regular polytope, such as hypercubes or the icosahedron, we see that their 1-subdivisions are weakly norming.

We prove Theorem 1.1 as a corollary of a more general result showing how weakly norming graphs arise from finite reflection groups. A *finite reflection group* is a finite subgroup of $\mathbf{GL}(n, \mathbb{R})$ generated by a set of reflections across hyperplanes passing through the origin. Those readers who are not familiar with these groups may temporarily assume that, given a finite reflection group W , there exists a distinguished set of generators S known as simple reflections (for further details, see Section 4.1). Fixing a generating set of simple reflections S in a finite reflection group W , let S_1 and S_2 be subsets of S and let W_1 and W_2 be the subgroups of W generated by S_1 and S_2 , respectively. Consider the bipartite graph between the (left-)cosets of W_1 and W_2 , where wW_1 and wW_2 are adjacent for every $w \in W$. We call this graph the $(S_1, S_2; S, W)$ -reflection graph and say that a graph H is a *reflection graph* if it is isomorphic to an $(S_1, S_2; S, W)$ -reflection graph for some suitable choice of parameters. With these definitions, we may now state our main result.

Theorem 1.2. *Reflection graphs are weakly norming.*

This class includes the (k, r) -incidence graphs of regular polytopes, but also provides other simple examples, such as the hypercube (rather than its subdivision) and the graph obtained by replacing each edge of an octahedron with a cycle of length 4 (see Example 4.16 for a more formal

description). It also opens the door to more exotic examples, coming from the exceptional reflection groups E_6 , E_7 , and E_8 . An analogous result also holds for norming graphs, though under a slightly stronger condition.

Theorem 1.3. *Let W be a finite reflection group and let S be a generating set of simple reflections. Then, for any disjoint subsets S_1 and S_2 of S , the $(S_1, S_2; S, W)$ -reflection graph is norming.*

When proving that $\|\cdot\|_{r(H)}$ is a norm, all of the difficulties lie in proving the triangle inequality. Hatami's work in [12] started from the observation that a Hölder-like inequality is equivalent to the triangle inequality for $\|\cdot\|_{r(H)}$. To state his condition, we have to introduce some notation generalising (2) and (3). Let $m = |E(H)|$ and let $\chi : E(H) \rightarrow [m]$ be a (not necessarily proper) edge colouring of H . Consider a family $\mathcal{F} = \{f_1, f_2, \dots, f_m\}$ of bounded measurable functions on $[0, 1]^2$ and define

$$\langle \mathcal{F}; \chi \rangle_H := \int \prod_{e=ij \in E(H)} f_{\chi(e)}(x_i, x_j) d\mu^{|V(H)|}.$$

Note that if $f_i = |f|$ for all $i = 1, 2, \dots, m$, then $\langle \mathcal{F}; \chi \rangle_H = \|f\|_{r(H)}^{|E(H)|}$, while if $f_i = f$, then $|\langle \mathcal{F}; \chi \rangle_H| = \|f\|_H^{|E(H)|}$. Hatami's result now says that the triangle inequality holds for $\|\cdot\|_{r(H)}$ if and only if

$$\langle \mathcal{F}; \chi \rangle_H \leq \prod_{e \in E(H)} \|f_{\chi(e)}\|_{r(H)} \quad (4)$$

for all choices of \mathcal{F} and χ . Furthermore, $\|\cdot\|_H$ is a semi-norm if and only if the analogous inequality obtained by replacing $\|f_{\chi(e)}\|_{r(H)}$ with $\|f_{\chi(e)}\|_H$ holds.

We will think of (4) in the following terms: regard the functions f_1, \dots, f_m as m distinct graphs on the same vertex set¹ and imagine each edge of f_i has the colour i . Then $\langle \mathcal{F}; \chi \rangle_H$ is the number of (homomorphic) copies of H which are coloured according to χ , i.e., each edge $e \in E(H)$ receives the colour $\chi(e)$. In particular, if χ is a one-to-one map then $\langle \mathcal{F}; \chi \rangle_H$ counts the number of 'rainbow' copies of H , while $\|f_i\|_{r(H)}^{|E(H)|}$ counts the number of monochromatic copies of H in colour i . Thus, (4) is equivalent to the statement that the number of rainbow copies of H is bounded above by the geometric mean of the number of monochromatic copies in each colour.

The proof that (4) holds for reflection graphs has two steps. In the first step, discussed in Section 3, we show that any graph whose edges percolate in an appropriate way under the action of a certain natural family of automorphisms is weakly norming. This statement, Theorem 3.3, is more general than Theorem 1.2, and may be of independent interest, but the resulting condition needs to be verified by hand for any particular graph or class of graphs. Accordingly, the second step in our proof, discussed in Section 4, is to find a general argument that verifies this condition for all reflection graphs. It is here that we use results from the theory of finite reflection groups.

Suppose now that H is a weakly norming graph, f is a bounded measurable function on $[0, 1]^2$, and e^* is an edge of H . If we put $f_1 = |f|$, $f_2 = f_3 = \dots = f_m = 1$, and make χ one-to-one with $\chi(e^*) = 1$, then $\langle \mathcal{F}; \chi \rangle_H = \|f\|_{r(K_2)}$, where K_2 is just a single edge, so (4) implies that

$$\|f\|_{r(K_2)} \leq \|f\|_{r(H)}. \quad (5)$$

¹When the f_i are non-negative, as they will be when studying weakly norming graphs, there are large graphs approximating f_1, \dots, f_m by the limit theory of dense graphs [18].

That is, when H is weakly norming, H satisfies Sidorenko’s conjecture [22, 23], which says exactly that for any bipartite graph H and any bounded measurable function f , an inequality of the form (5) holds. Sidorenko’s conjecture is one of the major open problems in extremal graph theory, and there has been much recent work [4, 6, 14, 16, 24] verifying the conjecture for a widening class of graphs. As noted above, all of the weakly norming graphs found in this paper also satisfy Sidorenko’s conjecture. However, this is not the only application of our results to Sidorenko’s conjecture. By applying the entropy techniques developed in [16, 14, 24, 6], we will prove that weakly norming graphs can also be used as building blocks for constructing new graphs that satisfy the conjecture. We refer the reader to Section 5.3 for more details.

Finally, also in Section 5, we will discuss generalisations of our results to hypergraphs, with the main result being that a suitably defined family of reflection hypergraphs are weakly norming. We will then show that every norm defined in this manner is equivalent, in some well-defined sense, to a corresponding cut-norm and then to an appropriate variant of the octahedral norms introduced by Gowers [8, 9] in his work on hypergraph regularity. Our methods also allow us to compare the relative strengths of these variants. These results generalise earlier work by Gowers [8, 9] and by Conlon, Hàn, Person, and Schacht [5] on equivalences between cut-norms and octahedral norms and the notions of quasirandomness they define.

2 A motivating example

It is already non-trivial to show that (5) holds, even for graphs as simple as paths [2] or trees [22, 23], so it should not be surprising that it is more difficult to prove the strictly stronger inequality (4). To motivate what follows, we will prove that C_6 , the cycle of length six, is norming without invoking spectral graph theory or Schatten–von Neumann norms. Let $H = C_6$ be the graph with vertex set $\{1, 2, \dots, 6\}$, where i and $i + 1$ are adjacent for all i (with addition taken modulo 6), and let $\mathcal{F} = \{f_1, f_2, \dots, f_6\}$ be a family of six functions, each taking two variables.

To show that (4) holds, we may assume that χ is a one-to-one map, that is, a rainbow colouring, as it is in the most general case. Without loss of generality, put $\chi(e_i) = i$ if e_i is the edge between i and $i + 1$ modulo 6. Define functions g and h by

$$g(x_1, x_4) = \int f_1(x_1, x_2)f_2(x_2, x_3)f_3(x_3, x_4) dx_2dx_3, \text{ and}$$

$$h(x_1, x_4) = \int f_4(x_4, x_5)f_5(x_5, x_6)f_6(x_6, x_1) dx_5dx_6,$$

so that $\int gh dx_1dx_4 = \langle \mathcal{F}; \chi \rangle_{C_6}$. The function g gives the (normalised) count of rainbow walks from x_1 to x_4 which are coloured, in order, with the colours 1, 2, and 3. The function h can be similarly interpreted, but using the colours 6, 5, and 4. A simple application of the Cauchy–Schwarz inequality gives

$$\langle \mathcal{F}; \chi \rangle_{C_6}^2 \leq \int g^2 \int h^2, \tag{6}$$

where here and throughout the paper we suppress the variables of integration if they are clear from context. Observe now that

$$\int g^2 = \langle \mathcal{F}; \chi_L \rangle_{C_6},$$

where χ_L is the colouring obtained by doubling the ‘left half’ of the rainbow cycle, represented by the vector $(1, 2, 3, 3, 2, 1)$. A similar equation also holds for h , namely, $\int h^2 = \langle \mathcal{F}; \chi_R \rangle_{C_6}$, where

$\chi_R = (6, 5, 4, 4, 5, 6)$ is the colouring obtained by doubling the ‘right half’ of the rainbow cycle. We have therefore bounded $\langle \mathcal{F}; \chi \rangle_{C_6}$ from above by the geometric mean of two functions of the same form, but simpler in the sense that they both contain fewer colours. Repeating this procedure twice more, by first applying the Cauchy–Schwarz inequality to each of $\langle \mathcal{F}; \chi_L \rangle_{C_6}$ and $\langle \mathcal{F}; \chi_R \rangle_{C_6}$ with respect to the variables (x_2, x_5) and then applying the Cauchy–Schwarz inequality to each of the four resulting terms with respect to the variables (x_3, x_6) , we find that

$$\langle \mathcal{F}; \chi \rangle_{C_6} \leq \prod_{i=1}^8 \langle \mathcal{F}; \chi_i \rangle_{C_6}^{1/8},$$

where

$$\begin{aligned} \chi_1 &= (1, 1, 1, 1, 1, 1), & \chi_2 &= (2, 1, 1, 2, 2, 2), & \chi_3 &= (2, 2, 2, 2, 3, 3), & \chi_4 &= (3, 3, 3, 3, 3, 3), \\ \chi_5 &= (6, 6, 6, 6, 6, 6), & \chi_6 &= (5, 6, 6, 5, 5, 5), & \chi_7 &= (5, 5, 5, 5, 4, 4), & \chi_8 &= (4, 4, 4, 4, 4, 4). \end{aligned} \quad (7)$$

Recall that $\langle \mathcal{F}; \chi_1 \rangle_{C_6} = \|f_1\|_{C_6}^6$ and hence it gives one of the terms in the desired upper bound, though with an incorrect exponent.

We now iterate this 3-step process of applying the Cauchy–Schwarz inequality along different vertex cuts, noting that the more often we repeat the process the more monochromatic colourings appear in our upper bound. In particular, as in (7), after each repetition, the sum of the exponents of the non-monochromatic forms reduces by at least a half. Hence, if we iterate k times, we get the bound

$$\langle \mathcal{F}; \chi \rangle_{C_6} \leq \prod_{i=1}^{2^{3k}} \langle \mathcal{F}; \chi_{i,k} \rangle_{C_6}^{1/2^{3k}},$$

where at least a $1 - 2^{-k}$ proportion of the 2^{3k} colourings $\chi_{i,k}$ are monochromatic. We may rewrite this inequality as

$$\langle \mathcal{F}; \chi \rangle_{C_6} \leq \prod_{i=1}^{6^6} \langle \mathcal{F}; \chi'_i \rangle_{C_6}^{\alpha_{i,k}},$$

where each χ'_i represents one of the 6^6 possible edge-colourings of C_6 with 6 colours. In particular, we assume that $\chi'_1, \chi'_2, \dots, \chi'_6$ represent the monochromatic colourings in $1, 2, \dots, 6$, respectively. By the argument above, $\alpha_{1,k} + \dots + \alpha_{6,k} \geq 1 - 2^{-k}$ and it is also easy to see that $\alpha_{i,k}$ is non-decreasing in k for all $i = 1, 2, \dots, 6$. Therefore, taking the limit as k tends to infinity, we have

$$\langle \mathcal{F}; \chi \rangle_{C_6} \leq \prod_{i=1}^6 \langle \mathcal{F}; \chi'_i \rangle_{C_6}^{\alpha_i},$$

where $\alpha_1 + \alpha_2 + \dots + \alpha_6 = 1$. If $\alpha_1, \alpha_2, \dots, \alpha_6$ were equal, this would be the desired inequality. If they are not equal, we note that an analogous procedure, but applying the Cauchy–Schwarz inequalities first with (x_2, x_5) , then with (x_3, x_6) , and finally with (x_1, x_4) , allows one to prove the inequality

$$\langle \mathcal{F}; \chi \rangle_{C_6} \leq \prod_{i=1}^6 \langle \mathcal{F}; \chi'_i \rangle_{C_6}^{\beta_i},$$

where $\beta_i = \alpha_{i-1}$ (and addition is again taken modulo 6). Repeating the same idea four more times, we can cyclically permute the exponents in the inequality to all six possible positions. Taking the product of these six inequalities then completes the proof.

3 Cut involutions and Cauchy–Schwarz inequalities

There were three steps to the proof given in the previous section: firstly, we showed how to apply the Cauchy–Schwarz inequality along certain vertex cuts; then we showed that a monochromatic edge-colouring of H can be obtained through a sequence of such Cauchy–Schwarz inequalities; finally, we used a limiting argument and the edge transitivity of C_6 to complete the proof. In this section, we generalise these arguments. The first step will be generalised through the use of a natural class of graph automorphisms which we call cut involutions, while an appropriate generalisation of the last step is relatively straightforward. Generalising the second step, that is, finding a monochromatic edge-colouring, proves more difficult, so in this section we reduce it to a simpler question which we will resolve in Section 4 for the special case of reflection graphs.

To generalise the Cauchy–Schwarz inequality (6), we need to find a vertex cut which splits the graph into two isomorphic pieces, which we refer to, informally, as the left and right halves. To be more precise, given a connected graph H , say that an automorphism ϕ of H is a *cut involution* if it is an involution, i.e., $\phi = \phi^{-1}$, and the fixed point set $F_\phi = \{v \in V(H) : \phi(v) = v\}$ is a vertex cut of H . Let W_H be the subgroup of the automorphism group of H generated by the cut involutions and call it the *cut involution group* of H .

Since it is possible for $H \setminus F_\phi$ to have more than two components, the left and right halves of a cut involution may not be well-defined. In reflection graphs, we can avoid this ambiguity by focusing on certain special cut involutions. For now, we will simply assume that the left and right halves are somehow fixed for each cut involution ϕ . That is, we will assume that there is a fixed pair of disjoint subsets L_ϕ and R_ϕ of $V(H) \setminus F_\phi$ that are unions of connected components in $H \setminus F_\phi$ and which are mapped to each other under ϕ .

Define the *left-folding map* $\phi^+ : V(H) \rightarrow V(H)$ of a cut involution ϕ by

$$\phi^+(v) = \begin{cases} \phi(v) & \text{if } v \in R_\phi \\ v & \text{if } v \in L_\phi \cup F_\phi, \end{cases}$$

and, similarly, define the *right-folding map* ϕ^- of ϕ by swapping the roles of L_ϕ and R_ϕ . Let $m = |E(H)|$, $\mathcal{F} = \{f_1, f_2, \dots, f_m\}$ be a family of non-negative bounded measurable functions on $[0, 1]^2$, and χ be an edge-colouring of H with colours $\{1, 2, \dots, m\}$. Then the Cauchy–Schwarz inequality (6) can be generalised in terms of a cut involution ϕ by

$$\langle \mathcal{F}; \chi \rangle_H \leq \langle \mathcal{F}; \chi \circ \phi^+ \rangle_H^{1/2} \langle \mathcal{F}; \chi \circ \phi^- \rangle_H^{1/2}. \quad (8)$$

Here $\chi \circ \phi^+$ is the colouring of H where the colour of the edges in $R_\phi \cup F_\phi$ copies the corresponding edges in $L_\phi \cup F_\phi$. Similarly, $\chi \circ \phi^-$ is the colouring where the colour of the edges in $L_\phi \cup F_\phi$ copies the corresponding edges in $R_\phi \cup F_\phi$.

Note that in (8), the non-negativity of the functions in \mathcal{F} is used only if there is an edge inside F_ϕ . More precisely, if the fixed point set F_ϕ contains edges, then we are using a Cauchy–Schwarz inequality of the form

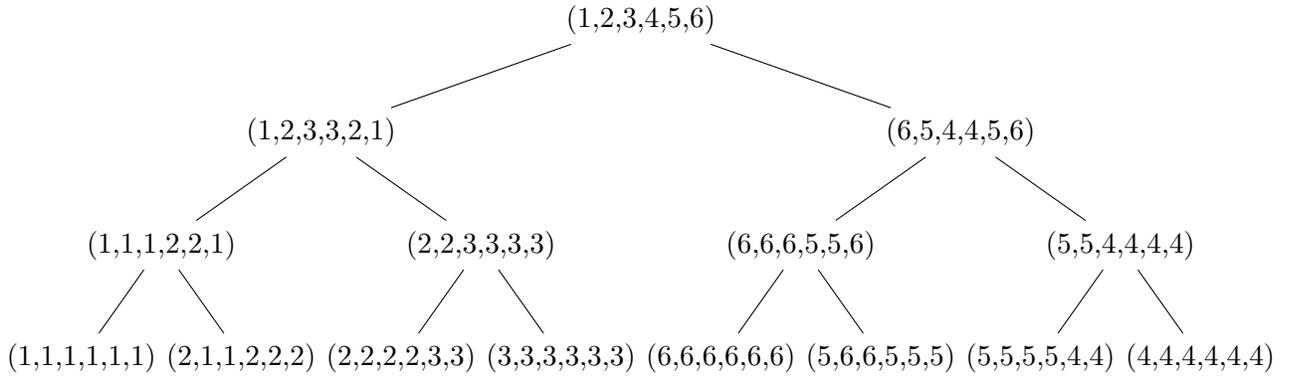
$$\int hg_1g_2 = \int (h^{1/2}g_1)(h^{1/2}g_2) \leq \int hg_1^2 \int hg_2^2,$$

where $h = \prod_{(u,v) \in H[F_\phi]} f_{\chi(uv)}(x_u, x_v)$. Conversely, if there is no edge of H fixed by ϕ , then the non-negativity of functions in \mathcal{F} is unnecessary. This observation leads naturally to the definition of a *stable involution*. This is a cut involution ϕ such that the fixed point set F_ϕ contains no edge of H , that is, it is an independent (or stable) set in H as well as a vertex cut. The subgroup of the

automorphism group of H generated by the stable involutions is then called the *stable involution group* of H and is denoted S_H . In what follows, we will show that graphs whose edges percolate under the action of the cut involution group are weakly norming. Using the simple observation above, it is possible to prove analogous results for the stronger norming property by substituting the stable involution group for the cut involution group. We will not always be explicit about this below, choosing to focus instead on cut involutions and the weakly norming property. Nevertheless, it is worth bearing in mind.

Consider now a sequence $\phi_1, \phi_2, \dots, \phi_d$ of cut involutions and suppose that we wish to apply a sequence of Cauchy–Schwarz inequalities of the form (8), first with $\phi = \phi_1$, then with $\phi = \phi_2$, and so on. Let \mathcal{T} be the rooted binary tree of depth d encoding which colourings have been obtained through such iterations: the vertices are labelled by m -edge-colourings of H , the root is labelled by the initial colouring χ_0 , and each vertex at depth $i < d$ labelled with χ , say, has two children with labels $\chi \circ \phi_{i+1}^+$ and $\chi \circ \phi_{i+1}^-$. We call this tree the *Cauchy–Schwarz tree* associated with $(\phi_i)_{i=1}^d$ rooted at χ_0 . As here, we will often abuse notation by identifying a vertex and its label.

Example 3.1. The 3-step process described in the previous section corresponds to the following Cauchy–Schwarz tree of depth 3:



Roughly speaking, a Cauchy–Schwarz tree shows how colours spread under applications of Cauchy–Schwarz inequalities of the form (8). With this terminology, the following theorem generalises part of what was proven in the previous section.

Theorem 3.2. *Let H be an edge-transitive graph with m edges and $(\phi_i)_{i=1}^d$ be a finite sequence of cut involutions of H . Let $\chi : E(G) \rightarrow [m]$ be a rainbow edge-colouring. If the Cauchy–Schwarz tree \mathcal{T} associated with $(\phi_i)_{i=1}^d$ rooted at χ contains a leaf labelled with a monochromatic colouring χ' , then H is weakly norming.*

Proof. Consider the Cauchy–Schwarz tree \mathcal{T} associated with $(\phi_i)_{i=1}^d$ rooted at χ , such that amongst its leaves $\chi_{1,d}, \dots, \chi_{2^d,d}$ there exists a leaf labelled with the monochromatic colouring χ' . By applying (8) repeatedly, we arrive at the upper bound

$$\langle \mathcal{F}; \chi \rangle_H \leq \prod_{i=1}^{2^d} \langle \mathcal{F}; \chi_{i,d} \rangle_H^{1/2^d}.$$

We now iterate the whole d -step process k times. Consider the Cauchy–Schwarz tree \mathcal{T}' of depth dk associated with the repeated sequence rooted at χ . For $i = 1, 2, \dots, m$, let χ'_i be the monochromatic colouring that only uses the colour i . Observe that all descendants of a vertex in \mathcal{T}' that is labelled

with χ'_i are again labelled with χ'_i . Moreover, by our main assumption, every iteration of the d -step process makes at least a $1/2^d$ proportion of the non-monochromatic leaves monochromatic. We may rearrange the upper bound obtained from \mathcal{T}' as

$$\langle \mathcal{F}; \chi \rangle_H \leq \prod_{i=1}^{m^m} \langle \mathcal{F}; \chi'_i \rangle_H^{\alpha_{i,k}},$$

where each χ'_i represents one of the m^m possible edge-colourings of H with m colours. By the remarks above, $\alpha_{1,k} + \dots + \alpha_{m^m,k} \geq 1 - (1 - 1/2^d)^k$ and each $\alpha_{i,k}$ with $1 \leq i \leq m^m$ is non-decreasing in k . Therefore, taking the limit as k tends to infinity, we get

$$\langle \mathcal{F}; \chi \rangle_H \leq \prod_{i=1}^m \langle \mathcal{F}; \chi'_i \rangle_H^{\alpha_i},$$

where $\alpha_1 + \dots + \alpha_m = 1$. It remains to prove that we may take $\alpha_i = 1/m$. Observe that if ψ is an automorphism of H then $\langle \mathcal{F}; \chi \rangle = \langle \mathcal{F}; \chi \circ \psi \rangle$, since the colouring $\chi \circ \psi$ may be seen as the same colouring but with the vertices relabelled. On the other hand, ψ can be regarded as a permutation of the set of colours $\{1, 2, \dots, m\}$: if an edge $e \in E(H)$ receives colour $j = \chi(e)$ under the rainbow colouring χ , then it receives $i = \chi(\psi(e))$ under $\chi \circ \psi$. Repeating the same argument as above, but with a Cauchy–Schwarz tree rooted at $\chi \circ \psi$, we have the inequality

$$\langle \mathcal{F}; \chi \rangle = \langle \mathcal{F}; \chi \circ \psi \rangle \leq \prod_{i=1}^m \langle \mathcal{F}; \chi'_i \rangle_H^{\beta_i},$$

where $\beta_i = \alpha_j$ if $i = \chi \circ \psi \circ \chi^{-1}(j)$. Taking the geometric mean of these upper bounds over all automorphisms ψ , we get

$$\langle \mathcal{F}; \chi \rangle \leq \prod_{i=1}^m \langle \mathcal{F}; \chi'_i \rangle_H^{\gamma_i},$$

where

$$\gamma_i = \frac{1}{|\text{Aut}(H)|} (\alpha_1 |A_{1 \rightarrow i}| + \alpha_2 |A_{2 \rightarrow i}| + \dots + \alpha_m |A_{m \rightarrow i}|).$$

Here $\text{Aut}(H)$ is the group of all automorphisms of H and $A_{j \rightarrow i}$ is the set of automorphisms sending the edge with colour j to the edge with colour i . Note that all $A_{j \rightarrow i}$ have the same size, as they are all cosets of the subgroup $A_{i \rightarrow i}$. Therefore, $\gamma_i = 1/m$ for all i , as required. \square

The theorem above shows that if we want to prove that a graph H is weakly norming, it is sufficient to show that there exists a Cauchy–Schwarz tree with a rainbow root and a monochromatic leaf. Suppose now that J is an edge subset of H and ϕ a cut involution of H . We define two edge sets $J^+(\phi)$ and $J^-(\phi)$ as follows:

$$J^+(\phi) = \{e \in E(H) : \phi^+(e) \in E(J)\} \text{ and } J^-(\phi) = \{e \in E(H) : \phi^-(e) \in E(J)\}.$$

That is, $J^+(\phi)$ is the graph formed by copying the edges of J from the left half onto the right half. Similarly, $J^-(\phi)$ copies the edges from the right half onto the left half. Let J_0, J_1, J_2, \dots be a sequence of edge subsets of H . We say that it is a *folding sequence* in H if, for each $i \geq 0$,

$$J_{i+1} = J_i^+(\phi) \text{ or } J_{i+1} = J_i^-(\phi)$$

for some cut involution ϕ . If a finite folding sequence J_0, J_1, \dots, J_N in a graph H starts from a set J_0 consisting of a single edge and ends with $J_N = E(H)$, then we call it a *percolating sequence*. With this terminology, we may rephrase Theorem 3.2.

Theorem 3.3. *Suppose that H is a graph which is edge-transitive under the cut involution group W_H . Then, if there exists a percolating sequence J_0, J_1, \dots, J_N , H is weakly norming.*

We say that a percolating sequence is a *strong percolating sequence* if every cut involution used in the sequence is a stable involution. The analogue of Theorem 3.3 for the full norming property is then as follows.

Theorem 3.4. *Suppose that H is a graph which is edge-transitive under the stable involution group S_H . Then, if there exists a strong percolating sequence J_0, J_1, \dots, J_N , H is norming.*

The key question now is whether it is possible to percolate over all edges starting from a single edge. Although there is always enough flexibility when choosing between $J_i^+(\phi)$ and $J_i^-(\phi)$ to guarantee that the number of edges does not decrease, this is far from guaranteeing that a percolating sequence exists. Indeed, as we update, we may lose as well as gain edges, making it difficult to keep track of the changes. However, for any known example of a weakly norming graph, one may easily check that there exists a percolating sequence. The main purpose of the next section is to find a common generalisation for these ad hoc arguments.

4 Euclidean embeddings of reflection graphs

4.1 Preliminaries on finite reflection groups

In this subsection, we state some preliminary facts about reflection groups, focusing on the case of finite groups to make the discussion more concise. For more details, we refer the reader to [1] and [13], while those familiar with the basics of Coxeter groups may safely skip this subsection.

Let W be a finite reflection group in $\mathbf{GL}(n, \mathbb{R})$ and let T be the family of reflections in W . Denote by Φ the set of unit vectors orthogonal to the reflection hyperplanes, where each hyperplane gives rise to two vectors α and $-\alpha$. This set of unit vectors Φ is called a *root system* and each element $\alpha \in \Phi$ is called a *root*. Fixing an ordered basis $\{u_1, u_2, \dots, u_n\}$ of \mathbb{R}^n , we say that a root α is *positive* if $\alpha = \sum c_i u_i$ and $c_k > 0$, where k is the smallest index i for which $c_i \neq 0$. Otherwise, a root is said to be *negative*. Clearly, the set Φ^+ of positive roots and Φ^- of negative roots partition Φ and are of equal size. Let Δ be a minimal subset of Φ^+ such that each $\alpha \in \Phi^+$ is a linear combination of positive roots in Δ with non-negative coefficients. Such a minimal subset always exists, since Φ^+ itself already satisfies the condition. We call this Δ a *simple system* and its elements are called *simple roots*. Given a simple root α , the hyperplane orthogonal to α is called a *simple reflection hyperplane*. We denote by s_α the reflection through the simple reflection hyperplane orthogonal to α and refer to such reflections as *simple reflections*.

The following theorem states some important facts about simple systems. For a proof, we refer the reader to Sections 1.3 and 1.5 of [13].

Theorem 4.1. *A simple system Δ has the following properties:*

- (i) *it is unique with respect to Φ^+ ;*
- (ii) *it consists of linearly independent vectors;*
- (iii) *the set of all simple reflections generates W .*

For each positive root α , let H_α be the reflection hyperplane orthogonal to α . Consider the collection \mathcal{C} of connected components of $\mathbb{R}^n \setminus \cup_{\alpha \in \Phi^+} H_\alpha$. Each component C in \mathcal{C} is called an *open chamber* of the reflection group W . Denote by $\bar{\mathcal{C}} = \{\bar{C} : C \in \mathcal{C}\}$ the set of closures of open

chambers, which we call *closed chambers*. Each open chamber consists of those vectors v in \mathbb{R}^n with a certain fixed sign for $\langle v, \alpha \rangle$ for each positive root α . Conversely, if we fix a sign for $\langle v, \alpha \rangle$ for each positive root α , then, provided these choices are consistent, the collection of such vectors is an open chamber. In particular, there is a unique open chamber consisting of all vectors v satisfying $\langle v, \alpha \rangle > 0$ for all positive roots α , since it is possible to have $\langle v, \alpha \rangle > 0$ for all simple roots α and any such vector must have a positive inner product with every positive root. We call this chamber C_0 (or its closure \bar{C}_0) the *fundamental open (or closed) chamber*. Note that the fundamental closed chamber \bar{C}_0 is a cone given by the intersection of closed half-spaces obtained from simple roots:

$$\bar{C}_0 = \bigcap_{\alpha \in \Delta} \{v \in \mathbb{R}^n : \langle v, \alpha \rangle \geq 0\}.$$

In other words, it is a closed cone surrounded by simple reflection hyperplanes, so a point $v \in \bar{C}_0$ must be contained in either a simple reflection hyperplane or the open chamber C_0 .

An important fact is that the action of W on \mathcal{C} is simply transitive. We refer the reader to Sections 1.6 and 1.7 of [13] for a proof.

Theorem 4.2. *The action of W on \mathcal{C} is simply transitive. In particular, the identity is the only element in W that fixes the fundamental open chamber.*

It immediately follows that the action of W on the set $\bar{\mathcal{C}}$ of all closed chambers is also simply transitive. For each reflection $t \in T$, let $H(t)$ be the hyperplane in \mathbb{R}^n defining t and let $D^+(t)$ be the open half-space $\{x \in \mathbb{R}^n : \langle x, \alpha \rangle > 0\}$, where α is the positive root orthogonal to $H(t)$. Let $S \subset T$ be the set of all simple reflections and I be a subset of S . Define the cone $C(I)$ by

$$C(I) := \left(\bigcap_{s \in I} H(s) \right) \cap \left(\bigcap_{s \in S \setminus I} D^+(s) \right).$$

Observe that the fundamental open chamber C_0 is $C(\emptyset)$ and $C(I)$ and $C(J)$ are disjoint whenever $I \neq J$. Moreover, the fundamental closed chamber can be expressed as

$$\bar{C}_0 = \bigcup_{I \subset S} C(I).$$

The simple transitivity of the W -action may be extended as follows:

Proposition 4.3. *The cone $C(I)$ is mapped into the fundamental closed chamber by $w \in W$ if and only if w is in the subgroup of W generated by I . In particular, the stabiliser of $C(I)$ is precisely the subgroup generated by I .*

Although we refer the reader to Theorem 3A8 in [19] for more information, we remark that one direction of the proposition above is easy to see: if $w = s_1 s_2 \cdots s_k$ for $s_i \in I$, then w fixes each point in $C(I)$ since the subspace $\bigcap_{s \in I} H(s)$ is fixed under each s_i . Note, therefore, that if $C(I)$ is mapped into the fundamental closed chamber by an element of W , it is in fact mapped to itself.

Example 4.4. Let $W \cong S_n$ be the reflection group in $\mathbf{GL}(n, \mathbb{R})$ with reflection hyperplanes

$$H_{ij} := \{(x_1, x_2, \dots, x_n) : x_i = x_j\}$$

for $1 \leq i < j \leq n$. Note that the reflection t_{ij} through H_{ij} is the map swapping x_i and x_j . With respect to the standard ordered basis $\{\mathbf{e}_1, \mathbf{e}_2, \dots, \mathbf{e}_n\}$, the set of positive roots is

$$\Phi^+ = \left\{ \frac{1}{\sqrt{2}}(\mathbf{e}_i - \mathbf{e}_j) : 1 \leq i < j \leq n \right\}.$$

The simple system Δ is then given by

$$\Delta = \left\{ \frac{1}{\sqrt{2}}(\mathbf{e}_i - \mathbf{e}_{i+1}) : 1 \leq i \leq n-1 \right\}$$

and each permutation $\sigma \in S_n$ corresponds to an open chamber

$$C_\sigma = \{(x_1, \dots, x_n) : x_{\sigma(1)} < x_{\sigma(2)} < \dots < x_{\sigma(n)}\}.$$

It is clear that W is a simply transitive action on the set $\mathcal{C} = \{C_\sigma : \sigma \in S_n\}$ of all open chambers.

Each $w \in W$ can be expressed as $w = s_1 s_2 \dots s_k$, where $s_1, s_2, \dots, s_k \in S$. Define the *length* $\ell(w)$ of $w \in W$ to be the minimum length k over all such expressions. This purely combinatorial concept of length has a geometric interpretation, which will play a crucial role in the next subsection.

Theorem 4.5. *Let $w \in W$ and $v \in \mathbb{R}^n$ be a point in the fundamental open chamber. For each positive root α and the reflection $t_\alpha \in T$ induced by α , $\ell(t_\alpha w) > \ell(w)$ holds if and only if $\langle w(v), \alpha \rangle$ is positive.*

The theorem above says that $t_\alpha w$ has a greater length than w if and only if $w(v)$ lies on the ‘positive side’ of H_α , i.e., the same side as the positive root α , whereas its image $t_\alpha w(v)$ under the reflection t_α lies on the other side. Observe now that a root α is positive if and only if $\langle v, \alpha \rangle > 0$ whenever v is chosen from the fundamental open chamber. Thus, the ‘positive side’ with respect to the hyperplane H_α is exactly the component of $\mathbb{R}^n \setminus H_\alpha$ containing the fundamental open chamber. To summarise, Theorem 4.5 implies that the length of $t_\alpha w$ becomes greater than w if and only if by applying t_α we move $w(v)$ from the same side of H_α as the fundamental open chamber to the opposite side. Theorem 4.5 is a key fact in the theory of Coxeter groups, so proofs can be found in many places: for instance, Sections 1.6 and 1.7 in [13] or Proposition 4.4.6 in [1].

4.2 The Euclidean embedding

We begin with a motivating example:

Example 4.6. Let H be the 1-subdivision of K_4 . This graph is *biregular*, that is, it is a bipartite graph such that all vertices on the same side of the bipartition have the same degree. In this case, all vertices on one side have degree three, while all vertices on the other side have degree two. One may check that the cut involution group W_H of H is isomorphic to the symmetric group S_4 . Note also that S_4 may be represented as the reflection group of the tetrahedron. There is a natural way to embed H into \mathbb{R}^4 that uses these observations: consider the map $\eta : V(H) \rightarrow \mathbb{R}^4$ which sends the vertices v_1, v_2, v_3 , and v_4 of degree 3 to $\mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_3$, and \mathbf{e}_4 , respectively, and the unique common neighbour u_{ij} of v_i and v_j , $i \neq j$, to the midpoint $\frac{1}{2}(\mathbf{e}_i + \mathbf{e}_j)$ of their images under η . Then every reflection in S_4 becomes a cut involution of H and, conversely, every cut involution is represented by a reflection.

This example shows that some graphs may be embedded in Euclidean space so that cut involutions are represented by genuine reflections. Our aim in this subsection will be to confirm the existence of a similar embedding for every reflection graph.

We will continue to use notation from Section 4.1. That is, we let $W \subset \mathbf{GL}(n, \mathbb{R})$ be a finite reflection group, T be the family of reflections in W , and $S \subset T$ be the set of simple reflections with respect to some fixed ordered basis of \mathbb{R}^n . Fix also subsets S_1 and S_2 of S and let W_1 and W_2 be the subgroups of W generated by S_1 and S_2 , respectively. For brevity in what follows, we will

write H for the $(S_1, S_2; S, W)$ -reflection graph. Recall that this is the bipartite graph between the (left-)cosets of W_1 and W_2 , where wW_1 and wW_2 are adjacent for every $w \in W$.

In Example 4.6, we mapped the vertices of H to points in \mathbb{R}^n so as to visualise the graph more easily. However, to avoid unnecessary ambiguity in the general case, we will instead construct a map η that sends vertices of H to cones in \mathbb{R}^n . Indeed, let η be the map from $V(H)$ to the family $\{wC(I) : w \in W, I \subset S\}$ of cones in \mathbb{R}^n such that, for $i = 1, 2$ and each $w \in W$,

$$\eta(wW_i) := wC(S_i).$$

In particular, $\eta(W_1) = C(S_1)$ and $\eta(W_2) = C(S_2)$, i.e., they are mapped into subcones of the closed fundamental chamber. Since $uW_i = wW_i$ if and only if $w^{-1}u \in W_i$ and the cone $C(S_i)$ is fixed by each $g \in W_i$, we have

$$\eta(uW_i) = uC(S_i) = w(w^{-1}u)C(S_i) = wC(S_i) = \eta(wW_i),$$

and hence the map η is well-defined. We call the map η the *Euclidean embedding* of H .

Example 4.7. It will be instructive to revisit the case where H is the 1-subdivision of K_4 . Following the notation in Example 4.4 with $n = 4$, let $W \cong S_4$ be the reflection group in $\mathbf{GL}(4, \mathbb{R})$ with the reflection t_{ij} swapping x_i and x_j , $1 \leq i < j \leq 4$. Let $S_1 = \{t_{23}, t_{34}\}$ and $S_2 = \{t_{12}, t_{34}\}$. Then the $(S_1, S_2; S, W)$ -reflection graph H is isomorphic to the 1-subdivision of K_4 . The Euclidean embedding η is then given by

$$\begin{aligned} \eta(W_1) &= \{(x_1, x_2, x_3, x_4) : x_1 > x_2 = x_3 = x_4\}, \\ \eta(W_2) &= \{(x_1, x_2, x_3, x_4) : x_1 = x_2 > x_3 = x_4\}, \end{aligned}$$

and the other values are given by permuting some coordinates of $\eta(W_1)$ and $\eta(W_2)$. Observe that W_1 and W_2 are the only pair of adjacent vertices such that both of their images under η are subsets of the fundamental closed chamber $\{(x_1, x_2, x_3, x_4) : x_1 \geq x_2 \geq x_3 \geq x_4\}$. One may also check that the points used in Example 4.6 to embed vertices are ‘typical’ points from the corresponding cones.

The key properties of the Euclidean embedding are spelled out in the following lemma.

Proposition 4.8. *Let H be the $(S_1, S_2; S, W)$ -reflection graph. Then the Euclidean embedding η of H has the following properties:*

- (i) *for two vertices a and a' on the same side of the bipartition of H , $\eta(a) = \eta(a')$ if and only if $a = a'$;*
- (ii) *every closed chamber \bar{C} contains both $\eta(a)$ and $\eta(b)$ as subsets for exactly one edge $ab \in E(H)$;*
- (iii) *there exists a closed chamber \bar{C} containing both $\eta(a)$ and $\eta(b)$ whenever ab is an edge in H .*

Proof. (i) Suppose that $a = wW_1$, $a' = uW_1$ and $\eta(a) = \eta(a')$. By the definition of η , we have $\eta(a) = wC(S_1)$ and $\eta(a') = uC(S_1)$. Since $\eta(a) = \eta(a')$, it follows that uw^{-1} fixes $C(S_1)$. By Proposition 4.3, this implies that $uw^{-1} \in W_1$, and hence $a = a'$.

- (ii) It follows from the construction of η that the fundamental closed chamber \bar{C}_0 contains $\eta(W_1)$ and $\eta(W_2)$. Again by Proposition 4.3, $\eta(W_1)$ is mapped to a cone in \bar{C}_0 by $w \in W$ if and only if $w \in W_1$ and hence every point in $\eta(W_1)$ is mapped to itself. Since η is one-to-one on each side of the bipartition by (i), W_1 and W_2 are the only vertices of H that are mapped to subsets of \bar{C}_0 . Moreover, each closed chamber $w\bar{C}_0$ contains exactly one pair of cones $(\eta(wW_1), \eta(wW_2))$ that are images of an adjacent pair (wW_1, wW_2) , since the action of W is simply transitive on \mathcal{C} and the fundamental closed chamber contains $\eta(W_1)$ and $\eta(W_2)$.

(iii) Since $ab \in E(H)$, there exists $w \in W$ such that $a = wW_1$ and $b = wW_2$. Then $\eta(a)$ and $\eta(b)$ must be contained in $w\bar{C}_0$, where \bar{C}_0 is the fundamental closed chamber. \square

The next corollary highlights the fact that we have identified a class of cut involutions in H , namely, those corresponding to genuine reflections in the Euclidean embedding.

Corollary 4.9. *Let H be the $(S_1, S_2; S, W)$ -reflection graph. Then every reflection $t \in T$ is a cut involution of H .*

Proof. It is straightforward to check that every $t \in T$ acts as an involutory graph automorphism on H . By part (iii) of the proposition above, every edge in H is mapped into a closed chamber by η , and hence there is no edge crossing a reflection hyperplane. \square

One may wonder if the converse of the corollary above also holds, that is, whether every cut involution of a reflection graph H becomes a reflection under the Euclidean embedding. However, this is not true in general, as evidenced by the following example:

Example 4.10. Suppose that $W \cong S_4$, $S = \{t_{12}, t_{23}, t_{34}\}$, and $S_1 = \{t_{23}, t_{34}\}$, as in Example 4.7, but take $S_2 = S$ instead. Then the $(S_1, S_2; S, W)$ -reflection graph is isomorphic to the star $K_{1,4}$. However, it is not possible to render all cut involutions so that they correspond to reflections. This is because $K_{1,4}$ has a cut involution which is a product of two other cut involutions and, therefore, the determinants of all three of the corresponding linear transformations cannot be -1 .

We have shown that every reflection $t \in T$ induces a cut involution on H , mapping the vertex wW_i to the vertex twW_i . Since T generates W , the graph H is edge-transitive under the action of these cut involutions. Thus, in order to apply Theorem 3.3, it remains to show that there exists a percolating sequence. This will be the topic of the next subsection.

4.3 Proof of Theorem 1.2

Proposition 4.8 allows us to lift the notion of folding sequences to a group theoretic setting. Here is an illustrative example:

Example 4.11. Recalling the example of C_6 on vertex set $\{1, 2, \dots, 6\}$ discussed in Section 2, suppose that we have oriented each cut involution in such a way that the component containing the edge $\{1, 2\}$ is on the left. Then the folding sequence starting from the single edge set $J_0 := \{\{1, 2\}\}$ with $J_{i+1} = J_i^+(\phi_i)$ at the i -th step, where the ϕ_i are chosen as in Section 2, is a percolating sequence. Now consider C_6 as the $(S_1, S_2; S, W)$ -reflection graph, where $W = I_2(3) \subset \mathbf{GL}(2, \mathbb{R})$ is the dihedral group, $S = \{s_1, s_2\}$ is a generating set of simple reflections of $I_2(3)$ with angle $\pi/3$ between the reflection lines defining s_1 and s_2 , $S_1 = \{s_1\}$, and $S_2 = \{s_2\}$. Figure 1 shows what η looks like if we label W_1 and W_2 with 1 and 2, respectively: we have chosen a basis of \mathbb{R}^2 in such a way that the fundamental open chamber C_0 , represented by the grey region, is the unique chamber whose closure contains $\eta(1)$ and $\eta(2)$ as subsets and each vertex is mapped to a ray from the origin as marked. The folding sequence described above naturally gives a folding sequence on chambers, allowing us to spread the grey colour on the fundamental chamber across all chambers. For instance, we may spread the colour on the grey chamber in Figure 1 to the chamber containing the edge $\{2, 3\}$ by applying the reflection through the horizontal line through 2 and 5. In this subsection, we will approach the general question in the opposite direction, using the algebraic tools built up over the previous subsections to show how to spread the grey colour across all chambers, and then using this to form a folding sequence for the original graph.

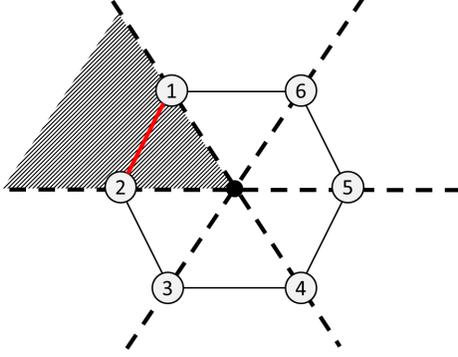


Figure 1: The Euclidean embedding of C_6 .

Let η be the Euclidean embedding of the $(S_1, S_2; S, W)$ -reflection graph H . For each reflection $t \in T$ and the associated positive root α , let t^+ and t^- be the left and right-folding maps defined by

$$t^+(x) := \begin{cases} t(x) & \text{if } \langle x, \alpha \rangle < 0 \\ x & \text{if } \langle x, \alpha \rangle \geq 0 \end{cases} \quad \text{and} \quad t^-(x) := \begin{cases} x & \text{if } \langle x, \alpha \rangle \leq 0 \\ t(x) & \text{if } \langle x, \alpha \rangle > 0, \end{cases}$$

respectively. Given a family \mathcal{J} of open chambers, we set

$$\mathcal{J}^+(t) := \{C \in \mathcal{C} : t^+(C) \in \mathcal{J}\} \quad \text{and} \quad \mathcal{J}^-(t) := \{C \in \mathcal{C} : t^-(C) \in \mathcal{J}\}.$$

Again, as for the graph case, we say that a sequence $\mathcal{J}_0, \mathcal{J}_1, \dots$ of families of open chambers is a *folding sequence* in \mathcal{C} if for each $i \geq 0$ there exists $t \in T$ such that

$$\mathcal{J}_{i+1} = \mathcal{J}_i^+(t) \quad \text{or} \quad \mathcal{J}_{i+1} = \mathcal{J}_i^-(t).$$

Similarly, we say that a finite folding sequence $\mathcal{J}_0, \mathcal{J}_1, \dots, \mathcal{J}_N$ is a *percolating sequence* if it starts with a single chamber \mathcal{J}_0 and ends with the set of all open chambers $\mathcal{J}_N = \mathcal{C}$.

We claim that the existence of a percolating sequence in \mathcal{C} implies the existence of a percolating sequence in the corresponding graph H . Suppose that there exists a folding sequence $\mathcal{J}_0, \dots, \mathcal{J}_N = \mathcal{C}$. By the transitivity of the W -action on \mathcal{C} , we may assume that \mathcal{J}_0 is the fundamental open chamber C_0 . At the i -th step, we say that an open chamber C or its closure \bar{C} is coloured if $C \in \mathcal{J}_i$. As chambers are coloured, we project the colouring down to edges of H by regarding an edge ab as coloured if both a and b are embedded in the closure of a coloured open chamber. This projection respects the folding operations: if the colour on a chamber C spreads to $t(C)$ under the reflection t , then the colour on the edge ab with $\eta(a), \eta(b) \subset \bar{C}$ spreads to the edge $t(a)t(b)$.

Formally, let $J_{\mathcal{J}} \subset E(H)$ be the set of edges ab such that $\eta(a), \eta(b) \subset \bar{C}$ for some $C \in \mathcal{J}$. Recall that each reflection $t \in T$ acts in two different ways: it acts on the set of chambers \mathcal{C} group theoretically and on the vertex set $V(H)$ consisting of all cosets of W_1 and W_2 as a cut involution. To distinguish these two actions, denote by ϕ_t the cut involution of H that corresponds to $t \in T$. We will show a correspondence

$$J_{\mathcal{J}^+(t)} = J_{\mathcal{J}}^+(\phi_t) \quad \text{and} \quad J_{\mathcal{J}^-(t)} = J_{\mathcal{J}}^-(\phi_t) \tag{9}$$

between edge sets and collections of open chambers, provided that each cut involution ϕ_t is oriented so as to guarantee the consistency of signs. By this correspondence, a percolating sequence $\{C_0\} =$

$\mathcal{J}_0, \dots, \mathcal{J}_N = \mathcal{C}$ induces a folding sequence J_0, \dots, J_N with $J_i = J_{\mathcal{J}_i}$. Recall that, by Proposition 4.8 (ii), (W_1, W_2) is the unique pair of vertices both of whose images under η are contained in the fundamental closed chamber. Thus, J_0 must consist of the single edge $\{W_1, W_2\}$. Since we also have $J_N = E(H)$ by Proposition 4.8 (iii), J_0, \dots, J_N is the desired percolating sequence.

It remains to prove the correspondence (9). We orient each cut involution ϕ_t in such a way that

$$L_{\phi_t} := \{x \in V(H) : \eta(x) \subset D^+(t)\} \text{ and } R_{\phi_t} := \{x \in V(H) : \eta(x) \subset D^-(t)\}.$$

For brevity, we write $L_t := L_{\phi_t}$ and $R_t := R_{\phi_t}$. Observe that L_t and R_t are disjoint sets which are mapped to each other by t and the remaining vertices $x \in V(H) \setminus (L_t \cup R_t)$ with $\eta(x) \subset H(t)$ are fixed points of t . With this orientation, the left-folding map ϕ_t^+ sends x to y if and only if t^+ maps $\eta(x)$ onto $\eta(y)$. Therefore, we have

$$\begin{aligned} J_{\mathcal{J}}^+(\phi_t) &= \{ab \in E(H) : \phi_t^+(a)\phi_t^+(b) \in E(J_{\mathcal{J}})\} \\ &= \{ab \in E(H) : t^+(\eta(a) \cup \eta(b)) \subset \bar{C} \text{ for some } C \in \mathcal{J}\} \\ &= J_{\mathcal{J}^+(t)} \end{aligned}$$

and, similarly, $J_{\mathcal{J}}^-(\phi_t) = J_{\mathcal{J}^-(t)}$.

We will now prove the existence of percolating sequences in \mathcal{C} by reducing to a purely group-theoretic framework. Let S be the set of simple reflections in W given by the fixed simple system Δ and let p be a point in the fundamental open chamber. There is a natural one-to-one correspondence between open chambers and elements of W that maps $C \in \mathcal{C}$ to the unique element $w \in W$ such that $w(p) \in C$. We will write $C = wC_0$ if $w(p) \in C$. For example, the identity element of W corresponds to the fundamental open chamber C_0 . For each reflection t , we may define the left and right-folding maps t_+ and t_- on W so that

$$(t_+w)C_0 = t^+(wC_0) \text{ and } (t_-w)C_0 = t^-(wC_0)$$

hold. Recall that $t^+(C) = t(C)$ if and only if $\langle x, \alpha \rangle < 0$ for all $x \in C$, where α is the positive root that induces t . In other words, $t_+w = tw$ if and only if for every $x \in wC_0$ we have $\langle x, \alpha \rangle < 0$. By Theorem 4.5, this is equivalent to the purely algebraic property $\ell(tw) < \ell(w)$. Thus, one may check that

$$t_+w = \begin{cases} tw & \text{if } \ell(tw) < \ell(w) \\ w & \text{if } \ell(tw) \geq \ell(w) \end{cases} \quad \text{and} \quad t_-w = \begin{cases} w & \text{if } \ell(tw) \leq \ell(w) \\ tw & \text{if } \ell(tw) > \ell(w) \end{cases} \quad (10)$$

hold.² For a subset $K \subset W$, let

$$K^+(t) := \{w \in W : t_+w \in K\} \text{ and } K^-(t) := \{w \in W : t_-w \in K\}.$$

For $\mathcal{J} \subset \mathcal{C}$, write $K_{\mathcal{J}}$ for the subset $\{w \in W : wC_0 \in \mathcal{J}\}$ of W . We have a natural correspondence $K_{\mathcal{J}^+(t)} = K_{\mathcal{J}}^+(t)$ and $K_{\mathcal{J}^-(t)} = K_{\mathcal{J}}^-(t)$. Again, we say that a sequence of subsets K_0, K_1, \dots is a folding sequence in W if and only if K_{i+1} equals $K_i^+(t)$ or $K_i^-(t)$ for each $i \geq 0$. Hence, it remains to find a folding sequence $\{e\} = K_0, K_1, \dots, K_N = W$. The existence of such a sequence follows from the next theorem. In fact, it shows something stronger: we may always choose $K^+(s)$ with s a simple reflection to update K .

²In fact, it is impossible to have $\ell(tw) = \ell(w)$. This follows from the simple fact that $w \mapsto \det(w) = (-1)^{\ell(w)}$ is a group homomorphism from W to the multiplicative group $\{\pm 1\}$.

Theorem 4.12. *Let W be a finite reflection group and e the identity of W . Then there exists a finite folding sequence $\{e\} = K_0, K_1, \dots, K_N = W$ such that for each $i = 0, 1, \dots, N - 1$ there is a simple reflection s_i for which $K_{i+1} = K_i^+(s_i)$.*

Our plan for proving Theorem 4.12 is to use induction on the length function ℓ . We say that a subset U of W is a *stack* if whenever $w \in U, s \in S$, and $\ell(sw) < \ell(w)$, sw is also in U .³ The following lemma is the key to proving Theorem 4.12.

Lemma 4.13. *Let U be a stack and $s \in S$ be a simple reflection. Then U is contained in $U^+(s)$.*

Proof. Let $u \in U$ be arbitrary. If $\ell(su) < \ell(u)$, the fact that U is a stack immediately implies that $su \in U$. By (10), we then have $s_+u = su$ and hence $s_+u \in U$, which means $u \in U^+(s)$. Otherwise, $s_+u = u \in U$, which again implies $u \in U^+(s)$. \square

Proof of Theorem 4.12. Let $S = \{s_1, \dots, s_k\}$ be the set of simple reflections, which is a generating set for W . We will prove by induction on L that there is a K_i which contains all $w \in W$ with $\ell(w) \leq L$. The only element with length 0 is the identity, so $K_0 = \{e\}$ satisfies the condition. Suppose now that U_L is the set of all $w \in W$ with $\ell(w) \leq L$ and K_{i_L} contains U_L . Let $U_{L+1,0} = W_{L+1,0} = U_L$ and define, for $i, j = 1, 2, \dots, k$,

$$U_{L+1,i} := \{w \in W : s_i w \in U_L\} \text{ and } W_{L+1,j} := \bigcup_{i=0}^j U_{L+1,i}.$$

Note that $W_{L+1,j}$ is always a stack, because elements of smaller length must be in U_L . Therefore, by Lemma 4.13, $W_{L+1,j}$ is a subset of $W_{L+1,j}^+(s_{j+1})$ for each $j = 0, 1, \dots, k - 1$. Moreover, $U_{L+1,j+1}$ is a subset of $W_{L+1,j}^+(s_{j+1})$ for each $j = 0, 1, \dots, k - 1$. To see this, note that $w \in U_{L+1,j+1}$ if and only if $s_{j+1}w \in U_L$. If $w \in U_L$, w is already contained in $W_{L+1,j}$. We may therefore assume that $w \notin U_L$ and so $\ell(s_{j+1}w) < \ell(w)$. This in turn shows that $(s_{j+1})_+w = s_{j+1}w \in W_{L+1,j}$, which implies that $w \in W_{L+1,j}^+(s_{j+1})$. Putting everything together, we see that $W_{L+1,j+1}$ is a subset of $W_{L+1,j}^+(s_{j+1})$ for each $j = 0, 1, \dots, k - 1$. Since $W_{L+1,k}$ is the set $\{w \in W : \ell(w) \leq L + 1\}$, we may therefore take $i_{L+1} = i_L + k$, completing the induction. \square

Since we have now shown the existence of a percolating sequence in H , this also completes the proof of Theorem 1.2. Theorem 1.1, concerning incidence graphs of regular polytopes, follows as a simple application.

Proof of Theorem 1.1. Let \mathcal{P} be a (realised) regular polytope with symmetry group W and let Ψ be a fixed flag in \mathcal{P} , that is, a maximal chain of faces. It is a folklore fact (see, for example, Theorem 3D7 in [19]) that we may choose a basis of \mathbb{R}^n so that the set of simple reflections S can be enumerated as $\{s_1, \dots, s_n\}$ in such a way that the k -face F_k in Ψ is fixed under all reflections in S but s_{k+1} , i.e., s_{k+1} is the only active mirror for F_k in S . Since \mathcal{P} is regular, there is a one-to-one correspondence between the cosets of the parabolic subgroup generated by $S \setminus \{s_{k+1}\}$ and the k -faces of \mathcal{P} . Hence, it follows that the $(S_1, S_2; S, W)$ -reflection graph with $S_1 = S \setminus \{s_{k+1}\}$ and $S_2 = S \setminus \{s_{r+1}\}$ is isomorphic to the (k, r) -incidence graph of \mathcal{P} . \square

There are many more examples than those coming from regular polytopes. For instance, the vertex-edge incidence graphs of many quasiregular, i.e., vertex and edge-transitive, polytopes can also be written as $(S_1, S_2; S, W)$ -reflection graphs. Moreover, it is possible to obtain interesting examples by considering the orbits of faces of the same rank:

³That is, a stack is a subset of W closed under downward inclusion with respect to the Bruhat (or strong) order, which is the poset structure induced by length. We have not defined the Bruhat order, but we use it implicitly.

Example 4.14. Let D_n be the n -demicube group, i.e., the order 2 subgroup of the hypercube group obtained by removing the ‘perpendicular’ mirrors $x_i = 0$. The action of D_n on the set of vertices of an n -dimensional hypercube is not transitive, but there exist two orbits, each of which corresponds to a colour class of the hypercube when considered as a bipartite graph. There is a generating set of simple reflections $S = \{s_1, s_2, \dots, s_n\}$ for D_n such that the order of $s_1 s_3$ and $s_i s_{i+1}$ for $i \geq 2$ is 3, while all other $s_i s_j$, $i \neq j$, are of order 2. Then one may check that the n -dimensional hypercube is isomorphic to the $(S_1, S_2; S, D_n)$ -reflection graph with $S_1 = S \setminus \{s_1\}$ and $S_2 = S \setminus \{s_2\}$. It follows that the hypercube is weakly norming, a fact first proved in [12].

By Theorem 1.1, we already know that the vertex-edge incidence graph of an octahedron, i.e., the 1-subdivision of an octahedron, is weakly norming, but we may also prove that it is norming through a judicious choice of reflection group.

Example 4.15. Let D_3 be the 3-dimensional demicube group and S the same generating set described in Example 4.14. Take $S_1 = \{s_1, s_2\}$ and $S_2 = \{s_3\}$. Then $(S_1, S_2; S, D_3)$ -reflection graph is isomorphic to the 1-subdivision of an octahedron, as the latter is the face-edge incidence graph of the cube. Therefore, the 1-subdivision of the octahedron is norming.

Let us mention another example of a norming graph. Following [6], we say that a graph is a $K_{2,t}$ -replacement of H if each edge of H is replaced with a copy of $K_{2,t}$ by identifying the two vertices of the edge with the two vertices on the smaller side of $K_{2,t}$.

Example 4.16. Let B_3 be the cube group, that is, the symmetry group of the cube. Then there is a generating set $S = \{s_1, s_2, s_3\}$ such that $s_1 s_2$, $s_2 s_3$, and $s_1 s_3$ are of orders 4, 3, and 2, respectively. Let $S_1 = \{s_1, s_2\}$ and $S_2 = \{s_3\}$. Then one may check that the $(S_1, S_2; S, B_3)$ -reflection graph is isomorphic to the $K_{2,2}$ -replacement of the octahedron graph, and thus is norming.

Finally, we remark that one may presumably use the exceptional reflection groups E_6, E_7 , and E_8 to build some more exotic (weakly) norming graphs, though we have not pursued this further.

5 Generalisations and applications

5.1 Hypergraph norms

As noted in Hatami’s PhD thesis [11], the concepts of norming and weakly norming graphs generalise in the obvious way to hypergraphs, with Gowers’ octahedral norms [8, 9] serving as standard examples. In this short subsection, we discuss the appropriate generalisation of Theorem 1.2.

To obtain a suitable generalisation of Theorem 1.2 to k -uniform hypergraphs, or k -graphs for short, with the same proof strategy, we should first define cut involutions for hypergraphs. One naive way might be to say that an involutory automorphism is a cut involution if the fixed point set is again a vertex cut, i.e., deleting it makes the hypergraph disconnected. However, this is not a good choice. For example, consider the 3-graph H on vertex set $\{a, b, c, x, y, z\}$ with three edges abc, yca , and zab . Let ϕ be the involutory automorphism of H that fixes x and a and maps b and z to c and y , respectively. Deleting the fixed vertex set consisting of x and a obviously makes H disconnected, but we cannot use x and a as ‘pivots’ to apply the Cauchy–Schwarz inequality, since there is an edge crossing the cut. Furthermore, if H were a (weakly) norming hypergraph then it would also satisfy the hypergraph generalisation of Sidorenko’s conjecture, but H was shown to be a counterexample to this conjecture in [23].

To define cut involutions so as to avoid the difficulties discussed in the previous example, we replace each edge in a k -graph H by a clique of size k and write \tilde{H} for the resulting graph on

$V(H)$. Note that every automorphism ϕ of H is also an automorphism of \tilde{H} . We say that an automorphism ϕ of H is a cut involution of H if it is a cut involution of \tilde{H} . With this definition, all of our arguments generalise without difficulty.

Given a finite reflection group W , the set of all simple reflections S , and subsets S_1, S_2, \dots, S_k of S , let the $(S_1, \dots, S_k; S, W)$ -reflection hypergraph be the k -partite k -graph whose parts are the cosets of the subgroup W_i generated by S_i for each $i = 1, \dots, k$, with an edge for every k -tuple of the form $(wW_1, wW_2, \dots, wW_k)$ with $w \in W$. We say that a k -graph is a reflection hypergraph if it is isomorphic to the $(S_1, \dots, S_k; S, W)$ -reflection hypergraph for a suitable choice of parameters. We have the following hypergraph generalisation of Theorem 1.2:

Theorem 5.1. *A k -partite k -graph H is weakly norming whenever it is a reflection hypergraph. Moreover, if there is no mirror containing a hyperedge, i.e., $\bigcap_{i=1}^k S_i = \emptyset$, then H is norming.*

With this framework, we may easily recover Gowers' octahedral norms [8, 9].

Example 5.2. Let $W = W_1 \times W_2 \times \dots \times W_k$, where each W_i is a reflection group generated by a single reflection s_i . Then the set S of simple reflections consists of k orthogonal reflections. If we let $S_i = S \setminus \{s_i\}$ and H_k be the $(S_1, \dots, S_k; S, W)$ -reflection hypergraph, then Theorem 5.1 implies that H_k , consisting of the $(k-1)$ -faces of the k -dimensional octahedron, is a norming k -graph.

Another example of a hypergraph norm comes from the work on weak quasirandomness in [5].

Example 5.3. In [5], a k -graph M_k is constructed recursively as follows. Given a k -partite k -graph M on $A_1 \cup \dots \cup A_k$, we write $\text{db}_i(M)$ for the k -graph obtained by gluing two vertex-disjoint copies of M so that the corresponding vertices in each copy of A_i are identified. Starting from the graph M_0 with a single edge, define

$$M_k := \text{db}_k(\text{db}_{k-1}(\dots \text{db}_1(M_0) \dots)).$$

Letting $S_i = \{s_i\}$ and using the same reflection group W as in Example 5.2, Theorem 5.1 implies that the $(S_1, \dots, S_k; S, W)$ -reflection hypergraph is norming. Since this hypergraph is isomorphic to M_k , we see that M_k defines a semi-norm.

The examples above are very natural: both of them use the simplest reflection group with k generators and the resulting Cauchy–Schwarz trees are very symmetric and easy to analyse. Our new framework gives a much larger class of hypergraph norms, each of which defines a certain notion of quasirandomness. We will discuss the relations between these notions in the next subsection.

5.2 Domination between norms

In this subsection, we will be interested in relations between (hyper)graph norms. To begin, we will study the question of determining whether the (absolute) H -norm dominates another (absolute) J -norm. That is, for any bounded measurable function $f : [0, 1]^2 \rightarrow \mathbb{R}$, we would like to know if

$$\|f\|_J \leq \|f\|_H \quad \left(\text{or } \|f\|_{r(J)} \leq \|f\|_{r(H)} \right) \quad (11)$$

holds. This question remains valid even if H and J are not (weakly) norming, but weakly norming graphs can be regarded as local maxima for such comparisons. To be more precise, whenever H is (weakly) norming and J is a subgraph of H , the (absolute) H -norm dominates the (absolute) J -norm. To see this, let $f_{\chi(e)} = 1$ for all $e \notin E(J)$ in inequality (4).

An immediate consequence of this inequality is Sidorenko's conjecture (5), which essentially states that the absolute H -norm dominates the single-edge norm whenever H is a bipartite graph. Hence, if H is weakly norming, Sidorenko's conjecture holds for H . Similar arguments show that if H is weakly norming and contains a cycle, then it also satisfies the so-called forcing conjecture, a central problem in the study of quasirandom graphs. We refer the interested reader to [4] for further information about this conjecture and its relationship with graph norms.

We remark that one (hyper)graph norm may dominate another even when the first (hyper)graph does not contain the second. For example, if $m \geq n$, then

$$\|f\|_{C_{2m}} \leq \|f\|_{C_{2n}}.$$

Note that the cut involution groups of different even cycles are always non-isomorphic. In general, it is hard to compare (hyper)graph norms coming from non-isomorphic reflection groups, but we can say something if the two (hyper)graphs have isomorphic cut involution groups of a certain type.

Proposition 5.4. *Let $W = W_1 \times W_2 \times \cdots \times W_n$, where each W_i is a reflection group generated by a single reflection s_i , and let $S = \{s_1, \dots, s_n\}$ be the set of simple reflections in W . Suppose S_1, \dots, S_k and S'_1, \dots, S'_k are subsets of S such that $S'_i \subset S_i$ for each $i = 1, 2, \dots, k$. Let H and H' be the $(S_1, \dots, S_k; S, W)$ -reflection hypergraph and the $(S'_1, \dots, S'_k; S, W)$ -reflection hypergraph, respectively. Then both H and H' are norming and the H -norm dominates the H' -norm.*

Our proof is quite heavy on notation, so we postpone it until the end of this subsection and instead give an example that conveys the rough idea of the proof.

Example 5.5. Consider the special case of Proposition 5.4 where $n = 3$, $S_i = S \setminus \{s_i\}$ for $i = 1, 2, 3$, $S'_2 = S_2$, $S'_3 = S_3$, and $S'_1 = \{s_3\}$. We may assume that s_1, s_2 , and s_3 are reflections along the planes $x = 0$, $y = 0$, and $z = 0$, respectively, in \mathbb{R}^3 . We use the standard basis $\{\mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_3\}$ to make the fundamental open chamber the first octant. Then H is the octahedral 3-graph on $\{X_0, X_1, Y_0, Y_1, Z_0, Z_1\}$ with edges $\{X_i Y_j Z_k : i, j, k = 0, 1\}$, where

$$\begin{aligned} X_0 &= \{(x, 0, 0) : x > 0\}, & X_1 &= \{(x, 0, 0) : x < 0\}, \\ Y_0 &= \{(0, y, 0) : y > 0\}, & Y_1 &= \{(0, y, 0) : y < 0\}, \\ Z_0 &= \{(0, 0, z) : z > 0\}, & Z_1 &= \{(0, 0, z) : z < 0\}. \end{aligned}$$

As noted in Example 5.2, the H -norm is the Gowers' octahedral norm for 3-graphs. On the other hand, H' is the 3-graph on eight vertices $\{X_{00}, X_{01}, X_{10}, X_{11}, Y_0, Y_1, Z_0, Z_1\}$ with eight edges $\{X_{ij} Y_i Z_k : i, j, k = 0, 1\}$. Here Y_i, Z_j are exactly the same as in H since $S_2 = S'_2$ and $S_3 = S'_3$, and

$$\begin{aligned} X_{00} &= \{(x, y, 0) : x > 0, y > 0\}, & X_{01} &= \{(x, y, 0) : x < 0, y > 0\}, \\ X_{10} &= \{(x, y, 0) : x > 0, y < 0\}, & X_{11} &= \{(x, y, 0) : x < 0, y < 0\}. \end{aligned}$$

Observe that the induced subgraphs of H on $\{X_0, X_1, Y_i, Z_0, Z_1\}$, for $i = 0, 1$, and the induced subgraphs of H' on $\{X_{j0}, X_{j1}, Y_j, Z_0, Z_1\}$, for $j = 0, 1$, are all isomorphic to the kite-shaped 3-graph with 4 edges, say K . Moreover, H is obtained by gluing two vertex-disjoint copies of K , identifying the copies of X_0, X_1, Z_0 , and Z_1 in each graph. We can build H' similarly but now we only identify the copies of Z_0 and Z_1 . Therefore, setting

$$g(x_0, x_1, z_0, z_1) = \int \prod_{i,j=0,1} f(x_i, y, z_j) dy,$$

whose average is the 4th power of the K -norm, we have

$$\|f\|_H = \int g^2 dx_0 dx_1 dz_0 dz_1 \text{ and } \|f\|_{H'} = \int \left(\int g dx_0 dx_1 \right)^2 dz_0 dz_1.$$

Thus, the Cauchy–Schwarz inequality implies that $\|f\|_H \geq \|f\|_{H'}$. We shall generalise this idea later to prove Proposition 5.4.

Another way to compare (hyper)graph norms is to ask that they be polynomially related, in the sense that if one norm is large, say at least c , then the other norm is at least c^k for some appropriate k . Formally, we say that a (semi-)norm $\|\cdot\|$ *polynomially dominates* another (semi-)norm $\|\cdot\|_0$ if, for any $|f| \leq 1$,

$$\|f\| \leq c \text{ implies } \|f\|_0 \leq c^\delta,$$

where $\delta > 0$ is a constant independent from f . It is also possible for two norms to polynomially dominate each other, so in this case we say that the two norms are *polynomially equivalent*. For the rest of this subsection, we shall state our results only for norming graphs, but analogous statements for weakly norming graphs hold if we replace f by $|f|$ and graph norms by absolute graph norms. Following Gowers' approach [7, 8, 9] to quasirandomness, but borrowing notation from graph limit theory [17, 18], we define the cut-norm $\|f\|_\square$ for a bounded measurable function f on $[0, 1]^2$ by

$$\|f\|_\square := \sup_{u,v} \left| \int f(x,y)u(x)v(y)dx dy \right|,$$

where the supremum is taken over all measurable $u, v : [0, 1] \rightarrow [-1, 1]$. By applying inequality (4) with C_4 , we see that

$$\int f(x,y)u(x)v(y)dx dy \leq \|f\|_{C_4} \|u\|_{C_4} \|v\|_{C_4} \leq \|f\|_{C_4} \tag{12}$$

holds for all measurable $u, v : [0, 1] \rightarrow [-1, 1]$, and hence $\|f\|_\square \leq \|f\|_{C_4}$. Conversely, if

$$\|f\|_{C_4}^4 = \int f(x,y)f(x,y')f(x',y)f(x',y') > c^4,$$

then there exist some $x^*, y^* \in [0, 1]$ such that

$$\left| \int f(x,y)f(x,y^*)f(x^*,y)f(x^*,y^*)dx dy \right| \geq c^4. \tag{13}$$

Taking $u(x) = f(x^*, y^*)f(x, y^*)$ and $v(y) = f(x^*, y)$ gives $\|f\|_\square \geq c^4$, so the C_4 -norm and the cut-norm are polynomially equivalent. Similarly, using (4) to deduce (12) and averaging to obtain (13), it is easy to see that any H -norm is polynomially equivalent to the cut-norm whenever H is norming and contains a cycle. Hence, if two norming graphs H and J contain cycles, then the H -norm and the J -norm are always polynomially equivalent. Conversely, suppose H is a norming graph isomorphic to a tree. Then H must be bi-regular by Theorem 2.10 in [12] and, hence, is isomorphic to a star $K_{1,t}$. By taking f to be the balanced function of a non-quasirandom but regular graph, one may obtain an example with small $K_{1,t}$ -norm but large cut-norm. Hence, star-norms are polynomially dominated by the cut-norm, but not vice versa.

Our aim now is to decide whether two different H -norms, where each H is an $(S_1, \dots, S_k; S, W)$ -reflection hypergraph, are polynomially equivalent. The result below says that we may ‘forget’ the angles between the reflection hyperplanes of a reflection group W if H is an $(S_1, \dots, S_k; S, W)$ -reflection hypergraph and still obtain a polynomially equivalent norm:

Proposition 5.6. *Let W and W' be two reflection groups with the same number of simple reflections, i.e., $S = \{s_1, \dots, s_n\}$ and $S' = \{s'_1, \dots, s'_n\}$ are the set of simple reflections of W and W' , respectively. Suppose, for $S_1, \dots, S_k \subset S$ and $S'_1, \dots, S'_k \subset S'$, that the indices covered by S_i and S'_i are always the same. That is, $\{j : s_j \in S_i\} = \{j : s'_j \in S'_i\}$ for each $i = 1, 2, \dots, k$. Then the $(S_1, \dots, S_k; S, W)$ -reflection hypergraph H and the $(S'_1, \dots, S'_k; S', W')$ -reflection hypergraph H' give polynomially equivalent hypergraph semi-norms, provided that both hypergraphs are norming.*

For example, the proposition above allows us to prove that if a reflection hypergraph H consists of ‘tight triples’, then the H -norm is polynomially equivalent to Gowers’ octahedral norm for 3-graphs.

Example 5.7. In Example 4.7, the reflection group W is isomorphic to the symmetric group on 4 elements and the set of simple reflections $S = \{t_{12}, t_{23}, t_{34}\}$, where t_{ij} is a reflection swapping the i -th and j -th coordinates. Let $S_1 = S \setminus \{t_{12}\}$, $S_2 = S \setminus \{t_{23}\}$, $S_3 = S \setminus \{t_{34}\}$, and H be the $(S_1, S_2, S_3; S, W)$ -reflection hypergraph. Observe that H is a tripartite 3-graph isomorphic to the vertex-edge-face incidence 3-graph of a tetrahedron, i.e., a vertex, an edge, and a face form an edge if they can be extended to a flag. The fact that $S_1 \cap S_2 \cap S_3 = \emptyset$ implies that H is norming and, by Proposition 5.6 and Example 5.2, the H -norm is polynomially equivalent to Gowers’ octahedral norm for 3-graphs.

The key to proving Proposition 5.6 is to show a polynomial equivalence between hypergraph norms and suitably generalised cut-norms. For hypergraphs, Gowers [8, 9] proved that there is a polynomial equivalence between the octahedral norms defined in Example 5.2 and certain generalised cut-norms. Later, Conlon, Hàn, Person, and Schacht [5] showed that the norms discussed in Example 5.3 exist and are polynomially equivalent to certain weaker cut-norms. More recently, Reiher, Rödl, and Schacht [20] proposed further cut-norms lying between these two extremes. As we shall see, the result below gives a suitable H -norm which is polynomially equivalent to each of these cut-norms.

Let $\mathcal{M} = (M_1, M_2, \dots, M_n)$ be an n -tuple of subsets of $[k] = \{1, 2, \dots, k\}$ and define the hypergraph cut-norm $\|\cdot\|_{\square, \mathcal{M}}$ with respect to \mathcal{M} by

$$\|f\|_{\square, \mathcal{M}} := \sup_{u_1, \dots, u_k} \left| \int f(x_1, \dots, x_k) \prod_{i=1}^n u_i(x_{M_i}) \right|,$$

where f is a bounded measurable function on $[0, 1]^k$, x_I for $I \subset [k]$ is the vector $(x_i)_{i \in I}$ of variables with indices in I , and the supremum is taken over all measurable functions u_1, \dots, u_k taking values between $[-1, 1]$. For example, $n = 2$, $M_1 = \{1\}$, and $M_2 = \{2\}$ is the graph cut-norm. The cut-norms described by Gowers [8, 9] correspond to the case when $n = k$ and $M_i = [k] \setminus \{i\}$. In particular, for $n = 3$, we have

$$\|f\|_{\square, \mathcal{M}} = \sup_{u, v, w: [0, 1]^2 \rightarrow [-1, 1]} \left| \int f(x, y, z) u(x, y) v(y, z) w(z, x) \right|.$$

The lemma below states that whenever H is the $(S_1, \dots, S_k; S, W)$ -reflection hypergraph, the H -norm is polynomially equivalent to the generalised cut-norm with respect to a suitably chosen \mathcal{M} .

Lemma 5.8. *Let W be a reflection group with simple reflections $S = \{s_1, \dots, s_n\}$ and let S_1, \dots, S_k be subsets of S . Suppose $\mathcal{M} = (M_1, \dots, M_n)$ is the n -tuple of subsets of $[k]$ such that $M_i = \{j \in [k] : s_i \in S_j\}$ for $i = 1, \dots, n$. Then, whenever H is an $(S_1, \dots, S_k; S, W)$ -reflection hypergraph that is norming, the H -norm is polynomially equivalent to the cut-norm $\|\cdot\|_{\square, \mathcal{M}}$.*

Proof. Let $V_1V_2\cdots V_k$ be the edge of the $(S_1, S_2, \dots, S_k; S, W)$ -reflection hypergraph H contained in the fundamental closed chamber. Observe that M_i is the collection of indices j such that V_j is fixed under the simple reflection s_i , since each collection of inactive mirrors S_j defines the cone V_j . It follows that there is another edge of H , formed by reflecting the edge $V_1V_2\cdots V_k$ in s_i , that contains the set $\{V_j : j \in M_i\}$. The remainder of the proof is now similar to the C_4 case discussed earlier.

To see that $\|f\|_{\square, \mathcal{M}} \leq \|f\|_H$, observe that, since H is norming,

$$\int f(x_1, \dots, x_k) \prod_{i=1}^n u_i(x_{M_i}) \leq \|f\|_H \prod_{i=1}^n \|u_i\|_H \leq \|f\|_H$$

whenever $u_i : [0, 1]^{|M_i|} \rightarrow [-1, 1]$. To prove the converse, note that if $\|f\|_H > c$, then we may assign values to the variables corresponding to the vertices other than V_1, V_2, \dots, V_k to get

$$\left| \int f(x_1, \dots, x_k) \prod_{i=1}^n u_i(x_{M_i}) \right| \geq c^{|E(H)|},$$

as desired. □

Since the choice of \mathcal{M} is independent of the group structure of W , Proposition 5.6 follows as a corollary. Note that each M_i corresponds to the i -th column vector of the $k \times n$ incidence matrix between k sets S_1, S_2, \dots, S_k and the simple reflections s_1, s_2, \dots, s_n . Thus, given a cut-norm $\|\cdot\|_{\square, \mathcal{M}}$ with an ordered collection \mathcal{M} of index subsets of k , we may always construct a k -graph norm that is polynomially equivalent to it. To give some examples, we will now describe hypergraph norms that are polynomially equivalent to the weak cut-norms introduced in [20].

Example 5.9. In [20], two different cut-norms for 3-graphs were studied. For 3-graphs, we always have $k = 3$ and if we let $n = 2$, $M_1 = \{1\}$, and $M_2 = \{2, 3\}$, we get the cut-norm

$$\|f\|_{\square, \mathcal{M}} = \sup_{u, v, w: [0, 1]^2 \rightarrow [-1, 1]} \left| \int f(x, y, z) u(x) v(y, z) \right|.$$

To ensure that $M_i = \{j : s_i \in S_j\}$ holds, let $S_1 = \{s_1\}$ and $S_2 = S_3 = \{s_2\}$. Considering the simplest reflection group $W = \langle s_1 \rangle \times \langle s_2 \rangle$ generated by s_1 and s_2 , the $(S_1, S_2, S_3; S, W)$ -reflection hypergraph is isomorphic to the 3-graph H with

$$V(H) = \{x_0, x_1, y_0, y_1, z_0, z_1\} \text{ and } E(H) = \{x_i y_j z_j : i, j = 0, 1\}.$$

The second cut-norm is the case when $n = 2$, $M_1 = \{1, 2\}$, and $M_2 = \{2, 3\}$, i.e.,

$$\|f\|_{\square, \mathcal{M}} = \sup_{u, v, w: [0, 1]^2 \rightarrow [-1, 1]} \left| \int f(x, y, z) u(x, y) v(y, z) \right|.$$

Then we have $S_1 = \{s_1\}$, $S_2 = \{s_1, s_2\}$, and $S_3 = \{s_2\}$. With the same group $W = \langle s_1 \rangle \times \langle s_2 \rangle$, the $(S_1, S_2, S_3; S, W)$ -reflection hypergraph is isomorphic to the kite-shaped 3-graph with 5 vertices denoted by K in Example 5.5.

By combining Proposition 5.6 and Proposition 5.4, we also have the following weak domination result between hypergraph norms with the same cut involution group.

Corollary 5.10. *Let W be a reflection group with $S = \{s_1, \dots, s_n\}$ the set of simple reflections and let $S_1, \dots, S_k, S'_1, \dots, S'_k$ be subsets of S such that $S'_i \subset S_i$ for each $i = 1, \dots, k$. Suppose H and H' are the $(S_1, \dots, S_k; S, W)$ -reflection hypergraph and the $(S'_1, \dots, S'_k; S, W)$ -reflection hypergraph, respectively, and suppose that they are both norming. Then the H -norm polynomially dominates the H' -norm.*

Getting back to the proof of Proposition 5.4, we first describe the structure of the given hypergraphs H and H' . Let $W \subset \mathbf{GL}(n, \mathbb{R})$ be a reflection group with simple reflections $S = \{s_1, s_2, \dots, s_n\}$ such that each s_i is the reflection along the hyperplane

$$H_i := \{(x_1, x_2, \dots, x_n) \in \mathbb{R}^n : x_i = 0\},$$

i.e., s_i flips the sign of the i -th coordinate. Then the $(S_1, \dots, S_k; S, W)$ -reflection hypergraph H is the k -partite graph with k -partition $\mathcal{A}_1 \cup \dots \cup \mathcal{A}_k$, where \mathcal{A}_i consists of the $2^{n-|S_i|}$ vertices of the form

$$v_i(\xi) := \{(x_1, \dots, x_n) : \xi_j x_j > 0 \text{ for all } j \notin S_i, x_j = 0 \text{ for all } j \in S_i\},$$

where $\xi = (\xi_1, \dots, \xi_n) \in \{\pm 1\}^n$. Here each $\xi \in \{\pm 1\}^n$ represents the closed chamber

$$\{(x_1, \dots, x_n) : \xi_j x_j \geq 0 \text{ for all } j = 1, 2, \dots, n\},$$

so a k -tuple of vertices in $\mathcal{A}_1 \times \mathcal{A}_2 \times \dots \times \mathcal{A}_k$ is an edge of H if and only if it can be written as $(v_1(\xi), v_2(\xi), \dots, v_k(\xi))$ with the same $\xi \in \{\pm 1\}^n$. Let H^+ be the subgraph of H induced on cones in the half-space $\{(x_1, \dots, x_n) : x_n \geq 0\}$ induced by the hyperplane H_n . Using the fact that s_n is a cut involution of H , we have

$$\|f\|_H^{|E(H)|} = \int \left(\int g \, dy_{V(H^+) \setminus F(H)} \right)^2 dy_{F(H)}, \quad (14)$$

where $g = \prod_{(u_1, \dots, u_k) \in E(H^+)} f(y_{u_1}, \dots, y_{u_k})$, $y_I = (y_i)_{i \in I}$ for $I \subset V(H)$, and $F(H)$ is the set of vertices $\{v_j(\xi) \in V(H) : s_n \in S_j, \xi \in \{\pm 1\}^n\}$ that are fixed by s_n . Furthermore, H^+ is isomorphic to the $(S_1^+, \dots, S_k^+; S, W)$ -reflection hypergraph with $S_i^+ = S_i \cup \{s_n\}$ for all i , since adding s_n to all S_i makes the sign of x_n fixed.

Proof of Proposition 5.4. To show that H is always norming, we may assume that there exists a simple reflection $s \in \bigcap_{i=1}^k S_i$. Observe that the $(S_1^-, \dots, S_k^-; S, W)$ -reflection hypergraph H^- , where $S_i^- = S_i \setminus \{s\}$ for all i , consists of two vertex-disjoint copies of H mapped to each other by s . Then the H^- -norm takes exactly the same value as the H -norm and, hence, by induction on $|\bigcap_{i=1}^k S_i|$, we are done.

To prove that the H -norm dominates the H' -norm, we may assume $S'_i = S_i \setminus \{s_n\}$, $s_n \in S_i$, and $S'_j = S_j$ for $j \neq i$. Both H and H' have the same induced subgraph on cones in the half-space $\{(x_1, \dots, x_n) : x_n \geq 0\}$, because it must be isomorphic to the $(S_1^+, \dots, S_k^+; S, W)$ -reflection hypergraph with $S_j^+ = S_j \cup \{s_n\}$ for all j . When writing $\|f\|_{H'}$ as in (14), the fixed point set $F(H')$ under s_n becomes

$$F(H') = \{v_j(\xi) \in V(H') : s_n \in S'_j, \xi \in \{\pm 1\}^n\} = F(H) \setminus \mathcal{A}_i,$$

since s_n is removed from S_i . Therefore, the H' -norm can be expressed, with the same g as in (14), by

$$\|f\|_{H'}^{|E(H')|} = \int \left(\int g \, dy_{V(H^+) \setminus F(H)} dy_{\mathcal{A}_i} \right)^2 dy_{F(H) \setminus \mathcal{A}_i},$$

and the Cauchy–Schwarz inequality implies that this is less than or equal to (14). \square

5.3 Applications to Sidorenko's conjecture

We will now discuss how to apply our results to Sidorenko's conjecture. As mentioned in the introduction, there is the obvious fact that any weakly norming graph satisfies Sidorenko's conjecture. However, one can go beyond this by applying the entropy techniques that have been developed in recent years [15, 16, 14, 24, 6]. The key facts we will use about entropy are contained in the following lemma, though we refer the reader to [6] for further information on entropy, conditional entropy, and their use in relation to Sidorenko's conjecture. Throughout this subsection, logarithms will be understood to be base 2.

Lemma 5.11. *Let X, Y , and Z be random variables and suppose that X takes values in a set S , $H(X)$ is the entropy of X , and $H(X|Y)$ is the conditional entropy of X given Y . Then*

- (i) $H(X) \leq \log |S|$,
- (ii) $H(X|Y, Z) = H(X|Z)$ if X and Y are conditionally independent given Z .

As in [14, 6], we say that a bipartite graph H has Sidorenko's property if H satisfies (5) for all non-negative symmetric functions f , i.e., Sidorenko's conjecture holds for H .

To motivate what follows, suppose we wish to show that the graph obtained by gluing two copies H_1 and H_2 of the 1-subdivision of K_4 along induced 6-cycles has Sidorenko's property. Formally, there is a vertex set $J = V(H_1) \cap V(H_2)$ on which the induced subgraphs of both H_1 and H_2 are isomorphic to a 6-cycle and we are interested in proving that the graph H on vertex set $V(H_1) \cup V(H_2)$ with edges $E(H_1) \cup E(H_2)$ has Sidorenko's property. As in [6] and [24], we generate a random homomorphism in $\text{Hom}(H, G)$ and analyse its entropy to get a lower bound for $|\text{Hom}(H, G)|$. To this end, we consider the following way of generating a random element in $\text{Hom}(H, G)$:

- (i) Take a uniform random homomorphic copy of H_1 , i.e., choose an element in $\text{Hom}(H_1, G)$ uniformly at random.
- (ii) Choose another homomorphic copy of H_2 that extends the 6-cycle on J uniformly at random.

This algorithm always gives a random homomorphism \mathbf{w} in $\text{Hom}(H, G)$, so it remains to analyse the entropy. We regard \mathbf{w} as a random vector indexed by $V(H)$ and, for $I \subset V(H)$, write \mathbf{w}_I for the random vector $(\mathbf{w}(v))_{v \in I} \in V(G)^I$. Let $U = V(H_1) \setminus V(J)$ and $W = V(H_2) \setminus V(J)$. Then

$$\begin{aligned}
 \log |\text{Hom}(H, G)| &\geq H(\mathbf{w}) = H(\mathbf{w}_U, \mathbf{w}_J, \mathbf{w}_W) \\
 &= H(\mathbf{w}_W | \mathbf{w}_U, \mathbf{w}_J) + H(\mathbf{w}_U, \mathbf{w}_J) \\
 &= H(\mathbf{w}_W | \mathbf{w}_J) + \log |\text{Hom}(H_1, G)| \\
 &= H(\mathbf{w}_W, \mathbf{w}_J) - H(\mathbf{w}_J) + \log |\text{Hom}(H_1, G)| \\
 &= 2 \log |\text{Hom}(H_1, G)| - H(\mathbf{w}_J).
 \end{aligned} \tag{15}$$

If J were an independent set, we could upper bound $|\text{Hom}(J, G)|$ by $|G|^{|J|}$ and use Lemma 5.11 (i) to derive an upper bound on $H(\mathbf{w}_J)$. This is one of the main tricks used in [6] and [24], an elaboration of which also allows one to handle the case where the graph induced on J is a tree. In our case, the graph induced on J is a 6-cycle, so we have $H(\mathbf{w}_J) \leq \log |\text{Hom}(C_6, G)|$. Since H_1 is a weakly norming graph, the monotonicity property (11) implies that the homomorphism densities of H_1 and C_6 in G satisfy $t_{H_1}(G)^{1/|E(H_1)|} \geq t_{C_6}(G)^{1/|E(C_6)|}$. This in turn implies that

$$|\text{Hom}(H_1, G)|^{1/2} |V(G)| \geq |\text{Hom}(C_6, G)|.$$

Plugging this into (15), we have

$$\begin{aligned} \log |\text{Hom}(H, G)| &\geq \frac{3}{2} \log |\text{Hom}(H_1, G)| - \log |V(G)| \\ &\geq 18 \log \left(\frac{|\text{Hom}(K_2, G)|}{|V(G)|^2} \right) + 14 \log |V(G)|, \end{aligned}$$

where the last inequality follows from the fact that H_1 has Sidorenko's property. Therefore, H has Sidorenko's property.

More generally, we may iterate this gluing process in a tree-like way. To describe the resulting graph, we use the notion of tree decomposition introduced by Halin [10] and developed by Robertson and Seymour [21]. A *tree decomposition* of a graph H 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} .

Now, given a weakly norming graph N , an N -*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 each isomorphic to N .
2. For every pair $X, Y \in \mathcal{F}$ which are adjacent in \mathcal{T} , there is an isomorphism between the two copies of N that fixes $X \cap Y$.

We say that a bipartite graph is N -*decomposable* if it allows a N -decomposition, i.e., it can be obtained by gluing copies of the weakly norming graph N in a tree-like way. The main theorem in this subsection is as follows.

Theorem 5.12. *If N is weakly norming, any N -decomposable graph H has Sidorenko's property.*

Proof. Let $(\mathcal{F}, \mathcal{T})$ be an N -decomposition of H and let G be the target graph in which we wish to embed H . The randomised algorithm for generating a copy of H is a straightforward generalisation of that discussed in the example above: pick a root $R \in \mathcal{F}$ and choose a uniform random homomorphism from $\text{Hom}(H[R], G)$ and, for each child X of R , choose a uniform random (homomorphic) copy of $H[X]$ extending the embedded copy of $H[R \cap X]$. Repeating this process, we obtain a random homomorphic copy \mathbf{w} of H . Following the proof of Theorem 1.2 in [6], we arrive at the identity

$$H(\mathbf{w}) = \sum_{X \in \mathcal{F}} H(\mathbf{w}_X) - \sum_{XY \in E(\mathcal{T})} H(\mathbf{w}_{X \cap Y}).$$

Let K_{XY} be the induced subgraph $H[X \cap Y]$ for $XY \in E(\mathcal{T})$. We may bound $H(\mathbf{w})$ and $H(\mathbf{w}_{X \cap Y})$ from above by $\log |\text{Hom}(H, G)|$ and $\log |\text{Hom}(K_{XY}, G)|$, respectively. Moreover, since \mathbf{w}_X is distributed uniformly in $\text{Hom}(N, G)$, this gives

$$\log t_H(G) \geq |\mathcal{F}| \log t_N(G) - \sum_{XY \in E(\mathcal{T})} \log t_{K_{XY}}(G), \quad (16)$$

after rescaling the equation by subtracting $|V(H)| \log |V(G)|$ on each side. Since N is weakly norming, the monotonicity property (11) implies that

$$t_{K_{XY}}(G) \leq t_N(G)^{|E(K_{XY})|/|E(N)|}.$$

Plugging this bound into (16), it follows that

$$\log t_H(G) \geq \frac{1}{|E(N)|} \left(|\mathcal{F}| |E(N)| - \sum_{XY \in E(\mathcal{T})} |E(K_{XY})| \right) \log t_N(G).$$

Note that $|\mathcal{F}| |E(N)| - \sum_{XY \in E(\mathcal{T})} |E(K_{XY})|$ is exactly the number of edges in H and, hence, we have the bound $t_H(G) \geq t_N(G)^{|E(H)|/|E(N)|}$. As N has Sidorenko's property, we conclude that H does also. \square

The theorem above is intended as an example of how norming graphs may be used to build up new examples of graphs satisfying Sidorenko's property. We expect that additional refinements of the entropy method can be used to broaden this class further, but we have not attempted a comprehensive treatment here.

6 Concluding remarks

Characterising weakly norming graphs. Although Theorem 1.2 gives a fairly large class of weakly norming graphs, it is still an open problem to characterise them all. For our arguments to work on a particular graph H , it was necessary that H be edge-transitive under its cut involution group. We also insisted on the existence of a percolating sequence in H , but we suspect that this condition may not be needed.

Conjecture 6.1. *A bipartite graph H is weakly norming if it is edge-transitive under its cut involution group.*

Underlying this conjecture is a further suspicion that reflection graphs may constitute the entire class of graphs which are edge-transitive under their cut involution groups. Were this indeed the case, Conjecture 6.1 would follow immediately from Theorem 1.2.

In [12], Hatami showed that weakly norming graphs must be balanced and bi-regular and that the class of weakly norming graphs is closed under taking tensor products. There are analogues of these results for reflection graphs: any reflection graph must be bi-regular since it is vertex-transitive on each side, while the tensor product of two reflection graphs is again a reflection graph. It is not too difficult to prove that the same properties hold for graphs that are edge-transitive under the cut involution group. These coincidences suggest that the converse to Conjecture 6.1 may also be true, though this is significantly more tentative than Conjecture 6.1.

Conjecture 6.2. *A bipartite graph H is weakly norming only if it is edge-transitive under its cut involution group.*

It would already be very interesting to show that weakly norming graphs are necessarily edge-transitive.

Norming graphs versus weakly norming graphs. When can we guarantee that a weakly norming graph is also norming? Theorem 1.3 gives a partial answer to this natural question. However, we again suspect, in analogy with Conjectures 6.1 and 6.2, that edge-transitivity under the stable involution group is a necessary and sufficient property.

Conjecture 6.3. *A bipartite graph H is norming if and only if it is edge-transitive under its stable involution group.*

This is reminiscent of another conjecture, proposed in [3], about determining those graphs H for which $t_H(f)$ is always non-negative. It might be interesting to investigate the connection.

Conjecture 6.4 (Positive graph conjecture [3]). *$t_H(f)$ is non-negative for all $f : [0, 1]^2 \rightarrow \mathbb{R}$ if and only if there is a stable involution of H .*

Complex-valued functions. Our results on norming (hyper)graphs, Theorem 1.3 and Theorem 5.1, can also be used to define (hyper)graph norms for complex-valued functions. Suppose H is the $(S_1, \dots, S_k; S, W)$ -reflection hypergraph, where W is a reflection group and $\bigcap_{i=1}^k S_i = \emptyset$. Note that each reflection in W acts as a stable involution and there exists a one-to-one correspondence between the edges of H and the set \mathcal{C} of open chambers, since each edge is contained in exactly one closed chamber. By the one-to-one correspondence between W and \mathcal{C} , there also exists a one-to-one correspondence $\xi : W \rightarrow E(H)$ between W and $E(H)$. Let $W_+ := \{w \in W : \ell(w) \text{ is even}\}$ and $W_- := \{w \in W : \ell(w) \text{ is odd}\}$. We may write W_+ and W_- as the inverse images of 1 and -1 under the group homomorphism $w \mapsto \det(w) = (-1)^{\ell(w)}$ from W to the multiplicative group $\{\pm 1\}$. It follows that W_+ and W_- partition W into parts of equal size and each reflection $t \in W$ maps $w \in W_+$ to $tw \in W_-$ and vice versa. For a complex-valued function f on $[0, 1]^k$, define

$$\|f\|_H := \left| \int \prod_{e \in \xi(W_+)} f(x_e) \prod_{e' \in \xi(W_-)} \overline{f(x_{e'})} \right|^{1/|E(H)|},$$

where $x_e = (x_{u_1}, x_{u_2}, \dots, x_{u_k})$ for $e = (u_1, u_2, \dots, u_k)$. This is well-defined, even without taking the absolute value of the integral above, since

$$\int \prod_{e \in \xi(W_+)} f(x_e) \prod_{e' \in \xi(W_-)} \overline{f(x_{e'})} = \int g \bar{g},$$

where

$$g(x_{F_\phi}) = \int \prod_{e \in \xi(W_+) \cap (L_\phi \cup F_\phi)^k} f(x_e) \prod_{e' \in \xi(W_-) \cap (L_\phi \cup F_\phi)^k} \overline{f(x_{e'})} dx_{L_\phi}$$

for a stable involution ϕ that corresponds to a reflection in W . When $W = \langle s_1 \rangle \times \dots \times \langle s_n \rangle$, $s_i \in S$, this gives complex-valued versions of Gowers' octahedral norms, but it also allows some more exotic examples. For example, the hypergraph described in Example 5.7 can now easily be modified to define a norm on the vector space of complex-valued functions $f : [0, 1]^3 \rightarrow \mathbb{C}$.

Acknowledgements. Part of this work was carried out while the authors participated in the LMS-CMI Research School on Regularity and Analytic Methods in Combinatorics at the University of Warwick and also while the second author was visiting KIAS. The second author would like to thank Seung Jin Lee for suggestions of references and helpful discussions on algebraic combinatorics. We would also like to thank Alexander Sidorenko for some helpful remarks on an earlier version of this paper.

References

- [1] A. Björner and F. Brenti. *Combinatorics of Coxeter groups*, volume 231 of *Graduate Texts in Mathematics*. Springer, New York, 2005.
- [2] G. Blakley and P. Roy. A Hölder type inequality for symmetric matrices with nonnegative entries. *Proc. Amer. Math. Soc.*, 16:1244–1245, 1965.
- [3] O. A. Camarena, E. Csóka, T. Hubai, G. Lippner, and L. Lovász. Positive graphs. *European J. Combin.*, 52:290–301, 2016.
- [4] D. Conlon, J. Fox, and B. Sudakov. An approximate version of Sidorenko’s conjecture. *Geom. Funct. Anal.*, 20:1354–1366, 2010.
- [5] D. Conlon, H. Hàn, Y. Person, and M. Schacht. Weak quasi-randomness for uniform hypergraphs. *Random Structures Algorithms*, 40:1–38, 2012.
- [6] D. Conlon, J. H. Kim, C. Lee, and J. Lee. Some advances on Sidorenko’s conjecture. Preprint available at arXiv:1510.06533 [math.CO].
- [7] W. T. Gowers. A new proof of Szemerédi’s theorem. *Geom. Funct. Anal.*, 11:465–588, 2001.
- [8] W. T. Gowers. Quasirandomness, counting and regularity for 3-uniform hypergraphs. *Combin. Probab. Comput.*, 15:143–184, 2006.
- [9] W. T. Gowers. Hypergraph regularity and the multidimensional Szemerédi theorem. *Ann. of Math.*, 166:897–946, 2007.
- [10] R. Halin. S-functions for graphs. *J. Geom.*, 8:171–186, 1976.
- [11] H. Hatami. *On generalizations of Gowers norms*. PhD thesis, University of Toronto, 2009.
- [12] H. Hatami. Graph norms and Sidorenko’s conjecture. *Israel J. Math.*, 175:125–150, 2010.
- [13] J. Humphreys. *Reflection groups and Coxeter groups*, volume 29 of *Cambridge Studies in Advanced Mathematics*. Cambridge University Press, 1992.
- [14] J. H. Kim, C. Lee, and J. Lee. Two approaches to Sidorenko’s conjecture. *Trans. Amer. Math. Soc.*, 368:5057–5074, 2016.
- [15] S. Kopparty and B. Rossman. The homomorphism domination exponent. *European J. Combin.*, 32:1097–1114, 2011.
- [16] J. X. Li and B. Szegedy. On the logarithmic calculus and Sidorenko’s conjecture. To appear in *Combinatorica*.
- [17] L. Lovász. *Large networks and graph limits*. Amer. Math. Soc. Colloq. Publ. American Mathematical Society, Providence, RI, 2012.
- [18] L. Lovász and B. Szegedy. Limits of dense graph sequences. *J. Combin. Theory Ser. B*, 96:933–957, 2006.
- [19] P. McMullen and E. Schulte. *Abstract regular polytopes*, volume 92 of *Encyclopedia of Mathematics and its Applications*. Cambridge University Press, Cambridge, 2002.

- [20] C. Reiher, V. Rödl, and M. Schacht. On a Turán problem in weakly quasirandom 3-uniform hypergraphs. Preprint.
- [21] N. Robertson and P. D. Seymour. Graph minors III: Planar tree-width. *J. Combin. Theory Ser. B*, 36:49–64, 1984.
- [22] A. Sidorenko. A correlation inequality for bipartite graphs. *Graphs Combin.*, 9:201–204, 1993.
- [23] A. Sidorenko. Inequalities for functionals generated by bipartite graphs. *Discrete Math. Appl.*, 2:489–504, 1993.
- [24] B. Szegedy. An information theoretic approach to Sidorenko’s conjecture. Preprint available at arXiv:1406.6738 [math.CO].