

Counting Graph Homomorphisms

Christian Borgs*, Jennifer Chayes†, László Lovász‡,
Vera T. Sós§, Katalin Vesztegombi¶

February 2006

Abstract

Counting homomorphisms between graphs (often with weights) comes up in a wide variety of areas, including extremal graph theory, properties of graph products, partition functions in statistical physics and property testing of large graphs.

In this paper we survey recent developments in the study of homomorphism numbers, including the characterization of the homomorphism numbers in terms of the semidefiniteness of “connection matrices”, and some applications of this fact in extremal graph theory.

We define a distance of two graphs in terms of similarity of their global structure, which also reflects the closeness of (appropriately scaled) homomorphism numbers into the two graphs. We use homomorphism numbers to define convergence of a sequence of graphs, and show that a graph sequence is convergent if and only if it is Cauchy in this distance. Every convergent graph sequence has a limit in the form of a symmetric measurable function in two variables. We use these notions of distance and graph limits to give a general theory for parameter testing.

The convergence can also be characterized in terms of mappings of the graphs into fixed small graphs, which is strongly connected to important parameters like ground state energy in statistical physics, and to weighted maximum cut problems in computer science.

1 Introduction

For two finite graphs G and H , let $\text{hom}(G, H)$ denote the number of homomorphisms (adjacency-preserving mappings) from G to H . Counting homomorphisms between graphs has many interesting aspects.

(a) A large part of extremal graph theory can be expressed as inequalities between various homomorphism numbers. For example, Turán’s Theorem for triangles follows from the inequality (due to Goodman [30]):

$$\text{hom}(K_1, H)\text{hom}(K_3, H) \geq \text{hom}(K_2, H)(2\text{hom}(K_2, H) - \text{hom}(K_1, H)^2). \quad (1)$$

One may wish to obtain a characterization of such inequalities.

(b) Homomorphism numbers characterize graphs: it was proved in [37] that if two graphs G and G' have the property that $\text{hom}(F, G) = \text{hom}(F, G')$ for every finite graph F , then G and

*Microsoft Research, One Microsoft Way, Redmond, WA 98052; borgs@microsoft.com

†Microsoft Research, One Microsoft Way, Redmond, WA 98052; jchayes@microsoft.com

‡Microsoft Research, One Microsoft Way, Redmond, WA 98052; lovasz@microsoft.com

§Alfréd Rényi Institute of Mathematics, Budapest, Reáltanoda U. 13-15., H-1053 Budapest, Hungary; sos@renyi.hu; Research supported in part by OTKA grants No. TO32236, TO39210, TO42750.

¶Department of Computer Science, Eötvös Loránd University, Pázmány Péter sétány 1/C, H-1117 Budapest, Hungary; vezsterk@cs.elte.hu

G' are isomorphic (one actually needs one additional condition, the condition that both graphs are twin-free; see Section 2.1 for the definition of this notion). In other words, let us order all finite graphs in a sequence (F_1, F_2, \dots) , and assign to each graph G its *profile*, the (infinite) sequence $(\text{hom}(F_1, G), \text{hom}(F_2, G), \dots)$; then this sequence characterizes G . (This fact can be used to prove, for example, the cancellation property of strong multiplication of graphs).

It is often worthwhile to normalize the homomorphism numbers, and consider the *homomorphism densities*

$$t(F, G) = \frac{\text{hom}(F, G)}{|V(G)|^{|V(F)|}}. \quad (2)$$

(Thus $t(F, G)$ is the probability that a random map of $V(F)$ into $V(G)$ is a homomorphism.) Instead of the profile $(\text{hom}(F_1, G), \text{hom}(F_2, G), \dots)$, we can consider the *scaled profile* $(t(F_1, G), t(F_2, G), \dots)$ of the graph G . Such a normalization was first introduced in [23].

(c) Partition functions of many models in statistical mechanics can be expressed as graph homomorphism functions. For example, let G be an $n \times n$ grid, and suppose that every node of G (every “site”) can be in one of two states, “UP” or “DOWN”. The properties of the system are such that no two adjacent sites can be “UP”. A “configuration” is a valid assignment of states to each node. The number of configurations is the number of independent sets of nodes in G , which in turn can be expressed as the number of homomorphisms of G into the graph H consisting of two nodes, “UP” and “DOWN”, connected by an edge, and with an additional loop at “DOWN”. To define the thermodynamic functions in physical models, one needs to extend the notion of graph homomorphism to the case when the nodes and edges of H have weights (see Section 2.1).

(d) Suppose that G is a huge graph, and we know the numbers $\text{hom}(F, G)$ (exactly or approximately) for small graphs F . What kind of information can be derived about the global structure of G ? The long-standing Reconstruction Conjecture is equivalent to the assertion that it is enough to know all numbers $\text{hom}(F, G)$ with $|V(F)| < |V(G)|$ in order to recover the isomorphism type of G . The fact that this is unsolved shows the difficulty of this kind of question, but our interest here is in the case when much less is given: $\text{hom}(F, G)$ is only known for very small graphs F .

(e) This is closely related to an important area of computer science called *Property Testing*. In this model, we have a huge graph about which we can obtain information only by taking a small sample of the nodes and examine the subgraph induced by them. This is equivalent to knowing the homomorphism densities (2) for graphs F of small size. What makes the theory of Property Testing interesting is the fact that from such meager local information nontrivial properties and parameters of the graph can be inferred.

(f) Increasing sequences of (sparse) graphs generated by some specific random rule of growth have recently been used to model the Internet; see [8] and references therein. What are the limiting properties of such graph sequences? To what extent can these properties be derived from local observation (observing a neighborhood of bounded radius of a few randomly chosen nodes)? This question can be rephrased as follows: To what extent are these properties characterized by the homomorphism numbers of smaller graphs into the modeling sequence?

The main setup of our studies is the following. If we are given a (large, usually simple) graph G , we may try to study its local structure by counting the homomorphisms of various “small” graphs F into G ; and we can study its global structure by counting its homomorphisms into various small graphs H (called “softcore” weighted graphs; see Section 2.1 for a definition of softcore). So the scheme to keep in mind is

$$F \longrightarrow G \longrightarrow H. \quad (3)$$

According to the above discussion, in the scheme (3), the study of the local structure of G by “probing from the left with F ” is related to property testing, while the study of the global

structure of G by “probing from the right with H ” is related to statistical physics. As in statistical physics, the best choice of graphs H to “probe G from the right” is not simple unweighted graphs but weighted graphs. Furthermore, besides counting (weighted) homomorphisms into H , it is also useful to consider maximizing the weight of such homomorphisms, which is again related to well-studied questions both in statistical physics and graph property testing.

The number $\text{hom}(F, G)$, as a function of F with G fixed, is a graph parameter (a function of graphs F invariant under isomorphism). Graph parameters arising this way were characterized in [28]. The necessary and sufficient condition involves certain matrices, called *connection matrices*, associated with the graph parameter (see Section 3 for the definition): These matrices must be positive definite and must satisfy a rank condition. The semidefiniteness condition is in fact familiar from statistical physics, where it is called *reflection positivity*.

The scaled profile does not determine the graph: if we “blow up” every node into the same number of “twins”, then we get a graph with exactly the same scaled profile. It turns out that this is all: any two graphs with the same scaled profile are obtained from one and the same graph by blowing up its nodes in two different ways.

Now we come to our main question: What can be said about two graphs whose scaled profiles are approximately the same? Which properties of a graph G are determined if we only know a few of the numbers $\text{hom}(F, G)$, and even these are only known approximately? This question turns out to be very interesting and it leads to a number of results, and even more open problems, leading to quasirandom and generalized quasirandom graphs, and connecting to the “Property Testing” research in computer science and to statistical physics.

There is a way to measure the “distance” of two graphs so that they are close in this distance if and only if they have approximately the same scaled profile [16]. This distance has many nice properties. On the one hand, important parameters like the triangle density or the fraction of edges in the maximum cut are continuous (often even Lipschitz) functions in this metric. On the other hand, a sufficiently large random subgraph of an arbitrarily large graph will be close to the whole graph with large probability. This fact explains several results in the theory of Property Testing. Szemerédi’s Regularity Lemma (at least in its weaker but more effective form due to Frieze and Kannan [29]) can be rephrased as follows: for every $\varepsilon > 0$, all graphs with at most $2^{2/\varepsilon^2}$ nodes form an ε -net in the metric space of all graphs. In our context, an ε -net is defined as a set of weighted graphs such that, for every graph G , there exists a graph H in the set which has at most distance ε from G .

Once we make the set of all graphs into a metric space, we can make it complete. Is there any combinatorial meaning of the new points in the completion? Surprisingly, the answer is in the affirmative, and in fact in more than one way [44, 45]. These limit points can be described as symmetric measurable functions $W : [0, 1]^2 \rightarrow [0, 1]$ modulo measure-preserving transformations, as reflection positive graph parameters, as random graph models satisfying some natural compatibility conditions, or as probability distributions of countable graphs with natural symmetries (see Chapter 4).

It is an important property of this completion that it is compact, so that every infinite sequence of graphs has a convergent subsequence. Several arguments in extremal graph theory and elsewhere can be simplified by going to the limit and thereby getting rid of remainder terms.

We conclude this introduction by mentioning two related bodies of work. There are many interesting questions that concern the existence of graph homomorphisms rather than their number (an example is 4-colorability of a graph), and such questions have been studied quite extensively, especially by the Czech school. These results are described in the recent book by Hell and Nešetřil [34]. The set of all homomorphisms between two graphs can be endowed with a topological structure, which turns out to be an important tool in the study of chromatic number. See the book of Matoušek [51], and also the recent papers of Babson and Kozlov [5, 6].

This paper is organized as follows. In Section 2 we define homomorphism numbers, including homomorphisms into a measurable function, and describe the basic examples. In Section 3

we define connection matrices, study their rank, and use their semidefiniteness to characterize homomorphism functions. We also describe an analogous (but much more difficult) edge-coloring version. In Section 4 we define convergence of a sequence of dense graphs, and show that they have interesting limit objects, which can be described as measurable functions, reflection positive graph parameters, or very natural models of finite or countable random graphs. In Section 5, we introduce a metric on graphs that corresponds to the above notion of convergence, and show its connection with Szemerédi’s Regularity Lemma and graph property testing. Section 6 studies homomorphisms from large graphs, rather than into large graphs, which leads to quantities of both combinatorial and physical interest; we show that they too can be used to characterize convergent graph sequences. Section 7 contains some applications to extremal graph theory. In Section 8 we conclude with describing some analogous (but less complete) results for sequences of graphs with bounded degree.

2 Homomorphism numbers

2.1 Unweighted and weighted graphs

A graph is *simple* if it has no loops or parallel edges. A *graph parameter* is a function defined on finite graphs, invariant under isomorphisms. We’ll talk of a *simple graph parameter* if it is only defined on simple graphs. Sometimes it is convenient to think of a simple graph parameter as a function defined on all graphs with multiple edges (but no loops) that is invariant under adding parallel edges. A graph parameter t is called *multiplicative* if $t(G) = t(G_1)t(G_2)$ whenever G is the disjoint union of G_1 and G_2 . We say that a graph parameter is *normalized* if its value on K_1 , the graph with one node and no edge, is 1. (Note that if a graph parameter is multiplicative and not identically 0, then its value on K_0 , the graph with no nodes and edges, is 1. The graph parameter $t(\cdot, G)$ introduced in the introduction is multiplicative and normalized for every graph G .)

Recall that for two (finite) simple graphs F and G , $\text{hom}(F, G)$ denotes the number of homomorphisms (adjacency preserving maps) from F to G .

A *weighted graph* G is a graph with a weight $\alpha_i(G)$ associated with each node i and a weight $\beta_{ij}(G)$ associated with each edge ij . We’ll assume (unless otherwise stated) that the weights $\alpha_i(G)$ are positive. The weights $\beta_{ij}(G)$ will be real, and most often nonnegative. If the graph G is understood from the context, we will also use the notation α_i and β_{ij} .

An edge with weight 0 will play the same role as no edge between those nodes, so we could assume that we only consider weighted complete graphs with loops at all nodes (but this is not always convenient). A weighted graph is called *softcore* if it is a complete graph with loops at each node, and every edgeweight is strictly positive. An *unweighted graph* is a weighted graph where all the nodeweights and edgeweights are 1.

For a weighted graph G , we denote by $\alpha(G)$ the sum of its nodeweights. Often it will be useful to divide all nodeweights by $\alpha(G)$, to get a weighted graph \widehat{G} in which the sum of nodeweights is 1.

Let F and G be two weighted graphs. To every map $\phi : V(F) \rightarrow V(G)$, we assign the weight

$$\text{hom}_\phi(F, G) = \prod_{uv \in E(F)} [\beta_{\phi(u)\phi(v)}(G)]^{\beta_{uv}(F)} \quad (4)$$

(here $0^0 = 1$). We then define

$$\text{hom}(F, G) = \sum_{\phi: V(F) \rightarrow V(G)} \alpha_\phi \text{hom}_\phi(F, G), \quad (5)$$

where

$$\alpha_\phi = \prod_{u \in V(F)} [\alpha_{\phi(u)}(G)]^{\alpha_u(F)}. \quad (6)$$

(A little care is necessary, since the exponential $[\beta_{\phi(u)\phi(v)}(G)]^{\beta_{uv}(F)}$ may not be well defined; but it well defined e.g. if the edge weights are positive. This problem will not arise in the cases we consider.)

We'll use this definition most often in the case when F is a simple unweighted graph, so that

$$\alpha_\phi = \prod_{u \in V(F)} \alpha_{\phi(u)}(G)$$

and

$$\text{hom}_\phi(F, G) = \prod_{uv \in E(F)} \beta_{\phi(u)\phi(v)}(G).$$

An interesting case is when, in addition, $\alpha(G) = 1$. Then the nodeweights in G define a probability distribution on $V(G)$, and the coefficient α_ϕ is the probability of the random map ϕ (if the images of the nodes of F are chosen independently). So $\text{hom}(F, G)$ is the expectation of $\text{hom}_\phi(F, G)$.

We can extend homomorphism densities (2) to the case when G is a weighted graph with nodeweights α_i and edgeweights β_{ij} by replacing $|V(G)|$ by $\alpha(G)$:

$$t(F, G) = \frac{\text{hom}(F, G)}{\alpha(G)^{|V(F)|}} = \text{hom}(F, \widehat{G}).$$

Let $\mathcal{F} = \{F_1, F_2, \dots\}$ denote the set of (isomorphism types of) all simple finite graphs. To every weighted graph G , we assign its *hom-profile* (or briefly *profile*), the (infinite) vector

$$h_G = (\text{hom}(F_1, G), \text{hom}(F_2, G), \dots) \in \mathbb{R}^{\mathcal{F}}.$$

Recall that we define the *scaled profile* of G as the (infinite) vector $t_G = (t(F_1, G), t(F_2, G), \dots) \in \mathbb{R}^{\mathcal{F}}$. The scaled profile does not determine the graph. For example, if G is an unweighted graph and G' is obtained from G by replacing every node by N independent nodes, then

$$t_{G'} = t_G.$$

More generally, let G be a weighted graph and let $u, v \in V(G)$ be *twins*, i.e., $\beta_{uw}(G) = \beta_{vw}(G)$ for every $w \in V(G)$ (note that $\alpha_u(G)$ may be different from $\alpha_v(G)$). Merging twins in a weighted graph does not change its scaled profile. Furthermore, if we multiply all nodeweights of a weighted graph by the same positive constant, then its scaled profile does not change. If we merge twins as long as we can, we say that we have performed *twin-reduction*.

Proposition 2.1 [39] *If two weighted graphs have the same scaled profile, then after twin-reduction one can be obtained from the other by multiplying all nodeweights by the same positive scalar.*

One of our main concerns will be: *if we only know a bounded number of entries of the scaled profile of a graph G , and even this only approximately, to what degree is the graph determined?*

From a graph-theoretic perspective, the following variations of homomorphism functions are perhaps more important: Let $\text{inj}(F, G)$ denote the number of homomorphisms that are injective on the nodes, and $\text{ind}(F, G)$, the number of embeddings as an induced subgraph. Finally, let $\text{surj}(F, G)$ denote the number of homomorphisms that are surjective on the nodes.

For weighted graphs, $\text{inj}(F, G)$ and $\text{surj}(F, G)$ are easily defined by restricting the sum in (5) to sums over injective and surjective maps, respectively, but the definition of $\text{ind}(F, G)$ requires

some care. Here we only consider the case where F is simple, and G is a weighted graph without loops. We then define

$$\text{ind}(F, G) = \sum_{\phi: V(F) \rightarrow V(G)} \alpha_{\phi} \text{ind}_{\phi}(F, G), \quad (7)$$

where the sum goes over all injective maps from $V(F)$ to $V(G)$ and

$$\text{ind}_{\phi}(F, G) = \prod_{uv \in E(F)} \beta_{\phi(u)\phi(v)}(G) \prod_{uv \in E(\bar{F})} (1 - \beta_{\phi(u)\phi(v)}(G)), \quad (8)$$

with \bar{F} denoting the complement of F (again defined to be a simple graph).

Analogously to the hom-profile, we can define the *inj-profile* and *ind-profile* of a graph G . It follows from the identities to be discussed below that any of the profiles determines the others. It is obvious that the ind-profile determines the graph (look at the largest non-zero entry). It follows that the inj-profile and hom-profile of a graph also determine the graph (up to isomorphism) [37]. For future reference, we also define the set of simple graph parameters which is the pointwise closure of the set of all hom-profiles:

$$\mathcal{T}_0 = \{t(\cdot) : \exists (G_n) \text{ s.t. } t(F) = \lim_{n \rightarrow \infty} t(F, G_n) \ \forall \text{ finite } F\}.$$

Here we restrict (G_n) to be a sequence of simple graphs. The reason for the subscript on \mathcal{T} is that, in our paper [16], we will consider a more general class of graph parameters, which will be denoted by \mathcal{T} . Note also that in [44], the authors used \mathcal{T} rather than \mathcal{T}_0 to denote the smaller class considered here.

2.2 Simple properties

There are some simple identities that hold for homomorphism numbers. If F is the disjoint union of two graphs F_1 and F_2 , then

$$\text{hom}(F, G) = \text{hom}(F_1, G) \text{hom}(F_2, G). \quad (9)$$

If F is connected, and G is the disjoint union of two graphs G_1 and G_2 , then

$$\text{hom}(F, G) = \text{hom}(F, G_1) + \text{hom}(F, G_2). \quad (10)$$

So in a sense it is enough to study homomorphisms between connected graphs.

For two simple graphs G_1, G_2 , we define their *categorical product* $G_1 \times G_2$ as the graph with node set $V(G_1) \times V(G_2)$, in which (i_1, i_2) is connected to (j_1, j_2) ($i_1, j_1 \in V(G_1)$, $i_2, j_2 \in V(G_2)$) if and only if $i_1 j_1 \in E(G_1)$ and $i_2 j_2 \in E(G_2)$. The definition can be extended to weighted graphs (possibly with loops) by defining the weight of the node (i_1, i_2) as the product of the weights of the nodes i_1 and i_2 , and the weight of the edge $(i_1, i_2)(j_1, j_2)$ as the product of the weights of the edges $i_1 j_1$ and $i_2 j_2$. For this product, we have the identity

$$\text{hom}(F, G_1 \times G_2) = \text{hom}(F, G_1) \cdot \text{hom}(F, G_2). \quad (11)$$

What about $\text{hom}(F_1 \times F_2, G)$? There is no identity for this number in terms of the notions introduced so far, but there is one if one also introduces the operation of exponentiation (which we do not discuss here; see [37]).

For a fixed graph G , identity (9) gives an algebraic relation between the entries of its hom-profile. It was proved by Whitney [62] that there are no other algebraic relations between these entries valid for all graphs G . A slightly stronger result was proved in [23]: *the projection of \mathcal{T}_0 to the coordinates corresponding to any finite set of connected graphs is full-dimensional*. This excludes any other kind of equations (e.g. exponential) between these numbers.

There are simple identities relating homomorphism numbers with the injective and induced versions. In order to spare the reader from separate provisos for each relation, we restrict ourselves to the case where both F and G are simple graphs. If Θ is any equivalence relation on $V(F)$, then we denote by F/Θ the graph obtained by identifying nodes that belong to the same class of Θ . Note that this may create loops and parallel edges.

We have some easy relations:

$$\text{hom}(F, G) = \sum_{\Theta} \text{inj}(F/\Theta, G) \quad (12)$$

and

$$\text{inj}(F, G) = \sum_{F' \supset F} \text{ind}(F', G), \quad (13)$$

where the sum runs over graphs $F' \supset F$ with the same node set.

From these, we can get reverse relations by Möbius inversion (or inclusion-exclusion):

$$\text{ind}(F, G) = \sum_{F' \supset F} (-1)^{|E(F') \setminus E(F)|} \text{inj}(F', G), \quad (14)$$

and

$$\text{inj}(F, G) = \sum_{\Theta} \mu(\Theta) \text{hom}(F/\Theta, G), \quad (15)$$

where the last sum runs over equivalence relations and

$$\mu(\Theta) = \prod_{A \in \Theta}^k \left((-1)^{|A|-1} (|A|-1)! \right),$$

with the product running over all classes $A \in \Theta$.

We have the following relations describing complementation (as an operation from simple graphs to simple graphs):

$$\text{ind}(F, \bar{G}) = \text{ind}(\bar{F}, G), \quad (16)$$

$$\text{inj}(F, \bar{G}) = \sum_{F' \subset F} (-1)^{|E(F')|} \text{inj}(F', G) \quad (17)$$

and

$$\text{hom}(F, \bar{G}) = \sum_{F' \subset F} (-1)^{|E(F')|} \text{hom}(F', G). \quad (18)$$

2.3 Examples of homomorphism functions

Example 2.2 (Stars and degrees) Let S_k denote the star with k nodes. Then for any graph G on n nodes,

$$\text{hom}(S_k, G) = \sum_{i=1}^n d_i^{k-1}, \quad (19)$$

where d_1, \dots, d_n are the degrees of G . Hence $\text{hom}(S_k, G)^{1/(k-1)}$ tends to the maximum degree of G as $k \rightarrow \infty$.

Example 2.3 (Cycles and eigenvalues) Let C_k denote the cycle on k nodes, and again let G be any graph on n nodes. Then

$$\text{hom}(C_k, G) = \sum_{i=1}^n \lambda_i^k, \quad (20)$$

where $\lambda_1, \dots, \lambda_n$ are the eigenvalues of the adjacency matrix of G . Hence $\text{hom}(C_{2k}, G)^{1/(2k)}$ tends to the largest eigenvalue of G as $k \rightarrow \infty$.

Example 2.4 (Independent sets) Let H be the graph on two nodes, with an edge connecting the two nodes and a loop at one of the nodes. Then for every simple graph G , $\text{hom}(G, H)$ is the number of independent sets of nodes in G .

Example 2.5 (Colorings) It is easy to see that $\text{hom}(G, K_q)$ is the number of colorings of the graph G with q colors. It is well known that for a fixed G , this number is a polynomial in q , called the *chromatic polynomial*. The chromatic polynomial defines a graph invariant for every complex number q , but this cannot be expressed as the number of homomorphisms into any graph unless q is a nonnegative integer [27, 28] (cf. Example 3.4).

It is often useful to consider homomorphisms into a fixed graph H as generalized colorings, where the colors are the nodes of H , and every edge of H imposes a constraint on the coloring that these two colors cannot be used at adjacent nodes.

Example 2.6 (Maximum cut) Let H denote the looped complete graph on two nodes, weighted as follows: the non-loop edge has weight 2; all other edges and nodes have weight 1. Then for every simple graph G with n nodes,

$$\log_2 \text{hom}(G, H) - n \leq \text{MaxCut}(G) \leq \log_2 \text{hom}(G, H).$$

where $\text{MaxCut}(G)$ denotes the size of the maximum cut in G . So unless G is very sparse, $\log_2 \text{hom}(G, H)$ is a good approximation of the maximum cut in G .

Example 2.7 (Random graphs) Let $G = G(n, p)$ be a random graph with n nodes and edge-density p . Then for every simple graph F with k nodes,

$$\mathbb{E}(\text{hom}(F, G)) = (1 + o(1))n^k p^{|E(F)|} \quad (n \rightarrow \infty).$$

By a straightforward application of high concentration results, it follows that $\text{hom}(F, G)$ is very close to its expectation with large probability.

Example 2.8 (Partition functions of the Ising model) Let G be any simple graph, and let $T > 0$, $h \geq 0$, and J be three real parameters. Let H be the looped complete graph on two nodes, denoted by $+$ and $-$, weighted as follows: $\alpha_+ = e^{h/T}$, $\alpha_- = e^{-h/T}$, $\beta_{++} = \beta_{--}$, $\beta_{+-} = \beta_{-+}$, and $\beta_{++}/\beta_{+-} = e^{2J/T}$. Then $\text{hom}(G, H)$ is the partition function of the Ising model on the graph G at temperature T with coupling J in external magnetic field h .

2.4 Homomorphisms into measurable functions

The following definition from [28] (which will play an important role later on) generalizes the homomorphism function. Every bounded function $W : [0, 1]^2 \rightarrow \mathbb{R}$ defines a graph parameter as follows: For a finite graph F on k nodes, let

$$t(F, W) = \int_{[0, 1]^k} \prod_{ij \in E(F)} W(x_i, x_j) dx_1 \dots dx_k.$$

(We can think of the interval $[0, 1]$ as the set of nodes, and of the value $W(x, y)$ as the weight of the edge xy .) While this definition is meaningful for all graphs F , we will mostly use it for simple graphs.

It is easy to see that for every weighted graph G , the graph parameter $t(\cdot, G)$ is a special case. We may assume that $V(G) = \{1, \dots, n\}$ and $\alpha(G) = 1$. Define a function $W_G : [0, 1]^2 \rightarrow [0, 1]$ as follows. For $(x, y) \in [0, 1]^2$, let a and b be determined by

$$\begin{aligned} \alpha_1(G) + \dots + \alpha_{a-1}(G) &\leq x < \alpha_1(G) + \dots + \alpha_a(G), \\ \alpha_1(G) + \dots + \alpha_{b-1}(G) &\leq y < \alpha_1(G) + \dots + \alpha_b(G), \end{aligned}$$

and let

$$W_G(x, y) = \beta_{ab}(G).$$

(Informally, W_G is obtained by replacing the (i, j) entry in the weighted adjacency matrix of G by a rectangle of size $\alpha_i \times \alpha_j$, and define the function value on this square as β_{ij} .) Then

$$t(F, G) = t(F, W_G)$$

for every finite simple graph G .

Example 2.9 For an undirected simple graph F , let $\text{eul}(F)$ denote the number of eulerian orientations of F (i.e., orientations in which every node has the same outdegree as indegree). By Euler's theorem, $\text{eul}(F) = 0$ if and only if F has a node with odd degree.

It can be shown [44] that this graph parameter can be represented in the form $t(\cdot, W)$, where

$$W(x, y) = 2 \cos(2\pi(x - y)).$$

On the other hand, it follows e.g. from Theorem 3.6 below that eul is not of the form $\text{hom}(\cdot, G)$ with any finite weighted graph G .

3 Connection matrices

3.1 The connection matrix of a graph parameter

A k -labeled graph ($k \geq 0$) is a finite graph in which k nodes are labeled by $1, 2, \dots, k$. Two k -labeled graphs are *isomorphic*, if there is a label-preserving isomorphism between them. We denote by K_k the k -labeled complete graph on k -nodes, and by O_k , the k -labeled graph on k nodes with no edges.

Let G_1 and G_2 be two k -labeled graphs. Their *product* $G_1 G_2$ is defined as follows: we take their disjoint union, and then identify nodes with the same label. Clearly this multiplication is associative and commutative. For 0-labeled graphs, this notation is in line with our notation for disjoint union.

The following construction is central to the theory of homomorphisms functions. Let f be any graph parameter. For every integer $k \geq 0$, we define the following (infinite) matrix $M(f, k)$. The rows and columns are indexed by isomorphism types of k -labeled graphs. The entry in the intersection of the row corresponding to G_1 and the column corresponding to G_2 is $f(G_1 G_2)$. We call the matrices $M(f, k)$ the *connection matrices* of the graph parameter f (see Figure 1).

For a simple graph parameter, the above construction causes trouble if we get multiple edges when gluing the two graphs. In this case, we suppress the edge multiplicities in $G_1 G_2$ when defining the entry corresponding to the pair (G_1, G_2) .

3.2 The rank of connection matrices

Connection matrices of a graph parameter are infinite matrices and their rank may be infinite. However, the rank is quite often finite, and if so, this fact has interesting consequences. Let us denote by $\text{rk}(f, k)$ the rank of the k -th connection matrix of the graph parameter f .

We start with several examples of graph parameters for which the rank of connection matrices is finite. The most important case for us will be when the graph parameter is defined as $\text{hom}(\cdot, H)$ for some fixed weighted graph H . This will be discussed in detail in the next section.

Example 3.1 (Edges) Let $\mathbf{e}(G) = |E(G)|$ denote the number of edges in G . Then $\mathbf{e}(G_1 G_2) = \mathbf{e}(G_1) + \mathbf{e}(G_2)$, and so $M(\mathbf{e}, k)$ is the sum of two matrices of rank 1. Thus $M(\mathbf{e}, k)$ has rank 2, so $\text{rk}(\mathbf{e}, k) = 2$ for all k .




Figure 1: A small part of the connection matrix for $k = 2$. The matrix entries are obtained by evaluating the parameter on the graph shown.

If we restrict $e(G)$ to simple graphs G to get a simple graph parameter e' , then the situation is more complicated: we have

$$e'(G_1G_2) = e'(G_1) + e'(G_2) - e'(G_1 \cap G_2).$$

Rewriting $e'(G_1 \cap G_2)$ as $x(G_1)^\top x(G_2)$ where $x(G)$ is the $\binom{k}{2}$ -dimensional vector with entries $x_{ij}(G) = 1$ if G contains an edge joining the labeled vertices i and j and $x_{ij}(G) = 0$ otherwise, we see that the matrix whose (G_1, G_2) entry is $e'(G_1 \cap G_2)$ has rank $\binom{k}{2}$, implying that $\text{rk}(e', k) \leq \binom{k}{2} + 2$. One can check that this is the exact value.

Example 3.2 (Subgraphs) Let $\text{subg}(G)$ denote the number of spanning subgraphs of G , i.e., $\text{subg}(G) = 2^{e(G)}$. Then $\text{subg}(G_1G_2) = \text{subg}(G_1)\text{subg}(G_2)$, and so $M(\text{subg}, k)$ has rank 1. Thus $\text{rk}(\text{subg}, k) = 1$ for all k .

Again, the version when we only consider simple graphs is more complicated: Let $\text{subg}'(G)$ denote this simple graph parameter. Then

$$\text{subg}'(G_1G_2) = \frac{\text{subg}'(G_1)\text{subg}'(G_2)}{\text{subg}'(G_1 \cap G_2)}.$$

The first two factors do not change the rank, and the rows of the matrix given by the second factor are determined by the edges induced by the labeled nodes, so it has only $2^{\binom{k}{2}}$ different rows, implying that $\text{rk}(\text{subg}', k) \leq 2^{\binom{k}{2}}$. Again one can check that this is the exact value.

Example 3.3 (Matchings) Let $\text{pmatch}(G)$ denote the number of perfect matchings in the graph G . It is trivial that $\text{pmatch}(G)$ is multiplicative. We claim that its node-rank-connectivity is exponentially bounded:

$$r_{\text{pmatch}}(k) \leq 2^k.$$

Let G be a k -labeled graph, let $X \subseteq [k] = \{1, \dots, k\}$, and let $\text{pmatch}(G, X)$ denote the number of matchings in G that match all the unlabeled nodes and the nodes with label in X , but not any of the other labeled nodes. Then we have for any two k -labeled graphs G_1, G_2

$$\text{pmatch}(G_1G_2) = \sum_{X_1 \cap X_2 = \emptyset, X_1 \cup X_2 = [k]} \text{pmatch}(G_1, X_1)\text{pmatch}(G_2, X_2).$$

This can be read as follows: The matrix $M(\text{pmatch}, k)$ can be written as a product $N^T W N$, where N has infinitely many rows indexed by k -labeled graphs, but only 2^k columns, indexed by subsets of $[k]$,

$$N_{G,X} = \text{pmatch}(G, X),$$

and W is a symmetric $2^k \times 2^k$ matrix, where

$$W_{X_1, X_2} = \begin{cases} 1 & \text{if } X_1 = [k] \setminus X_2, \\ 0 & \text{otherwise.} \end{cases}$$

Hence the rank of $M(\text{pmatch}, k)$ is at most 2^k (it is not hard to see that in fact equality holds).

Example 3.4 (Chromatic polynomial) We have seen that the number of q -colorings is a special case of homomorphism functions. This number is the evaluation of the chromatic polynomial $\text{chr}(G; x)$ at nonnegative integers q . What about evaluations at other values? It turns out that these evaluations violate both conditions in Theorem 3.6 below [27]. For every fixed x , this is a multiplicative graph parameter. To describe its rank-connectivity, we need the following notation. For $k, q \in \mathbb{Z}_+$, let $B_{k,q}$ denote the number of partitions of a k -element set into at most q parts. So $B_k = B_{k,k}$ is the k -th Bell number. With this notation,

$$\text{rk}(\text{chr}, k) = \begin{cases} B_{k,x} & \text{if } x \text{ is a positive integer,} \\ B_k & \text{otherwise.} \end{cases}$$

Note that this is always finite: if x is a positive integer, then it is bounded by x^k , but otherwise it grows faster than c^k for every c . Similar results can be derived for the Tutte polynomial, where the exceptional values are the hyperbolas in the Tutte plane for which $(x-1)(y-1)$ is a positive integer.

Let f be a graph parameter that is not identically 0. Then f is multiplicative if and only if $f(K_0) = 1$ (K_0 is the empty graph) and $\text{rk}(f, 0) = 1$. Every multiplicative graph parameter f satisfies the inequality

$$\text{rk}(f, k+l) \geq \text{rk}(f, k) \cdot \text{rk}(f, l). \quad (21)$$

In the most important special case (to be discussed below) when f is a homomorphism function, a stronger version of property (21) holds: the sequence $\text{rk}(f, k)$ is logconvex. We do not know if this property holds for more general graph parameters.

Finiteness of the rank connectivity function has interesting algorithmic consequences:

Theorem 3.5 [27] *If $r(f, k)$ is finite for some k , then f can be computed in polynomial time for graphs with treewidth at most k .*

3.3 Connection matrices of homomorphisms

Homomorphism functions, which are our main concern, provide the most important class of graph parameters for which connection matrices have finite rank. In fact, connection matrices can be used to characterize these parameters, as the following theorem of Freedman, Lovász and Schrijver shows.

Theorem 3.6 [28] *The graph parameter f , defined on graphs with multiple edges but no loops, is equal to $\text{hom}(\cdot, H)$ for some weighted graph H on q nodes if and only if*

- (a) $M(f, k)$ is positive semidefinite and
- (b) $\text{rk}(f, k) \leq q^k$ for all k .

In terms of statistical physics, this theorem can be viewed as a characterization of partition functions of models whose degrees of freedom sit on vertices (as opposed to the edge coloring models considered below). The property that $M(f, k)$ is positive semidefinite is related to the “reflection positivity” property in statistical physics, and we will call a graph parameter *reflection positive* if $M(f, k)$ is positive semidefinite for every k .

The proof of the necessity of the conditions in Theorem 3.6 is easy and it is instructive to present it here. (The sufficiency is more involved, and the proof is based on algebraic considerations.)

We need the following notation: For any k -labeled graph G and mapping $\phi : [k] \rightarrow V(H)$, let

$$\text{hom}_\phi(G, H) = \sum_{\substack{\psi: V(G) \rightarrow V(H) \\ \psi \text{ extends } \phi}} \frac{\alpha_\psi}{\alpha_\phi} \text{hom}_\psi(G, H), \quad (22)$$

so that

$$\text{hom}(G, H) = \sum_{\phi: [k] \rightarrow V(H)} \alpha_\phi \text{hom}_\phi(G, H). \quad (23)$$

For any two k -labeled graph G_1 and G_2 ,

$$\text{hom}_\phi(G_1 G_2, H) = \text{hom}_\phi(G_1, H) \text{hom}_\phi(G_2, H). \quad (24)$$

The decomposition (23) writes the matrix $M(f, k)$ as the sum of $|V(H)|^k$ matrices, one for each mapping $\phi : [k] \rightarrow V(H)$; (24) shows that these matrices are positive semidefinite and have rank 1.

In the presence of condition (a), condition (b) in Theorem 3.6 can be replaced by the following quite different type of condition. To formulate it, we need the notion of a quantum graph, defined as a formal linear combination of graphs with real coefficients; but the condition concerns only the existence of a single 2-labeled quantum graph [45]:

(c) *There is a 2-labeled quantum graph g_0 with the following property: if G is a 2-labeled graph having no edge between the labeled nodes, and G' denotes the graph obtained from G by identifying the two labeled nodes, then $f(g_0 G) = f(G')$.*

In other words, attaching g_0 at two nodes is effectively the same as identifying the two nodes (this is why it is called a “contractor” in [45]).

Let us conclude with a discussion of the independence of the two conditions in Theorem 3.6.

In Example 3.3 we saw that the number of perfect matchings in a graph provides an example for a graph parameter for which the rank of connection matrices grows simply exponentially. This parameter is also multiplicative, so for $k = 0$ the connection matrix is positive semidefinite. But it is easy to see that for $k = 1$, the submatrix indexed by K_1 and K_2 is

$$\begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}$$

which is not positive semidefinite. Thus the number of perfect matchings cannot be represented as a homomorphism function.

Recalling Examples 3.1 and 3.2, let us define the multigraph parameter f by $f(G) = 1/\text{subg}'(G) = 2^{-e(G')}$, where G' is obtained from G by removing duplicate edges. As in Example 3.2, the rank of the connection matrix $M(f, k)$ grows as $2^{\binom{k}{2}}$. It is further not hard to check that $M(f, k)$ is positive semidefinite. The graph parameter f is, in fact, the limit of parameters of the form $\text{hom}(\cdot, H)$: take homomorphisms into a random graph $H = \mathbf{G}(n, 1/2)$, with all nodeweights $1/n$ and all edge-weights 1. But the rank of its connection matrices is finite but superexponential, so the parameter is not of the form $\text{hom}(\cdot, H)$.

This example also illustrates the importance of the condition that f is defined on graphs with multiple edges for the validity of 3.6. Indeed, for a simple graph G (*i.e.*, if G has no multiple

edges), $f(G) = 2^{-e'(G)}$ can be represented as the number of homomorphisms into the graph $K_1(1/2)$, consisting of a single node with a loop, where the node has weight 1 and the loop has weight $1/2$.

The chromatic polynomial (Example 3.4) was another example whose connection matrices had superexponential rank growth if the variable x was not a nonnegative integer. Here the reflection positivity condition gives the same condition on x : the k -th connection matrix is positive semidefinite if and only if either x is a positive integer or $k \leq x + 1$. Thus $M(\text{chr}, k)$ is semidefinite for all k if and only if x is a positive integer.

3.4 The exact rank of connection matrices for homomorphisms

How good is the upper bound on the rank given in Theorem 3.6? It can be proved [39] that equality holds in the “generic” case.

One reason why the upper bound is not always reached are twins. As remarked earlier, twin-reduction in H does not change the numbers $\text{hom}(G, H)$, and so it does not change the connection matrices (but of course it decreases the upper bounds $|V(H)|^k$). So we may assume that H is twin-free.

The second reason for rank loss in connection matrices is that if H has a proper automorphism (a permutation of the nodes that preserves both the nodeweights and edgeweights), then in formula (22), any two terms defined by a mapping $\phi : [k] \rightarrow V(H)$ and $\phi\sigma$ ($\sigma \in \text{Aut}(H)$) are equal, so the sum of all such terms is still rank 1. So the rank of $M(\text{hom}(\cdot, H), k)$ is at most the number of orbits of the automorphism group of H on ordered k -tuples of its nodes.

Theorem 3.7 [39] *Assume that the target graph H is twin-free. Then for every k , $\text{rk}(\text{hom}(\cdot, H), k)$ is the number of orbits of the automorphism group of H on ordered k -tuples of its nodes.*

It is worthwhile to formulate two corollaries.

Corollary 3.8 *Let H be a weighted graph that has no twins and no automorphisms. Then $\text{rk}(\text{hom}(\cdot, H), k) = |V(H)|^k$ for every k .*

Swapping twins i and j is “almost” an automorphism: the only additional condition needed is that $\alpha_i = \alpha_j$. In particular, for unweighted graphs the condition that there are no automorphisms implies that there are no twins.

Our second corollary is in fact an equivalent reformulation of Theorem 3.7 in the framework of quantum graphs. To state it, we need the following notation: given a weighted graph H , a k -labeled quantum graph x and nodes $i_1, \dots, i_k \in V(H)$, we define $\text{hom}_{i_1 \dots i_k}(x, H)$ to be the number of homomorphisms from x to H such that the labeled nodes are mapped into $i_1 \dots i_k$.

Corollary 3.9 *Let H be a weighted graph that has no twins, and let $h : V(H)^k \rightarrow \mathbb{R}$. Then there exists a k -labeled quantum graph x such that $\text{hom}_{i_1 \dots i_k}(x, H) = h(i_1, \dots, i_k)$ for all $i_1, \dots, i_k \in V(H)$ if and only if h is invariant under the automorphisms of H .*

3.5 Extensions: directed graphs, hypergraphs, semigroups

In Theorem 3.6 we allowed parallel edges in the graphs G , but no loops. Indeed, the representation theorem is false if G can have loops: it is not hard to check that the graph parameter

$$\text{loop}(G) = 2^{-\#\text{loops}}$$

cannot be represented as a homomorphism function, even though its connection matrix $M(\text{loop}, k)$ is positive semidefinite and has rank 1. To get a representation theorem for graphs

with loops, each loop e in the target graph H must have two weights: one which is used when a non-loop edge of G is mapped onto e , and the other, when a loop of G is mapped onto e . With this modification, the theorem remains valid.

The constructions and results above are in fact more general; they extend to directed graphs and hypergraphs. One new element in the case of directed graphs is the following. For a directed graph, homomorphism functions can be defined by the same formulas (4) and (5), except that weights are now assigned to the directed edges of G and H . But there are (at least) two substantially different ways to define connection matrices.

(1) The easier way is to define k -labeled digraphs similarly, and glue them together just like we did in the undirected case. The related theorem and proof are precisely the same as above.

(2) The alternative generalization of Theorem 3.6 to directed graphs is a bit more interesting. (Example 3.13 below is a case when this second theorem applies.) For this, we consider weighted directed graphs in which the edgeweights can be complex. Such a graph H is called *Hermitian* if for every arc uv with (complex) weight β_{uv} , the arc vu is also present and has weight $\beta_{vu} = \overline{\beta_{uv}}$.

For any directed graph D , let D^* be the digraph obtained from D by reversing all arcs. For two k -labeled digraphs D and D' , let DD' denote, as before, their union with the labeled nodes identified. For any complex-valued digraph parameter f defined on loopless directed graphs and for each natural number k , we define the matrix $\tilde{M}(f, k)$ as follows: its rows and columns are indexed by k -labeled directed graphs, and the entry in position D_1, D_2 is $f(D_1^* D_2)$.

Theorem 3.10 [40] *Let f be a complex valued digraph parameter. Then $f = \text{hom}(\cdot, H)$ for some Hermitian weighted digraph H if and only if $f(K_0) = 1$ and there exists a $d \geq 0$ such that, for each $k \geq 0$, \tilde{M}_k is positive semidefinite and has rank at most d^k .*

Note that \tilde{M}_k is a complex valued matrix. The condition that it is positive semidefinite includes the condition that it is Hermitian.

There is a common formulation of these results, using semigroups; see [40] for details.

3.6 Edge coloring models

Let G be a finite graph. An *edge coloring model* or *edge model* is determined by a finite set C and a mapping $h : \mathbb{Z}_+^C \rightarrow \mathbb{R}_+$, which we call the *node evaluation function*. Here C is the set of possible edge colors; for any coloring of the edges, we think of $h(a)$ as the value of a node incident with $a(c)$ edges with the color c ($c \in C$). In terms of statistical physics, an edge coloring is a state of the system, and $\log h(a)$ is the contribution of a node (incident with $a(c)$ edges with the color c) to the energy of the state.¹

To be more precise, for an edge-coloring $\phi : E(G) \rightarrow C$ and node v , let $a_{\phi, v}(c)$ denote the number of edges e incident with v with $\phi(e) = c$. So $a_{\phi, v} \in \mathbb{Z}_+^C$ is the “local view” of node v . The weight of the assignment ϕ is defined by

$$w(\phi) = \prod_{v \in V(G)} h(a_{\phi, v}),$$

and the *edge coloring parameter*, by

$$\text{col}(G, h) = \sum_{\phi: E(G) \rightarrow C} w(\phi).$$

(It will be also useful to allow a single edge with no endpoints; we call this graph the *circle*, and denote it by \bigcirc . By definition, $\text{col}(\bigcirc, h) = |C|$.)

We can define edge-connection matrices that are analogous to the connection matrices defined before: Instead of gluing graphs together along nodes, we glue them together along edges. To be

¹For this reason, these models are usually called vertex models in the physics literature.

precise, we define a *k-broken graph* as a *k*-labeled graph in which the labeled nodes have degree one. (It is best to think of the labeled nodes not as nodes of the graph, but rather as points where the *k* edges sticking out of the rest of the graph are broken off.) We allow that both ends of an edge be broken off.

For two *k*-broken graphs G_1 and G_2 , we define $G_1^*G_2$ by gluing together the corresponding broken ends of G_1 and G_2 . These ends are not nodes of the resulting graph any more, so $G_1^*G_2$ is different from the graph G_1G_2 we would obtain by gluing together G_1 and G_2 as *k*-labeled graphs. One very important difference is that while G_1G_2 is *k*-labeled, $G_1^*G_2$ has no broken edges any more, and so it is not *k*-broken. This fact leads to considerable difficulties in the treatment of edge models.

For every graph parameter f and integer $k \geq 0$, we define the *edge-connection matrix* $M'(f, k)$ as follows. The rows and columns are indexed by isomorphism types of *k*-broken graphs. The entry in the intersection of the row corresponding to G_1 and the column corresponding to G_2 is $f(G_1^*G_2)$. Note that for $k = 0$, we have $M(f, 0) = M'(f, 0)$, but for other values of k , connection and edge-connection matrices are different. We say that f is *edge reflection positive*, if $M'(f, k)$ is positive semidefinite for every $k \geq 0$.

It is easy to see (similarly as in the case of homomorphism functions) that if $h : \mathbb{Z}_+^C \rightarrow \mathbb{R}_+$ and $f = \text{col}(\cdot, h)$ is an edge-coloring function, then

$$\text{rk}(M'(f, k)) \leq |C|^k,$$

and $M'(f, k)$ is positive semidefinite. Unlike in the case of node-connection matrices, these two properties are not independent any more:

Proposition 3.11 [58] *If f is a multiplicative graph parameter such that $M'(f, k)$ is positive semidefinite for every $k \geq 0$, then $f(\circ)$ is a nonnegative integer and $\text{rk}(M'(f, k)) \leq f(\circ)^k$.*

The analogue of Theorem 3.6 is even simpler to state (but much more difficult to prove):

Theorem 3.12 [58] *A graph parameter f can be represented as $f(\cdot) = \text{col}(\cdot, h)$ for some edge coloring model h if and only if it is multiplicative and edge-reflection positive.*

Just as for homomorphism functions, it is natural to ask what determines the rank of connection matrices of edge models. This question seems to lead to difficult algebraic questions in group representations, and is unanswered at this time.

3.7 Edge colorings and homomorphisms

The connection between homomorphism functions and edge coloring functions seems to go farther than analogy, but it is not well understood.

In one direction, edge coloring functions are more general than homomorphism functions. This connection is easy in the directed case. We can generalize the edge coloring model to directed graphs in the obvious way. It is easy to see that the directed edge model is more general than the directed homomorphism model: If we are given a pair $H = (a, B)$ ($a \in \mathbb{R}_+^q$, $B \in \mathbb{C}^{q \times q}$, with entries α_c and β_{uv} , resp.), then to every homomorphism $\phi : V(G) \rightarrow [q]$ we can assign an edge-coloring in which the edge ij is colored with the pair $\psi(ij) = (\phi(i), \phi(j))$. The evaluation function at a node v is given as follows: if there is a color c such that all the outgoing edges have the same first color c and all the incoming edges have the same second color c , then the value is

$$\alpha_c \prod_{u: uv \in E(G)} \beta_{\psi(uv)_1, \psi(uv)_2};$$

otherwise, the node evaluates to 0. It is easy to see that an edge-coloring has nonzero weight only if it comes from a homomorphism, and in that case, the weight of the edge-coloring is the same as the weight of the corresponding homomorphism.

It is not obvious, but it is true, that undirected edge coloring functions generalize undirected homomorphism functions [58], at least if complex values are allowed for h . It is not clear which (real) homomorphism functions can be obtained as edge coloring functions with a real valued h .

In the opposite direction, edge coloring models cannot be translated into node coloring models (homomorphisms) in general; but there are some nontrivial examples of important graph parameters that are defined as edge coloring functions, but that can also be represented as homomorphism functions in a nontrivial way. A general understanding of these examples would be very interesting.

Example 3.13 (Nowhere-zero flows) Let $\text{eul}(G) = 1$ if G is eulerian (i.e., all nodes have even degree), and $\text{eul}(G) = 0$ otherwise. To represent this function as a homomorphism function, let

$$a = \begin{pmatrix} 1/2 \\ 1/2 \end{pmatrix}, \quad B = \begin{pmatrix} 1 & -1 \\ -1 & 1 \end{pmatrix}$$

It was noted by de la Harpe and Jones [33] that for the weighted graph $H = (a, B)$ we have $\text{hom}(G, H) = \text{eul}(G)$.

This example can be generalized quite a bit. Let Γ be a finite abelian group and let $S \subseteq \Gamma$ be such that S is closed under inversion. For any graph G , fix an orientation of the edges. An S -flow is an assignment of an element of S to each edge such that for each node v , the product of elements assigned to edges entering v is the same as the product of elements assigned to the edges leaving v . Let $\text{sflo}(G)$ be the number of S -flows. This number is independent of the orientation.

The choice $\Gamma = \mathbb{Z}_2$ and $S = \mathbb{Z}_2 \setminus \{0\}$ gives the special case above (incidence function of eulerian graphs). If $\Gamma = S = \mathbb{Z}_2$, then $\text{sflo}(G)$ is the number of eulerian subgraphs of G . Perhaps the most interesting special case is when $|\Gamma| = t$ and $S = \Gamma \setminus \{0\}$, which gives the number of nowhere zero t -flows.

Surprisingly, this parameter (which is an edge coloring model) can be described as a homomorphism function. Let Γ^* be the character group of Γ . Let H be the complete directed graph (with all loops) on Γ^* . Let $\alpha_\chi := 1/|\Gamma|$ for each $\chi \in \Gamma^*$, and let

$$\beta_{\chi\chi'} := \sum_{s \in S} \chi^{-1}(s)\chi'(s),$$

for $\chi, \chi' \in \Gamma^*$. Using arguments related to duality transformations of models in statistical physics (see e.g. [22] and references therein) one can show [28] that this weighted graph H represents sflo in the sense that $\text{sflo}(\cdot) = \text{hom}(\cdot, H)$. The condition on S that it is closed under inversion can be dropped if we use homomorphisms of directed graphs (Section 3.5).

In statistical physics, negative—and more generally, complex—nodeweights correspond to complex magnetic fields, which arise in the well-known Lee-Yang theory of phase transitions [9, 36, 60]. The next example [58] shows that it is also interesting in our context to extend the definition of weighted graphs by allowing negative nodeweights (a direction we will not pursue here except for this example).

Example 3.14 (Matchings revisited) We have seen that the number $\text{pmatch}(G)$ of perfect matchings has exponential rank connectivity but is not reflection positive, and hence it is not a homomorphism function. However, consider the following weighted graph H_x : We take a looped complete graph on two nodes u and v , and define

$$\alpha(u) = \frac{1}{x}, \quad \alpha(v) = -\frac{1}{x},$$

and

$$\beta(uu) = x + 1, \quad \beta(uv) = \beta(vv) = 1.$$

Then the following surprising fact holds:

$$\lim_{x \rightarrow 0} \text{hom}(G, H_x) = \text{pmatch}(G)$$

for every graph G .

4 Convergence and limit

4.1 Quasi-random graphs

Quasirandom (also called pseudorandom) graphs were introduced by Thomason [59] and Chung, Graham and Wilson [19]. These graphs have many properties that true random graphs have.

A sequence $(G_n : n = 1, 2, \dots)$ of graphs is called *quasirandom with density p* (where $0 < p < 1$), if for every simple finite graph F ,

$$t(F, G_n) = (1 + o(1))p^{|E(F)|}. \quad (25)$$

(this is the asymptotic number of labeled copies of F in a random graph with edge probability p). The definition is usually formulated in terms of the number of injections (labeled copies) of F into G , but the two differ only in lower order terms, which are swallowed by the $o(1)$ in the definition.

It turns out that (25) implies many other properties that are familiar from the theory of random graphs; for example, almost all degrees are about pn , almost all codegrees are about p^2n etc. Many of these properties characterize quasirandom graphs, and so these provide many equivalent ways to define a quasirandom sequence [19, 59]. Quasirandomness is closely related to Szemerédi's lemma [56, 57]. One of the most surprising facts proved in [19] is that it is enough to require the condition about the number of copies of F for just two graphs, namely K_2 (which just defines the edge density p) and the 4-cycle C_4 . This fact can be stated and proved in a simpler way using the "limit set" \mathcal{T}_0 defined in Section 2.1:

Theorem 4.1 *If $t \in \mathcal{T}_0$ satisfies $t(C_4) = t(K_2)^4$, then for every simple graph H ,*

$$t(H) = t(K_2)^{|E(H)|}.$$

In other words, $t(H)$ is the expected profile of a random graph $G(n, p)$ with $p = t(K_2)$; it is also the profile of the weighted graph consisting of a single node and a loop with weight p (the weight of the node does not matter).

To illustrate the power of reflection positivity, we give a proof of this theorem (the proof goes along the lines of the original, just the details are simpler).

Proof. Let $p = t(K_2)$. We first prove the conclusion for stars $K_{1,j}$:

$$t(K_{1,j}) = p^j. \quad (26)$$

Starting with $t(K_{1,2})$, let us first consider the connection matrix $M(t, 1)$ and its 2×2 submatrix formed by the rows and columns corresponding to the graph K_1 and K_2 . Positive semidefiniteness of this matrix gives $t(K_{1,2}) \geq t(K_2)^2 = p^2$. On the other hand $t(K_{1,2}) \leq \sqrt{t(C_4)} = p^2$ by positive semidefiniteness of the 2×2 submatrix of $M(t, 2)$ indexed by $K_{1,2}$ (with its endpoints labeled) and \bar{K}_2 , the empty graph on two nodes. The above two inequalities give $t(K_{1,2}) = p^2$, which proves (26) for $j = 2$. To prove the identity for $j > 2$, we again consider the connection matrix $M(t, 1)$. By the identity we just established, its 2×2 submatrix formed by the rows and columns corresponding to the graph K_1 and K_2 has 0 determinant. By positive semidefiniteness,

the corresponding two rows of the whole connection matrix $M(t, 1)$ are proportional. But this means that

$$t(K_{1,j+1}) = pt(K_{1,j})$$

for every j , from which (26) follows by induction.

Next we show that for all complete bipartite graphs $K_{2,j}$:

$$t(K_{2,j}) = p^{2j}. \tag{27}$$

Since $t(K_{2,2}) = t(C_4) = p^4$ by assumption, this is true for $j = 2$. For general j , it follows just like (26) from the positive semidefiniteness of the matrix $M(t, 2)$.

Now we prove the general case by a similar induction. Let us view the graph H as glued together from a star $K_{1,d}$ and a graph F on one fewer nodes, along the set T of the leaves of the star, and suppose that we know the assertion for F . Consider the matrix $M(t, d)$ and its 2×2 submatrix formed by the rows and columns indexed by \overline{K}_d (the graph on d labeled nodes with no edge) and $K_{1,d}$ (with the leaves labeled). By (26) and (27), this submatrix is singular, and hence these two rows of the whole matrix are proportional, the second row is p^d times the first. But the graphs F and H define two elements of these rows above each other, so

$$t(H) = p^d t(F) = p^d p^{|E(F)|} = p^{|E(H)|},$$

which proves the theorem. □

4.2 Convergent sequences

Let (G_n) be a sequence of unweighted simple graphs, and assume again that $|V(G_n)| = n$. We say that this sequence is *convergent*, if the sequence $t(F, G_n)$ has a limit for every simple graph F . Note that it would be enough to assume this for connected graphs F .

In the definition, we could replace the homomorphism function by the number of embeddings (injective homomorphisms), with appropriate normalization. Indeed, the difference between the number of homomorphisms and embeddings is the number of non-injective homomorphisms, which is of lower order, so it tends to 0 when divided by $n^{|V(F)|}$.

We could also replace the homomorphism function by the number of embeddings as induced subgraphs. Indeed, the number of embeddings can be obtained by summing the numbers of induced embeddings over all supergraphs (on the same set of nodes). Conversely, the number of induced embeddings can be expressed in terms of the numbers of embeddings of supergraphs by inclusion-exclusion.

Example 4.2 Let $\mathbf{G}(n, p)$ be a random graph on n nodes with edge-density p ; the sequence $(\mathbf{G}(n, p), n = 1, 2, \dots)$ is convergent with probability 1. The limiting simple graph parameter is given by $t(F) = p^{|E(F)|}$.

By definition, every quasirandom graph sequence with density p is also convergent, and the homomorphism densities into it tend to the same value.

4.3 Finite limits, a.k.a. generalized quasirandom graphs

A *generalized random graph* $\mathbf{G}(n; H)$ is defined by the number n of its nodes and by a weighted “model” graph H . We assume that $V(H) = [q]$ and set $\alpha_i = \alpha_i(H)$ ($i = 1, \dots, q$) and $\beta_{ij} = \beta_{ij}(H)$ ($i, j = 1, \dots, q$). We partition $[n]$ into q classes V_1, \dots, V_q , by putting each $u \in [n]$ into V_i with probability α_i and connect each pair $u \in V_i$ and $v \in V_j$ with probability β_{ij} (all these decisions are made independently).

A *generalized quasirandom graph sequence* (G_n) with model graph H (or briefly *H-quasirandom sequence*) is defined by the property that for every fixed finite graph F ,

$$t(F, G_n) \longrightarrow t(F, H) \quad (n \rightarrow \infty).$$

In other words, the number of homomorphisms of F into G_n is approximately the same as the expected number of homomorphisms of F into a generalized random graph $\mathbf{G}(N, H)$ on $N = |V(G_n)|$ nodes.

This definition suggests that we should consider the graph H as the “limit” of the H -quasirandom sequence. The definition of a quasirandom sequence of graphs (with edge-density p) is equivalent to saying that the sequence converges to $K_1(p)$. (Warning: not every convergent sequence will have a limit of this form!)

In view of the theory of quasirandom graphs, we can ask the following two basic questions concerning generalized quasirandom graphs:

(a) Is it enough to require the condition concerning the number of copies of F for a finite set of graphs F_i (depending on α and β)?

(b) Is the structure of a generalized quasirandom graph G_n similar to a generalized random graph?

To be more precise, we want that the nodes of G_n can be partitioned into q classes U_1, \dots, U_q of sizes $\alpha_1 n, \dots, \alpha_q n$ so that the graph spanned by U_i is quasirandom with density β_{ii} , and the bipartite graph formed by the edges between U_i and U_j is quasirandom with density β_{ij} .

The answer to the first two questions is in the affirmative. More precisely, the following theorems hold.

Theorem 4.3 [43] *Let H be a weighted graph with $V(H) = [q]$, nodeweights $(\alpha_i : i = 1, \dots, q)$ and edgeweights $(\beta_{ij} : i, j = 1, \dots, q)$. Let $(G_n, n = 1, 2, \dots)$ be an H -quasirandom sequence of unweighted simple graphs. Then for every n there exists a partition $V(G_n) = \{U_1, \dots, U_q\}$ such that*

$$(a) \quad \frac{|U_i|}{|V(G_n)|} \rightarrow \alpha_i \quad (i = 1, \dots, q),$$

(b) *the subgraph of G_n induced by U_i is a quasirandom graph sequence with edge density β_{ii} , and*

(c) *the bipartite subgraph between U_i and U_j is a quasirandom bipartite graph sequence with density β_{ij} .*

Theorem 4.4 [43] *Let H be a weighted graph with $V(H) = [q]$. A sequence $(G_n, n = 1, 2, \dots)$ is H -quasirandom if and only if*

$$t(F, G_n) \longrightarrow t(F, H) \quad (t \rightarrow \infty)$$

for every graph F with at most $(10q)^q$ nodes.

4.4 The general limit object

4.4.1 Limits as reflection positive parameters

We now turn to describing limits of general convergent graph sequences.

Let $\tilde{\mathcal{T}}_0$ be the set of homomorphism density functions $t(\cdot, G)$ defined on simple (unweighted) graphs, where G is any simple unweighted target graph. Let \mathcal{T}_0 be the set of all graph parameters that are pointwise limits of graph parameters in $\tilde{\mathcal{T}}_0$ (i.e., its closure in the product topology on $\mathbb{R}^{\mathcal{F}}$). It is not hard to see that \mathcal{T}_0 would not change if we allowed weighted target graphs with edge weights between 0 and 1.

The characterization of homomorphism functions (Theorem 3.6) extends to the limit, at least for simple graphs:

Theorem 4.5 [44] *A simple graph parameter f is in \mathcal{T}_0 if and only if f is normalized, multiplicative and reflection positive.*

4.4.2 Limits as measurable functions

Graph parameters in the set \mathcal{T}_0 can be represented as homomorphism functions into measurable functions [44]. Let \mathcal{W} denote the set of all bounded measurable functions $W : [0, 1]^2 \rightarrow \mathbb{R}$ such that $W(x, y) = W(y, x)$ for all $x, y \in [0, 1]$. We also introduce the set

$$\mathcal{W}_0 = \{W \in \mathcal{W} : 0 \leq W \leq 1\}.$$

Theorem 4.6 [44] *A simple graph parameter f is in \mathcal{T}_0 if and only if there is a function $W \in \mathcal{W}_0$ such that $f = t(\cdot, W)$.*

This function W is not unique: for example, $W(1-x, 1-y)$ will define the same graph parameter. More generally, if $\phi : [0, 1] \rightarrow [0, 1]$ is a measure preserving map (not necessarily bijective), then

$$W^\phi(x, y) = W(\phi(x), \phi(y))$$

defines the same parameter. The following theorem says that this is all: Let us call functions $W_1, W_2 \in \mathcal{W}$ *equal up to measure preserving transformation* if there is a third function $W \in \mathcal{W}$ and measure preserving maps $\phi_1, \phi_2 : [0, 1] \rightarrow [0, 1]$ such that $W_i = W^{\phi_i}$.

Theorem 4.7 [14] *Two functions $W_1, W_2 \in \mathcal{W}$ define the same simple graph parameter if and only if they are equal up to measure preserving transformation.*

4.4.3 Limits as distributions over finite and countable graphs

Given any function $W \in \mathcal{W}_0$ and an integer $n > 0$, we can generate a random graph $\mathbf{G}(n, W)$, called a *W-random graph*, on node set $[n]$ as follows. We generate n independent numbers X_1, \dots, X_n from the uniform distribution on $[0, 1]$, and then connect nodes i and j with probability $W(X_i, X_j)$.

As a special case, if W is the identically p function, we get “ordinary” random graphs $\mathbf{G}(n, p)$. More generally, if $W = W_H$ for a (finite) weighted graph H , then $\mathbf{G}(n, W_H)$ is a quasirandom graph with model H .

Theorem 4.8 [44] *With probability 1, the graph sequence $\mathbf{G}(n, W)$ is convergent, and its limit is the function W .*

Let us define a *random graph model* as a distribution \mathbf{G}_n on simple graphs on $[n]$, for every $n \in \mathbb{Z}_+$. The random graph model $\mathbf{G}(n, W)$ defined above has the following three obvious properties:

- (i) The distribution of \mathbf{G}_n is invariant under relabeling nodes;
- (ii) If we delete node n from \mathbf{G}_n , the distribution of the resulting graph is the same as the distribution of \mathbf{G}_{n-1} ;
- (iii) for every $1 < k < n$, the subgraphs of \mathbf{G} induced by $[k]$ and $\{k+1, \dots, n\}$ are independent (as random variables).

It turns out that these three properties characterize the model $\mathbf{G}(n, W)$:

Theorem 4.9 [44] *A random graph model is of the form $\mathbf{G}(n, W)$ for some function $W \in \mathcal{W}_0$ if and only if it satisfies conditions (i), (ii) and (iii). Furthermore, two functions $W_1, W_2 \in \mathcal{W}_0$ define the same random graph model if and only if they are equal up to measure preserving transformation.*

We have seen that the limit of ordinary random graphs $\mathbf{G}(n, 1/2)$ is the function $W \equiv 1/2$. It is, however, quite natural to think that the limit of ordinary random graphs should be the Rado graph (the countable random graph). It turns out that this is also true in the following sense: For every $W \in \mathcal{W}_0$, we can define a countable random graph $\mathbf{G}(\omega, W)$ on \mathbb{Z}_+ , by choosing an infinite sequence (X_0, X_1, \dots) of independent uniform samples from $[0, 1]$, and connecting i and j with probability $W(X_i, X_j)$. A *countable random graph model* is a probability distribution on graphs on \mathbb{Z}_+ (with the σ -algebra generated by cylinders consisting of all graph containing a given edge).

Theorem 4.10 [48] *A countable graph model is of the form $\mathbf{G}(\omega, W)$ if and only if it satisfies (i) and (iii) above. Furthermore, the countable graph model $\mathbf{G}(\omega, W)$ determines W up to measure preserving transformation.*

Thus it is justified to say that with probability 1, $\mathbf{G}(n, 1/2)$ converges to the Rado graph $\mathbf{G}(\omega, 1/2)$. A word of caution is warranted here: as an unlabeled graph, $\mathbf{G}(\omega, 1/2)$ is isomorphic to $\mathbf{G}(n, 1/3)$ with probability 1. So viewing the Rado graph as an unlabeled graph would not contain enough information to characterize the limit; we have to view it as a probability distribution over graphs on a fixed countable set of nodes.

4.4.4 Examples

Example 4.11 Consider the *half-graphs* $H_{n,n}$: they are bipartite graphs on $2n$ nodes $\{1, \dots, n, 1', \dots, n'\}$, where i is connected to j' if and only if $i \leq j'$. It is easy to see that this sequence is convergent. Indeed, let F be a simple graph with k nodes; we show that the limit of $t(F, H_{n,n})$ exists. We may assume that F is connected. If F is non-bipartite, then $t(F, H_{n,n}) = 0$ for all n , so suppose that F is bipartite; let $V(F) = V_1 \cup V_2$ be its (unique) bipartition. Then every homomorphism of F into H preserves the 2-coloring, and so the homomorphisms split into two classes: those that map V_1 into $\{1, \dots, n\}$ and those that map it into $\{1', \dots, n'\}$. By the symmetry of the half-graphs, these two classes have the same cardinality.

Now F defines a partial order P on $V(F)$, where $u \leq v$ if and only if $u = v$ or $u \in V_1$, $v \in V_2$, and $uv \in E$. With respect to this partial order, $\frac{1}{2}\text{hom}(F, H_{n,n})$ is just the number of order-preserving maps from $V(F)$ to the chain $\{1, \dots, n\}$, and so

$$2^{k-1}t(F, H_{n,n}) = 2^{k-1} \cdot \frac{\text{hom}(F, H_{n,n})}{(2n)^k} = \frac{\frac{1}{2}\text{hom}(F, H_{n,n})}{n^k}$$

is the probability that a random map of $V(F)$ into $\{1, \dots, n\}$ is order-preserving. As $n \rightarrow \infty$, the fraction of non-injective maps tends to 0, and hence it is easy to see that $2^{k-1}t(F, H_{n,n})$ tends to a number $2^{k-1}t(F)$, which is the probability that a random ordering of $V(F)$ is compatible with P . In other words, $k!2^{k-1}t(F)$ is the number of linear extensions of P .

However, the half-graphs do not converge to any finite weighted graph. To see this, let S_k denote the star on k nodes, and consider the (infinite) matrix M defined $M_{k,l} = t(S_{k+l-1})$. If $t(F) = t(F, G_0)$ for some finite weighted graph G_0 , then it follows from the characterization of homomorphism functions in [28] that this matrix has rank at most $|V(G_0)|$; on the other hand, it is easy to compute that

$$M_{k,l} = \frac{1}{2^{k+l-2}(k+l-1)},$$

and this matrix (up to row and column scaling, the *Hilbert matrix*) has infinite rank (see e.g [18]).

One can say, however, that in the limit, we are considering order-preserving maps of the poset P into the interval $[0, 1]$; equivalently, adjacency-preserving maps of F into the infinite

graph with node set $[0, 1]$ and edge-set $\{xy : x \leq 1/2, y > 1/2, x \leq y - 1/2\}$. More precisely, the limit is given by the function

$$W(x, y) = \begin{cases} 1, & \text{if } x \geq y + \frac{1}{2} \text{ or } y \geq x + \frac{1}{2}, \\ 0, & \text{otherwise.} \end{cases}$$

Example 4.12 (Preferential attachment graphs) Preferential attachment as a paradigm of generating power law distributions goes back to Yule [61] in 1923. In the context of networks and graph theory, the idea is usually credited to Barabasi and Albert [7], even though it can be already found in [53]. For the simplest case of preferential attachment trees, the rigorous mathematical analysis of these models was first carried out in [50], while the rigorous analysis of the more general model of Barabasi and Albert was first carried out in [10]. For more general models of undirected preferential attachment graphs see [20], and for models with directed edges, see [12]. All these model are sparse models with bounded average degree.

Here we define a *preferential attachment graph* $\text{PAG}(n, m)$ as the random graph with n nodes and m edges obtained by the following procedure. Fix a set of n nodes, and let $v_1 \dots v_n$ be any ordering of the nodes. We extend this sequence one by one by picking an element of the current sequence randomly and uniformly, and append a copy of it at the end. We repeat this until $2m$ further elements have been added. So we get a sequence $v_1 \dots v_n v_{n+1} \dots v_{n+2m}$.

Now we construct G by connecting nodes v_{n+2k-1} and v_{n+2k} for $k = 1, 2, \dots, m$, to get $G(n, m)$. (Note that G may have multiple edges and loops, which we have to live with for the time being).

Another way of describing this construction is to view it as adding edges one by one, where the probability of adding an edge connecting u and v is proportional to the product of the “degrees”. To be more precise, the probability that the $(k + 1)$ -st edge connects u and v is

$$\begin{cases} \frac{2(d(u) + 1)(d(v) + 1)}{(n + 2k)(n + 2k + 1)} & \text{if } u \neq v, \\ \frac{(d(u) + 1)(d(u) + 2)}{(n + 2k)(n + 2k + 1)} & \text{if } u = v, \end{cases}$$

where $d(u)$ is the current degree of the node (adding 1 to the degree is needed to start the procedure at all; adding 2 to the second factor in the case when $u = v$ makes everything come out nicer).

It can be shown that the limit of preferential attachment graphs $\text{PAG}(n, cn^2)$, with probability 1, is the function $W_c(x, y) = c(\log x)(\log y)$. It is interesting to note that the graphs $\mathbf{G}(n, W_c)$ form another (different) sequence of random graphs tending to the same limit W_c with probability 1.

5 The metric space of graphs

5.1 Distances of graphs

In this section, we assume that the graph G (which we probe by F from the left and by H from the right) is dense, i.e., the number of edges of G is $\Omega(n^2)$, where n is the number of nodes (the results are valid but mostly vacuous for sparse graphs). We’ll discuss analogous questions for sparse graphs (specifically, for graphs with bounded degree) in Section 8.

5.1.1 Matrix norms

For an $n \times n$ matrix A , the *rectangle norm* (also called the *cut norm*) is defined as

$$\|A\|_{\square} = \max_{u, v \in \{0, 1\}^n} |u^T A v|. \quad (28)$$

This norm is closely related to $\ell_\infty \rightarrow \ell_1$ norm, which can be defined by

$$\|A\|_{\infty \rightarrow 1} = \max_{u,v \in [-1,1]^n} |u^\top Av| = \max_{u,v \in \{-1,1\}^n} u^\top Av; \quad (29)$$

in fact,

$$\frac{1}{4}\|A\|_{\infty \rightarrow 1} \leq \|A\|_{\square} \leq \|A\|_{\infty \rightarrow 1}. \quad (30)$$

For symmetric matrices A , the norm

$$\|A\|'_{\square} = \max_{u \in \{0,1\}^n} |u^\top Au| \quad (31)$$

is simpler to define and is also equivalent to the rectangle norm:

$$\|A\|'_{\square} \leq \|A\|_{\square} \leq 2\|A\|'_{\square}. \quad (32)$$

A constant factor approximation of the rectangle norm can be computed in polynomial time, using semidefinite optimization and Grothendieck's inequality in functional analysis (see [3]).

5.1.2 Labeled graphs on the same set of nodes

Let G and G' be two graphs on the same set of n nodes. We want to define a notion of distance between them that reflects structural similarity. A first attempt is to define

$$d_1(G, G') = \frac{1}{n^2} |E(G) \Delta E(G')|.$$

(Here the division by n^2 is just a convenience, so that the distance of two graphs is always between 0 and 1.) However, this notion is too restrictive: For example, the distance of two random graphs with the same density is of constant order (with large probability), even though two random graphs are structurally very similar.

For our purposes, the following distance function will be more useful. For a graph G and sets $S, T \subseteq V(G)$, let $e_G(S, T)$ denote the number of edges in G with one endnode in S and the other in T (the endnodes may also belong to $S \cap T$; so $e_G(S, S)$ is twice the number of edges spanned by S). We define:

$$d_{\square}(G, G') = \frac{1}{n^2} \max_{S, T \subseteq V(G)} |e_G(S, T) - e_{G'}(S, T)|.$$

Note that we are dividing by n^2 and not by $|S| \times |T|$, so the contribution of a pair S, T is at most $|T| \times |S|/n^2$. Thus small sets of size $o(n)$ play no role when measuring the distance. In terms of the adjacency matrices A and A' of G and G' , respectively, this can be expressed as

$$d_{\square}(G, G') = \frac{1}{n^2} \|A - A'\|_{\square}.$$

Note that the definition can be extended to the case when G and G' have edgeweights. Furthermore, by (30) and (32), we could replace the $\|\cdot\|_{\square}$ norm in the definition by one of the other matrix norms defined above without distorting the distance by more than a constant factor.

We need to extend this notion to weighted graphs on the same set of nodes. Let G and G' be weighted graphs with $V(G) = V(G')$. We assume that G and G' both have total nodeweight 1, but the weights of individual nodes in G and G' may be different. Then we define

$$\begin{aligned} d_{\square}(G, G') = & \sum_i |\alpha_i(G) - \alpha_i(G')| \\ & + \max_{S, T \subseteq V(G)} \left| \sum_{\substack{i \in S \\ j \in T}} (\alpha_i(G)\alpha_j(G)\beta_{ij}(G) - \alpha_i(G')\alpha_j(G')\beta_{ij}(G')) \right|. \end{aligned} \quad (33)$$

If the two graphs do not both have total nodeweights of one, then we simply define the distance d_{\square} in terms of the corresponding “normalized” graphs, i.e., we replace $\alpha_i(G)$ with $\alpha_i(G)/\alpha(G)$, and similarly for G' .

5.1.3 Unlabeled graphs with the same number of nodes

Now assume that G and G' are unlabeled unweighted graphs on n nodes. It is natural to define

$$\widehat{\delta}_{\square}(G, G') = \min_{\tilde{G}, \tilde{G}'} d_{\square}(\tilde{G}, \tilde{G}'), \quad (34)$$

where \tilde{G} and \tilde{G}' range over all labelings of G and G' by $1, \dots, n$, respectively (of course, we could fix the labeling of one of the graphs).

Consider any labeling that attains the minimum in the definition of $\widehat{\delta}_{\square}$, and identify the nodes of G and G' with the same label. In this case, we say that G and G' are *optimally overlaid*.

5.1.4 Unlabeled graphs with different number of nodes

To define the distance (in any of the above senses) of two unlabeled graphs with different number of nodes, say G with n nodes and G' with n' nodes, a first idea is to blow up each node of G into n' nodes, and each node of G' into n nodes, so that both graphs now will have nn' nodes. An improved version of this idea is to match up the nodes “fractionally”. This also allows us to extend the notion of distance to weighted graphs.

Let G and G' be weighted graphs with (say) $V(G) = [n]$, $V(G') = [n']$, and assume that the sum of nodeweights is 1 (just scale the nodeweights of each graph). Let X be a nonnegative $n \times n'$ matrix such that

$$\sum_{u=1}^{n'} X_{iu} = \alpha_i(G)$$

and

$$\sum_{i=1}^n X_{iu} = \alpha_u(G').$$

We think of X_{iu} as the portion of node i that is mapped onto node u . We call such a matrix X a *fractional overlay* of G and G' . Let $\mathcal{X}(G, G')$ denote the set of all fractional overlays. Note that for every $X \in \mathcal{X}(G, G')$,

$$\sum_{i=1}^n \sum_{u=1}^{n'} X_{iu} = \sum_{i=1}^n \alpha_i(G) = \sum_{u=1}^{n'} \alpha_u(G') = 1.$$

(If we view α_G and $\alpha_{G'}$ as probability distributions, then every $X \in \mathcal{X}(G, G')$ is a *coupling* of these distributions.)

For each fractional overlay, we construct the following two weighted graphs. The nodes of $G[X]$ are all pairs (i, u) where $1 \leq i \leq n$ and $1 \leq u \leq n'$. The weight of the node (i, u) is X_{iu} , and the weight of the edge $((i, u), (j, v))$ is β_{ij} . The other graph $G'[X^{\top}]$ is defined similarly, except that the roles of i and u are interchanged. Now the node sets of $G[X]$ and $G'[X^{\top}]$ are labeled by the same set of pairs (i, u) , so their distances are well defined.

Thus we can define the distance of two weighted unlabeled graphs G and G' (with total nodeweight 1):

$$\delta_{\square}(G, G') = \min_{X \in \mathcal{X}(G, G')} d_{\square}(G[X], G'[X^{\top}]).$$

We can express this distance in terms of the original graphs G and G' by the following formula:

$$\delta_{\square}(G, G') = \min_{X \in \mathcal{X}(G, G')} \max_{S, T \subseteq V \times V'} \left| \sum_{\substack{(i,u) \in S \\ (j,v) \in T}} X_{iu} X_{jv} (\beta_{ij}(G) - \beta_{uv}(G')) \right|. \quad (35)$$

Of course, this definition also applies if G and G' have the same number of nodes; however, it may give a different value than (34). It is proved in [16] that there is a constant $c > 0$ such that

$$\delta_{\square}(G, G') \leq \widehat{\delta}_{\square}(G, G') \leq c \delta_{\square}(G, G')^{1/4} \quad (36)$$

(the lower bound is trivial; we do not have an example showing that the exponent $1/4$ is needed in the upper bound). While the definition of $\widehat{\delta}_{\square}$ is more straightforward, the distance δ_{\square} will be easier to work with, and we will use mostly the latter distance.

A very special weighted graph is $K_1(p)$: a single node with a loop with weight p . For the random graph $\mathbf{G} = \mathbf{G}(n, p)$, we have a.s.

$$\delta_{\square}(\mathbf{G}, K_1(p)) \longrightarrow 0 \quad (n \rightarrow \infty).$$

Let (G_n) be a sequence of simple graphs. It follows by standard results on quasirandom graphs that

Proposition 5.1 [19] *A sequence (G_n) of graphs is quasirandom if and only if $\delta_{\square}(G_n, K_1(p)) \rightarrow 0$ as $n \rightarrow \infty$.*

The following result connects this distance to homomorphism functions.

Lemma 5.2 [44] *For any three simple graphs F , G and G'*

$$|t(F, G) - t(F, G')| \leq |E(F)| \cdot \delta_{\square}(G, G').$$

5.2 Szemerédi's Lemma

For a graph $G = (V, E)$ and two subsets $U, W \subset V$ we define the “irregularity” of the pair U, W as the quantity

$$\text{irreg}_G(U, W) = \max_{X \subseteq U, Y \subseteq W} |e_G(X, Y) - d|X| \cdot |Y||,$$

where d is the density $d = e_G(U, W)/(|U| \cdot |W|)$. Let $\text{twr}(\varepsilon)$ denote the $\lceil 1/\varepsilon^2 \rceil$ times iterated exponential function (the “tower”). With this notation, we can state one version of the Regularity Lemma:

Lemma 5.3 (Szemerédi Regularity Lemma) *For every $\varepsilon > 0$ and every graph $G = (V, E)$ there is a partition \mathcal{P} of V into $k \leq \text{twr}(\varepsilon)$ classes V_1, \dots, V_k such that*

$$\sum_{1 \leq i < j \leq k} \text{irreg}_G(V_i, V_j) \leq \varepsilon |V|^2.$$

(While this form is perhaps easiest to state and prove, there are equivalent forms that are more suited for applications. We refer to [35] for a survey, to [55] for connections with quasirandom graphs, and to [47] for analytic aspects of the Regularity Lemma.)

The appearance of the tower function in the Lemma forbids practical applications (and unfortunately this bound on the number of parts is not far from best possible, as it was shown by Gowers [32]). A more reasonable threshold was proved by Frieze and Kannan [29], at the cost of using a weaker measure of irregularity.

Given a graph $G = (V, E)$ and a partition $\mathcal{P} = \{V_1, \dots, V_k\}$ of V , we define a weighted graph $G_{\mathcal{P}}$ on V by setting $\alpha_u(G_{\mathcal{P}}) = 1$ and $\beta_{uv}(G_{\mathcal{P}}) = d_{i(u)i(v)}$, where $i(u)$ is the index of the set V_i containing u and $d_{ij} = e_G(V_i, V_j)/(|V_i||V_j|)$. The edgeweight matrix of $G_{\mathcal{P}}$ is thus obtained from the adjacency matrix of G by replacing each entry in the block $V_i \times V_j$ by the average over the block. In this notation, the result of Frieze and Kannan [29] can be formulated as follows:

Lemma 5.4 (Weak Regularity Lemma) *For every $\varepsilon > 0$ and every graph $G = (V, E)$, there exists a partition \mathcal{P} of V into $k \leq 2^{2/\varepsilon^2}$ classes such that $d_{\square}(G, G_{\mathcal{P}}) \leq \varepsilon$.*

The bound on the number of partition classes is still rather large (exponential), but at least not a tower. Frieze and Kannan show that the partition can be obtained as an “overlay” of only $1/\varepsilon^2$ sets, so it has a description that is polynomial in $1/\varepsilon$, which in some applications leads to polynomial time algorithms (see e.g. [2]).

The Weak Regularity Lemma immediately implies the following slight variant:

Lemma 5.5 [16] *For every $\varepsilon > 0$ and every graph $G = (V, E)$, there is a weighted graph H with at most $\lceil 2^{1/\varepsilon^2} \rceil$ nodes such that $\delta_{\square}(G, H) \leq \varepsilon$.*

We note that other versions strengthen the conclusion (of course, at the cost of replacing the tower function by an even more huge value). Such a “super-strong” Regularity Lemma was proved and used by Alon and Shapira [4]. It would be interesting to fit the original Regularity Lemma or one of its applications into this framework.

5.3 Sampling from a graph

The following important fact connecting sampling and graph distance follows from a result in [2]; see also [16] for a simple proof of (a):

Theorem 5.6 *Let G_1 and G_2 be two graphs on the same set of nodes V , let $\varepsilon = d_{\square}(G_1, G_2)$, $\delta > 0$, and let S be a random k -subset of V .*

(a) *If $k \geq \frac{300}{\varepsilon^2} \log(\frac{2}{\delta})$, then with probability at least $1 - \delta$,*

$$2^{-10}\varepsilon \leq d_{\square}(G_1[S], G_2[S]) \leq 4\varepsilon^{1/4}.$$

(b) *If $k \geq 10^{10} \log(2/\varepsilon)/(\varepsilon^4\delta^5)$, then with probability at least $1 - \delta$,*

$$2^{-10}\varepsilon \leq d_{\square}(G_1[S], G_2[S]) \leq 10^7 \frac{\varepsilon}{\sqrt{\delta}}.$$

Using this bound and the (weak) Regularity Lemma, it is not hard to prove the following theorem, which is the key to several further results.

Theorem 5.7 [16] *Let G be a (possibly weighted) graph, $\varepsilon > 0$ and $k \geq 2^{c_1/\varepsilon^8}$. Let S be a random subset of $V(G)$ of size k . Then with probability at least $1 - \varepsilon$, we have $\delta_{\square}(G, G[S]) < \varepsilon$.*

Informally, if we take a sample of k points, and blow up each node of this subgraph into $|V(G)|/k$ twins, then the resulting graph can be overlaid with G so that the d_{\square} -distance will be small.

This theorem can be thought of as a strengthening of the (weak) Regularity Lemma in two directions. First, it says that the approximating weighted graph can be required to be unweighted; second, that it can be obtained just by drawing a random sample.

From this theorem, it is easy to deduce the following converse of Lemma 5.2:

Theorem 5.8 [16] *Let G, G' be simple graphs, and let $\varepsilon > 0$. Set $k = \lceil 2^{c_2/\varepsilon^8} \rceil$, and assume that for every simple graph F on at most k nodes, we have $|t(F, G) - t(F, G')| < 2^{-2k^2}$. Then $\delta_{\square}(G, G') \leq \varepsilon$.*

These results allow us to characterize convergent graph sequences:

Theorem 5.9 [16] *A graph sequence is convergent if and only if it is Cauchy in the δ_{\square} metric.*

Let \mathcal{F} denote the metric space of all finite, simple graphs with the δ_{\square} metric. It follows from the above that the completion \mathcal{X} of \mathcal{F} can be described as follows. Consider the space of all functions in \mathcal{W}_0 , with the distance

$$d_{\square}(U, W) = \sup_{S, T \subseteq [0, 1]} \left| \int_{S \times T} W(x, y) dx dy \right|.$$

Define a new metric by

$$\delta_{\square}(U, W) = \inf d_{\square}(U^{\phi}, W^{\psi}),$$

where ϕ and ψ range over all measure preserving maps $\phi, \psi : [0, 1] \rightarrow [0, 1]$. Then the elements of \mathcal{X} can be obtained by identifying functions that are at distance 0, and the δ_{\square} metric between these classes extends the δ_{\square} metric on graphs.

The above results give various other descriptions of this completion: for example, \mathcal{X} is isomorphic to \mathcal{T}_0 with the metric

$$\delta(t_1, t_2) = \sup_F \frac{1}{|E(F)|} |t_1(F) - t_2(F)| \quad (t_1, t_2 \in \mathcal{T}_0).$$

Szemerédi's Lemma can be used to show that \mathcal{X} is compact.

5.4 Testing huge graphs

Imagine that we have a huge graph G ; this graph is so large that we cannot describe it completely in any way. All we can do is sample a bounded number of nodes of G and look at the subgraph that is induced by them. What can we learn about G ?

There are two related, but slightly different ways of asking this question.

5.4.1 Parameter testing

Parameter testing is easier to state. We may want to determine some parameter of G ; say what is the edge density? How large is the density of the maximum cut? Of course, we will not be able to determine the exact value of this parameter; the best we can hope for is that if we take a sufficiently large sample, we can find the approximate value of the parameter with large probability. To be precise, a graph parameter f is *testable*, if for every $\varepsilon > 0$ there is a positive integer k such that if G is a graph with at least k nodes and we select a set X of k independent uniform random nodes of G , then from the subgraph induced by them we can compute an estimate $\tilde{f}(G[X])$ of f such that

$$P(|f(G) - \tilde{f}(G[X])| > \varepsilon) < \varepsilon.$$

It is an easy observation that we can always use $\tilde{f}(G[X]) = f(G[X])$.

Using the notions of graph distance and convergence introduced above, we can give a number of characterizations of testable parameters.

Proposition 5.10 [16] *A simple graph parameter is testable if and only if any of the following equivalent conditions holds.*

(a) *For every convergent graph sequence (G_n) , the limit of $f(G_n)$ exists as $n \rightarrow \infty$ (continuity at infinity).*

(b) *For every $\varepsilon > 0$ there is an integer k_0 such that for every $k > k_0$ and every graph G on at least k nodes, a random set X of k nodes of G satisfies*

$$|f(G) - \mathbb{E}(f(G[X]))| < \varepsilon.$$

(c) *f is “essentially” uniformly continuous with respect to the δ_{\square} distance in the following sense: For every $\varepsilon > 0$ there is an $\varepsilon_0 > 0$ and a positive integer n_0 so that if G_1 and G_2 are two graphs with $|V(G_i)| \geq n_0$ and $\delta_{\square}(G_1, G_2) < \varepsilon_0$, then $|f(G_1) - f(G_2)| < \varepsilon$.*

(d) *There exists a functional $\hat{f}(W)$ on \mathcal{W}_0 that is continuous in the rectangle norm, and extends f in the sense that $|\hat{f}(W_G) - f(G)| \rightarrow 0$ if $|V(G)| \rightarrow \infty$.*

If we want to use (c) to prove that a certain invariant is testable, then the complicated definition of the δ_{\square} distance may cause a difficulty. So it is useful to show that (c) can be replaced by a weaker condition, which consists of three special cases of (c):

Supplement 5.11 [16] *The following three conditions together are also equivalent to testability:*

(c.1) *For every $\varepsilon > 0$ there is an $\varepsilon' > 0$ such that if G and G' are two simple graphs on the same node set and $d_{\square}(G, G') \leq \varepsilon'$ then $|f(G) - f(G')| < \varepsilon$.*

(c.2) *For every simple graph G , $f(G(m))$ has a limit as $m \rightarrow \infty$, where $G(m)$ denotes the graph obtained from G by replacing each node by m twins.*

(c.3) *$f(G(m)) - f(G) \rightarrow 0$ if $|V(G)| \rightarrow \infty$.*

Some of the implications between conditions (a)-(d) in the Theorem are easy, some others follow from the general theory sketched above. To illustrate the use of this theorem, let us consider the density of the maximum cut:

$$f(G) = \max_{S \subseteq V(G)} \frac{e_G(S, V(G) \setminus S)}{\binom{n}{2}}.$$

This parameter is testable: this fact is nontrivial, and its first proof by Goldreich, Goldwasser and Ron [31] was one of the first important results in Property Testing. Of the conditions above, (a) and (b) are more or less a reformulation of testability.

Condition (c), on the other hand, is easy to verify in this case. Let $\varepsilon > 0$, and let G_1 and G_2 be two graphs for which $\delta_{\square}(G_1, G_2) < \varepsilon$. Let us blow up the points of each graph so that the new graphs G'_1 and G'_2 have the same number N of points and they can be overlaid so that $d_{\square}(G'_1, G'_2) < \varepsilon$. For any subset $S \subseteq V(G'_1) = V(G'_2)$, we have

$$|e_{G'_1}(S, V(G) \setminus S) - e_{G'_2}(S, V(G) \setminus S)| < \varepsilon N^2,$$

and hence

$$|f(G'_1) - f(G'_2)| < \varepsilon.$$

To complete the proof, one must argue that $|f(G'_i) - f(G_i)|$ is small, which is not hard (and is not given here).

Condition (d) can also be directly verified: we can extend the definition of a maximum cut to functions $W \in \mathcal{W}_0$ in a natural way:

$$\hat{f}(W) = \sum_{S \subseteq [0,1]} \int_S \int_{[0,1] \setminus S} W(x, y) dx dy.$$

Then it is easy to check that this functional is continuous in the norm $\|\cdot\|_{\square}$, and extends f .

5.4.2 Property testing

Instead of estimating a numerical parameter, we may want to determine some property of G : Is G 3-colorable? Is it connected? Does it have a triangle? The answer will of course have some uncertainty. A precise definition was given by Goldreich, Goldwasser and Ron [31], who also proved several fundamental results about this problem. There are in fact several ways to formalize this question. For this exposition, we take the following.

As for parameter testing, we specify an $\varepsilon > 0$ and want to find a positive integer k (depending on ε) with the following property. We select k independent uniform random nodes of G , and from the subgraph induced by them we compute a guess $X \in \{YES, NO\}$. Ideally, we want that if the graph does have the property, our guess should be YES with large probability, and if the graph does not have the property, then we should guess NO with large probability. But this is too much to ask. Suppose that we have two graphs that can be obtained from each other by changing a very tiny fraction of the edges, but one has the property, and the other does not. Then a sample induced subgraph from one graph will have almost the same distribution as a sample (of the same size) from the other, and so our guess for the two graphs will be almost the same.

A graph property \mathcal{P} is *testable*, if our guess satisfies the following: if a graph has the property in a robust way so that changing at most εn^2 edges in any way it still has the property, we must guess YES with probability at least $1 - \varepsilon$; similarly, if changing at most εn^2 edges in any way the obtained graph does not have the property, then we must guess NO with probability at least $1 - \varepsilon$; in the grey area inbetween, we can guess arbitrarily. In other words, whatever we guess, we should be able to change at most εn^2 edges to make our guess right.

Remark 5.12 While changing a small number of edges is the most natural way to formalize that there is a “nearby” graph with the property, we have seen that the rectangular distance is often better behaved. One is tempted to define that a property is *weakly testable*, if for every $\varepsilon > 0$ there is a k such that for every graph G on at least k nodes we can make a guess based on a sample induced subgraph of size k such that with probability at least $1 - \varepsilon$, there is a graph G' such that $d_{\square}(G, G') < \varepsilon$ and our guess is right for G' . But this notion is not very interesting due to the fact that *every graph property is weakly testable*. This is an easy application of Theorem 5.7 above.

Among the many results on graph property testing, let us quote a surprisingly general recent result of Alon and Shapira [4]. A graph property is called *hereditary*, if it is inherited by induced subgraphs.

Theorem 5.13 [4] *Every hereditary graph property is testable.*

They in fact obtain a stronger result, which can be cast in the framework of parameter testing. For a hereditary graph property \mathcal{P} , define the *distance from the property* as $d(G, \mathcal{P}) = D(G, \mathcal{P})/|V(G)|^2$, where $D(G, \mathcal{P})$ is the minimum number of edges we need to change in G to obtain a graph with property \mathcal{P} . Alon and Shapira proved:

Theorem 5.14 [4] *The distance from a hereditary graph property is testable.*

This theorem has a reasonably short proof using graph limits, see [47].

6 Partitions and homomorphisms into small graphs

6.1 Ground state energy

In Section 4 we defined convergence of a graph sequence G_n in terms of the homomorphism numbers from small graphs F to G_n (“convergence from the left”). But there are many applications

where one wants to study homomorphisms from G_n into a small graph H . This naturally raises the question whether suitably normalized homomorphism numbers $\text{hom}(G_n, H)$ converge if G_n is convergent from the left.

Consider a graph G on n nodes, and a softcore graph H on q nodes. (Recall that H is called softcore if all edge weights are strictly positive.) Then $\log \text{hom}(G, H)$ typically grows like the number of edges in G . For dense graphs, it therefore seems natural to consider the quantity

$$\frac{1}{n^2} \log \text{hom}(G, H). \quad (37)$$

This quantity is closely related to a weighted maximum cut problem on G . Indeed, let $B = (B_{ij})_{1 \leq i, j \leq q}$ be a symmetric matrix with real entries. We then define the *ground state energy* of the “model” B on the graph G as

$$\mathcal{E}(G, B) = \frac{1}{n^2} \max_{\phi: V(G) \rightarrow [q]} \sum_{uv \in E(G)} B_{\phi(u)\phi(v)}. \quad (38)$$

For large n , the quantity defined in (37) is well approximated by the ground state energy.

Lemma 6.1 *Let G be an unweighted graph with n nodes, and let H be a softcore graph with $\alpha(H) = 1$. Let $\alpha = \min_i \alpha_H(i)$. Let B be the matrix of logarithms of the edgeweights of H . Then*

$$\mathcal{E}(G, B) - \frac{\log(1/\alpha)}{n} \leq \frac{\log \text{hom}(G, H)}{n^2} \leq \mathcal{E}(G, B).$$

(Note that the upper bound on $\log \text{hom}(G, H)/n^2$ does not depend on the nodeweights of H , and in the lower bound, only the error term does.)

Proof. Note that

$$\max_{\phi: V(G) \rightarrow V(H)} \text{hom}_{\phi}(G, H) = e^{n^2 \mathcal{E}(G, B)}.$$

Thus we have

$$\text{hom}(G, H) = \sum_{\phi} \alpha_{\phi} \text{hom}_{\phi}(G, H) \leq \sum_{\phi} \alpha_{\phi} e^{n^2 \mathcal{E}(G, B)} = e^{n^2 \mathcal{E}(G, B)},$$

and

$$\text{hom}(G, H) \geq \max_{\phi} \alpha_{\phi} \text{hom}_{\phi}(G, H) \geq \alpha^n e^{n^2 \mathcal{E}(G, B)}.$$

From these bounds the Lemma follows. \square

Example 6.2 As a special case, consider the graph H consisting of two nodes of weight $1/2$, with loops of weight 1 at each, and connected by an edge of weight e (the base of the natural logarithm). Then $n^2 \mathcal{E}(G, H')$ is the size of the maximum cut in G , which we denote by $\text{MAXCUT}(G)$. By Lemma 6.1,

$$\log \text{hom}(G, H) - n \leq \text{MAXCUT}(G) \leq \log \text{hom}(G, H).$$

Since $\text{MAXCUT}(G) \geq |E(G)|/2$, this gives a very good approximation of the maximum cut. More generally, for a fixed H , computing $\mathcal{E}(G, H)$ is a weighted multiway cut problem.

Our next theorem states that convergence from the left implies convergence of the ground state energies. Its proof uses the notion of fractional partitions, a notion we will need at several places in this section: a *fractional partition* of a set V into q classes (briefly, a fractional q -partition) is a q -tuple $\rho = (\rho_1, \dots, \rho_q)$ of functions from V to $[0, 1]$ such that for all $x \in V$, we have $\rho_1(x) + \dots + \rho_q(x) = 1$. (Later in this section, we will apply this definition also to the case when $V = [0, 1]$, when we tacitly assume that the functions ρ_i are measurable.) We will use the notation Pd_q for the set of probability distributions on $[q]$, i.e., the set of vectors $a = (a_1, \dots, a_q)$ such that $a_i \geq 0$ and $\sum_i a_i = 1$, and the notation Sym_q for the set of $q \times q$ symmetric matrices.

Theorem 6.3 [16] *Let q be a positive integer, let $B \in \text{Sym}_q$, and let G_n be a convergent sequence of simple graphs. Then $\mathcal{E}(G_n, B)$ is a convergent sequence.*

As a simple illustration of the usefulness of Proposition 5.10 and its Supplement 5.11, we sketch the proof. Indeed, let us consider the quantity

$$\mathcal{E}_\phi(G, B) = \frac{1}{n^2} \sum_{uv \in E(G)} B_{\phi(u)\phi(v)}$$

where ϕ is a map from $V(G)$ to $[q]$. Identifying these maps with partitions $\mathcal{P} = (V_1, \dots, V_q)$ of $V(G)$ and using the definition of the d_\square metric, we immediately see that

$$|\mathcal{E}_\phi(G, B) - \mathcal{E}_\phi(G', B)| \leq q^2 \|B\|_\infty d_\square(G, G')$$

whenever G and G' are simple graphs on the same set of nodes, which verifies the condition (c.1) of Supplement 5.11.

The condition (c.2) is not hard to verify either: Define the energy of a fractional q -partition ρ of $V(G)$ as

$$\mathcal{E}_\rho(G, B) = \frac{1}{n^2} \sum_{uv \in E(G)} \sum_{i, j \in [q]} \rho_i(u) \rho_j(v) B_{ij}. \quad (39)$$

Then $\mathcal{E}(G(k), B) = \max_\rho \mathcal{E}_\rho(G, B)$ where the maximum runs over all fractional partitions such that all $\rho_i(u)$'s are multiples of $1/k$. We claim that the maximum is attained for fractional partitions which are $\{0, 1\}$ valued, so that $\mathcal{E}(G(k), B) = \mathcal{E}(G, B)$ for all k . It is clear from the above description that $\mathcal{E}(G(k), B) \geq \mathcal{E}(G, B)$, so the only thing we need to show is a matching upper bound on $\mathcal{E}(G(k), B)$. Consider a fractional partition ρ maximizing $\mathcal{E}_\rho(G, B)$, and a fixed node $u \in V(G)$. Then $\mathcal{E}_\rho(G, B)$ is a linear function of the vector $(\rho_1(u), \dots, \rho_q(u))$, implying that the maximum over all these vectors is obtained at a vertex of the simplex Pd_q . Applying this procedure to all nodes $u \in V(G)$, this gives the desired inequality $\mathcal{E}(G(k), B) \leq \mathcal{E}(G, B)$, and hence the equality of $\mathcal{E}(G(k), B)$ and $\mathcal{E}(G, B)$ for all k . The condition (c.2) therefore holds trivially. The verification of condition (c.3) is even easier.

Let us finally note that the limiting ground state energy of a convergent sequence can be expressed explicitly. Indeed, for $W \in \mathcal{W}_0$ and a symmetric $q \times q$ matrix B , let

$$\mathcal{E}(W, B) = \max_\rho \frac{1}{2} \sum_{i, j=1}^q B_{ij} \int_{[0,1]^2} \rho_i(x) \rho_j(y) W(x, y) dx dy, \quad (40)$$

where the maximum runs over fractional q -partitions of $[0, 1]$. Then we have the following theorem.

Theorem 6.4 [16] *Let (G_n) be a convergent sequence, and let $W \in \mathcal{W}_0$ be its limit. Let H be a softcore graph, and let B be the matrix of logarithms of the edgeweights of H . Then*

$$\lim_{n \rightarrow \infty} \frac{\log \text{hom}(G_n, H)}{|V(G_n)|^2} = \lim_{n \rightarrow \infty} \mathcal{E}(G_n, B) = \mathcal{E}(W, B).$$

6.2 Entropy and free energy

In addition to the ground state energy and the closely related quantity (37), we will consider the so-called ‘‘pressure’’ or ‘‘free energy’’ of the model H on G , defined as

$$\widehat{\mathcal{P}}(G, H) = \frac{1}{n} \log \text{hom}\left(\frac{1}{n}G, H\right) \quad (41)$$

where n is the number of nodes in G and $\frac{1}{n}G$ is obtained from G by multiplying the edge weights of G by $\frac{1}{n}$.

To discuss the convergence of the free energy, we need the notion of the entropy of fractional partitions. Let $\rho = (\rho_1, \dots, \rho_q)$ be a fractional partition of a finite set V . Then the entropy of ρ is defined as

$$\mathcal{H}(\rho) = -\frac{1}{|V|} \sum_{u \in V} \sum_{i=1}^q \rho_i(u) \log \rho_i(u).$$

Theorem 6.5 [16] *Let q be a positive integer, let H be a softcore weighted graph, and let (G_n) be a convergent sequence of simple graphs. Then $\widehat{\mathcal{P}}(G_n, H)$ is a convergent sequence.*

If G_n converges to a function $W \in \mathcal{W}_0$, the limiting free energy can again be expressed as an explicit function of W . To this end, let us define the entropy of a fractional q -partition ρ of $[0, 1]$ as

$$\mathcal{H}(\rho) = -\int_0^1 \sum_{i=1}^q \rho_i(x) \log \rho_i(x) dx,$$

and the entropy of the H colorings of the limit graph W as

$$\mathcal{P}(W, H) = \max_{\rho} \left(\mathcal{H}(\rho) + \sum_i h_i \int_{[0,1]} \rho_i(x) dx + \frac{1}{2} \sum_{i,j} B_{ij} \int_{[0,1]^2} W(x,y) \rho_i(x) \rho_j(y) dx dy \right), \quad (42)$$

where the maximum goes over all fractional q -partitions of $[0, 1]$, $h_i = \log \alpha_i(H)$ and $B_{ij} = \log \beta_{ij}(H)$.

Theorem 6.6 [16] *Let (G_n) be a convergent sequence of simple graphs, and let $W \in \mathcal{W}_0$ be its limit. Let H be a weighted softcore graph. Then*

$$\lim_{n \rightarrow \infty} \frac{1}{n} \log \text{hom}\left(\frac{1}{n}G_n, H\right) = \mathcal{P}(W, H).$$

This theorem can be proved using Proposition 5.10 and its Supplement 5.11, in a way similar to (but more involved than) the proof of Theorem 6.3.

6.3 Factor graphs

Let G be a weighted graph and let $\mathcal{P} = (V_1, \dots, V_q)$ be a partition of $V(G)$. The *factor graph* (or briefly *factor*) $H(G, \mathcal{P})$ is the weighted graph on $[q]$ with nodeweights

$$\alpha_i(H(G, \mathcal{P})) = \frac{\alpha(G[V_i])}{\alpha(G)} = \frac{\sum_{u \in V_i} \alpha_u(G)}{\alpha(G)},$$

and edges weights

$$\beta_{ij}(H(G, \mathcal{P})) = \frac{\sum_{u \in V_i, v \in V_j} \alpha_u(G) \alpha_v(G) \beta_{uv}(G)}{\alpha(G[V_i]) \alpha(G[V_j])}.$$

Note that $H(G, \mathcal{P})$ is invariant under scaling the nodeweights of G .

In the special case when G is unweighted, this definition specializes to

$$\alpha_i(H(G, \mathcal{P})) = \frac{|V_i|}{|V(G)|},$$

and

$$\beta_{ij}(H(G, \mathcal{P})) = \frac{e_G(V_i, V_j)}{|V_i| \cdot |V_j|}.$$

Here $e_G(V_i, V_j)$ denotes the number of edges $uv \in E(G)$ with $u \in V_i$ and $v \in V_j$; note that we allow $i = j$, in this case $e(V_i, V_i)$ is *twice* the number of edges spanned by V_i .

We denote by $\widehat{\mathcal{S}}_q(G)$ the set of factors of G with q nodes. Note that by our definition the factors are labeled graphs, but since permuting the nodes of a factor also gives a factor, we would not lose information by forgetting the labeling. We can consider $\widehat{\mathcal{S}}_q(G)$ as a subset of $\mathbb{R}^{q \times (q+1)}$.

We extend these definitions to functions $U \in \mathcal{W}_0$. Let \mathcal{P} be a q -partition of $[0, 1]$; we then define a weighted graph $H(U, \mathcal{P})$ on $V(H(U, \mathcal{P})) = [q]$, where node i has weight

$$\alpha_i(H(U, \mathcal{P})) = \lambda(V_i),$$

and edge ij has weight

$$\beta_{ij}(H(U, \mathcal{P})) = \frac{1}{\lambda(V_i)\lambda(V_j)} \int_{V_i \times V_j} U(x, y) dx dy.$$

(If $\lambda(V_i)\lambda(V_j) = 0$, then we define $\beta_{ij}(H(U, \mathcal{P})) = 0$.) We call the graph $H(U, \mathcal{P})$ a *factor* of U , and use the symbol $\widehat{\mathcal{S}}_q(U)$ to denote the set of all factors of U with q nodes.

Note that the knowledge of the factors of G is enough to recover the ground state energies of G . Indeed, in terms of the factors of G , the ground state energy defined in (38) can be expressed as

$$\mathcal{E}(G, B) = \max_{(a, X) \in \widehat{\mathcal{S}}_q(G)} \sum_{i,j=1}^q a_i a_j B_{ij} X_{ij}, \quad (43)$$

where $a = (a_1, \dots, a_q)$ and $X = (X_{ij})_{1 \leq i, j \leq q}$.

6.4 Fractional factor graphs

The set $\widehat{\mathcal{S}}_q(G)$ is typically a very large finite set, which makes it difficult to work with. It will be convenient to introduce a fractional version of factors.

Let G be a weighted graph. For every fractional partition $\rho = (\rho_1, \dots, \rho_q)$ of $V(G)$, we define the *fractional factor* G_ρ as the graph with nodeweights

$$\alpha_i(G_\rho) = \frac{\sum_{u \in V(G)} \rho_i(u) \alpha_u(G)}{\alpha(G)},$$

and edgeweights

$$\beta_{ij}(G_\rho) = \frac{\sum_{u, v \in V} \rho_i(u) \rho_j(v) \alpha_u(G) \alpha_v(G) \beta_{uv}(G)}{\alpha_i(G_\rho) \alpha_j(G_\rho)}.$$

Let $\mathcal{S}_q(G)$ denote the set of all fractional factors of G with q nodes. Then $\mathcal{S}_q(G)$ is a closed set; it is not convex in general.

We also extend these notions to functions. Let $U \in \mathcal{W}_0$ and let $\rho = (\rho_1, \dots, \rho_q)$ be a fractional partition of $[0, 1]$. Set $\alpha_i(\rho) = \int_0^1 \rho_i(x) dx$. Then we define

$$\alpha_i(U_\rho) = \alpha_i(\rho),$$

and

$$\beta_{ij}(U_\rho) = \frac{1}{\alpha_i(\rho) \alpha_j(\rho)} \int_{[0,1]^2} \rho_i(x) \rho_j(y) U(x, y) dx dy.$$

Let $\mathcal{S}_q(U)$ denote the set of all fractional factors of U with q nodes. For a weighted graph G and a vector $a \in \text{Pd}_q$, let $\widehat{\mathcal{B}}_a(G)$ denote the set of all weighted adjacency matrices of all factors of G with nodeweights a_1, \dots, a_q . So $\mathcal{S}_q(G)$ is the set of all pairs (a, B) with $B \in \widehat{\mathcal{B}}_a(G)$. The sets $\mathcal{B}_a(G)$, $\widehat{\mathcal{B}}_a(U)$ and $\mathcal{B}_a(U)$ are defined analogously.

We clearly have that $0 \leq \beta_{ij}(U_\rho) \leq 1$ whenever $U \in \mathcal{W}_0$. Furthermore, it is not hard to see that for $U \in \mathcal{W}_0$, the set $\mathcal{S}_q(U)$ is closed in the obvious topology of weighted labeled graphs on q nodes.

Lemma 6.7 [16] *For every $U \in \mathcal{W}_0$, the set $\mathcal{S}_q(U)$ is the closure of $\widehat{\mathcal{S}}_q(U)$.*

(The two sets are not equal in general.)

For every weighted graph G , $\mathcal{S}_q(G)$ is again a closed connected set. Obviously, $\mathcal{S}_q(G)$ contains $\widehat{\mathcal{S}}_q(G)$, but it is not its closure in general (since the latter is a finite set). It is not hard to see that for every weighted graph G ,

$$\mathcal{S}_q(G) = \mathcal{S}_q(W_G) = \widehat{\mathcal{S}}_q(W_G). \quad (44)$$

Clearly $\widehat{\mathcal{S}}_q(G)$ is a finite subset of these infinite sets. But it can be shown that it is not much smaller:

Lemma 6.8 [16] *If c is the largest nodeweight in G , then for every $H \in \mathcal{S}_q(G)$ there is an $H' \in \widehat{\mathcal{S}}_q(G)$ such that $\delta_{\square}(H, H') \leq 4q^2\sqrt{c}$.*

Most of the time, we will work with the fractional versions, which are much easier to handle.

6.5 Microcanonical ground state energy, a.k.a. multiway cut problems

We have seen that convergence of the sequence (G_n) implies convergence of the free energies and ground state energies. But the converse does not hold. To get convergence of (G_n) from the convergence of the ground state energies, we will need a finer measure than ground state energy, where the sizes of the partition classes are also taken into account (in physics terms, this could be called the “microcanonical version” of the ground state energy). We will see that this quantity also contains a number of frequently studied graph parameters.

For a finite set S , $q \geq 1$ and $a \in \text{Pd}_q$, let a^S denote the set of all maps $\phi : S \rightarrow [q]$ such that

$$|\phi^{-1}(i) - a_i|S| \leq 1 \quad (45)$$

for every $i \in [q]$. (In other words, we prescribe the proportions of elements of S mapped onto each $i \in [q]$, as closely as possible.) For a simple graph G and a softcore graph H with node weights one, we then introduce microcanonical homomorphism numbers

$$\text{hom}_a(G, H) = \sum_{\phi \in a^{V(G)}} \text{hom}_{\phi}(G, H).$$

The *microcanonical ground state energies* and *free energies* of the model B on a graph G with n nodes are then defined as

$$\widehat{\mathcal{E}}_a(G, B) = \frac{1}{n^2} \max_{\phi \in a^{V(G)}} \sum_{ij \in E(G)} B_{\phi(i), \phi(j)}, \quad (46)$$

and

$$\widehat{\mathcal{P}}_a(G, B) = \frac{1}{n} \log \text{hom}_a\left(\frac{1}{n}G, H\right) \quad (47)$$

respectively.

The ground state energy $\widehat{\mathcal{E}}_a$ contains a number of important graph parameters as special cases. If

$$q = 2, \quad a = \begin{pmatrix} \alpha \\ 1 - \alpha \end{pmatrix} \quad \text{and} \quad B = \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix} \quad (48)$$

then $\mathcal{E}_a(G, B)$ corresponds to the *densest subgraph* on $\alpha|V(G)|$ nodes. If

$$q = 2, \quad a = \begin{pmatrix} 1/2 \\ 1/2 \end{pmatrix} \quad \text{and} \quad B = \begin{pmatrix} 0 & -1 \\ -1 & 0 \end{pmatrix} \quad (49)$$

then $-\mathcal{E}_a(G, B)$ is the *minimal bisection*; if we replace here B by $-B$, then $\mathcal{E}_a(G, B)$ is the *maximal bisection*. For $q > 2$, we get *multiway cut problems* in a similar way.

The free energy $\hat{\mathcal{P}}_a$ is a finer measure. For example, in the case (48), $\mathcal{P}_a(G, B)$ will pick out the density of an induced subgraph that is not necessarily the maximum, but for which the number of induced subgraphs with this density is large, at the cost of some loss in density.

The use of (45) is cumbersome and in some cases it leads to unpleasant discontinuities; it will be much more convenient to work with a fractional version. We formulate our definition for a weighted graph G . Then for all $a \in \text{Pd}_q$ and $B \in \text{Sym}_q$ we define

$$\mathcal{E}_a(G, B) = \max_{\rho} \frac{1}{2} \sum_{u, v \in V(G)} \alpha_u(G) \alpha_v(G) \beta_{uv}(G) \sum_{i, j \in [q]} \rho_i(u) \rho_j(v) B_{ij}, \quad (50)$$

where ρ ranges over all fractional q -partitions of $V(G)$ such that

$$\sum_{u \in V(G)} \alpha_u(G) \rho_i(u) = a_i.$$

We also extend the notion of microcanonical ground state energy to functions. For every $W \in \mathcal{W}_0$, $a \in \text{Pd}_q$, and $B \in \text{Sym}_q$, we define

$$\mathcal{E}_a(W, B) = \max_{\rho: \alpha_{\rho} = a} \frac{1}{2} \sum_{i, j=1}^q B_{ij} \int_{[0,1]^2} \rho_i(x) \rho_j(y) W(x, y) dx dy, \quad (51)$$

where ρ ranges over fractional q -partitions of $[0, 1]$ with $\alpha_i(\rho) = a_i$ for all i . It is easy to see that

$$\mathcal{E}_a(G, B) = \mathcal{E}_a(W_G, B).$$

Using a result from [2], it follows easily that for every simple graph G on n nodes and every matrix B ,

$$|\hat{\mathcal{E}}_a(G, B) - \mathcal{E}_a(G, B)| \leq \frac{6q^3}{n} \|B\|_{\infty}. \quad (52)$$

The notion of the microcanonical free energy for functions is defined analogously: for every $W \in \mathcal{W}_0$, $a \in \text{Pd}_q$, and $B \in \text{Sym}_q$, we define

$$\mathcal{P}_a(W, B) = \max_{\rho: \alpha(\rho) = a} \left(\mathcal{H}(\rho) + \frac{1}{2} \sum_{i, j} B_{ij} \int_{[0,1]^2} W(x, y) \rho_i(x) \rho_j(y) dx dy \right), \quad (53)$$

where ρ ranges over fractional q -partitions of $[0, 1]$ with $\alpha_i(\rho) = a_i$ for all i .

6.6 Convergence from the right

Our goal is to give a characterization of convergent graph sequences in terms of partition information, specifically the sets $\mathcal{S}_q(G)$ and the functions $\mathcal{E}_a(G, B)$. Recall that for two sets A, B in a metric space (X, d) , the *Hausdorff distance* is defined as

$$d^{\text{Hf}}(A, B) = \max(\sup_{x \in A} \inf_{y \in B} d(x, y), \sup_{x \in B} \inf_{y \in A} d(x, y)).$$

It turns out that the following three types of information about two graphs G, G' are equivalent: (1) G and G' are close in the δ_{\square} distance; (2) $\mathcal{S}_q(G)$ and $\mathcal{S}_q(G')$ are close in the d^{Hf} distance for every q up to a certain bound, and (3) $|\mathcal{E}_a(G, B) - \mathcal{E}_a(G', B)|$ is small for all $a \in \text{Pd}_q$ and $B \in \text{Sym}_q$ for every q up to a certain bound.

The exact statement is the following:

Theorem 6.9 [16] (a) For two simple graphs G, G' and $a \in \text{Pd}_q$,

$$d_{\square}^{\text{Hf}}(\mathcal{S}_q(G), \mathcal{S}_q(G')) \leq \delta_{\square}(G, G')$$

(b) For two simple graphs G, G' , $a \in \text{Pd}_q$ and $B \in \text{Sym}_q$,

$$|\mathcal{E}_a(G, B) - \mathcal{E}_a(G', B)| \leq q^2 \delta_{\square}^{\text{Hf}}(\mathcal{S}_q(G), \mathcal{S}_q(G')).$$

(c) Let G, G' be two simple graphs, and suppose that

$$|\mathcal{E}_a(G, B) - \mathcal{E}_a(G', B)| \leq c_3 \frac{\varepsilon^2}{q^2}$$

for all $q \leq 4 \cdot 2^{c_4/\varepsilon^2}$, $a \in \text{Pd}_q$ and $B \in \text{Sym}_q$. Then $\delta_{\square}(G, G') < \varepsilon$.

Similar statements hold for the microcanonical free energy, see [16].

6.7 Summary: convergence criteria

The following theorem summarizes several equivalent conditions for a sequence of (dense) graphs to be convergent:

Theorem 6.10 [16] Let (G_n) be a sequence of simple graphs with $|V(G_n)| \rightarrow \infty$. Then the following are equivalent:

- (a) For every simple graph F , $t(F, G_n)$ is convergent.
- (b) The sequence (G_n) is Cauchy in the δ_{\square} metric.
- (c) For every $q \geq 1$, the sequence $\mathcal{S}_q(G_n)$ is Cauchy with respect to the Hausdorff metric d_1^{Hf} .
- (d) For every $q \geq 1$, $a \in \text{Pd}_q$ and $B \in \text{Sym}_q$, the sequence $\mathcal{E}_a(G_n, B)$ is a Cauchy sequence.
- (e) For every $q \geq 1$, $a \in \text{Pd}_q$ and $B \in \text{Sym}_q$, the sequence $\widehat{\mathcal{P}}_a(G_n, B)$ is a Cauchy sequence.

We can also characterize convergent graph sequences in terms of Szemerédi partitions.

Supplement 6.11 [16] The following two conditions are also equivalent to conditions (a)–(e) in Theorem 6.10:

(f) For every $k \geq 1$ there is an $n_k \geq 1$ such that if $n, m > n_k$, then G_n and G_m have weak Szemerédi k -partitions \mathcal{P} and \mathcal{P}' with error $2/\sqrt{\log k}$ such that $d_{\square}(H(G_n, \mathcal{P}), H(G_m, \mathcal{P}')) < 2/\sqrt{\log k}$.

(g) For every $k \geq 1$ there is an $n_k \geq 1$ such that if $n, m > n_k$, then G_n and G_m have strong Szemerédi k -partitions \mathcal{P} and \mathcal{P}' with error $1/\log^* k$ such that $d_{\square}(H(G_n, \mathcal{P}), H(G_m, \mathcal{P}')) < 1/\log^* k$.

A convergent sequence has a limit $W \in \mathcal{W}_0$ by Theorem 4.6. The conditions in Theorem 6.10 can be rephrased to characterize the convergence to this limit:

Theorem 6.12 [16] For a sequence (G_n) of simple graphs with $|V(G_n)| \rightarrow \infty$, and for any $W \in \mathcal{W}_0$, the following are equivalent:

- (a) For every simple graph F , $t(F, G_n) \rightarrow t(F, W)$.
- (b) $\delta_{\square}(G_n, W) \rightarrow 0$.
- (c) For every $q \geq 1$, $\mathcal{S}_q(G_n) \rightarrow \mathcal{S}_q(W)$ in the Hausdorff metric d_1^{Hf} .
- (d) For every $q \geq 1$, $a \in \text{Pd}_q$ and $B \in \text{Sym}_q$, $\mathcal{E}_a(G_n, B) \rightarrow \mathcal{E}_a(W, B)$.
- (e) For every $q \geq 1$, $a \in \text{Pd}_q$ and $B \in \text{Sym}_q$, $\widehat{\mathcal{P}}_a(G_n, B) \rightarrow \widehat{\mathcal{P}}_a(W, B)$.

One can define Szemerédi partitions for the limit objects $W \in \mathcal{W}_0$ (see [45]), and then formulate analogues of (f) and (g) in Supplement 6.11 describing the convergence to the limit.

7 Homomorphisms and extremal graph theory

7.1 Inequalities between homomorphism numbers

We have mentioned in the introduction that many results in extremal graph theory can be expressed as algebraic inequalities between homomorphism densities. Every algebraic inequality that holds for all finite graphs also holds for simple graph parameters in the closure \mathcal{T}_0 (and of course vice versa). So for example, Goodman's Theorem (1) is equivalent to saying that for every simple graph parameter $t \in \mathcal{T}_0$,

$$t(K_3) \geq t(K_2)(2t(K_2) - 1). \quad (54)$$

By multiplicativity, every algebraic inequality is equivalent to a linear inequality. For example, (1) is equivalent to

$$t(K_3) \geq 2t(K_2K_2) - t(K_2).$$

The positive semidefiniteness of the connection matrix (reflection positivity) implies linear and nonlinear inequalities between the values of a simple graph parameters $t \in \mathcal{T}_0$ for “small” graphs. In fact, Theorem 4.5 says that every (say, algebraic) inequality between the values of simple graph parameters $t \in \mathcal{T}_0$ is a consequence of multiplicativity and reflection positivity.

Let us describe some of these derivations. Fix a simple graph parameter $t \in \mathcal{T}_0$. If $v = (v_G)$ is any real vector with finite support, indexed by k -labeled graphs, then

$$v^\top M(t, k)v \geq 0 \quad (55)$$

gives a linear inequality between the values $t(F)$.

Probably all linear inequalities can be obtained by taking nonnegative linear combinations of inequalities (55). However, the (infinitely many) inequalities (55) and their consequences are not so easy to understand, and we formulate some special inequalities.

Trivial linear inequalities are that if F_1 is a subgraph of F_2 (not necessarily on the same set of nodes), then

$$t(F_1) \geq t(F_2). \quad (56)$$

For any two graphs $F_1 \subseteq F_2$ on the same set of nodes, the expression

$$\sum_{F_1 \subseteq F \subseteq F_2} (-1)^{|E(F_2) \setminus E(F)|} \text{hom}(F, G)$$

counts (by inclusion-exclusion) the number of homomorphisms of F_1 into G that map the edges in $E(F_2) \setminus E(F_1)$ onto non-adjacent pairs. This number is nonnegative, which implies that

$$\sum_{F_1 \subseteq F \subseteq F_2} (-1)^{|E(F_2) \setminus E(F)|} t(F, G) \geq 0. \quad (57)$$

If we fix $V(F) = V$ with $|V| = k$, then $t(F)$ can be considered as a setfunction on $\binom{k}{2}$ elements. Then (57) can be used to show that this setfunction is supermodular, i.e., it satisfies

$$t(F_1 \cup F_2) + t(F_1 \cap F_2) \geq t(F_1) + t(F_2). \quad (58)$$

We leave it to the reader as an exercise to derive these inequalities from (55).

Semidefiniteness of the connection matrix can also be formulated in terms of nonlinear inequalities (nonnegativity of certain determinants). One of these is worth mentioning. Let G be a k -labeled graph, then for $t \in \mathcal{T}_0$, we have

$$t(GG) \geq t(G)^2.$$

As special cases, we mention that for the path P_3 on 3 nodes,

$$t(P_3) \geq t(K_2)^2 \tag{59}$$

and

$$t(C_4) \geq t(P_3)^2. \tag{60}$$

7.2 Re-proving some results in extremal graph theory

Let us start with deriving (54). By (57),

$$t(K_3) - 2t(P_3) + t(K_2K_1) \geq 0,$$

Using multiplicativity, we have $t(K_2 + K_1) = t(K_2)t(K_1) = t(K_2)$. Furthermore, by (59) we have $t(P_3) \geq t(K_2)^2$. This implies (54).

We can derive the following theorem of Moon and Moser [52] in a similar way: Let G be a graph with n nodes, and let N_r denote the number of complete subgraphs with r nodes. Then

$$\frac{N_{r+1}}{N_r} \geq \frac{1}{r^2 - 1} \left(r^2 \frac{N_r}{N_{r-1}} - n \right). \tag{61}$$

Using that $N_r = t(K_r, G)n^r/r!$, this inequality can be expressed in the following simpler form: For every $t \in \mathcal{T}_0$,

$$r \frac{t(K_r)}{t(K_{r-1})} \leq (r-1) \frac{t(K_{r+1})}{t(K_r)} + 1. \tag{62}$$

This shows that (54) is a special case, and the derivation of (54) above can be extended.

The (simplest, asymptotic) case of the Kruskal-Katona theorem,

$$t(K_3)^2 \leq t(K_2)^3, \tag{63}$$

also follows. One uses multiplicativity to write it as $t(K_3)^2 \leq t(K_2)t(K_2K_2)$; by monotonicity, it suffices to prove the stronger inequality $t(K_3)^2 \leq t(K_2)t(C_4)$; which then follows by considering the following submatrix of M_2 :

$$\begin{pmatrix} t(K_2) & t(K_3) \\ t(K_3) & t(C_4) \end{pmatrix}.$$

7.3 A 2-dimensional projection

Perhaps the most basic question about inequalities between homomorphism densities concerns edges and triangles. What are the possible pairs $(t(K_2, G), t(K_3, G))$? If we disregard number theoretic conditions like these numbers must be rational, we can ask: What are the possible pairs $(t(K_2), t(K_3))$, where $t \in \mathcal{T}_0$? In other words, what is the projection \mathcal{T}'_0 of \mathcal{T}_0 onto the 2-dimensional plane determined by the density of edges and density of triangles?

The answer to this questions turns out quite complicated and is only partially solved. Figure 2 describes the set \mathcal{T}'_0 . The upper boundary curve is given the edge-density is given by the Kruskal-Katona Theorem (63):

$$x^3 = y^2$$

(and this is tight for all densities). The lower bound is more complicated. Inequality (54) gives a parabola that is a lower bound. However, this is only tight for special values of the edge-density: $t(K_2) = 1 - 1/k$, $k = 1, 2, \dots$. For these values, a Turán graph (complete k -partite graph with equal color classes) gives equality. Bollobás [11] proved that in the intervals between these values, the lower boundary is above the chord. A conjecture for the curve was formulated in [42]: An extremal graph in the interval $1 - 1/k \leq t(K_2) \leq 1 - 1/(k+1)$ is a complete $(k+1)$ -partite

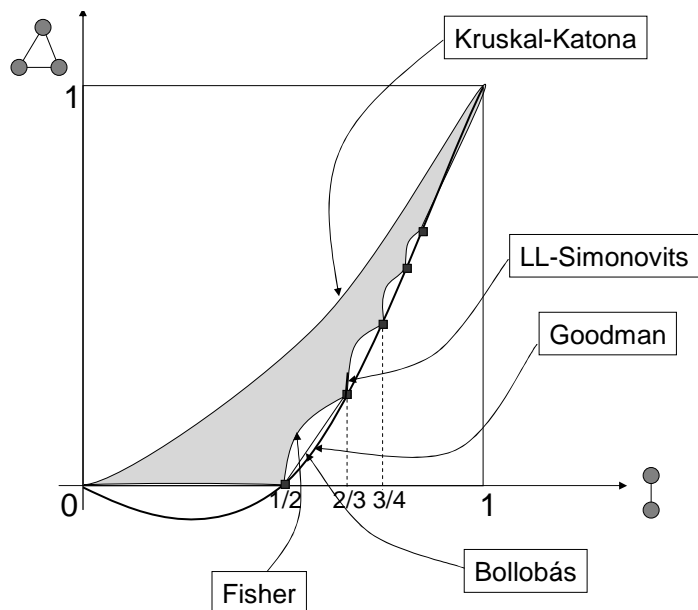


Figure 2: The region of possible edge densities and triangle densities. (The concave arcs are distorted to make the qualitative properties more visible.)

graph with $k - 1$ equal color classes. The size of the two special color classes must be optimized to make sure that the density of triangles is minimized, subject to the given edge density. For the interval $1/2 \leq t(K_2) \leq 2/3$, the optimization yields the cubic curve

$$9y^2 - 18xy + 8x - 3x^2 + 6x^3 = 0.$$

This case was proved by Fisher [24]; a recent proof by Razborov [54] uses methods quite closely related to those described in this paper. The conjecture was proved in [42] if $t(K_2)$ was in a small neighborhood of any one of the special values $1 - 1/k$. The general case is open.

One should add that this conjecture is all that is needed to describe \mathcal{T}'_0 : it is easy to see that between two points of \mathcal{T}'_0 the plane on a vertical line, the whole interval is contained in \mathcal{T}'_0 .

8 Graphs with bounded degree

Fix a positive integer D , and consider (as the middle graph G in (3)) only graphs in which all degrees are at most D . The normalization $t(F, G)$ we used before does not make sense any more, since (if G is large) the probability that a random mapping from a given F is a homomorphism is very small.

How to normalize $\text{hom}(F, G)$ in this case? Let $n = |V(G)|$ and $k = |V(F)|$. If F is not connected, then $\text{hom}(F, G)$ is just the product of the numbers $\text{hom}(F', G)$, where F' is a connected component of F , so we may restrict our attention to the case when F is connected; then there is the following rather obvious upper bound:

$$\text{hom}(F, G) \leq n \cdot D^{k-1}.$$

So it makes sense to consider

$$\tau(F, G) = \frac{\text{hom}(F, G)}{n} \tag{64}$$

(Since we consider D and F fixed, it does not seem to help to divide by D^{k-1} .)

8.1 Convergence for graphs with bounded degree

Let (G_1, G_2, \dots) be a sequence of graphs whose degrees are uniformly bounded by D . We say that this sequence is *locally convergent*, or just *convergent*, if $\tau(F, G_n)$ tends to a limit for every connected graph F .

There is another way to define this. Given a graph G with all degrees bounded by D , and a positive integer r , let $S_G(v, r)$ denote the neighborhood of node v with radius r . We consider $S_G(v, r)$ as a rooted graph (where v is its root). For fixed D and r , there is a finite number of possible neighborhoods, and we can make a statistic of these: we denote by $\nu_G(N, r)$ the fraction of nodes $v \in V(G)$ for which $S_G(v, r) \cong N$. Then ν_G is a probability distribution on all possible r -neighborhoods.

For bounded degrees, it is easy to see that the probability distributions $\nu_{G_n}(\cdot, r)$ tend to a limit distribution for all $r > 0$ if and only if G_n converges in the sense defined above. The convergence of the probability distributions $\nu_{G_n}(\cdot, r)$ can therefore be used as an alternative characterization of convergent sequences (G_1, G_2, \dots) of graphs with degrees bounded by D .

This notion of convergence was introduced implicitly by Aldous [1], and explicitly by Benjamini and Schramm [17]. It was extended to the case of bounded average degree by Lyons [49]. The limit object has several descriptions, the strongest is due to Elek [21]: A *continuous graphing* is an infinite graph on $[0, 1]$ that has the following structure: we take a finite number of continuous measure preserving involutions $\phi_1, \dots, \phi_N : [0, 1] \rightarrow [0, 1]$, and connect every $x \in [0, 1]$ to every $\phi_i(x)$ ($i = 1, \dots, N$) such that $x \neq \phi_i(x)$ by an edge.

Every graphing defines a probability distribution on countable graphs with bounded degree with a specified root: we pick a uniform random point x in $[0, 1]$, and consider the connected component of the graphing containing x , with root x . This yields the description of the limit by Benjamini and Schramm.

8.2 Left and right convergence

The theory of convergence of graphs with bounded degree is much less satisfactory than the analogous theory of dense graphs described above. In particular, the “right” notion of distance and a powerful analogue of Szemerédi’s Lemma are missing. But there are some nontrivial facts, relating homomorphisms into a large graph to homomorphisms from it, which can be thought of as analogues of some results in Section 6.

Let us start with an example.

Example 8.1 Consider the sequence of cycles C_n . It is trivial that this is convergent, but the numbers $\text{hom}(C_n, K_2)$ alternate between 0 and 2.

But in a sense parity is all that goes wrong. Let us consider the subsequence C_{2n} of even cycles. Then for any graph H ,

$$\text{hom}(C_{2n}, H) = \sum_{i=1}^k \lambda_i^{2n},$$

where $\lambda_1 \geq \dots \geq \lambda_k$ are the eigenvalues of H . Hence

$$\text{hom}(C_{2n}, H)^{1/(2n)} \rightarrow \lambda_1.$$

It is perhaps more usual to take the logarithm here, to get the sequence $\frac{1}{2n} \log \text{hom}(C_{2n}, H)$ (in statistical physics, this parameter is called the free energy or pressure of the H -colorings of C_{2n}). This sequence is convergent for every H , and the limiting parameter is $\log \lambda_1(H)$.

The sequence of odd cycles behaves similarly, but the limit parameter $a(H) = \lim_{n \rightarrow \infty} \frac{1}{2n+1} \log \text{hom}(C_{2n+1}, H)$ is a bit more complicated to describe: $a(H)$ is the logarithm of largest eigenvalue of its non-bipartite components ($a(H) = -\infty$ if H is bipartite).

Based on this example, let us consider a convergent sequence (G_n) of graphs with bounded degree, and ask for which graphs H the numbers

$$\widehat{\mathcal{P}}(G_n, H) = \frac{1}{|V(G_n)|} \log \text{hom}(G_n, H)$$

converge.

Before discussing this further, let us rewrite the right hand side using Möbius inversion. Let

$$\psi(G, H) = \sum_{V \subset V(G)} (-1)^{|V(G)|-|V|} \log \text{hom}(G[V], H).$$

Then

$$\log \text{hom}(G, H) = \sum_{V \subset V(G)} \psi(G[V], H) = \sum_F \frac{\text{ind}(F, G)}{|V(F)|!} \psi(F, H), \quad (65)$$

where the sum goes over all finite simple graphs F and $\text{ind}(F, G)$ is the number of embedding of F into G as an induced subgraph. We thus have shown that

$$\widehat{\mathcal{P}}(G, H) = \sum_F \frac{\tau_{\text{ind}}(F, G)}{|V(F)|!} \psi(F, H) \quad (66)$$

where

$$\tau_{\text{ind}}(F, G) = \frac{1}{|V(G)|} \text{ind}(F, G). \quad (67)$$

Recall that the numbers $\text{ind}(F, G)$ can be obtained from the homomorphism numbers $\text{hom}(F, G)$ by Möbius inversion (see (14) and (15)). If G_n is convergent in the sense that the normalized homomorphism numbers $\tau(F, G_n)$ are convergent for all F , then the numbers $\tau_{\text{ind}}(F, G_n)$ are convergent as well; let $\tau_{\text{ind}}(F, G_n) \rightarrow \tau(F)$. One might therefore hope that convergence of the sequence G_n implies convergence of the free energies $\widehat{\mathcal{P}}(G_n, H)$, with the limit given by the infinite sum

$$\mathcal{P}(\tau, H) = \sum_F \frac{\tau(F)}{|V(F)|!} \psi(F, H). \quad (68)$$

It is clear, however, that this cannot be true in general; an easy counterexample is the example $G_n = C_n$ discussed above. But it turns out that we can prove the convergence under suitable conditions on H and the maximal degree D of the graphs G_n .

Theorem 8.2 [15] *If (G_n) is a sequence of graphs with all degrees bounded by a constant D , and (G_n) is convergent, then $\widehat{\mathcal{P}}(G_n, H)$ has a limit for every unweighted target graph H in which all degrees are at least $(1 - \frac{1}{2D})|V(H)|$.*

The conceptually most transparent proof of the theorem (with a slightly worse constant) proceeds by proving uniform convergence of the expansion (68) using the method of cluster expansions. As stated, the theorem can be proven using Dobrushin's uniqueness theorem for uniform H -colorings on G_n .

The cluster expansion proof also gives an analogue of Theorem 8.2 for weighted graphs (see [15]) and allows to prove that convergence from the right for weighted graphs implies convergence from the left. In fact, we only need convergence from the right for graphs H that are small perturbations of completely looped complete graphs with edge weights 1 to conclude that G_n is convergent from the left.

Theorem 8.3 [15] *If (G_n) is a sequence of graphs with all degrees bounded by a constant D , and $\mathcal{P}(G_n, H)$ is convergent for every weighted graph H which is a looped complete graph with all edgeweights arbitrarily close to 1, then (G_n) is convergent.*

References

- [1] D.J. Aldous: Tree-valued Markov chains and Poisson-Galton-Watson distributions, in: *Microsurveys in Discrete Probability* (D. Aldous and J. Propp, editors), DIMACS Ser. Discrete Math. Theoret. Comput. Sci. **41** (1998) Amer. Math. Soc., Providence, RI. (1998), 1–20.
- [2] N. Alon, W. Fernandez de la Vega, R. Kannan and M. Karpinski: Random sampling and approximation of MAX-CSPs, *J. Comput. System Sci.* **67** (2003), 212–243.
- [3] N. Alon and A. Naor: Approximating the Cut-Norm via Grothendieck’s Inequality, preprint: <http://research.microsoft.com/research/theory/naor/homepage%20files/cutnorm.pdf>
- [4] N. Alon and A. Shapira, Every monotone graph property is testable, Proc. of the 37th ACM STOC, Baltimore, ACM Press (2005), to appear: <http://www.math.tau.ac.il/~nogaa/PDFS/MonotoneSTOC.pdf>
- [5] E. Babson and D. Kozlov: Complexes of graph homomorphisms, arXiv: <http://lanl.arxiv.org/abs/math.CO/0310056>
- [6] E. Babson and D. Kozlov: Proof of the Lovász Conjecture, arXiv: <http://lanl.arxiv.org/abs/math.CO/0402395>
- [7] A. Barabasi and R. Albert: Emergence of scaling in random networks, *Science* **286** (1999) 509–512.
- [8] B. Bollobás and O. Riordan: Mathematical results on scale-free random graphs, in: *Handbook of graphs and networks*, 1–34, Wiley-VCH, Weinheim, 2003.
- [9] M. Biskup, C. Borgs, J.T. Chayes, L. Kleinwaks and R. Kotecký: Partition function zeros at first-order phase transitions: A general analysis, *Commun. Math. Phys.* **251** (2004) 79–131.
- [10] B. Bollobás, O.M. Riordan, J. Spencer, and G. Tusnády: The degree sequence of a scale-free random graph process, *Random Structures and Algorithms* **18** (2001), 279–290.
- [11] B. Bollobás: Relations between sets of complete subgraphs, in: *Combinatorics*, Proc. 5th British Comb. Conf. (ed. C.St.J.A. Nash-Williams, J. Sheehan), Utilitas Math. (1975), 79–84.
- [12] B. Bollobás, C. Borgs, J. T. Chayes, and O. Riordan: Directed scale-free graphs, in: *Proceedings of the 14th ACM-SIAM Symposium on Discrete Algorithms*, 132–139, 2003.
- [13] C. Borgs: Statistical Physics Expansion Methods for Combinatorics and Computer Science, CBMS lecture notes (in preparation).
- [14] C. Borgs and J. Chayes and L. Lovász: Unique limits of dense graph sequences (in preparation).
- [15] C. Borgs, J. Chayes, J. Kahn, L. Lovász and V.T. Sós: Convergent sequences for sparse graphs (in preparation).
- [16] C. Borgs, J. Chayes, L. Lovász, V.T. Sós and K. Vesztergombi: Convergent sequences of dense graphs (in preparation).
- [17] I. Benjamini and O. Schramm: Recurrence of Distributional Limits of Finite Planar Graphs, *Electronic J. Probab.* **6** (2001), no. 23, 1–13.

- [18] M.-D. Choi: Tricks or Treats with the Hilbert Matrix, *Amer. Math. Monthly* **90** (1983), 301–312.
- [19] F. Chung, R.L. Graham and R.M. Wilson: Quasi-random graphs, *Combinatorica* **9** (1989), 345–362.
- [20] C. Cooper and A. Frieze: On a general model of web graphs, *Rand. Struct. Alg.* **22** (2003) 311–335.
- [21] G. Elek: Graphings and graph sequences (preprint).
- [22] K. Drühl and H. Wagner: Algebraic formulation of duality transforms for abelian lattice models, *Anal. of Physics* **141** (1982), 225–253.
- [23] P. Erdős and L. Lovász, J. Spencer: Strong independence of graphcopy functions, in: *Graph Theory and Related Topics*, Academic Press (1979), 165–172.
- [24] D.C. Fisher: Lower bounds on the number of triangles in a graph, *J. Graph Theory* **13** (1989), 505–512.
- [25] D.C. Fisher and J. Ryan: Conjectures on the number of complete subgraphs, in: *Proc. of the 20-th Southeastern Conf. on Comb., Graph Theory, and Computing*, Congr. Numer. **70** (1990), 217–219.
- [26] D.C. Fisher and A. Solow: Dependence polynomials, *Discrete Math.* **82** (1990), 251–258.
- [27] M. Freedman and L. Lovász, D. Welsh (unpublished).
- [28] M. Freedman, L. Lovász and A. Schrijver: Reflection positivity, rank connectivity, and homomorphism of graphs (MSR Tech Report # MSR-TR-2004-41).
<ftp://ftp.research.microsoft.com/pub/tr/TR-2004-41.pdf>
- [29] A. Frieze and R. Kannan: Quick approximation to matrices and applications, *Combinatorica* **19** (1999), 175–220.
- [30] A.W. Goodman: On sets of acquaintances and strangers at any party, *Amer. Math. Monthly* **66** (1959) 778–783.
- [31] O. Goldreich, S. Goldwasser and D. Ron: Property testing and its connection to learning and approximation, *J. ACM* **45** (1998), 653–750.
- [32] W.T. Gowers: Lower bounds of tower type for Szemerédi’s Uniformity Lemma, *Geom. Func. Anal.* **7** (1997), 322–337.
- [33] P. de la Harpe and V.F.R. Jones: Graph Invariants Related to Statistical Mechanical Models: Examples and Problems, *Journal of Combinatorial Theory B* **57** (1993), 207–227.
- [34] P. Hell and J. Nešetřil: *Graphs and Homomorphisms*, Oxford University Press, 2004.
- [35] J. Komlós and M. Simonovits: Szemerédi’s Regularity Lemma and its applications in graph theory, in: *Combinatorics, Paul Erdos is Eighty* (D. Miklos et. al, eds.), Bolyai Society Mathematical Studies **2** (1996), pp. 295–352.
- [36] T.D. Lee and C.N. Yang, Statistical theory of equations of state and phase transitions: II. Lattice gas and Ising model, *Phys. Rev.* **87** (1952) 410–419.
- [37] L. Lovász: Operations with structures, *Acta Math. Hung.* **18** (1967), 321–328.

- [38] L. Lovász: Direct product in locally finite categories, *Acta Sci. Math. Szeged* **23** (1972), 319-322.
- [39] L. Lovász: The rank of connection matrices and the dimension of graph algebras, *Europ. J. Combin.*, to appear:
<http://research.microsoft.com/~lovasz/homdim.pdf>
- [40] L. Lovász and A. Schrijver: Graph parameters and semigroup functions (manuscript).
- [41] L. Lovász and M. Simonovits: On the number of complete subgraphs of a graph (M. Simonovits), in: *Combinatorics*, Proc. 5th British Comb. Conf. (ed. C.St.J.A.Nash-Williams, J.Sheehan), Utilitas Math. (1976) , 439-441.
- [42] L. Lovász and M. Simonovits: On the number of complete subgraphs of a graph II, in: *Studies in Pure Math.*, To the memory of P. Turán (ed. P. Erdős), Akadémiai Kiadó (1983), 459-495.
- [43] L. Lovász and V.T. Sós: Generalized quasirandom graphs, preprint:
<http://research.microsoft.com/~lovasz/quasirandom4.pdf>
- [44] L. Lovász and B. Szegedy: Limits of dense graph sequences, Microsoft Research Technical Report MSR-TR-2004-79:
<ftp://ftp.research.microsoft.com/pub/tr/TR-2004-79.pdf>
- [45] L. Lovász and B. Szegedy: Szemerédi's Lemma for the analyst, Microsoft Research Technical Report MSR-TR-2005-90:
<ftp://ftp.research.microsoft.com/pub/tr/TR-2005-90.pdf>
- [46] L. Lovász and B. Szegedy: Contractors and connectors of graph algebras Microsoft Research Technical Report TR-2005-91:
<ftp://ftp.research.microsoft.com/pub/tr/TR-2005-91.pdf>
- [47] L. Lovász and B. Szegedy: Graph limits and testing hereditary graph properties, Microsoft Research Technical Report MSR-TR-2005-110:
<ftp://ftp.research.microsoft.com/pub/tr/TR-2005-110.pdf>
- [48] L. Lovász and B. Szegedy: Moments of 2-variable functions (in preparation).
- [49] R. Lyons: Asymptotic Enumeration of Spanning Trees, preprint:
<http://mypage.iu.edu/~rdlyons/pdf/est.pdf>
- [50] H. Mahmoud, R. Smythe and J. Szymański: On the structure of plane-oriented trees and their branches. *Random Struct. Alg.* **3** (1993) 255–266.
- [51] J. Matoušek: *Using the Borsuk-Ulam Theorem: Lectures on Topological Methods in Combinatorics and Geometry* Springer, 2003.
- [52] J.W. Moon, L. Moser: On a problem of Turán, *Magyar Tud. Akad. Mat. Kutat Int. Kzl.* **7** (1962), 283–286.
- [53] D. J. de S. Price: A general theory of bibliometric and other cumulative advantage processes. *J. Amer. Soc. Inform. Sci.*, **27** (1976) 292–306.
- [54] A. Razborov (unpublished).

- [55] M. Simonovits and V.T. Sós: Szemerdi's partition and quasirandomness, *Random Structures Algorithms* **2** (1991), 1–10.
- [56] M. Simonovits and V.T. Sós: Hereditary extended properties, quasi-random graphs and induced subgraphs, *Combinatorics, Probability and Computing* **12** (2003), 319–344.
- [57] M. Simonovits and V.T. Sós: Hereditarily extended properties, quasi-random graphs and not necessarily induced subgraphs. *Combinatorica* **17** (1997), 577–596.
- [58] B. Szegedy: Edge models and reflection positivity, preprint:
<http://arxiv.org/abs/math.CO/0505035>
- [59] A. Thomason: Pseudorandom graphs, in: *Random graphs '85* North-Holland Math. Stud. **144**, North-Holland, Amsterdam, 1987, 307–331.
- [60] C.N. Yang and T.D. Lee: Statistical theory of equations of state and phase transitions: I. Theory of condensation, *Phys. Rev.* **87** (1952), 404–409.
- [61] G. U. Yule: A mathematical theory of evolution, based on the conclusions of Dr. J. C. Willis. *Philos. Trans. Roy. Soc. London*, Ser. B **213** (1924) 21–87.
- [62] H. Whitney: The coloring of graphs, *Ann. of Math.* **33** (1932), 688–718.

AMS Subject Classification: Primary 05C99, Secondary 05C35, 05C80, 82B99

Keywords: Graph homomorphism, partition function, extremal graphs, Szemerédi Regularity Lemma, connection matrix, convergent graph sequence, limits of graphs.