



An analytic approach to stability

Oleg Pikhurko*

Department of Mathematical Sciences, Carnegie Mellon University, Pittsburgh, PA 15213, United States

ARTICLE INFO

Article history:

Received 1 December 2008

Received in revised form 2 July 2010

Accepted 4 July 2010

Available online 3 August 2010

Keywords:

Graph limits

Stability property

Erdős–Simonovits Stability Theorem

Turán problem

ABSTRACT

The stability method is very useful for obtaining exact solutions of many extremal graph problems. Its key step is to establish the *stability property* which, roughly speaking, states that any two almost optimal graphs of the same order n can be made isomorphic by changing $o(n^2)$ edges.

Here we show how the recently developed theory of graph limits can be used to give an analytic approach to stability. As an application, we present a new proof of the Erdős–Simonovits stability theorem.

Also, we investigate various properties of the edit distance. In particular, we show that the combinatorial and fractional versions are within a constant factor from each other, thus answering a question of Goldreich, Krivelevich, Newman, and Rozenberg.

© 2010 Elsevier B.V. All rights reserved.

1. Introduction

The notion of the left convergence of graph sequences was introduced by Borgs, Chayes, Lovász, Sós, and Vesztegombi (2003, unpublished) and was developed in [4,6–9,12,16,23,28,29,31] and other papers. Benjamini and Schramm [1] introduced convergence for graphs of bounded maximum degree. Tardos [38] defined limits of trees. Lovász [27] presents a nice survey of this area.

It is possible that graph limits will become a very powerful tool, especially in extremal graph theory. The left limits are closely related to the (Weak) Regularity lemma, see Lovász and Szegedy [29], which is a very important and useful result. The algebraic characterization of Lovász and Szegedy [28, Theorem 2.2] of possible limiting subgraph densities seems to have great potential. Although these developments are very recent, Razborov [35,36] has already used graph limits to obtain a spectacular progress on the long-standing Rademacher–Turán problem. Also, graph limits have proved helpful for property and parameter testing, see Benjamini et al. [2], Borgs et al. [5], Elek [11], Lovász and Szegedy [30], and other.

Here is an example of how graph limits may be applied to extremal graph problems.

Suppose that the convergence on graphs is encoded by a compact metric space (\mathcal{X}, δ) and a map that corresponds to each graph G a point $A(G)$ of \mathcal{X} and respects graph isomorphism (that is, $A(G) = A(H)$ whenever $G \cong H$). Then we say that a sequence of graphs $(G_n)_{n \in \mathbb{N}}$ converges if the sequence $(A(G_n))_{n \in \mathbb{N}}$ is Cauchy in the metric δ . In this case, the limit of $(G_n)_{n \in \mathbb{N}}$ is the (unique) limiting point of the sequence $(A(G_n))_{n \in \mathbb{N}}$ in (\mathcal{X}, δ) , which exists since (\mathcal{X}, δ) is compact.

Suppose that we are given a graph parameter f , that is, a function on graphs that respects graph isomorphism, and a graph property \mathcal{P} , that is, a family of graphs closed under isomorphism. Let $\mathcal{P}_n = \{G \in \mathcal{P} : v(G) = n\}$ consist of all graphs in \mathcal{P} with n vertices. The corresponding extremal (f, \mathcal{P}) -problem is to determine for each n

$$\begin{aligned} \text{ex}_f(n, \mathcal{P}) &= \max\{f(G) : G \in \mathcal{P}_n\}, \\ \mathcal{E}\mathcal{X}_f(n, \mathcal{P}) &= \{G \in \mathcal{P}_n : f(G) = \text{ex}_f(n, \mathcal{P})\}, \end{aligned}$$

* Tel.: +1 412 268 9782; fax: +1 412 268 6380.

E-mail address: pikhurko@andrew.cmu.edu.

URL: <http://www.math.cmu.edu/~pikhurko>.

the maximum of $f(G)$ over all graphs from \mathcal{P}_n as well as the set of *extremal graphs*, i.e. graphs that achieve this maximum. For example, if we let $h(G)$ be the maximum size of a *homogeneous set* (a clique or an independent set) in a graph G , $f(G) = -h(G)/\log_2 v(G)$ be its scaled version, and \mathcal{P} be the family of all graphs, then we obtain the inverse problem for the diagonal Ramsey numbers. Many extremal graph problems can be represented this way.

Let us try to formulate some approximation (the “limiting” case) of the problem as $n \rightarrow \infty$. We suggest the following definition. Let the *limit set* $\text{LIM}(f, \mathcal{P})$ consist of those $x \in \mathcal{X}$ for which there is an infinite increasing sequence of indices $n_1 < n_2 < n_3 < \dots$ and graphs $G_{n_i} \in \mathcal{P}_{n_i}$ such that

$$\lim_{i \rightarrow \infty} (f(G_{n_i}) - \text{ex}_f(n_i, \mathcal{P})) = 0 \tag{1}$$

and the sequence $(G_{n_i})_{i \in \mathbb{N}}$ converges to x , that is,

$$\lim_{i \rightarrow \infty} \delta(A(G_{n_i}), x) = 0.$$

Although we are ultimately interested in $\mathcal{E}\mathcal{X}_f(n, \mathcal{P})$, we do not require that $G_{n_i} \in \mathcal{E}\mathcal{X}_f(n_i, \mathcal{P})$ here. One of the reasons is that we often know $\text{ex}_f(n, \mathcal{P})$ asymptotically but not exactly, in which case one can test if (1) holds but not the membership in $\mathcal{E}\mathcal{X}_f(n, \mathcal{P})$.

Now, we can try to study the set $\text{LIM}(f, \mathcal{P})$, which is independent of n . If we succeed in completely describing it, then we might be able to discover some information about extremal graphs. Indeed, if we select arbitrary extremal graphs $G_n \in \mathcal{E}\mathcal{X}_f(n, \mathcal{P})$ for infinitely many n , then, by the compactness of (\mathcal{X}, δ) , there always is a convergent subsequence, whose limit belongs to $\text{LIM}(f, \mathcal{P})$. Suppose that this convergence implies some structural statement (in purely graph theoretical terms) that necessarily occurs for infinitely many of the selected extremal graphs. Then one can conclude that the statement fails only for finitely many extremal graphs overall.

One can call this approach the *limit method*. It applies in principle to very general settings. For example, the families \mathcal{P}_n need not be related to each other for different n nor the graph parameter f has to behave well with respects to taking limits: the above definitions make perfect sense for arbitrary f and \mathcal{P} (and $\text{LIM}(f, \mathcal{P}) \neq \emptyset$ provided infinitely many of \mathcal{P}_n 's are non-empty). Also, the definition of the limit set may be modified to work with other extremal problems, those which are indexed by a different parameter than the order of a graph.

Since the limit method deals only with some approximation of the extremal problem, one would hope to obtain only the asymptotic of $\text{ex}_f(n, \mathcal{P})$ at best. However, this approach might work well together with the so-called *stability method* that has proved very useful in solving many extremal problems exactly (including the description of $\mathcal{E}\mathcal{X}_f(n, \mathcal{P})$) for all large n .

The *stability method* proceeds as follows. Suppose that we know the value of $\text{ex}_f(n, \mathcal{P})$ asymptotically and that we have some set \mathcal{C}_n believed to be exactly the set $\mathcal{E}\mathcal{X}_f(n, \mathcal{P})$ for all large n . Assume that $\mathcal{C}_n \subseteq \mathcal{P}_n$ and f is constant on \mathcal{C}_n . (Of course, these assumptions are necessary for $\mathcal{C}_n = \mathcal{E}\mathcal{X}_f(n, \mathcal{P})$ and, usually, they are easy to check.) Given \mathcal{C}_n , we have to prove first that for any almost extremal graph $G \in \mathcal{P}_n$ (i.e. $G \in \mathcal{P}_n$ satisfying $f(G) = \text{ex}_f(n, \mathcal{P}) - o(1)$) there is $H \in \mathcal{C}_n$ such that $\hat{\delta}_1(G, H) = o(1)$, where

$$\hat{\delta}_1(G, H) = \frac{2}{n^2} \min\{|E(G)\Delta\sigma(E(H))|: \text{bijective } \sigma : V(H) \rightarrow V(G)\} \tag{2}$$

is the *edit distance* between two graphs of the same order n : it is $2/n^2$ times the minimum number of adjacencies that one has to change in G to make it isomorphic to H . Next, pick an arbitrary $G \in \mathcal{E}\mathcal{X}_f(n, \mathcal{P})$ for a sufficiently large n . By the above, we know that G is close in the distance $\hat{\delta}_1$ to the graph property \mathcal{C}_n . In order to complete the proof, it is enough to argue that G is necessarily in \mathcal{C}_n . Here we can use various arguments, such as applying “local improvements” to G or arguing that every “wrong” adjacency in G bears too much penalty. Knowing all but $o(n^2)$ edges of G greatly helps in this task; this is what makes this method so successful. This approach was pioneered by Simonovits [37] in the late 1960s. It has been used to obtain exact solutions for an impressive array of problems since then.

The term “stability” refers to the property that every almost extremal graph has structure almost the same as some extremal graph. A class of extremal problems for which this method seems to be particularly suited is when there is only one pattern independent of n for all almost extremal graphs. In order to state this property formally, we have to define a version of edit distance for arbitrary pairs of graphs. Namely, the δ_1 -distance, denoted by $\delta_1(G, H)$, between graphs G and H on vertex sets $\{x_1, \dots, x_m\}$ and $\{y_1, \dots, y_n\}$ respectively is the minimum over all non-negative $m \times n$ -matrices $A = (\alpha_{i,j})$ with row sums $1/m$ and column sums $1/n$ of

$$\delta_1(G, H, A) = \sum_{(i,j,g,h) \in \Delta} \alpha_{i,g} \alpha_{j,h}, \tag{3}$$

where Δ consists of all quadruples $(i, j, g, h) \in [m]^2 \times [n]^2$ such that exactly one of the following two relations holds: either $\{x_i, x_j\} \in E(G)$ or $\{y_g, y_h\} \in E(H)$. Informally speaking, we view G and H as uniformly vertex-weighted graphs of total weight 1 while $\alpha_{i,j}$ tells what fraction of vertex x_i is mapped into vertex y_j . It is not hard to show (see Section 3) that this defines a *pre-metric* on the set of graphs, that is, δ_1 is symmetric, non-negative and satisfies the Triangle Inequality (but may assume value zero on distinct graphs: e.g. $\delta_1(K_{m,m}, K_{n,n}) = 0$ for any $m, n > 0$).

Note that, for graphs G_1 and G_2 of the same order, we trivially have $\hat{\delta}_1(G_1, G_2) \geq \delta_1(G_1, G_2)$. This inequality is in general strict (see Arie Matsliah’s example presented in the technical report [20, Appendix B] or Example 13 here). However, we prove in Lemma 14 that

$$\hat{\delta}_1(G_1, G_2) \leq 3\delta_1(G_1, G_2), \tag{4}$$

answering in the affirmative an open question posed by Goldreich et al. [20, Section 6] (see [21] for the journal version).

Now, let us say that the extremal (f, \mathcal{P}) -problem is *stable* if for every $\varepsilon > 0$ there are $\varepsilon' > 0$ and n_0 such that for every $n_1, n_2 \geq n_0$ and every two graphs G_1, G_2 with $G_i \in \mathcal{P}_{n_i}$ and $f(G_i) \geq \text{ex}_f(n_i, \mathcal{P}) - \varepsilon'$, for $i = 1, 2$, we necessarily have $\delta_1(G_1, G_2) < \varepsilon$. Theorem 15 here gives an alternative characterization of stable extremal problems. However, we postpone the exact statement as well as the proof until Section 5 after we define graph limits in Section 2 and extend the distance δ_1 to them in Section 3.

For example, our approach applies to the *Turán problem* that asks for the maximum size of an \mathcal{F} -free graph of order n . This is a central question of extremal graph theory that was introduced by Turán [39]. Its scaled version can be represented in our notation as $\text{ex}_\rho(n, \text{Forb}(\mathcal{F}))$, where $\rho(G) = 2e(G)/(v(G))^2$ denotes the *edge density* of G and $\text{Forb}(\mathcal{F})$ consists of all \mathcal{F} -free graphs. By applying our Theorem 15, we obtain a new proof of the following celebrated result in Section 6.

Theorem 1 (The Erdős–Simonovits Stability Theorem [13,37]). *For every (possibly infinite) family \mathcal{F} of non-empty graphs the extremal $(\rho, \text{Forb}(\mathcal{F}))$ -problem is stable.*

It is well-known that $\text{ex}_\rho(n, \text{Forb}(\mathcal{F})) = \frac{r-1}{r} + o(1)$, where

$$r = \min\{\chi(F) : F \in \mathcal{F}\} - 1 \geq 1, \tag{5}$$

and the lower bound is given by the *Turán graph* $T_r(n) \in \text{Forb}(\mathcal{F})$, the complete r -partite graph on $[n]$ with parts of size $\lfloor n/r \rfloor$ or $\lceil n/r \rceil$. Thus, by (4), Theorem 1 can be reformulated in the more familiar form that for any $\varepsilon > 0$ there are $\varepsilon' > 0$ and n_0 such that every \mathcal{F} -free graph with $n \geq n_0$ vertices and at least $(\frac{r-1}{r} - \varepsilon') \binom{n}{2}$ edges can be made isomorphic to $T_r(n)$ by changing at most $\varepsilon \binom{n}{2}$ edges.

Theorem 1 was first applied by Simonovits [37] to determine the exact value of the Turán function $\text{ex}(n, F)$ for various forbidden graphs F . This theorem has a huge number of applications. For example, Theorem 1 turned up quite a few times in the author’s research alone: see the papers with Jiang [24], Lazebnik and Woldar [25], Loh and Sudakov [26], Mubayi [32], Yilma [34]. Another proof of Theorem 1 was recently discovered by Füredi [18].

2. Graph limits

Here we present the main definitions of “dense” graph limits. This notion of convergence (also called the *left convergence* in [8, Section 2.2]) will be of main interest for this paper. We refer the reader to e.g. [8] for further details.

Until recently, the measure-theoretic methods were rare in discrete mathematics (if compared with, for example, linear algebra or topological tools). Bearing in mind a combinatorialist reader who does not use real analysis in research, we decided to take an extra care with measure theoretic concepts and to give references or detailed explanations whenever feasible (even of some fairly standard results). For example, the result of Lemma 11 is stated in [30, page 5] without proof; here we carefully fill in all missing details. All analytical terms that we do not define can be found in the book by Folland [15].

Let \mathbb{R} denote the set of reals and $I \subseteq \mathbb{R}$ denote the closed unit interval $[0, 1]$. For $Y \subseteq \mathbb{R}^n$, let $\mathcal{L}_Y = \{A \cap Y : A \in \mathcal{L}\}$ denote the restriction of the σ -algebra \mathcal{L} of Lebesgue measurable subsets of \mathbb{R}^n to Y . If $Y \subseteq \mathbb{R}^n$ is Lebesgue measurable, then μ_Y denotes the restriction of the Lebesgue measure μ to \mathcal{L}_Y . Let $\mathcal{B}_Y = \{A \cap Y : A \in \mathcal{B}\}$ be the restriction of the σ -algebra \mathcal{B} of Borel subsets of \mathbb{R}^n to Y . When the set Y is clear from the context, we write \mathcal{L}, μ , and \mathcal{B} for \mathcal{L}_Y, μ_Y , and \mathcal{B}_Y respectively. We say that some property holds *almost everywhere* (abbreviated as *a.e.*) if the set of x for which it fails has Lebesgue measure 0. A measurable function is called *simple* if it assumes only finitely many values.

A function $W : I^2 \rightarrow \mathbb{R}$ is called *symmetric* if $W(x, y) = W(y, x)$ for every $x, y \in I$. Let \mathcal{W} consist of all symmetric bounded measurable functions $W : (I^2, \mathcal{L}) \rightarrow (\mathbb{R}, \mathcal{B})$. Following [8], we call the elements of \mathcal{W} *graphons*. Let \mathcal{W}_I consist of those graphons $W \in \mathcal{W}$ such that $0 \leq W(x, y) \leq 1$ for every $x, y \in I$.

A function $\phi : (I, \mathcal{L}, \mu) \rightarrow (I, \mathcal{L}, \mu)$ is called *measure preserving* if it is measurable and $\mu(\phi^{-1}(A)) = \mu(A)$ for any $A \in \mathcal{L}_I$. Let Φ consist of all such functions. Note that $\phi \in \Phi$ may be very far from being invertible as e.g. $\phi(x) = 2x - \lfloor 2x \rfloor$ shows. Let Φ_0 consist of bijections $\phi : I \rightarrow I$ such that both ϕ and ϕ^{-1} belong to Φ . Clearly, each of Φ and Φ_0 is closed under taking compositions of functions. For $\phi \in \Phi$ and $W \in \mathcal{W}$, let W^ϕ be defined by $W^\phi(x, y) = W(\phi(x), \phi(y))$. It is easy to see that $W^\phi \in \mathcal{W}$ and for any $\psi \in \Phi$, we have

$$(W^\phi)^\psi = W^{(\phi \circ \psi)}. \tag{6}$$

A few remarks are in order. It is standard (see e.g. [15, page 44]) to consider the σ -algebra \mathcal{B} of *Borel* sets whenever (a subset of) \mathbb{R}^n is the range of a function from some measure space. This has many advantages: we can add or multiply such functions [15, Proposition 2.6], take pointwise limits [15, Proposition 2.7], etc., with the resulting function being measurable. In particular, by [15, Theorem 6.6], the vector space

$$L^1 := L^1(I^2, \mathcal{L}, \mu) = \{\text{integrable } W : (I^2, \mathcal{L}, \mu) \rightarrow (\mathbb{R}, \mathcal{B})\} / \sim, \tag{7}$$

where we write $U \sim W$ iff $U = W$ a.e., is a Banach space with respect to the ℓ_1 -norm

$$\|W\|_1 = \int_{I \times I} |W(x, y)| d\mu(x, y). \tag{8}$$

On the other hand, DiBenedetto [10, Section 14.1] demonstrates that the set of measurable functions from (I, \mathcal{L}) to (I, \mathcal{L}) is not closed under taking pointwise limits (nor under multiplication, nor under addition, even if we take the interval $[0, 2]$ as the new range, as some easy modifications of his example can show). Note that, by definition, the set Φ consists of Lebesgue-to-Lebesgue measurable functions (so that, e.g., for every $W \in \mathcal{W}_I$ and $\phi \in \Phi$, we have $W^\phi \in \mathcal{W}_I$).

One can show that for any $W \in \mathcal{W}$ there is $U \in \mathcal{W}$ such that $W = U$ a.e. and U is measurable as a function $(I^2, \mathcal{B}) \rightarrow (\mathbb{R}, \mathcal{B})$. (Indeed, by writing the values of W in base 2, represent $W = \sum_{i \in \mathbb{Z}} 2^i \mathbb{1}_{X_i}$ as a linear combination of the indicator functions of Lebesgue sets $X_i \in \mathcal{L}_{I^2}$ and then replace each X_i by some Borel set $Y_i \in \mathcal{B}_{I^2}$ with $\mu(X_i \Delta Y_i) = 0$.) This allows some flexibility in the definitions above. Still, in order to eliminate any ambiguity, we decided to specify the corresponding σ -algebras whenever the measurability of functions may matter.

Also, note that every graphon $W \in \mathcal{W}$, as a bounded measurable function on the finite measure space (I^2, \mathcal{L}, μ) , is integrable (see [15, Section 2.2]), that is, $W \in L^1$.

Finally, let us remark that the standard definition of L^1 allows functions to assume values $\pm\infty$. (This is convenient in the statements of many theorems of real analysis.) Since any integrable function assumes value $\pm\infty$ on a set of measure 0 and we identify a.e. equal functions, we can restrict ourselves in (7) to functions with values in \mathbb{R} only.

For any integrable function $W : (I^2, \mathcal{L}, \mu) \rightarrow (\mathbb{R}, \mathcal{B})$ (in particular, for any graphon W), define its *cut-norm* (also called the *box-norm*, *rectangle-norm*, etc.) by

$$\|W\|_\square = \sup_{S, T \in \mathcal{L}_I} \left| \int_{S \times T} W(x, y) d\mu(x, y) \right|. \tag{9}$$

The *cut-distance* $\delta_\square(U, W)$ between $U, W \in \mathcal{W}$ is the infimum of $\|U - W^\phi\|_\square$ over all $\phi \in \Phi_0$. See [8, Lemma 3.5] for other equivalent ways to define this distance. For any $S, T \in \mathcal{L}_I, \phi \in \Phi_0$, and an integrable function $W : (I^2, \mathcal{L}, \mu) \rightarrow (\mathbb{R}, \mathcal{B})$, we have

$$\int_{S \times T} W(x, y) d\mu(x, y) = \int_{\phi^{-1}(S) \times \phi^{-1}(T)} W^\phi(x, y) d\mu(x, y), \tag{10}$$

which is easiest to see from the definition of the Lebesgue integral by approximating W by simple functions [15, Section 2.2]. It follows that $\|U - W^\phi\|_\square = \|U^{\phi^{-1}} - W\|_\square$ and that δ_\square is a pre-metric on \mathcal{W}_I (see the argument leading to (18)).

For a graphon $W \in \mathcal{W}$ we consider its equivalence class

$$[W] = \{U \in \mathcal{W} : \delta_\square(U, W) = 0\}.$$

Let

$$\mathcal{X} = \{[W] : W \in \mathcal{W}_I\} \tag{11}$$

consist of those equivalence classes that have a representative in \mathcal{W}_I . We call elements of \mathcal{X} *graph limits*. The pre-metric δ_\square induces a metric on \mathcal{X} , which we still denote by the same symbol δ_\square .

Usually, it is more convenient to operate with graphons, understanding equivalence classes implicitly. But here we try to be as explicit as it is reasonably possible. Since the words “graph” and “limit” are frequently used in this paper in various contexts, we will use (in the absence of a better name) the term *graphit* when referring to an equivalence class $[W]$ with $W \in \mathcal{W}_I$. (One might view terms “graphon” and “graphit” as abbreviations of “graph function” and “graph limit”.)

For a graph G on vertices $\{x_1, \dots, x_n\}$, the corresponding element of \mathcal{X} is $A(G) = [W_G]$, where $W_G \in \mathcal{W}_I$ is defined by

$$W_G(x, y) = \begin{cases} 1, & \text{if } (x, y) \in \left[\frac{k-1}{n}, \frac{k}{n} \right) \times \left[\frac{l-1}{n}, \frac{l}{n} \right) \text{ and } \{x_k, x_l\} \in E(G), \\ 0, & \text{for all other } (x, y) \in I^2, \end{cases} \tag{12}$$

that is, we encode the adjacency matrix of G by a function $W_G \in \mathcal{W}_I$. Clearly, the graphit $A(G)$ does not depend on the labeling of $V(G)$ (while the graphon W_G does in general).

We have completely defined the metric space $(\mathcal{X}, \delta_\square)$ and the special points $A(G)$. This determines the promised convergence on graphs. Let us give some brief pointers to the main properties of this construction.

Lovász and Szegedy [29, Theorem 5.1] proved that the metric space $(\mathcal{X}, \delta_\square)$ is compact. Also, they showed [28, Theorem 2.2] that the set $\{[W_G] : G \text{ is a graph}\}$ is dense in $(\mathcal{X}, \delta_\square)$, that is, every graphit $[W]$ with $W \in \mathcal{W}_I$ is a limit of some sequence of graphs.

Any graph sequence G_n with $e(G_n) = o(v(G_n)^2)$ as $n \rightarrow \infty$, converges to the graphit $[\text{Const}(0)]$, where for $\alpha \in I, \text{Const}(\alpha) \in \mathcal{W}_I$ is the constant function that assumes the value α . This is why the phrase “convergence of dense graphs” is often used.

The graphon W_G can be viewed as a version of the adjacency matrix of a graph G . However, a better informal interpretation of a general graphon $W \in \mathcal{W}_I$ is as a continuous version of the matrix that encodes densities between parts of a (weak) regularity partition, see [29, Section 5]. This also hints why, although we start with 0/1-valued functions W_G , we have to allow general real-valued functions when we pass to limits. Having this data for the graph, one can approximate, for example, the value of a max-cut: for graphons the corresponding computation is the supremum of the integral in (9) over disjoint measurable $S, T \subseteq I$.

For graphs F and G the *density* $t(F, G)$ of F in G is the probability that a random (not necessarily injective) map $V(F) \rightarrow V(G)$ induces a homomorphism from F into G .

As it turns out, the subgraph densities behave well with respect to the δ_\square -distance. In combinatorial terms, this says, roughly speaking, that if for two graphs G and H on $[n]$ we have

$$|e(G[A, B]) - e(H[A, B])| = o(n^2), \quad \text{for every } A, B \subseteq [n], \tag{13}$$

then for every fixed graph F we have $|t(F, G) - t(F, H)| = o(1)$. We refer the reader to [28, Lemma 4.1] or [8, Theorems 2.3 and 3.7] for the precise statements and proofs. This may be viewed as a version of the Counting Lemma: if we know the pairwise densities in a regularity partition $V(G) = V_1 \cup \dots \cup V_k$ of a graph G , and generate the corresponding k -partite random graph H on $V(G)$, then as $v(G)$ and k tend to infinity, with high probability (13) holds, and we can approximate subgraph densities in G by those in H . This greatly motivates why the cut-norm is chosen to define the distance on graphons. The role of ϕ in the definition of δ_\square is, in the discrete language, to overlay fractionally the vertex sets of two graphs, cf. (3) here and [8, Section 5.1].

It is natural to define the *density* of a graph F on $[k]$ in a graphit $[W]$ by picking an arbitrary graph sequence (G_n) convergent to $[W]$ and letting

$$t(F, [W]) = \lim_{n \rightarrow \infty} t(F, G_n). \tag{14}$$

This is well-defined and does not depend on the choice of (G_n) . In fact, by writing $t(F, G_n)$ as a k -fold sum and approximating it by a k -fold integral, one can show (see [28, Lemma 4.1] or [8, Theorem 3.7.a]) that

$$t(F, [W]) = \int_{I^k} \prod_{\{i,j\} \in E(F)} W(x_i, x_j) d\mu(x_1, \dots, x_k). \tag{15}$$

Furthermore, neither of these definitions depends on the choice of $W \in [W]$, so we can write $t(F, W)$ in place of $t(F, [W])$. Also, we have $t(F, G) = t(F, W_G)$.

More generally, in terms of graphons, [28, Lemma 4.1] (see also [8, Theorem 3.7.a]) implies that the induced function $t(F, -) : (\mathcal{X}, \delta_\square) \rightarrow I$ is continuous for any F . Thus if $(W_n)_{n \in \mathbb{N}}$ is δ_\square -Cauchy, then the sequence $(t(F, W_n))_{n \in \mathbb{N}}$ of reals is Cauchy for every fixed graph F . The converse of this also holds, by a result of Borgs et al. [8, Theorem 3.7.b]. Thus for $W, W_1, W_2, \dots \in \mathcal{W}_I$,

$$\lim_{n \rightarrow \infty} \delta_\square(W_n, W) = 0 \quad \text{if and only if} \quad \forall \text{ graph } F \quad \lim_{n \rightarrow \infty} t(F, W_n) = t(F, W). \tag{16}$$

It follows that each graphit $[W]$ is uniquely determined by its “moments function” $t(-, W)$. An algebraic characterization of all possible functions $t(-, W)$ realizable by some $W \in \mathcal{W}_I$ is given by Lovász and Szegedy [28, Theorem 2.2].

Let us also say a few words about graph limits and property testing. (See [19] for a precise definition of property testing and several fundamental results.) In the most restrictive sense (the *oblivious* or *order independent testing*), we have a (very big) unknown graph G and are told the subgraph $G[X]$ induced by a random m -set X of vertices, where m is a fixed number. It is known that with probability at least $1 - \varepsilon$ we have $\delta_\square(W_{G[X]}, W_G) \leq \varepsilon$, provided $m \geq m_0(\varepsilon)$ (see [28, Theorem 2.5] or [8, Theorem 3.7]). This means that we can learn a good δ_\square -approximation to the graph G . The objective of *property testing* is to approximate with high probability how far G is from a given property \mathcal{P} , but the *edit* distance $\hat{\delta}_1$ is to be used here. Graphons seem to provide very convenient tools and language for dealing with this problem (which essentially amounts to relating the $\hat{\delta}_1$ and δ_\square distances from an arbitrary graph to the given property), see [5,30].

3. Extending the δ_1 -distance to graph limits

Here we show how to extend the distance δ_1 from graphs to graphits. This definition is standard but it seems that no formal proofs of some of its properties have appeared in the literature. Therefore we give careful proofs of all claims (or references to them). The author thanks László Lovász for pointing out that Lemma 11 can be deduced from the results in [4, 30], which is the proof presented here.

Here is the definition of δ_1 for graphits. First, we define δ_1 on \mathcal{W} , the set of graphons. For $U, W \in \mathcal{W}$, let

$$\delta_1(U, W) = \inf \{ \|U - W^\phi\|_1 : \phi \in \Phi_0 \}, \tag{17}$$

where $\|U - W^\phi\|_1$ is the standard ℓ_1 -norm of $U - W^\phi$ as defined by (8).

Clearly, δ_1 is non-negative. It is symmetric by (10). Also, δ_1 satisfies the Triangle Inequality. Indeed, for every $U, V, W \in \mathcal{W}$ and $\varepsilon > 0$ we can choose $\phi, \psi \in \Phi_0$ such that $\|U^\phi - V\|_1 \leq \delta_1(U, V) + \varepsilon$ and $\|V - W^\psi\|_1 \leq \delta_1(V, W) + \varepsilon$. Now, by the Triangle Inequality for the ℓ_1 -norm,

$$\begin{aligned} \delta_1(U, W) &\leq \|U - (W^\psi)^\phi\|_1 = \|U^\phi - W^\psi\|_1 \\ &\leq \|U^\phi - V\|_1 + \|V - W^\psi\|_1 \leq \delta_1(U, V) + \delta_1(V, W) + 2\varepsilon. \end{aligned} \tag{18}$$

Since $\varepsilon > 0$ was arbitrary, the claim follows. Hence, δ_1 is a pre-metric on \mathcal{W}_I .

We will present an equivalent definition of δ_1 in Lemma 9 and will conclude in Corollary 12 that δ_1 gives a metric on \mathcal{X} . Let us state a few auxiliary or related results first.

Lemma 2. *Let an integrable $W : (I^2, \mathcal{L}, \mu) \rightarrow (\mathbb{R}, \mathcal{B})$ satisfy $\|W\|_\square = 0$. Then $W = 0$ a.e. In particular, for any $U, W \in \mathcal{W}$, $\|U - W\|_\square = 0$ implies that $\|U - W\|_1 = 0$.*

Proof. Let Z be the Lebesgue set of the function W , which can be defined as the set of those (x, y) in the interior of I^2 such that

$$\lim_{\substack{c \rightarrow 0 \\ c > 0}} \frac{1}{\mu(R_{x,y,c})} \int_{(x',y') \in R_{x,y,c}} |W(x', y') - W(x, y)| d\mu(x', y') = 0, \tag{19}$$

where $R_{x,y,c}$ is the open rectangle $(x - c, x + c) \times (y - c, y + c)$.

The Lebesgue Differentiation Theorem ([15, Theorem 3.21]) implies that $\mu(Z) = 1$. If $W(x, y) \neq 0$ for some $(x, y) \in Z$, then by (19) there is $c > 0$ such that

$$\left| W(x, y) - \frac{1}{4c^2} \int_{R_{x,y,c}} W \right| < \frac{|W(x, y)|}{2}.$$

Thus $\|W\|_\square \geq \int_{R_{x,y,c}} |W| \geq 2c^2 |W(x, y)| > 0$, a contradiction. Thus $W = 0$ a.e. \square

A function $U : I^2 \rightarrow \mathbb{R}$ is called an *interval step function* if there is a partition $I = I_1 \cup \dots \cup I_k$ into finitely many intervals such that U is constant on each rectangle $I_i \times I_j$. Any interval step function is a simple function. Of course, such U is necessarily measurable, even in the strongest sense as a function from (I^2, \mathcal{B}) to $(\mathbb{R}, 2^\mathbb{R})$.

Lemma 3. *For any $\varepsilon > 0$ and any integrable function $W : (I^2, \mathcal{L}, \mu) \rightarrow (\mathbb{R}, \mathcal{B})$ there is an interval step function U such that $\|W - U\|_1 < \varepsilon$. Moreover, if $W \in \mathcal{W}_I$, then we can also require that $U \in \mathcal{W}_I$.*

Proof. The first part of the lemma follows from [15, Theorem 2.41] (see also [8, Lemma 3.2]). Let us establish the second part. Let $W \in \mathcal{W}_I$ and U_0 be the interval step function with $\|W - U_0\|_1 < \varepsilon$, given by the first part. Let $U_1(x, y) = g(U_0(x, y))$, where $g(z) = \max(0, \min(1, z))$ maps $z \in \mathbb{R}$ to the nearest point from I . Since for every $z' \in I$ and $z \in \mathbb{R}$ we have $|g(z) - z'| \leq |z - z'|$, we conclude that $\|U_1 - W\|_1 \leq \|U_0 - W\|_1 \leq \varepsilon$. Finally, we take $U(x, y) = (U_1(x, y) + U_1(y, x))/2$. Then the new interval step function U belongs to \mathcal{W}_I . Also, in view of inequality $|a - c| + |b - c| \geq 2|\frac{a+b}{2} - c|$ valid for any $a, b, c \in \mathbb{R}$, we have $\|W - U\|_1 \leq \|W - U_1\|_1 \leq \varepsilon$, as desired. \square

Remark. This approximation reminds the one given by the Weak Regularity lemma of Frieze and Kannan [17] (see also [29, Section 2]) with respect to the cut-norm, except we cannot bound the number of parts in Lemma 3 in terms of ε only. This is an important distinction between the cut-norm and the ℓ_1 -norm, giving another motivation for taking δ_\square as the distance between graphons. This allows one to construct a finite ε -net for the metric space $(\mathcal{X}, \delta_\square)$. Namely, let $n = n(\varepsilon)$ be large and take all interval steps functions with steps $[\frac{i}{n}, \frac{i+1}{n})$ that assume values in $\{\frac{1}{n}, \dots, \frac{n}{n}\}$; there are at most $n^{n^2} < \infty$ such functions. Thus $(\mathcal{X}, \delta_\square)$ is totally bounded, which is one of the ingredients needed for compactness. See [29, Theorem 5.1] for more details.

Lemma 4. *Let $X, Y \in \mathcal{L}_I$ have measure 1 and let ψ be a bijection from X onto Y such that for any interval $J \subseteq I$ the sets $\psi(J \cap X)$ and $\psi^{-1}(J \cap Y)$ are Lebesgue measurable with $\mu(\psi(J \cap X)) = \mu(\psi^{-1}(J \cap Y)) = \mu(J)$. Then there is $\phi \in \Phi_0$ such that $\phi = \psi$ a.e.*

Proof. Suppose first that $|I \setminus X| = |I \setminus Y| = c$, that is, the cardinality of both $I \setminus X$ and $I \setminus Y$ is continuum. Let ϕ be an arbitrary bijection between $I \setminus X$ and $I \setminus Y$ while $\phi(x) = \psi(x)$ if $x \in X$. Then $\phi = \psi$ a.e. Also, for any interval $J \subseteq I$, the pre-image $\phi^{-1}(J)$ differs from $\psi^{-1}(J \cap X) \in \mathcal{L}$ on a set of measure 0, so it is Lebesgue measurable of measure $\mu(J)$. Since \mathcal{B} is generated by intervals as a σ -algebra ([15, Theorem 1.6]), it follows (e.g. by application of the uniqueness claim of [15, Theorem 1.14]) that ϕ is a measure preserving function from (I, \mathcal{L}) to (I, \mathcal{B}) . But a subset of I is Lebesgue measurable set if and only if it can be sandwiched between two Borel sets of the same measure ([15, Theorem 1.19]). This easily implies that ϕ is a measure preserving map from (I, \mathcal{L}) to (I, \mathcal{L}) , that is, $\phi \in \Phi$. Likewise, $\phi^{-1} \in \Phi$, giving $\phi \in \Phi_0$ as required.

Finally, suppose that, for example, $|I \setminus X| < \epsilon$. Let $C \subseteq I$ be the Cantor set, which has measure 0 and cardinality continuum [15, Proposition 1.22]. Let $X' = X \setminus C$ and $Y' = Y \setminus \psi(X \cap C)$. Then ψ maps X' bijectively onto Y' . Also, $\mu(\psi(X \cap C)) = 0$. Indeed, for every $\epsilon > 0$, we can find a set $J \supseteq C$ which is the union of finitely many intervals of total length at most ϵ that covers C . By the assumption of the lemma, $\psi(X \cap J)$ has measure at most ϵ . Since $\epsilon > 0$ was arbitrary, $\mu(\psi(X \cap C)) = 0$. Thus $\mu(X \setminus X') = \mu(Y \setminus Y') = 0$ and the restriction $\psi|_{X'}$ satisfies the assumptions of the lemma. Since $|I \setminus X'| = |I \setminus Y'| = \epsilon$, we already know how to find the required $\phi \in \Phi_0$ for $\psi|_{X'}$. The very same function ϕ works for ψ as well. \square

Let us call a point x lying inside a Lebesgue set $A \subseteq \mathbb{R}$ a *density point* of A if

$$\lim_{\substack{c \rightarrow 0 \\ c > 0}} \frac{\mu(A \cap (x - c, x + c))}{2c} = 1,$$

or equivalently, if x belongs to the Lebesgue set (as defined by the 1-dimensional version of (19)) of the characteristic function $\mathbb{I}_A : \mathbb{R} \rightarrow \{0, 1\}$ of A . Again, Theorem 3.21 in [15] implies that almost every point of A is a density point.

The arithmetic operations and the linear order on $I = [0, 1]$ play no role in the definition of graphons; see [4, Section 2.1] for a more general point of view. The following simple lemma suffices for our purposes.

Lemma 5. *For every partition of $I = A_1 \cup \dots \cup A_k$ into Lebesgue measurable sets A_i there are a partition $I = I_1 \cup \dots \cup I_k$ into intervals and $\psi \in \Phi_0$ such that $\mu(\psi(A_i) \Delta I_i) = 0$ for each $i \in [k]$.*

Proof. It is enough to prove the case $k = 2$ with the general claim following by a simple induction on k .

Let $a_1 = 0$, $a_2 = \mu(A_1)$, $I_1 = [0, a_2]$, and $I_2 = I \setminus I_1$. Assume that $0 < a_2 < 1$ (for otherwise any $\psi \in \Phi_0$ works).

Let $i = 1$ or 2 . Let $X_i \subseteq A_i$ be the set of density points of A_i . For $x \in X_i$ let

$$\psi(x) = a_i + \mu(A_i \cap [0, x]).$$

Then $\psi(X_i)$ lies in the interior of I_i . Indeed, if, for example, $\psi(x) = a_i$, then $\mu(A_i \cap (-\infty, x)) = 0$, so x cannot be a density point for A_i . Likewise, if $y \in X_i \setminus \{x\}$ is another density point of A_i , then $\psi(y) \neq \psi(x)$. Let $Y_i = \psi(X_i)$. The pre-image under ψ of any open interval $J = (a_i, a_i + b) \subseteq I_i$ is the intersection of the interval $(0, c)$ with X_i , where

$$c = \sup\{x \in I : \mu(A_i \cap [0, x]) < b\} = \sup\{x \in I : \mu(X_i \cap [0, x]) < b\}.$$

Since $b \leq \mu(X_i)$ and the measure μ is continuous from below ([15, Theorem 1.8.c]), we conclude that $\mu(\psi^{-1}(J)) = b = \mu(J)$. Also, for any open interval $J = (b, c) \subseteq I$, the image under ψ of $X_i \cap J$ is $Y_i \cap J_i$, where

$$J_i = (a_i + \mu(A_i \cap [0, b]), a_i + \mu(A_i \cap [0, c]))$$

is a subinterval of I_i with $\mu(J_i) = \mu(J \cap X_i)$.

Let $X = X_1 \cup X_2$ and $Y = Y_1 \cup Y_2$. It routinely follows that all assumptions of Lemma 4 with respect to the bijection $\psi : X \rightarrow Y$ are satisfied. The element $\phi \in \Phi_0$ returned by Lemma 4 has the required properties. \square

Lemma 6. *For every interval step function $U \in \mathcal{W}$ and $\phi \in \Phi$, there is $\psi \in \Phi_0$ such that $(U^\phi)^\psi = U$ a.e.*

Proof. Let $I = I_1 \cup \dots \cup I_k$ be a partition into intervals such that U is constant on each rectangle $I_i \times I_j$. For $i, j \in [k]$, let $\alpha_{i,j} = \mu(A_{i,j})$, where $A_{i,j} = I_j \cap \phi^{-1}(I_i)$. Since ϕ is measure preserving, $\sum_{j=1}^k \alpha_{i,j} = \mu(I_i)$ for every $i \in [k]$. Partition the interval $I_i = I_{i,1} \cup \dots \cup I_{i,k}$ into intervals of lengths respectively $\alpha_{i,1}, \dots, \alpha_{i,k}$. By Lemma 5 find $\eta \in \Phi_0$ such that $\mu(\eta(A_{i,j}) \Delta I_{i,j}) = 0$. The element $\psi = \eta^{-1} \in \Phi_0$ has the required properties by (6) because for a.e. $x \in I_{i,j}$ we have $\psi(x) \in A_{i,j}$ and $\phi(\psi(x)) \in I_i$. \square

Lemmas 3 and 6 easily imply the following result.

Corollary 7. *For any $U, W \in \mathcal{W}$ and $\phi \in \Phi$, we have $\delta_1(U, W) = \delta_1(U^\phi, W)$. \square*

Theorem 8. *For $U, W \in \mathcal{W}_1$, the following are equivalent.*

- (a) *For every graph F , we have $t(F, U) = t(F, W)$.*
- (b) $\delta_{\square}(U, W) = 0$.
- (c) *There are $\phi, \psi \in \Phi$ such that $U^\phi = W^\psi$ a.e.*

Proof. The equivalence of (a) and (b) follows from (16) (i.e. from [28, Lemma 4.1] and [8, Theorem 3.7]). The equivalence of (a) and (c) is proved by Borgs et al. [4, Corollary 2.2]. \square

Lemma 9. *For any $U, W \in \mathcal{W}$, we have*

$$\delta_1(U, W) = \inf_{\phi, \psi \in \Phi} \|U^\phi - W^\psi\|_1. \tag{20}$$

Proof. Since Φ_0 is a subset of Φ and Φ_0 contains the identity function $\text{Id} : I \rightarrow I$, the “ \geq ”-inequality in (20) easily follows. Let us show the converse.

Let $U, W \in \mathcal{W}$ and $\varepsilon > 0$. By Lemma 3 we can find interval step functions U_0 and W_0 lying within ε from respectively U and W in the ℓ_1 -norm. For any $\phi, \psi \in \Phi$, we have by (10)

$$\|U^\phi - W^\psi\|_1 \geq \|U_0^\phi - W_0^\psi\|_1 - \|U^\phi - U_0^\phi\|_1 - \|W^\psi - W_0^\psi\|_1 \geq \|U_0^\phi - W_0^\psi\|_1 - 2\varepsilon.$$

Likewise, $\|U - W^\phi\|_1 \leq \|U - W_0^\phi\|_1 + 2\varepsilon$. Since $\varepsilon > 0$ was arbitrary, it is enough to prove (20) on the additional assumption that U and W are interval step functions.

Again, let $\varepsilon > 0$. Let $\phi, \psi \in \Phi$ be such that $\|U^\phi - W^\psi\|_1 - \varepsilon$ is at most the right-hand side of (20). By Lemma 6 choose $\eta \in \Phi_0$ such that $(W^\psi)^\eta = W$ a.e. Then, by (6),

$$\|U^\phi - W^\psi\|_1 = \|(U^\phi)^\eta - (W^\psi)^\eta\|_1 = \|U^{(\phi \circ \eta)} - W\|_1. \tag{21}$$

Again, by Lemma 6 applied to U and $\phi \circ \eta \in \Phi$, find $\nu \in \Phi_0$ such that $(U^{(\phi \circ \eta)})^\nu = U$ a.e. From (21) we conclude that $\|U^\phi - W^\psi\|_1 = \|U - W^\nu\|_1$, which is at least the right-hand side of (17). Since ε was arbitrary, the lemma follows. \square

Lemma 10. For any two graphs G and H , the δ_1 -distance $\delta_1(G, H)$ defined by (3) is equal to $\delta_1(W_G, W_H)$, where W_G and W_H are defined by (12).

Proof. Let $V(G) = \{x_1, \dots, x_m\}$ and $V(H) = \{y_1, \dots, y_n\}$. For $\phi \in \Phi_0$, $\|W_G^\phi - W_H\|_1$ is equal to the expression in (3) with $\alpha_{i,j} = \mu(I_i \cap \phi^{-1}(J_j))$, $I_i = (\frac{i-1}{m}, \frac{i}{m})$ and $J_j = (\frac{j-1}{n}, \frac{j}{n})$. Conversely, given numbers $\alpha_{i,j}$ such the matrix $(\alpha_{i,j})_{i,j \in [n]}$ has row sums $1/m$ and column sums $1/n$, one can easily construct $\phi \in \Phi_0$ giving these $\alpha_{i,j}$ as above. \square

Lemma 11. Let $U, W \in \mathcal{W}$ satisfy $\delta_\square(U, W) = 0$. Then $\delta_1(U, W) = 0$.

Proof. By Theorem 8, there are $\phi, \psi \in \Phi$ such that $U^\phi = W^\psi$ a.e. The claim follows by using the equivalent definition of δ_1 from Lemma 9. \square

Corollary 12. The function δ_1 induces a metric on the set \mathcal{X} of graphits, extending the δ_1 -distance from graphs. \square

Remark. Let us point out that the convergence with respect to the cut-distance does not generally imply the convergence with respect to δ_1 . For example, the infinite sequence of random graphs $G_n \in G_{n,1/2}$ converges in the δ_\square -distance with probability 1 to the graphit $[\text{Const}(1/2)]$ by [28, Corollary 2.6] while no graph sequence whatsoever can converge in the δ_1 -distance to $[\text{Const}(1/2)]$ by Theorem 17 here.

4. Comparing the discrete and fractional δ_1 -distances

Clearly, for graphs G and H of the same order we have $\hat{\delta}_1(G, H) \geq \delta_1(G, H)$, where $\hat{\delta}_1$ is defined by (2). The distances $\hat{\delta}_1$ and δ_1 do not coincide in general as Example 13 demonstrates. Independently, Arie Matsliah (see [20, Appendix B]) presented another construction that achieves ratio $6/5$. Although our ratio is smaller (only $11/10$), the ideas behind our construction are different from those of Matsliah and might be useful in the quest for better ratios. Hence, we decided to keep this example in the paper.

Example 13. There are graphs G and H such that $v(G) = v(H)$ but

$$\hat{\delta}_1(G, H) \geq \frac{11}{10} \delta_1(G, H) > 0.$$

Proof. Fix an integer $n \geq 24$. Pick disjoint sets $X = \{x_1, \dots, x_4\}$, $M = M_1 \cup \dots \cup M_4$, and $N = N_1 \cup \dots \cup N_5$ with each M_i having 4 elements and each N_i having n elements.

Let $V(G) = V(H) = N \cup M \cup X$. It will be the case that $N \cup M$ spans the same subgraph in both G and H . Namely, N spans the complete graph while, for $i \in [4]$, we put the complete bipartite graph between M_i and $\cup_{j=1}^i N_j$. These are all edges inside $M \cup N$.

Fix another partition $M = L_1 \cup \dots \cup L_4$ such that each L_i has 4 elements and $|L_i \cap M_i| = |L_{i+1} \cap M_i| = 2$ for $i \in [4]$, where we agree that $L_5 = L_1$.

In G , the edges incident to X are as follows: $\{x_i, x_j\}$ for $1 \leq i < j \leq 4$ with $j - i$ even plus all pairs $\{x_i, y\}$ for $i \in [4]$ and $y \in M_i$. In H , the edges incident to X are as follows: $\{x_i, x_j\}$ for $1 \leq i < j \leq 4$ with $j - i$ odd plus all pairs $\{x_i, y\}$ for $i \in [4]$ and $y \in L_i$.

We have

$$|E(G) \Delta E(H)| = \sum_{i=1}^4 |M_i \Delta L_i| + \binom{|X|}{2} = 22. \tag{22}$$

Let us show that this is smallest possible. Pick an optimal bijection $\sigma : V(G) \rightarrow V(H)$. In each of G and H , every vertex in N has degree at least $5n - 1$ while any vertex in $M \cup X$ has degree at most $4n + 1$. Hence, if σ does not preserve N , then the number of discrepancies will be at least $(5n - 1) - (4n + 1) \geq 22$. So, assume that $\sigma(N) = N$. Likewise, we have

$\sigma(M_i) = M_i$, for otherwise the number of discrepancies (between M and N) is at least $n > 22$. Finally, consider the action of σ on X . For every $x, y \in X$, their neighborhoods in M with respect to G and H differ by at least 4. If σ does not map some x_i into $\{x_i, x_{i+1}\}$, where $x_5 = x_1$, then the neighborhoods $N_G(x_i)$ and $N_H(\sigma(x_i))$ in M are disjoint and this vertex alone creates at least 8 discrepancies. Moreover, since X spans 2 and 4 edges in G and H respectively, the total number of discrepancies is at least $8 + 3 \times 4 + 2 = 22$ and we cannot improve (22). Thus let us assume that $\sigma(x_i) \in \{x_i, x_{i+1}\}$ for every $i \in [4]$. This implies that either σ is constant on X or shifts indices by 1. In either case, this gives the same bound as in (22).

Hence, $\hat{\delta}_1(G, H) \geq \frac{2 \cdot 22}{(5n+20)^2}$. Let us establish an upper bound on $\delta_1(G, H)$ now.

Let $G[2]$ be the 2-fold blow-up of G , where each vertex x is replaced by two vertices x', x'' and each edge $\{x, y\}$ by the complete bipartite graph with parts $\{x', x''\}$ and $\{y', y''\}$. For $Y \subseteq V(G)$, let $Y[2] = \{y', y'' : y \in Y\}$. Consider the following bijection σ between the vertex sets of $G[2]$ and $H[2]$. It is the identity bijection on $M[2] \cup N[2]$. For $i \in [4]$, let $\sigma(x'_i) = x'_i$ and $\sigma(x''_i) = x''_{i+1}$. Easy checking shows that σ , when restricted to $X[2]$, mismatches only 16 adjacencies (versus $4 \times \binom{4}{2} = 24$ if σ were the identity). The number of discrepancies between $X[2]$ and $M[2]$ is 4×16 . We have

$$\delta_1(G, H) \leq \hat{\delta}_1(G[2], H[2]) \leq \frac{2}{4(5n+20)^2} (4 \times 16 + 16) \leq \frac{10}{11} \hat{\delta}_1(G, H). \quad \square$$

Lemma 14. For any two graphs G and H on the same vertex set $[n]$, we have

$$\hat{\delta}_1(G, H) \leq 3\delta_1(G, H).$$

Proof. If $G \cong H$, then $\delta_1(G, H) = \hat{\delta}_1(G, H) = 0$, so assume $G \not\cong H$. Let $\ell = n^2 \hat{\delta}_1(G, H)/2$ be the smallest number of adjacencies we have to change in G to make it isomorphic to H .

Let $A = (\alpha_{i,j})_{i,j \in [n]}$ be an optimal overlay matrix as in (3), where we assume $x_i = i$ and $y_j = j$. (Thus nA is doubly-stochastic.)

Although nA can be represented as a convex combination of permutation matrices by Birkhoff's theorem [3], we find it more convenient to work with an approximation where all coefficients are equal. (Thus some permutation matrices may be repeated more than once.) Such an approximation is easy to find as follows.

Pick a large $m > m_0(A)$. Inductively on i , we construct permutation matrices P_i as follows. Suppose that $i \geq 0$ and we have already found P_1, \dots, P_i such that $P' = P_1 + \dots + P_i \leq mnA$ (where matrix inequalities are meant component-wise). If there is a permutation matrix P_{i+1} such that $P' + P_{i+1} \leq mnA$, take it and repeat the step.

Suppose that no such P_{i+1} exists. Let $B = (\beta_{f,g})_{f,g \in [n]} = mnA - P'$. This is a non-negative matrix with row/column sums $m - i$. By Hall's Marriage theorem [22], there is a set $R \subseteq [n]$ of r rows and a set $S \subseteq [n]$ of $n - r + 1$ columns such that each entry of the $R \times S$ -submatrix of B is less than 1. Hence,

$$\begin{aligned} (m - i)r &= \sum_{f \in R} \sum_{g=1}^n \beta_{f,g} = \sum_{f \in R} \sum_{g \in S} \beta_{f,g} + \sum_{f \in R} \sum_{g \in [n] \setminus S} \beta_{f,g} \\ &\leq r(n + 1 - r) + (m - i)(n - (n - r + 1)), \end{aligned}$$

and therefore $m - i \leq r(n + 1 - r) \leq (n + 1)^2/4$. Let P_{i+1}, \dots, P_m be arbitrary permutation matrices and $P = \frac{1}{mn}(P_1 + \dots + P_m)$. It follows that

$$\|A - P\|_\infty \leq 2 \times \frac{(n + 1)^2}{4mn} = \frac{(n + 1)^2}{2mn}.$$

Since m is arbitrarily large, in order to prove the lemma it is enough to show that

$$\hat{\delta}_1(G, H) \leq 3\delta_1(G, H, P), \tag{23}$$

where $\delta_1(G, H, P)$ is defined by (3).

Let $\sigma_1, \dots, \sigma_m : [n] \rightarrow [n]$ be the permutations encoded by P_1, \dots, P_m respectively. As it was defined after (3), Δ is the set of all quadruples $(x, y, x', y') \in [n]^4$ such that exactly one of the relations $\{x, y\} \in E(G)$ and $\{x', y'\} \in E(H)$ holds. Note that we allow $x = y$ or $x' = y'$ but both equalities cannot hold simultaneously by the definition of Δ .

For $(i, j) \in [m]^2$, let $\Delta(i, j)$ consist of $(x, y) \in [n]^2$ such that $(x, y, \sigma_i(x), \sigma_j(y)) \in \Delta$. For $(x, y, x', y') \in \Delta$, let $I(x, y, x', y')$ consist of all pairs $(i, j) \in [m]^2$ such that $\sigma_i(x) = x'$ and $\sigma_j(y) = y'$. Also, for $X \subseteq [m]$, define

$$S_X = \sum_{i,j \in X: i < j} |\Delta(i, j)|.$$

We have

$$\delta_1(G, H, P) = \sum_{(x,y,x',y') \in \Delta} P_{x,x'} P_{y,y'}$$

$$\begin{aligned}
 &= \sum_{(x,y,x',y') \in \Delta} \left(\sum_{i:\sigma_i(x)=x'} \frac{1}{mn} \right) \left(\sum_{j:\sigma_j(y)=y'} \frac{1}{mn} \right) \\
 &= \frac{1}{m^2 n^2} \sum_{(x,y,x',y') \in \Delta} |I(x,y,x',y')| \\
 &= \frac{1}{m^2 n^2} \sum_{i,j \in [m]} |\Delta(i,j)| = \frac{2S_{[m]} + \sum_{i=1}^m |\Delta(i,i)|}{m^2 n^2}. \tag{24}
 \end{aligned}$$

Let us show that for any $1 \leq g < i < j \leq m$ we have

$$|\Delta(g,i)| + |\Delta(j,i)| + |\Delta(j,g)| \geq |\Delta(g,g)|. \tag{25}$$

Start with any $(x,y) \in \Delta(g,g)$. Let us transform (x,y) into $(\sigma_g(x), \sigma_g(y))$ in three steps, where we consecutively apply (σ_g, σ_i) , $(\sigma_j^{-1}, \sigma_i^{-1})$, and (σ_j, σ_g) :

$$(x,y) \rightarrow (\sigma_g(x), \sigma_i(y)) \rightarrow (\sigma_j^{-1}(\sigma_g(x)), y) \rightarrow (\sigma_g(x), \sigma_g(y)).$$

Since $(x,y, \sigma_g(x), \sigma_g(y)) \in \Delta$, at least one of these three steps changes adjacency. Depending on the number of the step when this happens, we get respectively that $(x,y) \in \Delta(g,i)$, $(\sigma_j^{-1}(\sigma_g(x)), y) \in \Delta(j,i)$, or $(\sigma_j^{-1}(\sigma_g(x)), y) \in \Delta(j,g)$. Conversely, suppose that we are given the resulting conclusion of the form $(u,v) \in \Delta(a,b)$ with distinct $a, b \in \{i,j,g\}$. The pair (a,b) determines the number $k \in \{1, 2, 3\}$ of the step. This k , when combined with (u,v) , easily allows us to reconstruct the ordered pair (x,y) . Thus no element in the left-hand side of (25) is doubly counted. This proves (25).

By (25) (and $|\Delta(a,b)| = |\Delta(b,a)|$) we conclude that $S_{\{g,i,j\}} \geq |\Delta(g,g)| \geq 2\ell$. A simple averaging over all choices of $\{i, g, h\} \in \binom{[m]}{3}$ implies that $S_{[m]} \geq 2\ell \binom{m}{2} / \binom{3}{2} = \ell m(m-1)/3$. By (24), we have

$$\delta_1(G, H, P) \geq \frac{2\ell m(m-1)/3 + 2\ell m}{m^2 n^2} \geq \frac{2\ell}{3n^2} = \frac{\hat{\delta}_1(G, H)}{3},$$

finishing the proof of Lemma 14. \square

Remark. The author thanks Alexander Razborov for the remarks that simplified the original proof of Lemma 14.

The interesting problem of finding the best possible constant in Lemma 14 remains open. At the moment, we know only that it is between $6/5$ (see [20, Appendix B]) and 3 .

The situation for the cut-distance is somewhat similar: the discrete version $\hat{\delta}_\square$ of δ_\square , as defined by [8, Equation (2.6)], is not always equal to the δ_\square -distance ([8, Section 5.1]) while for any two graphs G and H of the same order we have

$$\delta_\square(G, H) \leq \hat{\delta}_\square(G, H) \leq 32(\delta_\square(G, H))^{1/67}$$

([8, Theorem 2.3]). It is open whether $\hat{\delta}_\square(G, H)$ can be bounded from above by a linear function of $\delta_\square(G, H)$, see e.g. [8, page 1830].

5. Characterization of stability

Recall that in the Introduction we defined when an extremal (f, \mathcal{P}) -problem is stable. Here we give an alternative characterization. Since stability deals with relating the δ_1 and δ_\square distances, it is not surprising that the methods developed by Lovász and Szegedy [30] in the context of property testing apply here.

Theorem 15. *Let \mathcal{P} be an arbitrary graph property with $\mathcal{P}_n \neq \emptyset$ for infinitely many n and let f be a graph parameter. Then the extremal (f, \mathcal{P}) -problem is stable if and only if $\text{LIM}(f, \mathcal{P})$ consists of a single graphit $[W]$, where moreover $W \in \mathcal{W}_1$ can be chosen to assume values 0 and 1 only.*

The rest of this section is dedicated to proving Theorem 15, in the course of which we observe an interesting dichotomy result (Theorem 17).

We will need the following result, which is a special case of [30, Lemma 2.2].

Lemma 16. *Let $W, W_1, W_2, \dots \in \mathcal{W}$ be such that $\|W_n - W\|_\square \rightarrow 0$ as $n \rightarrow \infty$. Let $S \in \mathcal{L}_2$. Then $\int_S W_n d\mu \rightarrow \int_S W d\mu$ as $n \rightarrow \infty$.*

Sketch of Proof. If S is a rectangle, then the conclusion follows from the definition of the cut-norm. A general $S \in \mathcal{L}_2$ can be approximated within any $\varepsilon > 0$ by a finite union of disjoint rectangles, cf. Lemma 3. \square

Theorem 17. *Let $W \in \mathcal{W}_1$ and let $W_1, W_2, \dots \in \mathcal{W}_1$ be an arbitrary sequence such that $\delta_\square(W_n, W) \rightarrow 0$ as $n \rightarrow \infty$.*

If $\mu(W^{-1}(\{0, 1\})) = 1$ (that is, W assumes only values 0 and 1 a.e.), then the sequence $(W_n)_{n \in \mathbb{N}}$ is necessarily convergent to W in the δ_1 -distance.

If $\mu(W^{-1}(\{0, 1\})) < 1$ and each W_n is a.e. $\{0, 1\}$ -valued, then the sequence $(W_n)_{n \in \mathbb{N}}$ does not contain any Cauchy subsequence with respect to the δ_1 -distance.

Proof. Suppose first that W is $\{0, 1\}$ -valued a.e. Let $S = W^{-1}(0) \in \mathcal{L}_{I^2}$. For each $n \in \mathbb{N}$ choose $\phi_n \in \Phi_0$ such that $\|W_n^{\phi_n} - W\|_{\square} \leq \delta_{\square}(W_n, W) + 1/n$. Clearly, $\|W_n^{\phi_n} - W\|_{\square}$ tends to 0, so by Lemma 16 we have

$$\begin{aligned} \delta_1(W_n, W) &\leq \|W_n^{\phi_n} - W\|_1 = \int_S W_n^{\phi_n} d\mu + \int_{I^2 \setminus S} (1 - W_n^{\phi_n}) d\mu \\ &\rightarrow \int_S W d\mu + \int_{I^2 \setminus S} (1 - W) d\mu = 0. \end{aligned}$$

Now, suppose that $\mu(W^{-1}(\{0, 1\})) < 1$ and that the second part of the theorem is false. By choosing a subsequence and relabeling, we can assume that $(W_n)_{n \in \mathbb{N}}$ itself is a Cauchy sequence with $\delta_1(W_m, W_n) \leq 1/2^m$ for every $m \leq n$. Let $\phi_1 : I \rightarrow I$ be the identity map and $U_1 = W_1$. Inductively on $n = 2, 3, \dots$, do the following. By induction, we assume that we have $U_{n-1} = W_{n-1}^{\phi_{n-1}}$ with $\phi_{n-1} \in \Phi_0$. By Corollary 7,

$$\delta_1(U_{n-1}, W_n) = \delta_1(W_{n-1}^{\phi_{n-1}}, W_n) = \delta_1(W_{n-1}, W_n) \leq \frac{1}{2^{n-1}}.$$

Thus there is $\phi_n \in \Phi_0$ such that, letting $U_n = W_n^{\phi_n}$, we have

$$\|U_{n-1} - U_n\|_1 \leq \frac{1}{2^{n-2}}. \tag{26}$$

The sequence $(U_n)_{n \in \mathbb{N}}$ is Cauchy with respect to the ℓ_1 -norm: for $m \leq n$ we have

$$\|U_n - U_m\|_1 \leq \sum_{i=m+1}^n \|U_i - U_{i-1}\|_1 \leq \sum_{i=m+1}^n \frac{1}{2^{i-2}} < \frac{1}{2^{m-2}}.$$

Since the normed space L^1 defined by (7) is complete ([15, Theorem 6.6]), the sequence $(U_n)_{n \in \mathbb{N}}$ has a limit $U \in L^1$:

$$\lim_{n \rightarrow \infty} \|U_n - U\|_1 = 0. \tag{27}$$

We have $\int_{I^2} |U(x, y) - U(y, x)| d\mu(x, y) = 0$ because it is at most

$$2\|U - U_n\|_1 + \int_{I^2} |U_n(x, y) - U_n(y, x)| d\mu(x, y) = 2\|U - U_n\|_1 \rightarrow 0.$$

Thus U is symmetric a.e. on I^2 by e.g. [15, Proposition 2.16]. Likewise, $0 \leq U(x, y) \leq 1$ a.e. By changing U on a subset of I^2 of measure zero, we can assume that $U \in \mathcal{W}_I$. By the Triangle Inequality,

$$\delta_{\square}(U, W) \leq \delta_{\square}(U, U_n) + \delta_{\square}(U_n, W) \leq \delta_1(U, U_n) + \delta_{\square}(W_n^{\phi_n}, W).$$

This tends to 0 as $n \rightarrow \infty$. Thus $\delta_{\square}(U, W) = 0$ and by Theorem 8, $U^{\psi} = W^{\phi}$ a.e. for some $\psi, \phi \in \Phi$. Thus U is not $\{0, 1\}$ -valued a.e.

For $m \in \mathbb{N}$, let

$$A_m = \{(x, y) \in I^2 : 1/m < U(x, y) < 1 - 1/m\}.$$

Each A_m is Lebesgue measurable since U is measurable. Also, $Z = \cup_{m \in \mathbb{N}} A_m = \{z \in I^2 : U(z) \notin \{0, 1\}\}$ has positive measure c . By the continuity from below [15, Theorem 1.8.c] of the measure μ , there is $m \in \mathbb{N}$ with $\mu(A_m) > c/2$. Since each $U_n = W_n^{\phi_n}$ is $\{0, 1\}$ -valued by assumption, we have $\|U_n - U\|_1 \geq c/2m$. This contradicts (27), and finishes the proof of the lemma. \square

Remark. The first part of Theorem 15 can also be deduced from [30, Lemma 2.9].

Corollary 18. Let a sequence of graphs G_1, G_2, \dots converge in the δ_{\square} -distance to a graph $[W]$. Then the sequence $(G_n)_{n \in \mathbb{N}}$ converges to $[W]$ in the δ_1 -distance if and only if W is $\{0, 1\}$ -valued a.e. \square

Proof of Theorem 15. Suppose first that the extremal (f, \mathcal{P}) -problem is stable, as defined in Section 1. Let $[U], [W] \in \text{LIM}(f, \mathcal{P})$. Choose witnesses of this, that is, sequences of almost extremal graphs $(G_{m_i})_{i \in \mathbb{N}}$ and $(H_{n_i})_{i \in \mathbb{N}}$ with $G_{m_i} \rightarrow U$ and $H_{n_i} \rightarrow W$ in the cut-distance as $i \rightarrow \infty$. By stability, $\delta_1(G_{m_i}, H_{n_i}) \rightarrow 0$. Hence,

$$\delta_{\square}(U, W) \leq \delta_{\square}(U, G_{m_i}) + \delta_{\square}(G_{m_i}, H_{n_i}) + \delta_{\square}(H_{n_i}, W) \leq \delta_1(G_{m_i}, H_{n_i}) + o(1) = o(1).$$

Thus $\delta_{\square}(U, W) = 0$. Since $[U], [W] \in \text{LIM}(f, \mathcal{P})$ were arbitrary, the limit set $\text{LIM}(f, \mathcal{P})$ consists of a single graphit $[W]$. Since (G_{m_i}) is Cauchy with respect to the δ_1 -distance, we conclude by Theorem 17 that W is $\{0, 1\}$ -valued a.e., proving one direction of the theorem.

Conversely, suppose that $\text{LIM}(f, \mathcal{P}) = \{[W]\}$ for a $\{0, 1\}$ -valued $W \in \mathcal{W}_I$. Suppose on the contrary that the extremal problem is not stable. This implies that there is some $\varepsilon > 0$ such that for every $i \in \mathbb{N}$ there are $m_i, n_i \geq i, G_{m_i} \in \mathcal{P}_{m_i}, H_{n_i} \in \mathcal{P}_{n_i}$ such that $f(G_{m_i}) \geq \text{ex}_f(m_i, \mathcal{P}) - 1/i, f(H_{n_i}) \geq \text{ex}_f(n_i, \mathcal{P}) - 1/i$, and

$$\delta_1(G_{m_i}, H_{n_i}) \geq \varepsilon. \tag{28}$$

By choosing a subsequence and relabeling, we can additionally assume that for every $i < j$ we have $m_i \leq n_i < m_j \leq n_j$.

By the compactness of $(\mathcal{X}, \delta_{\square})$ we can find a sequence $i_1 < i_2 < \dots$ such that $(G_{m_{i_k}})_{k \in \mathbb{N}}$ is convergent in the δ_{\square} -distance. Since $(G_{m_{i_k}})_{k \in \mathbb{N}}$ is a sequence of almost optimal graphs with increasing orders, its limit is necessarily $[W]$, the unique element of $\text{LIM}(f, \mathcal{P})$. Likewise, we can find a subsequence $j_1 < j_2 < \dots$ of $(i_k)_{k \in \mathbb{N}}$ such that the graph sequence $(H_{n_{j_k}})_{k \in \mathbb{N}}$ converges to $[W]$ in δ_{\square} . Clearly, the intertwined sequence $(G_{m_{j_1}}, H_{n_{j_1}}, G_{m_{j_2}}, H_{n_{j_2}}, \dots)$ still converges to $[W]$. By Corollary 18, the last sequence is Cauchy with respect to the δ_1 -distance. This contradicts (28) and finishes the proof of Theorem 15. \square

6. The Erdős–Simonovits Stability Theorem

In this section, we will prove Theorem 1. For this purpose, we adopt the nice proof of Erdős [14] that every K_{r+1} -free graph G is dominated by some r -partite graph H , that is, $V(H) = V(G)$ and $d_H(x) \geq d_G(x)$ for every $x \in V(G)$, where e.g. $d_H(x)$ denotes the degree of x in H . In order to prove this, Erdős [14] uses induction on r as follows. The case $r = 1$ is trivially true. Let x be a vertex of maximum degree in G and V' be the set of neighbors of x . Then, $G[V']$ is K_r -free, so by the induction assumption we can find an $(r - 1)$ -partite graph H' that dominates $G[V']$. Let H be the r -partite graph obtained from H' by adding a new part on $V(G) \setminus V'$. It is not hard to check that H is the required graph, see [14] for details.

Unfortunately, our proof of the graphon version of this degree-domination result (Theorem 20 in the previous version [33] of this manuscript) is quite long and complicated. Later, during a discussion with Peter Keevash, it was realized that if one is content to prove just Theorem 1, then the arguments dealing with graphons can be shortened. Here we present the shorter proof, referring the interested reader to [33] for the more general result.

Since we are going to apply the Fubini Theorem a few times, we state it here. For a function $W : I^2 \rightarrow \mathbb{R}$ and $x \in I$, let the section functions $W_x, W^x : I \rightarrow \mathbb{R}$ be defined by $W_x(y) = W(x, y)$ and $W^x(y) = W(y, x)$. Let $W_*(x) = \int_I W_x(y) d\mu(y)$ and $W^*(x) = \int_I W^x(y) d\mu(y)$ (and let it be arbitrary if the integral is undefined). Clearly, for a symmetric W , we have $W_x = W^x$ and $W_* = W^*$. Since $(I^2, \mathcal{L}_{I^2}, \mu_{I^2})$ is not the product $(I, \mathcal{L}_I, \mu_I) \times (I, \mathcal{L}_I, \mu_I)$ but its completion, we have to use the Fubini Theorem for Complete Measures ([15, Theorem 2.39]) which easily follows from the standard Fubini Theorem ([15, Theorem 2.37.a]), with the derivation being described in [15, Exercise 2.49].

Theorem 19 (The Fubini Theorem for the Lebesgue Measure). *If $W \in L^1(I^2, \mathcal{L}_{I^2}, \mu_{I^2})$, then $W_x, W^x \in L^1(I, \mathcal{L}_I, \mu_I)$ for a.e. $x \in I$. Furthermore, $W_*, W^* \in L^1(I, \mathcal{L}_I, \mu_I)$ and*

$$\int_{I^2} W(x, y) d\mu(x, y) = \int_I W_*(x) d\mu(x) = \int_I W^*(x) d\mu(x). \quad \square$$

Let $W \in \mathcal{W}_I$ and F be a graph on $[n]$. We call W F -free if for every (not necessarily distinct) $x_1, \dots, x_n \in I$ there is a pair $\{i, j\} \in E(F)$ such that $W(x_i, x_j) = 0$. Equivalently, W is F -free if and only if $W(x, x) = 0$ for every $x \in I$ and there is no homomorphism from F to the infinite (uncountable) graph with vertex set I in which x, y are connected if $W(x, y) > 0$.

If $W \in \mathcal{W}_I$ is F -free, then $t(F, W) = 0$. The converse is not true: for example, fix distinct $x_1, \dots, x_n \in I$ and let $W(x, y) = 0$ except $W(x_i, x_j) = 1$ for all distinct $i, j \in [n]$. However, please note the following Lemma 20, which is a rewording of a special case of a result of Elek and Szegedy [12, Lemma 3.4].

Lemma 20 (The Infinite Removal Lemma). *For every $W \in \mathcal{W}_I$ there is $U \in \mathcal{W}_I$ such that $W = U$ a.e. and for every graph F either $t(F, U) > 0$ or U is F -free.*

Sketch of Proof. Let Z be the Lebesgue set of W , as defined by (19). Clearly, $Z \subseteq I^2$ is symmetric. Let $U(x, y) = W(x, y)$ if $(x, y) \in Z$ and $U(x, y) = 0$ otherwise. Since $\mu(Z) = 1$, $U = W$ a.e. Also, if x_1, \dots, x_n give an F -subgraph in U , then there is $c > 0$ such that for any $\{i, j\} \in E(F)$, the measure of

$$\{(x, y) \in (x_i + c, x_i - c) \times (x_j - c, x_j + c) : W(x, y) > W(x_i, x_j)/2 > 0\}$$

is, for example, at least $(1 - n^{-2}) \cdot 4c^2$. It follows that $t(F, W) > 0$. \square

Remark. Note that $W = U$ a.e. implies that $t(F, U) = t(F, W)$ for every graph F .

For the rest of the section, fix an arbitrary family \mathcal{F} of graphs. Recall that $\rho(G) = 2e(G)/(v(G))^2$ denotes the edge density and $\text{Forb}(\mathcal{F})$ consists of all \mathcal{F} -free graphs. For a graphit $[W]$, define $\rho([W]) = t(K_2, [W])$. For convenience, we just write

$\rho(W)$. This is compatible with the previous definition in the sense that for every graph G we have $\rho(G) = \rho(W_G)$. Define r by (5) and assume that $r \geq 1$.

Let \mathcal{A} consist of those graphits $[W]$ that maximize $\rho(W)$ given that $t(F, W) = 0$ for every $F \in \mathcal{F}$. By the compactness of $(\mathcal{X}, \delta_{\square})$ and the continuity of each function $t(F, -)$, the maximum is attainable. Denote this maximum value by a .

Lemma 21. $\text{LIM}(\rho, \text{Forb}(\mathcal{F})) = \mathcal{A}$.

Proof. Let $[W] \in \text{LIM}(\rho, \text{Forb}(\mathcal{F}))$. Pick a sequence of almost extremal graphs (G_{n_i}) convergent to $[W]$. Since each G_{n_i} is \mathcal{F} -free, we have $t(F, W) = 0$ for each $F \in \mathcal{F}$ by the first definition (14) of $t(F, W)$. We conclude that $\rho(W) \leq a$.

Thus, in order to prove the lemma, it is enough to construct for every $[W] \in \mathcal{A}$ a sequence of \mathcal{F} -free graphs $(G_n)_{n \in \mathbb{N}}$ convergent to $[W]$ with $\rho(G_n) \geq a - o(1)$. Let $U \in [W]$ be obtained from W by applying Lemma 20. For each integer n we generate a random graph G_n on $[n]$ as follows. Pick uniformly at random n elements $x_1, \dots, x_n \in I$ and let a pair $\{i, j\}$ be an edge of G_n with probability $U(x_i, x_j)$, with all $n + \binom{n}{2}$ random choices being mutually independent. With probability 1 we have that the sequence (G_n) converges to $[U] = [W]$ (see [8, Theorem 4.5] or [28, Corollary 2.6]). Thus at least one such sequence (G_n) exists. In particular, we have $\lim_{n \rightarrow \infty} \rho(G_n) = \rho(U) = a$. Also, since U does not contain any copy of $F \in \mathcal{F}$, each G_n is (surely) \mathcal{F} -free, as desired. \square

Remark. The above proof, which is applicable to many other extremal problems, gives another justification why it is better not to restrict ourselves to extremal graphs when defining the limit set $\text{LIM}(f, \mathcal{P})$.

Lemma 20 implies that

$$\mathcal{A} = \{[W]: \forall F \in \mathcal{F} F \not\subseteq W \text{ and } \rho(W) \text{ is maximum}\}. \tag{29}$$

Since W_* is the analytic analog of the degree sequence, the following lemma can be informally rephrased that extremal graphons are degree-regular. The combinatorial interpretation of the proof is that if we have too much discrepancy between degrees in an almost extremal graph G , then by deleting $\varepsilon v(G)$ vertices of smaller degree and cloning $\varepsilon v(G)$ vertices of larger degree, we would substantially increase the size of G , which would be a contradiction (provided we do not create any forbidden subgraph).

Lemma 22. For every $[W] \in \mathcal{A}$ we have $W_*(x) = a$ for a.e. $x \in I$.

Proof. The Fubini Theorem implies that if $W = U$ a.e., then $W_* = U_*$ a.e. (Indeed, if e.g. $W_* > U_*$ on a set $X \subseteq I$ of positive measure, then $\int_{X \times I} (W - U) = \int_X (W_* - U_*) > 0$, a contradiction.) Hence we can assume by Lemma 20 that W is \mathcal{F} -free.

Suppose on the contrary that the lemma is false. Let $X_n = \{x \in I: W_*(x) \leq a - 1/n\}$ and $Y_n = \{y \in I: W_*(y) \geq a + 1/n\}$. Note that e.g. $\cup_{n \in \mathbb{N}} X_n = \{x \in I: W_*(x) < a\}$. Since $\cup_{n \in \mathbb{N}} (X_n \cup Y_n)$ has positive measure, there is some n with $\mu(X_n \cup Y_n) > 0$. Assume, for example, that $\mu(Y_n)$ is positive. By the Fubini Theorem (and $\int_I W = a$), we conclude that $\mu(\cup_{m \in \mathbb{N}} X_m) > 0$. By increasing n , assume that $c = \min(\mu(X_n), \mu(Y_n))$ is positive.

Let $\varepsilon = \min(c, 1/3n)$. By Lemma 5, we can find $\phi \in \Phi_0$ such that $\mu(\phi([0, \varepsilon]) \setminus X_n) = 0$ and $\mu(\phi([\varepsilon, 2\varepsilon]) \setminus Y_n) = 0$. Let $U = W^\phi$. Then $U \in [W]$ is still \mathcal{F} -free while $U_*(x)$ is at most $a - 1/n$ (resp. at least $a + 1/n$) for a.e. x in the interval $[0, \varepsilon]$ (resp. $[\varepsilon, 2\varepsilon]$).

For $x \in I$, let $\psi(x) = x$ if $x \geq \varepsilon$ and $\psi(x) = x + \varepsilon$ if $x < \varepsilon$. Let $V = U^\psi \in \mathcal{W}_I$. (Although ψ is not measure preserving, this definition makes perfect sense.) Note that V is \mathcal{F} -free: if $x_1, \dots, x_m \in I$ induce a copy of F in V , then $\psi(x_1), \dots, \psi(x_m)$ induce a copy of F in U . Moreover,

$$\rho(V) = \int_I V_* \geq \int_I U_* - \int_{[0, \varepsilon]} U_* + \int_{[\varepsilon, 2\varepsilon]} U_* - (2\varepsilon)^2 \geq a - \varepsilon(a - 1/n) + \varepsilon(a + 1/n) - 4\varepsilon^2 > a.$$

This contradicts the maximality of a . \square

For disjoint measurable sets $A_1, \dots, A_r \subseteq I$, the complete r -partite graphon K_{A_1, \dots, A_r} is the simple function from I^2 to $\{0, 1\}$ that assumes value 1 on $\cup_{i \in [r]} \cup_{j \in [r] \setminus \{i\}} A_i \times A_j$ and 0 on the remaining part of I^2 . (In other words, $W(x, y) = 1$ if x, y come from two different sets A_i and 0 otherwise.) Clearly, K_{A_1, \dots, A_r} is K_{r+1} -free.

Next, we prove that the graphon problem has the unique solution when we forbid the clique K_{r+1} only.

Lemma 23. If $\mathcal{F} = \{K_{r+1}\}$ and $[W] \in \mathcal{A}$, then there is a partition $I = A_1 \cup \dots \cup A_r$ into sets of measure $1/r$ such that $W = K_{A_1, \dots, A_r}$ a.e.

Proof. We use induction on r with the case $r = 1$ being trivially true.

Let $r \geq 2$. The \mathcal{F} -free graphon W_{K_r} demonstrates that $a \geq (r - 1)/r$. Let $[W] \in \mathcal{A}$. Assume that W is K_{r+1} -free by (29) (that is, by Lemma 20). Pick $u \in I$ such that $W_*(u) = a$ which exists by Lemma 22. Let $B = \{w \in I: W(u, w) > 0\}$ and $A_1 = I \setminus B$. Let $b = \mu(B)$. Since $W \leq 1$, we have $b \geq a$. We are free to replace W by W^ϕ with any $\phi \in \Phi_0$; thus we can assume by Lemma 5 that $\mu(B \Delta [0, b]) = 0$. The graphon $U(x, y) = W(bx, by)$ is K_r -free: if $x_1, \dots, x_r \in I$ induce K_r in U , then $bx_1, \dots, bx_r, u \in I$ induce K_{r+1} in W , a contradiction. Note that, by the Fubini Theorem,

$$a = \rho(W) = \int_{B^2} W + 2 \int_{A_1} W_* - \int_{A_1^2} W = b^2 \rho(U) + 2(1 - b)a - \int_{A_1^2} W. \tag{30}$$

The inductive assumption implies that $\rho(U) \leq (r-2)/(r-1)$. Thus

$$\frac{r-1}{r} \leq a \leq \frac{r-2}{r-1} b^2 + 2(1-b)a \leq \frac{r-2}{r-1} b^2 + 2(1-b)b.$$

Routine algebra implies that $a = b = (r-1)/r$ and all inequalities are in fact equalities. Thus $W(x, y) = 0$ for a.e. $(x, y) \in A_1^2$. Since $W_*(x) = a = 1 - \mu(A_1)$ for almost every $x \in A_1$, we have by the Fubini Theorem that $W(x, y) = 1$ for a.e. $(x, y) \in A_1 \times B$. Furthermore, by the uniqueness part of the induction assumption, $U = K_{B_2, \dots, B_r}$ a.e. for some equitable partition $I = B_2 \cup \dots \cup B_r$. Letting $A_i = \{bx : x \in B_i\}$, we get $W = K_{A_1, \dots, A_r}$ a.e., as required. \square

Proof of Theorem 1. By Theorem 15, it suffices to show that any $[W] \in \text{LIM}(\rho, \text{Forb}(\mathcal{F}))$ we have $\delta_{\square}(W, K_r) = 0$. By Lemma 21 and (29), we can assume that W is \mathcal{F} -free.

Let us show that W is K_{r+1} -free. Suppose on the contrary that $x_1, \dots, x_{r+1} \in I$ induce K_{r+1} in W . Select $F \in \mathcal{F}$ of chromatic number $\chi(F) = r+1$ and fix a proper coloring $c : V(F) \rightarrow [r+1]$. Then the map $f : V(F) \rightarrow I$ with $f(u) = x_{c(u)}$ shows that $F \subseteq W$, a contradiction.

Turán graphs $T_r(n)$ (or the graphon W_{K_r} and Lemma 21) show that $\rho(W) \geq (r-1)/r$. By Lemma 23 we have that $\rho(W) \leq \rho(K_r) \leq (r-1)/r$. Thus $[W]$ is extremal for the $(\rho, \text{Forb}(\{K_{r+1}\}))$ -problem and (again by Lemma 23) is equal to $[W_{K_r}]$, as required. \square

Acknowledgements

The author thanks Boris Bukh, Tomasz Gdala, Peter Keevash, Po-Shen Loh, László Lovász, Alexander Razborov, Asaf Shapira, and Benny Sudakov for helpful discussions and/or comments. Also, the author is grateful to the anonymous referee for careful reading. The author was partially supported by the National Science Foundation, Grant DMS-0758057, and the Berkman Faculty Development Fund, CMU.

References

- [1] I. Benjamini, O. Schramm, Recurrence of distributional limits of finite planar graphs, *Electron. J. Probab.* 6 (2001) 13.
- [2] I. Benjamini, O. Schramm, A. Shapira, Every minor-closed property of sparse graphs is testable, *Adv. Math.* 223 (2010) 2200–2218.
- [3] G. Birkhoff, Three observations on linear algebra, *Univ. Nac. Tucumán. Rev. Ser. A* 5 (1946).
- [4] C. Borgs, J. Chayes, L. Lovász, Moments of two-variable functions and the uniqueness of graph limits, *Geom. Funct. Anal.* 19 (2010) 1597–1619.
- [5] C. Borgs, J. Chayes, L. Lovász, V.T. Sós, B. Szegedy, K. Vesztegombi, Graph limits and parameter testing, in: *Proc. of 38th ACM Symp. Theory of Comp.*, STOC, 2006, pp. 261–270.
- [6] C. Borgs, J. Chayes, L. Lovász, V.T. Sós, K. Vesztegombi, Counting graph homomorphisms, in: *Topics in Discrete Mathematics*, in: *Algorithms Combin.*, vol. 26, Springer Verlag, 2006, pp. 315–371.
- [7] C. Borgs, J. Chayes, L. Lovász, V.T. Sós, K. Vesztegombi, Convergent graph sequences II: multiway cuts and statistical physics, Preprint, 2008.
- [8] C. Borgs, J. Chayes, L. Lovász, V.T. Sós, K. Vesztegombi, Convergent sequences of dense graphs I: subgraph frequencies, metric properties and testing, *Adv. Math.* 219 (2008) 1801–1851.
- [9] P. Diaconis, S. Janson, Graph limits and exchangeable random graphs, *Rend. Mat.* 28 (2008) 33–61.
- [10] E. DiBenedetto, *Real Analysis*, Birkhäuser, 2002.
- [11] G. Elek, A regularity lemma for bounded degree graphs and its applications: parameter testing and infinite volume limits, Version 2, 2007. E-print: arxiv.org:0711.2800.
- [12] G. Elek, B. Szegedy, A measure-theoretic approach to the theory of dense hypergraphs, 2008. E-print: arxiv.org:0810.4062.
- [13] P. Erdős, Some recent results on extremal problems in graph theory. Results, in: *Theory of Graphs*, Internat. Sympos., Rome, 1966, Gordon and Breach, New York, 1967, pp. 117–123 (in English); pp. 124–130 (in French).
- [14] P. Erdős, On the graph theorem of Turán, *Mat. Lapok* 21 (1970) 249–251 (in Hungarian).
- [15] G.B. Folland, *Real Analysis*, 2nd ed., Wiley-Intersci. Publ., 1999.
- [16] M. Freedman, L. Lovász, A. Schrijver, Reflection positivity, rank connectivity, and homomorphism of graphs, *J. Amer. Math. Soc.* 20 (2006) 37–51.
- [17] A. Frieze, R. Kannan, Quick approximation to matrices and applications, *Combinatorica* 19 (1999) 175–220.
- [18] Z. Füredi, A proof of the stability of extremal graphs. Talk at (Your Regular) Workshop on Extremal Graphs and Hypergraphs, CMU, May 5–6, 2007.
- [19] O. Goldreich, S. Goldwasser, D. Ron, Property testing and its connection to learning and approximation, *J. ACM* 45 (1998) 653–750.
- [20] O. Goldreich, M. Krivelevich, I. Newman, E. Rozenberg, Hierarchy theorems for property testing, Technical Report TR08-097, Electronic Colloquium on Computational Complexity, 2008.
- [21] O. Goldreich, M. Krivelevich, I. Newman, E. Rozenberg, Hierarchy Theorems for Property Testing, in: *Lecture Notes in Comp. Sci.*, vol. 5687, Springer Verlag, 2009, pp. 504–519.
- [22] P. Hall, On representatives of subsets, *J. Lond. Math. Soc.* 10 (1935) 26–30.
- [23] S. Janson, Connectedness in graph limits, 2008. E-print: arxiv.org:0802.3795.
- [24] T. Jiang, O. Pikhurko, Anti-Ramsey numbers of doubly edge-critical graphs, *J. Graph Theory* 61 (2009) 210–218.
- [25] F. Lazebnik, O. Pikhurko, A. Woldar, Maximum number of colorings of $(2k, k^2)$ -graphs, *J. Graph Theory* 56 (2007) 135–148.
- [26] P.-S. Loh, B. Sudakov, O. Pikhurko, Maximizing the number of q -colorings. *Proc. Lond. Math. Soc.*, in press ([doi:10.1112/plms/pdp041](https://doi.org/10.1112/plms/pdp041)).
- [27] L. Lovász, Very large graphs, *Curr. Dev. Math.* 2008 (2009) 67–128.
- [28] L. Lovász, B. Szegedy, Limits of dense graph sequences, *J. Combin. Theory Ser. B* 96 (2006) 933–957.
- [29] L. Lovász, B. Szegedy, Szemerédi’s lemma for the analyst, *Geom. Funct. Anal.* 17 (2007) 252–270.
- [30] L. Lovász, B. Szegedy, Testing properties of graphs and functions, Version 2, 2008. E-print: arxiv.org:0803.1248.
- [31] L. Lovász, B. Szegedy, Finitely forcible graphons, 2009. E-print: arxiv.org:0901.0929.
- [32] D. Mubayi, O. Pikhurko, A new generalization of Mantel’s theorem to k -graphs, *J. Combin. Theory Ser. B* 97 (2007) 669–678.
- [33] O. Pikhurko, An analytic approach to stability, 2008. E-print: arxiv.org:0812.0214.
- [34] O. Pikhurko, Z. Yilma, The maximum number of K_3 -free and K_4 -free edge-colorings, 2010, 24 pages (submitted for publication).
- [35] A. Razborov, Flag algebras, *J. Symbolic Logic* 72 (2007) 1239–1282.
- [36] A. Razborov, On the minimal density of triangles in graphs, *Combin. Probab. Comput.* 17 (2008) 603–618.
- [37] M. Simonovits, A method for solving extremal problems in graph theory, stability problems, in: *Theory of Graphs*, Proc. Colloq., Tihany, 1966, Academic Press, 1968, pp. 279–319.
- [38] G. Tardos, Limits of trees. Talk at the Workshop “Topics in Graphs and Hypergraphs”, IPAM, November 2–6, 2009.
- [39] P. Turán, On an extremal problem in graph theory, *Mat. Fiz. Lapok* 48 (1941) 436–452 (in Hungarian).