Pseudorandomness and regularity in graphs I

David Conlon

January 17, 2017

Three regimes

The study of pseudorandomness in graphs splits into three different regimes:

Three regimes

The study of pseudorandomness in graphs splits into three different regimes:

- The sparse case - $d = O(1)$

Three regimes

The study of pseudorandomness in graphs splits into three different regimes:

- The sparse case - $d = O(1)$
- The dense case - $d = \Omega(n)$

Three regimes

The study of pseudorandomness in graphs splits into three different regimes:

- The sparse case - $d = O(1)$
- The dense case - $d = \Omega(n)$
- The intermediate case

Three regimes

The study of pseudorandomness in graphs splits into three different regimes:

- The sparse case - $d = O(1)$
- The dense case - $d = \Omega(n)$
- The intermediate case

In this talk, we will focus on discussing the second and third cases.

The dense case

The dense case

p-quasirandom sequence

A sequence of graphs $(G_n)_{n\in\mathbb{N}}$ with $|G_n| = n$ is said to be p-quasirandom if

$$|e(X, Y) - p|X||Y|| = o(pn^2)$$

for all $X, Y \subseteq V(G_n)$.

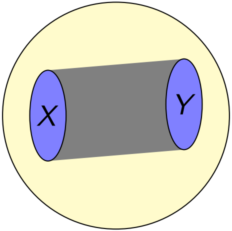

The dense case

p-quasirandom sequence

A sequence of graphs $(G_n)_{n\in\mathbb{N}}$ with $|G_n| = n$ is said to be p-quasirandom if

$$|e(X,Y) - p|X||Y|| = o(pn^2)$$

for all $X, Y \subseteq V(G_n)$.

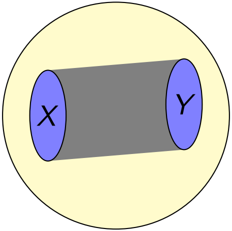

Definition motivated by Szemerédi's regularity lemma.

The adjacency matrix

The adjacency matrix A of a graph G on vertex set $\{1, 2, \ldots, n\}$ is the $n \times n$ matrix with entries given by

$$A_{uv} = \begin{cases} 0 & \text{if } uv \notin E(G); \\ 1 & \text{if } uv \in E(G). \end{cases}$$

The adjacency matrix

The adjacency matrix A of a graph G on vertex set $\{1, 2, \ldots, n\}$ is the $n \times n$ matrix with entries given by

$$A_{uv} = \begin{cases} 0 & \text{if } uv \notin E(G); \\ 1 & \text{if } uv \in E(G). \end{cases}$$

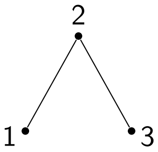

$$\begin{pmatrix} 0 & 1 & 0 \\ 1 & 0 & 1 \\ 0 & 1 & 0 \end{pmatrix}$$

The adjacency matrix

The adjacency matrix A of a graph G on vertex set $\{1, 2, \ldots, n\}$ is the $n \times n$ matrix with entries given by

$$A_{uv} = \begin{cases} 0 & \text{if } uv \notin E(G); \\ 1 & \text{if } uv \in E(G). \end{cases}$$

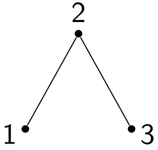

$$\begin{pmatrix} 0 & 1 & 0 \\ 1 & 0 & 1 \\ 0 & 1 & 0 \end{pmatrix}$$

Denote the eigenvalues of the adjacency matrix of A by $\lambda_1 \geq \lambda_2 \geq \cdots \geq \lambda_n$.

The Chung–Graham–Wilson theorem

Theorem (Chung–Graham–Wilson, 1989)

For any fixed $0 < p < 1$ and any sequence of graphs $(G_n)_{n \in \mathbb{N}}$ of density p with $|G_n| = n$, the following properties are equivalent:

(i) for all $X, Y \subseteq V(G_n)$, $|e(X, Y) - p|X||Y|| = o(pn^2)$;

The Chung–Graham–Wilson theorem

Theorem (Chung–Graham–Wilson, 1989)

For any fixed $0 < p < 1$ and any sequence of graphs $(G_n)_{n \in \mathbb{N}}$ of density p with $|G_n| = n$, the following properties are equivalent:

(i) for all $X, Y \subseteq V(G_n)$, $|e(X, Y) - p|X||Y|| = o(pn^2)$;

(ii) $\lambda_1(G_n) = (1 + o(1))pn$ and $\lambda_i(G_n) = o(pn)$ for all $i \neq 1$;

The Chung–Graham–Wilson theorem

Theorem (Chung–Graham–Wilson, 1989)

For any fixed $0 < p < 1$ and any sequence of graphs $(G_n)_{n \in \mathbb{N}}$ of density p with $|G_n| = n$, the following properties are equivalent:

(i) for all $X, Y \subseteq V(G_n)$, $|e(X, Y) - p|X||Y|| = o(pn^2)$;

(ii) $\lambda_1(G_n) = (1 + o(1))pn$ and $\lambda_i(G_n) = o(pn)$ for all $i \neq 1$;

(iii) for all fixed graphs H, the number of labelled copies of H is $(1 + o(1))p^{e(H)}n^{v(H)}$;

The Chung–Graham–Wilson theorem

Theorem (Chung–Graham–Wilson, 1989)

For any fixed $0 < p < 1$ and any sequence of graphs $(G_n)_{n \in \mathbb{N}}$ of density p with $|G_n| = n$, the following properties are equivalent:

(i) for all $X, Y \subseteq V(G_n)$, $|e(X, Y) - p|X||Y|| = o(pn^2)$;

(ii) $\lambda_1(G_n) = (1 + o(1))pn$ and $\lambda_i(G_n) = o(pn)$ for all $i \neq 1$;

(iii) for all fixed graphs H, the number of labelled copies of H is $(1 + o(1))p^{e(H)}n^{v(H)}$;

(iv) the number of labelled copies of C_4 is $(1 + o(1))p^4 n^4$.

The Chung–Graham–Wilson theorem

Theorem (Chung–Graham–Wilson, 1989)

For any fixed $0 < p < 1$ and any sequence of graphs $(G_n)_{n \in \mathbb{N}}$ of density p with $|G_n| = n$, the following properties are equivalent:

(i) for all $X, Y \subseteq V(G_n)$, $|e(X, Y) - p|X||Y|| = o(pn^2)$;

(ii) $\lambda_1(G_n) = (1 + o(1))pn$ and $\lambda_i(G_n) = o(pn)$ for all $i \neq 1$;

(iii) the number of labelled copies of C_4 is $(1 + o(1))p^4 n^4$.

The Chung–Graham–Wilson theorem

Theorem (Chung–Graham–Wilson, 1989)

For any fixed $0 < p < 1$ and any sequence of graphs $(G_n)_{n \in \mathbb{N}}$ of density p with $|G_n| = n$, the following properties are equivalent:

- DIS for all $X, Y \subseteq V(G_n)$, $|e(X, Y) - p|X||Y|| = o(pn^2)$;
- EIG $\lambda_1(G_n) = (1 + o(1))pn$ and $\lambda_i(G_n) = o(pn)$ for all $i \neq 1$;
- CYC the number of labelled copies of C_4 is $(1 + o(1))p^4 n^4$.

Expander Mixing Lemma

If G is a graph for which all eigenvalues of the adjacency matrix, save the largest, have absolute value at most λ, then

$$|e(X, Y) - p|X||Y|| \leq \lambda\sqrt{|X||Y|}$$

for all $X, Y \subseteq V(G)$.

EIG $\implies$ DIS

Expander Mixing Lemma

If G is a graph for which all eigenvalues of the adjacency matrix, save the largest, have absolute value at most λ, then

$$\big|e(X,Y) - p|X||Y|\big| \leq \lambda\sqrt{|X||Y|}$$

for all $X, Y \subseteq V(G)$.

Applying this lemma with $\lambda = o(pn)$ gives the required result.

Since the number of cycles of length 4 is $(1 + o(1))p^4 n^4$, the same is true for the number of circuits of length 4.

CYC $\implies$ EIG

Since the number of cycles of length 4 is $(1+o(1))p^4 n^4$, the same is true for the number of circuits of length 4. But the number of circuits of length 4 is equal to

$$Tr(A^4) = \sum_{i=1}^{n} \lambda_i^4.$$

Since the number of cycles of length 4 is $(1+o(1))p^4n^4$, the same is true for the number of circuits of length 4. But the number of circuits of length 4 is equal to

$$Tr(A^4) = \sum_{i=1}^{n} \lambda_i^4.$$

Since $\lambda_1 = (1+o(1))pn$, we must have $\lambda_i = o(pn)$ for all $i \neq 1$.

DIS $\implies$ CYC

Counting Lemma

If G is a graph which satisfies DIS with p fixed, then, for any fixed graph H, the number of labelled copies of H in G is

$$(1 + o(1))p^{e(H)}n^{v(H)}.$$

DIS $\implies$ CYC

Counting Lemma

If G is a graph which satisfies DIS with p fixed, then, for any fixed graph H, the number of labelled copies of H in G is

$$(1 + o(1))p^{e(H)}n^{v(H)}.$$

The required result is just the case where $H = C_4$.

DIS $\implies$ CYC

To prove the counting lemma when $H = K_3$, write

$$e(S, T) = \sum_{s \in S, t \in T} 1_G(s, t),$$

where 1_G is the indicator function for edges of G.

DIS $\implies$ CYC

To prove the counting lemma when $H = K_3$, write

$$e(S, T) = \sum_{s \in S, t \in T} 1_G(s, t),$$

where 1_G is the indicator function for edges of G. By assumption,

$$|e(S, T) - p|S||T|| = o(n^2)$$

for all $S, T \subseteq V(G)$. Rewriting this conclusion, we have

$$|\sum_{s \in S, t \in T} (1_G(s, t) - p)| = o(n^2).$$

DIS $\implies$ CYC

To prove the counting lemma when $H = K_3$, write

$$e(S, T) = \sum_{s \in S, t \in T} 1_G(s, t),$$

where 1_G is the indicator function for edges of G. By assumption,

$$|e(S, T) - p|S||T|| = o(n^2)$$

for all $S, T \subseteq V(G)$. Rewriting this conclusion, we have

$$|\sum_{s \in S, t \in T} (1_G(s, t) - p)| = o(n^2).$$

In turn, this implies that for any functions $u, v : V(G) \to [0, 1]$,

$$|\sum_{x, y \in V(G)} (1_G(x, y) - p) u(x) v(y)| = o(n^2).$$

This follows since the function we wish to optimise is linear in $u(x)$ and $v(y)$ for each x and y, so the maximum occurs when u and v are $\{0, 1\}$-valued, i.e., indicator functions.

By telescoping, the deviation between the number of labeled triangles in a set $S \subseteq V(G)$ and its expected value is

$$\sum_{x,y,z \in S} (1_G(x,y)1_G(y,z)1_G(z,x) - p^3)$$

$$= \sum_{x,y,z \in S} (1_G(x,y) - p)1_G(y,z)1_G(z,x) + \sum_{x,y,z \in S} p(1_G(y,z) - p)1_G(z,x) + \sum_{x,y,z \in S} p^2(1_G(z,x) - p).$$

DIS $\implies$ CYC

By telescoping, the deviation between the number of labeled triangles in a set $S \subseteq V(G)$ and its expected value is

$$\sum_{x,y,z \in S} (1_G(x,y)1_G(y,z)1_G(z,x) - p^3)$$

$$= \sum_{x,y,z \in S} (1_G(x,y) - p)1_G(y,z)1_G(z,x) +$$

$$\sum_{x,y,z \in S} p(1_G(y,z) - p)1_G(z,x) + \sum_{x,y,z \in S} p^2(1_G(z,x) - p).$$

Each term on the right-hand side of this equation may be written as a sum over terms of the form $\sum_{x,y \in V(G)} (1_G(x,y) - p)u(x)v(y)$ for some appropriate u and v, thus implying that the deviation we are interested in is $o(n^3)$.

Forcing graphs

Forcing graph

A graph H is said to be forcing if any sequence of graphs $(G_n)_{n \in \mathbb{N}}$ of density p with $|G_n| = n$ containing $(1 + o(1))p^{e(H)}n^{v(H)}$ labelled copies of H is p-quasirandom.

Forcing graphs

Forcing graph

A graph H is said to be forcing if any sequence of graphs $(G_n)_{n\in\mathbb{N}}$ of density p with $|G_n| = n$ containing $(1 + o(1))p^{e(H)}n^{v(H)}$ labelled copies of H is p-quasirandom.

Examples

- C_4

Forcing graphs

Forcing graph

A graph H is said to be forcing if any sequence of graphs $(G_n)_{n\in\mathbb{N}}$ of density p with $|G_n| = n$ containing $(1+o(1))p^{e(H)}n^{v(H)}$ labelled copies of H is p-quasirandom.

Examples

- C_4
- C_{2k} - Chung–Graham–Wilson, 1989

Forcing graphs

Forcing graph

A graph H is said to be forcing if any sequence of graphs $(G_n)_{n\in\mathbb{N}}$ of density p with $|G_n| = n$ containing $(1+o(1))p^{e(H)}n^{v(H)}$ labelled copies of H is p-quasirandom.

Examples

- C_4
- C_{2k} - Chung–Graham–Wilson, 1989
- $K_{s,t}$ - Skokan–Thoma, 2004

Forcing graphs

Forcing graph

A graph H is said to be forcing if any sequence of graphs $(G_n)_{n\in\mathbb{N}}$ of density p with $|G_n| = n$ containing $(1+o(1))p^{e(H)}n^{v(H)}$ labelled copies of H is p-quasirandom.

Examples

- C_4
- C_{2k} - Chung–Graham–Wilson, 1989
- $K_{s,t}$ - Skokan–Thoma, 2004

Negative results

H is not forcing if it is

Forcing graphs

Forcing graph

A graph H is said to be forcing if any sequence of graphs $(G_n)_{n\in\mathbb{N}}$ of density p with $|G_n| = n$ containing $(1+o(1))p^{e(H)}n^{v(H)}$ labelled copies of H is p-quasirandom.

Examples

- C_4
- C_{2k} - Chung–Graham–Wilson, 1989
- $K_{s,t}$ - Skokan–Thoma, 2004

Negative results

H is not forcing if it is

- non-bipartite

Forcing graphs

Forcing graph

A graph H is said to be forcing if any sequence of graphs $(G_n)_{n\in\mathbb{N}}$ of density p with $|G_n| = n$ containing $(1+o(1))p^{e(H)}n^{v(H)}$ labelled copies of H is p-quasirandom.

Examples

- C_4
- C_{2k} - Chung–Graham–Wilson, 1989
- $K_{s,t}$ - Skokan–Thoma, 2004

Negative results

H is not forcing if it is

- non-bipartite or
- a tree.

The forcing conjecture

The forcing conjecture

A graph H is forcing if it is bipartite and contains a cycle.

The forcing conjecture

The forcing conjecture

A graph H is forcing if it is bipartite and contains a cycle.

Progress on this conjecture closely parallels the progress on Sidorenko's conjecture (Conlon–Fox–Sudakov, Szegedy–Li, Kim–Lee–Lee, Szegedy, Conlon–Kim–Lee–Lee).

The forcing conjecture

A graph H is forcing if it is bipartite and contains a cycle.

Progress on this conjecture closely parallels the progress on Sidorenko's conjecture (Conlon–Fox–Sudakov, Szegedy–Li, Kim–Lee–Lee, Szegedy, Conlon–Kim–Lee–Lee).

Example - Conlon–Fox–Sudakov; Szegedy–Li

Every bipartite graph which is not a tree and has a vertex which is complete to the other side is forcing.

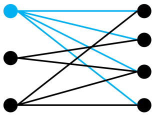

A result of Simonovits and Sós

Theorem (Chung–Graham–Wilson, 1989)

For any fixed $0 < p < 1$ and any sequence of graphs $(G_n)_{n \in \mathbb{N}}$ of density p with $|G_n| = n$, the following properties are equivalent:

- **DIS** for all $X, Y \subseteq V(G_n)$, $|e(X, Y) - p|X||Y|| = o(pn^2)$;
- **EIG** $\lambda_1(G_n) = (1 + o(1))pn$ and $\lambda_i(G_n) = o(pn)$ for all $i \neq 1$;
- **CYC** the number of labelled copies of C_4 is $(1 + o(1))p^4 n^4$;

A result of Simonovits and Sós

Theorem (Chung–Graham–Wilson, 1989)

For any fixed $0 < p < 1$ and any sequence of graphs $(G_n)_{n \in \mathbb{N}}$ of density p with $|G_n| = n$, the following properties are equivalent:

- DIS for all $X, Y \subseteq V(G_n)$, $|e(X, Y) - p|X||Y|| = o(pn^2)$;
- EIG $\lambda_1(G_n) = (1 + o(1))pn$ and $\lambda_i(G_n) = o(pn)$ for all $i \neq 1$;
- CYC the number of labelled copies of C_4 is $(1 + o(1))p^4 n^4$;
- HER for every $S \subseteq V(G_n)$, the number of labelled triangles in $G_n[S]$ is equal to $p^3|S|^3 + o(p^3 n^3)$.

Hereditary quasirandomness

Theorem (Simonovits–Sós, 1997)

For any $0 < p < 1$ and any $\epsilon > 0$, there exists $\delta > 0$ such that if the number of labelled triangles in $G[S]$ is equal to $p^3|S|^3 \pm \delta n^3$, then

$$|e(X, Y) - p|X||Y|| \leq \epsilon n^2$$

for all $X, Y \subseteq V(G)$.

Hereditary quasirandomness

Theorem (Simonovits–Sós, 1997)

For any $0 < p < 1$ and any $\epsilon > 0$, there exists $\delta > 0$ such that if the number of labelled triangles in $G[S]$ is equal to $p^3|S|^3 \pm \delta n^3$, then

$$|e(X, Y) - p|X||Y|| \leq \epsilon n^2$$

for all $X, Y \subseteq V(G)$.

The proof uses the regularity lemma and so gives a bound of the form

$$\delta^{-1} \leq 2^{2^{\cdot^{\cdot^{\cdot^{2}}}}}\Big\}_{\epsilon^{-c}}$$

Hereditary quasirandomness

Theorem (Simonovits–Sós, 1997)

For any $0 < p < 1$ and any $\epsilon > 0$, there exists $\delta > 0$ such that if the number of labelled triangles in $G[S]$ is equal to $p^3|S|^3 \pm \delta n^3$, then

$$|e(X, Y) - p|X||Y|| \leq \epsilon n^2$$

for all $X, Y \subseteq V(G)$.

The proof uses the regularity lemma and so gives a bound of the form

$$\delta^{-1} \leq 2^{2^{\cdot^{\cdot^{\cdot^{2}}}}} \Big\} \epsilon^{-c}$$

Theorem (C.–Fox–Sudakov, 2016+)

One may take $\delta = \Omega(\epsilon)$.

Hereditary quasirandomness

Theorem (Simonovits–Sós, 1997)

For any $0 < p < 1$ and any $\epsilon > 0$, there exists $\delta > 0$ such that if the number of labelled triangles in $G[S]$ is equal to $p^3|S|^3 \pm \delta n^3$, then

$$|e(X, Y) - p|X||Y|| \leq \epsilon n^2$$

for all $X, Y \subseteq V(G)$.

The proof uses the regularity lemma and so gives a bound of the form

$$\delta^{-1} \leq 2^{2^{\cdot^{\cdot^{\cdot^{2}}}}} \Big\} \epsilon^{-c}$$

Theorem (C.–Fox–Sudakov, 2016+)

One may take $\delta = \Omega(\epsilon)$.

An alternative regularity-free proof given by Reiher and Schacht.

An application

Ramsey's theorem

For any graph H, there exists a natural number n such that if the edges of the complete graph K_n are two-coloured, there is always a monochromatic copy of H.

The Ramsey number $r(H)$ is the smallest n for which this holds.

An application

Ramsey's theorem

For any graph H, there exists a natural number n such that if the edges of the complete graph K_n are two-coloured, there is always a monochromatic copy of H.

The Ramsey number $r(H)$ is the smallest n for which this holds.

Bounds for complete graphs

$$\sqrt{2}^t \leq r(K_t) \leq 4^t$$

An application

Theorem (Erdős–Szekeres, 1935)

$$r(t) \leq \binom{2t-2}{t-1} = O\left(\frac{4^t}{\sqrt{t}}\right).$$

An application

Theorem (Erdős–Szekeres, 1935)

$$r(t) \leq \binom{2t-2}{t-1} = O\left(\frac{4^t}{\sqrt{t}}\right).$$

Theorem (C., 2009)

There exists a positive constant c such that

$$r(t) \leq \frac{4^t}{t^{c \log t / \log \log t}}.$$

An application

Theorem (Erdős–Szekeres, 1935)

$$r(t) \le \binom{2t-2}{t-1} = O\left(\frac{4^t}{\sqrt{t}}\right).$$

Theorem (C., 2009)

There exists a positive constant c such that

$$r(t) \le \frac{4^t}{t^{c \log t / \log \log t}}.$$

Proof idea:

- Show that any colouring with no monochromatic K_t is quasirandom.

An application

Theorem (Erdős–Szekeres, 1935)

$$r(t) \leq \binom{2t-2}{t-1} = O\left(\frac{4^t}{\sqrt{t}}\right).$$

Theorem (C., 2009)

There exists a positive constant c such that

$$r(t) \leq \frac{4^t}{t^{c \log t / \log \log t}}.$$

Proof idea:

- Show that any colouring with no monochromatic K_t is quasirandom.
- Show that the colouring contains the correct number of monochromatic K_p for $p \approx \log t / \log \log t$.

An application

Theorem (Erdős–Szekeres, 1935)

$$r(t) \le \binom{2t-2}{t-1} = O\left(\frac{4^t}{\sqrt{t}}\right).$$

Theorem (C., 2009)

There exists a positive constant c such that

$$r(t) \le \frac{4^t}{t^{c \log t / \log\log t}}.$$

Proof idea:

- Show that any colouring with no monochromatic K_t is quasirandom.
- Show that the colouring contains the correct number of monochromatic K_p for $p \approx \log t / \log\log t$.
- Show that this contradicts another property of the colouring.

Sparser graphs

Problem

Does the Chung–Graham–Wilson theorem hold when $p = o(1)$?

Sparser graphs

Problem

Does the Chung–Graham–Wilson theorem hold when $p = o(1)$?

YES: EIG $\Rightarrow$ DIS

Sparser graphs

Problem

Does the Chung–Graham–Wilson theorem hold when $p = o(1)$?

YES: EIG $\Rightarrow$ DIS

Apply the Expander Mixing Lemma.

Sparser graphs

Problem

Does the Chung–Graham–Wilson theorem hold when $p = o(1)$?

YES: EIG $\Rightarrow$ DIS

Apply the Expander Mixing Lemma.

NO: DIS $\not\Rightarrow$ EIG

Sparser graphs

Problem

Does the Chung–Graham–Wilson theorem hold when $p = o(1)$?

YES: EIG $\Rightarrow$ DIS

Apply the Expander Mixing Lemma.

NO: DIS $\not\Rightarrow$ EIG

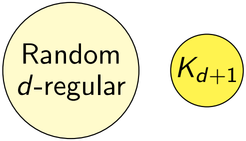

Some positive results

Theorem (Bilu–Linial, 2006)

Suppose that G is a d-regular graph with n vertices such that

$$|e(X,Y) - \frac{d}{n}|X||Y|| \leq \eta\sqrt{|X||Y|}$$

for all $X, Y \subseteq V(G)$. Then every eigenvalue of the adjacency matrix, save the largest, has absolute value $O(\eta \log d)$.

Some positive results

Theorem (Bilu–Linial, 2006)

Suppose that G is a d-regular graph with n vertices such that

$$|e(X, Y) - \frac{d}{n}|X||Y|| \leq \eta\sqrt{|X||Y|}$$

for all $X, Y \subseteq V(G)$. Then every eigenvalue of the adjacency matrix, save the largest, has absolute value $O(\eta \log d)$.

Theorem (Alon–Coja-Oghlan–Hàn–Kang–Rödl–Schacht, 2010)

Suppose that $(G_n)_{n\in\mathbb{N}}$ with $|G_n| = n$ is a sequence of graphs such that

$$|e(X, Y) - p|X||Y|| = o(pn^2)$$

for all $X, Y \subseteq V(G_n)$. Then one may remove a $o(1)$-fraction of the vertices to find a sequence of graphs $(G'_n)_{n\in\mathbb{N}}$ satisfying EIG.

Cayley graphs

Cayley graph

Suppose that G is a group and S is a subset of G satisfying $S = S^{-1}$. The Cayley graph $\mathrm{Cay}(G, S)$ is the graph with vertex set G and edge set $\{(sg, g) : g \in G, s \in S\}$.

Cayley graph

Suppose that G is a group and S is a subset of G satisfying $S = S^{-1}$. The Cayley graph $\mathrm{Cay}(G, S)$ is the graph with vertex set G and edge set $\{(sg, g) : g \in G, s \in S\}$.

$\mathrm{Cay}(G, S)$ is an n-vertex d-regular graph with $n = |G|$ and $d = |S|$.

Cayley graph

Suppose that G is a group and S is a subset of G satisfying $S = S^{-1}$. The Cayley graph $\mathrm{Cay}(G, S)$ is the graph with vertex set G and edge set $\{(sg, g) : g \in G, s \in S\}$.

$\mathrm{Cay}(G, S)$ is an n-vertex d-regular graph with $n = |G|$ and $d = |S|$.

Examples

- Paley graph - $G = \mathbb{Z}_p$, $S = \{x^2 : x \in \mathbb{Z}_p \setminus \{0\}\}$

Cayley graphs

Cayley graph

Suppose that G is a group and S is a subset of G satisfying $S = S^{-1}$. The Cayley graph $\text{Cay}(G, S)$ is the graph with vertex set G and edge set $\{(sg, g) : g \in G, s \in S\}$.

$\text{Cay}(G, S)$ is an n-vertex d-regular graph with $n = |G|$ and $d = |S|$.

Examples

- Paley graph - $G = \mathbb{Z}_p$, $S = \{x^2 : x \in \mathbb{Z}_p \setminus \{0\}\}$
- Lubotzky–Phillips–Sarnak - $G = PSL(2, q)$, $S = \ldots$

DIS $\Rightarrow$ EIG in Cayley graphs

Theorem (Kohayakawa–Rödl–Schacht, 2016+)

If G is an abelian group and $\mathrm{Cay}(G, S)$ satisfies

$$\left|e(X, Y) - \frac{d}{n}|X||Y|\right| \leq \epsilon dn$$

for all $X, Y \subseteq V(G)$, then all eigenvalues of $\mathrm{Cay}(G, S)$, save the largest, have absolute value $O(\epsilon d)$.

DIS ⇒ EIG in Cayley graphs

Theorem (Kohayakawa–Rödl–Schacht, 2016+)

If G is an abelian group and Cay(G, S) satisfies

$$|e(X, Y) - \frac{d}{n}|X||Y|| \leq \epsilon dn$$

for all $X, Y \subseteq V(G)$, then all eigenvalues of Cay(G, S), save the largest, have absolute value $O(\epsilon d)$.

Theorem (C.–Zhao, 2016+)

If G is any group and Cay(G, S) satisfies

$$|e(X, Y) - \frac{d}{n}|X||Y|| \leq \epsilon dn$$

for all $X, Y \subseteq V(G)$, then all eigenvalues of Cay(G, S), save the largest, have absolute value $O(\epsilon d)$.

DIS $\Rightarrow$ EIG in Cayley graphs

Theorem (Kohayakawa–Rödl–Schacht, 2016+)

If G is an abelian group and Cay(G, S) satisfies

$$\left|e(X, Y) - \frac{d}{n}|X||Y|\right| \leq \epsilon dn$$

for all $X, Y \subseteq V(G)$, then all eigenvalues of Cay(G, S), save the largest, have absolute value $O(\epsilon d)$.

Theorem (C.–Zhao, 2016+)

If G is any group and Cay(G, S) satisfies

$$\left|e(X, Y) - \frac{d}{n}|X||Y|\right| \leq \epsilon dn$$

for all $X, Y \subseteq V(G)$, then all eigenvalues of Cay(G, S), save the largest, have absolute value $O(\epsilon d)$.

The intermediate case

Jumbledness

A graph G on vertex set V is (p, β)-jumbled if, for all vertex subsets $X, Y \subseteq V(G)$,

$$|e(X, Y) - p|X||Y|| \leq \beta\sqrt{|X||Y|}.$$

The intermediate case

Jumbledness

A graph G on vertex set V is (p, β)-jumbled if, for all vertex subsets $X, Y \subseteq V(G)$,

$$|e(X, Y) - p|X||Y|| \leq \beta\sqrt{|X||Y|}.$$

Example - random graphs

Let $p = p(n) \leq 0.99$. Then, asymptotically almost surely, the binomial random graph $G(n, p)$ has the following property. For any two subsets $X, Y \subseteq V(G)$,

$$|e(X, Y) - p|X||Y|| = O(\sqrt{pn|X||Y|}).$$

The intermediate case

Jumbledness

A graph G on vertex set V is (p, β)-jumbled if, for all vertex subsets $X, Y \subseteq V(G)$,

$$|e(X, Y) - p|X||Y|| \leq \beta\sqrt{|X||Y|}.$$

Example - random graphs

Let $p = p(n) \leq 0.99$. Then, asymptotically almost surely, the binomial random graph $G(n, p)$ has the following property. For any two subsets $X, Y \subseteq V(G)$,

$$|e(X, Y) - p|X||Y|| = O(\sqrt{pn|X||Y|}).$$

This is close to best possible, in that we must have $\beta = \Omega(\sqrt{pn})$.

(n, d, λ)-graphs

(n, d, λ)-graphs

A graph G is said to be an (n, d, λ)-graph if it has n vertices, every vertex has degree d and, if $\lambda_1 \geq \lambda_2 \geq \cdots \geq \lambda_n$ are the eigenvalues of the adjacency matrix of G, $|\lambda_i| \leq \lambda$ for all $i \geq 2$.

(n, d, λ)-graphs

A graph G is said to be an (n, d, λ)-graph if it has n vertices, every vertex has degree d and, if $\lambda_1 \geq \lambda_2 \geq \cdots \geq \lambda_n$ are the eigenvalues of the adjacency matrix of G, $|\lambda_i| \leq \lambda$ for all $i \geq 2$.

By the expander mixing lemma,

$$\left| e(X, Y) - \frac{d|X||Y|}{n} \right| \leq \lambda \sqrt{|X||Y|},$$

so that (n, d, λ)-graphs are $(d/n, \lambda)$-jumbled.

Strongly regular graphs

Strongly regular graphs

A strongly regular graph $srg(n, d, \eta, \mu)$ is a d-regular graph on n vertices in which every pair of adjacent vertices have exactly η common neighbours and every pair of nonadjacent vertices have exactly μ common neighbours.

Strongly regular graphs

Strongly regular graphs

A strongly regular graph $srg(n, d, \eta, \mu)$ is a d-regular graph on n vertices in which every pair of adjacent vertices have exactly η common neighbours and every pair of nonadjacent vertices have exactly μ common neighbours.

Theorem

For $|\eta - \mu| = O(\sqrt{d})$, the strongly regular graph $srg(n, d, \eta, \mu)$ is an (n, d, λ)-graph with $\lambda = O(\sqrt{d})$.

Strongly regular graphs

Strongly regular graphs

A strongly regular graph $srg(n, d, \eta, \mu)$ is a d-regular graph on n vertices in which every pair of adjacent vertices have exactly η common neighbours and every pair of nonadjacent vertices have exactly μ common neighbours.

Theorem

For $|\eta - \mu| = O(\sqrt{d})$, the strongly regular graph $srg(n, d, \eta, \mu)$ is an (n, d, λ)-graph with $\lambda = O(\sqrt{d})$.

Paley graphs

Suppose $q \equiv 1 \pmod 4$ is prime. The Paley graph P_q is the graph with vertex set $\mathbb{Z}_q$, where x and y are joined if and only if $x - y$ is a quadratic residue. This graph is strongly regular, with parameters $(q, (q-1)/2, (q-5)/4, (q-1)/4)$, and so is an (n, d, λ)-graph with $n = q$, $d = (q-1)/2$ and $\lambda = O(\sqrt{q})$.

Further examples

Let q be a prime power and $PG(q,t)$ the projective space of dimension t, that is, each element is an equivalence class of non-zero vectors of length $t+1$ over the finite field of order q, where two vectors are taken as equivalent if one is a multiple of the other by an element in the field. This set has $n = (q^{t+1}-1)/(q-1)$ elements.

Further examples

Let q be a prime power and $PG(q, t)$ the projective space of dimension t, that is, each element is an equivalence class of non-zero vectors of length $t + 1$ over the finite field of order q, where two vectors are taken as equivalent if one is a multiple of the other by an element in the field. This set has $n = (q^{t+1} - 1)/(q - 1)$ elements.

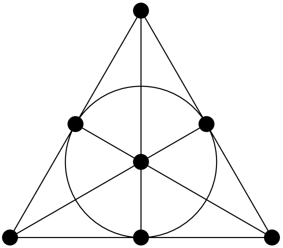

Further examples

We define a graph G whose vertices are the points of $PG(q,t)$ and where two vertices $x = (x_0, x_1, \ldots, x_t)$ and $y = (y_0, y_1, \ldots, y_t)$ are adjacent if and only if

$$x_0 y_0 + x_1 y_1 + \cdots + x_t y_t = 0.$$

Further examples

We define a graph G whose vertices are the points of $PG(q,t)$ and where two vertices $x = (x_0, x_1, \ldots, x_t)$ and $y = (y_0, y_1, \ldots, y_t)$ are adjacent if and only if

$$x_0 y_0 + x_1 y_1 + \cdots + x_t y_t = 0.$$

It is straightforward to check that G is d-regular with $d = (q^t - 1)/(q-1)$, though there may be some loops. However, there are $O(q^{t-1})$ vertices with loops. To calculate the eigenvalues of G, let I be the $n \times n$ identity matrix and J the $n \times n$ all-one matrix. Then

$$A^2 = \mu J + (d - \mu) I,$$

where $\mu = (q^{t-1} - 1)/(q-1)$. This easily implies that all eigenvalues apart from the largest have absolute value $\sqrt{d - \mu}$.

Further examples

Theorem (Alon, 1994)

There is a triangle-free (n, d, λ)-graph with $n = 2^{3k}$, $d = 2^{k-1}(2^{k-1} - 1)$ and $\lambda = O(2^k) = O(\sqrt{d})$.

Further examples

Theorem (Alon, 1994)

There is a triangle-free (n, d, λ)-graph with $n = 2^{3k}$, $d = 2^{k-1}(2^{k-1} - 1)$ and $\lambda = O(2^k) = O(\sqrt{d})$.

Note that this has $d = \Omega(n^{2/3})$, which is significantly denser than the triangle-free graphs obtained from modifications of random graphs, which have $d = \tilde{\Omega}(\sqrt{n})$.

Some properties

Lemma

The independence number $\alpha(G)$ of a (p, β)-jumbled graph G satisfies

$$\alpha(G) \leq \beta / p.$$

Some properties

Lemma

The independence number $\alpha(G)$ of a (p, β)-jumbled graph G satisfies

$$\alpha(G) \leq \beta/p.$$

To see this, note that if U is an independent set in G, then, by the definition of (p, β)-jumbledness,

$$p|U|^2 = |2e(U) - p|U|^2| \leq \beta|U|.$$

Lemma

The independence number $\alpha(G)$ of a (p, β)-jumbled graph G satisfies

$$\alpha(G) \leq \beta/p.$$

To see this, note that if U is an independent set in G, then, by the definition of (p, β)-jumbledness,

$$p|U|^2 = |2e(U) - p|U|^2| \leq \beta|U|.$$

Tightness

In the Paley graph P_q, where $q = p^2$ is the square of a prime, every element of the subfield $\mathrm{GF}(p)$ is a quadratic residue in $\mathrm{GF}(p^2)$ since $\mathrm{GF}(p^2)$ is a splitting field for $X^2 - a$ for all $a \in \mathrm{GF}(p)$. To form an independent set of order $\sqrt{q}$, consider $\beta\mathrm{GF}(p)$, where β is a quadratic non-residue.

Lemma

The chromatic number $\chi(G)$ of a (p, β)-jumbled graph G with n vertices satisfies

$$\chi(G) \geq pn/\beta.$$

Some properties

Lemma

The chromatic number $\chi(G)$ of a (p, β)-jumbled graph G with n vertices satisfies

$$\chi(G) \geq pn/\beta.$$

The Paley graph P_q with $q = p^2$ again shows that this is tight.

Some properties

Lemma

For any $\epsilon > 0$ and any graph H with maximum degree Δ, there exists $c > 0$ such that if G is a (p, β)-jumbled graph with $\beta \leq cp^{\Delta}n$, then the number of labelled copies of H in G is equal to

$$(1 \pm \epsilon)p^{e(H)}n^{v(H)}.$$

Some properties

Lemma

For any $\epsilon > 0$ and any graph H with maximum degree Δ, there exists $c > 0$ such that if G is a (p, β)-jumbled graph with $\beta \leq cp^{\Delta}n$, then the number of labelled copies of H in G is equal to

$$(1 \pm \epsilon)p^{e(H)}n^{v(H)}.$$

For triangles, Alon's example shows that this is tight.

An application

Problem

What is the MAXCUT of an m-edge triangle-free graph?

An application

Problem

What is the MAXCUT of an m-edge triangle-free graph?

For an ordinary graph with m edges, MAXCUT is $m/2 + \Omega(\sqrt{m})$

An application

Problem

What is the MAXCUT of an m-edge triangle-free graph?

For an ordinary graph with m edges, MAXCUT is $m/2 + \Omega(\sqrt{m})$

Theorem (Alon, 1996)

The MAXCUT of an m-edge triangle-free graph is $m/2 + \Omega(m^{4/5})$.

An application

Problem

What is the MAXCUT of an m-edge triangle-free graph?

For an ordinary graph with m edges, MAXCUT is $m/2 + \Omega(\sqrt{m})$

Theorem (Alon, 1996)

The MAXCUT of an m-edge triangle-free graph is $m/2 + \Omega(m^{4/5})$.

Tightness

The MAXCUT of an (n, d, λ)-graph with $m = dn/2$ edges is at most $m/2 - \lambda_n n/4$. In particular, the MAXCUT of Alon's example is $m/2 + O(m^{4/5})$.

Thank you for listening!