Summer '05

LIFTS OF GRAPHS

NATI LINIAL

- Hebrew University
- Jerusalem, Israel
- http://www.cs.huji.ac.il/~nati/
- Based on joint papers with: Alon Amit, Yonatan Bilu, Yotam Drier, Jirka Matousek, and Eyal Rozenman.

Fahrplan (don't worry about the German, The rest of the talk is in Hebrew...)

We have several territories to visit:

- Random lifts
- Topological considerations
- Extremal problems
- computational aspects

Let's start with $G(n, p)$

The theory of random graphs is one of the greatest achievements of modern combinatorics. The most developed (nearly complete?) part of this theory is the Erdős-Rényi $G(n, p)$ model.
In this model you start with n vertices. Each pair of vertices x, y are joined by an edge with probability p. These choices are done independently for each pair of vertices.
Such graphs have been used in a wide spectrum of contexts.

Random graphs in $G(n, p)$...

- ... have taught us that small graphs that we can draw on paper and inspect with our eyes often give you the wrong impression about the behavior of large graphs. (E.g. the existence of graphs with high chromatic number and no short cycles).
- ... are used to construct graphs with desired extremal properties. (E.g. Ramsey Graphs).
- ... can model various natural phenomena (E.g. in statistical mechanics).
- ... are very simple and flexible model to work with.

But (on the mathematical side)....

There are many things that we may want from random graphs which $G(n, p)$ cannot do for us:

- These graphs are almost never regular, so we need a theory of random regular graphs. Indeed such a theory exists, and is still being developed.
- We sometimes want even higher degrees of regularity. Rudiments of a theory of random Cayley graphs already exist, but much remains to be studied.
- What if we want to understand random higher dimensional complexes, random knots, random maps....?

Wish list (cont.)

There are still quite a few graph theoretic problems where random constructions based on $G(n, p)$ do not provide the answer:

- Graphs of maximum girth.
- Optimal error-correcting codes.
- Tight estimates for Ramsey numbers.

Some computational desiderata

Even though $G(n, p)$ is a versatile and useful tool in many computational problems, it does leave a lot to be desired.

- For computational purposes we often want a model that requires fewer random bits. Many questions arise about explicit constructions of graphs.
- There is a very tight connection between the P *vs.* NP Problem and questions about explicit constructions. Any reduction in randomness is interesting.

Some computational desiderata (contd.)

- There are important naturally-occurring random graphs that $G(n, p)$ certainly does not capture well: The graph of the Internet; Graphs of social connections; Biological molecular pathways etc.
- More generally - We often need models of random graphs which we can better design for the purpose at hand.

Covering maps - a bit of background

- A fundamental object of study in topology.
- The most famous case - covering map $x \to e^{ix}$ from the real line to the unit circle.
- But a graph is also a topological space (a one-dimensional simplicial complex), so covering maps can be defined and studied for graphs.

Definition 1. A map $\varphi : V(H) \to V(G)$ where G, H are graphs is a covering map if for every $x \in V(H)$, the neighbor set $\Gamma_H(x)$ is mapped $1:1$ onto $\Gamma_G(\varphi(x))$.
When such a mapping exists, we say that H is a lift of G.

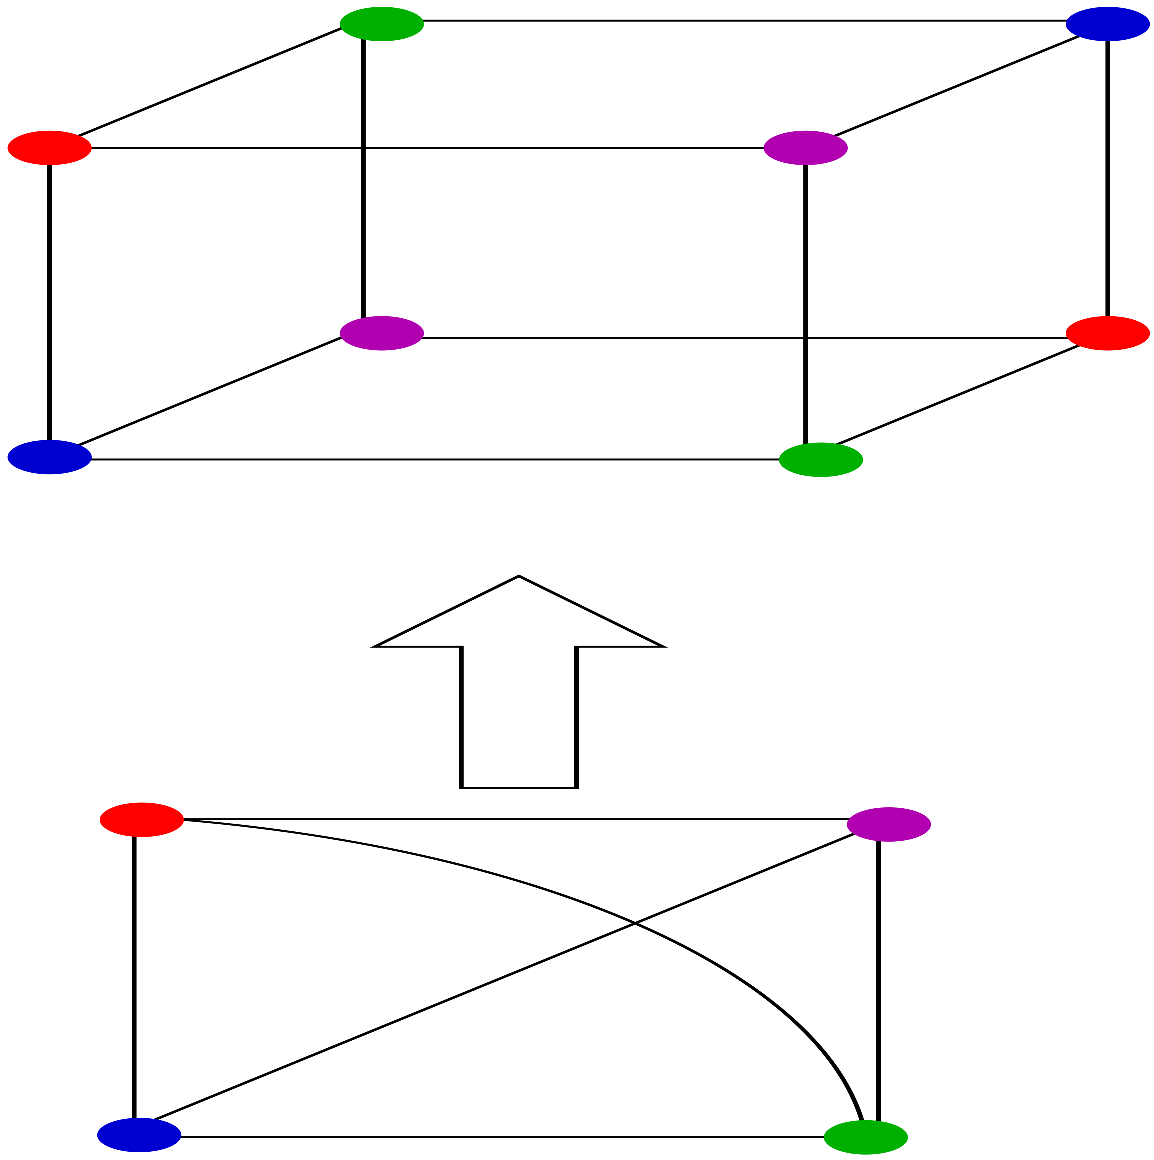

Figure 1: The 3-dimensional cube is a 2-lift of K_4

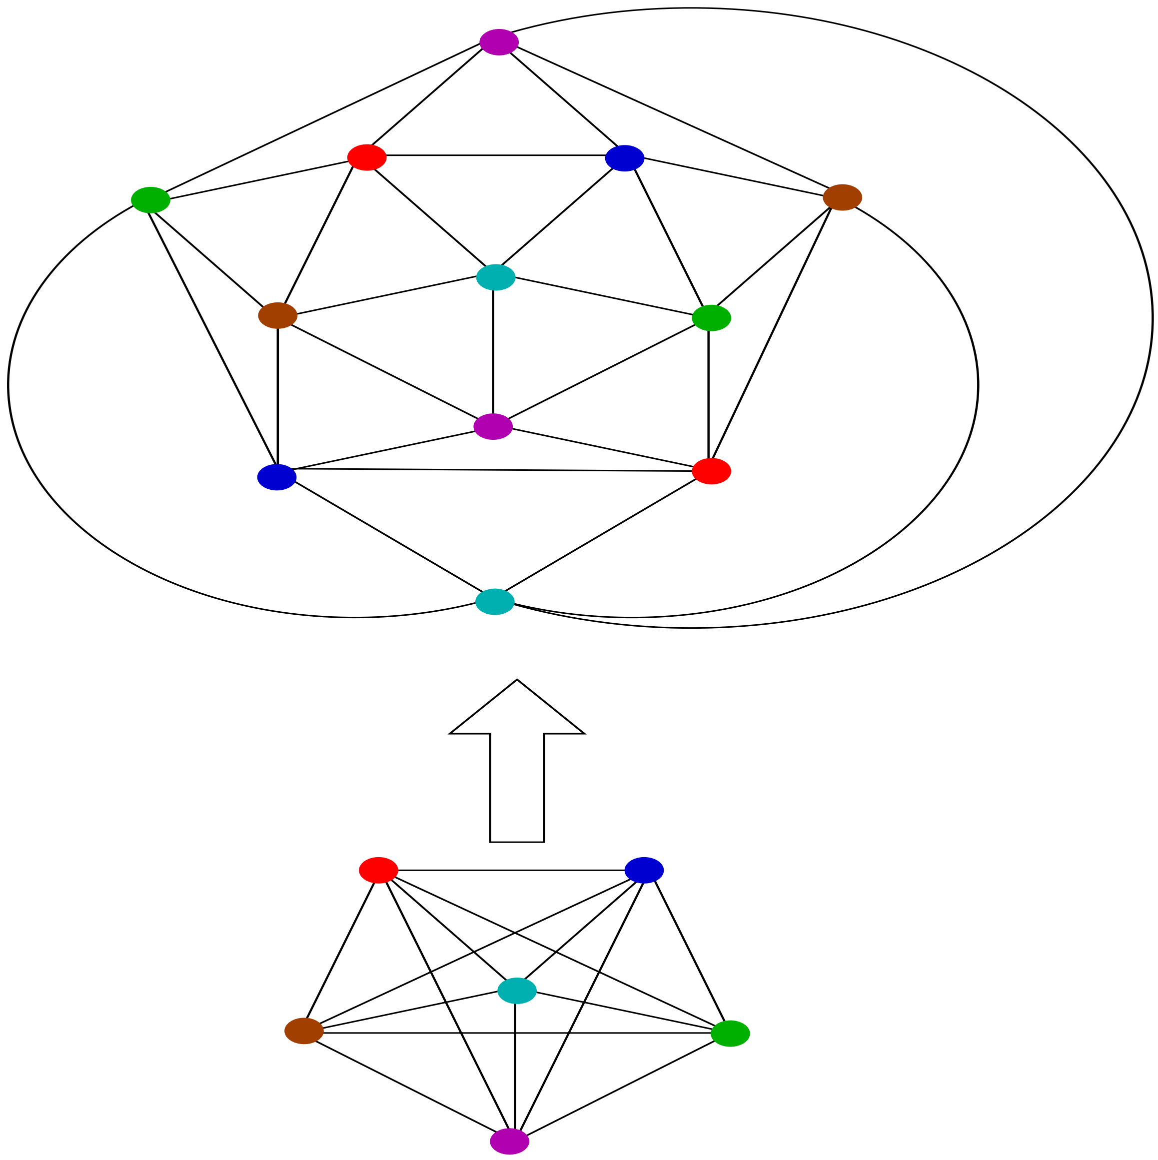

Figure 2: The icosahedron graph is a 2-lift of K_6

How should we think of it in concrete terms?

We see in the previous examples that the covering map φ is $2:1$.

- The 3-cube is a 2-lift of K_4.
- The graph of the icosahedron is a 2-lift of K_6.

In general, if G is a connected graph, then every covering map $\varphi : V(H) \to V(G)$ is $n:1$ for some integer n (easy).

Fold numbers etc.

- We call n the fold number of φ.
- We say that H is an n-lift of G.
- Sometime we say that H is an n-cover of G.

A direct, constructive perspective

The set of those graphs that are n-lifts of G is called $L_n(G)$.

- Every $H \in L_n(G)$, has vertex set $V(H) = V(G) \times [n]$.
- We call the set $F_x = \{x\} \times [n]$ the fiber over x.
- For every edge $e = xy \in E(G)$ we have to select some perfect matching between the fibers F_x and F_y, i.e., a permutation $\pi = \pi_e \in S_n$ and connect (x, i) with $(y, \pi(i))$ for $i = 1, \ldots, n$.
- This set of edges is denoted by F_e, the fiber over e.
- We refer to G as our base graph.

Random lifts - A new versatile class of random graphs

- When the permutations π_e are selected at random, we call the resulting graph a random n-lift of G.

- This allows us to "design" random graphs to suit our purposes.

We will see (quite a bit) more about such graphs below.

What do we want to know about these graphs?

- Use them for to solve problems in extremal graph theory.
- In random lifts - Determine the lift's typical properties and how they reflect the properties of the base graph.
- Study graph-algorithmic problems for such graphs.
- Make connections with topology (random higher-dimensional complexes etc.)

Let's get down to business...

We will always assume (unless otherwise stated):
Convention 1. The base graph G is connected.

Let H be an n-lift of G. We say that an edge $e = xy \in E(G)$ is flat if the corresponding permutation $\pi_e \in S_n$ is the identity. In other words, the edges between the fibers F_x and F_y connect (x, i) to (y, i) for every i.
We note that by renaming vertices, if necessary, we may assume

Convention 2. Let G be a graph and let T be a spanning tree of G. When we deal with lifts of G, we may assume without loss of generality that all edges of T are flat.

So what do we know about random lifts? A warm-up exercise.

Perhaps the simplest graph property is connectedness. We will henceforth always assume that our base graph G is connected, and ask:

Question 1. *For which (connected) graphs G is it the case that graphs $H \in L_n(G)$ are likely to be connected?*

When is a lift connected? (cont.)

Case $e_G - v_G = -1$: When G is a tree, all n-lifts of G are trivial, and every $H \in L_n(H)$ simply consists of n disjoint copies of G.
To conclude: No n-lift of G is connected.

Case $e_G - v_G = 0$: G is monocylic. By the above remarks it's enough to consider the case where $G = C$ is a cycle. All edges in C but one e are flat. The lift is connected iff π_e is a cyclic permutation. This happens with probability $\frac{1}{n} = o(1)$.
To conclude: Almost no n-lift of G is connected.

Case $e_G - v_G > 0$: Standard random-graphs arguments can be used to conclude: Almost every n-lift of G is connected.

The degree of connectivity

What is G's degree of connectivity? I.e. What is the least number of vertices whose removal makes G disconnected? We need some simple facts:

1. Lifts preserve vertex degrees. If $x \in V(G)$ has d neighbors in G, then so is the case with all vertices in the fiber F_x.

2. The degree of connectivity of a graph cannot exceed the smallest vertex degree in the same graph.

Consequently,
Observation 3. If $\delta = \delta(G)$ is the smallest vertex degree in G, then every graph in $L_n(G)$ is at most δ-connected.

The degree of connectivity

Theorem 1 (Amit, Linial). *If $\delta = \delta(G)$ is the smallest vertex degree in G and if $\delta \geq 3$, then almost every graph in $L_n(G)$ is δ-connected. (And, of course, each of these graphs is at most δ-connected.)*

A word on asymptotics: I am using here (and below) a somewhat sloppy asymptotic terminology. For those not accustomed to this terminology, here is the same theorem stated accurately:

Theorem 2. *Let G be a finite connected graph and let $\delta = \delta(G)$ be the smallest vertex degree in G. If $\delta \geq 3$, then a graph H drawn at random from $L_n(G)$ is δ-connected with probability $1-\epsilon_n$ where $\epsilon_n \to 0$ as $n \to \infty$.*

Perfect matchings in lifts - I

For a given G we consider the typical size of the largest matching in graphs $H \in L_n(G)$. In particular, does H tend to have a perfect matching? For simplicity we assume for this problem that n is even.
First easy case: If G has a perfect matching (=PM) M, then every lift of

G has a PM - Just lift M.

Perfect matchings in lifts - II

If a lift $H \in L_n(G)$ contains a perfect matching P, let

$$f(e) = |P \cap F_e|/n$$

be the fraction of the n edges in the fiber F_e that belong to the perfect matching P.
What are the properties of this f?

Fractional perfect matchings

- $f : E(G) \to \mathbb{R}$.
- $f(e) \geq 0 \quad \forall e \in E(G)$.
- $\sum_{e \ni x} f(e) = 1 \quad \forall x \in V(G)$.

Such an f is called a fractional perfect matching (=fpm).
Second easy case: If G has no fpm, then in every lift of G, every matching misses at least a fraction $\lambda = \lambda_G > 0$ of the vertices.

Perfect matchings in lifts - III

An edge $e \in E(G)$ is called essential if $f(e) > 0$ for some fpm f. Clearly we can eliminate all inessential edges from G and consider each connected component in the reaming graph separately.
We only need to analyze the class $\mathcal{G}$ of graphs that are:

- Connected.
- Have no perfect matching.
- Have a fractional perfect matching.
- Have no inessential edges.

Perfect matchings in lifts - IV

Third case: It can be shown (and this is not easy) that every graph $K \in \mathcal{G}$ is either

1. An odd cycle and then in almost every n-lift of the graph the largest matching misses $\Theta(\log n)$ vertices.

2. For every other graph $K \in \mathcal{G}$, almost every lift of K has a perfect matching.

Zero-one laws

Theorem 3 (Amit, Linial). *For every graph G, almost all lifts of G have the same degree of connectivity.*
Theorem 4 (Linial, Rozenman). *For every graph G either:*

- *Almost every lift of G contains a perfect matching, or*
- *Almost no lift of G contains a perfect matching.*

Question 2. *Which graph properties exhibit this kind of zero/one laws?*

Note: Not all graph properties behave this way, e.g. containing a triangle.

...and speaking of open problems...

There are lots of open problems about random lifts. For example, for which graphs G do the lifts of G tend to have Hamiltonian Circuit? Perhaps the same conditions that guarantee the almost sure existence of a perfect matching in fact yield this stronger conclusion. Specifically:

Problem 1. *Is there a zero-one law for Hamiltonicity? Namely, is it true that for every G almost every or almost none of the graphs in $L_n(G)$ have a Hamiltonian Circuit?*
Problem 2. *Let G be a d-regular graph with $d \geq 3$. Is it true that almost every graph in $L_n(G)$ has a Hamiltonian Circuit?*

There are some preliminary results this direction by Frieze and his students.

A few words about chromatic numbers

It is easy to see that if H is a lift of G, then

$$\chi(H) \leq \chi(G)$$

(just lift any coloring of G to a coloring of H).
Some reverse inequalities are known (and are quite surprising, I think).

Chromatic numbers (cont.)

Theorem 5 (Amit, Linial, Matousek). *For almost every* $H \in L_n(G)$*,*

- $\chi(H) \geq \Omega\left(\sqrt{\frac{\chi(G)}{\log \chi(G)}}\right)$
- $\chi(H) \geq \Omega\left(\frac{\chi_f(G)}{\log^2 \chi_f(G)}\right)$

where $\chi_f(G)$ *is* G*'s fractional chromatic number.*

I tend to suspect that, in fact:

Conjecture 1.

$$\chi(H) \geq \Omega\left(\frac{\chi(G)}{\log \chi(G)}\right)$$

for almost every $H \in L_n(G)$*.*

Some lift-free comments on coloring

Given a graph G, let

$$L(G) = \min\left\{k \,\middle|\, \exists\, S_1, \ldots, S_k \subseteq V(G) \text{ such that } \sum_{i: x\in S_i} \frac{1}{\mathsf{dgn}(S_i)+1} \geq 1 \text{ for every } x \in V(G)\right\}. \quad (1)$$

- $\chi(G) \geq L(G) \geq \sqrt{\chi(G)/2}$ (Easy/not hard).
- Pyatkin found a (fairly complicated) example of a graph with

$$L(G) = 3 < \chi(G) = 4$$

.

Problem 3. *Is the lower bound on $L(G)$ asymptotically tight?*

What are lifts good for?

The main challenge is to turn lifts and random lifts into tools in the study of questions in computational complexity and discrete mathematics. Lifts can be used to construct regular graphs with large spectral gap. This property implies that the graph is a good expander.

A quick review on expansion, spectral gap etc.

A graph $G = (V, E)$ is said to be ϵ-edge-expanding if for every partition of the vertex set V into X and $X^c = V \setminus X$, where X contains at most a half of the vertices, the number of cross edges

$$e(X, X^c) \geq \epsilon |X|.$$

In words: in every cut in G, the number of cut edges is at least proportionate to the size of the smaller side.

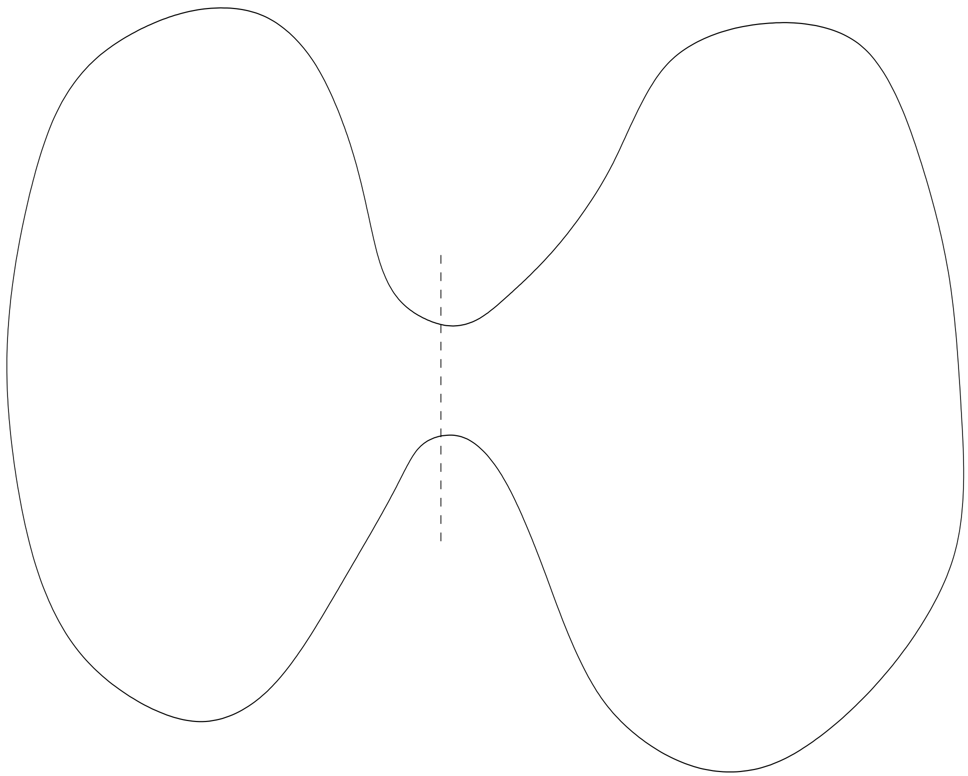

Figure 3: Expansion means - No bottlenecks in the graph

Expansion and spectral gap

Associated with every graph G is its adjacency matrix $A = A_G$ where $a_{ij} = 1, 0$ according to whether vertex i and vertex j are adjacent or not.
Let $\lambda_1 \geq \lambda_2 \geq \ldots \lambda_n$ be the eigenvalues of A (they are real since A is real symmetric).
If G is d-regular, then $\lambda_1 = d$.
It is known that if $\lambda_2 << d$ then G has large expansion. "A large spectral gap implies high expansion".
This is easy to prove and the quantitative statement is very tight.
It is also known that Expansion implies a spectral gap. This is harder to prove, and the bounds that can be proved on this are rather weak.

What's a "large" spectral gap?

How small can λ_2 be in a d-regular graph? (i.e., how large can the spectral gap get)?
This was answered as follows:

Theorem 6 (Alon, Boppana).

$$\lambda_2 \geq 2\sqrt{d-1} - o(1)$$

Here the $o(1)$ is a quantity that goes to zero as the graph become big. Note that the term is negative. Indeed, small Ramanujan Graphs are plenty and easy to construct. The difficulty is to find, for d fixed arbitrarily large Ramanujan Graphs.

Large spectral gaps and the number $2\sqrt{d-1}$

A good approach to problems in extremal combinatorics is to guess where the extremum is attained (the ideal example), and show that there are no better instances.
What, then, is the ideal expander?
Funny answer: It's the infinite d-regular tree.
In fact, something like eigenvalues (spectrum) can also be defined for infinite graphs. It turns out that the supremum of the spectrum for the d-regular infinite tree is....

$$2\sqrt{d-1}$$

In fact...

It's an easy fact that the infinite d-regular tree is the universal covering space of any d-regular graph.

...and so we ask

Problem 4. *Are there d-regular graphs with second eigenvalue*

$$\lambda_2 \leq 2\sqrt{d-1} \quad ?$$

When such graphs exist, they are called Ramanujan Graphs.

Some of what we know about Ramanujan Graphs

Margulis; Lubotzky-Phillips-Sarnak: d-regular Ramanujan Graphs exist when $d-1$ is a prime power. The construction is easy, but the proof uses a lot of heavy mathematical machinery.

Friedman: If you are willing to settle for $\lambda_2 \leq 2\sqrt{d-1}+\epsilon$, they exist. Moreover, almost every d-regular graph satisfies this condition.

Experimentally, Novikov-Sarnak: For n large, about 52 percent of all 3-regular graphs are Ramanujan. (No proof for this in sight...)

...and many things that we do not know...

- Do d-regular Ramanujan Graph exist for every $d \geq 3$? (Currently known only for $d = p^{\alpha} + 1$).
- Can they be constructed using combinatorial/probabilistic methods?
- Can their existence be established in a reasonably short manner?

Large spectral gap through 2-lifts

Basic idea: Start with a small d-regular Ramanujan Graph (e.g. the complete graph K_{d+1}) and proceed to create large Ramanujan graphs by applying a sequence of 2-lifts.
Why should this work? If H is a 2-lift of G, What is the connection between the eigenvalues of G and those of H?
It is not hard to see that every eigenvalue of G is also an eigenvalue of H (simply reproduce the same eigenvector on each of H's two "levels").
OK, so if G has n vertices, H should have $2n$ eigenvalues. There are n eigenvalues inherited from G, so what are the n new eigenvalues?

The New Eigenvalues

This question has a simple and satisfying answer.
There is a very natural encoding for G's 2-lifts using its adjacency matrix A_G. Recall that π_e is the identity we say that e is flat. Otherwise (when $\pi_e = (12)$) we say that e is switched. We encode the lift through a signing of the matrix A_G. It's the matrix S where $s_{ij} = a_{ij}$ for all i, j, except that $s_{ij} = -1$ (whereas $a_{ij} = 1$) if $e = ij$ is a switched edge.
Claim 4. *The n new eigenvalues of a 2-lift of G are exactly the eigenvalues of the signed matrix S.*

(Pretty easy to prove: If $Sx = \lambda x$, associate x with the lower level of H and $-x$ with the upper level).

Basic idea (cont.)

Start with a small d-regular Ramanujan graph. At each step apply a 2-lift such that the eigenvalues of the signed matrix S all fall into the interval

$$[-2\sqrt{d-1}, 2\sqrt{d-1}]$$

It is clear that if we can prove the following conjecture, then we can achieve our goals about constructing Ramanujan Graphs.

Conjecture 2. *For every d-regular Ramanujan Graph G there is a signing S of A_G such that all the eigenvalues of S are in $[-2\sqrt{d-1}, 2\sqrt{d-1}]$.*

We even dare and state:

Conjecture 3. *For every d-regular graph G there is a signing S of A_G such that all the eigenvalues of S are in $[-2\sqrt{d-1}, 2\sqrt{d-1}]$.*

What do we know about this conjecture?

- It cannot be improved in general. It is tight for graphs in which at least one vertex does not belong to any cycle of bounded length.
- There are interesting infinite families of graphs (e.g. planar grid) where it holds with an even better bound.
- Substantial numerical evidence.
- Just random signing is not enough.
- Experimentally, at least, fairly simple heuristics can find such a signing.

The main result about the signing conjecture

Theorem 7 (Bilu, Linial). *For every d-regular graph there is a signing of* A_G *with spectral radius* $O(\sqrt{d} \cdot \log^{3/2} d)]$*. Moreover, such a signing can be found efficiently.*

A few words about the proof

It uses a probabilistic argument (Local Lemma) followed by derandomization. Another important ingredient is a lemma showing that Discrepancy and spectral gap are essentially equivalent.
Recall what we said before:

- *A large spectral gap implies high expansion*. This is easy to prove and the quantitative statement is very tight.
- It is also known that *expansion implies a spectral gap*. This is harder to prove, and the bounds that can be proved on this are rather weak.

Capturing the linear-algebraic condition with combinatorics

In view of this state of affairs, it is natural to ask:

Problem 5. *Are there combinatorial conditions that capture the spectral gap more accurately?*

The Expander Mixing Lemma and its (new) Converse

Here is the most common way expansion is being used:

Theorem 8 (The expander mixing lemma). *Let $G = (V, E)$ be a d-regular graph on n vertices, and let $\lambda = \lambda_2(G)$ be the second largest eigenvalue of (the adjacency matrix of) G.*
Then for any two disjoint sets of vertices $X, Y \subseteq V$, the number of edges connecting X and Y satisfies

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

Discrepancy and spectral gap are essentially equivalent

Theorem 9 (Bilu, Linial). *Let $G = (V, E)$ be a d-regular graph on n vertices and suppose that for any two disjoint sets of vertices $X, Y \subseteq V$, the number of edges connecting X and Y satisfies*

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

for some $\alpha > 0$. Then the second eigenvalue of G is at most $O(\alpha \log(\frac{d}{\alpha}))$. The bound is tight

Two EXTREMAL questions on lifts - Minors

Recall: The infinite d-regular tree $\mathbf{T}_{\mathrm{d}}$ is the ideal expander. We seek finite graphs with similar combinatorial and spectral properties.
Likewise, $\mathbf{T}_{\mathrm{d}}$ has no nontrivial minors. How well can this be emulated in finite graphs?
Given a finite G, which of its minors M is persistent? I.e., M is a minor of *every* lift of G?
Open Problem 1. *Are there lifts of* K_n *with Hadwiger number* $o(n)$*?*

The other EXTREMAL question on lifts - Large girth

Just a word on this fascinating subject:
A "joke" proof that for every d and g there exists a d-regular graph Γ of girth $\geq g$ (Γ has no cycles shorter than g).
Let G have girth γ and let k be the number of γ-cycles in G. Consider H, a random 2-lift of G and let X be the r.v. that counts the number of γ-cycles in H.

The obligatory proof...

- The girth of H is $\geq \gamma$. Moreover,
- a γ-cycle in H must be a lift of a γ-cycle in G.
- A γ-cycle in G yields
  1. A 2γ-cycle in H, (with probability $1/2$), or
  2. Two γ-cycles in H, (with probability $1/2$).
- Consequently, $EX = k$.

Now forbid the identity lift to conclude that there is a 2-lift H with fewer than k γ-cycles.

A role for lifts in computational complexity

Definition 2. A section in a lift $H \in L_n(G)$ is a subset $S \subseteq V(H)$ such that S contains exactly one vertex in every fiber ($\forall x \in V(G) \ \ |S \cap F_x| = 1$).

A recent and most interesting problem (originally stated in different terms):

Conjecture 4 (Khot's unique games conjecture). *The following problem in NP-hard:*
Input: *A lifted graph $H \in L_n(G)$.*
Output: *A section S with as many edges as possible ($\max_S |E_H(S)|$).*

The unique game conjecture

Clearly, $|E_H(S)| \leq e_G$ for every $H \in L_n(G)$ and every section S in H. An even stronger conjecture states:

Conjecture 5 (Khot). *The following problem in NP-hard:*
Input: *A lifted graph $H \in L_n(G)$.*
Output: *To decide whether*

- *There exists section S with $|E_H(S)| \geq (1 - \epsilon) \cdot e_G$ or*
- *$|E_H(S)| \leq \epsilon \cdot e_G$ for every section.*

The unique game conjecture (cont.)

It is known that the unique games conjecture implies many interesting corollaries. In particular, many famous problems can be shown to be hard to approximate.
On the other hand, there is also some evidence that this computational problem may not be so difficult (work of Trevisan and Charikar-Makarychev-Makarychev).