EXPANDER GRAPHS - ARE THERE ANY MYSTERIES LEFT?

NATI LINIAL

- Hebrew University
- Jerusalem, Israel
- http://www.cs.huji.ac.il/~nati/
- See survey w/Hoory & Wigderson

Expansion -
A relative notion of connectivity

- Graph connectivity is a basic concept in graph theory that is usually taught in your discrete math and algorithms classes.
- Namely, how many vertices/edges must be removed to make the graph disconnected. (How small can S be so that $G \setminus S$ is disconnected).
- However, in graph connectivity you do not care whether the connected components of the graph $G \setminus S$ are small or large.
- In contrast, G may still have large expansion even though there is a small set S for which $G \setminus S$ is disconnected. That is, if $G \setminus S$ consists of one very large component and several small ones.

- We only forbid bottlenecks. Concretely, we rule out the possibility that there is a small set S for which $G \setminus S$ has at least two large connected components.

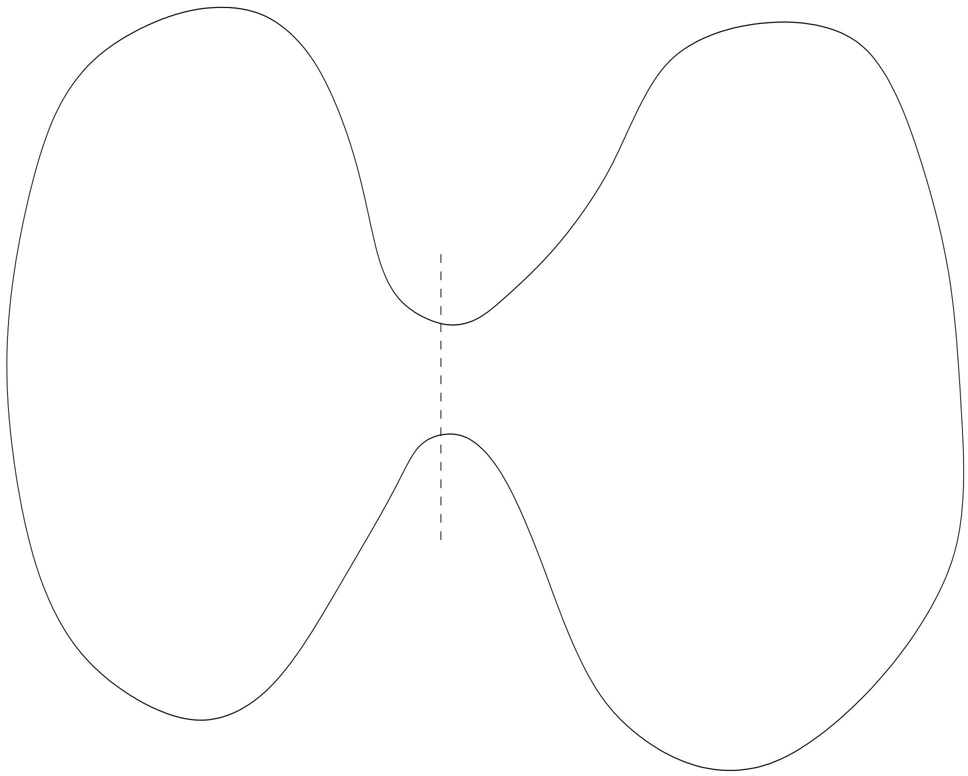

Some basic questions

- What are expanders good for?
- Do they exist in abundance, or are they rare?
- How can you tell if a given graph is an expander?
- Can you explicitly construct them?
- How expanding can you get? (and does it matter)?

The formal definition

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.

Do expanders actually exist?

It should be quite clear that for every finite connected graph G there is an $\epsilon > 0$, such that G is ϵ-edge-expanding. What is not obvious is whether we can keep ϵ bounded away from zero when G gets large.

This, too, is not hard to do, provided that G has a lot of edges. So, the essence of the problem is whether we can keep ϵ bounded away from zero, while the vertex degrees remain bounded from above by a constant. (e.g. every vertex has exactly d neighbors, in which case we say that G is d-regular.)

Do expanders actually exist? (cont'd.)

Here is our first concrete nontrivial question about expanders: Do there exist an integer d and $\epsilon > 0$, and infinitely many graphs $G_n = (V_n, E_n)$ so that:

- G_n get arbitrarily large ($|V_n| \to \infty$).
- All graphs G_n are d-regular (every vertex has exactly d neighbors).
- All the G_n are ϵ-edge-expanding.

ANSWER: Yes. This can be shown using the probabilistic method.

Do expanders actually exist? (conclusion)

- They do. In fact, they exist in abundance. They are the rule rather than the exception.
- The probabilistic argument not only shows that a regular graph can be an expander. Rather, that almost every regular graph is an expander.
- If so, why don't we see them all over the place?
- Because we are too shortsighted and because our computational resources are too limited.

The computational hardness of expansion

Theorem 1. *The following computational problem is co-NP hard*

Input: *A graph* G, *and* $\epsilon > 0$.

Output: *Is* G *an* ϵ*-edge-expander?*

The probabilistic method as an observational aid

- We are incapable of getting a good view of large graphs. The probabilistic method is our major tool for overcoming this deficiency.

- It works particularly well when (as is often the case) a property that we seek is, in fact, very common.

- The probabilistic method can help us detect such situations.
  More refined applications of the method (e.g. The Lovász Local Lemma) can even help you find needles in haystacks.

But this seems rather bad....

- In the present case, we are able to show that expanding graphs are all around us, but it's quite nontrivial to construct them.
(This is the problem of finding hay in a haystack...)

- What's worse, we cannot even recognize such a graph (even if it falls on our head....)

- ... and besides: why do we really care? Let's address this first.

What are expander graphs good for? I

They yield good error-correcting codes. A bipartite graph $H = (A, B, E)$ that's a good expander, yields a good error-correcting code that comes with efficient encoding and decoding algorithms. Roughly, if A has n vertices, there is a $1:1$ correspondence between subsets of A and n-bit messages. The vertices in B correspond to parity check bits. It turns out that if H is a good expander, then not only is the resulting (linear) code good, it can be provably corrected by a simple belief-propagation scheme applied to H.

What are expander graphs good for? II

They offer good deterministic emulation of random behavior. Or, as we often say, they exhibit certain desired pseudo random properties. More on this - later.

What are expander graphs good for? III

A bridge between computer science and mathematics: The explicit construction of expander graphs helps recruit a lot of deep mathematics into the area. Expander graphs were recently used in resolving several old problems in mathematics....

How to save on random bits with expanders

Suppose that you have a randomized algorithm that uses k random bits and succeeds with probability $\frac{3}{4}$. If you can tolerate only error rates $< \delta$. What can you do to reach such a low probability of error?

The natural answer would be to repeat the experiment enough times to bring the probability of failure below δ. To carry out each single run of the test you must generate k random bits for a total of $k \log \frac{1}{\delta}$ random bits. So in this straightforward solution, you will generate k random bits each time you want to reduce the error probability by a factor of $\frac{1}{4}$.

Seems unavoidable, right?

Well, expander graphs can often be used to do things that seem "too good to be true", and this is one of them.

Randomized decision algorithms - A quick refresher

A language $\mathcal{L}$ is a collection of strings. In the decision problem for $\mathcal{L}$, we are given a string x and we have to determine whether x belongs to $\mathcal{L}$ or not.

What is a randomized algorithm that uses k random bits and has success probability $\geq \frac{3}{4}$? This is a (computationally feasible) boolean function $f(\cdot,\cdot)$. The first variable is x. The second one is r, a k-bit string. When $f(x,r) = 1$ we know with certainty that x is in $\mathcal{L}$. However, $f(x,r) = 0$ only means that x is apparently not in $\mathcal{L}$.

Randomized decision algorithms (cont'd.)

So, if we ever find an r for which $f(x, r) = 1$, our search is over, and we know that x is in $\mathcal{L}$. However, if we repeatedly sample r's with $f(x, r) = 0$, we are collecting more and more statistical evidence, that apparently x is not in $\mathcal{L}$.
In other words, if x is not in $\mathcal{L}$, then $f(x, r) = 0$ for every binary string r. However, if x is in $\mathcal{L}$, then at least $\frac{3}{4}$ of all strings r are witness to that. Namely, if $x \in \mathcal{L}$, then there is a subset W comprising at least $\frac{3}{4}$ of the 2^k binary strings r of length k such that $f(x, r) = 1$.

The trouble is, we only know that W is large, but have no idea how to search for its members.

A classical example of a randomized decision algorithms Primality testing.

Here $\mathcal{L}$ is the set of all composite (=non-prime) integers.
Fermat's Little Theorem says that if n is prime and $1 \leq a < n$, then $a^{n-1} - 1$ is divisible by n.
So if n is your input and if you found an integer a for which $a^{n-1} - 1$ is indivisible by n, then a is a witness to the fact that n is composite.

(Warning: Here I am oversimplifying things a bit....)
Usually, if n is composite, then at least a half of the integers $1 \leq a < n$ are witnesses to that.

Good. There are lots of witnesses, but where are they and how will we find even one?

How to hit a hiding elephant?

In some huge universe (the collection of all 2^k k-bit strings; the range $1 \leq a < n$) at least $\frac{3}{4}$ of the elements are good. We need to find at least one good element, but we must succeed with overwhelming probability. If we repeatedly sample the space, say m times, then our probability of failure will indeed be very low, only $(\frac{1}{4})^m$. However, to do this, we will have to generate many random bits at each of the m times that we run our statistical test. It should seem quite convincing (but this impression is false...) that this cannot be achieved with fewer random bits.

How to hit a hiding elephant? (cont'd.)

Let us reformulate the expansion property for a bipartite $G = (L, R, E)$ graph as follows. Fix any set $B \subseteq R$ of "bad" vertices on the right side of G. If B is not too big relative to R, then the following set $S \subseteq L$ is tiny. Namely, S is the set of those vertices $x \in L$ on the left side all of whose neighbors belong to B.

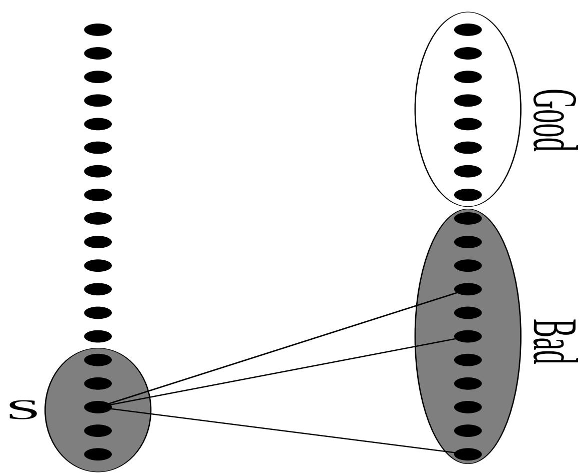

Hiding elephants III

So here is an efficient way of finding at least one member in W with overwhelming probability. We use a bipartite expander $H = (L, R, E)$ where R is identified with the set of random strings r.
The algorithm is extremely simple: We sample a random vertex v in L (for this we need to generate only $\log|L|$ random bits) and go over the list of x's neighbors in R. Remember that these vertices correspond to some strings r. We run the randomized algorithm with these strings.
These samples are NOT independent, and still this test is very powerful.

Hiding elephants IV

Since W is large, its complement $B := R \setminus W$ is not too large. As we said above, there is only a tiny set S of vertices v which would fail us. Namely, only those all neighbors of which are in the bad set B. With proper choice of parameters we can reduce the probability of failure to the desired level without having to generate lots of random bits.

Hiding elephants V

- Even better error reduction can be achieved through a short random walk on an expander graph. This even works for randomized algorithms with a two sided error.

- To select the next step in a random walk on a d-regular graph, only $\log d$ random bits are required.

- Unlike the previous construction, and in order to counter the possibility of a two-sided error, we make our decision by taking a majority vote among the different runs of the algorithm.

The algebraic perspective

A standard way of representing a graph G is through its adjacency matrix $A = A_G$. Namely, $a_{ij} = 1, 0$ according to whether the i-th and j-th vertex are adjacent or not.

Can we use the eigenvectors, eigenvalues of A to study G's expansion properties?

Reminder - Eigenvalues, Rayleigh quotients etc.

Since the matrix A is a real and symmetric, all its eigenvalues $\lambda_1 \geq \lambda_2 \geq \ldots \geq \lambda_n$ are real.
If A is a matrix of a d-regular graph, then $\lambda_1 = d$, and the corresponding eigenvector is $\mathbf{1}$, the all-ones vector.

$$\lambda_2 = \max \frac{x^t A x}{\|x\|^2}$$

where the maximum is over vectors x that are perpendicular to $\mathbf{1}$. (This expression is called a Rayleigh quotient).

What does this have to do with expansion?

Let $G = (V, E)$ be an n-vertex d-regular graph, and let $S \subseteq V$ be a set with $\leq n/2$ vertices. We'd like to know whether there are many edges connecting S to $V \setminus S$. Here is a way of deriving this information using Rayleigh quotients. Let

$$x = |V \setminus S|\chi_S - |S|\chi_{V \setminus S}$$

where χ_Y is the characteristic function of the set Y (i.e. $\chi_Y(u) = 1, 0$ depending on whether u is in Y or not.)

Rayleigh quotients and expansion

Here are some easy facts:

- x is perpendicular to $\mathbf{1}$ (obvious).
- $\|x\|^2 = n|S|(n - |S|)$ (one line of calculation).
- $x^t A x$ depends only (and simply) on the number of edges between S and its complement. (a three-line calculation).

So this is telling us that

A large spectral gap implies high expansion, and in fact, expansion is just a special case of spectral gap between $\lambda_1 = d$ and λ_2. (But wait till you see the reverse implication).

The second eigenvector reveals bottlenecks in G: This is only the beginning of another fascinating story that we won't touch on that goes under "the nodal region theorem".

Spectral gap, expansion et. al.

What should come as a bigger surprise is that there is also an implication in the opposite direction. Namely, expansion implies spectral gap.

Why is this surprising? Spectral gap means a small Rayleigh quotient for all (uncountably many) vectors $x \perp \mathbf{1}$, whereas expansion means a small Rayleigh quotient for a finite number of such x. How can you draw such a strong conclusion from such a weak assumption?

Not so fast....

- The earliest version of this implication was proved in a geometric context (Cheeger).
- Expansion can be considered a discrete analogue of the classical isoperimetric problem.
- The graph-theoretic analogue of Cheeger's proof was found by N. Alon. The proof is quite difficult and the bounds are pretty weak.
- This is unavoidable, though. A graph can be an excellent expander and still have only a very meager spectral gap.

High expansion and poor spectral gap

- The line graph of $G = (V, E)$ is a graph L whose vertex set is $E(G)$.
- Two vertices in L are adjacent if the corresponding edges in G share a vertex.
- If G is d-regular, then L is $(2d - 2)$-regular.
- If G is an excellent expander, then so is L.
- The first eigenvalue of L is $2d - 2$ and the second one is $> d$. This spectral gap is poor as we'll see soon.

Is there a combinatorial equivalent to spectral gap?

Granted, it is fairly surprising that expansion (a combinatorial condition) and spectral gap (a linear-algebraic condition) are qualitatively equivalent. However, for a long time I was wondering if the spectral gap can also be quantitatively captured in combinatorial terms.
This was solved in joint work with Yonatan Bilu. It is based on the notion of discrepancy.

Recall what we said that expanders often behave as if they were random graphs. One major realization of this principle is:

Theorem 2 (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 3 (Yonatan Bilu, L.). *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

What's a "large" spectral gap?

An instinct that most combinatorialists have is to ask about any new parameter they encounter "how large/small can it be?"
In this spirit we ask:
Problem 1. *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 4 (Alon, Boppana).

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

The largest possible spectral gap and the number $2\sqrt{d-1}$

A good approach to extremal combinatorics and questions of the form "How large/small can this parameter get" is to guess a/the extremal case (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. Do you see why?
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....

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

Can we emulate phenomena from infinite graphs in the finite realm?

Part of the business of asking extremal problems is to ask for any bound you find whether it is tight.

Problem 2. *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.

We also want to understand the typical behavior.

Problem 3. *How likely is a (large) random d-regular graph to be Ramanujan?*

A few words about Ramanujan Graphs

Margulis; Lubotzky-Phillips-Sarnak; Morgenstern: 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.

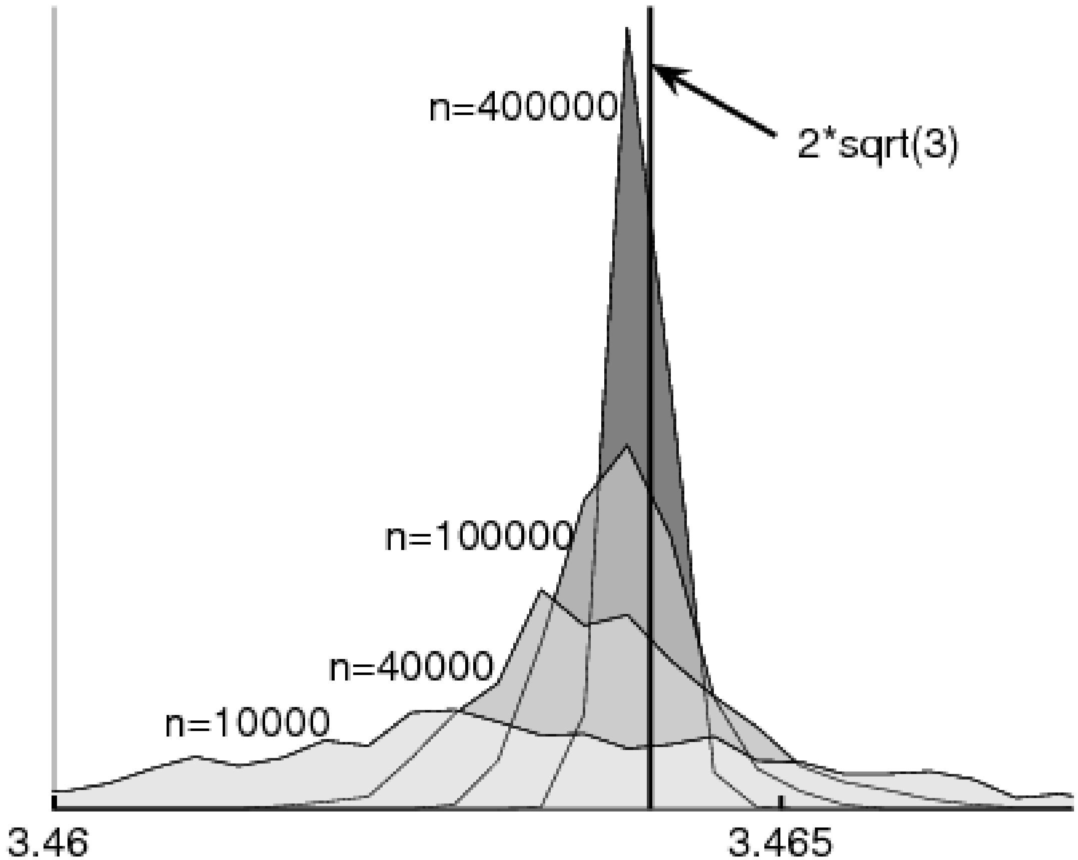
n=400000
2*sqrt(3)
n=100000
n=40000
n=10000
3.46
3.465

Some open problems

- Can we construct arbitrarily large d-regular Ramanujan Graphs for every d? Currently no one seems to know. $d = 7$ is the first unknown case.

- Can we find elementary methods of proof for graphs with large spectral gap (or even Ramanujan)? Some recent progress: constructions based on random lifts of graphs (Bilu-L.) can get $\lambda_2 \leq O(\sqrt{d}\log^{3/2} d)$.

The signing conjecture

The following, if true, would prove the existence of arbitrarily large d-regular Ramanujan graphs for every $d \geq 3$.

Conjecture 1. *Every d regular graph G has a signing with spectral radius $\leq 2\sqrt{d-1}$.*

A signing is a symmetric matrix in which some of the entries in the adjacency matrix of G are changed from $+1$ to -1. The spectral radius of a matrix is the largest absolute value of an eigenvalue.

This conjecture, if true, is tight.

Expansion at various size scales

- We have defined expansion in terms of worst-case. In particular a d-regular graph can have edge expansion at most $(1 - o(1))\frac{d}{2}$. On the other hand, in almost every d-regular graph, sets S of size $|S| \leq \epsilon n$ have at least $(d - 2 - \delta)|S|$ neighbors.
  (For every $d \geq 3$ and for every $\delta \geq 0$ there is an $\epsilon > 0$ such that in a random d-regular graph the following holds with probability $\geq 1 - \delta$: Every set S of $\leq \epsilon n$ vertices has $(d - 2 - \delta)|S|$ outside neighbors.)

- By a result of Kahale such conclusions cannot be derived from spectral information. Can we explicitly construct such graphs? The zig-zag method by Reingold-Vadhan-Wigderson goes a long way in this direction, but the problem is still far from being resolved.

- This is also closely related to the girth problem: What is the largest $g = g(n, d)$ such that there exist d-regular graphs on n vertices with no cycles of length $< g$?

- Even more generally this leads to local-global problems on graphs. How can a large graph look at the micro level? How does this affect the graph's global properties?

... and more and more and more ...

Problem 4. *How well can you approximate the expansion rate of a graph?*

Problem 5. *Which NP-hard computational problem on graph remain hard when restricted to expander graphs? Is the situation similar to what you see with random graphs?*

Something for analysts

Problem 6. *What do you see when you examine the second eigenvector of a random d-regular graph?*

Some preliminary calculations suggest that there is a limit distribution to the entries that is neither normal nor Tracy-Widom. Is this true?
Are we encountering here anything universal?

Something for algebraists

- The notion of expansion for Cayley graphs has been the focus of a lot of research in the last decades.
- There were several exciting developments in the very recent past: Typical behavior as well as extremal problems about expansion were solved for S_n and various families of Lie groups.
- Deep and interesting connections with Number Theory and Representation Theory (e.g. Kazhdan's Property (T)).
- Zig-zag products were used to resolve some problems in this area.

Something for geometers

Problem 7. *Is there a theory of expander hypergraphs?*

More concretely:

Problem 8. *How do you generate random 3-uniform hypergraphs that cover each edge d times? (Random Steiner Triple Systems.)*

Something for probabilists

Expansion is also equivalent to rapid mixing of the random walk on the graph. Progress is evaluated by measuring the distance between the current and limit (uniform) distribution.
If we evaluate this distance in l_2 we arrive at the usual theory of expansion/spectral gap etc.
What if we use l_1? Or entropy?

Problem 9. *What are the best quantitative connections between these different measures of progress (combinatorial, algebraic, probabilistic)?*

... and for computer scientists

Expansion is closely related to fundamental problems/constructions in derandomization/extraction of randomness.

There are many nice properties that expanders may have which we are still unable to accomplish.

Problem 10. *Can we construct expanders on which can efficiently solve:*

- *Shortest path problems.*
- *Routing problems.*