Introduction to Expanders

Styopa Zharkov

1 Introduction

Expanders are graphs that are sparse and well-connected. We'll allow for parallel edges and loop edges here. Those terms "sparse" and "well-connected" are very vague. We'll define what we mean by them in this paper, but in reality every application of expanders uses a different definition of sparseness and well-connectedness.

Expanders were first studied to understand how to build efficient networks, but there are now applications in many areas. Expanders can be used to create efficient pseudorandom generators [XZ07], they, they can be used to create good error-correcting codes called expander codes [HOW15], and they have found a shelf within the mathematical community.

2 Sparseness

First, let's talk about what it means for a graph to be sparse. Our goal will be to build expanders, and we want to be able to build them to be arbitrarily large. So, it makes sense for us to defines sparseness as something asymptotic. We can say "our construction generates sparse graphs" instead of "this graph is sparse".

Definition 4: We say a collection of graphs is *sparse* if the number of edges in each graph is linear in the number of vertices. ◁

Note that this definition is specific to this introduction. We will be more interested in a specific kind of sparse graphs.

Definition 5: We say an undirected graph G is a *max-degree-d* graph if all vertices have at most d edges coming out. Similarly, we say G is a *d-regular* graph if every vertex has exactly d edges coming out. ◁

We can see that for any constant d, every max-degree-d graph, and thus every d-regular graph, with n vertices has at most $nd/2$ edges, so these graphs are indeed sparse. If we imagine our graphs to be networks of servers, it makes sense that we don't want to put too much pressure on one server, so there should be a maximum on the number of connections it has. So, can see some intuition for why most expander applications people care about are restricted to max-degree-d graphs. Plus, we'll soon see an example that will show that without this restriction, constructing good expanders is easy.

3 Well-Connectedness

Now, lets look at the second property we care about, well-connectedness. Before we can arrive at the definition, we need a few preliminary definitions.

Definition 6: Let A and B be disjoint sets of vertices. We denote by $E(A, B)$ the set of vertices that have one endpoint in A and another endpoint in B. ◁

We will be specifically interested in $E(S, \bar{S})$. These are the edges that "cross the border" of S. Let's use this notation to define conductance.

Definition 7: The *conductance* of a nonempty set of vertices S in a graph $G = (V, E)$ is the ratio of the fraction of edges with exactly one endpoint in S to the fraction of vertices in S. We denote this by

$$\Phi(S) = \frac{|E(S, \bar{S})|/|E|}{|S|/|V|}.$$

◁

For d-regular graphs, $|V|/|E|$ is a constant, so we can throw out the $|V|$ and the $|E|$ and think of $\Phi(S)$ as something proportional to the ratio of the number of edges crossing into S to the number of vertices in S. Intuitively, $\Phi(S)$ is big if S is small, but connected to the rest of the graph well.

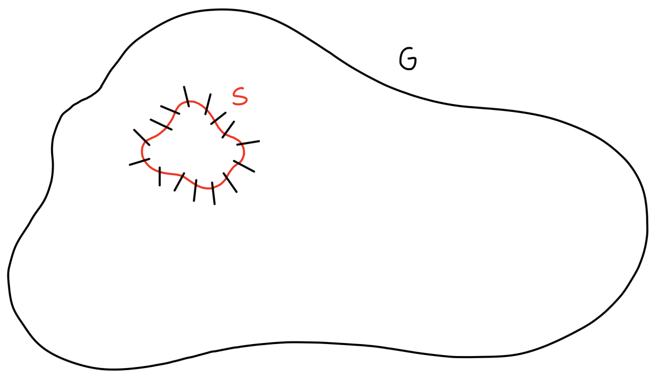

From this notion, we can build a graph property.

Definition 8: Let $G = (V, E)$ be a graph with at least two vertices. The *minimum conductance* of G is the minimum of the conductances of all subsets of no more than half the vertices. This is denoted by

$$\Phi_G = \min_{\substack{S \subseteq V \\ 0 < |S| \leq |V|/2}} \Phi(S).$$

◁

It makes sense that we restrict ourselves to only sets of at most half the vertices because the conductance of the entire graph is 0. A disconnected graph will have a minimum conductance of 0 because the conductance of any connected component is 0. Any connected graph will always have a positive minimum conductance since any nonempty S that isn't too large will have edges coming out. We can see that the higher the minimum conductance is, the more connected the graph feels. In fact, that's how we'll define well-connectedness.

Definition 9: We'll say that a collection of graphs is well-connected if the minimum conductance is at least 0.001 for each one. ◁

Of course, the choice to use 0.001 is arbitrary, we really just want the minimum conductance to be $\Omega(1)$. In a sense, any set S can "expand" well into the rest of the graph by attaching its neighbors, hence the name expanders.

Consider the collection of star graphs, where it's $n - 1$ vertices all connected to one central vertex and there are no other edges.

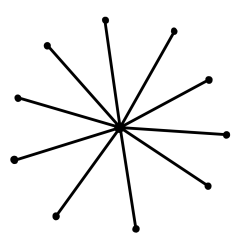

We can see that this collection of graphs is sparse. The graphs are also well-connected. We can see that for any set S of less than half the vertices, either the center vertex is in S or it isn't. If the center is in S, then at least half the edges are crossing into S.

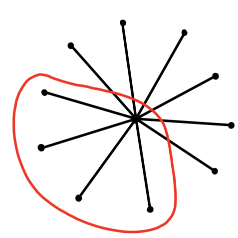

If the center is outside of S, then every vertex in S has an edge leading out of it.

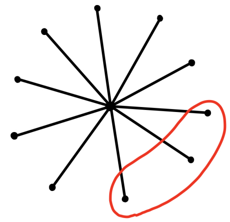

In either case, the number of edges crossing S is at least the number of vertices in S, so

$$\Phi(S) = \frac{|E(S,\bar{S})|/|E|}{|S|/|V|} \geq \frac{|S|/|E|}{|S|/|V|} = \frac{|V|}{|E|} = \frac{n}{n-1} \geq 0.001.$$

This means that $\Phi_G \geq 0.001$, so star graphs are indeed expanders. However, star graphs have a center node with many edges. Star graphs are not max-degree-d for any d. The problem becomes significantly more difficult if we bound the degree of the vertices.

4 Some Examples

Now, we can look as some more examples and nonexamples of expanders. This time, we will require that our graphs are d-regular.

First, consider 1-regular graphs. Every vertex only has one neighbor, so any 1-regular graph is simply many disconnected 2-point, 1-edge components. These cannot be expanders because setting S to be any disconnected component will give a conductance of 0.

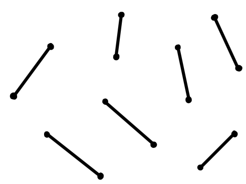

Bumping d up to 2, the situation is similar. A 2-regular graph is just one or more disconnected cycles.

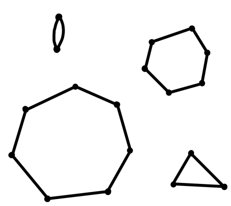

If there is more than one cycle, then at least one cycle has less than half the vertices. Setting S to be the set of vertices in that cycle will give a conductance of 0, so the graph isn't an expander. If there is only one big cycle, then we can set S to be some half of the cycle and there are only two edges going out of S, but the number of vertices in S grows with n, so there is no lower bound on the conductance of the graph, so 2-regular graphs aren't expanders.

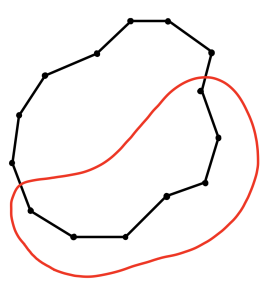

The story changes entirely when we set d to 3. Not only do 3-regular expanders exist, but most 3 regular graphs are expanders (i.e. a random 3-regular graph is an expander with high probability) [AS00]. The proof of this statement isn't complex, but it's tedious, so we'll leave it out. The idea is to follow the probabilistic method.

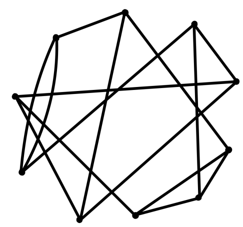

Of course, this means that for any $d \geq 3$, most d-regular graphs are expanders as well.

5 Explicit is Better

So we know that d-regular expanders exist, but the probabilistic method isn't constructive, so we don't know how to build these things. In many applications, the whole point is to construct them. So, we can set our goal to be explicitly defined expanders. Although "explicit" seems like something that doesn't need to be defined, not having a rigorous definition will always lead to trouble. For example, we could technically have the following construction for expanders.

Iterate through all d-regular graphs and check if it's an expander until you have reached one.

We know this method works because d-regular expanders exist. However, it's not clear if we should count this as an explicit construction. After all, it does take very long to search through all expanders. So, let's define explicit formally.

Definition 10: A construction of a graph is *explicit* if there exists a poly-time algorithm that constructs the adjacency matrix of the graph we want. ◁

Note that the is poly-time in n. We can similarly define a stronger condition.

Definition 11: A construction of a graph is *strongly explicit* if there is a polytime algorithm that takes in a vertex and an n as input and outputs all of its neighbors. ◁

Note that here the algorithm is poly-time in the length of the description of a vertex and n, which is $O(\log n)$. So, the algorithm runtime must actually be $\text{polylog}(n)$. This essentially lets us construct the graph vertex by vertex using these "neighbor queries", even if the size of the graph is exponential.

6 Strongly Explicit d-regular Expanders

We now have all of the preliminaries to appreciate the Margulis-Gabber-Galil construction of strongly regular 8-regular expanders. The idea for the construction was presented by Margulis in 1973 [Mar73], and Gabber and Galil [GG81] completed the idea in 1981. This construction only works when n is a perfect square.

Here's the construction. Let m be a positive integer and let $n = m^2$. Let's create n vertices, and assign an element of $\mathbb{Z}_m \times \mathbb{Z}_m$ to each one. Next, for every vertex (x, y), create an edge connecting it

1. to $(x \pm y, y)$,
2. to $(x \pm (y + 1), y)$,
3. to $(x, y \pm x)$,
4. and to $(x, y \pm (x + 1))$.

So, there are 8 edges from every vertex. Note that if a vertex v is supposed to be connected to u, then vertex u is supposed to be connected to v, so we don't run into any problems with that. Here's an example for $m = 3$:

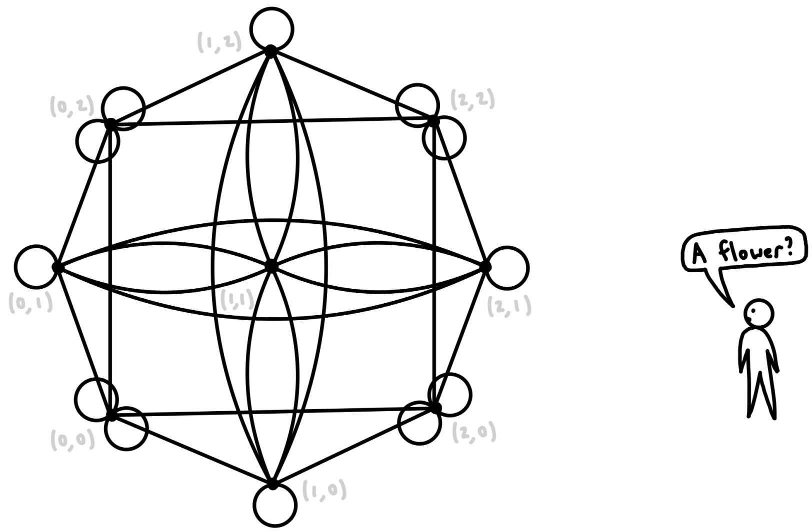

We will not prove the correctness of this construction here. In fact, there are many different constructions for expanders, all of which are relatively simple, but the proofs of correctness usually rely on a combination of results from many different areas of math.

7 Related Work

One recent important result is the Zig Zag construction of expanders presented by Reingold, Vadhan, and Widgerson in [RVW02]. The construction is not as simple—it uses an iterative technique to build big expanders from constant-size small expanders. However, the analysis is simpler than of most examples. The idea relies on a somewhat simple fact about matrices and their eigenvalues.

This result was a crucial step in proving Omer Reingold's famous result that places Undirected ST-Connectivity in Logspace [Rei08] in 2005.

An even more recent important result is presented by Noga Alon in 2020. He shows how to construct explicit d-regular expanders for any d and any n. For $d = p + 2$, where p is a prime that is 1 mod 4, he gives a strongly explicit construction.

The construction is a modification of already known constructions using something called Ramanujan graphs. The analysis and proof are a complicated mix of known results from spectral graph theory [Alo20].

References

[Alo20] Noga Alon. Explicit expanders of every degree and size, 2020.

[AS00] N. Alon and J.H. Spencer. *The Probabilistic Method.* John Wiley & Sons, Inc., 2 edition, 2000.

[GG81] Ofer Gabber and Zvi Galil. Explicit constructions of linear-sized superconcentrators. *Journal of Computer and System Sciences*, 22(3):407–420, 1981.

[HOW15] Brett Hemenway, Rafail Ostrovsky, and Mary Wootters. Local correctability of expander codes, 2015.

[Mar73] G. A. Margulis. Explicit constructions of concentrators. *Probl. Peredachi Inf.*, 9(4):71–80, 1973.

[Rei08] Omer Reingold. Undirected connectivity in log-space. *J. ACM*, 55(4), September 2008.

[RVW02] Omer Reingold, Salil Vadhan, and Avi Wigderson. Entropy waves, the zig-zag graph product, and new constant-degree expanders. *Annals of Mathematics*, 155(1):157–187, 2002.

[XZ07] Kristin E. Lauter Denis Charles X and Eyal Goren Zvi. Pseudorandom number generation with expander graphs, July 2007.