Jonathan Sun

jsun@cs.ucsb.edu

CS 290H

Expander Graphs: Theory and Applications

Contents

Introduction

Graph Properties and Motivation

Many research problems in graph theory study the abstract properties of graphs that originate from real-world structures and intuitive motivations. Formalizing the connection between these properties and their meaning in real-world applications is often just as involved as defining and computing the property itself, since researchers often have to choose between several competing or ambiguous definitions that each have their own advantages. Furthermore, some graph properties that arise naturally from real-world data prove to be intractably difficult to compute or even approximate, resulting in a mathematical model that is of little use. One important, widely applicable graph property used in network theory is the idea of *expansion*, which is closely related to the concept of graph connectivity. Although expansion can be defined naturally and seems to describe a wide variety of network phenomena, computing the expansion for a graph or even describing good expanders proved to be difficult for many years. Fortunately, research results linking expander graphs to spectral theory have yielded useful tools for discussing the expansion of a graph.

Expansion as an Intuitive Graph Property

Similar to centrality, the expansion factor of a node in a graph measures how well connected the node is to every other node by looking for bottlenecks in the graph; intuitively, a node with high expansion should have little or no bottlenecks to every other node. In real world applications, this can be interpreted roughly as the rate of information spread from a node, since nodes with high expansion should be able to propagate information quickly to the rest of the graph. In general, however, this definition is used to describe the entire graph; a graph has high expansion if there are no bottlenecks anywhere in the graph. Aside from being a mathematical graph property, the notion of graph expansion has wide applications in other fields of computing, including areas like error-correcting codes, embeddings, and random walks. [9]

Formal Definitions of Expansion

One reason expansion is so widely applicable is because it is an intuitive property of large networks that can still be described well in formal mathematical terms. For any set of nodes S in a graph G, define its *boundary d(S)* as the set of vertices that neighbors S in G. In other words:

$$d(S) = \{v \in S' \mid (u, v) \in E, u \in S\}$$

where S' is the complement of S in G, and E is the set of edges in G. In other words, the boundary of S is the set of nodes that can be reached from S in one "hop". Clearly, the boundary can also be defined as the set of *edges* neighboring S rather than the vertices [9];

each definition has its own advantages and properties. We can now define the *vertex expansion factor h(G)* for the graph *G* to be:

$$h(G) = \min_S \frac{d(S)}{|S|} \text{ where } |S| \leq \frac{|V|}{2}$$

The corresponding alternative formulation for the edge expansion factor uses the edge boundary rather than the vertex boundary. With the expansion factor formally defined, there are three overarching ways to interpret this measure in terms of real-world data. Combinatorially speaking, the expansion factor is small if the graph has a small min-cut, since each side of the cut has a large size and very small boundary (corresponding to the cut size). Conversely, a well-connected graph with no small cuts should have a large expansion. Geometrically, a graph with large expansion factor should yield, for every node, a large neighborhood (in terms of number of nodes) at all scales. Finally, a major result in spectral graph theory reveals that large expansion is equivalent to having a large spectral gap in the corresponding adjacency matrix; a similar result exists for the Laplacian matrix as well. [2] The first two interpretations of expansion follow simply from the definition; the algebraic correlation, however, provides a nontrivial link between graph theory and the powerful tools of linear algebra and scientific computing.

Properties of Expander Graphs

Matrix View of Graphs

Alone, the definition of expansion provides a well-defined measure on all graphs, but the notion of an *expander graph* is still unclear, since clearly all connected graphs have positive expansion, and the complete graph should have maximal expansion for its size. A useful constraint, therefore, is to require a constant vertex degree in the graph as the size of the graph scales. In other words, we define a graph *G* to be *d-regular* if all its nodes have degree *d*, and focus on the relationship between *d* and *h(G)* as the size of *G* grows. [9] Although the definition of expansion is mathematically simple and intuitive, computation and detection of bounded-degree expander graphs remains a difficult open problem. Detecting whether an arbitrary graph is an ε-expander (that is, has an expansion factor at least ε) is co-NP-hard. Equivalently, computing the expansion factor for a graph is co-NP-hard even if the problem is restricted to subsets of constant size *k*, and combinatorial approximation algorithms for expansion are lacking. [9] So far, the most prevalent research results in the field have come in the form of linear algebra, as the expansion factor of a graph can be tightly linked to the eigenvalues of its associated matrices—the *adjacency matrix A*, which contains ones at adjacent node indices, and the *Laplacian matrix L*, which contains vertex degrees on the diagonal and negative ones at adjacent node indices. These eigenvalues, which can be computed/approximated much more efficiently using known

linear algebra methods, provide much of the theoretical foundation behind expander graphs and their applications.

Spectral Gap as an Estimator of Expansion

The basic objects of study in spectral graph theory are the adjacency matrix A and the Laplacian matrix L of a graph G. Since both A and L are real symmetric n-by-n matrices (where n is the number of vertices in G), they each have a complete set of n real eigenvalues and eigenvectors, including multiplicity. The *spectrum* of G is obtained by ordering the eigenvalues $\{\lambda_1, \lambda_2 \ldots \lambda_n\}$ of A in decreasing order. Conveniently, in a d-regular graph, the value of λ_1 is exactly d. Furthermore, using the spectrum of G, the expansion factor of a graph is bounded by the following theorem [9]:

$$\frac{\lambda_1 - \lambda_2}{2} \leq h(G) \leq \sqrt{2\lambda_1(\lambda_1 - \lambda_2)}$$

The key quantity here, $\lambda_1 - \lambda_2$, is equivalent to $d - \lambda_2$ for d-regular graphs, and is called the *spectral gap* as it quantifies the "gap" between the largest and second-largest eigenvalues of the adjacency matrix. Therefore, writing the spectral gap as Δ, the above theorem can be thought of in simpler terms:

$$\frac{\Delta}{2} \leq h(G) \leq \sqrt{2d\Delta}$$

Clearly, the expansion factor is directly related to the spectral gap of the graph. More importantly, the inequality gives some sense of how much the expansion can vary with the vertex degree—for all graphs of bounded vertex degree d, the expansion is roughly characterized by the second-largest eigenvalue of the graph's adjacency matrix. [2] This inequality is a discrete extension of the theorem known as Cheeger's inequality from analysis—the main difference is that the form originally proved by Cheeger involves continuous manifolds rather than discrete graphs.

Correlation with the Laplacian Matrix Spectrum

Along with the expansion bounds given by the spectral gap of the adjacency matrix is an alternative formulation that uses the closely related Laplacian matrix. The Laplacian matrix contains some useful properties not immediately visible in the adjacency matrix; for example, the eigenvalues of the Laplacian matrix are all nonnegative, and the multiplicity of the zero-eigenvalue is precisely the number of connected components in the graph. If λ is the smallest *non-zero* eigenvalue of the Laplacian matrix, then it is the case that $\lambda = \Delta$, which allows the Cheeger inequality to be written in terms of the Laplacian eigenvalues instead. Furthermore, since the minimum eigenvalue of a matrix can be expressed using Rayleigh quotients [13], this quantity can also be written as:

$$\Delta = \lambda = \min_{x \perp \mathbf{1}} \frac{x^T L x}{x^T x}$$

where $\mathbf{1}$ is the unit vector of all ones, and L is the Laplacian matrix.

Examples of Expander Graphs

Simple Extremal Cases

Although the notion of graph expansion is simple to define, and the spectral gap allows an easy estimation of the magnitude of the expansion factor, the existence and construction of large, sparse graphs with good expansion remained a difficult problem years after the notion was defined. Clearly the complete graph has optimal expansion factor, since every subset of nodes has the whole set as its boundary, but the vertex degree of the complete graph grows linearly in the size of the graph; an ideal expander graph combines large expansion factor with a constant vertex degree. A few other simple examples demonstrate that the bounds given by the Cheeger inequality are, in fact, tight. In the case of the lower bound, the *d*-dimensional hypercube graph on 2^d nodes yields an expansion factor of 1 and a spectral gap of 2, while in the case of the upper bound, the *n*-cycle graph yields an expansion factor of $\frac{1}{n}$ with a spectral gap of $\frac{1}{n^2}$. For bounded-degree graphs, if infinite graphs are considered, the infinite *d*-tree has the optimal expansion factor, but graphs of infinite size are usually impractical to consider in real-world applications. [9]

Margulis Construction and Generalizations

The first explicit example of an infinite family of finite constant-degree expanders was due to Margulis, who gave a non-constructive existence proof, and Gabber and Galil, who gave an explicit, elementary construction. [1] While the proofs of expansion themselves are much more complex, the construction itself uses only the elements of $\mathbb{Z}_n \times \mathbb{Z}_n$ as nodes and adds edges by transforming each element into its 8 neighbors, resulting in a bipartite 8-regular graph. A few years later, this result was generalized in the form of *Ramanujan* graphs, which are *d*-regular graphs for which $\lambda \leq 2\sqrt{d-1}$, where λ is the second largest eigenvalue of the adjacency matrix. [3] The discovery that Ramanujan graphs for arbitrarily large values of *d* exist and could be constructed was another major step in describing explicit expanders. Furthermore, Friedman showed using the probabilistic method that:

$$P\left(\lambda \leq 2\sqrt{d-1} + \varepsilon\right) \to 0 \text{ as } n \to \infty \text{ for all } \varepsilon > 0.$$

In other words, a randomly generated *d*-regular graph of size *n* will be Ramanujan with nonzero probability as *n* approaches infinity, so Ramanujan graphs are actually quite "common" in this sense [5].

Zigzag Products of Graphs

Although the first explicit constructions of expander graphs took years to discover and required deep proofs, since then, many new techniques have appeared that use various techniques to generate expander graphs of different types. One important example is the *zig-zag* product, which allows expander graphs to be "multiplied" together to form new, distinct expanders. First described by Reingold, Vadhan, and Wigderson [6], the zig-zag product takes a d-regular graph G_1 on n vertices, combines it with a c-regular graph G_2 on d vertices, and outputs a c^2-regular graph with c × n vertices. From the definition, it is clear that G_2 is expected to be much smaller than G_1, and c to be much smaller than d. Yet the product graph is similar in size to G_1, has the desired properties of bounded degree, and is proven to have an expansion factor that is similar to the original graphs. In other words, the output of a zig-zag product on two expander graphs is yet another expander. Aside from being used to generate arbitrary families of expander graphs, the zig-zag product has further theoretical applications in the study of graph operations.

Applications of Expander Graphs

Expander Graphs for Error-Correcting Codes

Although much of the developed theory on expander graphs involves graph theory and linear algebra, the original motivations for constructing expanders came out of communication theory, in the form of *error-correcting codes*. In a communication setting with signal noise, a sender and receiver cannot guarantee the fidelity of sent messages—they can only ensure that a proportion of the data bits sent are accurate. However, using error correcting codes, the sender can still deliver the correct message to the receiver by sending extra, redundant bits with the original message. Generally speaking, in a error correcting model these bits essentially "vote" for the correct message content, allowing the receiver to deduce the original message even if some of the bits were corrupted in transit. In formal terms, an error correcting code C is a set of n-bit binary strings (though more complex languages can be used) that, ideally, maximizes two parameters:

$$d(C) = \min_{x \neq y \in C} d(x, y) \text{ and } r(C) = \frac{\log|C|}{n}$$

where $d(x, y)$ is the Hamming distance (number of bits different) between two words. The first parameter, called the *distance* of C, is a measure of the breadth of messages and errors that C can account for, while the second parameter, called the *rate*, inversely measures the data overhead required per message sent. Error correcting codes and algorithms form their own deep topic of study, but Pinsker and Bassalygo realized that high quality codes for large messages could be formed using sparse bipartite graphs—one particular form of expanders. [9] Furthermore, later results proved that the codes formed from such graphs

would allow for efficient encoding and decoding of messages, including linear-time algorithms from Sipser and Spielman. [4] These results were much of the motivation for finding explicit constructions of expander graph families, as opposed to merely proving their existence.

Expansion Properties of Social Network Graphs

A second, more recent application of expander graphs involves the study of large social networks that have been made available by the growth of the internet and the interest in analyzing social data. Social network graphs are unique in that they combine bounded vertex degree with large size, yet tend to have very small diameters (an observation colloquially called the "small world" phenomenon). Aside from the basic graph properties, however, one important concept of study on social network graphs is that of *community structure*, which intuitively refers to the identification of highly-connected subsets of the social network. [7] For a social network graph where nodes are individual people, a *community* is generally interpreted as either a group of tightly knit individuals (like a circle of friends) or a set of highly similar people (by shared interests, etc.). In either case, social network analysis seeks to detect and identify such groups using properties of the underlying graph. However, due to the complexities of quantifying social interactions and data, no universal definition of a community exists, at least in the mathematical sense. The most commonly used measure is called *modularity* and is computed as follows [11]:

$$Q = \sum_{i=1}^{k} e_i - a_i^2$$

Given a graph G, and a set of communities $c_1, c_2 \ldots c_k$ that *partition* G (in other words, no node belongs to two communities, and the union of all communities is G), the quantity e_i is the proportion of edges of G with both ends in c_i, while the quantity a_i is the proportion of edges with at least one end in c_i. Intuitively speaking, a community partition for G exhibits high modularity if many edges lie in the communities (so each community is highly connected) and few edges lie across communities (so each community is tightly bounded). Several community detection algorithms proceed by finding a community partition with high modularity (optimal modularity maximization is NP-hard). While the modularity measure has some shortcomings of its own, it does have an important correlation to expansion factor. Since graphs with high expansion have, for all vertex subsets, a large boundary, any community partition of such a graph should have low modularity due to the large number of inter-community edges. Conversely, if a graph does have a partition with high modularity, then there exist subsets of the graph that do not expand well—precisely the communities. If social network graphs exhibit strong community structure, then their expansion factors should be correspondingly small. In reality, however, the results are mixed. Estrada experimentally verified a small expansion factor for several small real-

world data sets [10], but Malliaros and Megalooikonomou later demonstrated that many large-scale social graphs actually exhibit surprisingly good expansion behavior [12]. Studying the expansion properties of social network graphs becomes increasingly difficult with scale, as even computing eigenvalues on graphs the size of the internet becomes infeasible. Therefore, whether large social graphs are in general expanders still remains an open question.

Open Questions on Expander Graphs

Because both the expansion factor and expander graphs have historically been difficult to compute in general, many open questions still exist regarding their nature. Since computing the expansion factor exactly is known to be in co-NP-hard, the major theoretical question remains whether it is possible to efficiently *approximate* the expansion up to a constant factor. Without knowledge of the spectral gap, recent results have only yielded $\sqrt{\log n}$ approximations on the expansion factor. [8] Various other problems of a similar nature remain unsolved, most involving optimizing the vertex/edge-boundaries of graph subsets. On the application side, many algorithms, including error-correcting codes, make use of explicitly constructed infinite families of expander graphs, and finding simpler, faster constructions is always an open research field. Finally, as the graphs of very large scale become increasingly common in graph research (most notably coming from the internet and social networks), an important problem is to study their expansion properties efficiently. Furthermore, while some probabilistic results for expansion are known for random graphs in general, many open questions exist regarding the expansion factor for large *social* random graph models.

Conclusion

The study of expander graphs arose in a convoluted manner, since continuous results (on manifolds and Laplace operators) have existed for decades before graph theory was a common research topic. The invention of error-correcting codes and their subsequent link to expander graphs pushed forward a large research effort to verify whether bounded-degree expanders exist, and, if so, how to construct them for arbitrary size and degree. Formalizing the concept of expansion to general graphs and correlating the expansion factor to the spectral gap was the major result that now underpins much of the theory on expander graphs. Finally, as researchers begin to apply spectral graph theory to a wide range of applications, in social networks as mentioned, but also fields like biology, number theory, and physics, studying the expansion factor becomes an important part of analyzing a graph's properties.

References

[1] O. Gabber and Z. Galil. Explicit constructions of linear-sized superconcentrators. J. Comput. System Sci., 22(3):407–420, 1981.

[2] N. Alon. Eigenvalues and expanders. Combinatorica, 6(2):83–96, 1986. Theory of Computing (Singer Island, FL, 1984).

[3] Moshe Morgenstern. 1994. Ramanujan Diagrams. SIAM J. Discret. Math. 7, 4 (November 1994), 560-570.

[4] M. Sipser and D. A. Spielman. Expander codes. IEEE Trans. Inform. Theory, 42(6, part 1):1710–1722, 1996. Codes and complexity.

[5] Friedman, Joel. 'A Proof Of Alon's Second Eigenvalue Conjecture And Related'. CoRR cs.DM/0405020 (2004): n. pag. Print.

[6] O. Reingold, S. Vadhan, and A. Wigderson. Entropy waves, the zig-zag graph product, and new constant-degree expanders. Ann. of Math. (2), 155:157–187, 2002.

[7] Newman, Mark EJ, and Juyong Park. "Why social networks are different from other types of networks." Physical Review E 68.3 (2003): 036122.

[8] S. Arora, S. Rao, and U. Vazirani. Expander flows, geometric embeddings and graph partitioning. In Proceedings of the 36th Annual ACM Symposium on Theory of Computing, pages 222–231, 2004.

[9] Hoory, Shlomo, Nathan Linial, and Avi Wigderson. "Expander graphs and their applications." Bulletin of the American Mathematical Society 43.4 (2006): 439-561.

[10] Estrada, E. "Spectral scaling and good expansion properties in complex networks." EPL (Europhysics Letters) 73.4 (2006): 649.

[11] Newman, Mark EJ. "Modularity and community structure in networks." Proceedings of the National Academy of Sciences 103.23 (2006): 8577-8582.

[12] Malliaros, Fragkiskos D., and Vasileios Megalooikonomou. "Expansion properties of large social graphs." Database Systems for Adanced Applications. Springer Berlin Heidelberg, 2011. 311-322.

[13] Makarychev, Konstantin, and Yury Makarychev. 'How To Play Unique Games On Expanders'. Springer Berlin Heidelberg 6534 (2011): 190-200.