

## Lecture 1

*In which we describe what this course is about.*

### 1 Overview

This class is about the following topics:

1. *Approximation algorithms for graph partitioning problems.* We will study approximation algorithms for the *sparsest cut* problem, in which one wants to find a cut (a partition into two sets) of the vertex set of a given graph so that a minimal number of edges cross the cut compared to the number of pairs of vertices that are disconnected by the removal of such edges.

This problem is related to estimating the edge expansion of a graph and to find *balanced separators*, that is, ways to disconnect a constant fraction of the pairs of vertices in a graph after removing a minimal number of edges.

Finding balanced separators and sparse cuts arises in *clustering* problems, in which the presence of an edge denotes a relation of similarity, and one wants to partition vertices into few clusters so that, for the most part, vertices in the same cluster are similar and vertices in different clusters are not. For example, sparse cut approximation algorithms are used for *image segmentation*, by reducing the image segmentation problem to a graph clustering problem in which the vertices are the pixels of the image and the (weights of the) edges represent similarities between nearby pixels.

Balanced separators are also useful in the design of divide-and-conquer algorithms for graph problems, in which one finds a small set of edges that disconnects the graph, recursively solves the problem on the connected components, and then patches the partial solutions and the edges of the cut, via either exact methods (usually dynamic programming) or approximate heuristic. The sparsity of the cut determines the running time of the exact algorithms and the quality of approximation of the heuristic ones.

We will study three approximation algorithms:

- (a) The *Spectral Partitioning Algorithm*, based on linear algebra;
- (b) The *Leighton-Rao Algorithm*, based on a linear programming relaxation;
- (c) The *Arora-Rao-Vazirani Algorithm*, based on a semidefinite programming relaxation.

The three approaches are related, because the continuous optimization problem that underlies the Spectral Partitioning algorithm is a relaxation of the ARV semidefinite programming relaxation, and so is the Leighton-Rao relaxation. Rounding the Leighton-Rao and the Arora-Rao-Vazirani relaxations raise interesting problems in metric geometry, some of which are still open.

2. *Explicit Constructions of Bounded-Degree Expanders.* Expander graphs are graphs with very strong connectivity and “pseudorandomness” properties. Constructions of constant-degree expanders are useful in a variety of applications, from the design of data structures, to the derandomization of algorithms, from efficient cryptographic constructions to being building blocks of more complex quasirandom objects.

There are two families of approaches to the explicit (efficient) construction of bounded-degree expanders. One is via algebraic constructions, typically ones in which the expander is constructed as a Cayley graph of a finite group. Usually these constructions are easy to describe but rather difficult to analyze. The study of such expanders, and of the related group properties, has become a very active research program, involving mostly ergodic theorists and number theorists. There are also combinatorial constructions, which are somewhat more complicated to describe but considerably simpler to analyze.

3. *Bounding the Mixing Time of Random Walks and Approximate Counting and Sampling.* If one takes a random walk in a regular graph that is connected and not bipartite, then, regardless of the starting vertex, the distribution of the  $t$ -th step of the walk is close to the uniform distribution over the vertices, provided that  $t$  is large enough. It is always sufficient for  $t$  to be quadratic in the number of vertices; in some graphs, however, the distribution is near-uniform even when  $t$  is just poly-logarithmic. Among other applications, the study of the “mixing time” (the time that it takes to reach the uniform distribution) of random walks has applications to analyzing the convergence time of certain randomized algorithms.

The design of approximation algorithms for *combinatorial counting* problems, in which one wants to count the number of solutions to a given NP-type problem, can be reduced to the design of *approximately uniform sampling* in which one wants to approximately sample from the set of such solutions. For example, the task of approximately counting the number of perfect matchings can be reduced to the task of sampling almost uniformly from the set of perfect matchings of a

given graph. One can design approximate sampling algorithms by starting from an arbitrary solution and then making a series of random local changes. The behavior of the algorithm then corresponds to performing a random walk in the graph that has a vertex for every possible solution and an edge for each local change that the algorithm can choose to make. Although the graph can have an exponential number of vertices in the size of the problem that we want to solve, it is possible for the approximate sampling algorithm to run in polynomial time, provided that a random walk in the graph converges to uniform in time poly-logarithmic in its size.

The study of the mixing time of random walks in graphs is thus a main analysis tool to bound the running time of approximate sampling algorithms (and, via reductions, of approximate counting algorithms).

These three research programs are pursued by largely disjoint communities, but share the same mathematical core.

One direction of *Cheeger's inequality*, for example, a basic result in algebraic graph theory, is useful in the construction of expanders because it establishes that good edge-expansion (the property that one is usually looking for, but that is a coNP-complete, and thus rather hard to certify) is implied by good spectral expansion (a property that is usually easier to establish and that is in P and thus has short certificates); the other direction of Cheeger's inequality, that good edge expansion implies good spectral expansion, is often used in the study of random walks, because spectral expansion is the property that controls the mixing time of random walks, and in some cases it is easier to prove that a graph has good spectral expansion by proving that it has good edge expansion. Both directions are equivalent to the statement that the nearly-linear-time *spectral partitioning algorithm* achieves a non-trivial approximation for the sparsest cut problem.

In this course we will study these three research programs back-to-back, emphasizing the connections, and providing, when necessary, a “dictionary” to translate the ways the same mathematical facts are thought about in the three communities.

## 2 Expander Graphs and Sparse Cuts

Before giving the definition of expander graph, it is helpful to consider examples of graphs that are not expanders, in order to gain intuition about the type of “bad examples” that the definition is designed to avoid.

Suppose that a communication network is shaped as a path, with the vertices representing the communicating devices and the edges representing the available links. The clearly undesirable feature of such a configuration is that the failure of a single

edge can cause the network to be disconnected, and, in particular, the failure of the middle edge will disconnect half of the vertices from the other half.

This is a situation that can occur in reality. Most of Italian highway traffic is along the highway that connect Milan to Naples via Bologna, Florence and Rome. The section between Bologna and Florence goes through relatively high mountain passes, and snow and ice can cause road closure. When this happens, it is almost impossible to drive between Northern and Southern Italy. Closer to California, I was once driving from Banff, a mountain resort town in Alberta which hosts a mathematical institute, back to the US. Suddenly, traffic on Canada’s highway 1 came to a stop. People from the other cars, after a while, got out of the cars and started hanging out and chatting on the side of the road. We asked if there was any other way to go in case whatever accident was ahead of us would cause a long road closure. They said no, this is the only highway here. Thankfully we started moving again in half an hour or so.

Now, consider a two-dimensional  $\sqrt{n} \times \sqrt{n}$  grid. The removal of an edge cannot disconnect the graph, and the removal of a constant number of edges can only disconnect a constant number of vertices from the rest of the graph, but it is possible to remove just  $\sqrt{n}$  edges, a  $1/O(\sqrt{n})$  fraction of the total, and have half of the vertices be disconnected from the other half.

A  $k$ -dimensional hypercube with  $n = 2^k$  is considerably better connected than a grid, although it is still possible to remove a vanishingly small fraction of edges (the edges of a dimension cut, which are a  $1/k = 1/\log_2 n$  fraction of the total number of edges) and disconnect half of the vertices from the other half.

Clearly, the most reliable network layout is the clique; in a clique, if an adversary wants to disconnect a  $p$  fraction of vertices from the rest of the graph, he has to remove at least a  $p \cdot (1 - p)$  fraction of edges from the graph.

This property of the clique will be our “gold standard” for reliability. The expansion and the sparsest cut parameters of a graph measure how worse a graph is compared with a clique from this point.

**Definition 1 (Sparsest Cut)** *Let  $G = (V, E)$  be a graph and let  $(S, V - S)$  be a partition of the vertices (a cut). Then the sparsity of the cut is*

$$\phi(S) := \frac{E(S, V - S)}{|E|} \cdot \left( \frac{|S| \cdot |V - S|}{|V|^2/2} \right)^{-1}$$

where  $E(S, V - S)$  is the number of edges in  $E$  that have one endpoint in  $S$  and one endpoint in  $V - S$ .

The sparsest cut is, given a graph, to find the set of minimal sparsity. The sparsity of a graph  $G = (V, E)$  is

$$\phi(G) := \min_{S \subseteq V: S \neq \emptyset, S \neq V} \phi(S)$$

That is, we are looking at the ratio between the fraction of edges that need to be removed in order to disconnect  $S$  from  $V - S$  and the fraction of pairs of vertices that would be so disconnected.

It is more common to define the sparsity as

$$\frac{E(S, V - S)}{|S| \cdot |V - S|}$$

without the normalizing factor ( $V^2/2|E|$ ); the normalized definition used above yields simpler formulas in some of the applications that we will discuss later.

Note that if  $G$  is a  $d$ -regular graph, then

$$\phi(S) := \frac{E(S, V - S)}{\frac{d}{|V|} \cdot |S| \cdot |V - S|}$$

In a  $d$ -regular graph, the *edge expansion* of a cut  $(S, V - S)$  is the related quantity

$$h(S) := \frac{E(S, V - S)}{d \cdot \min\{|S|, |V - S|\}}$$

in which we look at the ratio between the number of edges between  $S$  and  $V - S$  and the obvious upper bound given by the total number of edges incident on the smaller side of the cut.

The edge expansion  $h(G)$  of a graph is the minimum of  $h(S)$  over all non-trivial partitions  $(S, V - S)$ .

(It is common to define the edge expansion without the normalizing factor of  $d$  in the denominator.)

We note that for every regular graph  $G$  we have that, for every set  $S$ ,

$$\phi(S) \leq h(S) \leq 2 \cdot \phi(S)$$

and hence

$$\phi(G) \leq h(G) \leq 2 \cdot \phi(G)$$

A family of constant degree expanders is a family of (multi-)graphs  $\{G_n\}_{n \geq d}$  such that each graph  $G_n$  is a  $d$ -regular graph with  $n$  vertices and such that there is an absolute constant  $h > 0$  such that  $h(G_n) \geq h$  for every  $n$ .

Constant-degree graphs of constant expansion are sparse graphs with exceptionally good connectivity properties. For example, we have the following observation.

**Lemma 2** *Let  $G = (V, E)$  be a regular graph of expansion  $h$ . Then, after an  $\epsilon < h$  fraction of the edges are adversarially removed, the graph has a connected component that spans at least  $1 - \epsilon/2h$  fraction of the vertices.*

PROOF: Let  $d$  be the degree of  $G$ , and let  $E' \subseteq E$  be an arbitrary subset of  $\leq \epsilon|E| = \epsilon \cdot d \cdot |V|/2$  edges. Let  $C_1, \dots, C_m$  be the connected components of the graph  $(V, E - E')$ , ordered so that  $|C_1| \geq |C_2| \geq \dots \geq |C_m|$ . We want to prove that  $|C_1| \geq |V| \cdot (1 - 2\epsilon/h)$ . We have

$$|E'| \geq \frac{1}{2} \sum_{i \neq j} E(C_i, C_j) = \frac{1}{2} \sum_i E(C_i, V - C_i)$$

If  $|C_1| \leq |V|/2$ , then we have

$$|E'| \geq \frac{1}{2} \sum_i d \cdot h \cdot |C_i| = \frac{1}{2} \cdot d \cdot h \cdot |V|$$

but this is impossible if  $h > \epsilon$ .

If  $|C_1| \geq |V|/2$ , then define  $S := C_2 \cup \dots \cup C_m$ . We have

$$|E'| \geq E(C_1, S) \geq d \cdot h \cdot |S|$$

which implies that  $|S| \leq \frac{\epsilon}{2h} \cdot |V|$  and so  $|C_1| \geq (1 - \frac{\epsilon}{2h}) \cdot |V|$ .  $\square$

In words, this means that, in a  $d$ -regular expander, the removal of  $k$  edges can cause at most  $O(k/d)$  vertices to be disconnected from the remaining “giant component.” Clearly, it is always possible to disconnect  $k/d$  vertices after removing  $k$  edges, so the reliability of an expander is essentially best possible.