CS 860: Modern Topics in Graph Algorithms **University of Waterloo: Winter 2024**

Lecture 6: Expanders and Expander Decompositions

February 27, 2024

Instructor: Sepehr Assadi *Scribes: Cicely Henderson, Hidde Koerts, Massimo Vicenzo*

Topics of this Lecture

1 Basic Expander Principles

1.1 Definitions and Intuition

In this lecture, we would like to analyze the properties of "well-connected" graphs, and how we can use these properties to efficiently solve certain problems.

Firstly, what do we mean by "well-connected"? The radius and diameter of a graph are both measures of being connected, however they give a very global view of the graph, only telling us about the distance between any two vertices in the graph. Similarly k-connectivity is also a global measure of connectedness, though it does give us more information about the degrees of the vertices. In this lecture, we consider "expansion" as a measure of well-connectedness.

Notation. To continue, let us set up some notation. For a graph $G = (V, E)$ and subsets $A, B \subseteq V$, define:

$$E(A, B) := \{(u, v) \in E \mid u \in A, v \in B\}.$$

For a set $S \subseteq V$, $(S, V \setminus S)$ is a **cut** in G and $E(S, V \setminus S)$ is its **cut edges**. When clear from the context, we use $\overline{S} := V \setminus S$. We also define the **volume** of a set S as:

$$\mathrm{vol}(S) := \sum_{v \in S} \deg_G(v).$$

Notice that we always have,

$$|E(S,V)| \leqslant \mathrm{vol}(S) \leqslant 2 \cdot |E(S,V)| \qquad \text{(as the sum counts every edge inside } S \text{ twice and other edges once)}$$
$$\mathrm{vol}(S) + \mathrm{vol}(\overline{S}) = 2m. \qquad \text{(by the handshaking lemma)}$$

We are now ready for our key definitions in this lecture.

Definition 1. Let $G = (V, E)$ be any undirected graph. For any $S \subseteq V$, the **conductance** of the cut $(S, \overline{S})$, denoted by $\phi_G(S)$, is defined as:

$$\phi_G(S) := \frac{|E(S, \overline{S})|}{\min\big(\mathrm{vol}(S), \mathrm{vol}(\overline{S})\big)}.$$

Furthermore, the **conductance of a graph** G, denoted by $\phi(G)$, is

$$\phi(G) := \min_{\emptyset \subsetneq S \subsetneq V} \phi(S).$$

We sometimes (informally) refer to the conductance of G as the **expansion** of G also[a].

[a]This is with a bit of abuse of notation given there are multiple definitions of "expansion" which, even though related, are not identical to conductance.

The conductance is what gives us our measure for "well connectivity": informally speaking, the conductance of a set S tells us what proportion of the edges incident on S is part of the cut-edges of $(S, \overline{S})$ (this is assuming S is the "smaller" side of the cut in terms of the volume). A large conductance tells us that S and $\overline{S}$ are "well connected", as there are many edges between the two sets, rather that within them. By the same token, the conductance of the graph is then a measure of the least "well connected" cut in G. This intuition is what finally gives us our definition of *expanders*.

Definition 2. Let $G = (V, E)$ be any undirected graph and $\varphi \in [0, 1]$ be a parameter. A φ**-sparse cut** in G is a cut $(S, \overline{S})$ such that $\phi_G(S) < \varphi$. G is called a φ**-expander** iff it has no φ-sparse cuts, or equivalently $\phi(G) \geqslant \varphi$.

Note that sometimes a graph is described as an "expander" where φ is not specified. Though this does not have a rigorous definition, this usually means that φ is quite large. In general, graph G is considered to be a "good" expander if $\phi = \Omega(1)$ or $\phi \geqslant \frac{1}{\mathrm{polylog}(n)}$. Most of the results in this lecture do not require assumptions about the size of φ. Also note that in some settings we want to work with sparse expanders specifically. However, this is not always the case and is not a consideration for the results in this lecture.

1.2 Examples

Let us view a couple of example graphs and their conductance. In each example we would like to find the cut with the minimum conductance, which means minimizing the number of cut edges $E(S, \overline{S})$, and maximizing $\min(\mathrm{vol}(S), \mathrm{vol}(\overline{S}))$, which is equivalent to making $\mathrm{vol}(S)$ and $\mathrm{vol}(\overline{S})$ as equal as possible.

Example 1: Paths. If $G = P_n$, the path on n vertices, then clearly any minimum cut will use only one edge, and so the minimum cut will be when S is the first half of the vertices. This gives a conductance of $O(1/n)$ and so P_n can only be an $O(1/n)$-expander.

Example 2: Stars. If $G = K_{1,t}$, the star on $t + 1$ vertices then one side of the cut must only contain leaves of the graph, say this is S. Then the number of cut edges is exactly $|S|$ and similarly, the volume of

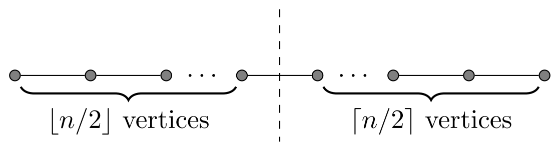

S would also be $|S|$ while $\mathrm{vol}(\overline{S}) = t + (t - |S|) \geqslant |S|$. Therefore the conductance of any cut S in a star is $|S| / \min(|S|, 2t - |S|) \geqslant 1$ and so $K_{1,t}$ is a 1-expander.

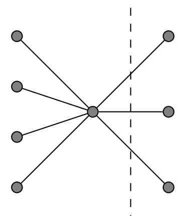

Example 3: Cycles. Similar to the path, if $G = C_n$ the cycle on n vertices, at minimum we need two cut edges. Then by splitting the vertices evenly we obtain the minimum conductance $O(1/n)$. So C_n is also an $O(1/n)$-expander.

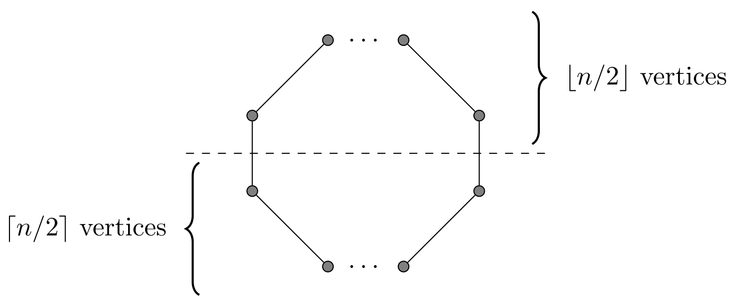

Example 4: Cliques. Let $G = K_n$, the complete graph on n vertices. Then for any subset of the vertices S, the number of cut edges is $|S||\overline{S}|$, and $\mathrm{vol}(S) = 2\binom{|S|}{2} + |S||\overline{S}|$. Assuming that $|S| \leqslant |\overline{S}|$, this gives a conductance of

$$\phi(S) = \frac{|S||\overline{S}|}{|S|(|S|-1) + |S||\overline{S}|} = \frac{|S|(n-|S|)}{|S|(|S|-1) + |S|(n-|S|)}.$$

Noticing that $|S| \leqslant n/2$ it can be shown that $\phi(S) \geqslant 1/2$ and so K_n is a $(1/2)$-expander.

Example 5: Dumbbell graphs. Lastly, if G is a dumbbell graph on $2n$ vertices (two disjoint copies of K_n with a single edge connecting them), then the obvious smallest cut is the single edge between the two cliques. This gives us a conductance of $O(1/n^2)$ meaning that dumbbell graphs are only $O(1/n^2)$-expanders.

2 Some Structural Properties of Expanders

2.1 Characterizing Expanders

Think about what the examples of graphs of high conductance in Section 1.2 have in common. We might conjecture that expanders are "well-connected," but it's not clear what notion of connectedness we want. For example, we saw that the star $K_{1,t}$ and the complete graph K_n are both graphs of high conductance, while the dumbbell graph has low conductance. It's easy to intuitively argue that K_n is "well-connected,"

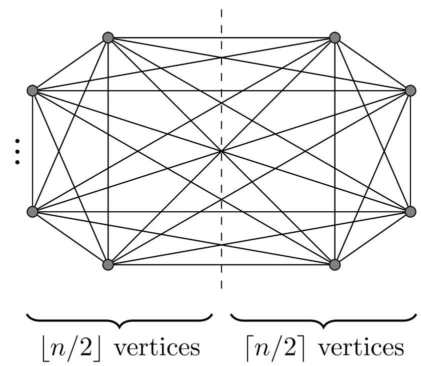

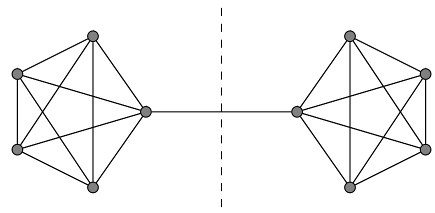

but it's less clear how to do so for $K_{1,t}$. For instance, there is a short path between any two vertices in a star, but this is also the case in the dumbbell graph.

Instead, we consider a property which we will show is almost equivalent to expansion. Given a graph, we ask if it is possible to remove a limited number of edges and get disconnected components each with a large proportion of edges. The dumbbell graph satisfies this property because we can remove a single edge (the one between the cliques), and get two disconnected, high-density components. Our high conductance graphs, however, do not satisfy this property. We can remove a single edge to disconnect $K_{1,t}$, but only one component has all of the remaining edges. Similarly, in K_n we must remove a large number of edges to disconnect the graph at all.

The next result formalizes this intuition and gives a sense of the "robustness" of expanders.

Theorem 3 (**Expander Characterization via "Robustness"**). *Let a connected graph $G = (V, E)$ and a value $\varphi \in [0, 1]$ be given. Then G is a φ-expander if and only if for all subset $D \subseteq E$ of edges, the following holds. Let $H = G \setminus D$ be the graph obtained by removing edges of D from G, and let $C_1, \ldots, C_k$ be the connected components of H. We say that component C_i for $i \in [k]$ is* small *if $vol_G(C_i) \leqslant vol_G(V)/2$. Then,*

$$\sum_{\substack{i\in[k] \\ C_i \text{ is small}}} vol_G(C_i) = O\left(\frac{|D|}{\varphi}\right).$$

In other words, Theorem 3 says that G is a φ-expander if and only if disconnecting the graph into high-density components costs a lot of edges. For the purpose of this lecture, we will treat Theorem 3 as a "practical definition" of an expander. We should emphasize that using the Oh-notation in this "characterization" is an abuse of notation, but we will not get into the precise details here.

Proof of the "if" direction of Theorem 3. Assume G is a φ-expander and consider any choice of D, $H = G\backslash D$ and the resulting connected components $C_1, \ldots, C_k$. For $i \in [k]$ where C_i is small, we have, $\phi_G(C_i) \geqslant \varphi$ (since G is a φ-expander) and so we have $\text{vol}(C_i) \leqslant \varphi^{-1} \cdot |E(C_i, \overline{C}_i)|$ (since C_i is the smaller side of $(C_i, \overline{C}_i)$).

Thus, we have,

$$\sum_{\substack{i\in[k]\\ C_i \text{ is small}}} \mathrm{vol}_G(C_i) \leqslant \sum_{\substack{i\in[k]\\ C_i \text{ is small}}} \varphi^{-1}\cdot|E(C_i,\overline{C}_i)| = O(\frac{|D|}{\varphi}),$$

where the last equation is because making C_i a connected component in H requires deleting all edges of the cut $E(C_i,\overline{C}_i)$ and these edges can be counted at most twice in the sum. □

Proof of the "only if" direction of Theorem 3. Suppose G is not a φ-expander and let $(S,\overline{S})$ be a φ-sparse cut in G. Let $D=E(S,\overline{S})$ be the cut edges of S. Thus, assuming S is the smaller side of the cut, we can remove $|D|<\varphi\cdot\mathrm{vol}(S)$ edges and separate $\mathrm{vol}(S)\leqslant\mathrm{vol}(V)/2$ volume from the rest of the graph. But this means we now have more than $|D|/\varphi$ many edges separated by deleting only $|D|$ edges. □

2.2 Diameter of an Expander

Another notion of graph connectedness comes in the form of the graph diameter.

Definition 4. The **diameter** of a (connected) graph $G=(V,E)$, is defined as follows:

$$\mathrm{diam}(G)=\max_{u,v\in V}\mathrm{dist}_G(u,v).$$

In other words, the diameter of a graph is the length of the longest shortest uv-path over all choices of vertices u and v. In Section 1.2, we saw that the concept of diameter is not enough to characterize high conductance graphs (since both a star and a dumbbell graph have low diameters). However, for expanders we do have the following result[1]:

Proposition 5. *For any φ-expander $G=(V,E)$, we have that $diam(G)=\mathcal{O}(\log(m)/\varphi)$.*

This result shows that in an expander, all vertices are "close" to each other. The proof uses a "ball growing" technique that is helpful for many problems that concern with distance in a graph.

Proof of Proposition 5. We start with the following notation. Given a graph $G=(V,E)$ and $v\in V$, for any integer $r\geqslant 0$, the *ball of radius r centered at vertex v* is defined as follows:

$$B(v,r)=\{u\in V\mid \mathrm{dist}_G(u,v)\leqslant r\}.$$

Let G be a connected φ-expander. For any $v\in V$, define $r_v\geqslant 1$ such that

$$\mathrm{vol}(B(v,r_v-1))\leqslant\frac{\mathrm{vol}(V)}{2} \text{ and } \mathrm{vol}(B(v,r_v))>\frac{\mathrm{vol}(V)}{2}.$$

Note that since G is connected, r_v is well-defined for every v. Let $s,t\in V$.

Claim 6. $B(s,r_s)\cap B(t,r_t)\neq\emptyset$.

Proof of Claim 6. Assume $B(s,r_s)\cap B(t,r_t)=\emptyset$. This implies

$$\mathrm{vol}(B(s,r_s))+\mathrm{vol}(B(t,r_t))\leqslant\mathrm{vol}(V).$$

However, by our choice of r_s and r_t we have

$$\mathrm{vol}(B(s,r_s))+\mathrm{vol}(B(t,r_t))>\mathrm{vol}(V),$$

a contradiction. □

[1]As always, in the following, we have $m=|E|$ and $n=|V|$.

Therefore, there exists some $v \in V$ such that $v \in B(s, r_s) \cap B(t, r_t)$. This means that there is an sv-path of length at most r_s and a vt-path of length at most r_t in G. Hence, $\mathrm{dist}(s,t) \leqslant r_s + r_t$.

Claim 7. *For every $v \in V$, we have $r_v \leqslant \frac{2 \log m}{\phi}$.*

Proof of Claim 7. Let $r < r_v$. Note that by the definition of r_v, we know $\mathrm{vol}(B(v,r)) \leqslant \frac{\mathrm{vol}(V)}{2}$. Consider the cut $(B(v,r), V \setminus B(v,r))$. Since G is a φ-expander, we know

$$|E(B(v,r), V \setminus B(v,r))| \geqslant \varphi \cdot \mathrm{vol}(B(v,r)).$$

This means that if we expand the radius of the ball centered at v, we will add all of the cut edges $E(B(v,r), V \setminus\ B(v,r))$ to the new ball. Thus,

$$\mathrm{vol}(B(v, r+1)) \geqslant (1+\varphi) \cdot \mathrm{vol}(B(s,r)).$$

Each time we increase the radius r by 1, we increase the volume of the ball centered at v by a factor of $(1+\varphi)$. Hence, we have,

$$\mathrm{vol}(B(v,r)) \geqslant (1+\varphi)^r \cdot 1 \geqslant \exp\left(\frac{\varphi}{2}\right),$$

where we used the inequality $1 + x \geqslant e^{x/2}$ for $x \in (0,1)$. Thus, after $r \geqslant 2 \log m / \varphi$, we will get $\mathrm{vol}(B(v,r)) \geqslant m$, which implies $r_v \leqslant 2 \log m / \varphi$ as desired. $\square$

The proof of Proposition 5 now follows because

$$\mathrm{dist}_G(s,t) \leqslant r_s + r_t \leqslant \frac{2 \log m}{\varphi} + \frac{2 \log m}{\varphi} = O(\frac{\log m}{\varphi}).$$

Since this is true for every $s, t \in V$, we know $\mathrm{diam}(G) = O(\frac{\log m}{\varphi})$, as desired. $\square$

Remark. The converse direction of Proposition 5 does not hold, as in, we can have graphs with a very small diameter that are not all good expanders. The class of dumbbell graphs forms one counterexample.

3 A Simple Algorithmic Application of Expanders

To show case the "power" of using expanders and expander decompositions (to be defined shortly), let us see one simple example in this lecture in the context of the minimum cut problem.

Problem 1. Given an undirected simple graph $G = (V, E)$, return a **minimum cut** $(S, \overline{S})$ in G.

We will see a deterministic algorithm with runtime $m^{1+o(1)}$ for this problem due to [Sar21]. Let us start by seeing why expanders may help in solving the minimum cut problem.

Notation. Throughout the rest of this section, let $\delta(G)$ denote the minimum degree of G and $\lambda(G)$ denote the size of a minimum cut in G. Notice that $\lambda(G) \leqslant \delta(G)$ as we can always consider the singleton cut consisting of the minimum degree vertex on one side. We say that a cut is **non-singleton** if it consists of at least two vertices on each side.

3.1 Minimum Cut in in Expanders

Let $G = (V, E)$ be a simple graph and $\delta = \delta(G)$ be its minimum degree. Suppose G is a $(2/\delta)$-expander (so, not even really a "very good" expander if δ is large). How can we find a minimum cut of G? The following simple result shows that the structure of minimum cuts in such expanders is extremely basic.

Proposition 8. *Let $G = (V, E)$ be a simple $(2/\delta(G))$-expander; all minimum cuts of G are singleton.*

Proof. Let $C \subseteq V$ be a vertex set inducing a minimum cut in G. That is, $|E(C, V \setminus C)| = \lambda(G)$. Consider the conductance of the cut induced by C. By the definitions of conductance of cuts and graphs respectively,

$$\varphi(G) \leqslant \varphi_G(C) = \frac{|E(C, V \setminus C)|}{\min\{\text{vol}(C), \text{vol}(V \setminus C)\}} = \frac{\lambda(G)}{\min\{\text{vol}(C), \text{vol}(V \setminus C)\}}.$$

Without loss of generality, assume that $|C| \leqslant |V \setminus C|$. Then, it follows that

$$\begin{aligned}
\delta(G) &\geqslant \lambda(G) && \text{(as argued in the "notation" part above)}\\
&\geqslant \varphi(G) \cdot \min\{\text{vol}(C), \text{vol}(V \setminus C)\} && \text{(by the above inequality)}\\
&= \varphi(G) \cdot \min\{\sum_{v \in C} \deg_G(v), \sum_{v \in V \setminus C} \deg_G(v)\} && \text{(by the definition of volume)}\\
&\geqslant \varphi(G) \cdot \min\{|C| \cdot \delta(G), |V \setminus C| \cdot \delta(G)\} && \text{(as } \delta \text{ is the minimum degree)}\\
&= \varphi(G) \cdot |C| \cdot \delta(G) && \text{(by the assumption that } |C| \leqslant |V \setminus C|)\\
&\geqslant \frac{2}{\delta(G)} \cdot |C| \cdot \delta(G) && \text{(by the expansion of } G)\\
&= 2 \cdot |C|.
\end{aligned}$$

As such, we have that every minimum cut C should have at most $\delta(G)/2$ vertices at most. We now argue that this means that actually C can only be of size one.

Suppose by contradiction that $|C| \geqslant 2$. Consider any vertex $v \in C$. We have,

$$|E(v, V \setminus C)| \geqslant \delta(G) - (|C| - 1) > \delta(G)/2,$$

where the first inequality is because G is a simple graph and thus only $|C| - 1$ edges of v can be inside C, and the second is because $|C| \leqslant \delta(G)/2$. But the conclusion is now that if we have more than one vertex in C, then we will have,

$$|E(C, V \setminus C)| > |C| \cdot \cdot \delta(G)/2 \geqslant \delta(G) \geqslant \lambda(G),$$

which means C cannot be a minimum cut (given the strict inequality). Hence, the only minimum cuts of C are singletons, concluding the proof. ☐

Notice that a direct corollary of Proposition 8 is that we can find a minimum cut in any $(2/\delta(G))$-expander in linear time – it suffices to check the values of the cuts induced by each of the vertices and return the singleton cut consisting of a minimum degree vertex.

3.2 Expander Decompositions or "Expanders in Non-Expander Graphs"

Hopefully, at this point in the lecture, you are convinced that expanders are pretty cool and useful. But of course, not every graph we work with is an expander. Can we "lift" some of the insights we have about expanders to non-expander graphs also? It turns out the answer is *Yes* through the notion of *expander decompositions*, a technique that has been extremely fruitful in the design of "modern" graph algorithms[2].

[2] Although this technique has been around for quite some time now, it got revamped in a couple of recent work in the last five years or so, and has been instrumental in various recent breakthroughs on graph algorithms.

Definition 9 (**Expander decomposition**)**.** Let $G = (V, E)$ be any arbitrary undirected graph. An **expander decomposition** of G with parameters $\varphi, M > 0$ is a partitioning of the vertex set V into sets $X_1, \ldots, X_k$ such that

1. $G[X_i]$ is a φ-expander for all $i \in [k]$; and
2. $\sum_{i=1}^{k} |E(X_i, V \setminus X_i)| \leqslant M$.

In other words, G can be partitioned into a collection of vertex disjoint φ-expanders by deleting at most M edges from G.

One intuitive approach to construct expander decompositions is the following; for our desired φ, if G is not a φ-expander, then G must have a φ-sparse cut. Delete the edges of this sparse cut and recursively find expander decompositions for the two vertex sets inducing this cut, and combine them to obtain the expander decomposition for G. Given we are cutting edges of a sparse cut, we hope that we will not be deleting many edges of G and use this to bound the number of deleted edges in total. This idea is formalized in the following algorithm;

Algorithm 1. Decompose(G, φ):

1. If G is a φ-expander, return G.
2. Otherwise, let $(S, \overline{S})$ be a φ-sparse cut in G.
3. Return Decompose($G[S]$, φ) $\sqcup$ Decompose($G[\overline{S}]$, φ).

Lemma 10. *Given a graph $G = (V, E)$ and $\varphi > 0$, Algorithm 1 returns a (φ, M)-expander decomposition of G for the parameter $M = O(\varphi \cdot m \log m)$.*

Proof. Property 1 of the expander decomposition follows directly from the base case of the algorithm. Thus, we only focus on proving Property 2.

We proceed by a charging argument. Let S induce a φ-sparse cut as given in step 2 of Algorithm 1. Moreover, without loss of generality, assume that $\mathrm{vol}_G(S) \leqslant \mathrm{vol}_G(V \setminus S)$. Then, by the definitions of φ-expanders and φ-sparse cuts, we obtain that $|E(S, V \setminus S)| \leqslant \varphi \cdot \mathrm{vol}_G(S)$.

We *charge* each vertex $u \in S$ with $\varphi \cdot \deg_G(u)$ charge. Moreover, let charge(v) denote a "charge" that we define at a vertex $v \in V$. Then, using that $|E(S, V \setminus S)| \leqslant \varphi \cdot \mathrm{vol}_G(S)$, we obtain

$$|E(S, V \setminus S)| \leqslant \phi \cdot \mathrm{vol}_G(S) = \varphi \cdot \sum_{v \in S} \deg_G(v) = \sum_{v \in S} \mathrm{charge}(v).$$

Hence, we can assign each edge in $E(S, V \setminus S)$ to a specific unit of charge. In other words, the charge assigned to vertices in this step is at least as large as the number of deleted edges of the cut $(S, \overline{S})$.

We note that we only assign charge to the vertices in tho component of smaller volume for any given cut. Therefore, each vertex receives charge in at most $\log(2m)$ iterations of the algorithm (as the volume of the unique set containing the vertex drops by half each time and the volume starts at $2m$). Therefore, the final total charge assigned to any vertex $v \in V$ is at most $\varphi \cdot \log 2m \cdot \deg_G(v)$.

Then, since each edge that is contained in a φ-sparse cut in any iteration of the algorithm is assigned to a specific unit of charge, we obtain that

$$\sum_{i \in [k]} |E(X_i, V \setminus X_i)| \leqslant \sum_{v \in V} \mathrm{charge}(v) \leqslant \sum_{v \in V} \varphi \cdot \log(2m) \cdot \deg_G(v) = O(\varphi \cdot m \log m),$$

as desired. This concludes the proof. $\square$

Unfortunately, Algorithm 1 does not lead to an efficient (i.e., poly-time) algorithm for expander decomposition given that finding a sparse cut is $\mathcal{NP}$-hard in general. In the next lecture, we will study efficient algorithms for expander decomposition. For this lecture however, we will use the following efficient implementation of expander decompositions due to [CGL+20].

Proposition 11. *There exists a deterministic algorithm that given any undirected graph $G = (V, E)$ and parameter $\varphi \in (0, 1)$, in $m^{1+o(1)}$ time outputs a partitioning of V into sets $X_1, \ldots, X_k$ such that:*

1. *For every $i \in [k]$ and every $S \subseteq X_i$, we have $|E(S, X_i \setminus S)| \geqslant \varphi \cdot \min\{vol_G(S), vol_G(X_i \setminus S)\}$;*
2. *$\sum_{i=1}^{k} |E(X_i, V \setminus X_i)| = \varphi \cdot m^{1+o(1)}$.*

Two remarks about Proposition 11 are in order.

Firstly, the guarantee in Property 1 of this result is even stronger than the original expander decomposition; we not only have each $G[X_i]$ is a φ-expander but in fact the conductance of each cut in $G[X_i]$ is at least φ even if we divide by the volumes inside G and not only $G[X_i]$ (you are encouraged to reprove Lemma 10 but with such a guarantee instead; the current proof in fact already does this given that we simple use a charge proportional to the *original* degrees of vertices and not their degrees in the "recursed" graph[3]).

Secondly, while this algorithm is quite efficient in terms of runtime and has an *almost linear time*, the guarantee it provides in terms of number of deleted edges is worse than the existential bounds of Lemma 10 by a $m^{o(1)}$ factor (we will see some reason behind this loss in the next lecture).

3.3 Minimum Cuts in General Graphs via Expander Decompositions

We now show how to extend the simple ideas behind finding minimum cuts in expanders, to all arbitrary graphs using expander decompositions. The idea is to construct a smaller graph G' with fewer edges than G while preserving *non-singleton* minimum cuts of G in G' as well. We will then be able to run conventional min-cut algorithms on this reduced graph to efficiently determine a minimum cut for G. Specifically, we will use the following algorithmic min-cut result of [Gab95] for the last step.

Proposition 12 ([Gab95, Theorem 2.2])**.** *There exists an algorithm that finds a minimum cut in a graph $G = (V, E)$ in $O(m \cdot \lambda(G))$ time.*

Vertex contraction. To talk about our algorithm, we need to recall a simple operation on graphs. Let $G = (V, E)$ be a graph, and let $S \subseteq V$ be a set of vertices in G. We say that a graph $G' = (V', E')$ is obtained from G by **contracting** set S if

- $V' = V \setminus S \cup \{v_S\}$; and
- $E' = \{(u, v) \in E \,:\, u, v \notin S\} \cup \{(u, v_S) \,:\, u \in V \setminus S,\, v \in S\}$.

Equivalently, G' may be obtained by iteratively identifying all vertices in S. Note that, if G is simple, G' may contain parallel edges, but not self-loops. See Section 3.3 for an illustration.

The following proposition examines the effect of contraction to minimum cuts.

[3] In fact, one can say something stronger. Any algorithm that can find a (φ, M)-decomposition can also be used to find a decomposition of the type in Proposition 11 by only increasing M by a constant factor. The idea is to add $\deg(v)$ self-loops to each vertex $v \in V$ first and then run the (φ, M)-decomposition on this new graph. It is easy to see that this leads to the type of the stronger guarantee we require, while only increasing the number of deleted edges by a constant factor.

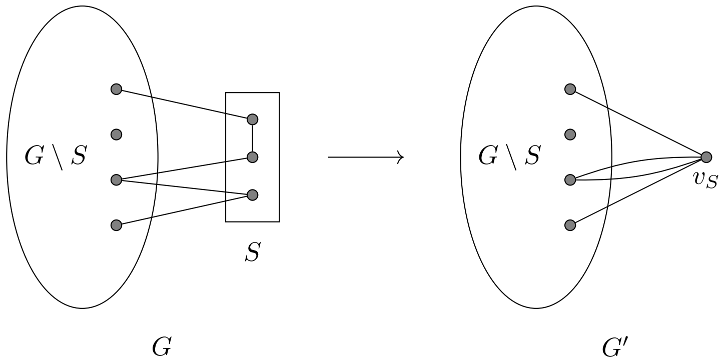

Figure 1: Illustration of vertex set contraction.

Proposition 13. *Let $(C, V \setminus C)$ be a minimum cut in a graph $H = (V_H, E_H)$. Suppose S is a set of vertices such that either $S \subseteq C$ or $S \cap C = \emptyset$. Then, the size of minimum cut in the graph obtained by contracting S in H is the same as the minimum cut size in H. Also, regardless of the choice of S, minimum cut size does not decrease after contraction.*

Proof. This proposition follows directly from the definition of vertex set contraction, and since the cuts after contraction can be mapped to the cuts of the original graph by "expanding" the contracted vertex. $\square$

The algorithm. Equipped with Proposition 13, we can now design our minimum cut algorithm. If G has any singleton minimum cut, finding the minimum degree vertex only solves the minimum cut problem, so, without loss of generality, we focus on non-singleton cuts only.

Let $(C, V \setminus C)$ be a non-singleton minimum cut of G. The goal of the algorithm is to find a sequence of vertex set contractions such that after those, in the resulting graph $G' = (V', E')$, we have,

1. $|E'| = m^{1+o(1)}/\delta(G)$; and,
2. In each contraction with a set S, we have $S \subseteq C$ or $S \cap S = \emptyset$, or in other words,
$$\min \{|S \cap C|, |S \setminus C|\} = 0.$$

The second property, together with Proposition 13 ensures that $(C, V \setminus C)$ remains a minimum cut in G' as well. But, now, running Proposition 12 on the graph G', finds its minimum cut (and by expanding the contracted vertices, we can find a minimum cut of G as well). Moreover, the factor $\delta(G)/m^{o(1)}$ saving in the number of edges in G' implies that the runtime of the final algorithm is going to be $O(m^{1+o(1)}/\delta(G) \cdot \delta(G)) = m^{1+o(1)}$ time, as desired.

Of course, the entire challenge in the design of the algorithm is that we do not the cut C a priori and thus the contractions should achieve the above properties with no knowledge of C. This is where expander decompositions come handy. In the following, we first run Proposition 11 with parameter $\varphi = 40/\delta(G)$ and thus start the algorithm given a set $X_1, \ldots, X_k$ of expanders with the properties mentioned in Proposition 11. We then use two "cleaning up" procedures on each expander as follows.

Algorithm 2. TRIM(G, S): given a graph $G = (V, E)$ and a set $S \subseteq V$:

- While there exists a vertex $v \in S$ with $\deg_{G[S]}(v) < 2/5 \cdot \deg_G(v)$, update $S \leftarrow S \setminus \{v\}$; return S.

Algorithm 3. $\textsc{Shave}(G, S)$: given a graph $G = (V, E)$ and a set $S \subseteq V$:

1. Return $\{v \in S : \deg_{G[S]}(v) > \deg_G(v)/2 + 1\}$

The full algorithm for reducing G to the smaller graph G' is as follows.

Algorithm 4. $\textsc{MinCutReduction}(G)$:

1. Let $X_1, \ldots, X_k$ be the sets returned by Proposition 11 with the parameter $\varphi = 40/\delta(G)$.
2. For every $i \in [k]$, let $X_i' \leftarrow \textsc{Trim}(X_i)$ and $X_i'' \leftarrow \textsc{Shave}(X_i')$, and contract X_i''.
3. Return the contracted graph as G'.

Fix a non-singleton minimum cut $(C, V \setminus C)$. Our goal is to show that no contraction "cuts" across the cut C, i.e., for every $i \in [k]$, we have

$$\min\{|X_i'' \cap C|, |X_i'' \setminus C|\} = 0,$$

so that we can apply Proposition 13 and argue G' and G have the same minimum cut value. We will then bound the number of edges remaining in G' also and conclude the proof. These two parts are done in the following two lemmas.

Lemma 14. *For any non-singleton minimum cut $(C, V \setminus C)$ and every $i \in [k]$,*

$$\min\{|X_i'' \cap C|, |X_i'' \setminus C|\} = 0.$$

Proof. We prove this in three step, each corresponding to the intersection sizes of X_i, X_i' and X_i'' with C, respectively.

Step 1. We first prove that

$$\min\{|X_i \cap C|, |X_i \setminus C|\} \leqslant \lambda(G)/40. \tag{1}$$

We have,

$$\begin{aligned}
\lambda(G) &\geqslant |E(X_i \cap C, X_i \setminus C)| && \text{(as all the edges counted here cross the minimum cut)}\\
&\geqslant 40/\delta(G) \cdot \min\{\mathrm{vol}_G(X_i \cap C), \mathrm{vol}_G(X_i \setminus C)\} \\
& && \text{(by the Property 1 of Proposition 11 with } \varphi = 40/\delta(G))\\
&= 40/\delta(G) \cdot \min\{\sum_{v \in X_i \cap C} \deg_G(v), \sum_{v \in X_i \setminus C} \deg_G(v)\}\\
&\geqslant 40/\delta(G) \cdot \min\{\delta(G) \cdot |X_i \cap C|, \delta(G) \cdot |X_i \setminus C|\}\\
&= 40 \cdot \min\{|X_i \cap C|, |X_i \setminus C|\}.
\end{aligned}$$

Hence, $\min\{|X_i \cap C|, |X_i \setminus C|\} \leqslant \lambda(G)/40$.

Step 2. We then show that

$$\min\{|X_i' \cap C|, |X_i' \setminus C|\} \leqslant 2. \tag{2}$$

Assume by symmetry that $|X_i \cap C| \leqslant |X_i \setminus C|$ in Eq (1). We have,

$$\delta(G) \geqslant \lambda(G) \qquad \text{(as argued earlier in the lecture)}$$

$$
\begin{aligned}
&\geqslant |E(X_i' \cap C, X_i' \setminus C| && \text{(as all these edges cross the minimum cut)}\\
&= |E(X_i' \cap C, X_i')| - 2 \cdot |E(X_i' \cap C, X_i' \cap C)|\\
& && \text{(by partitioning the edges between } X_i' \cap C \text{ and } X_i' \text{ accordingly)}\\
&\geqslant |E(X_i' \cap C, X_i')| - 2 \cdot |X_i' \cap C|^2,
\end{aligned}
$$

where the $2 \cdot |X_i' \cap C|^2$ term is an upper bound on the number of edges with both endpoints in $X_i' \cap C$. Thus,

$$
\begin{aligned}
\delta(G) &\geqslant |E(X_i' \cap C, X_i')| - 2 \cdot |X_i' \cap C|^2\\
&\geqslant \frac{2\delta(G)}{5} \cdot |X_i' \cap C| - 2|X_i' \cap C|^2 && \text{(by the termination condition for } \textsc{Trim}(G,\, X_i))\\
&\geqslant \frac{2\delta(G)}{5} \cdot |X_i' \cap C| - 2|X_i \cap C| \cdot |X_i' \cap C| && \text{(by a crude upper bound of } |X_i' \cap C| \leqslant |X_i \cap C|)\\
&\geqslant \frac{2\delta(G)}{5} \cdot |X_i' \cap C| - \frac{\delta(G)}{20} \cdot |X_i' \cap C|. && \text{(by Eq (1) and } \lambda(G) \leqslant \delta(G))
\end{aligned}
$$

This implies that

$$
|X_i' \cap C| \leqslant \frac{20}{7},
$$

and since the LHS is an integer, we obtain $|X_i' \cap C| \leqslant 2$.

Step 3. Finally, we prove that

$$
\min\{|X_i'' \cap C|, |X_i'' \setminus C|\} \leqslant 2. \tag{3}
$$

Again, assume by symmetry that $|X_i' \cap C| \leqslant |X_i' \setminus C|$ in Eq (2).

Suppose for the sake of contradiction that $\min\{|X_i'' \cap C|, |X_i'' \setminus C|\} > 0$. Then, let $u \in X_i'' \cap C$. By the operation defined by $\textsc{Shave}(G,\, X_i')$, we note that $\deg_{G[X_i']}(u) > \deg_G(u)/2 + 1$. Moreover, since $|X_i' \cap C| \leqslant 2$, we find that u has strictly more than $\deg_G(u)/2$ neighbors in $X_i' \setminus C$. However, then the cut $(C \setminus \{u\}, V \setminus C \cup \{u\})$ contradicts the minimality of cut $(C, V \setminus C)$ (in general, any vertex in a non-singleton minimum cut cannot have more neighbors across the cut vs inside it, as otherwise we can move the vertex to the other side and decrease the size of the cut[4]).

By this contradiction, we can conclude that $\min\{|X_i'' \cap C|, |X_i'' \setminus C|\} = 0$, as desired. □

Lemma 15. *We have,*

$$
\sum_{i=1}^{k} |E(X_i'', V \setminus X_i'')| = m^{1+o(1)}/\delta(G).
$$

Proof. We again prove this in three steps corresponding to the choices of each X_i, X_i', and X_i'', respectively.

Step 1. We first have,

$$
\sum_{i=1}^{k} |E(X_i, V \setminus X_i)| = m^{1+o(1)}/\delta(G). \tag{4}
$$

This follows immediately from Property 2 of Proposition 11. In the next two steps, we show that since the vertices removed from X_i during the algorithm already contributed "many" edges to outside X_i, adding their remaining edges to the mix also does not increase the number of remaining edges by a lot.

[4]Notice that this is the only place we use the fact that the cut is non-singleton; obviously, we cannot apply the same "switching" argument with a singleton cut as we will end up with an empty cut.

Step 2. We then have,

$$\sum_{i=1}^{k} |E(X'_i, V \setminus X'_i)| = m^{1+o(1)}/\delta(G). \tag{5}$$

We claim

$$\sum_{i=1}^{k} |E(X'_i, V \setminus X'_i)| \leqslant 2 \cdot \sum_{i=1}^{k} |E(X_i, V \setminus X_i)|,$$

which immediately implies the desired bound by Eq (4) (the leading constant of 2 will be absorbed in the $m^{o(1)}$ term). Let v be the first vertex $v \in X_i$ which is deleted from X'_i. We have,

$$|E(X_i \setminus \{v\}, V \setminus X_i \cup \{v\})| \leqslant |E(X_i, V \setminus X_i)| - \deg_G(v)/5,$$

because v had at least $3\deg_G(v)/5$ edges going out of X_i and now those edges are no longer counted in the LHS, but the $\leqslant 2\deg(v)/5$ edges of v inside X_i will now be counted. Thus, each deletion of v from X'_i, decreases the total number of edges out of $E(X_i, V \setminus X_i)$ by at least $\deg_G(v)/5$, while contributing at most $2\deg_G(v)/5$ edges to the remaining edges. Hence, size of the latter set can only be twice as large as the former one, concluding the proof.

Step 3. Finally, we have,

$$\sum_{i=1}^{k} |E(X''_i, V \setminus X''_i)| = m^{1+o(1)}/\delta(G). \tag{6}$$

We again claim that

$$\sum_{i=1}^{k} |E(X''_i, V \setminus X''_i)| \leqslant (1 + o(1)) \cdot \sum_{i=1}^{k} |E(X'_i, V \setminus X'_i)|,$$

which immediately concludes the proof of the lemma using Eq (5) (and absorbing the $(1 + o(1))$ factor in the $m^{o(1)}$ term). We can assume without loss of generality that $\delta(G) = \omega(1)$ as otherwise the entire lemma vacuously holds. For any vertex v removed from X'_i to obtain X''_i, v already contributed $\deg(v)/2 - 1 \geqslant (1 - o(1)) \cdot \deg(v)/2$ edges to outside of X'_i. By removing v, we can at most add another $(1 + o(1)) \cdot \deg(v)/2$ more edges to the mix, which increased the number of remaining edges by only a $(1 + o(1))$ factor. Thus, the claim follows immediately from this. This concludes the proof of the lemma. □

Combining all the above, we can prove the following result of [Sar21].

Theorem 16 ([Sar21])**.** *There is a deterministic $m^{1+o(1)}$ time algorithm for the minimum cut problem in undirected simple graphs.*

References

[CGL+20] Julia Chuzhoy, Yu Gao, Jason Li, Danupon Nanongkai, Richard Peng, and Thatchaphol Saranurak. A deterministic algorithm for balanced cut with applications to dynamic connectivity, flows, and beyond. In Sandy Irani, editor, *61st IEEE Annual Symposium on Foundations of Computer Science, FOCS 2020, Durham, NC, USA, November 16-19, 2020*, pages 1158–1167. IEEE, 2020. 9

[Gab95] H.N. Gabow. A matroid approach to finding edge connectivity and packing arborescences. *Journal of Computer and System Sciences*, 50(2):259–273, 1995. 9

[Sar21] T. Saranurak. A simple deterministic algorithm for edge connectivity. In *Symposium on Simplicity in Algorithms (SOSA)*, pages 80–85. SIAM, 2021. 6, 13

[SW19] T. Saranurak and D. Wang. Expander decomposition and pruning: Faster, stronger, and simpler. *Proceedings of the 2019 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 2616–2635, 2019.