Computer Science 2: Advanced Algorithms Fall 2026

Lecture 5 — September 8, 2026

Prof. Jarosław Błasiok *Scribe: Tommaso Giacomello*

1 Overview

In the last lecture we studied a graph from three different points of view: the *geometric* one (cuts and their boundaries), the *spectral* one (eigenvalues of the adjacency and Laplacian matrices) and the *dynamical* one (the random walk on the graph). We then proved Cheeger's inequality, which relates these three points of view to one another.

In this lecture we put that machinery to work on a concrete computational problem. The problem is **undirected** s–t **connectivity** (USTCON): given an undirected graph G and two vertices s, t, decide whether they lie in the same connected component. BFS and DFS solve it in linear time, but both have to remember which vertices have already been visited. Our question is instead:

*how much **memory** does USTCON need, if we are allowed as much time as we want?*

2 Math things

Here we collect the notation, the definitions and the facts from the previous lecture that will be used throughout these notes.

2.1 Setting: the matrices A, P, L_G and $\mathcal{L}_G$

Assumption 1 (D-regularity)**.** *Unless stated otherwise, $G = (V, E)$ is an undirected, D-regular multigraph on n vertices $V = \{1, \dots, n\}$, with parallel edges allowed and counted with multiplicity.*

Regularity is a convenience rather than a restriction: in Section 3.3 we show that an arbitrary graph can be made 3-regular in log space without changing its connected components. D-regularity gives good spectral properties to all the matrices we are about to consider, since the degree matrix is simply DI: the four matrices are simultaneously diagonalizable, and the uniform distribution is stationary for the walk.

Definition 1 (Adjacency, walk and Laplacian matrices)**.** *Let $A \in \mathbb{R}^{n \times n}$ be the **adjacency matrix**, $A_{ij} := \#\{\text{edges between } i \text{ and } j\}$. We write*

$$P := \tfrac{1}{D}A, \qquad L_G := DI - A, \qquad \mathcal{L}_G := \tfrac{1}{D}L_G = I - P.$$

*P is the **random-walk matrix**, $P_{ji} = \mathbb{P}[X_{t+1} = j \mid X_t = i]$; L_G is the **combinatorial Laplacian** and $\mathcal{L}_G$ the **normalized Laplacian**.*

Remark 1. *The symbol L_G is reserved for $DI - A$, as in Lecture 4. The normalized Laplacian $I - P$ is always written $\mathcal{L}_G$; it is the matrix whose eigenvalues define the gap γ.*

All these matrices are symmetric, hence orthogonally diagonalizable with real eigenvalues by the spectral theorem, and they share the same eigenvectors. We constantly use two elementary facts:

- P is doubly stochastic ($P\mathbf{1} = \mathbf{1}$ and $\mathbf{1}^T P = \mathbf{1}^T$), so it maps distributions to distributions and the uniform distribution $\pi := \mathbf{1}/n$ is stationary;
- $\mathbb{P}[X_t = i \mid X_0 = j] = (P^t)_{ij}$, so the t-step behaviour of the walk is read off from powers of P.

The combinatorial Laplacian L_G has a useful "energy" interpretation:

$$f^T L_G f \;=\; \sum_{\{i,j\}\in E} (f_i - f_j)^2 \;\geq\; 0, \tag{1}$$

so $L_G \succeq 0$ (positive semidefinite matrix), and (1) vanishes exactly when f is constant on every connected component.

2.2 Eigenvalues: range and meaning

Let $u_1, \ldots, u_n$ be the common orthonormal eigenbasis, with $u_1 = \mathbf{1}/\sqrt{n}$. The spectral properties of all these matrices are collected in the following table.

matrix	eigenvalues	relation	range	trivial eigenpair
A	$D = \theta_1 \geq \cdots \geq \theta_n$	—	$\theta_i \in [-D, D]$	$A\mathbf{1} = D\mathbf{1}$
$P = \frac{1}{D}A$	$1 = \mu_1 \geq \cdots \geq \mu_n$	$\mu_i = \theta_i/D$	$\mu_i \in [-1, 1]$	$P\mathbf{1} = \mathbf{1}$
$L_G = DI - A$	$0 = \lambda_1 \leq \cdots \leq \lambda_n$	$\lambda_i = D(1 - \mu_i)$	$\lambda_i \in [0, 2D]$	$L_G\mathbf{1} = 0$
$\mathcal{L}_G = I - P$	$0 = \gamma_1 \leq \cdots \leq \gamma_n$	$\gamma_i = 1 - \mu_i = \lambda_i/D$	$\gamma_i \in [0, 2]$	$\mathcal{L}_G\mathbf{1} = 0$

The four trivial eigenpairs, given by θ_1, μ_1, λ_1 and γ_1, are four points of view on the same trivial eigenpair, which is present for *every* D-regular graph.

The two ends of the spectrum have the following meaning.

- *The multiplicity of* $\mu = 1$ is the number of connected components. In particular G is connected if and only if $\mu_2 < 1$, i.e. $\gamma > 0$.
- $\mu_n = -1$ happens exactly when some connected component of G is bipartite. A random walk on a bipartite graph never converges to the uniform distribution, since it keeps oscillating between the two sides; the standard assumption that removes this obstruction is to make the walk *lazy* (Section 2.4).

Hence the trivial eigenpair is used to count connected components, and all the finer structure lies in the first nontrivial eigenvalue.

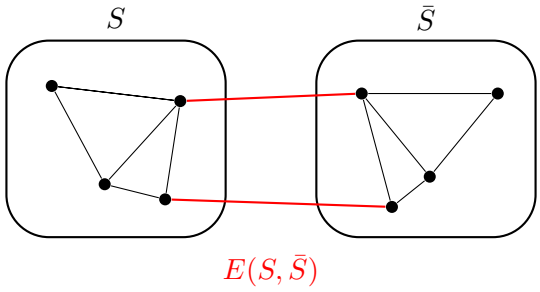
Definition 2 (Spectral gap)**.** *The* ***normalized spectral gap*** *of G is*

$$\gamma \;:=\; \gamma_2 \;=\; \frac{\lambda_2}{D} \;=\; 1-\mu_2.$$

When several graphs are in play we write $\gamma(G)$ for the gap of G, keeping subscripts on γ for eigenvalue indices.

2.3 Cuts, conductance and Cheeger's inequality

For $\emptyset \neq S \subseteq V$ let $E(S,\bar{S})$ denote the set of edges with exactly one endpoint in S and the other in $\bar{S}$, where $\bar{S} = V \setminus S$.



Definition 3 (Conductance)**.** *For $\emptyset \neq S \subseteq V$ with $|S| \leq n/2$,*

$$\phi(S) \;:=\; \frac{|E(S,\bar{S})|}{D|S|} \;\in [0,1], \qquad\qquad \phi_G \;:=\; \min_{\emptyset\neq S,\ |S|\leq n/2} \phi(S).$$

The normalization by $D|S|$ is what makes ϕ a probability: starting the walk at a uniformly random vertex of S, i.e. from $p_0 = \mathbf{1}_S/|S|$,

$$\mathbb{P}[X_1 \notin S] \;=\; \frac{\langle \mathbf{1}_{\bar{S}},\, P\mathbf{1}_S\rangle}{|S|} \;=\; \phi(S),$$

so conductance is the one-step escape probability from S. The minimizer of ϕ_G is the sparsest cut, that is, the separation with the fewest edges relative to its size, and equivalently the one with the smallest one-step escape probability. If ϕ_G is very small, then a random walk starting inside the corresponding set S has a low probability of leaving it and exploring the whole graph, and hence it will take longer to mix.

Theorem 1 (Cheeger's inequality)**.** *For every D-regular graph G with normalized spectral gap γ:*

$$\frac{\gamma}{2} \;\leq\; \phi_G \;\leq\; \sqrt{2\gamma}, \qquad \textit{equivalently} \qquad \frac{\lambda_2}{2D} \;\leq\; \phi_G \;\leq\; \sqrt{\frac{2\lambda_2}{D}}.$$

Rearranged,

$$\frac{\phi_G^2}{2} \;\leq\; \gamma \;\leq\; 2\phi_G.$$

We proved this in the last lecture and use it here as a black box. The top form is the convenient one when we have a spectral gap and want to conclude that the conductance is constant. The bottom form is the one we will actually use in Section 3.4, where we know a combinatorial fact about the cuts of the graph and want to deduce a lower bound on γ.

2.4 Spectral gap implies mixing

If G is connected and bipartite then $\mu_n = -1$ and $P^t p_0$ never converges. The standard fix is to make the walk **lazy**: with probability $1/2$ stay put, with probability $1/2$ take a step. Its matrix is

$$\tilde{P} \;:=\; \tfrac{1}{2}\,(I + P)\,, \qquad \tilde{\mu}_i = \tfrac{1+\mu_i}{2} \in [0,1], \qquad \tilde{\mu}_2 = 1 - \tfrac{\gamma}{2}.$$

All nontrivial eigenvalues now lie in $[0,1)$, so writing $p_0 = \pi + \sum_{i\ge 2} \alpha_i u_i$ in the eigenbasis (where u_i is a common eigenvector and α_i its coefficient in the linear combination), the trivial component survives and everything else decays geometrically, by the same mechanism as power iteration:

$$\|\tilde{P}^t p_0 - \pi\|_2 \;\le\; \left(1 - \tfrac{\gamma}{2}\right)^t \|p_0 - \pi\|_2 \;\le\; e^{-\gamma t/2}\, \|p_0 - \pi\|_2. \tag{2}$$

Bound (2) is stated in the ℓ_2 norm, whereas the distance between two distributions is usually measured in the following way.

Definition 4 (Total variation distance)**.** *Let p, q be two probability distributions on V. Their **total variation distance** is*

$$\|p - q\|_{TV} \;:=\; \tfrac{1}{2}\|p - q\|_1 \;=\; \frac{1}{2}\sum_{i \in V} |p_i - q_i| \;=\; \max_{S \subseteq V} |p(S) - q(S)|\,.$$

In words, $\|p - q\|_{TV}$ is the largest possible discrepancy between the two distributions on a single event. By Cauchy–Schwarz, $\|p - q\|_{TV} \le \frac{\sqrt{n}}{2}\|p - q\|_2$, which is how (2) is converted into the statement we actually use.

Corollary 1 (Mixing time)**.** *Using $\|p-q\|_{TV} \le \frac{\sqrt{n}}{2}\|p-q\|_2$, the lazy walk on a connected D-regular graph reaches total variation distance ε from uniform after*

$$t_{\mathrm{mix}}(\varepsilon) \;=\; O\!\left(\frac{\log n + \log(1/\varepsilon)}{\gamma}\right)$$

steps, from any starting distribution.

So the gap is the rate of mixing: $\gamma = \Omega(1)$ gives $O(\log n + \log(1/\varepsilon))$ steps, while γ polynomially small forces polynomially many steps.

The converse: a sparse cut traps the walk. For completeness we record the direction which shows that the implication cannot be improved. For $S \subseteq V$ define the *t-step escape probability* of the ordinary walk,

$$\phi_t(S) \;:=\; \mathbb{P}[X_t \notin S \mid X_0 \sim \mathrm{Unif}(S)] \;=\; \frac{\langle \mathbf{1}_{\bar{S}},\, P^t \mathbf{1}_S \rangle}{|S|},$$

so that $\phi_1(S) = \phi(S)$.

Lemma 1 (Escape lemma)**.** *For every $S \subseteq V$ and every integer $t \ge 1$, $\phi_t(S) \le t \cdot \phi(S)$.*

Consequently, if $|S| \le n/2$ then for $t \le 1/(4\phi(S))$ the walk is still in S with probability at least $3/4$, while $\pi(S) \le 1/2$, so the distribution at time t is at total variation distance at least $1/4$ from uniform: a set of conductance $\phi(S)$ delays mixing for $\approx 1/\phi(S)$ steps.

Summary of the relations. For a D-regular graph with gap γ:

1. every S with $|S| \le n/2$ satisfies $\phi(S) \ge \gamma/2$;
2. some S with $|S| \le n/2$ satisfies $\phi(S) \le \sqrt{2\gamma}$;
3. the lazy walk mixes in $O((\log n + \log(1/\varepsilon))/\gamma)$ steps;
4. a set of small conductance prevents mixing for $t \ll 1/\phi(S)$ steps.

2.5 Expander graphs

Graphs on which the random walk mixes very well are called expander graphs. They have a constant spectral gap and no bottleneck in which the walk can get stuck.

Definition 5 (Expanders)**.** *A D-regular graph G is an α-**edge expander** if $\phi(S) \ge \alpha$ for every nonempty S with $|S| \le n/2$, equivalently $\phi_G \ge \alpha$, equivalently $|E(S, \bar{S})| \ge \alpha D|S|$. It is an α-**spectral expander** if $\frac{\gamma}{2} = \frac{\lambda_2}{2D} \ge \alpha$.*

By Theorem 1, spectral expansion implies edge expansion, and α-edge expansion implies $\gamma \ge \alpha^2/2$.

A single connected graph is always an α-expander for some $\alpha > 0$, since μ_1 has multiplicity 1 and so $\mu_2 < 1$ and γ is nonzero.

2.6 Random regular graphs are excellent expanders

Expanders are not rare, and in fact a random D-regular graph is essentially optimal. Recall from the table above that $\theta_i = D\mu_i$, so a bound on the θ_i is a bound on the μ_i after dividing by D.

Theorem 2 (Friedman [2])**.** *Fix $D \ge 3$ and let $D = \theta_1 \ge \theta_2 \ge \cdots \ge \theta_n$ be the adjacency eigenvalues of a random D-regular graph on n vertices. Then, with probability tending to 1 as $n \to \infty$,*

$$\max_{i \ge 2} |\theta_i| \;\le\; 2\sqrt{D-1} + o_n(1).$$

Here $o_n(1)$ denotes a quantity that tends to 0 as $n \to \infty$ with D fixed, so the bound is $2\sqrt{D-1}$ up to an error that becomes negligible for large graphs. (Friedman proves the statement in the standard permutation model of a random regular graph, in the form: for every fixed $\varepsilon > 0$ the bound $2\sqrt{D-1}+\varepsilon$ holds with probability tending to 1.) In particular, for n large enough the bound holds with probability at least $3/4$, which is all we need if we want to *sample* an expander. In the normalized language of Lecture 4 this reads $\max_{i \ge 2} |\mu_i| \le \frac{2\sqrt{D-1}}{D} + o_n(1)$, so the corresponding two-sided spectral gap is $1 - \frac{2\sqrt{D-1}}{D}$, a constant depending only on D and approaching 1 as D grows. Friedman's proof is delicate; a much easier trace-moment argument already gives the weaker bound $\max_{i \ge 2} |\theta_i| = O(\sqrt{D \log D})$ with high constant probability, which is enough for most applications.

That $2\sqrt{D-1}$ cannot be improved is the content of the following phenomenon.

Theorem 3 (Alon–Boppana phenomenon; see [4])**.** *If a D-regular graph has diameter m, then*

$$\theta_2 \;\ge\; 2\sqrt{D-1} - O_D\!\left(\tfrac{1}{m}\right).$$

The subscript in $O_D(1/m)$ means that the hidden constant is allowed to depend on the degree D, but on nothing else: the error term is at most c_D/m for some constant c_D, and in particular it vanishes as the diameter grows. So Friedman's bound is essentially optimal. Graphs attaining $\max_{i\geq 2}|\theta_i| \leq 2\sqrt{D-1}$ exactly are called *Ramanujan* [3].

2.7 Explicit expanders

The previous fact about random graphs proves the *existence* of expanders, but it is of no use to an algorithm that is only allowed $O(\log n)$ bits of memory: we can neither store a randomly sampled graph nor afford the randomness needed to draw it. What we need instead is an *explicit* family, that is, a function neighbour(v, k) that computes a single neighbour on demand. Here v is a vertex represented with $O(\log n)$ bits and $k \in [D]$ is an edge label; the function returns the k-th neighbour of v and runs in $O(\log n)$ space.

The adjacency list may have $\Theta(n)$ entries, but the local rule needs only a few machine words. Here is the classical example.

Definition 6 (Margulis–Gabber–Galil graph [5])**.** *The vertex set is $V_M = \mathbb{Z}_M \times \mathbb{Z}_M$, and the eight neighbours of (x, y) are*

$$(x \pm 2y,\, y), \quad (x \pm (2y+1),\, y), \quad (x,\, y \pm 2x), \quad (x,\, y \pm (2x+1)) \pmod{M}.$$

Theorem 4. *There is a universal constant $c > 0$ such that every graph of this family is an 8-regular c-spectral expander on $n = M^2$ vertices.*

Margulis gave the construction and proved that a spectral gap exists, without an explicit bound; Gabber and Galil [5] later supplied an explicit value of c. Only modular addition and multiplication by 2 are needed, and these are computable with $O(\log M) = O(\log n)$ bits, so the family is explicit. We will use such families as ready-made components in Section 4.

2.8 Notation introduced in this lecture

Everything above is Lecture 4 notation. The symbols below are new, and none of them collides with a symbol of Lecture 4:

- N, the number of vertices *after* the regularization of Section 3.3; since $N = O(n+|E|) = O(n^2)$ we have $\log N = \Theta(\log n)$;
- $\rho(G) := \max_{i\geq 2}|\mu_i|$ and $\mathrm{gap}(G) := 1 - \rho(G)$, the *two-sided* spectral parameters of Section 4; for a lazy walk all nontrivial eigenvalues are nonnegative, so $\rho = \mu_2$ and $\mathrm{gap}(G) = \gamma(G)$;
- Rot_G, the rotation map, and $\textcircled{z}$, the zig-zag product.

3 Randomized logarithmic-space connectivity

3.1 The problem and the model

Problem (USTCON). Given an undirected graph $G = (V, E)$ on n vertices and $s, t \in V$, decide whether s and t belong to the same connected component.

The input is stored on a read-only tape and accessed through two oracles, degree(v) and neighbour(v, j) for $1 \le j \le \text{degree}(v)$; the output is one bit; and the algorithm may only write on a workspace of $O(\log n)$ bits. That budget is enough to hold a constant number of vertex names, edge labels and counters, and nothing else. In particular it is *not* enough to store the set of visited vertices, which is what rules out BFS and DFS. Recall that L (resp. RL) is the class of problems decidable by a deterministic (resp. randomized, with one-sided error and polynomial time) log-space algorithm.

3.2 The random-walk algorithm

A random walk is a log-space object: it only remembers where it currently is.

Algorithm WalkConnect(G, s, t, M)

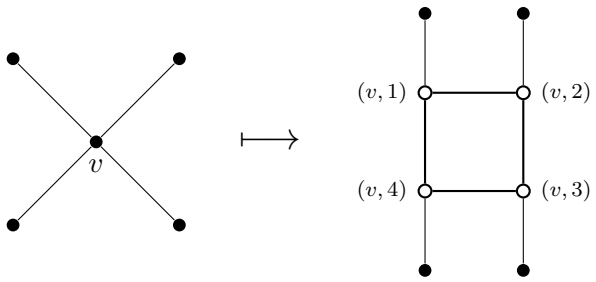
1. $X \leftarrow s$.
2. Repeat M times:
   - sample $j \sim \text{Unif}\{1, \dots, \text{degree}(X)\}$;
   - $X \leftarrow \text{neighbour}(X, j)$;
   - if $X = t$, **return** TRUE.
3. **return** FALSE.

If s and t are disconnected the walk can never reach t, so the algorithm never accepts: the error is one-sided, and the only thing to prove is that M can be chosen polynomially large so that a *reachable* t is hit with probability at least $2/3$. This is where the spectral machinery enters. Two adjustments are needed first. The analysis below runs the walk on the 3-regular graph built in Section 3.3 rather than on G itself, since on a non-regular graph the stationary distribution is proportional to the degrees and not uniform; and the walk is made lazy as in Section 2.4, so that bipartite components do not obstruct convergence.

3.3 Regularization

All our spectral tools assume regularity, so we first replace G by a 3-regular multigraph with the same connectivity structure.

Definition 7 (Port construction)**.** *Replace every vertex v of degree $d(v)$ by a cycle of $d(v)$* ***ports*** *$(v, 1), \dots, (v, d(v))$. Each port is incident to two cycle edges inside the gadget and to one external edge, namely the corresponding original edge of G.*



The resulting multigraph G' is 3-regular and has $\sum_v d(v) = 2|E|$ vertices. Contracting each port cycle back to a single vertex recovers G, so components are preserved:

$$s \leftrightarrow t \text{ in } G \qquad \Longleftrightarrow \qquad \text{port}(s) \leftrightarrow \text{port}(t) \text{ in } G',$$

where $\text{port}(v)$ is any port of v. (Degree-1 and degree-2 gadgets are fixed using parallel edges and self-loops, and isolated vertices are handled separately: $s = t$ is trivial and an isolated $s \neq t$ is trivially rejected.)

The reduction is *local*: given a port (v, j) and an edge label, computing the previous port, the next port, or the port at the other end of the external edge requires only vertex names, port indices and small counters, i.e. $O(\log n)$ bits. We never write G' down; we compute inside it on demand. From now on we work with a D-regular graph on N vertices with $D = O(1)$ and $N = O(n + |E|)$, and we make the walk lazy.

3.4 The spectral gap of a connected component

Lemma 2. *Let G be a connected D-regular graph on N vertices. Then $\phi_G \geq \frac{2}{DN}$ and consequently $\gamma \geq \frac{2}{D^2N^2}$. For constant D, every component satisfies $\gamma = \Omega(N^{-2})$.*

Proof sketch. Connectivity means every nonempty S with $|S| \leq N/2$ has at least one boundary edge, so $\phi(S) \geq \frac{1}{D|S|} \geq \frac{2}{DN}$. Now use the *upper* bound $\phi_G \leq \sqrt{2\gamma}$ of Theorem 1, rearranged as $\gamma \geq \phi_G^2/2$, which gives $\gamma \geq 2/(D^2N^2)$. □

This bound is worst-case and it is attained: a path or a cycle really does have $\gamma = \Theta(N^{-2})$, which corresponds to the familiar fact that a random walk on a cycle needs $\Omega(N^2)$ steps to travel distance N.

3.5 Analysis of the randomized algorithm

Theorem 5. *USTCON* $\in$ RL*: there is a randomized algorithm using $O(\log N) = O(\log n)$ workspace and $O(N^3 \log N)$ time which always rejects disconnected pairs and accepts connected pairs with probability at least $2/3$.*

Proof sketch. Work in the regularized graph and let G be a connected D-regular graph of size N. By Lemma 2, $\gamma(G) = \Omega(N^{-2})$, so by Corollary 1 the lazy walk started at s reaches total variation

distance $\frac{1}{4N}$ from the uniform distribution on G after

$$T \;=\; t_{\mathrm{mix}}\!\left(\tfrac{1}{4N}\right) \;=\; O\!\left(\frac{\log N}{\gamma(G)}\right) \;=\; O(N^2 \log N)$$

steps. Hence for every t in the same component, $\mathbb{P}[X_T = t] \geq \frac{1}{N} - \frac{1}{4N} \geq \frac{1}{2N}$.

Now run $3N$ consecutive blocks of T lazy steps each, and check at the end of every block whether the current vertex is t. Corollary 1 holds from any starting distribution, so each block ends at t with probability at least $\frac{1}{2N}$ whatever happened in the previous blocks, and the conditional probabilities multiply:

$$\mathbb{P}[\,t \text{ is never hit}\,] \;\leq\; \left(1 - \tfrac{1}{2N}\right)^{3N} \;<\; e^{-3/2} \;<\; \tfrac{1}{3}.$$

For the resources: we store the current vertex, a step counter and a block counter, each $O(\log N)$ bits, plus one random edge label, which after regularization is $O(1)$ bits. The total number of steps is $3N \cdot T = O(N^3 \log N)$. □

The randomness used is enormous in total ($O(N^3 \log N)$ random bits), but it is consumed *on the fly*, a few bits at a time and immediately forgotten. This is why the algorithm fits in logarithmic space, and also why derandomizing it is delicate.

Question: Can the randomness be removed without increasing the space?

4 Deterministic log space: squaring and the zig-zag product

4.1 Reingold's theorem and the strategy

Theorem 6 (Reingold [6])**.** *Undirected s–t connectivity can be decided deterministically in $O(\log n)$ workspace:*

$$\mathit{USTCON} \in \mathsf{L}.$$

The strategy has three steps:

1. regularize the input graph (Section 3.3);
2. transform every connected component into an *expander*;
3. deterministically enumerate all short walks.

The difficulty lies entirely in step 2: the transformation must preserve connected components, must increase the spectral gap, and must remain *locally computable* in log space, since we can never afford to write down an intermediate graph.

4.2 On expanders, connectivity is easy

Let us first see why step 2 would suffice.

Proposition 1. *Suppose G has constant degree D and every connected component has $\gamma = \Omega(1)$. Then for every pair s, t in the same component there is a walk from s to t of length $T = O(\log N)$.*

Proof sketch. By Corollary 1 the lazy walk mixes to error below $\frac{1}{2N}$ in $T = O(\log N)$ steps, hence $(\tilde{P}^T)_{ts} > 0$, and a positive transition probability is precisely the existence of a walk. The lazy walk is a walk in the graph with one self-loop added at every vertex, so deleting the loop steps turns it into a walk of length at most T in G. □

Such a walk is described by the string of edge labels it follows: a walk of length T in a D-regular graph is a string $(a_1, \dots, a_T) \in [D]^T$, using $T \log D = O(\log N)$ bits when $D = O(1)$ and $T = O(\log N)$. So we can afford to *enumerate all of them*: there are $D^T = N^{O(1)}$ such strings, and for each one we start at s, follow the prescribed neighbours one label at a time, and accept if we ever reach t. At any moment we store one current vertex and one string, which is $O(\log N)$ bits. Connectivity on a disjoint union of constant-degree expanders is therefore in L.

The real task. An arbitrary graph is nothing like an expander: by Lemma 2 its gap can be as small as $\Theta(N^{-2})$, in which case the guaranteed walks have polynomial length and the string describing one no longer fits in $O(\log N)$ bits. We need an operation $G \mapsto G'$ that

- preserves connected components;
- increases the spectral gap by a constant factor;
- keeps the degree manageable;
- keeps the number of vertices polynomial;
- has a neighbour function computable locally, from the neighbour function of G.

4.3 Squaring: good for the spectrum, terrible for the degree

The natural candidate is powering.

Definition 8 (Square)**.** *G^2 is the multigraph with one edge for each walk of length two in G. If G is D-regular then G^2 is D^2-regular and $P_{G^2} = P_G^2$.*

Since squaring the matrix squares the eigenvalues, a nontrivial eigenvalue $\mu = 1 - \gamma$ becomes μ^2, so

$$1 - \mu_2^2 \;=\; 1 - (1-\gamma)^2 \;=\; 2\gamma - \gamma^2 :$$

for small γ the gap essentially *doubles*. One caveat is that squaring maps a very negative eigenvalue to a very positive one, so the relevant parameter is now two-sided.

Definition 9 (Two-sided spectral parameter)**.**

$$\rho(G) := \max_{i \geq 2} |\mu_i(P_G)|, \qquad \mathrm{gap}(G) := 1 - \rho(G).$$

This is the quantity that already governed convergence of the ordinary walk in Lecture 4, and it is the ratio ρ of the power-iteration analysis of Lecture 3, since $\mu_1 = 1$. For a lazy walk all nontrivial eigenvalues are nonnegative, so $\rho = \mu_2$ and $\mathrm{gap}(G) = \gamma(G)$: the two notions of gap agree.

Because $P_{G^2} = P_G^2$, the spectral decomposition $P^t = \sum_r \mu_r^t u_r u_r^T$ of Lecture 4 gives $\rho(G^2) = \rho(G)^2$ and therefore $\mathrm{gap}(G^2) = 1 - (1 - \mathrm{gap}(G))^2$, while laziness (or self-loops) guarantees $\rho < 1$ on every component. So the gap practically doubles. Iterating k times gives $P \mapsto P^{2^k}$ and $\rho \mapsto \rho^{2^k}$, which is the spectral progress we want, but the degree becomes D^{2^k}. After $O(\log N)$ squarings, which is what Lemma 2 forces, an edge label no longer fits in logarithmic space. **Squaring must be paired with a degree-reduction operation.**

4.4 Rotation maps

To compose local descriptions of graphs we need a slightly richer interface than neighbour(v, i): we also want to know *under which label we arrive.*

Definition 10 (Rotation map)**.** *Let G be D-regular and label the D edges incident to each vertex by $[D]$. If the i-th edge at v leads to u and is the j-th edge at u, define*

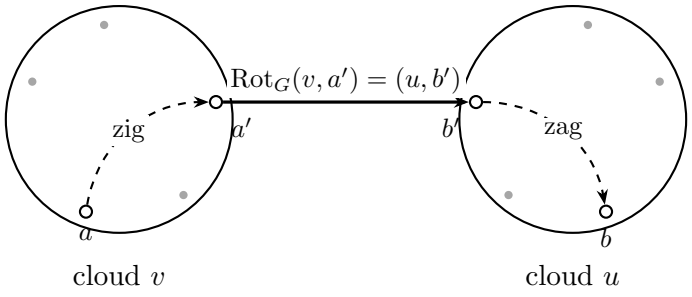
$$\mathrm{Rot}_G(v, i) := (u, j).$$

Note that Rot_G is an involution on $V \times [D]$, i.e. a map that is its own inverse: applying it twice returns the starting pair, $\mathrm{Rot}_G(\mathrm{Rot}_G(v, i)) = (v, i)$. This is just the statement that traversing an edge and then traversing the same edge backwards brings us back where we started. The reverse label j is what lets the next operation decide where to go, and it is why all the constructions below are described by rotation maps rather than by adjacency lists.

4.5 The zig-zag product

Definition 11 (Zig-zag product [7])**.** *Let G be D_1-regular on N vertices and let H be D_2-regular on the* label set $[D_1]$*. The zig-zag product $G \textcircled{z} H$ has vertex set $V(G) \times [D_1]$: a vertex (v, a) records a* cloud *v together with a port label a inside it. One step consists of a small step, a big step and a small step. Starting from (v, a):*

1. ***zig**: pick an H-neighbour a' of a;*
2. ***jump**: apply $\mathrm{Rot}_G(v, a') = (u, b')$;*
3. ***zag**: pick an H-neighbour b of b';*
4. ***output** (u, b).*



What matters here is the degree count: the two free choices are the zig and the zag, so

$$\deg(G \circledz H) = D_2^2,$$

which depends only on the *small* graph H and not at all on the large degree D_1. The number of vertices is multiplied by D_1, which we will keep polynomial. Connectivity is preserved in the right sense: the only way to move between clouds is through an edge of G, so two vertices of $G \circledz H$ lying in clouds of different G-components remain disconnected, while if H is connected the whole cloud of v lies in the component of v.

Theorem 7 (Zig-zag spectral theorem [7])**.**

$$\mathrm{gap}(G \circledz H) \;\geq\; \mathrm{gap}(G) \cdot \mathrm{gap}(H)^2.$$

We use this without proof; the proof decomposes the walk operator of $G \circledz H$ into a "cloud-constant" part, on which the product behaves like a step of G, and a "cloud-orthogonal" part, which is contracted by the two H-steps. In other words, if H is an excellent expander, say $\mathrm{gap}(H) \geq 1-\epsilon$, then

$$\mathrm{gap}(G \circledz H) \;\geq\; (1-\epsilon)^2\,\mathrm{gap}(G),$$

so the zig-zag product costs us only a factor $(1-\epsilon)^2$ of the gap, while collapsing the degree from D_1 to the constant D_2^2. Since H is a *constant-size* graph, we may take it from an explicit family (Section 2.7) or even find it by brute force once and for all.

4.6 Combining squaring and the zig-zag product

Now we alternate the two operations: squaring pushes the gap up, the zig-zag product pushes the degree back down.

Fix a constant D and a D-regular expander H on D^4 vertices with $\mathrm{gap}(H) \geq 1 - \epsilon$. The graph produced in Section 3.3 can be turned into a D^2-regular graph G_0 by adding self-loops, which also guarantees $\rho(G_0) < 1$ on every component; suppose then that G_s is D^2-regular with $\mathrm{gap}(G_s)$. One round is:

1. **square:** $G'_s := G_s^2$, which is D^4-regular with $\mathrm{gap}(G'_s) = 2\mathrm{gap}(G_s) - \mathrm{gap}(G_s)^2$;
2. **reduce the degree:** $G_{s+1} := G'_s \circledz H$, which is again D^2-regular, on $|V(G_{s+1})| = D^4\,|V(G_s)|$ vertices.

The two operations fit together because H lives on the label set of G'_s, which has size D^4.

While $\mathrm{gap}(G_s) \le \frac{1}{10}$ we have $2\mathrm{gap}(G_s) - \mathrm{gap}(G_s)^2 \ge 1.9\,\mathrm{gap}(G_s)$, so choosing ϵ small enough that $(1-\epsilon)^2 \ge \frac{1.8}{1.9}$, Theorem 7 gives

$$\mathrm{gap}(G_{s+1}) \;\ge\; (1-\epsilon)^2\left(2\mathrm{gap}(G_s) - \mathrm{gap}(G_s)^2\right) \;\ge\; 1.8\,\mathrm{gap}(G_s). \tag{3}$$

Corollary 2. *Starting from* $\mathrm{gap}(G_0) \ge N^{-O(1)}$*, after* $O(\log N)$ *rounds the gap is* $\Omega(1)$*. Along the way the degree stays constant* (D^2) *and the number of vertices is multiplied by* D^4 *per round, hence remains* $N^{O(1)}$*.*

This is the transformation demanded in Section 4.2: it preserves components, it drives the gap up geometrically by (3), and both the degree and the vertex count stay under control. Each round replaces a vertex by a cloud of D^4 copies of it, and by the remark above the whole cloud stays in the component of its centre, so it is enough to test connectivity between one designated copy of s and one of t in the final graph; both are computable on demand from the rotation maps. Combined with Proposition 1, every component of the final graph is an expander in which short walks suffice, and these can be enumerated deterministically.

5 Summary

1. For a D-regular graph, A, $P = \frac{1}{D}A$, $L_G = DI - A$ and $\mathcal{L}_G = I - P$ share eigenvectors; the trivial eigenpair carries no information and the first nontrivial one defines the gap $\gamma = \lambda_2/D = 1 - \mu_2$.

2. Cheeger's inequality, $\frac{\gamma}{2} \le \phi_G \le \sqrt{2\gamma}$, identifies a small gap with a sparse cut, up to a quadratic loss.

3. The lazy walk mixes to error ε in $O((\log n + \log(1/\varepsilon))/\gamma)$ steps; conversely a set S of conductance $\phi(S)$ traps the walk for $\approx 1/\phi(S)$ steps (Lemma 1).

4. Random D-regular graphs are essentially optimal expanders (Friedman), and $2\sqrt{D-1}$ is the right threshold (Alon–Boppana); for algorithms we nevertheless need *explicit* families.

5. A random walk decides USTCON in randomized log space: after regularization every component has $\gamma = \Omega(N^{-2})$, so $O(N^3 \log N)$ steps suffice.

6. On a disjoint union of constant-degree expanders, connectivity is in *deterministic* log space: enumerate all walks of length $O(\log N)$.

7. Squaring doubles the gap but squares the degree; the zig-zag product restores a constant degree at the cost of a factor $\mathrm{gap}(H)^2$. Alternating them multiplies the gap by 1.8 per round, so $O(\log N)$ rounds turn any component into an expander.

AI use disclosure. Claude Opus 5 was used to generate the TikZ figures and to check grammar and internal consistency.

References

[1] Salil P. Vadhan. The unified theory of pseudorandomness: guest column. *SIGACT News*, 38(3):39–54, September 2007. doi:10.1145/1324215.1324225.

[2] Joel Friedman. A Proof of Alon's Second Eigenvalue Conjecture and Related Problems. *Memoirs of the American Mathematical Society*, 195(910), viii+100 pages, 2008. doi:10.1090/memo/0910.

[3] Alexander Lubotzky, Ralph S. Phillips, Peter Sarnak. Ramanujan graphs. *Combinatorica*, 8(3):261–277, September 1988. doi:10.1007/BF02126799.

[4] Alon Nilli. On the second eigenvalue of a graph. *Discrete Mathematics*, 91(2):207–210, 1991. doi:10.1016/0012-365X(91)90112-F.

[5] Ofer Gabber, Zvi Galil. Explicit constructions of linear-sized superconcentrators. *Journal of Computer and System Sciences*, 22(3):407–420, 1981. doi:10.1016/0022-0000(81)90040-4.

[6] Omer Reingold. Undirected Connectivity in Log-Space. *Journal of the ACM*, 55(4), Article 17, 24 pages, September 2008. doi:10.1145/1391289.1391291.

[7] Omer Reingold, Salil Vadhan, Avi Wigderson. Entropy waves, the zig-zag graph product, and new constant-degree expanders. *Annals of Mathematics*, 155(1):157–187, 2002. doi:10.2307/3062153.