Undirected s-t Connectivity in Logarithmic Space

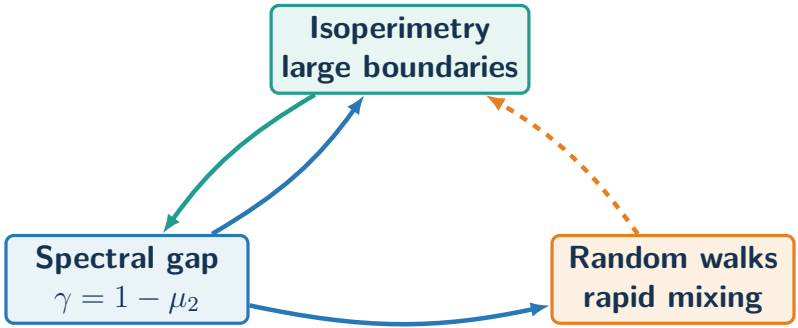
Expansion, random walks, zig-zag products, and derandomized squaring

Can connectivity be decided while storing only a few vertices?

Last lecture recap

Expansion, spectral gap, and mixing

Three phenomena describe the same bottleneck



A sparse cut simultaneously creates

a small boundary, a slowly varying eigenvector, a slowly mixing random walk.

Normalization from the previous lecture

Let G be an undirected D-regular graph with adjacency matrix A. Define

$$P := \frac{1}{D}A, \qquad \mathcal{L} := I - P.$$

Write the eigenvalues of P as

$$1 = \mu_1 \geq \mu_2 \geq \cdots \geq \mu_n \geq -1.$$

The first nontrivial normalized Laplacian eigenvalue is

$$\gamma := \lambda_2(\mathcal{L}) = 1 - \mu_2.$$

For a nonempty set S with $|S| \leq n/2$, its conductance is

$$\phi(S) := \frac{|E(S, \overline{S})|}{D|S|}.$$

Cheeger's inequality equates expansion and spectral gap

Define

$$\phi_G := \min_{\substack{\varnothing \neq S \subseteq V \\ |S| \leq n/2}} \phi(S).$$

Cheeger's inequality says

$$\boxed{\frac{\gamma}{2} \leq \phi_G \leq \sqrt{2\gamma}.}$$

Equivalently,

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

Thus a family has constant conductance exactly when it has a constant normalized spectral gap, up to the quadratic loss in the bounds.

Spectral gap implies mixing by the power-method argument

For the lazy walk

$$\widetilde{P} := \frac{1}{2}(I + P),$$

the nontrivial eigenvalues lie in $[0, 1)$, and

$$\widetilde{\mu}_2 = 1 - \frac{\gamma}{2}.$$

If $\pi = \mathbf{1}/n$, then

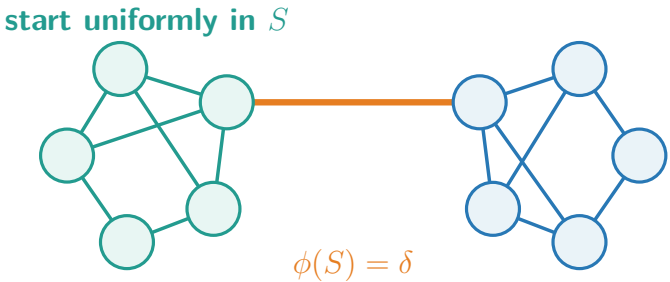
$$\left\| \widetilde{P}^t p_0 - \pi \right\|_2 \leq \left(1 - \frac{\gamma}{2}\right)^t \|p_0 - \pi\|_2 \leq e^{-\gamma t/2} \|p_0 - \pi\|_2 .$$

Consequently, mixing to total-variation error ε takes

$$O\left(\frac{\log n + \log(1/\varepsilon)}{\gamma}\right)$$

steps.

Why should rapid mixing force expansion?



If only a δ-fraction of the edges incident to S leave it, then a walk should remain trapped for roughly $1/\delta$ steps.

The t-step escape probability

For $A \subseteq V$, define

$$\phi_t(A) := \left\langle \mathbf{1}_{\overline{A}}, P^t \frac{\mathbf{1}_A}{|A|} \right\rangle.$$

Equivalently,

$$\phi_t(A) = \Pr[X_t \notin A \mid X_0 \text{ is uniform on } A].$$

Notice that

$$\phi_1(A) = \phi(A).$$

Escape lemma

For every integer $t \geq 1$,

$$\boxed{\phi_t(A) \leq t\,\phi(A).}$$

Begin with a stationary random walk

Let

$$X_0 \sim \mathrm{Uniform}(V),$$

and let X_{i+1} be a uniformly random neighbor of X_i.
Since G is regular, the uniform distribution is stationary. Therefore

$$X_i \sim \mathrm{Uniform}(V) \qquad \text{for every } i.$$

Conditioning on $X_0 \in A$ makes X_0 uniform on A. Hence

$$\begin{aligned}\Pr[X_0 \in A,\ X_t \notin A] &= \Pr[X_0 \in A]\ \Pr[X_t \notin A \mid X_0 \in A] \\ &= \frac{|A|}{n}\phi_t(A).\end{aligned}$$

Ending outside requires crossing the boundary

If

$$X_0 \in A \qquad \text{and} \qquad X_t \notin A,$$

then the walk must leave A at some time before t.
Therefore

$$\{X_0 \in A,\ X_t \notin A\} \subseteq \bigcup_{i=0}^{t-1} \{X_i \in A,\ X_{i+1} \notin A\}.$$

Taking probabilities and applying a union bound,

$$\Pr[X_0 \in A,\ X_t \notin A] \leq \sum_{i=0}^{t-1} \Pr[X_i \in A,\ X_{i+1} \notin A].$$

The walk may leave and re-enter A; the union bound does not require these events to be disjoint.

Stationarity makes every crossing probability identical

At time i, every directed edge is traversed with probability

$$\frac{1}{n} \cdot \frac{1}{D} = \frac{1}{nD}.$$

There are $|E(A, \overline{A})|$ directed edges from A to $\overline{A}$. Thus

$$\begin{aligned} \Pr[X_i \in A,\ X_{i+1} \notin A] &= \frac{|E(A, \overline{A})|}{nD} \\ &= \frac{|A|}{n} \frac{|E(A, \overline{A})|}{D|A|} \\ &= \frac{|A|}{n} \phi(A). \end{aligned}$$

This value is independent of i.

The escape lemma follows after cancellation

Combining the previous estimates,

$$\begin{aligned}\frac{|A|}{n}\phi_t(A) &= \Pr[X_0 \in A,\ X_t \notin A] \\ &\leq \sum_{i=0}^{t-1} \Pr[X_i \in A,\ X_{i+1} \notin A] \\ &= t\frac{|A|}{n}\phi(A).\end{aligned}$$

Since $A \neq \varnothing$, cancel $|A|/n$:

$$\boxed{\phi_t(A) \leq t\phi(A).}$$

This proof used only stationarity and the fact that every path from A to $\overline{A}$ must cross the boundary.

Small conductance keeps most probability inside

Let

$$q_t := P^t \frac{\mathbf{1}_A}{|A|}.$$

The escape lemma gives

$$q_t(\overline{A}) = \phi_t(A) \leq t\phi(A).$$

Therefore

$$q_t(A) = 1 - q_t(\overline{A}) \geq 1 - t\phi(A).$$

If

$$t \leq \frac{1}{4\phi(A)},$$

then

$$q_t(A) \geq \frac{3}{4}.$$

A walk starting uniformly in A still places at least three quarters of its probability

The walk is still far from uniform

Suppose $|A| \leq n/2$. The uniform distribution π satisfies

$$\pi(A) = \frac{|A|}{n} \leq \frac{1}{2}.$$

For $t \leq 1/(4\phi(A))$,

$$q_t(A) \geq \frac{3}{4}.$$

Total variation distance admits the characterization

$$\|q_t - \pi\|_{\mathrm{TV}} = \max_{B \subseteq V} |q_t(B) - \pi(B)|.$$

Taking $B = A$,

$$\|q_t - \pi\|_{\mathrm{TV}} \geq q_t(A) - \pi(A) \geq \frac{3}{4} - \frac{1}{2} = \boxed{\frac{1}{4}}.$$

The recap closes the three-way circle

For a D-regular graph with normalized spectral gap γ:

1. Every S with $|S| \le n/2$ satisfies
$$\phi(S) \ge \frac{\gamma}{2}.$$
2. Some S with $|S| \le n/2$ satisfies
$$\phi(S) \le \sqrt{2\gamma}.$$
3. The lazy walk mixes in
$$O\left(\frac{\log n + \log(1/\varepsilon)}{\gamma}\right)$$
steps.
4. A set S with small conductance prevents mixing for
$$t \ll \frac{1}{\phi(S)}$$
steps.

Part I

Expander graphs

Expanders have no sparse cuts

A D-regular graph is an α-edge expander if

$$\phi(S) \geq \alpha$$

for every nonempty $S \subseteq V$ with $|S| \leq n/2$.
Equivalently,

$$|E(S, \overline{S})| \geq \alpha D|S|.$$

Under the convention of the previous lecture, G is an α-spectral expander if

$$\frac{\gamma}{2} \geq \alpha.$$

Cheeger's inequality implies:

- spectral expansion implies edge expansion;
- α-edge expansion implies spectral gap

$$\gamma \geq \frac{\alpha^2}{2}.$$

Random regular graphs are excellent expanders

Let

$$D = \theta_1 \geq \theta_2 \geq \cdots \geq \theta_n$$

be the adjacency eigenvalues of a random D-regular graph.

Friedman's theorem

For fixed $D \geq 3$, with probability tending to 1,

$$\max_{i \geq 2} |\theta_i| \leq 2\sqrt{D-1} + o_n(1).$$

In particular, for all sufficiently large n, the conclusion holds with probability at least $3/4$.

The corresponding two-sided random-walk spectral gap is approximately

$$1 - \frac{2\sqrt{D-1}}{D}.$$

Random regular graphs are close to spectrally optimal

A substantially easier argument gives the weaker estimate

$$\max_{i \geq 2} |\theta_i| = O\left(\sqrt{D \log D}\right)$$

with high constant probability.

Friedman's bound is close to the best possible:

Alon–Boppana phenomenon

If a D-regular graph has diameter m, then

$$\theta_2 \geq 2\sqrt{D-1} - O_D(1/m).$$

Thus large D-regular graphs cannot have every nontrivial adjacency eigenvalue substantially smaller than $2\sqrt{D-1}$.

Explicit expanders avoid storing the entire graph

Random graphs prove existence, but algorithms often need a deterministic and locally computable family.

An explicit expander family provides a function

$$\mathrm{neighbor}(v, k)$$

such that:

- v is represented with $O(\log n)$ bits;
- $k \in [D]$ is an edge label;
- the k-th neighbor of v is computable in $O(\log n)$ space;
- the graph has constant degree and constant spectral gap.

The full adjacency list may have $\Theta(n)$ entries, but the local rule requires only a few machine words.

A discrete-torus expander

One standard Margulis–Gabber–Galil construction has vertex set

$$V_M = \mathbb{Z}_M \times \mathbb{Z}_M.$$

From (x, y), its eight neighbors are

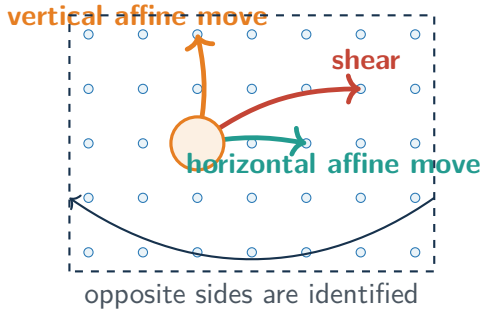
$$(x \pm 2y, y), \quad (x \pm (2y+1), y),$$
$$(x, y \pm 2x), \quad (x, y \pm (2x+1)),$$

with all arithmetic modulo M.

Expansion theorem

There is a universal constant $c > 0$ such that every graph in this family is an 8-regular c-spectral expander on M^2 vertices.

The torus rule combines local shifts and shears



Modular addition and multiplication by 2 are computable using $O(\log M) = O(\log n)$ bits.

Part II

Randomized logarithmic-space connectivity

The input graph is accessed locally

In undirected s-t connectivity, the input consists of

$$G = (V, E), \qquad s, t \in V,$$

and we must decide whether s and t lie in the same component.
We may access G through two read-only oracles:

$$\text{degree}(v),$$

$$\text{neighbor}(v, j), \qquad 1 \leq j \leq \text{degree}(v).$$

The workspace is restricted to

$$O(\log n)$$

bits: enough to store only a constant number of vertex indices, counters, and edge labels.

A random walk gives the natural one-sided algorithm

Random-walk connectivity test

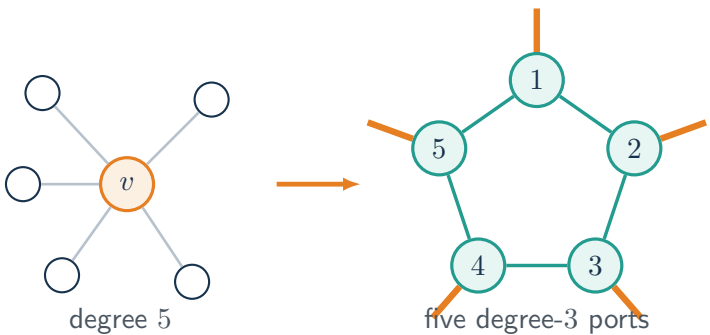
1. Set $X \leftarrow s$.
2. Repeat for M steps:
   2.1 choose
   $$j \sim \mathrm{Uniform}\{1, \ldots, \mathrm{degree}(X)\};$$
   2.2 set
   $$X \leftarrow \mathrm{neighbor}(X, j);$$
   2.3 if $X = t$, return True.
3. Return False.

If s and t are disconnected, the algorithm never accepts. The only possible error is failing to visit t when it is reachable.

Regularization lets us reuse the spectral theory

Replace every original vertex v of degree $d(v)$ by a cycle of $d(v)$ ports:

$$(v, 1), (v, 2), \ldots, (v, d(v)).$$



The port-cycle reduction preserves connectivity

Associate the j-th incident edge of v with port (v, j).
Each port has:

- two cycle edges inside the gadget;
- one external edge representing an original edge.

Hence the resulting multigraph is 3-regular and has

$$\sum_{v} d(v) = 2|E|$$

vertices.
Contracting every port cycle recovers the original graph, so

$$s \leftrightarrow t \text{ in } G \quad \Longleftrightarrow \quad \pi(s) \leftrightarrow \pi(t) \text{ in } G'.$$

Parallel edges or self-loops handle degree-1 and degree-2 gadgets. Isolated vertices are handled separately.

The reduction is computable in logarithmic space

A transformed vertex is represented by

$$(v, j), \qquad 1 \le j \le d(v).$$

Given (v, j) and one of its three edge labels, we can compute:

- the preceding port $(v, j-1)$;
- the following port $(v, j+1)$;
- the port at the opposite endpoint of the j-th original edge.

Only vertex names, port indices, and small counters are required. Each uses $O(\log n)$ bits.

This is why reductions that are routine in polynomial time require extra care in a log-space proof.

Every connected regular graph has some spectral gap

Let G be a connected D-regular graph on N vertices.
Every nonempty S with $|S| \leq N/2$ has at least one boundary edge:

$$|E(S, \overline{S})| \geq 1.$$

Therefore

$$\phi_G \geq \frac{1}{D|S|} \geq \frac{2}{DN}.$$

Cheeger's upper bound

$$\phi_G \leq \sqrt{2\gamma}$$

implies

$$\gamma \geq \frac{\phi_G^2}{2} \geq \boxed{\frac{2}{D^2N^2}}.$$

For constant D, every connected component has $\gamma = \Omega(N^{-2})$.

An arbitrary connected component mixes in polynomial time

For a lazy walk on a constant-degree component,

$$\gamma = \Omega(N^{-2}).$$

To mix to total-variation error $1/(4N)$, it is enough to take

$$T = O\left(\frac{\log N}{\gamma}\right) = O(N^2 \log N)$$

steps.

From any starting vertex, after T steps every target vertex in the same component has probability at least

$$\frac{1}{N} - \frac{1}{4N} \geq \frac{1}{2N}.$$

Thus every block of T steps has an inverse-polynomial chance of ending at t.

Repeating mixing blocks finds the target

Starting from the current position, run $T = O(N^2 \log N)$ lazy steps and check whether the endpoint is t.
If s and t are connected, then conditionally on the past, every block succeeds with probability at least

$$\frac{1}{2N}.$$

After $3N$ blocks,

$$\Pr[\text{never hit } t] \leq \left(1 - \frac{1}{2N}\right)^{3N} < e^{-3/2} < \frac{1}{3}.$$

Hence the algorithm accepts connected pairs with probability greater than $2/3$, and never accepts disconnected pairs.

The randomized algorithm uses only logarithmic space

At any moment the algorithm stores:

- the current vertex: $O(\log N)$ bits;
- a step counter: $O(\log N)$ bits;
- a block counter: $O(\log N)$ bits;
- one random edge label: $O(1)$ bits after regularization.

Therefore the workspace is

$$O(\log N).$$

The running time from this spectral analysis is

$$O(N^3 \log N).$$

The central remaining question is whether randomness can be removed without increasing the space.

Part III

From randomized walks to deterministic log space

Reingold's theorem removes randomness

Theorem

Undirected s-t connectivity can be decided deterministically using

$$O(\log n)$$

workspace.

In complexity-class notation,

$$\boxed{\text{USTCON} \in \text{L}.}$$

The high-level strategy is:

1. regularize the input graph;
2. transform every connected component into an expander;
3. deterministically enumerate short walks.

The transformation must preserve components and remain locally computable in logarithmic space.

Connectivity is easy when every component is an expander

Suppose G has constant degree and every connected component has spectral gap

$$\gamma = \Omega(1).$$

Then the lazy walk mixes in

$$T = O(\log N)$$

steps.
Taking the error below $1/(2N)$ ensures that, for any connected s, t,

$$(\widetilde{P}^T)_{ts} > 0.$$

Hence there exists a walk from s to t of length at most $O(\log N)$.

Short walks can be enumerated in logarithmic space

In a D-regular graph, a length-T walk is specified by

$$(a_1, \ldots, a_T) \in [D]^T.$$

If $D = O(1)$ and $T = O(\log N)$, the entire instruction string uses

$$T \log D = O(\log N)$$

bits.

Deterministically enumerate all $D^T = N^{O(1)}$ strings:

1. begin at s;
2. follow the prescribed neighbor at each step;
3. accept if t is reached.

Thus connectivity in a union of constant-degree expanders is deterministic log-space computable.

The real task is to remove bottlenecks

An arbitrary connected graph may have gap

$$\gamma = \Theta(N^{-2}),$$

so paths guaranteed by mixing may have polynomial length.
We need a graph transformation that:

- preserves connected components;
- increases the spectral gap;
- keeps the degree manageable;
- keeps the number of vertices polynomial;
- has a locally computable neighbor function.

Ordinary graph powering improves the gap, but its degree grows too fast.

Part IV

Squaring and the zig-zag product

Squaring rapidly improves the spectrum

Let G^2 contain one edge for every length-two walk in G. If G is D-regular, then G^2 is D^2-regular.
Its random-walk matrix is

$$P_{G^2} = P_G^2.$$

If a nontrivial eigenvalue of P_G is μ, then the corresponding eigenvalue of P_G^2 is

$$\mu^2.$$

If $\mu_2 = 1 - \gamma$, then

$$1 - \mu_2^2 = 1 - (1 - \gamma)^2 = 2\gamma - \gamma^2.$$

For small γ, the gap nearly doubles.

For powering, use the two-sided spectral parameter

Negative eigenvalues can become positive after squaring. Define

$$\rho(G) := \max_{i \geq 2} |\mu_i(P_G)|$$

and

$$\operatorname{gap}(G) := 1 - \rho(G).$$

Then

$$\rho(G^2) = \rho(G)^2,$$

so

$$\operatorname{gap}(G^2) = 1 - (1 - \operatorname{gap}(G))^2.$$

We may add laziness or self-loops initially to ensure that each connected component has $\rho < 1$.

From this point on, "spectral gap" refers to the two-sided gap unless stated otherwise.

Repeated squaring makes the degree unusable

After k squarings,

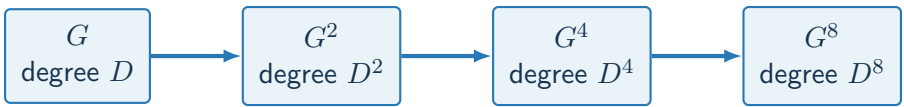
$$P \longmapsto P^{2^k}.$$

The nontrivial spectral radius becomes

$$\rho^{2^k},$$

but the degree becomes

$$D^{2^k}.$$



The spectrum improves, but even one neighbor label soon requires far too many bits.

Edge labels define a rotation map

Let G be D_1-regular. Label the edges incident to every vertex by $[D_1]$.
If the i-th edge at v leads to u, and that edge has label j at u, define

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



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

The reverse port label is essential for composing locally specified graph operations.

The zig-zag product replaces one costly step by three

Let

$$G : D_1\text{-regular on } N \text{ vertices},$$

$$H : D_2\text{-regular on the port set } [D_1].$$

The zig-zag product $G \circledast H$ has vertex set

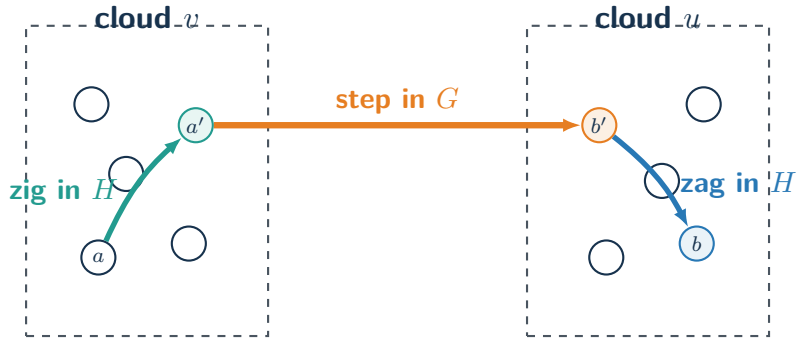
$$V(G) \times [D_1].$$

A vertex (v, a) records:

- a large-graph vertex v, called a *cloud*;
- a port label a inside that cloud.

One product step consists of a small step, a large step, and another small step.

A zig-zag step moves within, across, and within



The zig-zag step in port notation

Starting from (v, a):

1. **Zig:** choose an H-neighbor a' of a.
2. **Jump:** apply
$$\mathrm{Rot}_G(v, a') = (u, b').$$
3. **Zag:** choose an H-neighbor b of b'.
4. Output (u, b).

There are D_2 choices for the zig and D_2 choices for the zag. Hence

$$\deg(G \circledast H) = D_2^2.$$

Crucially, the output degree depends on the small graph H, not on the much larger degree D_1 of G.

The zig-zag product preserves most of the gap

Zig-zag spectral theorem

A convenient form of the zig-zag theorem gives

$$\operatorname{gap}(G \circledast H) \geq \operatorname{gap}(G)\, \operatorname{gap}(H)^2.$$

Thus, if H is an excellent expander with

$$\operatorname{gap}(H) \geq 1 - \varepsilon,$$

then

$$\operatorname{gap}(G \circledast H) \geq (1 - \varepsilon)^2 \operatorname{gap}(G).$$

The zig and zag randomize the port label inside each cloud; the middle step then captures almost all of the large graph's expansion.

We use this theorem without proving its full operator decomposition.

Squaring and zig-zag have complementary effects

Fix a D-regular expander H on D^4 vertices.
Suppose G_s is D^2-regular.

1. Square:

$$G'_s := G_s^2.$$

   Then G'_s is D^4-regular and

$$\mathrm{gap}(G'_s) = 2\gamma_s - \gamma_s^2.$$

2. Reduce the degree:

$$G_{s+1} := G'_s \circledast H.$$

   Then G_{s+1} is D^2-regular and has

$$|V(G_{s+1})| = D^4|V(G_s)|.$$

Each round multiplies a small gap by a constant

If $\gamma_s \leq 1/10$, then

$$2\gamma_s - \gamma_s^2 \geq 1.9\gamma_s.$$

If H is chosen so that $(1-\varepsilon)^2$ is sufficiently close to 1, then

$$\gamma_{s+1} \geq (1-\varepsilon)^2(2\gamma_s - \gamma_s^2) > 1.8\gamma_s.$$

Starting from

$$\gamma_0 \geq N^{-O(1)},$$

after

$$O(\log N)$$

rounds, the gap becomes $\Omega(1)$.

The degree remains constant while the number of vertices grows only polynomially.

Zig-zag gives a high-level route to deterministic connectivity

After $O(\log N)$ rounds:

- every original connected component becomes an expander;
- distinct components remain disconnected;
- the degree remains $D^2 = O(1)$;
- the transformed graph has $N^{O(1)}$ vertices;
- transformed neighbors are computable from rotation maps.

We can then enumerate all walks of length $O(\log N)$ in the final graph using $O(\log N)$ bits.

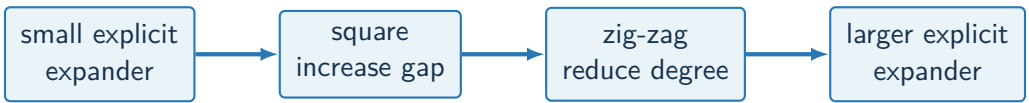
Implementing the recursively composed rotation maps in strict logarithmic space requires careful bookkeeping, but no graph is ever written down explicitly.

A bonus: zig-zag constructs expanders combinatorially

The same iteration yields explicit infinite families of

constant-degree, constant-gap

expanders.



Unlike many earlier constructions, this route is primarily combinatorial: local graph operations replace heavy algebraic descriptions.

Part V

Derandomized squaring

Derandomized squaring imitates two random steps

An ordinary two-step walk uses two independent random edge labels:

$$a \sim [D], \qquad b \sim [D].$$

This requires D^2 possible choices and produces degree D^2.
Derandomized squaring replaces the second independent label by a pseudorandom label obtained from an expander on $[D]$:

$$a \longrightarrow \text{incoming label } a' \xrightarrow{H} b.$$

If H mixes labels well, then b behaves almost as though it were independently uniform.

Ordinary squaring inserts a clique around each midpoint

Every length-two walk

$$x \longrightarrow v \longrightarrow y$$

creates an edge between x and y in G^2.
For a fixed midpoint v, all pairs of incident ports can be combined. Thus, as a multigraph,

$$G^2 = \sum_{v \in V} K_{N(v)},$$

where $K_{N(v)}$ denotes the complete averaging graph on the D incident ports at v.
The complete graph supplies perfect randomness for the second edge label—but has degree D inside every neighborhood.

The derandomized-square walk

Let G be Δ-regular, with rotation map

$$\mathrm{Rot}_G : V \times [\Delta] \to V \times [\Delta].$$

Let H be a d-regular expander on the label set $[\Delta]$.

Starting from v:

1. choose $a \in [\Delta]$ and compute
$$\mathrm{Rot}_G(v, a) = (u, a');$$
2. choose an H-neighbor b of a';
3. compute
$$\mathrm{Rot}_G(u, b) = (w, b');$$
4. output w.

Denote the resulting graph by

$$\mathrm{DSq}(G, H).$$

Derandomized squaring controls the degree

There are

$$\Delta$$

choices for the first G-edge and

$$d$$

choices for the label-expander edge.
Therefore

$$\boxed{\deg(\mathrm{DSq}(G, H)) = \Delta d.}$$

Compare this with ordinary squaring:

$$\deg(G^2) = \Delta^2.$$

When $d \ll \Delta$, the derandomized square is dramatically sparser.

Semidefinite order compares all quadratic forms

For symmetric matrices M, N, write

$$M \preceq N$$

if

$$N - M \succeq 0.$$

Equivalently,

$$x^{\mathsf{T}} M x \leq x^{\mathsf{T}} N x \qquad \text{for every } x.$$

This order is:

- reflexive: $M \preceq M$;
- transitive;
- additive:

$$M_1 \preceq N_1,\ M_2 \preceq N_2 \Longrightarrow M_1 + M_2 \preceq N_1 + N_2.$$

Spectral approximation preserves every energy

For PSD matrices M, N, write

$$M \approx_\varepsilon N$$

if

$$\boxed{\frac{1}{1+\varepsilon} M \preceq N \preceq (1+\varepsilon) M.}$$

Equivalently, for every vector x,

$$\frac{1}{1+\varepsilon} x^{\mathsf{T}} M x \leq x^{\mathsf{T}} N x \leq (1+\varepsilon) x^{\mathsf{T}} M x.$$

Spectral approximation is much stronger than merely approximating individual matrix entries.

A good expander spectrally approximates averaging

Let H be a regular graph on m vertices, and define

$$J_m := \frac{1}{m}\mathbf{1}\mathbf{1}^\mathsf{T}, \qquad \Pi_m := I - J_m.$$

The matrix Π_m is the normalized Laplacian of the complete averaging walk:

$$\Pi_m \mathbf{1} = 0, \qquad \Pi_m x = x \quad \text{for } x \perp \mathbf{1}.$$

Suppose all nontrivial eigenvalues of P_H lie in

$$[-\varepsilon, \varepsilon].$$

Then

$$\boxed{(1-\varepsilon)\Pi_m \preceq I - P_H \preceq (1+\varepsilon)\Pi_m.}$$

Proof: compare the two operators on $1^{\perp}$

Decompose any vector as

$$x = a\mathbf{1} + x_{\perp}, \qquad x_{\perp} \perp \mathbf{1}.$$

Both Π_m and $I - P_H$ vanish on $\mathbf{1}$. Hence

$$x^{\mathsf{T}}\Pi_m x = \|x_{\perp}\|_2^2 .$$

On $\mathbf{1}^{\perp}$, the eigenvalues of P_H lie in $[-\varepsilon, \varepsilon]$, so those of $I - P_H$ lie in

$$[1-\varepsilon, 1+\varepsilon].$$

Therefore

$$(1-\varepsilon)\,\|x_{\perp}\|_2^2 \le x^{\mathsf{T}}(I - P_H)x \le (1+\varepsilon)\,\|x_{\perp}\|_2^2 .$$

Local spectral approximations add together

If

$$M_v \approx_\varepsilon N_v \qquad \text{for every } v,$$

then additivity of semidefinite order gives

$$\sum_v M_v \approx_\varepsilon \sum_v N_v.$$

Ordinary squaring inserts a complete averaging graph on every intermediate-vertex port set.
Derandomized squaring replaces each such complete graph by a copy of a label expander.
Therefore, after matching the normalizations,

$$\boxed{\mathcal{L}_{\mathrm{DSq}(G,H)} \approx_{O(\varepsilon)} \mathcal{L}_{G^2}.}$$

Spectral approximation preserves the second eigenvalue

Suppose

$$L_G \approx_\varepsilon L_{G'}.$$

For every $x \perp \mathbf{1}$,

$$\frac{1}{1+\varepsilon} x^\mathsf{T} L_G x \le x^\mathsf{T} L_{G'} x \le (1+\varepsilon) x^\mathsf{T} L_G x.$$

Taking the minimum Rayleigh quotient over $x \perp \mathbf{1}$,

$$\boxed{\frac{1}{1+\varepsilon} \lambda_2(L_G) \le \lambda_2(L_{G'}) \le (1+\varepsilon) \lambda_2(L_G).}$$

Thus replacing local cliques by expanders preserves the global gap up to a controlled multiplicative factor.

Derandomized squaring nearly doubles a small gap

Ordinary squaring transforms a gap γ into

$$2\gamma - \gamma^2.$$

If the local label expander gives an ε-spectral approximation, then

$$\gamma_{\mathrm{DSq}} \geq \frac{2\gamma - \gamma^2}{1+\varepsilon}.$$

For sufficiently small constant ε, and while $\gamma \leq 1/10$,

$$\gamma_{\mathrm{DSq}} \geq 1.8\gamma.$$

The operation improves the gap geometrically while multiplying the degree only by the constant degree of the label expander.

Part VI

From a weak gap to a nearly complete graph

Phase I boosts the gap to a constant

Begin with a constant-degree connected component satisfying

$$\gamma_0 \geq N^{-O(1)}.$$

Repeatedly apply derandomized squaring using constant-quality label expanders. While the gap is small,

$$\gamma_{s+1} \geq 1.8\gamma_s.$$

After

$$T_1 = O(\log N)$$

rounds,

$$\gamma_{T_1} = \Omega(1).$$

The degree has grown by a constant factor per round, so

$$\Delta' = N^{O(1)}.$$

A first deterministic attempt uses too much space

Once the graph is an expander, any two vertices in the same component are connected by a walk of length

$$T_2 = O(\log N).$$

But the degree is now

$$\Delta' = N^{O(1)},$$

so one edge label requires

$$\log \Delta' = O(\log N)$$

bits.

A full length-T_2 instruction string therefore requires

$$T_2 \log \Delta' = O(\log^2 N)$$

bits.

This is deterministic, but it misses the $O(\log N)$ space target.

Instead of storing a path, create one final edge

Enumerating all walks of length $O(\log N)$ is equivalent to examining edges of a large graph power.
In a rapidly mixing graph, a sufficiently high power is almost the complete averaging graph:

$$P^t \approx \frac{1}{N}\mathbf{1}\mathbf{1}^{\mathsf{T}}.$$

Rather than storing $O(\log N)$ large-degree edge labels, continue derandomized squaring until every pair of vertices in the same component is joined by a single transformed edge.
Then only one final edge label must be stored.

Phase II squares the nontrivial spectrum

Suppose the nontrivial eigenvalues of the current random-walk matrix satisfy

$$|\mu_i| \le \varepsilon_s.$$

Ordinary squaring would give

$$|\mu_i(P^2)| \le \varepsilon_s^2.$$

Choose a label expander whose nontrivial spectral radius is at most

$$\varepsilon_s^2.$$

The derandomized-square operator then satisfies the standard estimate

$$\rho_{s+1} \le \rho_s^2 + \rho(H_s) \le 2\varepsilon_s^2.$$

Set

$$\boxed{\varepsilon_{s+1} := 2\varepsilon_s^2.}$$

Why the eigenvalue errors add

The walk operator acts on a lifted vertex–port space.

Decompose any nontrivial vector into:

- a component constant inside each port cloud;
- a component orthogonal to the constant vector inside each cloud.

On the cloud-constant component:

two large-graph steps contract by at most $\rho(G)^2$.

On the cloud-orthogonal component:

the label expander contracts by at most $\rho(H)$.

The triangle inequality yields

$$\rho(\mathrm{DSq}(G, H)) \leq \rho(G)^2 + \rho(H).$$

This is the operator-level reason that pseudorandom second steps suffice.

The error becomes tiny after $O(\log \log N)$ rounds

Define

$$a_s := \log \frac{1}{\varepsilon_s}.$$

The recurrence

$$\varepsilon_{s+1} = 2\varepsilon_s^2$$

becomes

$$a_{s+1} = 2a_s - \log 2.$$

Once ε_0 is a sufficiently small constant, a_s grows geometrically.
Therefore, after

$$T_2 = O(\log \log N)$$

rounds, we obtain

$$\varepsilon_{T_2} < N^{-100}.$$

Repeated squaring has turned constant contraction into superpolynomially small

High-accuracy label expanders increase the degree

To obtain label-expander error at most ε_s^2, use an explicit expander of degree

$$d_s = \left(\frac{1}{\varepsilon_s}\right)^{O(1)}.$$

A derandomized square multiplies the current degree by d_s. Hence

$$\Delta_{\text{final}} = \Delta' \prod_{s<T_2} d_s.$$

Taking logarithms,

$$\log \Delta_{\text{final}} = \log \Delta' + \sum_{s<T_2} \log d_s.$$

The total degree remains polynomial

Since

$$\log d_s = O\left(\log \frac{1}{\varepsilon_s}\right) = O(a_s),$$

we have

$$\log \Delta_{\text{final}} \leq O(\log N) + O\left(\sum_{s<T_2} a_s\right).$$

The sequence a_s grows geometrically, so its sum is dominated by its final term:

$$\sum_{s<T_2} a_s = O(a_{T_2}) = O(\log N).$$

Consequently,

$$\boxed{\log \Delta_{\text{final}} = O(\log N)}$$

and

$$\Delta_{\text{final}} = N^{O(1)}.$$

Tiny nontrivial spectrum forces every entry to be positive

Let P be the random-walk matrix of a connected regular graph on M vertices. Suppose every nontrivial eigenvalue satisfies

$$|\mu_i| \le \frac{1}{2M}.$$

Lemma

Every entry of P is positive:

$$P_{ij} > 0 \qquad \text{for all } i, j.$$

Therefore every ordered pair of vertices is connected by an edge in the corresponding regular multigraph.

Separate the uniform component from the error

Let

$$J := \frac{1}{M}\mathbf{1}\mathbf{1}^{\mathsf{T}}$$

and define

$$E := P - J.$$

The matrix J acts as the identity on the uniform eigenspace and as zero on $\mathbf{1}^{\perp}$. Since P is symmetric and every nontrivial eigenvalue has magnitude at most $1/(2M)$,

$$\|E\|_2 \leq \frac{1}{2M}.$$

For standard basis vectors e_i, e_j,

$$P_{ij} = e_i^{\mathsf{T}} P e_j = \frac{1}{M} + e_i^{\mathsf{T}} E e_j.$$

The spectral error cannot cancel the uniform entry

By the operator-norm inequality,

$$|e_i^{\mathsf{T}} E e_j| \le \|e_i\|_2 \, \|E\|_2 \, \|e_j\|_2 \le \frac{1}{2M}.$$

Therefore

$$P_{ij} \ge \frac{1}{M} - \frac{1}{2M} = \boxed{\frac{1}{2M} > 0}.$$

Positive transition probability means that the final graph has at least one edge from j to i.
Thus the transformed graph is complete inside each connected component.

Disconnected graphs become unions of complete components

All transformations preserve connected components:

- squaring never creates an edge between different components;
- local clique replacement never joins different components;
- zig-zag steps cross only edges inherited from the large graph.

Analyze each connected component C separately. If

$$\rho(P_C) \leq \frac{1}{2N},$$

where N bounds the total number of vertices, then

$$(P_C)_{ij} \geq \frac{1}{|C|} - \frac{1}{2N} > 0.$$

The final transformed graph is therefore a disjoint union of complete regular multigraphs.

The final connectivity test checks one edge

After the two phases:

- the transformed graph has the same connected components;
- every two vertices in one component are adjacent;
- the final degree is
$$\Delta_{\text{final}} = N^{O(1)};$$
- an edge label uses
$$\log \Delta_{\text{final}} = O(\log N)$$
bits.

Enumerate every edge label at s, compute the corresponding final neighbor, and accept exactly when one of them is t.
Only one vertex, one edge label, and a few counters are stored.

What remains hidden in the implementation

The conceptual proof uses three technical ingredients that must be uniform in logarithmic space:

1. explicit label expanders for many label-set sizes and accuracies;
2. rotation maps for regularized, squared, and product graphs;
3. recursive evaluation of a final neighbor without materializing any intermediate graph.

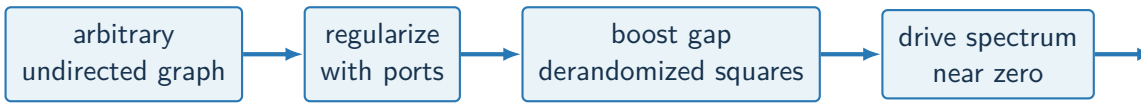
Reingold's construction arranges the transformations and labels so that only

$$O(\log N)$$

bits are live at once; discarded intermediate values can be recomputed when necessary.

The resulting algorithm has polynomial running time.

The proof transforms randomness into structure



Random walks suggested the algorithm. Expanders made short walks sufficient. Derandomized products encoded those walks without storing random choices.

The main ideas

1. A sparse cut traps a random walk:

$$\phi_t(S) \leq t\phi(S).$$

2. Cheeger's inequality relates conductance and spectral gap.
3. Random walks solve connectivity in randomized logarithmic space.
4. Connectivity in a union of constant-degree expanders is deterministically enumerable in logarithmic space.
5. Squaring improves the spectrum but explodes the degree.
6. Zig-zag products and derandomized squaring preserve spectral progress while controlling degree.
7. Driving the nontrivial spectrum below $1/(2N)$ makes every component complete.

Undirected connectivity needs only logarithmic workspace.

Expansion replaces long random exploration
by short deterministic structure.