The derandomized square operation (lecture notes)

Course: Derandomizing Space-Bounded Computation, Winter 2025, University of Chicago
Instructor: William Hoza (williamhoza@uchicago.edu)

In these notes, we prove the following theorem.

Theorem 0.1. *The undirected s-t connectivity problem is in L (deterministic log space).*

Theorem 0.1 was first proved by Reingold [Rei08], hence it is often called Reingold's theorem. We will present an alternative proof due to Rosenman and Vadhan [RV05].

1 A connectivity algorithm based on the INW generator

We begin by arguing that we can make various convenient assumptions without loss of generality.

Definition 1.1 (Consistent labeling)**.** Let G be a D-outregular directed multigraph on the vertex set $[N]$. We say that G is *labeled* if, for every vertex s, the outgoing edges from s have distinct labels in $[D]$. In this case, if an edge (s,t) has the label $x \in [D]$, then we write $G[s,x] = t$. That is, $G[s,x]$ is the x-th neighbor of s. Note that the labeling of G induces a labeling of G^k, namely, $G^{k+1}[s,xy] = G[G^k[s,x],y]$.

We say that G is *consistently labeled* if, for every vertex s, the incoming edges at s all have distinct labels. In other words, $G[s,x] = G[t,x]$ implies $s = t$. This is only possible if G is D-regular, i.e., every vertex has D incoming edges as well as D outgoing edges.

Lemma 1.2 (Reducing to the 4-regular consistently-labeled case)**.** *There is a deterministic log-space reduction from the undirected s_*-t_* connectivity problem to the problem of deciding s_*-t_* connectivity in a 4-regular consistently labeled directed multigraph in which every vertex has a self-loop.*

Proof. Let G be the given undirected graph. Without loss of generality, we assume that the vertex set is $[N]$. For each vertex $s \in [N]$, let the neighbors of s be $G[s,1] < G[s,2] < \cdots < G[s,\deg(s)]$. Our new graph G' is on the vertex set $\{(s,x) : s \in [N], x \in [\deg(s)]\}$. The edge set and the labels are defined by

$$\begin{aligned}
G'[(s,x),1] &= (s,x)\\
G'[(s,x),2] &= (s, x+1 \bmod \deg(s))\\
G'[(s,x),3] &= (s, x-1 \bmod \deg(s))\\
G'[(s,x),4] &= (G[s,x], y) \text{ where } G[G[s,x],y] = s.
\end{aligned}$$

One can verify that every vertex has in-degree 4 and the labeling is consistent. The new "s_*" is $(s_*,1)$, and the new "t_*" is $(t_*,1)$. □

Now let G be a 4-regular consistently labeled directed multigraph on N vertices. Our approach for solving s_*-t_* connectivity on G will be to design a pseudorandom generator $\mathsf{GEN} \colon \{0,1\}^r \to [4]^n$ for a suitable $n = \text{poly}(N)$. The algorithm: Accept iff there exists a seed x such that $G^n[s_*, \mathsf{GEN}(x)] = t_*$.

The generator GEN is, in fact, the INW generator, with its parameters tweaked so that it has seed length $r = O(\log N)$. In more detail, we have a family of generators $\mathsf{GEN}_i \colon \{0,1\}^{r_i} \to [4]^{2^i}$. We start with $r_0 = 2$ and $\mathsf{GEN}_0(x) = x$. Then we define

$$\mathsf{GEN}_{i+1}(x,y) = (\mathsf{GEN}_i(x), \mathsf{GEN}_i(H_i[x,y])),$$

where H_i is an ε_i-spectral expander on the vertex set $\{0,1\}^{r_i}$. Crucially, we will choose relatively "mild" expanders, i.e., the values ε_i will not be very small. Let us defer specifying the exact formula for ε_i until later, but for now we just mention that we will choose $\varepsilon_i = \Omega(1)$ for almost every i.

If there is no path from s_* to t_*, then clearly the algorithm correctly outputs "no." The more difficult case is when there is a path from s_* to t_*. In this case, by focusing on the connected component containing s_* and t_*, we may assume that the graph G is strongly connected. In this case, the algorithm's correctness proof is based on analyzing the *derandomized square* operation, defined next.

2 The derandomized square operation

Definition 2.1 (Derandomized square)**.** Let G be a consistently labeled D-regular graph on the vertex set $[N]$. Let H be a labeled d-regular graph on the vertex set $[D]$. The *derandomized square* $G \textcircled{s} H$ is a labeled (Dd)-outregular graph on the vertex set $[N]$ given by

$$(G \textcircled{s} H)[s,(x,y)] = G^2[s,(x,H[x,y])].$$

In Definition 2.1, we assume that G is consistently labeled. One can define the derandomized square without this assumption, but the "correct" definition is a bit subtle. Fortunately, the derandomized squaring property preserves consistent labeling, so we do not need to worry about graphs that are not consistently labeled.

Proposition 2.2 (Derandomized squaring preserves consistent labeling)**.** *If G is consistently labeled, then $G \textcircled{s} H$ is consistently labeled.*

Proof. If $(G \textcircled{s} H)[s,(x,y)] = (G \textcircled{s} H)[s',(x,y)]$, then $G[t,H[x,y]] = G[t',H[x,y]]$, where $t = G[s,x]$ and $t' = G[s',x]$. But G is consistently labeled, so $t = t'$ and hence $s = s'$. □

There is a close connection between the derandomized square operation and the INW generator. To see it, let G_i be the graph $G_i[s,x] = G^{2^i}[s,\mathsf{GEN}_i(x)]$. Then looking back through the definitions, we see that for every i, we have

$$G_{i+1} = G_i \textcircled{s} H_i.$$

From here, our analysis will be similar to our analysis of random walks in undirected graphs: we will show that $\lambda(G_i)$ rapidly goes to zero. The idea is that $G \textcircled{s} H$ approximates G^2. The true square G^2 satisfies $\lambda(G^2) = \lambda(G)^2$. Now we show that the derandomized square $G \textcircled{s} H$ satisfies a bound that is nearly as good.

Theorem 2.3 (Derandomized square approximates true square)**.** *Let G be any consistently labeled D-regular graph on the vertex set $[N]$. Let H be an ε-spectral expander on the vertex set $[D]$, i.e., $\lambda(H) \le \varepsilon$. Then*

$$\lambda(G \textcircled{s} H) \le (1-\varepsilon)\cdot\lambda(G)^2 + \varepsilon \le \max\{\lambda(G)^{1.5}, 4\varepsilon\}.$$

To prove the theorem, let us adopt the convenient convention of identifying each graph with its transition probability matrix. Let's think about what happens if we start at a vertex $s \in [N]$ and take a random step in the graph $G \textcircled{s} H$. We can break the random step into five substeps:

$$[N] \quad\to\quad [N]\times[D] \quad\to\quad [N]\times[D] \quad\to\quad [N]\times[D] \quad\to\quad [N]\times[D] \quad\to\quad [N].$$

The five substeps are as follows.

1. Step 1: Pick a random edge label $x \in [D]$ and move to (s,x). The corresponding transition probability matrix $L \in \mathbb{R}^{D\times ND}$ applies the map $\pi L = \pi \otimes u$.
2. Step 2: Move to $(s',x) = (G[s,x],x)$. Since G is consistently labeled, the corresponding transition "probability" matrix $A \in \mathbb{R}^{ND\times ND}$ is a permutation matrix, i.e., πA just permutes the coordinates of π.
3. Step 3: Pick a random edge label $y \in [d]$ and move to $(s',x') = (s',H[x,y])$. The corresponding transition probability matrix is the tensor product $I_N \otimes H$.
4. Step 4: Move to $(t,x') = (G[s',x'],x')$. This is another application of A.
5. Step 5: Delete the second coordinate, i.e., move to $t \in [N]$. The corresponding transition probability matrix $P \in \mathbb{R}^{ND\times D}$ applies the map $(\pi P)_t = \sum_{x'} \pi_{(t,x')}$.

Thus, the transition probability matrix of $G \circledS H$ is given by

$$G \circledS H = LA(I_N \otimes H)AP. \tag{1}$$

The next step is to apply the so-called Expander Decomposition Lemma.

Lemma 2.4 (Expander Decomposition Lemma)**.** *Let H be the transition probability matrix of an ε-spectral expander on the vertex set $[D]$. Let J_D denote the $D \times D$ matrix where every entry is $1/D$. There exists a matrix $E \in \mathbb{R}^{D \times D}$ such that $\|E\|_{\mathrm{op}} \le 1$[1] and $H = (1-\varepsilon) \cdot J_D + \varepsilon \cdot E$.*

Proof. Let $E = (1/\varepsilon) \cdot (H - (1-\varepsilon) \cdot J_D)$. Let v be any unit vector, and decompose it as $v = v^{\parallel} + v^{\perp}$, where $v^{\parallel}$ is parallel to u and $v^{\perp}$ is perpendicular to u. Then

$$\begin{aligned} \|vE\|_2^2 = \varepsilon^{-2} \cdot \|vH - (1-\varepsilon) \cdot vJ_D\|_2^2 = \varepsilon^{-2} \cdot \|\varepsilon v^{\parallel} + v^{\perp} H\|_2^2 &= \varepsilon^{-2} \cdot \left(\|\varepsilon v^{\parallel}\|_2^2 + \|v^{\perp} H\|_2^2 \right) \\ &\le \|v^{\parallel}\|_2^2 + \|v^{\perp}\|_2^2 \\ &= 1. \end{aligned}$$

□

Proof of Theorem 2.3. Applying the Expander Decomposition Lemma to Eq. (1), we get

$$G \circledS H = (1-\varepsilon) \cdot LA(I_N \otimes J_D)AP + \varepsilon \cdot LA(I_N \otimes E)AP.$$

The first term is the transition probability matrix of $G \circledS J$, i.e., the true square G^2. Therefore, if v is any unit vector orthogonal to the uniform distribution, we have

$$\begin{aligned} \|v(G \circledS H)\|_2 &= \|(1-\varepsilon) \cdot vG^2 + \varepsilon \cdot vLA(I_N \otimes E)AP\|_2 \\ &\le (1-\varepsilon) \cdot \lambda(G)^2 + \varepsilon \cdot \|LA(I_N \otimes E)AP\|_{\mathrm{op}} \\ &\le (1-\varepsilon) \cdot \lambda(G)^2 + \varepsilon \cdot \|L\|_{\mathrm{op}} \cdot \|A\|_{\mathrm{op}} \cdot \|I_N \otimes E\|_{\mathrm{op}} \cdot \|A\|_{\mathrm{op}} \cdot \|P\|_{\mathrm{op}}. \end{aligned}$$

Let us calculate each operator norm term.

- If v is any unit vector, then $\|vL\|_2 = \|v \otimes u\|_2 = \|v\|_2 \cdot \|u\|_2 = 1/\sqrt{D}$, so $\|L\|_{\mathrm{op}} = 1/\sqrt{D}$.
- Since A is a permutation matrix, we have $\|A\|_{\mathrm{op}} = 1$.
- The operator norm of a tensor product is the product of the operator norms, so $\|I_N \otimes E\|_{\mathrm{op}} = \|I_N\|_{\mathrm{op}} \cdot \|E\|_{\mathrm{op}} \le 1$.
- If v is any unit vector, then $\|vP\|_2^2 = \sum_t (\sum_{x'} v_{(t,x')})^2 \le D \sum_{t,x'} v_{(t,x')}^2 = D$, so $\|P\|_{\mathrm{op}} \le \sqrt{D}$.

It follows that

$$\lambda(G \circledS H) \le (1-\varepsilon) \cdot \lambda(G)^2 + \varepsilon.$$

Finally, to prove that $(1-\varepsilon) \cdot \lambda(G)^2 + \varepsilon \le \max\{\lambda(G)^{1.5}, 4\varepsilon\}$, let $\lambda = \lambda(G)$ for brevity's sake, and split into two cases. For the first case, suppose ε is small, namely

$$\varepsilon \le \lambda^{1.5} \cdot \frac{1-\sqrt{\lambda}}{1-\lambda^2}.$$

Then

$$(1-\varepsilon) \cdot \lambda^2 + \varepsilon = \lambda^2 + \varepsilon \cdot (1-\lambda^2) \le \lambda^2 + \lambda^{1.5} \cdot (1-\sqrt{\lambda}) = \lambda^{1.5}.$$

Now, for the second case, suppose ε is large, namely

$$\varepsilon > \lambda^{1.5} \cdot \frac{1-\sqrt{\lambda}}{1-\lambda^2} = \lambda^{1.5} \cdot \frac{1-\sqrt{\lambda}}{(1-\lambda)(1+\lambda)} = \lambda^{1.5} \cdot \frac{1-\sqrt{\lambda}}{(1-\sqrt{\lambda})(1+\sqrt{\lambda})(1+\lambda)} \ge \frac{\lambda^{1.5}}{4}.$$

[1] I.e., $\|vE\|_2 \le \|v\|_2$ for every $v \in \mathbb{R}^D$.

Then

$$(1-\varepsilon)\cdot\lambda^2+\varepsilon<(1-\varepsilon)\cdot(4\varepsilon)^{4/3}+\varepsilon.$$

It is clear that the expression above is $O(\varepsilon)$. To prove the specific bound of 4ε, let $p=1-\varepsilon\in[0,1]$. By taking a derivative with respect to p, one sees that p^3-p^4 is maximized at $p=3/4$, i.e., $p^3\cdot(1-p)\le 3^3/4^4$. Taking a cube root, we get $4^{4/3}\cdot p\cdot(1-p)^{1/3}\le 3$, i.e., $4^{4/3}\cdot(1-\varepsilon)\cdot\varepsilon^{1/3}\le 3$. Finally, adding one and multiplying by ε gives us $(1-\varepsilon)\cdot(4\varepsilon)^{4/3}+\varepsilon\le 4\varepsilon$. □

Proof of Theorem 0.1. Let G_0 be a 4-regular strongly connected directed multigraph in which every vertex has at least one self-loop. Our analysis of random walks on undirected graphs shows that there is a value $\lambda_0=1-1/O(N^2)$ such that $\lambda(G_0)\le\lambda_0$. Define $\lambda_i=\lambda_0^{1.5^i}$ and $\varepsilon_i=\frac{1}{4}\cdot\lambda_i^{1.5}$. We use ε_i as our expansion parameter for the expander graph H_i that we use to construct the PRG GEN_{i+1}. Our final generator is $\mathsf{GEN}=\mathsf{GEN}_{i_*}$, where i_* is the first value such that $\lambda_{i_*}<1/N$.

Define the graphs $G_1,G_2,\ldots$ as in Section 2, namely $G_i[s,x]=G^{2^i}[s,\mathsf{GEN}_i(x)]$. By Theorem 2.3 and induction, we have $\lambda(G_i)\le\lambda_i$. The fact that $\lambda(G_{i_*})<1/N$ implies that every two vertices in G_{i_*} are neighbors, hence our algorithm is correct.

Now let us analyze the seed length of the PRG. Since $\lambda_i=\lambda_0^{1.5^i}\le\exp(-1.5^i/O(N^2))$, we have $i_*=O(\log N)$. Furthermore, we can choose H_i to be an explicit expander with $\deg(H_i)=\mathrm{poly}(1/\varepsilon_i)$. Therefore, the seed length of GEN is given by

$$\begin{aligned}s_{i_*}=O\left(\sum_{i=1}^{i_*}\log(4/\lambda_0^{1.5^i})\right)=O\left(i_*+\log(1/\lambda_0)\cdot\sum_{i=1}^{i_*}1.5^i\right)&=O(i_*+\log(1/\lambda_0)\cdot 1.5^{i_*})\\&=O(i_*+\log(1/\lambda_{i_*}))\\&=O(\log N).\end{aligned}$$

Consequently, our algorithm only uses $O(\log N)$ bits of space. □