CSG399: Gems of Theoretical Computer Science. Lecs. 15-18. Feb. 27-Mar. 17, 2009.
Instructor: Emanuele Viola Scribe: Zhifeng Sun

Undirected reachability

In this lecture we study algorithms that solve the undirected reachability problem in $O(\log n)$-space. The undirected reachability problem is defined as follows.

Problem 1. *Given an undirected graph G with n vertices and two special vertices in the graph, s and t, decide whether s and t are connected or not.*

If we don't consider the space constraint, we can solve this problem using standard DFS or BFS algorithms, which have running time $O(|V| + |E|)$, but take space $O(|V|)$. We are interested in algorithms that use space $O(\log n)$ and run in time poly(n). In particular, we can only store $\Theta(1)$ vertices, since each vertex description takes $O(\log n)$ bits where n is the total number of vertices.

This lecture note is organized as follows. In Section 1, we are going to cover mathematical preliminaries. In Section 2 we give a randomized algorithm. In Section 3 we give a deterministic algorithm.

1 Preliminaries

In this section we cover some preliminaries. First, we list some basic definitions and properties from linear algebra.

- A vector $v = (v_1, v_2, \ldots, v_n) \in \mathbb{R}^n$.
- Inner product $< v, w >= \sum_i v_i \cdot w_i$.
- Two vectors are orthogonal, denoted $v \perp w$, if $< v, w >= 0$.
- The length of a vector is $\|v\| := \sqrt{\sum_i v_i^2} = \sqrt{< v, v >}$.
- $\|v + w\| \leq \|v\| + \|w\|$.
- If $v_1, v_2, \ldots, v_n$ are pairwise orthogonal, then they are independent, i.e. $\nexists (a_1, a_2, \ldots, a_n) \neq 0$ such that $\sum_i a_i v_i = 0$.

Lemma 1. *If $v \perp w$, then $\|v + w\|^2 = \|v\|^2 + \|w\|^2$.*

Proof.

$$\begin{aligned} \|v + w\|^2 &= < v + w, v + w > \\ &= \sum_i (v_i + w_i)^2 \end{aligned}$$

$$
\begin{aligned}
&= \sum_i \left(v_i^2 + 2v_iw_i + w_i^2\right) \\
&= \sum_i v_i^2 + \sum_i w_i^2 + 2\sum_i v_iw_i \\
&= \|v\|^2 + \|w\|^2 + 2\cdot 0 \\
&= \|v\|^2 + \|w\|^2.
\end{aligned}
$$

□

Now we view a vector as a probability distribution over the n vertices of G. For example, $u = (1/n, \ldots, 1/n)$ is the uniform distribution over all the vertices. For now, we assume that the graph G is *regular*, i.e. each vertex has degree d, and also that each vertex has a self-loop. We justify this assumption later. For each graph G, we have a normalized adjacency matrix A, where "normalized" means each entry is divided by the degree d. The following matrix is the adjacency matrix of the graph in Figure 1.

$$
A = \begin{bmatrix} 1/3 & 1/3 & 1/3 \\ 1/3 & 2/3 & 0 \\ 1/3 & 0 & 2/3 \end{bmatrix}.
$$

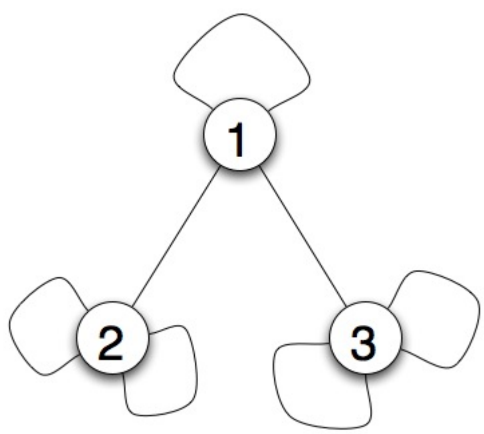

Figure 1: Graph example.

Let $v = (1, 0, 0)$. Then $Av = (1/3, 1/3, 1/3)$ is the probability distribution after doing one random step starting at v. Naturally, we are interested in $A^\ell v$ for large ℓ. It would be great if the act of multiplying by A corresponded to some simple behavior, like $Av = \lambda \cdot v$ for some scalar λ. Such λ and v are called *eigenvalue* and *eigenvector* respectively. Although every matrix has them, in general they need to be complex, as for example is needed if A corresponds to a rotation of the space. However, the matrices that arise from graphs have a special structure, in particular they are symmetric. One can show that in this case there are

n eigenvectors with corresponding eigenvalues all of which (both the eigenvectors and the eigenvalues) are real. Moreover, we can choose the eigenvectors to be orthonormal, i.e. length-1 vectors that have zero inner product. In particular, these vectors form a basis of the whole space. This allows us to write any vector v in this basis and see the act of multiplying the vector by the matrix A as simply multiplying each coordinate of v (w.r.t. the basis given by the eigenvectors) by the corresponding eigenvalues. We now state without proof this central fact from linear algebra, together with some other useful facts.

Theorem 2. *Let A be the normalized adjacency matrix of a connected, regular graph G on n vertices with a self-loop at each vertex. Then*

1. *there exist real vectors $v_1, v_2 \ldots, v_n \in \mathbb{R}^n$ such that $\forall i \ \|v_i\| = 1$, and $v_i \perp v_j$ if $i \neq j$,*
2. *there exist real numbers $\lambda_1, \ldots, \lambda_n \in \mathbb{R}$ such that $Av_i = \lambda_i v_i$,*
3. *$v_1 = (1/\sqrt{n}, \ldots, 1/\sqrt{n})$, $\lambda_1 = 1$, and $\forall i > 1$, $|\lambda_i| \leq 1 - \frac{1}{n^c}$ for some constant c. We let $1 = |\lambda_1| \geq |\lambda_2| \geq \ldots \geq |\lambda_n|$.*

The fact that $(1/\sqrt{n}, \ldots, 1/\sqrt{n})$ is an eigenvector with eigenvalue 1 is not hard to see using the fact that the graph is regular. That the other eigenvalues are bounded away from 1 by an inverse polynomial in n can be shown using the fact that the graph is connected.

The next lemma gives a useful characterization of the second largest eigenvalue.

Lemma 3. $|\lambda_2| = \max_{v \perp u} \frac{\|Av\|}{\|v\|}$ *where u is uniform distribution, $(1/n, \ldots, 1/n)$.*

Proof. Take any vector $v \in \mathbb{R}^n$. Write $v = c_1 v_1 + \ldots + c_n v_n$ ($v_1, \ldots, v_n$ are the eigenvectors of the matrix A, from Theorem 2). If $v \perp u$, then $c_1 = 0$ (note $v_1 = \sqrt{n} \cdot u$). So $v = c_2 v_2 + \ldots + c_n v_n$.

$$
\begin{aligned}
\|Av\|^2 &= \|\lambda_2 \cdot c_2 v_2 + \ldots + \lambda_n \cdot c_n v_n\|^2 \\
&= \|\lambda_2 c_2 v_2\|^2 + \ldots + \|\lambda_n c_n v_n\|^2 \quad \text{(By Lemma 1)} \\
&\leq \lambda_2^2 \left(c_2^2 \|v_2\|^2 + \ldots + c_n^2 \|v_n\|^2\right) \\
&= \lambda_2^2 \left(c_2^2 + \ldots + c_n^2\right) \\
&= \lambda_2^2 \|v\|^2.
\end{aligned}
$$

The equality is met for $v = v_2$. And this completes the proof of the lemma. □

The next useful lemma shows that any probability distribution v gets closer to uniform after multiplying by A.

Lemma 4. *Let v be a probability distribution, then $\|Av - u\| \leq \lambda_2 \|v - u\|$, where u is the uniform distribution, $(1/n, \ldots, 1/n)$.*

Proof. $\|Av - u\| = \|A(v-u)\|$. And we observe that $(v-u) \perp u$, because

$$
\begin{aligned}
< v-u, u > &= \sum_i (v_i - u_i) u_i \\
&= \frac{1}{n} \sum_i (v_i - u_i) \\
&= \frac{1}{n} \left(\sum_i v_i - \sum_i u_i \right) \\
&= \frac{1}{n}(1-1) \\
&= 0.
\end{aligned}
$$

Then by Lemma 3, we have $\|Av - u\| \leq \lambda_2 \|v - u\|$. □

2 Random walk algorithm

In this section we use the machinery from the previous section to give a randomized $O(\log n)$-space algorithm to solve the undirected reachability problem. The random walk algorithm is described in Algorithm 1.

```
input  : Graph G = (V, E), and vertices s and t.
output: Whether s and t are connected or not.

v ← s  repeat
    if v = t then
        return "connected"
    end
    v ← random neighbor of v
until l = n^100 times ;
return "not connected"
```

Algorithm 1: Random walk algorithm.

Theorem 5. *Random walk algorithm uses space* $O(\log n)$, *time* $\mathrm{poly}(n)$ *and satisfies* $\forall G, s, t$, *if* s *is connected to* t, *then* $\Pr[\text{algorithm returns "connected"}] \geq 1/2$; *if* s *is not connected to* t, *then* $\Pr[\text{algorithm returns "not connected"}] = 1$.

Proof. If s and t are not connected, then the random walk algorithm always returns "not connected." So we only need to show that when s and t are connected, the random walk algorithm returns "connected" with probability more than $1/2$. Let G be the connected graph that contains s and t. Take $l := n^c$ for some constant c.

Claim 1. *For any probability distribution* v, *a random walk of length* $\sqrt{l}$ *starting at* v *will end up in* t *with probability no less than* $1/2n$.

Proof.

$$\begin{aligned}
\|A^{\sqrt{l}}v - u\| &\le \lambda_2^{\sqrt{l}}\|v-u\| \quad \text{(By Lemma 4)}\\
&\le O\left(\lambda_2^{\sqrt{l}}\right)\\
&\le O\left(\left(1-\frac{1}{n^{c'}}\right)^{\sqrt{l}}\right) \quad \text{(By Theorem 2)}\\
&\le \frac{1}{n^2}. \quad \text{(for } l = n^c \text{ large enough)}
\end{aligned}$$

Therefore, $A^{\sqrt{l}}v$ puts mass at least $1/2n$ on t, because otherwise

$$\|A^{\sqrt{l}}v - u\| \ge \sqrt{\left(\frac{1}{2n} - \frac{1}{n}\right)^2} = \frac{1}{2n} \gg \frac{1}{n^2}$$

which is a contradiction. This completes the proof of the claim. □

Now consider the random walk after $j\sqrt{l}$ steps, $j = 1, 2, \ldots, \sqrt{l}$.

$$\begin{aligned}
& \Pr[\text{walk never touches } t]\\
\le\ & \Pr\left[\text{don't touch } t \text{ after } j\sqrt{l} \text{ steps } \forall j = 1, \ldots, \sqrt{l}\right]\\
\le\ & \left(1 - \frac{1}{2n}\right)^{\sqrt{l}} \quad \text{(By Claim 1)}\\
\le\ & \frac{1}{2}. \quad \text{(for sufficiently large } l = n^2\text{)}
\end{aligned}$$

And this completes the proof of this theorem. □

3 Deterministic algorithm

In this section we give a deterministic $O(\log n)$-space algorithm. Before we get into the deterministic algorithm for undirected reachability problem, we define a more "robust" version of the problem. The following claim also justifies our assumption that graphs are regular, which we used in Section 2.

Claim 2. *The problem of determining whether s is connected to t on an undirected graph G with n nodes and maximum degree n is $O(\log n)$-space reducible to the following problem: given (1) a graph G' on $\text{poly}(n)$ nodes which is 4-regular, and each node has a self-loop; (2) sets of nodes S and T (S is connected and T is connected), and $|S| \ge |G'|/3$, $|T| \ge |G'|/3$; decide whether S and T are in the same connected component.*

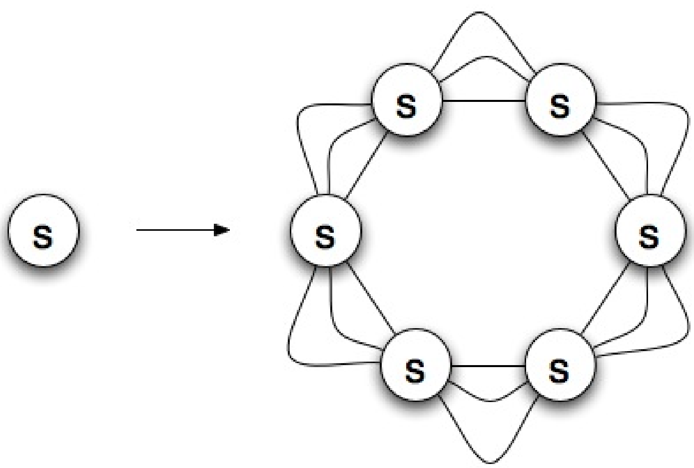

Figure 2: Duplicate s and add cycles.

Proof. Given the graph G, we construct G' in the following way. Add n copies of s and n copies of t. Let $S = \{s\} \cup \{\text{extra copies of } s\}$ and $T = \{t\} \cup \{\text{extra copies of } t\}$. Put $n/2$ copies of cycles on S, and similarly on T, as shown in Figure 2, so that the extra copies of s and t have degree n.

So far, this graph satisfies everything except the degree bound. Observe that the degree of each node is $\leq n$ except for s and t which may have degree as large as $2n$. To reduce the degree, replace each node of degree d with a 4-regular graph on d nodes, as shown in Figure 3. And call this final graph G'.

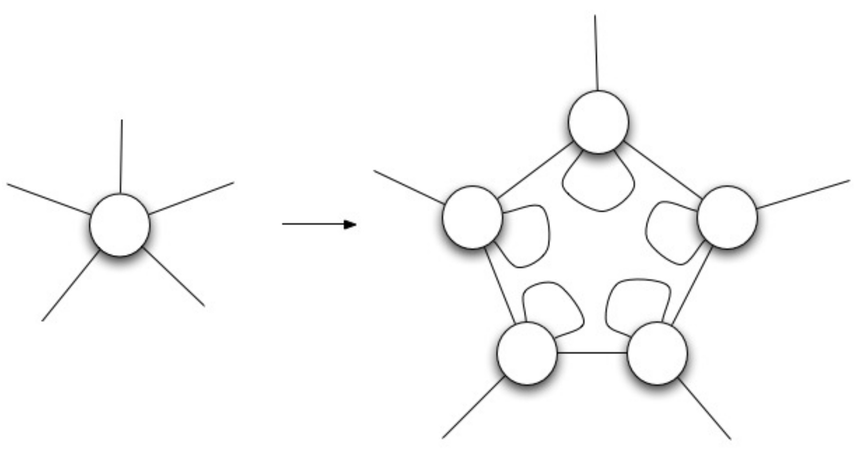

Figure 3: Replace each node with a 4-regular graph.

It is easy to observe that the number of nodes in G' is poly(n), that G' is 4-regular, and that $|S| \geq |G'|/3$, $|T| \geq |G'|/3$. The bound on $|S|, |T|$ follows because (1) this bound holds before reducing the degree, (2) the degree reduction blows up a node by its degree, and (3) every node in $S \cup T$ has degree at least as large as that of any other node in the graph. $\square$

We now show that this more "robust" versions is "easy" when the second eigenvalue is small. We do not restrict the degree: we will apply the following claim to graphs of polynomial degree. Jumping ahead, these will arise by modifying the 4-regular graphs given by Claim 2 using appropriate operations that reduce the eigenvalue bound.

Claim 3. *Let G be any graph where neighbors are log-space computable. Let S and T be connected sets and $|S| \geq |G|/3$, $|T| \geq |G|/3$. If $\lambda_2(G) \leq 1/10$, then*

$$\begin{aligned} & \textit{some node in } S \textit{ is } \textbf{\textit{connected}} \textit{ to some node in } T \\ \iff & \textit{some node in } S \textit{ is } \textbf{\textit{adjacent}} \textit{ to some node in } T \end{aligned}$$

Therefore, the problem of deciding whether S and T are in the same connected component can be solved deterministically in $O(\log n)$-space.

Proof. The "moreover" part follows easily from the first part of the claim by cycling over all nodes in S and their neighbors using logarithmic space. We now prove the first part of the claim. The "$\Longleftarrow$" direction is obvious. So in the following we focus on "$\Longrightarrow$" direction. The basic idea is: since S is large and λ_2 is small, S has many neighbors ($> 2n/3$), and one of them must be in T.

Let n be the number of nodes in the graph G, u represent the uniform distribution (i.e. $u = (1/n, 1/n, \ldots, 1/n)$), v represent the uniform distribution on S (i.e. $v = (3/n, \ldots, 3/n, 0, \ldots, 0)$ where the coordinates with mass $3/n$ are exactly those in S). And let $w = A \cdot v$ where A is the normalized adjacency matrix of G. Our goal is to show w has non-zero weight on some coordinate in T.

$$\begin{aligned} \|w - u\| &= \|A \cdot v - u\| \\ &\leq \lambda_2 \|v - u\| \qquad \text{(by lemma 4)} \\ &= \frac{1}{10}\sqrt{\frac{n}{3}\left(\frac{3}{n} - \frac{1}{n}\right)^2 + \frac{2n}{3}\left(0 - \frac{1}{n}\right)^2} \\ &= \frac{1}{10}\sqrt{\frac{2}{n}}. \end{aligned}$$

Now assume that w has zero in every coordinate in T, then

$$\|w - u\| = \sqrt{\sum_i (w_i - u_i)^2} \geq \sqrt{\sum_{i \in T} \frac{1}{n^2}} \geq \sqrt{\frac{n}{3} \cdot \frac{1}{n^2}} = \sqrt{\frac{1}{3n}} > \frac{1}{10}\sqrt{\frac{2}{n}}.$$

which is a contradiction that completes the proof of the claim. $\square$

We are now left with the task of reducing the eigenvalue bound of our graph. Before discussing this, let us make a remark on the notion of *degree* of a graph. It is convenient to work with the following definition of degree:

Definition 6. *The degree of a graph G is D if the graph can be specified by a neighbor function*

$$f : V \times \{1, 2, \ldots, D\} \to V,$$

which given a node and an edge index returns the corresponding neighbor.

Note the degree in this definition can be arbitrarily large, in particular larger than n.

3.1 An attempt to reduce the eigenvalue bound

One attempt to reduce the eigenvalue of a graph is by squaring. The squared graph is the graph in which edges correspond to paths of length 2 in the original graph; Figure 4 shows an example.

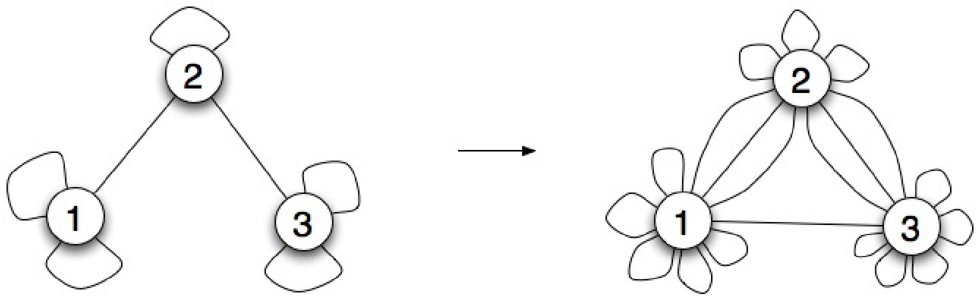

Figure 4: Squaring a graph.

Claim 4. *If G (with self-loop on each node) has n nodes, degree d, and $\lambda_2(G) = \alpha$, then G^2 has n nodes, degree d^2, and $\lambda_2(G^2) = \alpha^2$.*

Proof. It is easy to see G^2 has n nodes and degree d^2. So the following proof focuses on the eigenvalue. Let M be the normalized adjacency matrix of G, then M^2 is the normalized adjacency matrix of G^2. To bound $\lambda_2(G^2)$, let $v \perp u$ where u is the vector which represents the uniform distribution. Notice that $Mv \perp u$.

$$\lambda_2\left(G^2\right) = \max_{v \perp u} \frac{\|M^2 v\|}{\|v\|} = \max_{v \perp u} \frac{\|M^2 v\|}{\|Mv\|} \cdot \frac{\|Mv\|}{\|v\|} = \alpha^2.$$

□

The following is an attempt to show undirected reachability is in $O(\log n)$-space: given graph G and sets S, T, square the graph $l = O(\log n)$ times to obtain G^{2^l}, so that λ_2 becomes $1/10$.

$$\lambda_2\left(G^{2^l}\right) = \lambda_2(G)^{2^l} \le \left(1 - \frac{1}{n^c}\right)^{2^l} = \frac{1}{10}.$$

Although we obtain the desired eigenvalue bound, the degree of the graph G^{2^l} is $D^{2^l} = D^{\text{poly}(n)}$, which is exponential in n. This means we cannot apply Claim 3 to determine connectivity in logarithmic space, since the idea there was to cycle over all neighbors of nodes in S. So in order to apply claim 3, we need to give another operation that can decrease the value of λ_2 while at the same time keeping the degree small.

Note that an edge in G^{2^l} corresponds to a path of length poly(n) in G. With the new operation we will still have that an edge in the final graph corresponds to a path of that length in G. But the crucial difference is this: whereas in G^{2^l} we consider all, exponentially many paths, in the new graph we only consider a sparse, polynomial-size collection of paths. We will prove that this sparse collection has the same hitting properties of the collection of all paths, as measured by the eigenvalue bound.

3.2 Reducing the eigenvalue via derandomized graph squaring

Definition 7 (derandomized graph squaring)**.** *Let X be a k-regular graph on n nodes, and G be a d-regular graph on k nodes. $X \circledS G$ is a graph on n nodes with degree $k \cdot d$. The neighbors of v in $X \circledS G$ are $v[a][b]$ where b is a neighbor of a in graph G, i.e. $v[a][a[e]]$ where $a \in \{1, 2, \ldots, k\}$ and $e \in \{1, 2, \ldots, d\}$.*

Note that in the above definition we see a as both an edge index for X and a node in G. Whereas in graph squaring the neighbors of v are $v[a][b]$ for *every* a, b, in derandomized graph squaring the neighbors are $v[a][b]$ for *some* a, b.

Now we show that by applying derandomized graph squaring we decrease λ_2. We start with a useful lemma that shows that a random step in a graph G with $\lambda_2(G) = \lambda$ can be seen as going to the uniform distribution with probability $(1 - \lambda)$, and not doing too bad otherwise.

Lemma 8. *If $\lambda_2(G) = \lambda$ and A is the normalized adjacency matrix of G, then $A = (1 - \lambda) J_n + \lambda C$ where J_n is the $n \times n$ matrix with $1/n$ everywhere, and C satisfies $\forall v : \|Cv\| \leq \|v\|$.*

Proof. Let $C := \frac{1}{\lambda}[A - (1 - \lambda) J_n]$. We are going to show that for any v we have $\|Cv\| \leq \|v\|$. We can write $v = a \cdot u + w$ where a is a constant, u represents uniform distribution and $u \perp w$.

$$
\begin{aligned}
\|Cv\|^2 &= \|aCu + Cw\|^2 \\
&= \|a\frac{1}{\lambda}[A - (1 - \lambda) J_n] u + \frac{1}{\lambda}[A - (1 - \lambda) J_n] w\|^2 \\
&= \|a\frac{1}{\lambda}[Au - (1 - \lambda) J_n u] + \frac{1}{\lambda}[Aw - (1 - \lambda) J_n w]\|^2 \\
&= \|a\frac{1}{\lambda}[u - (1 - \lambda) u] + \frac{1}{\lambda}[Aw - (1 - \lambda) J_n w]\|^2 \qquad (Au = u) \\
&= \|au + \frac{1}{\lambda}[Aw - (1 - \lambda) J_n w]\|^2 \\
&= \|au + \frac{1}{\lambda}Aw\|^2 \qquad (J_n w = 0 \text{ because } u \perp w)
\end{aligned}
$$

$$\begin{aligned}
&= \|au\|^2 + \|\frac{Aw}{\lambda}\|^2 \quad (au \perp Aw) \\
&\leq \|au\|^2 + \|w\|^2 \quad \text{(by lemma 3)} \\
&= \|v\|^2.
\end{aligned}$$

This completes the proof of the lemma. □

To analyze the adjacency matrices arising from derandomized graph squaring we use the notion of tensor product of matrices.

Definition 9 (tensor product)**.** *Let A be a $n \times m$ matrix, and B be a $n' \times m'$ matrix. The tensor product of A and B is defined as*

$$A \otimes B = \begin{bmatrix} a_{11}B & \dots & a_{1n}B \\ \vdots & \dots & \vdots \\ a_{n1}B & \dots & a_{nn}B \end{bmatrix},$$

which is a $n \cdot n' \times m \cdot m'$ matrix.

Theorem 10. *If $\lambda_2(X) = \lambda$ and $\lambda_2(G) = \mu$, then $\lambda_2(X \circledS G) \leq (1-\mu)\lambda^2 + \mu$.*

Proof. Let A be the normalized adjacency matrix of X, and B be the normalized adjacency matrix of G. We can view a random step in $X \circledS G$ as

$$v \to (v,a) \to (v[a],a) \to (v[a],b) \to (v[a][b],b) \to v[a][b]$$

where a is a random node in G and b is a random neighbor of a in G.

We now define matrices that implement each of the above steps.

The first step is given by the "lift" matrix $L := I_n \otimes (1/k, \dots, 1/k)^T$ which is $n \cdot k \times n$.

The second step is given by $\tilde{A}$ which is a $n \cdot k \times n \cdot k$ matrix, where $\tilde{A}_{(u,a),(u',a')} = 1$ if and only if $a = a'$ and $u' = u[a]$. This matrix corresponds to taking a step in X *after the choice for the step has been made*. No entropy is added by $\tilde{A}$, which is a permutation matrix.

The third step is given by $\tilde{B} = I_n \otimes B$.

The fourth step is $\tilde{A}$ again.

Finally, the fifth step is given by the "projection" matrix $P := I_n \otimes (1, \dots, 1)$.

The adjacency matrix M of $X \circledS G$ satisfies

$$M = P\tilde{A}\tilde{B}\tilde{A}L.$$

By lemma 8, $B = (1-\mu) J_k + \mu C$ where $\|Cv\| \leq \|v\|$ for all v. It follows that

$$\tilde{B} = I_n \otimes B = (1-\mu) I_n \otimes J_k + \mu I_n \otimes C.$$

Plugging this into the expression for M one gets

$$M = (1-\mu) P\tilde{A} I_n \otimes J_k \tilde{A} L + \mu P \tilde{A} I_n \otimes C \tilde{A} L.$$

One can now observe the following:
(1) $I_n \otimes J_k = L \cdot P$, which is easy to check;
(2) $P \cdot \tilde{A} \cdot L = A$;
and (3) the matrix $D := P\tilde{A}I_n \otimes C\tilde{A}L$ satisfies that $\|Dv\| \leq \|v\|$ for every v. This can be shown also using the fact that C satisfies this property as we saw before. Therefore:

$$\begin{aligned} M &= (1-\mu)\, P\tilde{A} \cdot LP \cdot \tilde{A}L + \mu D \\ &= (1-\mu)\, A^2 + \mu D. \end{aligned}$$

Then, by lemma 3

$$\lambda_2\left(X \circledS G\right) = \max_{v \perp u} \frac{\|Mv\|}{\|v\|} \leq \frac{\|\left(1-\mu\right)A^2 v\|}{\|v\|} + \frac{\|\mu D v\|}{\|v\|} \leq \left(1-\mu\right)\lambda^2 + \mu.$$

And this completes the proof of this theorem. □

We can verify that if $\lambda_2\left(G\right) = \frac{1}{100}$ and $\lambda_2\left(X\right) = 1-\gamma \geq 1/10$, then $\lambda_2\left(X \circledS G\right) \leq 1 - \frac{12}{11}\gamma$. So repeating this operation $O(\log n)$ times will still give $\lambda_2 = 1/10$, qualitatively the same as graph squaring.

We will use a family of *expander graphs* for G, graphs on arbitrarily many nodes that have bounded degree and bounded second largest eigenvalue.

Fact 1 (Explicit expander graphs)**.** *For some constant $Q = 4^q$, $\forall m$ $\exists$ a Q-regular graph G_m on Q^m nodes with $\lambda_2 \leq \frac{1}{100}$. Given a node and an edge index we can compute the corresponding neighbor in logarithmic space.*

Many such constructions are available. One is due to Margulis (the needed expansion property was proved later). Here the vertex set of a graph G on N nodes is $Z_{\sqrt{N}} \times Z_{\sqrt{N}}$, where $Z_{\sqrt{N}}$ is the ring of the integers modulo $\sqrt{N}$. Each vertex v is a pair $v = (x, y)$ where $x, y \in Z_{\sqrt{N}}$. For matrices T_1, T_2 and vectors b_1, b_2 defined below, each vertex $v \in G_N$ is connected to $T_1 v, T_1 v + b_1, T_2 v, T_2 v + b_2$ and the four inverses of these operations. It can be shown that for the choices $T_1 := \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix}$, $T_2 := \begin{pmatrix} 1 & 0 \\ 1 & 1 \end{pmatrix}$, $b_1 := \begin{pmatrix} 1 \\ 0 \end{pmatrix}$ and $b_2 := \begin{pmatrix} 0 \\ 1 \end{pmatrix}$ the resulting graph is 8-regular with bounded second eigenvalue. By adding self-loops and taking powers of this graph one achieves the desired parameters.

Theorem 11. *The undirected reachability problem is solvable deterministically in $O(\log n)$-space.*

Proof. By Claim 2, we can focus on 4-regular graph with sets S and T. Call this 4-regular graph X. To make the degrees match, let $X_1 := X^q$. Then the degree of X_1 is $Q = 4^q$. Define

$$X_{i+1} = X_i \circledS G_i.$$

Note the degree of X_i is equal to the number of nodes in G_i which is equal to Q^i: every derandomized graph squaring increases the degree by a factor Q. Consider X_l for a suitable

$l = O(\log n)$. By Theorem 10 (cf. the observation right after its proof) we have $\lambda_2(X_l) \leq 1/10$. Then we can apply Claim 3, to solve the undirected reachability problem by going through all $s \in S$ and check if one of its neighbors in X_l lies in T, which can be done in $O(\log n)$-space. It only remains to verify that we can compute neighbors in X_l in $O(\log n)$-space.

The intuition is: if $v \in X_1$, then the neighbors are $v[a]$ where $a \in \{1, 2, \ldots, Q\}$; if $v \in X_2$, then the neighbors are $v[(a, b)] = v[a][a[b]]$ where $a, b \in \{1, 2, \ldots, Q\}$; if $v \in X_3$, then the neighbors are $v[(a, b, c)] = v[(a, b)][(a, b)[c]] = v[a][a[b]][(a', b')] = v[a][a[b]][a'][a'[b']]$ where $a, b, c \in \{1, 2, \ldots, Q\}$; etc.

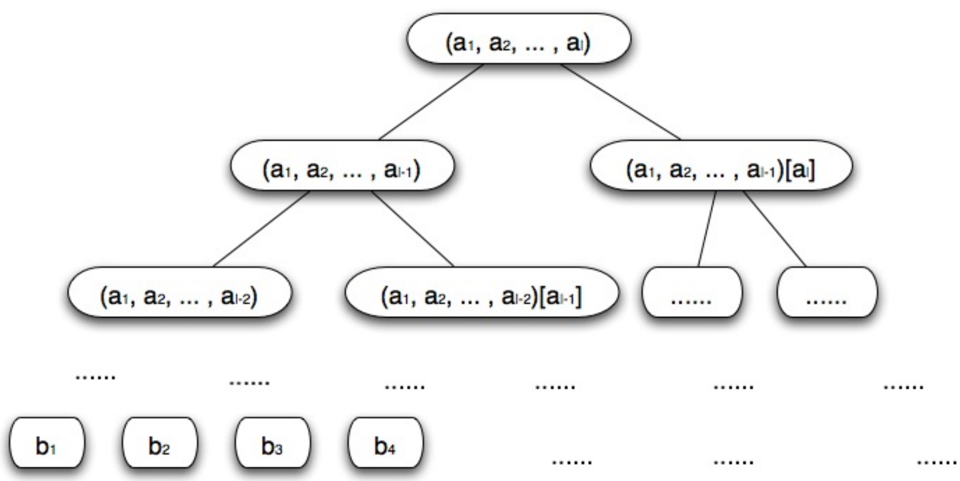

Figure 5: Compute b_i's from a_i's.

An edge in X_l is specified by $(a_1, a_2, \ldots, a_l)$ where $a_i \in \{1, 2, \ldots, Q\}$. To this edge there corresponds a path of length 2^{l-1} in the graph X with labels $(b_1, b_2, \ldots, b_{2^{l-1}})$ where $b_i \in \{1, 2, \ldots, Q\}$. The associated neighbor of v in X_l is $v[b_1][b_2] \ldots [b_{2^{l-1}}]$. So to compute $v[(a_1, \ldots, a_l)]$, we proceed in 2 steps:

1. compute $(b_1, \ldots, b_{2^{l-1}})$, and
2. compute $v[b_1][b_2] \ldots [b_{2^{l-1}}]$.

We must do this in space $l = O(\log n)$, and so we cannot afford to write down the output of the first step. Instead, the following shows given $(a_1, \ldots, a_l)$ and an index $i \leq 2^{l-1}$ how to compute $b_i \in \{1, 2, \ldots, Q\}$ in space $O(\log n)$. From this, one can easily compute $v[b_1][b_2] \ldots [b_{2^{l-1}}]$ in logarithmic space one step at the time.

Observe that the the indices b_i are obtained from the indices a_i as in Figure 5.

So to compute b_i, we just need to go from the root to the leaf b_i in the tree. The space needed for this is just the name of the node in the tree, which takes $O(l) = O(\log n)$ bits, plus the space needed to compute neighbors in the expander graphs which is also $O(\log n)$. So the total space is $O(\log n)$ and this completes the proof of the theorem. □

4 Notes

That undirected reachability is decidable deterministically using logarithmic space was first proved by Reingold. Reingold's proof is in the same spirit of the one we presented, but uses different graph operations to reduce the eigenvalue bound. The proof that we presented is due to Rozenman and Vadhan.