
Notes for Lecture 8

1 Undirected Connectivity

In the undirected $s - t$ connectivity problem (abbreviated ST-UCONN) we are given an undirected graph $G = (V, E)$ and two vertices $s, t \in V$, and the question is whether that is a path between s and t in G .

ST-UCONN is complete for the class **SL** of decision problems that are solvable by *symmetric* non-deterministic machines that use $O(\log n)$ space. A non-deterministic machine is symmetric if whenever it can make a transition from a global state s to a global state s' then the transition from s' to s is also possible. The proof of **SL**-completeness of ST-UCONN is identical to the proof of **NL**-completeness of ST-CONN except for the additional observation that the transition graph of a symmetric machine is undirected.

Reingold [Rei05] has given a log-space algorithm for ST-UCONN, thus showing that **SL** = **L**. This was a great breakthrough, and we will spend a good amount of time studying it and its context.

2 Randomized Log-space

We now wish to introduce *randomized* space-bounded Turing machine. For simplicity, we will only introduce randomized machines for solving decision problems. In addition to a read-only input tape and a read/write work tape, such machines also have a read-only *random* tape to which they have *one-way* access, meaning that the head on that tape can only move, say, left-to-right. For every fixed input and fixed content of the random tape, the machine is completely deterministic, and either accepts or rejects. For a Turing machine M , an input x and a content r of the random tape, we denote by $M(r, x)$ the outcome of the computation.

We say that a decision problem L belongs to the class **RL** (for *randomized* log-space) if there is a probabilistic Turing machine M that uses $O(\log n)$ space on inputs of length n and such that

- For every content of the random tape and for every input x , M halts.
- For every $x \in L$, $\Pr_r[M(r, x) \text{ accepts}] \geq 1/2$
- For every $x \notin L$, $\Pr_r[M(r, x) \text{ accepts}] = 0$.

Notice that the first property implies that M always runs in polynomial time. It is easy to observe that any constant bigger than 0 and smaller than 1 could be equivalently used instead of 1/2 in the definition above. It also follows from the definition that **L** \subseteq **RL** \subseteq **NL**.

The following result shows that, indeed, **L** \subseteq **SL** \subseteq **RL** \subseteq **NL**.

Theorem 1 *The problem ST-UCONN is in RL.*

The algorithm is very simple. Given an undirected graph $G = (V, E)$ and two vertices s, t , it performs a random walk of length $100 \cdot n^3$ starting from s . If t is never reached, the algorithm rejects.

```

input:  $G = (V, E), s, t$ 

 $v \leftarrow s$ 
for  $i \leftarrow 1$  to  $100 \cdot |E| \cdot |V|$ 
    pick at random a neighbor  $w$  of  $v$ 
    if  $w = t$  then halt and accept
     $v \leftarrow w$  reject

```

The analysis of the algorithm is based on the fact that if we start a random walk from a vertex s of an undirected vertex G , then each vertex in the connected component of s is likely to be visited at least once after $O(|V| \cdot |E|)$ steps. As we develop spectral graph theory in the next few lectures, we will see the proof of the weaker bound $O(|E|^2)$.

3 Eigenvalues and expanders

We now embark on the study of graph expansion and algebraic graph theory. Within the next lecture or two we will: (i) know about the equivalence of edge expansion and eigenvalue gap, (ii) understand spectral partitioning, (iii) know how to efficiently construct a family of bounded-degree expanders. We will then return to the question of the space complexity of ST-UCONN and see how Reingold's algorithm works by reducing ST-UCONN in arbitrary graphs to ST-UCONN in bounded-degree expanders, the latter problem having an easy log-space solutions.

We begin with the definition of (normalized) edge expansion. All graphs that we will talk about from now on will be regular.

Definition 1 (Edge-expansion of a graph) Let $G = (V, E)$ be a d -regular, then we define the normalized edge expansion of G as

$$h(G) := \min_{|S| \leq |V|/2} \frac{\text{edges}(S, V - S)}{d|S|}$$

In what follows, we consider $G = (V, E)$ to be a given d -regular graph and $A \in \mathbb{R}^{V \times V}$ its adjacency matrix, that is

$$A(u, v) := \text{number of edges between } u \text{ and } v \tag{1}$$

We denote by $M := \frac{1}{d}A$ the *random walk transition matrix* of G . The intuition for this definition is that if we take a one-step random walk starting at vertex u , then $M(u, v)$ is the probability that we reach vertex v .

Definition 2 (Eigenvalues and eigenvectors) If $M \in \mathbb{C}^{n \times n}$, $\lambda \in \mathbb{C}$, $x \in \mathbb{C}^n$ and $xM = \lambda x$ then λ is an eigenvalue of M and x is an eigenvector of M for the eigenvalue λ .

Example 1 Let M be the transition matrix of a regular graph. Then $(1, 1, \dots, 1) \cdot M = (1, 1, \dots, 1)$. Therefore, the vector $(1, 1, \dots, 1)$ is an eigenvector of M with corresponding eigenvalue 1.

Generally, $xM = \lambda x \Rightarrow x(M - \lambda I) = 0 \Rightarrow \det(M - \lambda I) = 0$. $\det(M - \lambda I)$ is a polynomial in λ over \mathbb{C} of degree n , and it has n roots (counting multiplicities). Therefore, λ is an eigenvalue of M iff it is a root of $\det(M - \lambda I)$ and so, counting multiplicities, M has n eigenvalues.

Theorem 2 If $M \in \mathbb{R}^{n \times n}$ is symmetric then the following properties hold:

1. all n eigenvalues $\lambda_1, \dots, \lambda_n$ are real
2. one can find an orthogonal set of eigenvectors x_1, \dots, x_n such that x_i has corresponding eigenvalue λ_i and $x_i \perp x_j$ for $i \neq j$.

We note that a multiple of an eigenvector is also an eigenvector and therefore we can assume w.l.o.g. that all the x_i have length one.

From now on we fix the convention that if we denote by $\lambda_1, \dots, \lambda_n$ the eigenvalues of M , then $\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_n$.

There are several equivalent characterizations of the eigenvalues of M ; the following one will be useful.

Lemma 3 Let $M \in \mathbb{R}^{n \times n}$ be symmetric. Then

$$\lambda_1 = \max_{x \in \mathbb{R}^n, \|x\|=1} xMx^T = \max_{x \in \mathbb{R}^n} \frac{xMx^T}{xx^T} \quad (2)$$

PROOF:

- (a) Assume $\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_n$. Then $x_1 M x_1^T = \lambda_1 x_1 x_1^T = \lambda_1$ therefore, $\max_{x \in \mathbb{R}^n, \|x\|=1} \{xMx^T\} \geq \lambda_1$.
- (b) Conversely, let x be any vector of length one, $x \in \mathbb{R}^n$, $\|x\| = 1$. Let $x = a_1 x_1 + a_2 x_2 + \dots + a_n x_n$.

$$\begin{aligned} xMx^T &= \sum_{i,j} x(i)x(j)M(i,j) = \left(\sum_i a_i x_i\right) M \left(\sum_i a_i x_i\right)^T = \\ &= \left(\sum_i \lambda_i a_i x_i\right) \left(\sum_j a_j x_j\right)^T = \sum_i \lambda_i a_i^2 \leq \max_i \lambda_i \sum_i a_i^2 = \lambda_1 \end{aligned}$$

Therefore $\max_{x \in \mathbb{R}^n, \|x\|=1} \{xMx^T\} \leq \lambda_1$.

□

We can also prove that

$$\lambda_2 = \max_{x \in \mathbb{R}^n, x \perp x_1} \frac{xMx^T}{xx^T} \quad (3)$$

For (a) use $x = x_2$, and conclude $\max_{x \in \mathbb{R}^n, x \perp x_1} \frac{xMx^T}{xx^T} \geq \lambda_2$.

For (b) take any $x \in \mathbb{R}^n, x \perp x_1$.

Let $x = a_2x_2 + \dots + a_nx_n$. Then

$$\frac{xMx^T}{xx^T} = \frac{\sum_{i=2}^n \lambda_i a_i^2}{\sum_{i=2}^n a_i^2} \leq \lambda_2 .$$

A similar argument shows that

$$\max\{|\lambda_2|, \dots, |\lambda_n|\} = \max_{x \perp x_1} \frac{|xMx^T|}{xx^T} \quad (4)$$

We now know enough to characterize the largest eigenvalue of the adjacency matrix of a regular graph.

Theorem 4 *Let G be a regular graph and M its adjacency matrix. Let $\lambda_1 \geq \dots \geq \lambda_n$ be its eigenvalues. Then $\lambda_1 = 1$.*

PROOF: Trivially, $\lambda_1 \geq 1$ because 1 is an eigenvalue.

Let $x \in \mathbb{R}^n, \|x\| = 1, xM = \lambda_1 x$

$$\begin{aligned} 0 &\leq \sum_{u,v} M(u,v)(x(u) - x(v))^2 = 2 \sum_v x(v)^2 - 2 \sum_{u,v} x(u)x(v)M(u,v) \\ &= 2xx^T - 2xMx^T = 2 - 2\lambda_1 \Rightarrow 1 \geq \lambda_1 \end{aligned}$$

Since $d \leq \lambda_1$ and $d \geq \lambda_1$ it follows $d = \lambda_1$. □

What about λ_2 ? An important theme in this theory is that the difference $1 - \lambda_2$ characterizes the expansion of the graph. The following is a simple special case which is a good warm-up example.

Claim 5 *Let G be a regular graph and M its adjacency matrix. Let $\lambda_1 \geq \dots \geq \lambda_n$ its eigenvalues. Then $\lambda_2 = 1$ if and only if the graph is disconnected.*

PROOF: Choose $x_1 = \frac{1}{\sqrt{n}}(1, 1, \dots, 1)$ and x_2 another eigenvector orthogonal to x_1 . x_2 should be $(x_2(1), \dots, x_2(n))$ with $\sum_i x_2(i) = 0$. Therefore, some entries should be positive and some others should be negative; in particular, the entries of x_2 are not all equal.

$$0 \leq \sum_{u,v} M(u,v)(x_2(u) - x_2(v))^2 = 2 - 2\lambda_2 = 0$$

Therefore, for x_2 any two adjacent vertices must have identical labels and, since the labels are not all equal, the graph has to be disconnected.

Conversely, if the graph is disconnected then let S and $V - S$ be a partition of the graph that is crossed by no edge. Let $p = \frac{|S|}{|V|}$, $q = \frac{|V-S|}{|V|}$. Assign

$$x(v) = \begin{cases} q & \text{if } v \in S \\ -p & \text{if } v \notin S \end{cases}$$

First, observe that $x \perp (1, 1, \dots, 1)$ since $\sum_v x(v) = q \cdot |S| - p \cdot |V - S| = qpn - pqn = 0$.

Second, look at $xM = \underbrace{(dq, dq, \dots, dq)}_{|S|}, \underbrace{(-pd, -pd, \dots, -pd)}_{|V-S|} = dx$.

Therefore, if the graph is disconnected we have $\lambda_2 = 1$. \square

References

The definition of **SL** is due to Lewis and Papadimitriou [LP82].

Prior to Reingold's algorithm [Rei05], an algorithm for ST-UCONN was known running in polynomial time and $O(\log^2 n)$ space (but the polynomial had very high degree), due to Nisan [Nis94]. There was also an algorithm that has $O(\log^{4/3} n)$ space complexity and super-polynomial time complexity, due to Armoni, Ta-Shma, Nisan and Wigderson [ATSWZ00], improving on a previous algorithm by Nisan, Szemeredy and Wigderson [NSW92].

The best known deterministic simulation of **RL** uses $O((\log n)^{3/2})$ space, and is due to Saks and Zhou [SZ95].

References

- [ATSWZ00] Roy Armoni, Amnon Ta-Shma, Avi Wigderson, and Shiyu Zhou. An $o(\log(n)^{4/3})$ space algorithm for (s, t) connectivity in undirected graphs. *Journal of the ACM*, 47(2):294–311, 2000. 5
- [LP82] Harry R. Lewis and Christos H. Papadimitriou. Symmetric space-bounded computation. *Theoretical Computer Science*, 19:161–187, 1982. 5
- [Nis94] N. Nisan. $RL \subseteq SC$. *Computational Complexity*, 4(1), 1994. 5
- [NSW92] N. Nisan, E. Szemeredy, and A. Wigderson. Undirected connectivity in $O(\log^{1.5} n)$ space. In *Proceedings of the 33rd IEEE Symposium on Foundations of Computer Science*, pages 24–29, 1992. 5
- [Rei05] Omer Reingold. Undirected ST-connectivity in log-space. In *Proceedings of the 37th ACM Symposium on Theory of Computing*, pages 376–385, 2005. 1, 5
- [SZ95] M. Saks and S. Zhou. $RSPACE(S) \subseteq DSPACE(S^{3/2})$. In *Proceedings of the 36th IEEE Symposium on Foundations of Computer Science*, pages 344–353, 1995. 5