6.895 Randomness and Computation	April 24, 2006

Lecture 18

Lecturer: Ronitt Rubinfeld *Scribe: Elena Grigorescu*

In this lecture we present the relatively recent fundamental result of O. Reingold [3] which establishes that undirected st-connectivity can be decided in deterministic logarithmic space.
Given an undirected graph G, two vertices s and t in G, USTCON(G, s, t) is the problem of deciding whether there exists a path in G connecting s and t. The more difficult, directed graph version of this problem is known to be NL-complete, and thus in L^2 by Savitch's theorem. Prior to the Reingold result, USTCON was known to be in RL [1], and later ([2]) in $\log^{4/3}$. Also, USTCON is a complete problem for the mysterious class SL (symmetric, non-deterministic, log space computations), and therefore USTCON$\in$ L implies SL=L.

We will start by introducing some notation and recalling some results from previous lectures/homeworks.

Definition 1 *A (N, D, λ)-graph is an undirected graph on N vertices, of degree D, and with the second largest eigenvalue bounded above by λ.*

The relationship between the second eigenvalue and that of vertex expansion is given in the following proposition.

Proposition 2 *Let $\lambda < 1$. Then $\exists \varepsilon > 0$ s.t. for any (N, D, λ)-graph G, $\forall S \subset V$, s.t. $|S| < \frac{N}{2}$ we have that $|N(S)| \geq (1 + \varepsilon)|S|$. In this case we say that G is an expander. Note: Elements of S may be in $N(S)$*

Corollary 3 *If $\lambda < 1$, then for every connected graph (N, D, λ), there exists a path of length $O(\log N)$ between any vertices s and t. In particular, G has $O(\log N)$ diameter.*

Proof By the vertex expansion property given in Proposition 2, it follows that starting at s and successively following the neighboring sets for $l = O(\log N)$ steps we must have covered $> N/2$ of the vertices. Repeating the process from vertex t, it must be the case that there is a vertex reached from both s and t in these l steps. ■

For and expander graph G of constant degree, one can now easily check st-connectivity in $\log N$ space as follows:

- Enumerate all D^l paths starting at s of length $l = O(\log N)$.
- If have reached node t then output 'connected', else output 'not connected'.

Notice that the space requirement of the algorithm is $O(\log N \log D$, which is $O(\log N)$ for constant D. We will reduce the problem of checking st-connectivity in a general graph, to that of checking st- connectivity in an expander graph with constant degree D.

Observation: The results shown in this lecture do not apply to general directed graphs. However, for the special case of directed graphs with constant out-degree equal to the in-degree at each vertex, similar results do hold.

A key fact that will be relevant to the main proof states that the spectral gap $(1 - \lambda(G))$ of any connected, non-bipartite graph G is large enough, namely at least inverse polynomial in $|G|$, as described next.

Theorem 4 *For any connected D-regular and non-bipartite graph G the following holds*

$$\lambda(G) \leq 1 - \frac{1}{DN^2}.$$

We further add to our toolkit a useful graph.

Theorem 5 *There exist a constant D_e and a $(D_e^{16}, D_e, \frac{1}{2})$-graph.*

Also, recall that the graph powering operation is a simple method of increasing the connectivity of a graph.

Proposition 6 *(Graph powering) If G is a (N, D, λ)-graph, then the power G^t of G is a (N, D^t, λ^t)-graph.*

Observation: Notice that graph powering increases a lot the degree of the new graph. In particular, for a good enough λ we will want $t = O(\log N)$, while we only aim for $t = constant$. Reingold's proof does use manipulations of graphs using this operation. However, his main ingredient is a new operation (the **zig-zag product**) that will bring down the degree of the graph, while increasing λ by only a small constant factor.

A simple example of a way to lower the degree of a graph is shown in the picture below, where a node is substituted by a cycle.

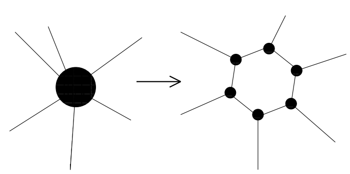

We next introduce two graph products that have the property that they reduce the degree of a graph while not increasing λ by too much either.

Replacement product:

Given graph an (N, D) graph G and a (D, d) graph H, the replacement product graph G' is obtained from G and H in the following way:

- Each node $v \in G$ is replaced by a copy of H, called H_v. Thus, there are N D nodes in G'.
- Each node $v' \in H_v$ corresponds to an edge from $v \in G$. Therefore, $v' \in H_v$ is adjacent to the d nodes in H_v and to a vertex in H_z, where $(v, z) \in E(G)$ and v' corresponds to z. Therefore, v' has degree $d + 1 << D$.

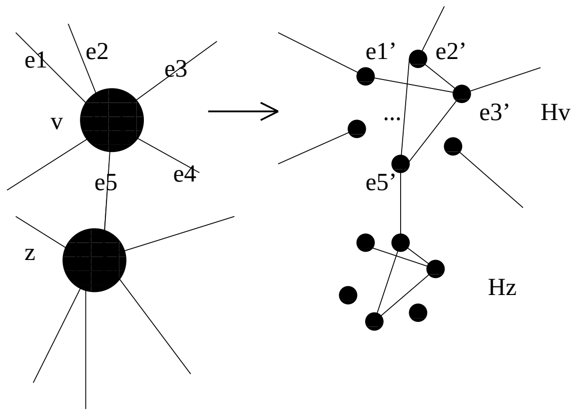

Zig-zag product:
The zig-zag product G'' of G and H (notation GzH), where G is D-regular and has N nodes and H is d-regular and has D nodes is constructed from the replacement product G' of G and H:

- $V(G'') = V(G')$.
- The edges of G'' are between vertices $x, w \in G'$ that are connected by a path of length three, that starts in $x \in H_v$, takes a step $(x, y) \in E(H_v)$, then follows the edge $(y, z) \in G$, where $z \in H_w$, and finally, takes a step $(z, w) \in E(H_w)$. Therefore, the degree of each vertex is d^2.

Next we state one of the main tools of the proof, namely the relationship between the second eigenvalues of the graphs G and H to the second eigenvalue of the zig-zag product of G and H.

Theorem 7 *([4]) Let G be a (N, D, λ)-graph, and H be a (D, d, α)-graph. Then GzH is a $(ND, d^2, f(\lambda, \alpha))$-graph, where $f(\lambda, \alpha) = \frac{1}{2}(1 - \alpha^2)\lambda + \frac{1}{2}\sqrt{(1 - \alpha^2)^2\lambda^2 + 4\alpha^2}$.*

In fact, a nicer form of this theorem will be enough for our purposes. The following corollary shows that the spectral gap of the zig-zag product is only larger by a small factor from the spectral gap of G.

Corollary 8 *For G, H as above,*

$$\frac{1}{2}(1 - \alpha^2)(1 - \lambda) \leq 1 - \lambda(GzH).$$

Main Transformation:
Given: G that is D^{16}-regular on N vertices,
and H that is D-regular on D^{16} vertices (given by Theorem 5),
Let l be the smallest integer s.t. $(1-\frac{1}{DN^2})^{2^l} < \frac{1}{2}$.
Let $G_0 \leftarrow G$
Let $G_i \leftarrow (G_{i-1} z H)^8$
Output: $\tau(G, H) = G_l$.

We are now ready for the final construction of the expander graph that we are after.

Observation: The number of nodes in G_l is $N(D^{16})^l = \text{poly}(N)$, since $l = O(\log N)$.

Lemma 9 *If H is an expander, then G_l is an expander. Equivalently, if $\lambda(H) \leq \frac{1}{2}$ and G is connected and not bipartite, then $\lambda(\tau(G, H)) \leq \frac{1}{2}$.*

Proof [Sketch] By Theorem 4 we have that $\lambda(G_0) \leq 1 - \frac{1}{DN^2}$. Notice that it is enough to show inductively that $\lambda(G_i) \leq \max\{\lambda(G_{i-1})^2, \frac{1}{2})\}$. This will imply that, for $l = O(\log(N)\log(D))$ we have that $\lambda(G_l) \leq \max\{\lambda(G_0)^{2^l}, \frac{1}{2}\} < \frac{1}{2}$. To prove the induction hypothesis, notice that $\lambda(H) \leq \frac{1}{2}$. By Corollary 8, we have that $\lambda(G_{i-1} z H) < 1 - 1/3\ (\lambda(G_{i-1})$. Therefore, $\lambda(G_i) < (1 - 1/3\ (\lambda(G_{i-1}))^8$. The conclusion follows then by elementary calculations, by considering the cases $\lambda(G_{i-1}) < 1/2$, and $\lambda(G_{i-1}) \geq 1/2$ ■

We next present an overview of the actual $\log N$ space algorithm solving st-connectivity.

1. Preprocessing stage: make the graph G D^{16} regular, preserving non-bipartiteness, and the connected components. This can be done by the replacement product of G with an N-cycle and then adding self-loops. Note that the number of nodes becomes N^2 but this operation is performed only once. Let G_e be the resulting graph.
2. Use the transformation τ described before on the preprocessed graph G_e and H the expander given by Proposition 5.
3. Run the expander algorithm on $\tau(G_e, H)$.

The only tricky part that one needs to verify now is how the walks are performed in log space. The choice of representing graphs G and H as so-called *rotation maps* turns out to be fortunate. This representation implies that each edge is labeled at both its endpoints, which gives a way of tracing a path back from any position by only remembering a constant number of labels.

Definition 10 *[3] For a D-regular undirected graph G, the* **rotation map** $Rot_G : [N] \times [D] \to [N] \times [D]$ *is defined as $Rot_G(v, i) = (w, j)$ if the i' th edge incident to v leads to w, and this edge is also the j'th edge incident to w.*

It follows easily that given the rotation map of G one can compute in $\log N$ space the rotation map of G_e. The heart of the problem is showing that the rotation map of $\tau(G_e, H)$ is computable in log space given the rotation maps of G and H. The proof uses an inductive argument and elementary techniques, and we do not attempt it here.

References

[1] Aleliunas, Karp, Lipton, Lovasz, Rackoff. Random walks, universal traversal sequences, and the complexity of maze problems. it Annual Symposium on Foundations of Computer Science. 1979.

[2] Armoni, Ta-Shma, Wigderson, Zhou. An $o(log(n)^{4/3})$ space algorithm for st-connectivity in undirected graphs. JACM, 2000.

[3] O. Reingold. Undirected st-connectivity in Log Space. 2004.

[4] O. Reingold, S. Vadhan, A. Wigderson. Entropy waves, the zig-zag graph product, and new constant-degree expanders. FOCS 2000.