CS 880: Pseudorandomness and Derandomization 3/6/2013

Lecture 13: Derandomized Squaring

Instructors: Holger Dell and Dieter van Melkebeek Scribe: Alexi Brooks

DRAFT

In this lecture, we introduce a new graph operation, the *derandomized squaring*, which affects the second-largest eigenvalue of a graph in a similar way as squaring and which is based on the replacement product. The advantage that derandomized squaring has over squaring is that the degree does not go up by too much. We can use this tool to show that the problem of Undirected Connectivity in deterministic Logspace.

As explained in the last lecture, we consider edge-colored graphs and we represent them formally as a *neighbor function* $\Gamma_G : V(G) \times [D] \to V(G)$. To describe the neighbor $v = \Gamma_G(u, \ell)$ that we reach from u if we follow the edge $\{u, v\}$ that has color ℓ under the coloring c. In this lecture, we will not require that the induced edge-coloring is proper or that the graph is undirected. Instead, we will consider "consistently labeled" graphs G, which are D-regular directed graphs with edge colors from $[D]$ such that any two outgoing edges have distinct colors and any two incoming edges have distinct colors at every vertex. We make this formal in the following definition.

Definition 1. *The edge-colored directed graph G represented by the neighbor function*

$$\Gamma_G : V(G) \times [D] \to V(G)$$

is called consistently labeled *if, for all $u, u' \in V(G)$ with $u \neq u'$ and all $\ell, \ell' \in [D]$ with $\Gamma_G(u, \ell) = \Gamma_G(u', \ell')$, we have $\ell \neq \ell'$.*

1 Replacement Product

We defined the replacement product $G \textcircled{r} H$ in the previous lecture for a D-regular N-vertex graph G and a d-regular D-vertex graph H: For every vertex u of G, we create a distinguished copy of H in $G \textcircled{r} H$. This copy is the "cloud" of v in $G \textcircled{r} H$. Each vertex in the cloud has a label from the set $V(H) = [D]$. Within a cloud, adjacencies correspond to adjacencies in H. Across clouds, adjacencies exist as they do in G; if we are at node (u, ℓ) in $G \textcircled{r} H$, then we have d edges to the vertex $(\Gamma_G(u; \ell), \ell)$. and for each edge we choose the vertices within the clouds that correspond to its edge label.

Transition Matrices in the Replacement Product. We recall a number of transition matrices related to the replacement product $G \textcircled{r} H$ that will be convenient in defining and analyzing the derandomized squaring:

1. A is the transition matrix of G
2. B is the transition matrix of H

3. $\hat{A}$ is the transition matrix corresponding to edges *between* the clouds of $G \textcircled{r} H$, that is, we have $\hat{A}_{(v,\ell),(u,\ell)} = 1$ if $\Gamma_G(u,\ell) = v$ and $\hat{A}_{(v,\ell'),(u,\ell)} = 0$, otherwise.

4. $\hat{B} = I_N \otimes B$ is the transition matrix for walking *within* any cloud of $G \textcircled{r} H$, that is, $\hat{B}_{(u,\ell'),(u,\ell)} = B_{\ell',\ell}$ and $\hat{B}_{(v,\ell'),(u,\ell)} = 0$ for all distinct $u, v \in V(G)$.

The transition matrix of the replacement product can be written as $M = \frac{1}{2}\hat{B} + \frac{1}{2}\hat{A}$.

2 Derandomized Squaring

Derandomized Squaring [RV05] is related to a certain random walk in the Replacement Product. To compute the derandomized square of a given D-regular directed graph G, we will thus again make use of a fixed constant-size expander H. In our construction, the number of vertices of H has to be equal to the degree D of G.

Conceptually, we generate the derandomized square $G \textcircled{s} H$ of G and H by "lifting" G onto $G \textcircled{r} H$, taking two steps between clouds, and then projecting back onto G. Since each step between clouds (a transition with respect to $\hat{A}$) is deterministic, we make sure we are located at any vertex in each cloud with uniform probability. We achieve this before the first across-cloud step by doing the lifting in such a way that we select a vertex in the corresponding cloud uniformly at random among the vertices in the cloud. We achieve randomness before the second across-cloud step by taking a step within the cloud, that is, according to $\hat{B}$. We exploit the fact that H is a good expander, so a single step in $\hat{B}$ should give us a reasonable approximation of a uniform distribution. Thus, the transition matrix $M(G \textcircled{s} H)$ of the derandomized square is defined as

$$M(G \textcircled{s} H) = P\hat{A}\hat{B}\hat{A}L\,.$$

Here L is the lifting operation and P is the corresponding projection. That is, L is the $ND \times N$ matrix and P is the $N \times ND$ matrix that satisfies

$$\begin{aligned}(Lx)_{(u,\ell)} &= x_u/D\,,\\ (Pz)_u &= \sum_{\ell} z_{(u,\ell)}\,,\end{aligned}$$

for all N-dimensional vectors x, all ND-dimensional vectors z, and all $u \in V(G)$ and $\ell \in V(H)$.

To understand this definition (cf. also Figure 2), recall that basic squaring generates the graph G^2 of degree D^2, and it connects vertices that can be reached in two steps. The degree of a graph is closely related to the amount of randomness needed to perform a random walk on it, and we therefore want it to be small. In the definition of $G \textcircled{s} H$, we only have two steps where randomness is needed — the two $\hat{A}$-steps and the projection step P are purely deterministic and do not contribute to the degree of $G \textcircled{s} H$. The steps that do require randomness are L and $\hat{B}$. The lifting L selects one of D vertices in a cloud uniformly at random, and the $\hat{B}$ steps select one of d neighbors uniformly at random. Thus, the degree of $G \textcircled{s} H$ is $D \cdot d$ as opposed to the degree $D \cdot D$ of G^2.

The transition matrix above suffices to do the eigenvalue analysis, but formally, we also need to specify the edge labels of the graph $G \textcircled{s} H$.

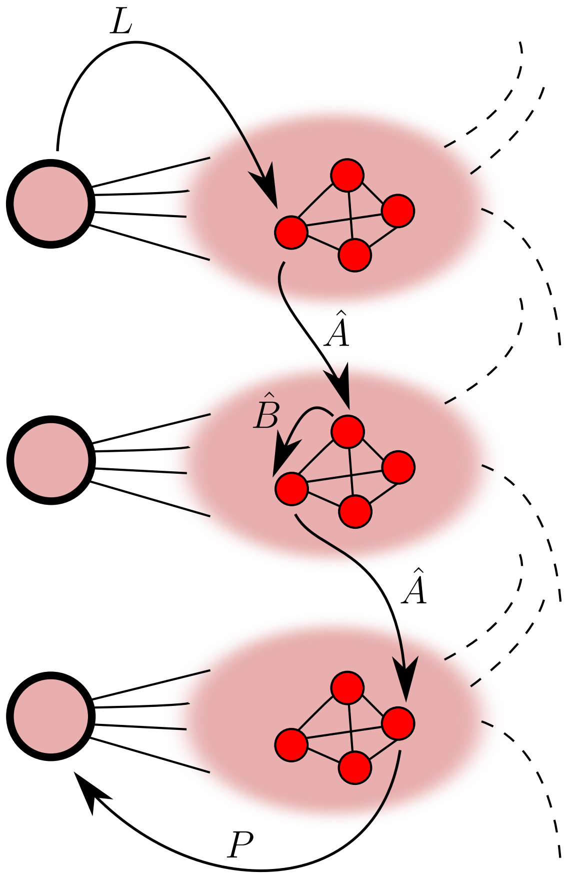

Figure 1: A pictorial representation of the five-step walk that makes up a single step in the derandomized square graph. The left side of the picture corresponds to the vertices of G and thus the vertices of $G \circledS H$. The right side of the picture corresponds to the graph $G \textcircled{r} H$; it is not actually part of the graph $G \circledS H$, but any edge in $G \circledS H$ can be though of as a five-step walk in the graph $G \dot{\cup} (G \textcircled{r} H)$: Starting at the top left vertex v, we first "lift" to $G \textcircled{r} H$ by selecting uniformly at random a vertex in the cloud H_v of v in $G \textcircled{r} H$. Next we take the deterministic step $\hat{A}$, which corresponds to switching the cloud, i.e., it is a step in G. Then we take a uniform step inside the second cloud, which corresponds to $\hat{B}$. We take another deterministic step $\hat{A}$ and arrive somewhere in a third cloud. Finally, we "project" back to the vertex of G whose cloud we just arrived at.

Definition 2. *The* derandomized square *of a degree-D graph G with respect to a degree-d graph H with vertex set $V(H) = [D]$ is the graph $G \circledS H$ whose vertex set is $V(G)$, whose degree is $D \cdot d$, and whose edges are defined by the neighbor function*

$$\begin{aligned} &\Gamma_{G \circledS H} : V(G) \times ([D] \times [d]) \to V(G) \\ &\Gamma_{G \circledS H}\big(u; (\ell, a)\big) = v, \textit{ where} \\ &\qquad (v, *) \doteq \Gamma_{G \circledR H}\Big(\Gamma_{G \circledR H}((u, \ell); d+1); a\Big) . \end{aligned}$$

In particular, if H is the complete graph, then $G \circledS H$ is equal to G^2. We now prove more generally that, if H is a good expander, then the second-largest eigenvalue of $G \circledS H$ is close to the one of G^2.

Lemma 1. *If G has second-largest eigenvalue λ and H has second-largest eigenvalue μ, then $G \circledS H$ has second-largest eigenvalue at most $(1-\mu)\lambda^2 + \mu$.*

Proof. The transition matrix of $G \circledS H$ is $M(G \circledS H) = P\hat{A}\hat{B}\hat{A}L$ as discussed above. By the matrix-decomposition view of second-largest eigenvalues, we can write $\hat{B} = I_N \otimes B = (1-\mu) \cdot I_N \otimes J_D + \mu \cdot I_N \otimes E$ for some error matrix E with $\|E\|_2 \leq 1$. We define $\tilde{J} = I_N \otimes J_D$ and $\tilde{E} = I_N \otimes E$. Then $\|\tilde{E}\|_2 \leq 1$ holds by the properties of the tensor product. The transition matrix $M = P\hat{A}\hat{B}\hat{A}L$ can then be decomposed into

$$M = (1-\mu) \cdot P\hat{A}\tilde{J}\hat{A}L + \mu \cdot P\hat{A}\tilde{E}\hat{A}L \,. \tag{1}$$

Note that $\tilde{J}$ corresponds to moving to uniformly at random vertex in the same cloud and the matrix product LP does the exact same thing, we have $\tilde{J} = LP$. Furthermore, note that $P\hat{A}L$ first moves to a uniform vertex in the cloud, crosses to a neighboring cloud deterministically, and then collapses the clouds to a single vertex. This is the same operation as one step in A, so we have $A = P\hat{A}L$. Then the first matrix product in (1) becomes $P\hat{A}LP\hat{A}L = (P\hat{A}L)^2 = A^2 = M(G)^2$. Then we can use the triangle inequality to bound $\lambda(M) \leq (1-\mu)\lambda(A^2) + \mu\lambda(P\hat{A}\tilde{E}\hat{A}L)$. Since $\lambda(A^2) = \lambda^2$ and $\lambda(P\hat{A}\tilde{E}\hat{A}L) \leq \|P\hat{A}\tilde{E}\hat{A}L\|_2 \leq \|P\|_2 \cdot \|\hat{A}\|_2 \cdot \|\tilde{E}\|_2 \cdot \|\hat{A}\|_2 \cdot \|L\|_2 = \sqrt{D} \cdot 1 \cdot \frac{1}{\sqrt{D}}$, we obtain the claimed inequality. ☐

In the next lecture, we will use the derandomized squaring operation to obtain a log-space algorithm for undirected connectivity.

References

[RV05] Eyal Rozenman and Salil Vadhan. Derandomized squaring of graphs. *Approximation, Randomization and Combinatorial Optimization. Algorithms and Techniques*, pages 611–611, 2005.