CS 388R: Randomized Algorithms Fall 2017

Lecture 21 — November 16, 2017

Prof. Eric Price *Scribe: Chenyang Wang*

NOTE: THESE NOTES HAVE NOT BEEN EDITED OR CHECKED FOR CORRECTNESS

1 Overview

In the last lecture we discussed the basics of network coding.

In this lecture we discuss the application of network coding to solving edge connectivity on a DAG, as well as introductory material on Markov chains and random walks.

2 Network Coding Analysis

Let s is the source vertex and t the destination vertex. Let $X_v \subset \mathbb{F}_q^{l'}$ be the vector space v knows about. If v is aware of $w \in \mathbb{F}_q^{l'}$, then there exists a basis vector $b \in X_v$ such that $\langle b, v\rangle = 0$. Let n be the number of vertices in the graph. Then after R rounds,

$$\mathbb{P}[t \text{ aware of } w] \geq 1 - (\frac{2e}{q})^{R|C_{s,t}|(1-\frac{n}{R})}$$

Then if $R \geq \frac{n}{\epsilon}$, and $q \geq 2^{1/\epsilon}$, then

$$\mathbb{P}[t \text{ aware of } w] \geq 1 - q^{-R|C_{s,t}|(1-O(\epsilon))}$$

If t is aware of all $w \in X_s$, then $X_t = X_s$. Therefore,

$$\mathbb{P}[X_t = X_s] \geq 1 - q^k q^{-R|C_{s,t}|(1-O(\epsilon))}$$

We want this probability to be greater than $1 - \frac{1}{q}$, which occurs when $k < R|C_{s,t}|(1 - O(\epsilon))$. If $l \geq \frac{k}{\epsilon}$, then this result will be optimal.

3 Edge Connectivity on a DAG

Our goal is to find the size of the minimum cut $|C_{s,t}|$ in $O(md^2)$ time, where m is the number of edges, and d is the max degree. The idea is to run one round of network coding on the DAG, then measure the throughput at t.

The algorithm is as follows: we let $X_s = \mathbb{F}_q^d$. Note that $d \geq |C_{s,t}|$. We then scan through the vertices of the DAG in topological sort order, and at each (u, v), we send $w \in X_v$, which is chosen uniformly at random. Finally, we compute the $\text{Dim}(X_t)$, which we claim is equal to $|C_{s,t}|$ w.h.p.

3.1 Analysis

To see that this is true, we first note that $\text{Dim}(X_t) \leq |C_{s,t}|$. Let S, S^c be the minimum $s - t$ cut in the graph, such that $s \in S$, $t \in S^c$. Then the number of unique vectors that are transmitted over the cut no greater than $|C_{s,t}|$, so the dimensionality of X_t can be at most $|C_{s,t}|$. With this in mind, it is sufficient to show that $\text{Dim}(X_t) \geq |C_{s,t}|$ w.h.p.

Next, we analyze $\mathbb{P}[t \text{ aware of } w]$ for all $w \in \mathbb{F}_q^d$. If t is aware of w, then that awareness must have been transmitted along some $s \to t$ path. The probability of awarness failing to pass along a given edge is $\frac{1}{q}$; then by the union bound, $\mathbb{P}[\text{awareness of } w \text{ passed along a given path}] \geq 1 - \frac{n}{q}$. Then the probability awareness passed on any path is

$$\begin{aligned}
\mathbb{P}[t \text{ aware of } w] &\geq 1 - (\frac{n}{q})^{|C_{s,t}|} \\
&= 1 - q^{-|C_{s,t}|(1-\log_q n)}
\end{aligned}$$

Let Z be the number of w that t is unaware of. Then $Z = q^{|X_t^{\perp}|} = q^{d-\dim(X_t)}$. From the above probability

$$\begin{aligned}
E[Z] &= q^d \, \mathbb{P}[t \text{ aware of } w] \\
&\geq q^{d-|C_{s,t}|(1-\log_q n)} \\
Z &\leq q^{d-|C_{s,t}|(1-\log_q n)+0.5} \text{ w.h.p} \\
\dim(X_t) &\geq |C_{s,t}|(1 - \log_q n) - 0.5
\end{aligned}$$

The second inquality is an application of Markov's inequality, and the third statement comes from the definition of Z. Let $q \geq n^{2d}$, and $\dim(X_t)$ will be exact.

3.2 Runtime

For each edge (u, v), the algorithm selects a $w \in X_u$, which takes $O(d^2)$ time, and orthogonalizes it with respect to X_v, also in $O(d^2)$ time. The total runtime is then $O(md^2)$. Using matrix operations, we can get this down to $O(md^{1.38})$.

4 Markov Chains

We now introduce the idea of Markov chains and random walks. Let n be the number of states, and P be the **transition matrix**, such that $P_{ij} = \mathbb{P}[i \to j]$ in the next round.

A **Markov chain** is a sequence $x_0, x_1, x_2, \ldots$ with $x_i \in [n]$, such that

$$\begin{aligned}\mathbb{P}[x_{t+1} = j | (x_1, x_2, \ldots, x_t)] &= \mathbb{P}[x_{t+1} = j | x_t = i] \\ &= P_{ij}\end{aligned}$$

At any time, t, the state of x_t can be expressed by a probability vector $q^{(t)} \in \mathbb{R}^n$, such that $\sum_i q_i^{(t)} = 1$. Note that $q^{(t+1)} = q^{(t)}P$.

A **stationary distribution** is a distrubution vector Π such that $\Pi = \Pi P$. In other words, Π is a eigenvector of P with eigenvalue 1.

The **hitting time** h_{ij}, is the expected number of steps required to reach state j from state i.

4.1 Fundamental Theorem of Markov Chains

A transition matrix P is **ergodic** if it satisfies the following:

- n is finite
- Irriducable: $\exists i \to j$ path for all $i, j \in [n]$.
- Aperiodic: $\forall$ states, gcd(loops) = 1.

An ergodic chain has the following characteristics:

- There exists a unique stationary distribution Π, $\Pi_i > 0$.
- All distributions q will eventually converge to Π.
- $h_{ii} = \Pi_i^{-1}$

4.2 Random Walks

Consider a random walk on an undirected, unweighted graph such that

$$P_{u,v} = \begin{cases} \frac{1}{d(u)} & \text{if } (u,v) \in E \\ 0 & \text{otherwise} \end{cases}$$

This random walk is ergodic iff the graph is connected (irriducible), and not bipartite (has an odd cycle $\to$ Aperiodic). If so, then $\Pi_v = \frac{d(v)}{2m}$, and $h_{vv} = \frac{2m}{d(v)}$.

4.3 More terms

Define $C_{u,v}$ as the expected commute time from u to v and back to u. In other words, $C_{u,v} = h_{u,v} + h_{v,u}$. Note that $h_{u,v}$ and $h_{v,u}$ are not necessarily equal; consider a lollipop graph.

Define $C_u(G)$ as the expected cover time starting at u; that is, the expected amount of time to visit all vertices in the graph.