CPSC 536N Randomized Algorithms (Winter 2014-15, Term 2)
Assignment 3

Due: Wednesday March 11th, in class.

Question 1: Streaming: improved dependence on δ

The algorithm of Lecture 14 estimates the ℓ_2-norm of the frequency vector with $(1+\epsilon)$-multiplicative error and failure probability δ. The space required is $O(\log(n)/\delta\epsilon^2)$ bits. In this problem, you must find a variant of this algorithm that uses only $O(\log(n)\log(1/\delta)/\epsilon^2)$ bits.

Use the following approach. Run the Lecture 14 algorithm with δ fixed to $1/4$, so that the space usage is only $O(\log(n)/\epsilon^2)$ bits. Now generate ℓ mutually independent estimates by running ℓ parallel copies of the algorithm. Combine those ℓ estimates using a trick from assignment 1.

Question 2: Variance of the ℓ_2-estimator

(a): Prove the claim in Lecture 14 about the variance of the estimate in the case $t = 1$.

Claim 1. Let $f \in \mathbb{R}^n$ be an arbitrary vector. Let L be a row vector of random signs that are 4-wise independent and with $\mathrm{E}[L] = 0$. Let $y = Lf$. Then

$$\mathrm{Var}\left[y^2\right] \;\leq\; \sum_{j_1,j_2,j_3,j_4\in[n]} \mathrm{E}\left[L_{j_1}L_{j_2}L_{j_3}L_{j_4}\right] f_{j_1}f_{j_2}f_{j_3}f_{j_4} \;\leq\; 3\,\|f\|_2^4\,.$$

(The left-hand inequality is already proven in Lecture 14.)

Hint: Consider each term in the sum separately. There are several cases. For example, what happens with the terms for which all $j_1, ..., j_4$ are distinct?

(b): OPTIONAL: Prove that actually $\mathrm{Var}\left[y^2\right] \leq 2\,\|f\|_2^4$.

(More on next page...)

Question 3: Sparsifiers

In this question, let us prove the following claim whose proof was omitted from Lecture 11. The proof is just an application of the Chernoff bound, but a bit fiddly.

Claim 2. Let $P \subseteq E_i$ be a projection of a cut. Then

$$\Pr\left[|w(P) - \mathrm{E}[w(P)]| > \frac{\epsilon \cdot \mathrm{sm}(P)}{\log n} \right] \leq 2 \exp\left(- \frac{\epsilon^2 \rho \cdot \mathrm{sm}(P)}{3 \cdot 2^i \log^2 n} \right)$$

To set up the proof, let $X_{j,e}$ be the random variable that is 1 if edge e is chosen during the j^{th} round of sampling. Then

$$w(P) = \sum_{j=1}^{\rho} \sum_{e \in P} \frac{k_e}{\rho} X_{j,e}.$$

We cannot directly apply the Chernoff bound to this sum because of the scaling factors k_e/ρ. But, with enough fiddling, the Chernoff bound can be applied.

Hints:

- The main properties that we need about each edge e are that $\Pr[X_{j,e} = 1] = 1/k_e$ and every $e \in P$ has $k_e \leq 2^i$.
- The main properties that we need about P are that $\mathrm{E}[w(P)] = |P|$ and $\mathrm{sm}(P) \geq |P|$.
- It may be useful to define the random variable $Y_{j,e} = (k_e/2^i) X_{j,e}$ and consider the sum $\sum_{j=1}^{\rho} \sum_{e \in P} Y_{j,e}$.
- When applying the Chernoff bound, you will need to separately handle the cases $\delta \leq 1$ and $\delta > 1$.