

Lecture 12: October 4

Instructor: Alistair Sinclair

Disclaimer: *These notes have not been subjected to the usual scrutiny accorded to formal publications. They may be distributed outside this class only with the permission of the Instructor.*

12.1 The Permanent

In this lecture we focus on a simple algorithm to approximate the permanent of a random matrix with 0,1 entries. We start with the definition of the permanent of a matrix.

Definition 12.1 *Let A be an $n \times n$ matrix such that each entry of A is either 0 or 1. The permanent of A , denoted by $\text{per}(A)$, is defined as*

$$\text{per}(A) = \sum_{\sigma} \prod_{i=1}^n a_{i,\sigma(i)},$$

where σ ranges over the permutations on $\{1, 2, \dots, n\}$.

Note the similarity between $\text{per}(A)$ and $\det(A)$, which is defined as the same sum but with each term weighted by $\text{sgn}(\sigma)$. However, while we can compute the determinant easily in polynomial time (say, by Gaussian elimination), computing the permanent is apparently very hard.

Computing the permanent of A is in fact equivalent to counting the number of *perfect matchings* in the bipartite graph G_A , which has n vertices on each side, and an edge connecting i to j iff $a_{ij} = 1$. It is an easy exercise to check that perfect matchings of G_A are in 1-1 correspondence with non-zero terms of $\text{per}(A)$.

While the problem of checking the existence of a perfect matching in a bipartite graph is easily solved in polynomial time by (e.g.) network flow techniques, counting the number of perfect matchings is $\#P$ -complete [Val79]. Hence computing $\text{per}(A)$ for a matrix A with 0,1 entries is also $\#P$ -complete, which means that (under standard complexity theoretic assumptions) it is not possible to obtain a polynomial time algorithm to compute the permanent. The focus has therefore shifted to efficient *approximation* algorithms with precise performance guarantees. In this lecture we will present a *fully polynomial randomized approximation scheme* for the permanent of a *randomly chosen* matrix (i.e., the algorithm works well with high probability over the choice of the input matrix, but may behave arbitrarily badly on a vanishing fraction of inputs).

12.2 An FPRAS for the Permanent of Random 0-1 Matrices

The goal of the lecture is to design a fully polynomial randomized approximation scheme (fpras) for *almost all* 0-1 matrices A . In other words, we will devise an algorithm that takes as input an $n \times n$ 0-1 matrix A and an accuracy parameter ϵ and outputs a random variable X_A such that

$$\Pr[(1 - \epsilon)\text{per}(A) \leq X_A \leq (1 + \epsilon)\text{per}(A)] \geq \frac{3}{4},$$

for almost all A (the meaning of “almost all” will be made precise later). Recall that “fully polynomial” requires that the run-time of the algorithm is polynomial in both $\frac{1}{\epsilon}$ and the size of the input, n . Note that the constant $\frac{3}{4}$ can be increased to $1 - \delta$ using only $O(\log \delta^{-1})$ trials by the median technique.

We remark that there is an fpras for the permanent of an *arbitrary* 0-1 matrix [JSV04], which we may see later in the course. Here, however, we focus on a much simpler algorithm, due to Rasmussen [R94], that works for almost all random matrices.

The algorithm is as follows:

Input: An $n \times n$ matrix A with 0-1 entries.
Output: A random variable X_A .
 if $n = 0$, then $X_A = 1$
 else let $W_A = \{j : a_{1j} = 1\}$ be the set of 1's in the first row
 Pick j from W_A u.a.r.
output $|W_A| \times X_{A_{1,j}}$,
 where $A_{1,j}$ is the $(1, j)$ -minor of A (i.e., row 1 and column j removed from A)

The above algorithm is essentially an iterative averaging scheme. Since $\text{per}(A) = \sum_j a_{1,j} \cdot \text{per}(A_{1,j})$, the algorithm works by assuming that the sub-permanent $\text{per}(A_{1,j})$ is the same for all j such that $a_{1,j}$ is 1.

We first argue that X_A is an *unbiased estimator* of $\text{per}(A)$ for any matrix A , i.e., that $E[X_A] = \text{per}(A)$. To see this, we can view the algorithm as a computation tree, where the root of the tree is the matrix A , with branches at the root to each minor $A_{1,j}$ with $a_{1j} = 1$ (thus the root has degree $|W_A|$); the computation tree is continued recursively from each child of the root. Every path from the root to a leaf l (at depth n) of the computation corresponds to a distinct (generalized) diagonal of A such that every entry of the diagonal is 1. Hence the number of leaves of the tree is exactly $\text{per}(A)$. Now observe that the algorithm reaches a particular leaf l with probability $\Pr[l] = \prod_{k=1}^n \frac{1}{|W_{A_k}|}$, where the W_{A_k} 's are the sets along the path from the root to l ; and, when it reaches l , it outputs precisely the reciprocal of this probability. Hence we have

$$E[X_A] = \sum_{\text{leaves } l} \Pr[l] \times \frac{1}{\Pr[l]} = \# \text{ of leaves} = \text{per}(A).$$

Remark 12.2 *The above argument is based on a technique of Knuth from the 1970's. The idea is that, given an arbitrary tree, to produce an unbiased estimator for the number of leaves of the tree it suffices to navigate the tree top-down while keeping track of the probabilities along the path from the root to the leaf, and output the reciprocal of the path probability.*

Proving that the algorithm works for most matrices amounts to showing that the random variable X_A is sufficiently concentrated about its mean; this is done by bounding the second moment.

Variance estimate. We first present an example that shows that the variance of the random variable X_A can be very bad in the worst case. Consider an upper triangular matrix A with 1's above the diagonal and 0's below it, as shown below:

$$A = \begin{bmatrix} 1 & 1 & 1 & \cdots & 1 \\ 0 & 1 & 1 & \cdots & 1 \\ 0 & 0 & 1 & \cdots & 1 \\ \vdots & \vdots & \vdots & \ddots & \\ 0 & 0 & 0 & & 1 \end{bmatrix}$$

There is only one generalized diagonal (the principal diagonal) all of whose entries are 1. Hence $\text{per}(A) = 1$, and there is only one leaf in the computation tree of the algorithm. The single computation path to this leaf is chosen with probability $\frac{1}{n!}$, which gives

$$X_A = \begin{cases} n! & \text{with prob } \frac{1}{n!} \\ 0 & \text{otherwise} \end{cases}$$

Hence the estimator is almost surely zero, and we would need a huge number of trials (on the order of $n!$) to get a decent estimate of $\text{per}(A)$.

For the rest of this lecture we focus on *random* matrices, and show that the above algorithm works with only polynomially many trials with high probability over the choice of matrix.

Definition 12.3 Let \mathcal{A}_n denote the probability space of $n \times n$ matrices such that every entry of the matrix is 0 or 1 with probability $\frac{1}{2}$ independently.

We will prove the following result.

Theorem 12.4 Let $A \in \mathcal{A}_n$ and let $\omega(n)$ be any function such that $\omega(n) \rightarrow \infty$. Then

$$\Pr_{\mathcal{A}_n} \left[\frac{\mathbb{E}[X_A^2]}{(\mathbb{E}[X_A])^2} > n \cdot \omega(n) \right] \rightarrow 0, \quad \text{as } n \rightarrow \infty.$$

Let us first interpret the result of Theorem 12.4. We can choose a function $\omega(n)$ that goes to infinity as slowly as we want. Since X_A is an unbiased estimator of $\text{per}(A)$, the number of trials of the algorithm required to ϵ -approximate $\text{per}(A)$ is $O(\frac{\mathbb{E}[X_A^2]}{\epsilon^2 \cdot (\mathbb{E}[X_A])^2})$, and the theorem says that with high probability this quantity is bounded by $O(\frac{n \cdot \omega(n)}{\epsilon^2})$.

Corollary 12.5 Given $A \in \mathcal{A}_n$, the above algorithm repeated $O(\frac{n \cdot \omega(n)}{\epsilon^2})$ times yields a fpras for $\text{per}(A)$ with probability tending to 1 over the choice of A .

Run-time analysis. The outer loop of the algorithm runs n times, and each iteration of the loop takes $O(n)$ time; hence each trial takes $O(n^2)$ time. Combining this with the bound on the number of trials gives an fpras for $\text{per}(A)$ that runs in time $O(\frac{n^3 \cdot \omega(n)}{\epsilon^2})$.

The proof of Theorem 12.4 will follow from a sequence of lemmas. We begin with the following claim:

Claim 12.6 The following assertions hold:

1. $\mathbb{E}_{\mathcal{A}_n}[\mathbb{E}[X_A]] = \mathbb{E}_{\mathcal{A}_n}[\text{per}(A)] = \frac{n!}{2^n}.$
2. $\mathbb{E}_{\mathcal{A}_n}[\mathbb{E}[X_A^2]] = \frac{1}{4^n} \prod_{i=1}^n (i^2 + i).$

Proof: An $n \times n$ matrix has $n!$ diagonals, and for each diagonal the probability that all entries are 1 is $\frac{1}{2^n}$. The first equation then follows by linearity of expectation.

We now present an alternative argument that proves both parts simultaneously. We may write

$$\mathbb{E}_{\mathcal{A}_n}[\text{per}(A)] = \mathbb{E} \left[\prod_{i=1}^n W_i \right],$$

where the W_i 's are independent and each $W_i \sim \text{Bin}(i, \frac{1}{2})$, i.e. is distributed as the # of heads in i tosses of a fair coin. Since

$$\mathbb{E}(W_i) = \frac{i}{2} \quad \text{and} \quad \mathbb{E}(W_i^2) = \frac{i^2 + i}{4},$$

we have

$$\begin{aligned} \mathbb{E}_{\mathcal{A}_n}[\mathbb{E}[X_A]] &= \prod_{i=1}^n \mathbb{E}[W_i] \quad (\text{by independence}) \\ &= \prod_{i=1}^n \frac{i}{2} = \frac{n!}{2^n}. \end{aligned}$$

For part 2, a straightforward induction (Exercise!) shows that $\mathbb{E}_{\mathcal{A}_n}[\mathbb{E}[X_A^2]] = \prod_{i=1}^n \mathbb{E}[W_i^2]$, from which we get

$$\begin{aligned} \mathbb{E}_{\mathcal{A}_n}[\mathbb{E}[X_A^2]] &= \prod_{i=1}^n \mathbb{E}[W_i^2] \\ &= \prod_{i=1}^n \frac{i^2 + i}{4} = \frac{1}{4^n} \prod_{i=1}^n (i^2 + i). \end{aligned}$$

■

Corollary 12.7

$$\frac{\mathbb{E}_{\mathcal{A}_n}[\mathbb{E}[X_A^2]]}{(\mathbb{E}_{\mathcal{A}_n}[\mathbb{E}[X_A]])^2} = \frac{\frac{1}{4^n} \prod_{i=1}^n (i^2 + i)}{\frac{1}{4^n} \prod_{i=1}^n i^2} = \prod_{i=1}^n \frac{i+1}{i} = n+1.$$

This Corollary shows that Theorem 12.4 holds “in expectation.” To turn this into a high-probability statement, we need to appeal to first and second moments (the first moment for the numerator and the second moment for the denominator). The second moment part is supplied by the following lemma. To simplify notation we will write $\mu(n) = \mathbb{E}_{\mathcal{A}_n}[\text{per}(A)]$.

Lemma 12.8 [Main Lemma] *For any $\omega(n) \rightarrow \infty$, we have $\Pr_{\mathcal{A}_n}[\text{per}(A) < \frac{\mu(n)}{\omega(n)}] \rightarrow 0$ as $n \rightarrow \infty$.*

We first prove Theorem 12.4 assuming Lemma 12.8.

Proof:(of Theorem 12.4). We handle the numerator and denominator of the expression in Theorem 12.4 as follows:

- *Numerator.* Markov's inequality gives

$$\Pr_{\mathcal{A}_n} \left[\mathbb{E}[X_A^2] \geq \omega(n) \cdot \mathbb{E}_{\mathcal{A}_n}[\mathbb{E}[X_A^2]] \right] \leq \frac{1}{\omega(n)} \rightarrow 0, \quad \text{as } n \rightarrow \infty.$$

- *Denominator.* Lemma 12.8 gives

$$\Pr_{\mathcal{A}_n} \left[\frac{1}{(\mathbb{E}[X_A])^2} > \frac{\omega(n)^2}{(\mathbb{E}_{\mathcal{A}_n}[\mathbb{E}[X_A]])^2} \right] \rightarrow 0.$$

The above line is a restatement of Lemma 12.8 with $\text{per}(A) = \mathbb{E}[X_A]$ and taking reciprocals and squaring.

Putting these together we have

$$\Pr_{\mathcal{A}_n} \left[\frac{\mathbb{E}[X_A^2]}{(\mathbb{E}[X_A])^2} > \omega(n)^3 \cdot \frac{\mathbb{E}_{\mathcal{A}_n}[X_A^2]}{(\mathbb{E}_{\mathcal{A}_n}[\mathbb{E}[X_A]])^2} \right] = \Pr_{\mathcal{A}_n} \left[\frac{\mathbb{E}[X_A^2]}{(\mathbb{E}[X_A])^2} > \omega(n)^3 \cdot (n+1) \right] \rightarrow 0,$$

where for the equality we used Corollary 12.7. Note that since $\omega(n)$ is an arbitrary function that goes to infinity, the same is true of $\omega(n)^3$. (Alternatively, we may replace $\omega(n)$ in the above argument by $\omega(n)^{1/3}$.) ■

Structure of proof of Main Lemma 12.8. We will consider generating a random matrix $A \in \mathcal{A}_n$ by first picking a number m according to the binomial distribution $\text{Bin}(n^2, \frac{1}{2})$, then distributing m 1's in the matrix uniformly at random, setting all other entries to 0.

Definition 12.9 We denote by $\mathcal{A}_{n,m}$ the probability space of random $n \times n$, 0-1 matrices where the number of 1's in the matrix is exactly m and the 1's are distributed uniformly at random in the matrix.

The reason we do this is that, for typical values of m (note that m will be sharply concentrated about its mean, $\frac{n^2}{2}$), $\text{per}(A)$ will be sharply concentrated about its mean in the model $\mathcal{A}_{n,m}$. This fact is expressed in the following lemma.

Lemma 12.10 Suppose $m = m(n)$ satisfies $\frac{m^2}{n^3} \rightarrow \infty$. Then for $A \in \mathcal{A}_{n,m}$ we have

1. $\mathbb{E}_{\mathcal{A}_{n,m}}[\text{per}(A)] = n! \cdot \left(\frac{m}{n^2}\right)^n \cdot \exp \left\{ -\frac{n^2}{2m} + \frac{1}{2} + O\left(\frac{n^3}{m^2}\right) \right\}.$
2. $\frac{\mathbb{E}_{\mathcal{A}_{n,m}}[\text{per}(A)^2]}{(\mathbb{E}_{\mathcal{A}_{n,m}}[\text{per}(A)])^2} = 1 + O\left(\frac{n^3}{m^2}\right).$

Observe that from part 2 of Lemma 12.10 it follows that given $\frac{m^2}{n^3} \rightarrow \infty$, we have $\frac{\text{Var}_{\mathcal{A}_{n,m}}[\text{per}(A)]}{(\mathbb{E}_{\mathcal{A}_{n,m}}[\text{per}(A)])^2} \rightarrow 0$, as $n \rightarrow \infty$. Hence the permanent is tightly concentrated in $\mathcal{A}_{n,m}$.

We now assume Lemma 12.10 and prove Lemma 12.8; to complete the entire analysis, we will then just need to go back and prove Lemma 12.10.

Proof of Lemma 12.8: We consider the following procedure to generate $A \in \mathcal{A}_n$:

- Pick M from $\text{Bin}(n^2, \frac{1}{2})$;
- pick $A \in \mathcal{A}_{n,M}$ u.a.r.

Let us denote by $\omega' = \omega'(n)$ an arbitrary function of n that goes to ∞ with n . We have the following inequalities:

- $\Pr[M < \frac{n^2}{2} - \omega'n] \rightarrow 0$; this follows by Chebyshev's inequality or the Central Limit Theorem because the standard deviation of M is $\Theta(n)$, so a deviation of $\omega'n$ is more than a constant times the s.d.
- For any $m = m(n)$ such that $\frac{m^2}{n^3} \rightarrow \infty$, we have

$$\begin{aligned} \Pr_{\mathcal{A}_{n,m}}[\text{per}(A) < \tfrac{1}{2} \mathbb{E}_{\mathcal{A}_{n,m}}[\text{per}(A)]] &< \frac{4 \text{Var}_{\mathcal{A}_{n,m}}[\text{per}(A)]}{(\mathbb{E}_{\mathcal{A}_{n,m}}[\text{per}(A)])^2} \quad (\text{by Chebyshev's inequality}) \\ &\rightarrow 0 \quad (\text{by part 2 of Lemma 12.10}). \end{aligned}$$

Since both of the above events happen with probability tending to 0, we will assume henceforth that neither of them happens: i.e., we assume that $M \geq \frac{n^2}{2} - \omega' \cdot n$ and $\text{per}(A) \geq \frac{1}{2} \mathbb{E}_{\mathcal{A}_{n,m}}[\text{per}(A)]$. Recall that $\mu(n) = \mathbb{E}_{\mathcal{A}_n}[\text{per}(A)]$, and define $\mu(n, m) = \mathbb{E}_{\mathcal{A}_{n,m}}[\text{per}(A)]$. We then have from part 1 of Lemma 12.10

$$\begin{aligned} \frac{\text{per}(A)}{\mu(n)} &\geq \frac{\frac{1}{2} \cdot \mu(n, \frac{n^2}{2} - \omega' \cdot n)}{\mu(n)} = \frac{1}{2} \cdot \frac{2^n}{n!} \cdot n! \cdot \left(\frac{\frac{n^2}{2} - \omega' \cdot n}{n^2} \right)^n \cdot \exp \left\{ -\frac{n^2}{n^2 - 2\omega' \cdot n} + \frac{1}{2} + O\left(\frac{1}{n}\right) \right\} \\ &\geq \frac{1}{2} \cdot \left(1 - \frac{2\omega'}{n} \right)^n \exp \left\{ -1 + O\left(\frac{1}{n}\right) \right\} \\ &\sim \frac{1}{2} \cdot \exp\{-2\omega' - 1\}. \end{aligned}$$

Finally, given an arbitrary function $\omega(n)$ such that $\omega(n) \rightarrow \infty$ as $n \rightarrow \infty$, we choose $\omega'(n) = \frac{1}{2} \cdot \log \frac{\omega(n)}{2e}$. Observe that $\omega'(n) \rightarrow \infty$ as $n \rightarrow \infty$, and from the above analysis we have $\frac{\text{per}(A)}{\mu(n)} \geq \frac{1}{2} \cdot \exp\{-2\omega' - 1\} = 1/\omega(n)$, as required to prove Lemma 12.8. ■

Finally, we go back and prove Lemma 12.10.

Proof of Lemma 12.10: The argument to prove Lemma 12.10 is graph-theoretic. Namely, we work with the interpretation of $\text{per}(A)$ (for an $n \times n$ 0-1 matrix A) as the number of perfect matchings in the associated graph G_A , as explained earlier. Given $A \in \mathcal{A}_{n,m}$, the graph G_A is a bipartite graph with n vertices on each side and exactly m edges distributed uniformly. Let H be a fixed labeled sub-graph of G_A with $t \leq 2n$ edges. Let $q(t) = \Pr[H \text{ is subgraph of } G_A]$. Then

$$q(t) = \frac{\binom{n^2-t}{m-t}}{\binom{n^2}{m}}.$$

To see this, note that $\binom{n^2-t}{m-t}$ is the number of possible ways of choosing G_A under the constraint that H is a subgraph of G_A (i.e., t edges are fixed), and $\binom{n^2}{m}$ is the number of possible ways of choosing m out of n^2 edges (i.e., choosing G_A). Hence we have

$$q(t) = \frac{m \cdot (m-1) \cdot \dots \cdot (m-t+1)}{n^2 \cdot (n^2-1) \cdot \dots \cdot (n^2-t+1)} = \left(\frac{m}{n^2} \right)^t \exp \left\{ -\frac{t^2}{2} \cdot \left(\frac{1}{m} - \frac{1}{n^2} \right) + O\left(\frac{n^3}{m^2}\right) \right\} \quad \text{for } t \leq 2n. \quad (12.1)$$

Exercise: Fill in the details of the above calculation. [Hint: Take logs, and use the approximation $\ln(1-x) = -x + O(x^2)$.]

Hence we have

$$\begin{aligned} \mathbb{E}_{\mathcal{A}_{n,m}}[\text{per}(A)] &= \sum_{\{H : H \text{ a perfect matching}\}} \Pr[H \text{ is subgraph of } G_A] \\ &= n! \cdot q(n) \\ &= n! \cdot \left(\frac{m}{n^2} \right)^n \cdot \exp \left\{ -\frac{n^2}{2m} + \frac{1}{2} + O\left(\frac{n^3}{m^2}\right) \right\}. \end{aligned} \quad (12.2)$$

This completes part 1 of Lemma 12.10.

We now prove part 2 of the Lemma. We have

$$\mathbb{E}_{\mathcal{A}_{n,m}}[\text{per}(A)^2] = \sum_{H, H'} \Pr[H, H' \text{ are subgraphs of } G_A],$$

where H, H' range over all pairs of perfect matchings in G_A . We first calculate the number of perfect matchings H and H' that overlap in exactly k edges. The expression for this is derived as follows: there are

$n!$ perfect matchings H , and given a perfect matching H the number of perfect matchings H' that overlap with H in exactly k edges is $\binom{n}{k} \times D(n-k)$, where $D(n-k)$ denotes the number of *derangements* of $(n-k)$ items (i.e., the number of permutations σ such that $\sigma(x) \neq x$ for all x). Hence the number of perfect matching pairs H, H' that have exactly k overlapping edges is given by $n! \cdot \binom{n}{k} \cdot D(n-k)$.

It is well known that $D(n) \sim \frac{n!}{e}$ (and the very small error in this estimate can be absorbed into our other error terms). Observe also that if H and H' overlap in exactly k edges, the union of H and H' has exactly $2n - k \leq 2n$ edges. Hence we have

$$\begin{aligned} \mathbb{E}_{\mathcal{A}_{n,m}}[\text{per}(A)^2] &= \sum_{k=0}^n n! \cdot \binom{n}{k} \cdot D(n-k) \cdot q(2n-k) \\ &\quad \vdots \\ &= n! \cdot \left(\frac{m}{n^2}\right)^{2n} \cdot \exp\left\{-\alpha n + O\left(\frac{n^3}{m^2}\right)\right\} \cdot \sum_{k=0}^n \binom{n}{k} \cdot D(n-k) \cdot \left(\frac{e^\alpha n^2}{m}\right)^k, \end{aligned}$$

where $\alpha = 2n(\frac{1}{m} - \frac{1}{n^2})$. [**Exercise:** Perform the algebraic manipulations to fill in the dots above, by plugging in the estimate (12.1) for $q(2n-k)$.]

We now compute the sum in the above expression:

$$\begin{aligned} \sum_{k=0}^n \binom{n}{k} \cdot D(n-k) \cdot \left(\frac{e^\alpha n^2}{m}\right)^k &\sim \sum_{k=0}^n \binom{n}{k} \cdot \frac{(n-k)!}{e} \cdot \left(\frac{e^\alpha n^2}{m}\right)^k \\ &= \frac{n!}{e} \cdot \sum_{k=0}^n \frac{1}{k!} \left(\frac{e^\alpha n^2}{m}\right)^k \\ &\leq \frac{n!}{e} \cdot \sum_{k=0}^{\infty} \frac{1}{k!} \left(\frac{e^\alpha n^2}{m}\right)^k \\ &= n! \cdot \exp\left\{\frac{e^\alpha n^2}{m} - 1\right\}. \end{aligned}$$

Since $e^\alpha = 1 + O(\frac{n}{m})$, the above expression simplifies to $n! \cdot \exp\{\frac{n^2}{m} - 1 + O(\frac{n^3}{m^2})\}$. Plugging this into the expression for $\mathbb{E}_{\mathcal{A}_{n,m}}[\text{per}(A)^2]$ we get

$$\mathbb{E}_{\mathcal{A}_{n,m}}[\text{per}(A)^2] = (n!)^2 \cdot \left(\frac{m}{n^2}\right)^{2n} \cdot \exp\left\{-\frac{n^2}{m} + 1 + O\left(\frac{n^3}{m^2}\right)\right\}. \quad (12.3)$$

Finally, we divide expression (12.3) for $\mathbb{E}_{\mathcal{A}_{n,m}}[\text{per}(A)^2]$ by the square of expression (12.2) for $(\mathbb{E}_{\mathcal{A}_{n,m}}[\text{per}(A)])^2$ and obtain the result claimed in part 2 of Lemma 12.10. [**Exercise:** check this!] ■

Exercise: Here is a trivial approximation algorithm for the permanent of a random matrix. On input $A \in \mathcal{A}_n$, let m be the number of 1's in A and simply output the approximation to $\mu(n, m)$ (the expected value of $\text{per}(A)$) given in part 1 of Lemma 12.10. Why is this significantly weaker than a fpras for $\text{per}(A)$?

References

- [JSV04] M. JERRUM, A. SINCLAIR and E. VIGODA “A polynomial-time approximation algorithm for the permanent of a matrix with nonnegative entries,” *Journal of the ACM* **51** (2004), pp. 671–697.

- [R94] L. E. RASMUSSEN, “Approximating the Permanent: a Simple Approach,” *Random Structures & Algorithms* **5** (1994), pp. 349–361.
- [Val79] L.G. VALIANT, “The complexity of computing the permanent,” *Theoretical Computer Science* **8** (1979), pp. 189–201.