A Monte-Carlo Algorithm
for Estimating the Permanent

N. Karmarkar
Bell Labs, Murray Hill, N.J.

R. Karp
University of California, Berkeley and
International Computer Science Institute, Berkeley, Calif.

R. Lipton
Princeton University

L. Lovász
Hungarian Academy of Sciences and Princeton University

M. Luby
International Computer Science Institute, Berkeley, Calif.

Abstract

Let A be an $n \times n$ matrix with 0-1 valued entries, and let per(A) be the permanent of A. We describe a Monte-Carlo algorithm which produces a "good in the relative sense" estimate of per(A) and has running time poly(n)$2^{n/2}$, where poly(n) denotes a function that grows polynomially with n.

1 Introduction

Let A be an $n \times n$ matrix with 0-1 valued entries, let $\det(A)$ denote the determinant of A and let $\text{per}(A)$ denote the permanent of A. The marked contrast between the computational complexity of computing $\det(A)$ versus that of computing $\text{per}(A)$, despite the deceiving similarity between the two tasks, has baffled researchers for years. One of the reasons for interest in computing $\text{per}(A)$ is that A can be viewed as the adjacency matrix of a bipartite graph, $H = (X, Y, E)$ where X corresponds to the rows in A, Y to the columns in A, and $A_{ij} = 1$ if there is and edge between X_i and Y_j. The quantity $\text{per}(A)$ is exactly the number of perfect matchings in H.

It is well known that $\det(A)$ can be computed in $\text{poly}(n)$ time. On the other hand, the fastest algorithm known for computing $\text{per}(A)$ runs in $n2^n$ time [Ryser]. Solid grounds for arguing that computing $\text{per}(A)$ is an inherently difficult problem were first provided in [Valiant], which shows that the problem is $\#P$-complete. One implication of this result is that if $P \neq NP$ then there is no $\text{poly}(n)$ time algorithm for computing $\text{per}(A)$.

Because of the apparent nonexistence of a $\text{poly}(n)$ time algorithm for computing $\text{per}(A)$ exactly, we focus our attention on finding an algorithm that produces a good estimate of $\text{per}(A)$ and has a small running time. An (ϵ, δ) *approximation algorithm* for $\text{per}(A)$ is a Monte-Carlo algorithm which accepts as input A and two positive parameters ϵ and δ. The output of the algorithm is an estimate Y of $\text{per}(A)$ which satisfies

$$\Pr[(1 - \epsilon)\text{per}(A) \leq Y \leq (1 + \epsilon)\text{per}(A)] \geq 1 - \delta.$$

The papers [JVV], [KL], [KLM] discuss (ϵ, δ) approximation algorithms for counting problems in greater detail. We develop an (ϵ, δ) approximation algorithm for $\text{per}(A)$ which runs in $2^{n/2} \frac{1}{\epsilon^2} \log(\frac{1}{\delta}) \text{poly}(n)$ time. For fixed ϵ and δ the running time of the approximation algorithm is essentially the square root of the running time for the fastest known algorithm that computes $\text{per}(A)$ exactly.

In [Broder], [JS1], [JS2] (ϵ, δ) approximation algorithms for $\text{per}(A)$ are given which run in $\text{poly}(n)$ time in the special case when each row and column in A contains at least $n/2$ $1's$. Whether or not there is an (ϵ, δ) approximation algorithm for $\text{per}(A)$ which runs in $\text{poly}(n)$ time for general A is still an open problem.

2 Some General Considerations Regarding (ϵ, δ) Approximation Algorithms

Suppose we would like to estimate some quantity Q and have available a stochastic experiment whose output is a random variable X such that $E[X] = Q$ and $E[X^2]$ is finite. Suppose further that we can repeat this experiment as many times as we wish, and that the outcomes of the successive trials will be independent and identically distributed, with the same distribution as X. Let X_i be the outcome of the i^{th} trial. A straightforward application of Chebychev's Inequality shows that, if we conduct N trials, where $N = \frac{E[X^2]}{E[X]^2} \frac{1}{\epsilon^2 \delta}$ then $\frac{\sum_{i=1}^{N} X_i}{N}$ gives an (ϵ, δ) approximation to Q.

We can improve the dependence of the number of trials on δ using a well-known trick. Setting $\delta = 1/4$, we find that, if $N = \frac{4E[Y^2]}{E[Y]^2} \frac{1}{\epsilon^2}$, the probability is at least $3/4$ that the average of the N trials will lie within ϵ of Q. To obtain an (ϵ, δ)-approximation algorithm, we repeat such an N-sample experiment K times, where K is an odd integer greater than a suitable constant c times $\log(1/\delta)$, and take as our estimator of Q the median of the estimators produced by the K experiments of N-samples each. Let us say that the outcome of a N-sample experiment is *good* if it lies within ϵ of Q. Then the median of the K outcomes will be good whenever the majority of the K outcomes are good. Using the fact that the outcomes are independent and identically distributed, and that each outcome has probability at least $3/4$ of being good, standard bounds on the tail of the binomial distribution [Chernoff] reveal that the median is good with probability at least $1 - \delta$. Thus, the number of trials required for an (ϵ, δ) approximation to Q is $O\left(\frac{E[X^2]}{E[X]^2} \frac{1}{\epsilon^2} \log(\frac{1}{\delta})\right)$.

3 The Godsil/Gutman Estimator of $\mathrm{per}(A)$

The discussion of the last section shows how an (ϵ, δ)-algorithm for approximating a quantity Q can be constructed from any computable stochastic experiment whose outcome is a random variable Y such that $E[Y] = Q$ and $E[Y^2]$ is finite. The efficiency of the algorithm will be based on the computational difficulty of performing the stochastic experiment, and on the ratio $\frac{E[Y^2]}{E[Y]^2}$. The rest of the paper is devoted to studying two particular stochastic experiments for estimating $\mathrm{per}(A)$. The first of these, which we call the Godsil/Gutman estimator, was suggested in [GG]; the second one is a variant of the Godsil/Gutman estimator which has a smaller second moment and thus leads to a more efficient algorithm.

3

The Godsil/Gutman estimator is defined as follows:

(1) An $n \times n$ matrix B is formed from A as follows:

For all i, j, $1 \le i, j \le n$,

If $A_{ij} = 0$ then $B_{ij} \leftarrow 0$

Elseif $A_{ij} = 1$ then randomly and independently choose $B_{i,j} \in \{-1, 1\}$,

each choice with probability $\frac{1}{2}$.

(2) $Y \leftarrow (\det(B))^2$.

This stochastic experiment can be executed in poly(n) time.

In what follows we introduce some technical language that is appropriate for all of the following analysis, show that $E[Y] = \text{per}(A)$ and derive an upper bound on $E[Y^2]$.

4   Terminology

Let P be the set of all $n!$ permutations of $1, \ldots, n$. For all $\sigma \in P$, let $\text{sgn}(\sigma) = -1^t$, where t is the number of transpositions to form σ.

Let $P(A) \subseteq P$ be the set of all permutations σ such that for $i = 1, \ldots, n$, $A_{i,\sigma(i)} = 1$. Then, $\text{per}(A) = \sum_{\sigma \in P(A)} 1 = |P(A)|$. For each $\sigma \in P(A)$, for $i = 1, \ldots, n$, we label $< i, \sigma(i) >$ with the symbol σ. Let $P^2(A) = P(A) \times P(A)$. For each $\dot{\sigma} = < \sigma_1, \sigma_2 > \in P^2(A)$, let $G(\dot{\sigma})$ be the unlabelled graph where there is an unlabelled node for each $< i, j >$ which is labelled with either or both of σ_1, σ_2, and where there is an edge between two distinct nodes $< i, j >$ and $< i', j' >$ iff $i = i'$ or $j = j'$. Each connected component of $G(\dot{\sigma})$ is an isolated node or an even length cycle. For each cycle in $G(\dot{\sigma})$, designate one of the nodes in the cycle as the root of the cycle. Let $D = \{G(\dot{\sigma}) : \dot{\sigma} \in P^2(A)\}$. For each graph $G \in D$, let $c(G)$ be the number of cycles in G. For each $G \in D$, let

$$\dot{\text{eq}}(G) = \{\dot{\sigma} \in P^2(A) : G(\dot{\sigma}) = G\}.$$

Proposition 1 $|\dot{\text{eq}}(G)| = 2^{c(G)}$.

PROOF: Let $\dot{\sigma} = < \sigma_1, \sigma_2 > \in \dot{\text{eq}}(G)$. Each isolated node in G is labelled with both σ_1 and σ_2. Let c be a cycle in G. If the label of the root in c is σ_1 then every

4

node at an even distance from the root in c must be labelled σ_1 and every node at an odd distance from the root in c must be labelled σ_2. The case when the root is labelled σ_2 is symmetric, interchanging the roles of σ_1 and σ_2. Thus, for each cycle there are two possible labellings and the total number of labellings of all cycles is then $2^{c(G)}$. ∎

Let $D' = \{G \in D : c(G) = 0\}$.

Proposition 2

(1) $G(<\sigma_1, \sigma_2>) \in D' \Leftrightarrow \sigma_1 = \sigma_2$

(2) $G(<\sigma_1, \sigma_1>) = G(<\sigma_2, \sigma_2>) \Leftrightarrow \sigma_1 = \sigma_2$

(3) $|D'| = \text{per}(A)$

(4) *For all* $G \in D', |\dot{\text{eq}}(G)| = 1$.

5 Analysis of the Godsil/Gutman Estimator

Theorem 1 (Godsil/Gutman) $E[Y] = \text{per}(A)$.

PROOF: For each $\dot{\sigma} =< \sigma_1, \sigma_2 >\in P^2(A)$, let

$$x(\dot{\sigma}) = \prod_{k=1}^{2} \left(\text{sgn}(\sigma_k) \prod_{i=1}^{n} B_{i,\sigma_k(i)} \right).$$

Then, since $\det(B) = \sum_{\sigma \in P(A)} \text{sgn}(\sigma) \prod_{i=1}^{n} B_{i,\sigma(i)}$,

$$Y = (\det(B))^2 = \sum_{\dot{\sigma} \in P^2(A)} x(\dot{\sigma}) = \sum_{G \in D} \sum_{\dot{\sigma} \in \dot{\text{eq}}(G)} x(\dot{\sigma})$$

$$= \sum_{\dot{\sigma} =<\sigma_1, \sigma_1> \in P^2(A)} x(\dot{\sigma}) \tag{1}$$

$$+ \sum_{G \in D - D'} \sum_{\dot{\sigma} \in \dot{\text{eq}}(G)} x(\dot{\sigma}). \tag{2}$$

For each $\dot{\sigma} =< \sigma_1, \sigma_1 >\in P^2(A)$, $x(\dot{\sigma}) = 1$ independent of the values chosen for B. Thus, part (1) is equal to $\text{per}(A)$ because the number of terms in the sum is

5

$|P(A)| = \text{per}(A)$. We show that the expected value of part (2) is equal to 0 as follows. Fix $G \in D - D'$ and $\dot{\sigma} = <\sigma_1, \sigma_2> \in \dot{\text{eq}}(G)$. We show that $\text{E}[x(\dot{\sigma})] = 0$, thus showing that the expected value of every term in part (2) is 0. Because $G \in D - D'$, G contains at least one cycle. Let $<i, j>$ be some node in some cycle of G. Because either $<i, j>$ is labelled with σ_1 and not with σ_2 or vice-versa, $x(\dot{\sigma})$ can be written as $y(\dot{\sigma})B_{i,j}$, where $y(\dot{\sigma})$ does not contain $B_{i,j}$. Because $B_{i,j}$ is independent of $y(\dot{\sigma})$,

$$\text{E}[x(\dot{\sigma})] = \text{E}[y(\dot{\sigma})]\text{E}[B_{i,j}].$$

Because $\text{E}[B_{i,j}] = 0$, $\text{E}[x(\dot{\sigma})] = 0$. ∎

Theorem 2 $\dfrac{\text{E}[Y^2]}{\text{E}[Y]^2} = \dfrac{\displaystyle\sum_{G \in D} 6^{c(G)}}{\displaystyle\sum_{G \in D} 2^{c(G)}}.$

PROOF:

$$\text{E}[Y]^2 = (\text{per}(A))^2 = \sum_{<\sigma_1,\sigma_2> \in P^2(A)} 1 = \sum_{G \in D} \sum_{<\sigma_1,\sigma_2> \in \dot{\text{eq}}(G)} 1 = \sum_{G \in D} 2^{c(G)},$$

where the last equality is from Proposition 1. Let $P^4(A) = P(A) \times P(A) \times P(A) \times P(A)$. For each $\ddot{\sigma} = <\sigma_1, \sigma_2, \sigma_3, \sigma_4> \in P^4(A)$, let $x(\ddot{\sigma}) = \prod_{k=1}^{4} \left(\text{sgn}(\sigma_k) \sum_{i=1}^{n} B_{i,\sigma_k(i)} \right)$. Then,

$$Y^2 = (\det(B))^4 = \sum_{\ddot{\sigma} \in P^4(A)} x(\ddot{\sigma})$$

and

$$\text{E}[Y^2] = \sum_{\ddot{\sigma} \in P^4(A)} \text{E}[x(\ddot{\sigma})].$$

Let ODD $= \{ \ddot{\sigma} \in P^4(A) :$ there is some $<i, j>$ which is labelled with an odd number of labels from $\{\sigma_1, \sigma_2, \sigma_3, \sigma_4\}\}$, and let EVEN $= P^4(A) - \text{ODD}$. For each $\ddot{\sigma} \in$ ODD, there is some $<i, j>$ such that $x(\ddot{\sigma})$ can be written as $y(\ddot{\sigma})B_{i,j}$ where $y(\ddot{\sigma})$ does not contain $B_{i,j}$. Thus, $\text{E}[x(\ddot{\sigma})] = \text{E}[y(\ddot{\sigma})]\text{E}[B_{i,j}]$. Because $\text{E}[B_{i,j}] = 0$, $\text{E}[x(\ddot{\sigma})] = 0$. For each $\ddot{\sigma} \in$ EVEN, for each row $i = 1, \ldots, n$, either there is a j such that $<i, j>$ is labelled with all four of $\{\sigma_1, \sigma_2, \sigma_3, \sigma_4\}$, or there is a j and a $j' \neq j$ such that $<i, j>$ is labelled with exactly two of $\{\sigma_1, \sigma_2, \sigma_3, \sigma_4\}$ and $<i, j'>$ is labelled with the other two. Let $G(\ddot{\sigma})$ be the graph where there is a node for each $<i, j>$ such that $<i, j>$ is labelled with at least one of $\{\sigma_1, \sigma_2, \sigma_3, \sigma_4\}$.

There is an edge between two distinct nodes $< i, j >$ and $< i', j' >$ iff $i = i'$ or $j = j'$. Then, $\{G(\ddot{\sigma}) : \ddot{\sigma} \in \text{EVEN}\} = D$, where D is as previously defined. For each $G \in D$, define

$$\ddot{\text{eq}}(G) = \{\ddot{\sigma} \in \text{EVEN} : G(\ddot{\sigma}) = G\}.$$

We claim that $|\ddot{\text{eq}}(G)| = 6^{c(G)}$. The reasoning is similar to that used for the proof of Proposition 1. Let $\ddot{\sigma} = < \sigma_1, \sigma_2, \sigma_3, \sigma_4 > \in \ddot{\text{eq}}(G)$. Each isolated node in G is labelled with all four of $\{\sigma_1, \sigma_2, \sigma_3, \sigma_4\}$. For each cycle c in G, the root of c and every node at an even distance from the root must be labelled with the same two elements of $\{\sigma_1, \sigma_2, \sigma_3, \sigma_4\}$ and every node at an odd distance from the root must be labelled with the remaining two. Thus, there are $\binom{4}{2} = 6$ possible labellings of each cycle, and consequently a total of $6^{c(G)}$ labellings.

It is not hard to see that for each $\ddot{\sigma} \in \text{EVEN}$, $x(\ddot{\sigma}) = 1$ independent of the values chosen for the entries in B, and thus $\text{E}[x(\ddot{\sigma})] = 1$.

Thus, $\text{E}[Y^2] = \sum_{\ddot{\sigma} \in \text{EVEN}} 1 = \sum_{G \in D} 6^{c(G)}$. Since $\text{E}[Y]^2 = \sum_{G \in D} 2^{c(G)}$, the proof is complete. ∎

Corollary 3 *The Godsil/Gutman estimator yields an (ϵ, δ)-approximation algorithm for estimating $\text{per}(A)$ which runs in time $\text{poly}(n) 3^{\frac{n}{2}} \frac{1}{\epsilon^2} \log(\frac{1}{\delta})$.*

PROOF: Each evaluation of the estimator can be performed in time $\text{poly}(n)$. Also, $\frac{\text{E}[Y^2]}{\text{E}[Y]^2} = \dfrac{\sum_{G \in D} 6^{c(G)}}{\sum_{G \in D} 2^{c(G)}} \leq \max_{G \in D} 3^{c(G)} \leq 3^{n/2}$, where the last inequality follows because there are at most $n/2$ cycles in any $G \in D$. ∎

6   A Better Estimator and Its Analysis

We now present a variant of the Godsil/Gutman estimator which yields a more efficient (ϵ, δ)-approximation algorithm for $\text{per}(A)$. Let

$$w_0 = 1, \ w_1 = -\frac{1}{2} + \frac{\sqrt{3}}{2}i, \ w_2 = -\frac{1}{2} - \frac{\sqrt{3}}{2}i$$

be the three cube roots of unity. If $y = a + bi$ is a complex number, then $\overline{y} = a - bi$ is the complex conjugate of y.

The estimator is computed as follows.

(1) An $n \times n$ matrix B is formed from A as follows:

For all $i, j,\ 1 \le i, j \le n$,

If $A_{i,j} = 0$ then $B_{i,j} \leftarrow 0$

Elseif $A_{i,j} = 1$ then randomly and independently choose

$B_{i,j} \in \{w_0, w_1, w_2\}$, each choice with probability $\frac{1}{3}$.

(2) $Z \leftarrow \det(B)\overline{\det(B)}$.

This estimator can be evaluated in poly(n) time.

Theorem 3 $\mathrm{E}[Z] = \mathrm{per}(A)$.

PROOF: Similar to the proof of Theorem 1. For each $\dot{\sigma} = <\sigma_1, \sigma_2> \in P^2(A)$, let

$$x(\dot{\sigma}) = \mathrm{sgn}(\sigma_1)\mathrm{sgn}(\sigma_2) \prod_{i=1}^{n} B_{i,\sigma_1(i)} \overline{B}_{i,\sigma_2(i)}.$$

Then,

$$Z = \det(B)\overline{\det(B)} = \sum_{\dot{\sigma} \in P^2(A)} x(\dot{\sigma})$$

$$= \sum_{\dot{\sigma} = <\sigma_1, \sigma_1> \in P^2(A)} x(\dot{\sigma}) \tag{3}$$

$$+ \sum_{G \in D - D'} \sum_{\dot{\sigma} \in \ddot{\mathrm{q}}(G)} x(\dot{\sigma}). \tag{4}$$

For each $\dot{\sigma} = <\sigma_1, \sigma_1> \in P^2(A)$, $x(\dot{\sigma}) = 1$ independent of the values chosen for B. Thus, part (3) is equal to per(A). Showing that the expected value of part (4) is equal to 0 is very similar to the corresponding portion of the proof in Theorem 1. The observation needed is again that for any $<i, j>$, $\mathrm{E}[B_{i,j}] = 0$. ∎

Theorem 4 $\dfrac{\mathrm{E}[Z^2]}{\mathrm{E}[Z]^2} = \dfrac{\displaystyle\sum_{G \in D} 4^{c(G)}}{\displaystyle\sum_{G \in D} 2^{c(G)}}.$

PROOF: The proof follows exactly the outline of the proof of Theorem 2. We only note the differences here.

8

For each $\ddot{\sigma} =< \sigma_1, \sigma_2, \sigma_3, \sigma_4 >\in P^4(A)$, let

$$x(\ddot{\sigma}) = \left(\prod_{k=1}^{2} \text{sgn}(\sigma_k) \prod_{i=1}^{n} B_{i,\sigma_k(i)}\right) \left(\prod_{k=3}^{4} \text{sgn}(\sigma_k) \prod_{i=1}^{n} \overline{B}_{i,\sigma_k(i)}\right).$$

Then, $\text{E}[Z^2] = \sum_{\ddot{\sigma}\in P^4(A)} \text{E}[x(\ddot{\sigma})]$. ¿From the definitions in the proof of Theorem 2,

$$P^4(A) = \text{ODD} \cup \text{EVEN} = \text{ODD} \cup \left(\bigcup_{G\in D}\left(\bigcup_{\ddot{\sigma}\in\ddot{\text{eq}}(G)} \dot{\sigma}\right)\right).$$

For $\ddot{\sigma} \in$ ODD there is some $< i,j >$ such that $x(\ddot{\sigma})$ can be written as $y(\ddot{\sigma})B_{i,j}$ or as $y(\ddot{\sigma})\overline{B}_{i,j}$, where in either case $y(\dot{\sigma})$ contains neither $B_{i,j}$ nor $\overline{B}_{i,j}$. Because $\text{E}[B_{i,j}] = \text{E}[\overline{B}_{i,j}] = 0$, $\text{E}[x(\ddot{\sigma})] = 0$. For each $G \in D$ we further partition $\ddot{\text{eq}}(G)$ as follows:

$$\ddot{\text{eq}}'(G) = \{\ddot{\sigma} \in \ddot{\text{eq}}(G) : \text{in each cycle } c \text{ in } G \text{ the root of } c \text{ is}$$
$$\text{labelled with exactly one of } \{\sigma_1,\sigma_2\} \text{ and with exactly}$$
$$\text{one of } \{\sigma_3,\sigma_4\}\}.$$

Because there are four possible labellings for each cycle in G, $|\ddot{\text{eq}}'(G)| = 4^{c(G)}$. For each $\ddot{\sigma} \in \ddot{\text{eq}}'(G)$, for each node $< i,j >$ in G there are an equal number of occurrences of $B_{i,j}$ and $\overline{B}_{i,j}$ in $x(\ddot{\sigma})$. Thus, $x(\ddot{\sigma}) = 1$ independent of the values chosen for B, and $\text{E}[x(\ddot{\sigma})] = 1$. For each $\ddot{\sigma} \in \ddot{\text{eq}}(G) - \ddot{\text{eq}}'(G)$, there is some node $< i,j >$ in G such that $B_{i,j}$ occurs twice in $x(\ddot{\sigma})$ and $\overline{B}_{i,j}$ does not occur at all. Thus, $x(\ddot{\sigma})$ can be written as $y(\ddot{\sigma})B_{i,j}^2$ where $y(\ddot{\sigma})$ contains no occurrences of $B_{i,j}$ or $\overline{B}_{i,j}$. Then, $\text{E}[x(\ddot{\sigma})] = \text{E}[y(\ddot{\sigma})]\text{E}[B_{i,j}^2]$. Because $\text{E}[B_{i,j}^2] = 0$, $\text{E}[x(\ddot{\sigma})] = 0$.

Putting this together, $\text{E}[Z^2] = \sum_{G\in D} 4^{c(G)}$, and thus

$$\frac{\text{E}[Z^2]}{\text{E}[Z]^2} = \frac{\sum_{G\in D} 4^{c(G)}}{\sum_{G\in D} 2^{c(G)}}.$$

∎

Corollary 4 *The estimator Z yields an (ϵ, δ)-approximation algorithm for estimating* $\text{per}(A)$ *which runs in time* $\text{poly}(n)2^{\frac{n}{2}}\frac{1}{\epsilon^2}\log(\frac{1}{\delta})$.

9

PROOF: Each evaluation of the estimator can be performed in time poly(n).

Also, $\frac{\mathrm{E}[Z^2]}{\mathrm{E}[Z]^2} = \dfrac{\sum\limits_{G \in D} 4^{c(G)}}{\sum\limits_{G \in D} 2^{c(G)}} \leq \max_{G \in D} 2^{c(G)} \leq 2^{n/2}$, where the last inequality follows

because there are at most $n/2$ cycles in any $G \in D$. ∎

7 Some Special Cases

We have introduced two unbiased estimators of per(A): the Godsil/Gutman estimator Y and a second estimator Z which refines the Godsil/Gutman technique by using cube roots of unity. We showed that

$$\frac{\mathrm{E}[Y^2]}{\mathrm{E}[Y]^2} = \frac{\sum\limits_{G \in D} 6^{c(G)}}{\sum\limits_{G \in D} 2^{c(G)}}$$

and

$$\frac{\mathrm{E}[Z^2]}{\mathrm{E}[Z]^2} = \frac{\sum\limits_{G \in D} 4^{c(G)}}{\sum\limits_{G \in D} 2^{c(G)}}.$$

We then obtained an upper bound of $3^{\frac{n}{2}}$ on $\frac{\mathrm{E}[Y^2]}{\mathrm{E}[Y]^2}$ and an upper bound of $2^{\frac{n}{2}}$ on $\frac{\mathrm{E}[Z^2]}{\mathrm{E}[Z]^2}$ using the trivial observation that, for all G, $c(G) \leq \frac{n}{2}$. Although this bound seems terribly pessimistic, the following example shows that there are cases in which it is close to the truth. Let n be even and let A be the $n \times n$ matrix such that for $i = 1, \ldots, \frac{n}{2}$, $A_{2i-1,2i-1} = A_{2i-1,2i} = A_{2i,2i-1} = A_{2i,2i} = 1$ and all other entries in A are 0. For the first algorithm the expected number of trials before there is even one trial where $Y \neq 0$ is $\Omega(2^{\frac{n}{2}})$, and for the second algorithm it is $\Omega((\frac{3}{2})^{\frac{n}{2}})$. There is a lot of room between $3^{\frac{n}{2}}$ and $2^{\frac{n}{2}}$ and between $(\frac{3}{2})^{\frac{n}{2}}$ and $2^{\frac{n}{2}}$; we are not sure which bound is closer to the worst case behavior for the respective algorithms.

Despite this bad example, one suspects that in most situations, these bounds are far too pessimistic, since $c(G)$ will "typically" be much smaller than $n/2$. As one heuristic indication of this phenomenon we note that, if $\dot{\sigma} =< \sigma_1, \sigma_2 >$, where σ_1 and σ_2 are independent random permutations, then $c(G(\dot{\sigma}))$ will be close to $\ln(n)$ with very high probability.

In this section we analyze a concrete example in which the bounds of $3^{n/2}$ for the Godsil/Gutman algorithm and $2^{n/2}$ for its refinement are provably too pessimistic: the $n \times n$ matrix C in which every element is equal to 1. The permanent of this matrix is, of course, $n!$.

Let the random variable Y be the estimator produced by the Godsil/Gutman algorithm applied to the matrix C, and let Z be the estimator produced by the refinement given in Section 6 applied to the matrix C. Our main result is as follows:

Theorem 5 *In the particular case of the $n \times n$ matrix C, $\frac{E[Y^2]}{E[Y]^2} \leq \frac{(n+1)(n+2)}{2}$ and $\frac{E[Z^2]}{E[Z]^2} \leq n + 1$.*

PROOF: Since the two estimators are unbiased, $E[Y] = E[Z] = n!$. In order to discuss the second moments of Y and Z we require a definition: for any permutation σ of $\{1, 2, \ldots, n\}$, let $d(\sigma)$ be the number of cycles of length at least two in the permutation σ. We shall show that

$$E[Y]^2 = n! \sum_\sigma 3^{d(\sigma)}$$

and

$$E[Z^2] = n! \sum_\sigma 2^{d(\sigma)},$$

where each summation extends over the permutations of $\{1, 2, \ldots, n\}$.

As in Section 4, let P be the set of all permutations of $\{1, 2, \ldots, n\}$, and let D be the set of graphs $G(\dot\sigma)$ associated with pairs of permutations in P. We know from Theorems 2 and 4 that $E[Y^2] = \sum_G 6^{c(G)}$ and $E[Z^2] = \sum_G 4^{c(G)}$.

We define a relation $R \subseteq D \times P$ as follows: let G be a graph in D and let σ be a permutation; then $GR\sigma$ if and only if there exists an ordered pair $\dot\sigma = <\sigma_1, \sigma_2>$ such that $G = G(\dot\sigma)$ and $\sigma = \sigma_1^{-1}\sigma_2$. We note the following properties:

- If $GR\sigma$ then $c(G) = d(\sigma)$.

- For each $\sigma \in P$ there are exactly $n!$ graphs G such that $GR\sigma$; these are exactly the graphs $G(\dot\sigma)$ such that $\dot\sigma = <\sigma_1, \sigma_2>$ where $\sigma = \sigma_1^{-1}\sigma_2$.

- For each G in D there are exactly $2^{c(G)}$ permutations σ such that $GR\sigma$.

It follows from these properties that $\sum_{G \in D} 6^{c(G)} = n! \sum_{\sigma \in P} 3^{d(\sigma)}$ and $\sum_{G \in D} 4^{c(G)} = n! \sum_{\sigma \in P} 2^{d(\sigma)}$ But it follows from [Lovász] (solution to Exercise 3.12, p. 203) that

11

$\sum_{\sigma \in P} 3^{d(\sigma)} \leq \frac{(n+2)!}{2}$ and $\sum_{\sigma \in P} 2^{d(\sigma)} \leq (n+1)!$. Hence, $\frac{E[Y^2]}{E[Y]^2} \leq \frac{(n+2)(n+1)}{2}$ and $\frac{E[Z^2]}{E[Z]^2} \leq n+1$. ∎

It follows that, in the special case of the matrix C, a quadratic number of trials of the Godsil/Gutman Monte-Carlo algorithm, or a linear number of trials of our refinement of the Godsil/Gutman algorithm, are sufficient for an (ϵ, δ)-approximation. We conjecture that similar claims hold true for almost all $n \times n$ zero-one matrices.

8   Comments, Generalizations, Refinements

For both algorithms presented here, once the values of B have been chosen, the value of the estimator can be computed exactly in poly(n) time. In the second algorithm, this requires that $\sqrt{3}$ be represented symbolically. The symbol $\sqrt{3}$ will not appear in the final answer, which is an integer.

Both Monte-Carlo algorithms can easily be modified to estimate per(A) when the entries in A are allowed to be arbitrary positive numbers. However, in this case there is an issue with precision; the estimators cannot be computed exactly in poly(n) time, although they can be closely approximated. The modification for the Godsil/Gutman estimator is to randomly choose $B_{i,j} \in \left\{ \sqrt{A_{i,j}}, -\sqrt{A_{i,j}} \right\}$, each choice with probability $1/2$. The modification for the second estimator is to randomly choose $B_{i,j} \in \left\{ \sqrt{A_{i,j}} w_0, \sqrt{A_{i,j}} w_1, \sqrt{A_{i,j}} w_2 \right\}$, each choice with probability $1/3$. In both cases, the expected value of the estimator is per(A) and the upper bounds on the number of trials to guarantee an (ϵ, δ) approximation algorithm stated in Corollaries 3 and 4 still apply, assuming that at each trial the estimator is computed exactly. The analysis would have to be modified to allow for truncation error.

Each of our Monte-Carlo algorithms consists of $O(\log(\frac{1}{\delta}))$ phases, with each phase consisting of some number N of independent trials. We required that each phase be an $(\epsilon, 1/4)$ approximation algorithm. Our analysis of the upper bound on the number of trials to guarantee this property was based on Chebychev's Inequality, and thus the analysis still holds if trials are pairwise independent. Consider running the first Monte-Carlo algorithm with a fixed ϵ and δ. As it is written it requires $\Theta(3^{\frac{n}{2}} n^2)$ random bits per phase in total, i.e. n^2 random bits per trial to randomly choose the values for B. This can be reduced to $O(n^3)$ random bits per phase using standard methods of generating pairwise independent unbiased random bits [ACGS], [CG], [Luby].

12

9 Open Questions

(1) Is there an (ϵ, δ) approximation algorithm for $\mathrm{per}(A)$ which runs in $\mathrm{poly}(n)$ time? One possible approach to solving this problem is based on the observation that the trials within a phase need only be pairwise independent, rather than completely independent. The result of each phase is $\frac{\sum_{i=1}^{\ell} Y_i}{\ell}$ where ℓ, the number of trials, is exponential in n. Pairwise independence permits the ℓ samples to be generated using only $O(n^3)$ random bits. Perhaps the rule for generating the samples from the random bits can be designed so that the quantity $\sum_{i=1}^{\ell} Y_i$ can be computed directly and efficiently from the random bits, without the need to calculate the result of each trial explicitly. Similar ideas have been used successfully in other contexts [ACGS], [Bach], [CG], [CW], [KR], [Luby].

(2) Is there a deterministic algorithm with running time $o(2^n)$ which accepts as input A and ϵ and which outputs Y such that

$$(1 - \epsilon)\mathrm{per}(A) \leq Y \leq (1 + \epsilon)\mathrm{per}(A)?$$

References

[ACGS] Alexi, W., Chor, B., Goldreich, O., Schnorr, C.P., "RSA Rabin Functions: Certain Parts Are As Hard As the Whole", *SIAM J. on Computing*, Vol. 17, 1988, pp. 194-209. A preliminary version appeared in 25^{th} *FOCS*, 1984, pp. 449-457.

[Bach] E. Bach, "Realistic Analysis of Some Randomized Algorithms", 19^{th} *STOC*, 1987, pp. 453-461.

[Broder] A. Broder, "How hard is it to marry at random (On the approximation of the permanent)", 18^{th} *STOC*, 1986, pp. 50-58.

[Chernoff] H. Chernoff, "A Measure of Asymptotic Efficiency for Tests of a Hypothesis Based on the Sum of Observations", *Annals of Math. Stat.*, Vol. 23, pp. 493-509, 1952.

[CG] B. Chor and O. Goldreich, "On the Power of Two-Points Based Sampling", to appear in Journal of Complexity.

[CW]   J. Carter and M. Wegman, "Universal classes of hash functions", *J. Comput. System Sci.*, 18, 1979, pp. 143-154.

[GG]   Godsil, C.D., and Gutman, I., "On the matching polynomial of a graph", *Algebraic Methods in Graph Theory*, I, L Lovász and V.T. Sós, editors, Colloq. Math. Soc. János Bolyai, 25, North-Holland, Amsterdam, 1981, pp. 241-249.

[JS1]   M. Jerrum and A. Sinclair, "Conductance and the rapid mixing property for Markov Chains: the approximation of the permanent resolved", *STOC* 1988, pp. 235-243.

[JS2]   M. Jerrum and A. Sinclair, "Approximating the permanent", *Internal Report CSR-275-88*, Department of Computer Science, University of Edinburgh, Submitted to *SIAM J. on Computing*.

[JVV]   M. Jerrum, L. Valiant and V. Vazirani, "Random generation of combinatorial structures from a uniform distribution", *Theoretical Computer Science*, 43, 1986, pp. 169-188.

[KL]   R. Karp and M. Luby, "Monte-Carlo algorithms for enumeration and reliability problems", 24th *FOCS*, 1983, pp. 56-64.

[KLM]   R. Karp, M. Luby and N. Madras, "Monte-Carlo Approximation Algorithms for Enumeration Problems", *J. of Algorithms*, Vol. 10., No. 3, Sept. 1989, pp. 429-448.

[KR]   H. Karloff and P. Raghavan, "Randomized Algorithms and Pseudorandom Numbers", 20th *STOC*, 1988, pp. 310-321.

[Lovász]   L. Lovász, "Combinatorial Problems and Exercises", North-Holland, 1979.

[Luby]   M. Luby, "A Simple Parallel Algorithm for the Maximal Independent Set Problem", *SIAM J. Computing*, Vol. 15, No. 4, November 1986, pp. 1036-1053.

[Ryser]   H. Ryser, "Combinatorial Mathematics", The Carus Mathematical Monographs, No. 14, the Mathematical Association of America, 1963.

[Valiant]   L. Valiant, "The complexity of computing the permanent", *Theoretical Computer Science*, 8, 1979, pp. 189-201.