Information Processing Letters 125 (2017) 20–25

Contents lists available at ScienceDirect

Information Processing Letters

www.elsevier.com/locate/ipl

Computing the permanent modulo a prime power

Andreas Björklund [a], Thore Husfeldt [a,b,*], Isak Lyckberg [a]

[a] Lund University, Box 118, 22100 Lund, Sweden
[b] ITU Copenhagen, Rued Langgaards Vej 7, 2300 København S, Denmark

ARTICLE INFO

Article history:
Received 16 May 2016
Received in revised form 17 April 2017
Accepted 28 April 2017
Available online 4 May 2017
Communicated by Ł. Kowalik

Keywords:
Algorithms
Graph algorithms
Randomized algorithms

ABSTRACT

We show how to compute the permanent of an $n \times n$ integer matrix modulo p^k in time $n^{k+O(1)}$ if $p = 2$ and in time $2^n/\exp\{\Omega(\gamma^2 n/p \log p)\}$ if p is an odd prime with $kp < n$, where $\gamma = 1 - kp/n$. Our algorithms are based on Ryser's formula, a randomized algorithm of Bax and Franklin, and exponential-space tabulation.

Using the Chinese remainder theorem, we conclude that for each $\delta > 0$ we can compute the permanent of an $n \times n$ integer matrix in time $2^n/\exp\{\Omega(\delta^2 n/\beta^{1/(1-\delta)} \log \beta)\}$, provided there exists a real number β such that $|\text{per}\, A| \le \beta^n$ and $\beta \le (\frac{1}{44}\delta n)^{1-\delta}$.

1. Introduction

The permanent of an $n \times n$-matrix $A = (a_{ij})$ is defined as

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

where the sum is over all permutations σ of the elements $1, \dots, n$. From the definition, per A can be computed in $O(n!n)$ arithmetic operations. Using Ryser's classic formula [9], per A can be computed in $O(2^n n)$ arithmetic operations. More recently it was shown that when the entries of A are $n^{O(1)}$-bit integers, then per A can be computed in time $2^n/\exp\{\Omega(\sqrt{n/\log n})\}$ [4].

These exponential running times seem particularly disappointing when compared to the polynomial-time algorithms for computing the determinant. This discrepancy was famously explained by Valiant's seminal result [12], which showed that the permanent is hard for the complexity class #P.

We present here some improved algorithms for computing per A modulo a prime power.

2. Results

Theorem 1. *Given an $n \times n$ integer matrix A and a positive integer k, the value* per A mod 2^k *can be computed in time $n^{k+O(1)}$ and $n^{O(1)}$ space.*

This result is established by Algorithm A in Section 4. This improves Valiant's algorithm [12] for per A mod 2^k, which runs in time $O(n^{4k-3})$. It is crucial here that computation is performed modulo a power of 2: There is little hope of finding, say, an algorithm for per A mod 3^k in time $n^{O(k)}$, since already the computation of per A mod 3 requires time $\exp(\Omega(n))$ under the randomized exponential time hypothesis [6].

Instead, for larger primes $p > 2$, we present an algorithm for per A mod p^k with running time $O\big((2-\epsilon_p)^n\big)$, where ϵ_p is positive and depends on p:

* Corresponding author at: Lund University, Box 118, 22100 Lund, Sweden.

E-mail address: thore.husfeldt@cs.lth.se (T. Husfeldt).

http://dx.doi.org/10.1016/j.ipl.2017.04.015
0020-0190/© 2017 Published by Elsevier B.V.

Theorem 2. *Given an $n \times n$ integer matrix A, a positive integer k, and a prime p such that $kp < n$, the value* per A mod p^k *can be computed in time within a polynomial factor of*

$$2^n \Big/ \exp\left\{\Omega(\gamma^2 n/p \log p)\right\} ,$$

where $\gamma = 1 - kp/n$.

In a setting where the product kp can be bounded away from n, say $kp \le \frac{99}{100}n$ for n sufficiently large, the term γ^2 can be absorbed in the Ω notation for a cleaner bound. This result is established by Algorithm B in Section 5.

Theorem 3 can be applied to permanents whose value is known to be small:

Theorem 3. *Given $\delta > 0$, an $n \times n$ matrix A of integers, and a real number $\beta \le (\frac{1}{44}\delta n)^{1-\delta}$ such that $|\operatorname{per} A| \le \beta^n$, the value* per A *can be computed in time within a polynomial factor of*

$$2^n \Big/ \exp\left\{\Omega(\delta^2 n/\beta^{1/(1-\delta)} \log \beta)\right\} .$$

In particular, if β is a constant, the bound can be given as $O\big((2-\epsilon_\beta)^n\big)$. This result is established by Algorithm C in Section 6.

An interesting special case is when the entries of A are restricted to $\{0,1\}$. Then the permanent equals the number of perfect matchings in the bipartite graph whose biadjacency matrix is A. For instance, assume that such a graph contains $\exp\{O(n)\}$ perfect matchings. Apply Theorem 3 with β constant. The resulting running time is $2^n/\exp\{\Omega(n)\}$.

We note that Theorem 3 can be applied even if no bound β is known, given that the input matrix contains only nonnegative integers. For such matrices, a celebrated randomized algorithm by Jerrum, Sinclair, and Vigoda [7] computes for given $\epsilon > 0$ in time polynomial in n and $1/\epsilon$ a value b such that $\Pr\big((1-\epsilon)\operatorname{per} A \le b \le (1+\epsilon)\operatorname{per} A\big) \ge \frac{1}{2}$. We can then take $\beta = b^{1/n}$, which is only a factor $(1+\epsilon)^{1/n}$ off the best possible bound. Provided that $\operatorname{per} A \le n^n$, the size restriction on β applies for all $\delta > 0$ and n sufficiently large, so we can apply Theorem 3.

An algorithm by Cygan and Pilipczuk [5] computes the permanent in time

$$2^n/\exp\{\Omega(n/d)\} ,$$

where d is the average number of nonzero entries per row. Their algorithm requires no bound on the size of the permanent. We can compare Theorem 3 to the result of [5] by looking at matrices over $\{-1,0,1\}$ with at most d nonzero entries per row. For such matrices, we have $|\operatorname{per} A| \le \prod_{i=1}^n \sum_{j=1}^n |a_{ij}| \le d^n$, so that Theorem 3 applies with $\beta = d$ for the weaker bound $2^n/\exp\{\Omega(n/d^{1/(1-\delta)} \log d)\}$. On the other hand, Theorem 3 outperforms [5] on families of matrices with many nonzero entries but small permanents. For instance, consider an $n \times n$ matrix A over $\{-1,0,1\}$ constructed by taking d nonzero random entries per row and picking the sign on each 1 uniformly at random. It is known that $|\operatorname{per} A| \le (\lambda_{\max})^n$, where $\lambda_{\max}$ is the spectral norm of the matrix A [1, Sec. 2]. By the Bai–Yin theorem [2, Thm. 2], the spectral norm of a random matrix whose elements have mean 0 and variance σ^2 is concentrated around $2\sigma\sqrt{n}$. Since the variance of the elements in A is $\sigma^2 = d/n$, the absolute value of the permanent of A is almost surely less than $(\frac{201}{100}\sqrt{d})^n$. Now Theorem 3 with $\beta = \frac{201}{100}\sqrt{d}$ gives $2^n/\exp\{\Omega(n/d^{1/(2-\delta)} \log d)\}$ for any $\delta > 0$.

3. Preliminaries

Our starting point is *Ryser's formula* [9] for the permanent. It is based on the principle of inclusion–exclusion and can be given as follows:

$$\operatorname{per} A = (-1)^n \sum_{x \in \{0,1\}^n} (-1)^{x_1+\cdots+x_n} \prod_{i=1}^n (Ax)_i . \qquad (2)$$

We now review an idea of Bax and Franklin [3].

Lemma 4. *Let A be an $n \times n$ integer matrix. Then for every vector $r \in \mathbf{Z}^n$,*

$$\operatorname{per} A = (-1)^{n+1} \sum_{x \in \{0,1\}^n} (-1)^{x_1+\cdots+x_n} \prod_{i=1}^n (Ax+r)_i . \qquad (3)$$

Proof. Define the matrix $A' \in \mathbf{Z}^{n+1 \times n+1}$ as

$$A' = \begin{pmatrix} a_{11} & \cdots & a_{1n} & r_1 \\ \vdots & \ddots & \vdots & \vdots \\ a_{n1} & \cdots & a_{nn} & r_n \\ 0 & \cdots & 0 & 1 \end{pmatrix} .$$

First, we observe $\operatorname{per} A = \operatorname{per} A'$, because in the Laplace expansion of the permanent of A' along the last row, all terms vanish except $a'_{n+1,n+1} \operatorname{per} A = 1 \cdot \operatorname{per} A$.

Now consider evaluating $\operatorname{per} A'$ with Ryser's formula (2). The factor

$$(A'x)_{n+1} = 0 \cdot x_1 + \cdots + 0 \cdot x_n + 1 \cdot x_{n+1} = x_{n+1}$$

vanishes unless $x_{n+1} = 1$. Thus, we can restrict our attention to vectors of the form $x' = (x_1, \ldots, x_n, 1)$. For such a vector, we have $A'x' = A(x_1, \ldots, x_n) + r$. Ryser's formula now gives

$$\operatorname{per} A' = (-1)^{n+1} \sum_{x \in \{0,1\}^n} (-1)^{x_1+\cdots+x_n} \prod_{i=1}^n (Ax+r)_i . \qquad \square$$

We turn to modular computation. Fix a positive integer k and let p be a prime. Let $\mathrm{GF}(p)$ denote the finite field of order p. Let X be the set of vectors $x \in \{0,1\}^n$ such that the vector $Ax + r$ has fewer than k zeros in $\mathrm{GF}(p)$. The crucial observation is that we can restrict our attention to X:

Lemma 5. *Let A be an $n \times n$ integer matrix. Then,*

$$\operatorname{per} A = (-1)^{n+1} \sum_{x \in X} (-1)^{x_1+\cdots+x_n} \prod_{i=1}^n (Ax+r)_i \pmod{p^k}.$$