Information and Computation 243 (2015) 133–144

Contents lists available at ScienceDirect

Information and Computation

www.elsevier.com/locate/yinco

Noncommutativity makes determinants hard ☆

Markus Bläser

Saarland University, Saarbrücken, Germany

ARTICLE INFO

Article history:
Received 15 September 2013
Available online 12 December 2014

Keywords:
Counting complexity
Determinant
Permanent
Associative algebras

ABSTRACT

We consider the complexity of computing the determinant over arbitrary finite-dimensional algebras. We first consider the case that A is fixed. In this case, we obtain the following dichotomy: If $A/\operatorname{rad} A$ is noncommutative, then computing the determinant over A is hard. "Hard" here means #P-hard over fields of characteristic 0 and $\mathsf{Mod}_p\mathsf{P}$-hard over fields of characteristic $p > 0$. If $A/\operatorname{rad} A$ is commutative and the underlying field is perfect, then we can compute the determinant over A in polynomial time.

We also consider the case when A is part of the input. Here the hardness is closely related to the nilpotency index of the commutator ideal of A.

Our work generalizes and builds upon previous papers by Arvind and Srinivasan (STOC 2010) [1] as well as Chien et al. (STOC 2011) [5].

1. Introduction

The determinant of a matrix $M = (m_{i,j}) \in k^{n\times n}$ is given by the well-known formula

$$\det M = \sum_{\sigma\in S_n} \operatorname{sgn}(\sigma) m_{1,\sigma(1)} \cdots m_{n,\sigma(n)}.$$

The determinant plays a central role in linear algebra. It can be efficiently computed; over fields, for instance, by Gaussian elimination. In fact, there are even efficient algorithms when the matrix M has entries from some commutative algebra, see [17] and the references given therein.

A related polynomial is the permanent of M, given by

$$\operatorname{per} M = \sum_{\sigma\in S_n} m_{1,\sigma(1)} \cdots m_{n,\sigma(n)}.$$

If M is $\{0,1\}$-valued, then $\operatorname{per} M$ is the number of perfect matchings of the bipartite graph with adjacency matrix M. While the determinant is easy over commutative algebras, the permanent is hard already over the rationals. Valiant [20] showed that evaluating the $\{0,1\}$-permanent over the rationals is at least as hard as counting the number of satisfying assignments of a formula in 3-CNF.

Since the determinant and the permanent have similar formulas, it is tempting to try to modify algorithms for the determinant and use them to compute the permanent. Godsil and Gutman [13] used the determinant to approximate the

☆ Work supported by DFG grant BL 511/10-1 and by the Indo-German Max-Planck Center for Computer Science (IMPECS). A preliminary version was presented at the 40th Int. Colloquium on Algorithms, Languages, and Programming (ICALP 2013).

E-mail address: mblaeser@cs.uni-saarland.de.

http://dx.doi.org/10.1016/j.ic.2014.12.010
0890-5401/

permanent: For simplicity, assume that the matrix M is $\{0,1\}$-valued. With probability $1/2$, we now replace a 1 in M by -1, where all coin flips are done independently. Let $\hat{M}$ be the resulting random matrix. In expectation, the square of the determinant of $\hat{M}$ is the permanent, that is,

$$\mathrm{E}\big[\det(\hat{M})^2\big] = \operatorname{per} M.$$

This is quite easy to see, $\det(\hat{M})^2$ consists of terms of the form

$$\operatorname{sgn}(\sigma)\hat{m}_{1,\sigma(1)}\cdots\hat{m}_{n,\sigma(n)}\cdot\operatorname{sgn}(\tau)\hat{m}_{1,\tau(1)}\cdots\hat{m}_{n,\tau(n)}.$$

If $\sigma = \tau$, then this is 1 iff $m_{1,\sigma(1)}\cdots m_{1,\sigma(n)} = 1$. If $\sigma \neq \tau$, then there is at least one variable that only appears once in the term. Its expected value is 0, therefore, the whole term evaluates to 0 in expectation. Since the determinant is easy but the permanent is hard, there must be a catch, namely, the variance is huge. This means that we need a lot of samples to get a good approximation of the permanent with high probability.

Karmarkar et al. [16] showed how to lower the variance by extending the underlying field to the complex numbers. Now a 1 is replaced by one of 1, -1, i, -1 with equal probability. Chien et al. [7], building upon the work by Barvinok [2], generalized this approach to so-called Clifford algebras. In particular it follows from their result that if one could compute the determinant of an $n \times n$-matrix the entries of which are themselves matrices of size $cn \times cn$ for some constant c, then there is a fully polynomial time randomized approximation scheme for the permanent of $\{0,1\}$-matrices. See [18] for further results in this direction. (Of course, there is a fully polynomial randomized approximation scheme based on Markov chains, see [15]. However, if we could evaluate noncommutative determinants as fast as commutative ones, then we would get much faster approximation schemes.)

Therefore, it is important to understand the complexity of the determinant over arbitrary finite-dimensional algebras, especially over noncommutative ones, and not only over fields or commutative algebras. The first to study this problem was Nisan [19]. He proved an exponential lower bound for the size of an algebraic branching program for computing the determinant over the free noncommutative algebra $k\langle X_{i,j}\rangle$. While the lower bound is strong, the setting is limited, because it only applies to a restricted circuit model and only to a very "powerful" algebra. Chien and Sinclair [6] extended these bounds to a wide range of "concrete" algebras by analyzing their polynomial identities, for instance to matrix algebras and the Hamiltonian quaternions, albeit only in the algebraic branching program model.

Recently Arvind and Srinivasan [1] showed that the noncommutative determinant cannot have small circuits unless the permanent has small circuits. Under the assumption that the permanent is hard, this result shows that the Godsil-Gutman approach is not suited to get a fully polynomial time randomized approximation scheme. Chien et al. [5] made further progress by proving the #P-hardness and $\mathsf{Mod}_p\mathsf{P}$-hardness of the determinant for odd p for large classes of algebras. Essentially they showed that Valiant's hardness proof of the permanent can be modified in such a way that it works for the determinant over theses algebras, too.

The fundamental question behind these results is: Which properties of the algebra makes the determinant hard? In this work, we prove that this is exactly noncommutativity.

1.1. A crash course on the structure of algebras

An associative algebra A over some field k is a k-vector space together with a bilinear mapping $\cdot : A \times A \to A$, the multiplication in A. Multiplication is associative and distributes over addition. If $\lambda \in k$, then $\lambda(x \cdot y) = (\lambda x) \cdot y = x \cdot (\lambda y)$ for all $x, y \in A$. We will always assume that A is finite-dimensional (as a vector space) and contains a unit element, which we denote by 1.

A left (right, twosided) *ideal* of an algebra is a vector space that is closed under multiplication with arbitrary elements of A from the left (right, both sides). If S is a subset of A, then the left (right, twosided) ideal of A generated by S is the intersection of all left (right, twosided) ideals that contain S. Alternatively, it can be defined as the linear span generated by all elements xs (sy, xsy) with $x, y \in A$ and $s \in S$.

If I is a twosided ideal, then the quotient space A/I becomes an algebra in the natural way by setting $(a + I)(b + I) = ab + I$. It is easy to check that this is well-defined, since I is a twosided ideal.

A left (right, twosided) ideal I is called *nilpotent*, if $I^s = \{0\}$ for some positive integer s. The nilpotency index of I is the smallest s such that $I^s = \{0\}$. If there is no such s, then the index is infinite.

The sum of all nilpotent left ideals of A is a nilpotent twosided ideal, which contains every nilpotent right ideal of A. This twosided ideal is called the *radical* of A and is denoted by $\operatorname{rad} A$. The quotient algebra $A/\operatorname{rad} A$ contains no nilpotent ideals other than the zero ideal. Since A is finite dimensional, we can alternatively define the radical of A as the intersection of all maximal twosided ideals. An ideal is maximal if it is not contained in any other ideal and is not equal to A. The radical of A is contained in every maximal twosided ideal of A. The algebras A and $A/\operatorname{rad} A$ have the same number of maximal twosided ideals.

We call an algebra A *semisimple*, if $\operatorname{rad} A = \{0\}$. By the above fact, $A/\operatorname{rad} A$ is semisimple. An algebra A is called *simple*, if there are no twosided ideals in A except the zero ideal and A itself. An algebra D is called a *division algebra*, if $D^\times = D \setminus \{0\}$. Here $D^\times$ is the set of all invertible elements in D. An algebra A is called *local*, if $A/\operatorname{rad} A$ is a division algebra.

The following fundamental theorem describes the structure of semisimple algebras.