3B2v7.51c GML4.3.1 YJCOM : 768 Prod.Type:FTP pp.1–29(col.fig.:Nil) ED:Kulkarni PAGN: tnn SCAN: Nil

ARTICLE IN PRESS

Available online at www.sciencedirect.com

SCIENCE DIRECT®

Journal of COMPLEXITY

ELSEVIER

Journal of Complexity ▮ (▮▮▮▮) ▮▮▮–▮▮▮

http://www.elsevier.com/locate/jco

Improved algorithms for computing determinants and resultants

Ioannis Z. Emiris[a,1] and Victor Y. Pan[b,*,2]

[a] *Department of Informatics & Telecommunications, National Kapodistrian University of Athens, 15771 Greece*

[b] *Mathematics and Computer Science Department, CUNY, Bronx, NY 10566, USA*

Received 10 January 2003; accepted 15 March 2004

Abstract

Our first contribution is a substantial acceleration of randomized computation of scalar, univariate, and multivariate matrix determinants, in terms of the output-sensitive bit operation complexity bounds, including computation modulo a product of random primes from a fixed range. This acceleration is dramatic in a critical application, namely solving polynomial systems and related studies, via computing the resultant. This is achieved by combining our techniques with the primitive-element method, which leads to an effective implicit representation of the roots. We systematically examine quotient formulae of Sylvester-type resultant matrices, including matrix polynomials and the *u*-resultant. We reduce the known bit operation complexity bounds by almost an order of magnitude, in terms of the resultant matrix dimension. Our theoretical and practical improvements cover the highly important cases of sparse and degenerate systems.

Keywords: Computer algebra; Randomized algorithms; Matrix determinant; Bit complexity; Structured matrix; Polynomial system solving

*Corresponding author.

E-mail addresses: emiris@di.uoa.gr (I.Z. Emiris), vpan@lehman.cuny.edu (V.Y. Pan).

[1] Partially supported by Project 70/4/6452 of the Research Council of the National Kapodistrian University of Athens, Greece.

[2] Partially supported by NSF Grant CCR 9732206 and PSC CUNY Awards 66383-0032, 64406-0033, and 65393-0034.

0885-064X/$ - see front matter © 2004 Published by Elsevier Inc.
doi:10.1016/j.jco.2004.03.003

1. Introduction

1.1. Our subjects and techniques

The classical problem of computing matrix determinants witnessed dramatic recent progress [EGV00,KV01,Sto03]. Our first step is to adapt the algorithm of Kaltofen and Villard [KV01] to the case of modular computation, where the input is generally a multivariate matrix. This shall be an important block for our algorithms, based on Chinese remaindering. We also contribute a new randomized *output-sensitive* algorithm, which improves the known bounds when the output is away from the available upper bounds. Such upper bounds (like Hadamard's) are notorious for being excessively high in general, as discussed later. These results are of independent interest, but become more important because of their application to polynomial system solving by the method of resultants and their matrices. Our focus on the block Wiedemann approach of Kaltofen and Villard [KV01] is motivated by this application where it is currently superior over the others in [EGV00,Sto03].

Computing the resultant R of a polynomial system, including the cases of scalar R and univariate or multivariate polynomial $R(x)$ or $R(u_1, \ldots, u_n)$ is a central practical and theoretical problem in the field of symbolic computation, leading to efficient methods for solving 0-dimensional systems, quantifier elimination, and deciding the theory of the reals. Our practical motivation is the real-time solution of systems in CAD, vision or robotics (which may give rise to matrices with dimension in the hundreds or even higher). Equally useful is the computation of the resultant polynomial, e.g., in modeling applications where an implicit representation of a curve or surface is obtained from the (perturbed) resultant, even in the presence of base points [DE01,MC93]. Resultant values and signs also capture important tests in computational geometry, including the case of infinitesimal perturbations of the input. These diverse applications are discussed in [BEPP99,Can88,CLO98,DE01, Man93].

A polynomial system has a solution iff $R = 0$, whereas computing isolated roots of the system can be reduced to factoring the polynomial $R(u_1, \ldots, u_n)$ into the product of linear factors. Moreover, solving the univariate polynomial equation $R(x) = 0$ projects all common isolated roots to one of their coordinates and is critical in several applications. Thus we arrive at our task of computing the resultants $R, R(x)$, and $R(u_1, \ldots, u_n)$, and in our approach we largely unify this computation. For computing the resultant, we exploit a recent major advance in the field, which generalizes the Macaulay formula to the toric case by expressing the (toric) resultant by a quotient formula $\det M/\det S$, where M is a Sylvester-type resultant matrix of a polynomial system, and S is one of its submatrices [DA02]. We also apply our methods to the perturbations for handling degenerate systems [Can90,DE01], as well as to Bézout-type matrices.

We focus on *bit operation complexity*; the computational model is the (Boolean, or logarithmic-cost) *random access machine*, unless otherwise stated [AHU74,BCS97]. Section 7 examines briefly straight-line programs and linear-complexity models, which we apply to capture a specific notion of sparseness. Our estimates do not

include the cost of generating the random quantities required in the algorithms, but this complexity is practically dominated at other stages. Randomized algorithms shall be of two types, namely *Las Vegas* or *Monte Carlo* depending on whether the estimated computational cost covers or not the cost of certifying correctness of the output. In the latter case, randomization may lead to a superset of the roots; in system solving applications, certification is very fast since it amounts to evaluating the input polynomials at the computed root values.

Since, we deal with highly important and long and well-studied problems, our progress had to employ and combine various known advanced techniques of algebraic computing as well as our novel techniques. In particular, the latter include multivariate output sensitive version of the recent block Wiedemann approach of Kaltofen and Villard [KV01] to the structured determinant evaluation (where we incorporate the Chinese remainder algorithm, and randomized Newton's interpolation to speed up the computation in the case where the determinant nearly vanishes, that is exactly in the case we need, although our progress should also have general appeal and independent importance) and acceleration of the resultant computation by exploiting quasi-Toeplitz matrix structure, rational formulae for the toric resultant (in order to bound the output size), and evaluation/interpolation techniques [Too63] (cf. also [Ber03b]).

1.2. The known and new results

Previous work on integer determinant evaluation includes [ABM99,EGV00,-Kal02,KV01,Pan02b]; other works such as [BEPP99,PY01] focus on sign determination but do not improve upon the current record complexity bounds (although the output-sensitive approach in [BEPP99, Section 4] was instrumental in improving these bounds). The best complexities for a general scalar matrix are due to the algorithms in [EGV00,KV01]; cf. also Theorem 2.1. Let $O^*(f)$ indicate that the factors polylogarithmic in f are omitted. Then the known randomized bit cost bounds are in $O^*(m^{3+1/5}L)$, implicit in [KV01] (see Appendix A), and $O^*(m^2\gamma^{2/3}L)$, where m denotes the dimension of integer matrix M, L its maximum 2-norm or column (infinity) norm (our choice), and γ the number of arithmetic operations for multiplying M by a vector of scalars. We bound the bit complexity of computing $(\det M) \bmod s$, for s being the product of random primes, by $O^*(m\gamma q)$, where $q = \log s$, zand by $O^*(m^{5/3}L^{2/3}\gamma^{2/3}q^{1/3})$. The recent works [Kal02] (which yields a bound of the order of m^3q) and [Sto03] do not apply efficiently to the quasi-Toeplitz structure and thus are not sufficient to support our improvements. In particular, [Sto03] considers univariate matrices and proposes a Las Vegas algorithm for their determinant computation with arithmetic complexity in $O^*(\mu d_1)$, where $O(\mu)$ bounds the arithmetic complexity of matrix multiplication and d_1 bounds the degree of the entries. Then again our estimates using structure are superior for resultant matrices. Similar comments apply to the algorithm in [EGV00], which yields the bound in $O^*(m^3L^2)$ for an average integer matrix but does not perform much faster for structured matrices.

For $s>2|\det M|$ we can immediately recover $\det M$ from its value modulo s, so our bit cost bounds turn into the new bit cost record output sensitive bounds for computing $\det M$. We also derive the bit cost bounds in terms of L, m, s only (without γ); this improves the record bounds of Kaltofen [Kal02] by the factor of $(q/L)^{3/10}$, where $q := \lg(|\det M|+2)$ denotes the bit size of the computed integer determinant. We extend all our results and new record bit cost estimates to uni- and multivariate matrices.

Regarding resultant operations, we reduce bit complexity by almost an order of magnitude in terms of the resultant matrix dimension, which is the largest quantity involved in the bounds. Roughly speaking, the bit complexity of existing approaches is proportional to m^3. These results, unbeaten for years, remained a clear challenge to the algorithm designers. In the present paper, we finally make a decisive step forward.

For evaluating a scalar resultant, the existing approaches exploit matrix structure [CKL89,EP02b] to arrive at bit complexities in $O^*(m^3n^2DL)$, where n is the number of eliminated variables and D denotes the total degree of the resultant R thought of as a polynomial in the input coefficients. We derive the bound $O^*(m^2nDL)$, where L is the bit length of these coefficients when specialized. Alternatively, a Monte Carlo bound depending on the bit size q of the resultant's value is $O^*(m^2nq)$. Analogous improvements are obtained in the harder and highly important case where the input degenerates and a perturbation is applied. This improves upon available approaches. We mention only one of those, because it uses a different and more direct method: The problem is reduced to computing a characteristic polynomial with Monte Carlo bit complexity proportional to $m^{3+1/5}$, up to polylog factors for a general $m\times m$ matrix [Pan02b]. Even an improvement to m^3 would not beat our estimate in the quasi-Toeplitz case.

For expanding a univariate resultant $R(x)$, the algorithm of [Man93] is based on Kronecker canonical forms. It yields a bit cost bound in $O((m+\deg(\det M(x)))^3)$, under the condition that M is nonsingular as a matrix polynomial in x. Otherwise, the algorithm has complexity proportional to m^4. The fastest general algorithms have bit complexity in $O^*(m^3n\deg(\det M(x)))$, typically relying on interpolation and matrix structure [CKL89,EP02b]. Our algorithms support a bit cost bound, for arbitrary inputs, in $O^*(m^2nDVL)$ where V is the actual degree of the univariate resultant. They also dramatically improve the known algorithms in the case of perturbed resultants. Similar improvements are obtained for expanding the n-variate u-resultant $R(u)$. This task has bit complexity proportional to m^2h, where h denotes the number of nonzero terms in $R(u)$. Previous methods have complexity cubic in m [CKL89,EP02b,Man93].

These results on resultant evaluation and expansion, coupled with the quotient formula of the (toric) resultant, immediately yield an improvement on the bit complexity of numerically approximating all isolated roots of a well-constrained algebraic system to a prescribed precision. Existing methods based on resultant formulations have *arithmetic* complexity cubic in m [CKL89,EP02b,Mou98]. We

employ the primitive-element method of Canny [Can88] (also known as rational univariate representation [Rou99]) leading to an implicit representation of the output and, by combining all our techniques, we obtain *bit operation complexity* bounds which are roughly cubic in m. The same complexity bounds hold when we wish to approximate the root projections to a single coordinate axis. In this context, all real roots are isolated with complexity quadratic in m multiplied by the fourth power of the resultant polynomial's degree.

The results of this paper appeared in preliminary form in [EP02a,EP03].

1.3. Paper organization

Besides the previous section, existing work is also covered in the sequel, in particular in Section 2 and the appendix. Sections 2 and 3 focus on computing the determinants of integer and polynomial matrices and improve the known complexity bounds for structured matrices and in the cases of the computation modulo a product of random primes and output-sensitive algorithms. Section 4 formalizes the notion of resultant matrices by supplying the necessary background information, and Section 5 accelerates their evaluation. Section 6 improves upon the existing algorithms for expanding univariate and multivariate resultants, and Section 7 does the same for polynomial system solving. The appendix provides some background to make the paper self-sufficient. Appendix A presents the main approach for integer matrix determinant computation, and Appendix B discusses random prime number generation.

2. Multivariate determinants preliminaries

This section introduces notation and details of the main algorithms for determinant evaluation on which our improvements shall be based. Our second base for improvement, concerning rational formulae for the resultant, shall be introduced later.

The algorithms for determinant evaluation use *block Wiedemann's algorithm with structured preconditioning* [Cop94,KP91,KS91,KV01,Wie86]; see the appendix. Our complexity bounds rely on fast univariate polynomial multiplication and/or on the fast Fourier transform (FFT), which is truly advantageous only for large inputs; otherwise, the classical or the Karatsuba divide-and-conquer algorithm may be preferable. On the other hand, none of the stated bounds relies on fast matrix multiplication, i.e., methods with complexity smaller than cubic in the dimension, which are hard to exploit in practice. Still, both the existing and our methods can be coupled with fast matrix multiplication in order to reduce the asymptotic bounds.

Let us fix notation for the entire paper. We assume integer matrices or matrix polynomials. We also assume Euclidean norms $||\sum_i c_i x^i||_2 = ||(c_i)_i||_2 = (\sum_i ||c_i||^2)^{1/2}$ for polynomials and vectors, and the consistent 2-norm $||(M_{i,j})_{i,j}||_2 = \max_{v:||v||=1}\{||Mv||_2\}$ for matrices, so

$$||M||^i \geqslant ||M^i||. \tag{1}$$

Here $||t|| = |t|$ for a scalar t. Alternatively, we may use the maximum or infinity norm for polynomials and vectors with $||\sum_i c_i x^i||_\infty = ||(c_i)_i||_\infty = \max_i\{|c_i|\}$ and the consistent column norm for matrices $||M||_\infty = \max_{v:||v||=1}\{||M_{i,j}v||_\infty\} = \max_i \sum_j |M_{i,j}|$. Yet another possibility is the 1-norm for polynomials and vectors, and the consistent row norm for matrices.

Let $\gamma(M)$ or γ denote the number of arithmetic operations required for multiplying a scalar vector by the matrix M. Unless otherwise stated, it suffices that γ bounds the complexity of vector multiplication on one side of the matrix. For a matrix polynomial $M = M(x)$, x denoting a variable or a set of variables, let $\gamma(M)$ denote the maximum $\gamma(M(\alpha))$ over the set of all values α of the variable(s) x. For an $m \times m$ matrix M, $m \leqslant \gamma \leqslant 2m^2 - m$ and $\gamma = o(m^2)$ for sparse and/or structured matrices.

We next present the known randomized algorithms and bit complexity bounds.

Theorem 2.1 (Pan [Pan02b, Theorem 5.1], cf. Kaltofen [KV01]). (a) *Let M be an $m \times m$ matrix whose entries are polynomials in k variables with respective degrees less than $d_1, \ldots, d_k, k \geqslant 0$, and with integer coefficients. Let $L := \log||M||, \Delta := d_1 \cdots d_k,\ g := 2/(k^2 + 4k + 5)$. Then* det M *can be computed by a Las Vegas algorithm with bit operation complexity bounded by*

$$O^*(m^{k+2+(k+3)g}\Delta L)$$

and

$$O^*(m^{k+2}\gamma^{2/(k+3)}\Delta L).$$

In particular, for scalar determinants we have the bounds $O^*(m^{16/5}L)$ and $O^*(m^2\gamma^{2/3}L)$, whereas for univariate determinants these bounds become $O^*(m^{19/5}d_1L)$ and $O^*(m^3\gamma^{1/2}d_1L)$.

(b) *For any fixed $\tau > 0$, the algorithm involves $O(m^2 \log m + Lm)$ random bits to ensure the upper bound τ on the failure probability. The generation of these bits is covered by the cost bounds in part* (a).

Let us supply some details for deriving the above cost bounds, which are instrumental in extending them below. Theorem 2.1 is supported by the algorithm in [Pan02b], which extends the one in [KV01, Section 3], outlined in Appendices A and B. The algorithm computes det M modulo sufficiently many distinct random primes $s_1, \ldots, s_u$ and then recovers det M by applying the *Chinese remaindering algorithm* (abbreviated as CRA); cf. [vzGG03,Zip93]. Under the requirement that $\log s > 1 + Lm > 1 + \log|\det M|$ for $s = s_1 \cdots s_u$, the algorithm either fails with a small probability or ensures computing the correct output at the claimed bit cost, dominated at the evaluation stage (see some further comments in Appendix A).

The bit cost of performing the entire algorithm is dominated by the bounds $O^*(m^3r^{k+1}\Delta L)$, $O^*((m^3/r)(m/t)^{k+1}\Delta L)$, and

$$O^*(m^{k+2}t^2\Delta L), \tag{2}$$

achieved at stages 2.2, 2.3, and 3, respectively (see Appendix B). We are based on (1) and on the simple observation that $\Delta(M^h) \leqslant \Delta(M)h^k$, for any positive $h \in \mathbb{Z}$.

Now, we consider the three bounds above for the parameters r and t of our choice, which have to satisfy $1 \leqslant r \leqslant m/t$, $m \geqslant t \geqslant 1$ since they define baby steps/giant steps and blocking in Wiedemann's algorithm, respectively. The overall cost bound $O^*(m^{k+2+(k+3)g}\Delta L)$ is obtained when the three summands are made asymptotically equal. The first two summands are equal for $r := (m/t)^{(k+1)/(k+2)}$, and the last two for $t := m^{(k+3)/(k^2+4k+5)}$. These specifications satisfy the requirements on r and t.

The bounds in terms of γ are obtained by trivializing the baby steps/giant steps from stage 2, i.e. writing $r = 1$ trivializes stage 2. Then the bit cost at stage 2.3 is bounded by

$$O^*\left(\gamma m\left(\frac{m}{t}\right)^{k+1}\Delta L\right). \tag{3}$$

Stage 3 is as before. So the overall complexity is bounded by the sum of (2) and (3) Then for $t := \gamma^{1/(k+3)}$, we obtain the overall bound $O^*(m^{k+2}\gamma^{2/(k+3)}\Delta L)$.

3. Improved determinant computation

The estimates in the following subsections are in γ and thus decrease for structured matrices. Only the last subsection deviates from this rule and presents output-sensitive bounds for unstructured matrices. We are going to extend Theorem 2.1 and for completeness partly repeat its derivation, which formally the reader is not required to read.

3.1. Computation modulo a product of random primes for structured matrices

Our first contribution (in Theorem 3.2) is the computation modulo a product of random primes which exploits matrix structure. Let M be an $m \times m$ matrix, with $||M(x)|| = 2^L$, whose entries are polynomials in $x_1, \ldots, x_k$ with respective degrees less than $d_1, \ldots, d_k$, $k \geqslant 0$, and with integer coefficients. Let $\Delta := d_1 \cdots d_k$ and set $\Delta := 1$ for $k = 0$.

Lemma 3.1. *For a multivariate matrix $M(x)$ as above, for $l = 2$ and ∞, we have*

$$||\det M(x)||_l \leqslant ||M||_2^m \leqslant 2^{Lm}.$$

Proof. We use the Kronecker substitution, namely $x_1 \to y, x_2 \to y^{d_1+1}, x_3 \to y^{d_1d_2+1}, \ldots$. So we have $\det M(x) = \sum_{i=0}^{N-1} c_i y^i$ where y is a single variable and $N = 1 + d_1 + d_1d_2 + \cdots + d_1 \cdots d_k$. Let ω be a primitive Nth root

of unity and let

$$\Omega = (1/\sqrt{N})(\omega^{ij})_{i,j=0}^{N-1}, \quad \Omega^* = (1/\sqrt{N})(\omega^{-ij})_{i,j=0}^{N-1}$$

be the scaled unitary matrices of the forward and inverse Fourier transforms, where $||\Omega^*||_l/\sqrt{N} \leqslant ||\Omega||_2 = ||\Omega^*||_2 = 1$. Then $(\det M(\omega^i))_{i=0}^{N-1} = \sqrt{N}\Omega(c_j)_{j=0}^{N-1}$, $(c_j)_{j=0}^{N-1} = (1/\sqrt{N})\Omega^*(\det M(\omega^i))_{i=0}^{N-1}$,

$$||(c_j)_{j=0}^{N-1}||_l = (1/\sqrt{N})||\Omega^*||_l||\,(\det M(\omega^i))_{i=0}^{N-1}||_l.$$

We have $||\Omega^*||_2 = 1$, $||(\det M(\omega^i))_{i=0}^{N-1}||_2 \leqslant \max_i|\det M(\omega^i)|\sqrt{N}$, so

$$||\det M(x)||_\infty \leqslant ||\det M(x)||_2 = ||(c_j)_{j=0}^{N-1}||_2 \leqslant \max_i |\det M(\omega^i)| \leqslant ||M||_2^m = 2^{Lm}. \quad \square$$

Alternatively, we may use the Goldstein–Graham bound [GG74], namely $||\det(M_{i,j}(x))||_2 \leqslant (\prod_i \sum_j |W_{i,j}|^2)^{1/2}$, where $W_{i,j} = ||M_{i,j}||_1$.

We shall operate modulo $s = s_1 \cdots s_u$, i.e. s is the product of u random primes sampled in the range $(w, 20w]$ with

$$w \geqslant m^2 \max\left\{||M||_2, \frac{m}{\tau}\log\frac{m}{\tau}\right\} \tag{4}$$

(see Appendix B), for a fixed upper bound $\tau > 0$ on the failure probability; we choose a sufficiently large parameter $u \leqslant m$ to have $q = \log s < 1 + u\log w$. We may increase the range if we need more primes. We assume that $\log w \leqslant c\log(mL) \leqslant q \leqslant cLm\gamma/2$, for a fixed constant c. This holds if $2u \leqslant Lm\gamma/\log(mL)$, and under this condition we may still choose $s = 2(mL)^{cLm\gamma/(2\log(mL))} > 2||\det M||$, for $c\gamma > 2$, due to the lemma.

The next theorem extends the second bound of Theorem 2.1(a) (in terms of γ) to computation of $(\det M) \bmod s$ where $\log w \leqslant q = \log s \leqslant cLm\gamma/2$.

Theorem 3.2. (a) *Under the above assumptions*, $(\det M) \bmod s$ *can be computed by a Las Vegas algorithm with bit operation complexity in*

$$O^*\left(m^{k+1}\gamma^{\frac{2}{k+2}}\Delta q\right)$$

and in

$$O^*\left(m^{k+1}\left((Lm\gamma)^2 q^{k+1}\right)^{\frac{1}{k+3}}\Delta\right),$$

where $q = \log s$.

(b) *The bounds in Theorem* 2.1(b) *apply for*

$$q = \sum_{i=1}^{u} \lceil \log s_i \rceil \leqslant u + \log s.$$

The first bound is smaller iff $q \leqslant 2Lm/\gamma^{1/(k+2)}$. The cost bounds cover the generation of random primes. By Lemma 3.1, $q = mL$ bounds $\log||\det M||$. On the

other hand, for q of the order of Lm, the second bound above is exactly the second bound in Theorem 2.1 (see Corollary 3.4).

Proof and algorithm. The theorem is supported essentially by the algorithm in [Pan02b] extending [KV01, Section 3] and adjusted to computing modulo s. The overall bit complexity is specified in [Pan02b] to be in

$$O^*\left(\gamma m\left(\frac{m}{t}\right)^k \min\left\{q, \frac{Lm}{t}\right\}\Delta + t^2 m^{k+1} q\Delta\right) \tag{5}$$

for $t \in \mathbb{Z}$ chosen in $[1, m]$. The first summand corresponds to stage 2.3 of the algorithm (see Appendix A), where the baby steps/giant steps have been trivialized in order to be able to bound the complexity in terms of γ, and Lm/t is replaced by $\min\{q, Lm/t\}$. The second summand expresses the complexity of stage 3 where we replace mL by q. Other stages are dominated.

To support these bounds, we need to operate modulo s in $O^*(q)$ bit operations per arithmetic operation. This is achieved by Schönhage–Strassen's algorithm or by CRA. Practically the latter is preferred because of the considerable overhead of the former. Theoretically one may combine the two approaches to compute $\det M$ modulo several larger coprimes $s_1, \ldots, s_v$, each a product of several random primes in $[a, am)$, and then recover $\det M \bmod s$, $s = s_1 \cdots s_v$ via the CRA. The recovery's cost is dominated.

Now distinguish between the two cases above depending on whether $\min\{q, Lm/t\}$ is q or Lm/t. Pick t to make the two summands in the bound equal up to a constant factor, such that $t^{k+2} := \gamma/2$ or $t^{k+3} := cLm\gamma/(2q)$, respectively. So the claimed bit cost bounds hold. Furthermore, in the first case, the hypothesis that $q \leqslant Lm/t$ implies that $q \leqslant 2Lm/\gamma^{1/(k+2)}$. In the second case, the hypothesis that $q > Lm/t$ implies $q > 2Lm/\gamma^{1/(k+2)}$, for the chosen value of t.

In both cases, the hypothesis $t \leqslant m$ is satisfied for $\gamma \leqslant 2m^{k+2}$, where the latter bound holds true even for $k = 0$. Similarly, $1 \leqslant t$ holds because $\gamma \geqslant 2$ and $2q \leqslant cLm\gamma$ by assumption.

In extending the bounds from the scalar to the multivariate case, it is important to bound the degree in the k variables of the product of M with a vector, repeated m/t times. Initially, the vector has only scalar entries; at the end its degree in the ith variable is at most $(m/t)d_i$. Hence, the product of degrees in the final vector product is $(m/t)^k\Delta$. □

Corollary 3.3. *The above bit operation cost bounds specialize to* $O^*(m\gamma q)$ *and* $O^*(m^{5/3}L^{2/3}\gamma^{2/3}q^{1/3})$ *for scalar determinants, and to* $O^*(m^2\gamma^{2/3}d_1 q)$ *and* $O^*(m^{5/2}L^{1/2}\gamma^{1/2}q^{1/2}d_1)$ *for univariate determinants. The former bounds are superior iff* $q \leqslant 2Lm/\gamma^{1/(k+2)}$.

Corollary 3.4. *If* $s > 2||\det M||$, *Theorem* 3.2 *can be applied to the evaluation of* $\det M$ *yielding the bit operation cost bound* $O^*(m^{k+2}L\gamma^{2/(k+3)}\Delta)$, *which is the second bound of Theorem* 3.2 *for* $q = cLm$ *and for a constant* c *and thus is exactly the second bound of*

Theorem 2.1. *In particular, for scalar determinants, if* $s>2(\sqrt{m}2^L)^m$, $q>1+m(L+\log(\sqrt{m}))>2Lm/\gamma^{1/(k+2)}$, *then* $s>2||\det M||$ *and the bound becomes* $O^*(m^2L\gamma^{2/3})$.

3.2. Output sensitive bit operation complexity for structured matrices

This section employs the previous approach, based on Theorems 2.1 and 3.2 and exploiting matrix structure, with the additional goal of achieving output-sensitive improvements of the known algorithms. In our iterative algorithms for polynomial systems, we may compute $\det M$ where it nearly vanishes, and then the output sensitive approach enables a substantial speedup, because the bit cost of determinant evaluation dominates the CRA cost. This leads to an improvement upon (output insensitive) algorithms whose complexity depends on static a priori bounds on $|\det M|$, e.g. [ABM99,KV01,Pan02b]. Static bounds on $|\det M|$ usually rely on Hadamard's inequality, which excessively high in general [ABM99], and even more so when $\det M$ may nearly vanish.

Practically, as we noted, $s_1, \ldots, s_u$ should be random primes sampled from a fixed range, then the number of residues required in the CRA is roughly proportional to a bound on the bit size of the output value v. Reconstructing v by Lagrange's deterministic scheme has bit complexity quasi-linear in this bound, versus $O(q^2)$ by Newton's incremental method, with q expressing the *actual* bit size of v, namely $q=\log_2(|v|+2)$. The latter method was applied in determinantal computations in [BEPP99]. Hence Newton's CRA in [BEPP99] is an output-sensitive algorithm (cf. [MC93]). It is a Monte Carlo algorithm because it stops when the computed value may still be different from the desired correct output but only with a low probability. More precisely, for primes s_i, write $v_k=\sum_{i=0}^{k} b_i s_1\cdots s_i, b_i=(v-v_i)(s_1\cdots s_i)^{-1} \bmod s_{i+1}$. By sampling s_i uniformly from a set S, we obtain $\text{Prob}\{v_k = v \text{ if } v_k=v_{k-1}\}\geqslant 1-1/|S|$, whereas the bound $s_1\cdots s_k>2|v|$ implies deterministically that $v_k=v$.

Let us use an a posteriori Monte Carlo bound in an output-sensitive way. Recall that the s_i are random primes from $(w,20w]$. Let $q_j:=\log s_j$ and $q^{(j)}:=\sum_{i=1}^{j} q_i$. We compute $\delta_j:=(\det M) \bmod s^{(j)}$, $s^{(j)}:=s_1\cdots s_j$, $j=1,2,\ldots,h$ for the minimum h such that $\delta_{h-1}=\delta_h$. To increase the probability of success, we may require that $\delta_{h-b}=\delta_h$ for some fixed $b>1$.

Corollary 3.5. *Assume the above computation of* $\delta_1, \ldots, \delta_h$ *such that* $\delta_{h-1}=\delta_h$ *for the minimum* h. *Then* $\delta_{h-1}=\det M$ *with a probability of at least* $1-1/|S|$, *using the above notation. Our bounds in Theorem* 2.1(b) *apply to the computation of* $\delta_1, \ldots, \delta_h$ *with the term* Lm *replaced with* $q^{(h)}+L+\log(m/r)+O(1)$. *The bit operation complexity bounds in Theorem* 3.2 *apply for* $q=q^{(h)}$, *and the first of these bounds is superior iff* $q^{(h)}\leqslant 2Lm/\gamma^{1/(k+2)}$.

Proof and algorithm. Combine the algorithm of Theorem 3.2 for $q = q^{(h)}$ with Newton's CRA to compute the δ_j. Concerning the precise probabilities and the complexity of CRA, see [BEPP99,Emi98,Kal02] and the discussion above.

A delicate point in our present proof is that $q^{(h)}$ is not known until we observe that $\delta_{h-b} = \delta_h$ for $b \geqslant 1$, and the value $q^{(h)}$ affects the choice of the parameter t. If we blindly extend the algorithm supporting Theorem 3.2, we may have to recompute $(\det M) \bmod s_i$ for the current primes s_i as soon as we change t. This could increase the bit cost by the factor of the order of u (the number of primes s_i), which can be of the order of Lm. A simple way out, however, is not to change t until the new value of q is doubled versus the last time it defined t. Then the overall number of distinct t is $O(\log(Lm))$ and the increase of the bit cost by the factor of $\log(Lm)$ is immaterial under the O^* notation. □

3.3. General output-sensitive bounds

Let us now apply the algorithm in [Pan02c], extending [KV01] in the multivariate case but with baby steps/giant steps. We do not express complexity in terms of γ here, hence the usefulness of the next theorem for structured resultant matrices is limited. Nonetheless, the result is presented for completeness, since it is of independent interest and achieves record output-sensitive complexity estimates.

Theorem 3.6. *Suppose* $q \geqslant L$ *and* $q \geqslant \log(c \log(mL))$ *with the above notation* (*cf. Theorem* 2.1). *Then* $\det M$ *can be computed by a Monte Carlo algorithm with the bit operation complexity bounded by*

(a)
$$O^*(m^{k+1+(2k+5)g} L^{(k+2)g} q^{1-(k+2)g} \Delta),$$
where $g := 2/(k^2 + 4k + 5)$;

(b)
$$O^*(m^{k+1+(k+4)h} L^h q^{1-h} \Delta),$$
where $h := 2/(k^2 + 3k + 4)$;

(c)
$$O^*(m^{k+1+(k+2)f} q \Delta),$$
where $f := 2/(k^2 + 2k + 2)$. *These bounds improve Theorem* 2.1 *for* $q \leqslant Lm$, $q \leqslant Lm^{k(k+2)/(k^2+3k+3)}$, $q \leqslant Lm^{k^2 f/2}$, *respectively. The bounds in Theorem* 2.1(b) *on the failure probability and the number of random bits can be applied.*

Proof. The estimates in Section 2, namely (2) and the ensuing discussion, rely on using the bit length of the orders of $Lr, Lm/t$, and Lm at stages 2.2, 2.3, and 3 of the algorithm, respectively. Using the output sensitive CRA (with dynamic choice of primes separately for each of stages 2.2, 2.3, and 3), we may replace the bit length

bounds by q wherever $q<Lr, q<Lm/t$, and $q<Lm$ and then we may still satisfy the constraints $1\leqslant r\leqslant m/t,\ 1\leqslant t\leqslant m$. The bit complexity of stages 2.2, 2.3, and 3 becomes $O^*(m^3r^k\min\{q,rL\}\Delta)$, $O^*((m^3/r)(m/t)^k\min\{q,Lm/t\}\Delta)$, and $O^*(m^{k+1}t^2q\Delta)$, respectively. As in the proof of Corollary 3.5 we change the bound on the bit length of q only when this at least doubles it, and thus we limit the number of changes to $O(\log(Lm))$.

To prove (a), we equalize the three stages' bounds by choosing

$$r := (m/t)^{(k+1)/(k+2)}\leqslant m/t, \quad t := (m^{(2k+5)/(k+2)}L/q)^{(k+2)g/2}.$$

Then $t\leqslant m\Leftrightarrow L/q\leqslant m^k$, which holds for $k\geqslant 0, L\leqslant q$. With no extra hypotheses, we may simply replace $\min\{q,rL\}$ by rL and $\min\{q,Lm/t\}$ by Lm/t, even if $q<rL, q<Lm/t$, thus obtaining the bounds encountered in the context of Theorem 2.1 for stages 2.2 and 2.3. The first two summands are equal due to our choice of r, and the last two summands become equal for our choice of t.

For (b), if we write $r := m^2/t^{k+2}, t := m^{(k+4)h/2}(L/q)^{h/2}$, then the hypothesis $q\leqslant Lm^{k(k+2)/(k^2+3k+3)}$ is equivalent to $q\leqslant Lm/t$. Therefore, we replace $\min\{q,Lm/t\}$ by q. We also replace $\min\{q,rL\}$ by rL, which is always possible. For the chosen values of r and t, we have $r\leqslant m/t$ because $r/(m/t)=m/t^{k+1}$, which equals

$$m^{1-(k+1)(k+4)/(k^2+3k+4)}\,(q/L)^{(h/2)(k+1)}.$$

This is bounded by $m^{-kh}(q/L)^{(k+1)h/2}$. For $k=0$, we have $q=L$, so $r/(m/t)\leqslant 1$. Otherwise, $k\geqslant 1$. Then substitute $q/L\leqslant m/t$ and obtain that

$$\frac{r}{m/t}\leqslant m^{(1-k)h/2}/t^{(k+1)h/2}<1.$$

Now observe that the choice of r makes the last two summands equal. Similarly, the choice of t makes the first and third summands equal. For this specialization, $t\leqslant m$ can be established by using $tq\leqslant Lm$ and $L\leqslant q$ and $t\geqslant 1$ by using $q\leqslant mL$.

For (c), let us write $r := (m/t)^{k/(k+1)}\leqslant m/t,\ t := m^{(k+2)f/2}\leqslant m$; these choices satisfy $t,r\geqslant 1$. Then the hypothesis $q\leqslant Lm^{k^2f/2}$ is equivalent to $q\leqslant rL$. Since $r\leqslant m/t, q\leqslant rL$ we have $q\leqslant mL/t$, so we can replace $\min\{q,rL\}$ and $\min\{q,Lm/t\}$ by q. The chosen value of r makes the first two summands in the overall bound equal. The choice of t makes the last two summands equal. □

Corollary 3.7. *The theorem's bounds specialize as follows. For* $k=0$: (a) $O^*(m^3L^{4/5}q^{1/5})$, (b) $O^*(m^3L^{1/2}q^{1/2})$ [Kal02], *and* (c) $O^*(m^3q)$. *For* $k=1$: (a) $O^*(m^{3+2/5}L^{3/5}q^{2/5}d_1)$ (b) $O^*(m^{3+1/4}L^{1/4}q^{3/4}d_1)$, (c) $O^*(m^{3+1/5}qd_1)$.

Proof. $k=0\Rightarrow g=\frac{2}{5}, h=\frac{1}{2}, f=1$. $k=1\Rightarrow g=\frac{1}{5}, h=\frac{1}{4}, f=\frac{2}{5}$. Substitute these values in the bounds of Theorem 3.6. □

4. Resultant matrices

This section reviews matrix formulae for the resultant. For further information see [CKL89,CLO98]. Consider a system of polynomials

$$f_0, \ldots, f_n \in K[x_1, \ldots, x_n],$$

i.e., in n affine variables, with indeterminate coefficients; typically $K = \mathbb{Z}$ (or $K = \mathbb{Q}$). Then their *resultant* R is an irreducible polynomial in these indeterminates, whose vanishing provides a necessary and sufficient condition for the existence of common roots of the system in a specified variety. For the classical resultant of homogeneous polynomials in $n+1$ variables, this variety is the projective space $\mathbb{P}^n_{\bar{K}}$. For the *toric resultant*, it is the toric variety obtained as the closure of the torus $(\bar{K} - \{0\})^n$ under the Veronese maps of certain monomials in a projective space of dimension usually larger than n, where $\bar{K}$ stands for the algebraic closure of K. In toric elimination theory, the polynomials can have integer exponents and each polynomial is characterized by its support in $\mathbb{Z}^n$ (exponent vectors of nonzero monomials) rather than its total degree. More precisely, each polynomial is characterized by the corresponding Newton polytope Q_i, defined as the convex hull of the support.

Let us think of the f_i having *symbolic* coefficients. The resultant R is a homogeneous polynomial in the (symbolic) coefficients of each f_i, with integer coefficients. Its degree in the (symbolic) coefficients of f_i equals the generic number of common roots of the other n polynomials in the corresponding variety. This is given by Bézout's number or the mixed volume $\mathrm{MV}_j(\cdot)$, where we consider mixed volumes of sets of j polynomials in j affine variables. Equivalently, mixed volume is also seen as a function on j convex polytopes in j-dimensional Euclidean space.

Mixed volume captures the sparsity of the equations, since it depends solely on their nonzero terms and, to be more precise, on their Newton polytopes. It is more general than Bézout's number in that it reduces to the latter for n dense homogeneous polynomials in $n+1$ variables [CLO98]. Hence, the total degree of R in the input coefficients is

$$D := \sum_{i=0}^{n} \mathrm{MV}_n(f_0, \ldots, f_{i-1}, f_{i+1}, \ldots, f_n). \tag{6}$$

There are two types of *resultant matrices* used to compute the resultant itself. Those of Bézout-type are discussed at the end of Section 6 briefly. Most of this paper focuses on *Sylvester-type* matrices, specified by means of generic polynomials $g_0, \ldots, g_n$ such that the map

$$[g_0, \ldots, g_n] \mapsto [g_0, \ldots, g_n]\ M = \left[\sum_{i=0}^{n} g_i f_i\right]$$

has two properties: first, it is surjective for generic f_i and, second, the dimensions of the domain and the range are equal. Generic polynomials can be thought of as having indeterminate coefficients. Then M is an $m \times m$ Sylvester-type resultant matrix such that $R | \det M$ and $\det M \neq 0$. In the case of dense homogeneous

polynomials, this is the classical Macaulay matrix. Matrix construction is independent of the coefficient values and can thus be conducted off-line; its complexity can be asymptotically smaller than that of manipulating the matrix for system solving [CE00].

M is *quasi-Toeplitz*, i.e., its entries depend only on $\alpha - \beta$, where α, β belong to two subsets of $\mathbb{Z}^n$ which index, respectively, the rows and columns of M. The most relevant property of such matrices is that, by Emiris and Pan [EP02b, Theorem 5.6] (cf. [CKL89]), for a vector v and M both filled with scalars, computing $v^{\mathrm{T}}M$ and Mv takes $O^*(mn + n\sqrt{d})$ arithmetic operations, where d is the maximum degree of f_i in any variable. Assuming $m \geqslant \sqrt{d}$, the time complexity becomes $O^*(mn)$; in practice, we usually have $m \gg d$. Hence, by Wiedemann's algorithm with structured preconditioning and [EP02b], the arithmetic complexity of computing $\det M$ is $O^*(m^2 n)$.

We may assume, $m > n$, otherwise the polynomial system degenerates and its resultant can be defined by fewer than $n+1$ polynomials. To bound the size of m and, eventually, of the complexity of constructing M, we recall the concept of Newton polytopes. Let $\sigma \in \mathbb{R}$, $\sigma \geqslant 1$ bound the amount of scaling necessary for the minimum-volume Newton polytope to enclose the largest one. Then we have, in the worst case,

$$m = O((\sigma e)^n D)$$

and the bit complexity of constructing M is in $O^*((\sigma e)^n D k^7)$, where e is the basis of natural logarithms and k bounds the number of vertices in any Newton polytope Q_i [CE00,EC95]. In fact, we typically have $m \gg D$.

Our algorithms heavily rely on the *quotient formula* of [DA02]:

$$R = (\det M)/(\det S),$$

which extends Macaulay's classical result to the toric case, with S being a submatrix of the Sylvester-type resultant matrix M. This formula gives a general means for computing R exactly, and leads to output-sensitive bounds. When we specialize the input coefficients in K, with $K = \mathbb{Z}$, the resultant is in $\mathbb{Z}$ even though it is a ratio of two determinants. Furthermore, $s = \dim S < m = \dim M$ and $m - s = D$; $m \gg D \Rightarrow s \ll m$. The specialized entries of M are integers of length $\leqslant \ell$.

To use resultants for finding all isolated roots of a well-constrained polynomial system with specialized coefficients, there are two approaches. The first considers $n+1$ input polynomials in $x_1, \ldots, x_{n+1}$, and regards this system as overconstrained over $K = \mathbb{Z}[x_{n+1}]$, where x_{n+1} is known as the *hidden variable*. No information is available a priori on the distribution of x_{n+1} in M [Man93]. Since, the input coefficients are specialized, $R(x_{n+1})$ is a univariate polynomial with degree

$$V := \mathrm{MV}_{n+1}(f_0, \ldots, f_n)$$

in x_{n+1}. Its coefficients have bit length $D\ell$, due to the quotient formula in [DA02] and Definition (6). There is no a priori relation between V and D but it is safe to assume $\log V = O(D)$. It is always the case that $V < m$. Some examples, for the cyclic-N family, are the following, where $N = n+1$ in our notation: for $N = 4, 5, 6, 7$, we

have $V = 16, 70, 156, 924$ and, respectively, $m = 25, 147, 851$ and more than 3000. See [EC95] for the definition of the cyclic-N systems and more examples on the relation of m and V.

The second way to arrive at an overconstrained system, when one is given n equations, is by adding a linear polynomial f_0 with indeterminate coefficients $u_1, \ldots, u_n$; they play the role of the hidden variables. Then R is a polynomial in these variables, where all coefficients are specialized, and is known as the *u-resultant*. In certain cases, we shall consider the homogenized u-resultant, with homogenizing variable u_0. Hence the hidden variables are $u = (u_1, \ldots, u_n)$ or $u = (u_0, \ldots, u_n)$. The u-resultant factorizes as $R(u) = C \prod_i L_i(u)^{e_i}$, where the L_i are linear,

$$\sum_i e_i = V_0 \coloneqq \mathrm{MV}_n(f_1, \ldots, f_n)$$

and C is independent of u; notice that $V_0 < D$. The polynomial $R(u)/C$ is also known as the Chow form of $f_1, \ldots, f_n$ [CLO98]. In $R(u)$, the coefficients' bit length is bounded by $D\ell$ as well as V_0 times the bit length of the roots (for the output-sensitive bound).

Degeneracies constitute the Achilles' heel of Sylvester-type matrices, including Macaulay's matrix, because these matrices are constructed for indeterminate coefficients. They must eventually be specialized to their input values, and this may make the determinant vanish identically. In addition, it is possible that the resultant polynomial vanishes, due to the existence of an infinite number of common roots, thus giving no information on the isolated roots. This is a degenerate case, since most applications still require determining all isolated solutions. The proposed linear perturbations [Can90,DE01,MC92] yield a *projection operator* as the trailing coefficient of $\det M(\varepsilon)/\det S(\varepsilon)$, which vanishes at the isolated roots, thus allowing their computation.

Since, the resultant can be obtained from one or more determinants, our results in the preceding sections apply to its computation. In the remaining sections, we specify the relevant estimates for the bit operation complexity which support a fixed bound on the error probability; we omit repeating the straightforward extensions for the number of random bits involved.

5. Resultant evaluation

This section studies evaluation of the scalar resultant. Let L stand for the matrix entries' bit length in general. This slightly differs from $L = \log||M||$ as used before, but by at most the additive term $\log m$, so we slightly abuse the notation for simplicity. When M contains integers, then $L = \ell$, the input data length. When M is univariate with d_1 bounding the degree of each entry, then, after specializing the variable to an (integer) value of fixed bit size, $L = O(d_1 + \ell)$. In the notation of Section 2, we consider only two cases: $k = 1, d_1 \geqslant 1$ and $k = n, d_1 = \cdots = d_n = 1$.

Problem 5.1. *Compute the value of the scalar resultant R. If the resultant is a polynomial in some hidden variables, then suppose these variables have been specialized. In either case, L stands for the matrix entries' bit length.*

Lemma 5.2. *If $\gamma(A)$ denotes the arithmetic cost of multiplication of a matrix A by a vector and A' is a submatrix of A, then $\gamma(A') \leqslant \gamma(A)$.*

Let $\mu_A(X)$ and $\mu_B(X)$ represent asymptotic arithmetic and bit complexities, ignoring polylog factors, for computing X. For instance, when $X = R, R(x_{n+1})$ or $R(u)$, the respective complexities denote those for computing the specialization of resultant R, or the uni- or multi-variate resultant polynomial $R(x_{n+1})$ or $R(u)$.

In particular, $\mu_A(R) = O^*(m^2 n)$, which follows simply since $\gamma(M) = O^*(mn)$ by Emiris and Pan [EP02b]. Notice that, by using Tellegen's reversion circuit theorem ([BCS97, Theorem 13.20]) or, directly, the transposition principle ([vzGG03, Notes 12.3]), $\gamma(M)$ bounds the asymptotic cost of pre- and post-multiplication under our computational model.

Theorem 5.3. *With the above notation, a Las Vegas algorithm solves Problem* 5.1 *with bit operation complexity $\mu_B(R) = O^*(m^2 nDL)$. A Monte Carlo algorithm solves the same problem with bit operation complexity $\mu_B(R) = O^*(m^2 nq)$, where $q := \lg(||R|| + 2)$ expresses the actual bit size of the specialized resultant R.*

Proof and algorithm. The algorithm is based on the CRA. It requires evaluation of $R \bmod s_i$, for $q+1$ primes s_i of logarithmic length. The discussion of the previous section implies $q = O(\log V_{(0)} + DL)$, where $V_{(0)}$ stands for V or V_0, depending on whether we deal with the case of a hidden variable x_{n+1} or the u-resultant. But $\log V = O(D)$ and $V_0 \leqslant D$, as explained in the previous section. Therefore, $q = O(DL)$.

Apply Corollary 3.3 with primes s_i such that $\log s_i = O(\log(m \log||M||))$ for scalar matrices M and S. The bit operation cost of each evaluation of $\det M$ is $O^*(m\gamma \log(m \log||M||)) = O^*(m^2 n \log\log||M||)$, by using the first part of the corollary for $q = \log s_i$. S is a submatrix of M, thus, by Lemma 5.2, its determinant is computed within the same complexity as M. So the bit cost of the evaluation phase is $O^*(m^2 nq)$.

Reconstructing the value R uses Lagrange's deterministic scheme with bit complexity quasi-linear in DL. This is dominated by the cost of evaluations since $DL = O^*(m^2 n)$. The latter equation follows from the fact that $D \ll m$, and we safely assume that L is not too large.

For the Monte Carlo version, $O(q)$ evaluations determine the overall complexity. The reconstruction phase has bit cost in $O(q^2)$ by Newton's interpolation with early termination [BEPP99]. □

The rest of this section covers linearly perturbed resultant matrices. The numerator is denoted by $M(\varepsilon) = M_0 + \varepsilon M_1$ where $\det M(\varepsilon) = \varepsilon^k D_k + \cdots + \varepsilon^m D_m$.

The bit length of the integer entries of M_1 is assumed the same as that of the input M_0, denoted L. The perturbation guarantees that $D_k \neq 0$ for some $m \geqslant k \geqslant 0$; typically, $m \gg k$. The perturbed denominator is $\det S(\varepsilon) = \varepsilon^s S_s + \cdots + \varepsilon^t S_t$, with $S_t \neq 0$ such that $s \geqslant t \geqslant 0$. By the divisibility of the (perturbed) determinants, $t \leqslant k$. The degree and coefficient length of the projection operator are bounded by the respective quantities of the resultant.

It is possible that M, S contain indeterminate(s) x_{n+1} or $u_1, \ldots, u_n$. In this section, all indeterminates are specialized such that $D_k S_t \neq 0$.

Problem 5.4. *Assume the resultant evaluates to zero. Compute the value of the trailing coefficient of the perturbed resultant, namely D_k/S_t, when both numerator and denominator matrices are perturbed with respect to $\varepsilon \to 0^+$. $D_k \neq 0$ and $S_t \neq 0$ are the trailing ε-coefficients in* $\det M(\varepsilon)$ *and* $\det S(\varepsilon)$, *respectively.*

It is known that k can be found by *binary search* in $[0, m]$ with bit complexity $m \log m \mu_A(\det M) = O^*(m \mu_A(\det M))$. [DE01, Lemma 5.1]. It is realistic to suppose $k = O(L\sqrt{m/n})$:

Lemma 5.5. *There is a Las Vegas randomized algorithm that determines k with bit operation complexity in $O^*(m^2 nDLk)$, and t in $O^*(m^2 nDLt)$. There is a Monte Carlo algorithm that determines k with bit operation complexity in $O^*(m^2 nk)$, and t in $O^*(m^2 nt)$.*

Proof and algorithm. The algorithm computes $O(\log k)$ univariate polynomials $\det M(\varepsilon) \bmod \varepsilon^h$, for $h = 1, 2, 4, 8, \ldots$. Each such computation has bit complexity bounded by that of computing the scalar $\det M$ multiplied by $O^*(h)$. Then it suffices to apply Theorem 5.3. The same algorithm works for computing t from matrix $S(\varepsilon)$.

In the Monte Carlo approach, first randomly specialize $\varepsilon \mapsto r_0$. Clearly, $h \leqslant k \Rightarrow (\det M(r_0)) \bmod r_0^h = 0$, whereas $h > k \Rightarrow (\det M(r_0)) \bmod r_0^h \neq 0$, with a high probability. Supposing r_0 is sufficiently small compared to $\det M(r_0)$, we deduce that $(\det M(r_0)) \bmod r_0^h$ is uniformly distributed in $[0, r_0^h)$ so that the probability of error is $1/r_0^h$. Therefore, $(\det M(r_0)) \bmod r_0^h \neq 0$ implies $h > k$, otherwise the algorithm decides that $h \leqslant k$; the test can be repeated with different random values of r_0 in order to decrease the error probability.

The algorithm performs a constant number of tests on whether $\det M(r_0) = 0 \bmod r_0^h$, for each h where $h = 1, 2, 4, 8, \ldots$ takes $O^*(\log k)$ values. Each test is performed modulo an integer s of bit size $h \leqslant k$. So we can apply the first part of Corollary 3.3, due to the bound on k. Hence the overall cost is $O(m\gamma k \lg k)$, where $\log s = \Theta(k)$. Finally, use $\gamma = O^*(mn)$.

The algorithm supports an analogous complexity bound for computing t, by replacing k by t. □

The CRA-based algorithm in [DE01] can be adapted to the quotient formula to support bit complexity $O^*((k-t)\mu_A(\det M)mL) = O^*(km^3nL)$ for Problem 5.4 provided k, t have been computed. Without knowing k, t, it is still possible to solve this problem by the approaches discussed in the next section, all of them with bit complexity proportional to m^3. We next show the solutions with a smaller bit cost and with no a priori knowledge of k, t.

Corollary 5.6. *There exists a Las Vegas algorithm for Problem* 5.4, *with bit operation complexity in* $O^*(km^2nDL)$. *There exists a Monte Carlo algorithm for the same problem, with bit operation complexity in* $O^*(m^2nq)$, *where* $q := \lg(|D_k/S_t| + 2)$.

Proof and algorithm. For the Las Vegas method, we perform the computation with symbolic ε, modulo ε^h; this adds the factor $O^*(h)$ to the complexity estimates for determinant evaluation, for each candidate $h = 1, 2, 4, 8, \ldots$, until h reaches or exceeds $k - t$, which is the degree corresponding to the trailing coefficient. There are about $\log(k-t)$ evaluations. Therefore, the bit complexity equals that of evaluating R, multiplied by $O^*(k-t)$. The claim follows by dropping the dependence on t for simplicity.

The Monte Carlo approach first computes k, t by applying the algorithm of the previous lemma separately to the numerator and denominator. We then apply the CRA with $O(q)$ evaluations. Then Corollary 3.3 gives an overall bound of $O^*(m^2nq)$ for the Monte Carlo algorithm. The latter cost dominates the complexity of computing k, t, because we may suppose $k = O^*(\log|D_k/S_t|)$. □

6. Resultant expansion

This section studies the problem of computing the resultant as a uni- or multivariate polynomial, assuming it is a nonzero polynomial.

Problem 6.1. *Assume the entries of resultant matrix* M *lie in* $\mathbb{Z}[x_{n+1}]$, *with coefficients of size* ℓ. *Then* $R(x_{n+1}) = \det M(x_{n+1})/\det S(x_{n+1})$. *Compute the polynomial* $R(x_{n+1})$, *of degree* V, *in the monomial basis.*

Corollary 6.2. *Problem* 6.1 *can be solved by a Las Vegas algorithm with bit operation complexity in* $O^*(m^2nVD\ell)$ *and by a Monte Carlo algorithm with bit operation complexity in* $O^*(m^2nVq)$, *where q expresses the actual bit size of the specializations of* $R(x_{n+1})$.

Proof and algorithm. We use Toom's standard evaluation–interpolation over $V + 1$ integer values (see [Ber03b,Too63]). Each resultant value is computed with bit operation complexity $\mu_B(R)$, given by Theorem 5.3, with $L = \ell$. The interpolation phase reduces to solving a transposed Vandermonde system, with arithmetic complexity $O^*(V)$ [Pan01], and its cost is dominated by the evaluation cost. If,

instead, we apply the Monte Carlo version of Theorem 5.3, we obtain the bound in terms of q. □

Next, consider the u-resultant. It is possible to factor out of $\det M(u)$ the minor $\det Q$, where Q is the largest square submatrix of $M(u)$ filled with scalars. In general, Q contains S as a proper submatrix.

Problem 6.3. *Compute* $R(u) = \det M(u)/\det Q$, *of total degree* V_0, *in the monomial basis.* Q *is a submatrix of* M *containing scalars, such that* $\dim Q = m - V_0$.

It is possible to apply Theorem 2.1 for $k = n + 1$ variables yielding a bound proportional at least to m^{n+3}. This bound grows as much as exponentially in n^2 because m may be exponential in n. Our bounds reduce this exponent of m.

Corollary 6.4. *If* h *is the actual support cardinality of* $R(u)$ *and* q *the maximum bit length of its specializations, then a Monte Carlo algorithm for Problem* 6.3 *has bit operation complexity in* $O^*(hnV_0\mu_B(R)) = O^*(hm^2n^2V_0q)$.

Proof and algorithm. Zippel's sparse interpolation algorithm [CKL89,Zip93] leads to an output-sensitive Monte Carlo solution: its cost depends on the *actual* support cardinality, denoted by h. The bottleneck is the evaluation stage, which requires $O^*(hnV_0)$ integer values, each computed with complexity $\mu_B(R)$ as in Theorem 5.3 by a Monte Carlo algorithm. The interpolation cost is dominated by the evaluations. □

The algorithms discussed above can be readily extended to the case that the resultant polynomial is identically zero. Then, one wishes to compute the trailing coefficient of the perturbed resultant in one or $n + 1$ hidden variable(s) when both numerator and denominator are perturbed by $\varepsilon \to 0^+$. The sought polynomial, known as a projection operator, equals $D_k(x_{n+1})/S_t(x_{n+1})$ or $D_k(u)/S_t(u)$, respectively. The precise algorithms and their complexities are based on those supporting Corollaries 5.6 and 6.4, but are omitted here for the sake of being concise.

In computing the Bézout-type resultant matrices, the bottleneck is expanding the determinant of an $(n+1)$-dimensional matrix, which contains the discrete differentials of the input polynomials f_i, $i = 0, \ldots, n$, in $2n$ variables [Mou98,Zip93]. The coefficients have bit length ℓ and, when all variables are specialized to scalars of fixed length, the bit size of the matrix entries is $L = O(dn + \ell)$, where d bounds the degree of f_i in each variable. The number of terms in the determinant is bounded by $n!d^n = O((nd/e)^n)$. Hence, a sparse interpolation method like those used above has bit complexity in $O^*((nd/e)^n(d + \ell))$ by Theorem 2.1.

7. Algebraic system solving

This section first examines methods for numerical approximation of a specific coordinate of the zeros, e.g. the $(n+1)$st coordinate, by solving the univariate

resultant $R(x_{n+1})$. The second problem under examination is to compute the zeros of a zero-dimensional polynomial system, which is equivalent to factoring the u-resultant $R(u)$. These problems signal the involvement of numerical approximation algorithms. This is typical for multivariate polynomial root-finding: symbolic techniques are used at the first stage, which can be viewed as preconditioning, followed by numerical approximation techniques at the final stage.

Let us consider the case of a hidden variable x_{n+1}. If we have computed the coefficients of $R(x_{n+1})$ with respect to the monomial basis, as in the algorithm supporting Corollary 6.2, we may compute all of its roots by a variety of available numerical methods (e.g., [McN93,McN97,Pan97,Pan02c,PH01] and their references). However, practically, one may prefer to rely on computing polynomial values rather than the coefficients; this shall be Problem 7.3.

For now, let us focus on exact-computation methods for real root isolation.

Problem 7.1. *Given a well-constrained polynomial system, isolate all real zeros of the univariate resultant polynomial $R(x_{n+1})$. This yields one real coordinate of all common roots of the system.*

Once we compute $R(x_{n+1})$ in the monomial basis, it is possible to apply any bisection or Sturm-based method for isolating all real roots. In practice, the best performance seems to be obtained from iterative methods relying on Descartes' rule of sign [CL82,RZ01].

Theorem 7.2. *A Las Vegas algorithm solves Problem* 7.1 by using $O^*(m^2V^4D^2\ell^2)$ *bit operations. A Monte Carlo algorithm for the same problem uses $O^*(m^2V^4q^2)$ bit operations, where $q := \lg(||R(x_{n+1})|| + 2)$.*

Proof and algorithm. By Corollary 6.2, we obtain $R(x_{n+1})$ in the monomial basis by a Las Vegas or a Monte Carlo algorithm with bit operation complexity in $O^*(m^2nVD\ell)$ or $O^*(m^2nVq)$, respectively. Polynomial $R(x_{n+1})$ has degree V and coefficient size bounded a priori by $D\ell$ and, in an output-sensitive manner, by q. Bisection methods based on Descartes' rule of sign isolate all real roots with bit operation complexity in $O^*(V^6k^2)$, where k bounds the coefficient size [CL82,RZ01]. Now it suffices to bound V by m in order to arrive at the sought bounds. □

We consider now the problem of isolating all roots of the univariate polynomial $R(x_{n+1})$ by relying on computing polynomial values rather than the coefficients. Values are computed by applying the algorithm of Theorem 5.3.

Problem 7.3. *When $R(x_{n+1}) = \det M(x_{n+1})/\det S(x_{n+1})$, approximate all (or some of) its zeros with output error tolerance $\tau = 2^{-b}$, without expanding it in the monomial basis. This yields one coordinate of all common roots of an algebraic system.*

Most popular polynomial root finders recursively update the current approximations to all (or one) root(s) of $R(x_{n+1})$ based on recursive evaluation of $R(x_{n+1})$ and possibly $R'(x_{n+1}) = dR/dx_{n+1}$ at these approximation points. Weierstrass' approach approximates all roots. It uses a multidimensional Newton iteration for solving Viete's system. Its variants include Durand–Kerner's, Aberth's, Farmer–Loizou's, Maehly's, and Werner's algorithms. They converge rapidly for any input polynomial as is testified by decades of extensive practical computations. At every iteration, the complexity is dominated by the cost of computing the values of R (in Durand–Kerner's) as well as R' (in the other algorithms) at $O(V)$ points, since V is the degree of $R(x_{n+1})$.

Alternatively, our resultant evaluation techniques can be combined with Jenkins–Traub's, modified Laguerre's, or modified Newton's algorithms for computing a single root. These algorithms can be recursively extended to the next roots via implicit deflation. In our next Theorem 7.4 and Corollary 7.5, we formally assume that all the above algorithms converge at least *quadratically* right from the start, so that the number of iterations by these algorithms is expressed by

$$O(\log b) \quad \text{where } \tau = 2^{-b}$$

denotes the root coordinates error tolerance.

To apply the results of previous sections, we need to scale the variable when the iterative procedure is near the root in order to make the matrix entries integral. To approximate the root within $\tau = 2^{-b}$, we have to scale the matrix by τ. Then the bit precision factor $\log|\det M|$ in our cost bounds grows by an additive term of at most mb. The l factor grows by an additive term b which is absorbed in the $O^*()$ notation.

Our assumption about at least quadratic convergence is more realistic than it may seem to be. It is proved for the input polynomials with no multiple roots (compare the effective treatment of multiple roots in [Yun76,Zen03]), provided that the initial approximations are either close enough to the roots ([McN93,McN97,PH01] and their references) or satisfy some other readily available conditions [Bin96,Kim88,-PHI98,PPI03,Sma86,ZW95]. Furthermore, immense statistics and extensive experiments show very rapid convergence of these algorithms (when properly implemented) for all input polynomials (including specially devised "hard" polynomials) even under the primitive customary choices of the initial approximations, e.g. uniformly distributed on a large circle centered at the origin [Bin96,For01].

Theorem 7.4. *Assume sufficiently close initial approximations enabling quadratic convergence of the Durand–Kerner algorithm. Then a Las Vegas iterative method, based on combining this algorithm and our determinant algorithms, solves Problem* 7.3 *for all roots by using* $O^*(m^2nDV(\ell + mb))$ *bit operations. A Monte Carlo algorithm for the same problem uses* $O^*(m^2nV(q + mb))$ *bit operations, where* $q := \lg(||R(x_{n+1})|| + 2)$. *If only a subset of* ρ *roots must be output, then the Jenkins–Traub algorithm yields a Las Vegas and a Monte Carlo method with the respective bit operation costs above multiplied by* ρ/V.

Proof and algorithm. The number of iterative steps of a rootfinder is $O(\log b)$. For evaluating R, we may use the Las Vegas algorithm supporting Theorem 5.3 with complexity $\mu_B(R)$, with $L = \ell + mb$. Then the overall cost is $O^*(V\mu_B(R)\log b)$ bit operations. This yields a bound of $O^*(m^2nDV(\ell \log b + mb))$ bit operations. Since $\ell \log b + mb = O^*(\ell + mb)$, we arrive at the claim.

We may also apply the Monte Carlo algorithm of Theorem 5.3. □

In order to use algorithms that require the derivative values, we may represent the Las Vegas algorithm of Theorem 5.3 by a *straight-line program* (SLP), i.e., without branching. SLPs define a polynomial by the black box for its (and its derivative) evaluation rather than its coefficients. Then, evaluating R' has the same cost (up to an extra factor of 4) as evaluating R [BCS97]. This is needed in all algorithms of the Weierstrass type except Durand–Kerner's, and those of Newton type, except Jenkins–Traub's. In fact, Laguerre's method requires second order derivatives, which multiplies the complexity by 16, but offers convergence of the 4th order.

Corollary 7.5. *All algorithms of the Weierstrass or Newton type mentioned above yield Las Vegas methods for solving the respective problems of the previous theorem. Their bit operation complexity in the SLP model is equal to that of the respective Las Vegas algorithm of the previous theorem.*

Now, we consider the u-resultant approach, supposing resultant matrices $M(u), S$ are available, where S contains only scalars. Recall that the factors of the u-resultant are in bijective correspondence with the vectors representing the common roots of the input algebraic system. Given its coefficients in the monomial basis, the u-resultant can be factorized by standard methods (e.g., [Zip93]). This can be rather expensive, therefore we wish to rely only on its values.

Problem 7.6. *With no expansion in the monomial basis, express the factors of* $R(u) = \det M(u)/\det S$, *where* $u = (u_0, \ldots, u_n)$.

The primitive-element method of Canny [Can88], now known as rational univariate representation [Rou99], expresses the factors of $R(u)$, and hence all affine roots, via the roots of a univariate polynomial and a set of n univariate rational expressions. The latter expressions, when specialized at the roots of the polynomial, yield the roots of the system. For completeness, we state Canny's lemma but use our own notation, the same as in the subsequent theorem.

Lemma 7.7 (Canny [Can88, Lemma 2.2]). *We consider an overconstrained system of* $n + 1$ *nonhomogeneous polynomials in n variables. There is a univariate polynomial* R^0 *of degree* V_0 *and n rational functions, such that every solution of the system not at infinity has as* ith *coordinate the value of the* ith *rational function at some root of* R^0. *All* $2n + 1$ *polynomials involved, including* R^0, *can be computed in polynomial space.*

This means of expressing the roots implicitly can be particularly useful, for instance, if they have to be compared to other algebraic numbers or if we wish to store all roots in order to compute only a small number of them upon demand later on.

Theorem 7.8. *There is a Monte Carlo algorithm that solves Problem* 7.6 *by the primitive-element method* (*i.e., by computing n rational expressions and a univariate polynomial as above*) *with bit operation complexity* $O^*(m^2n^2V_0 \log ||R(u_0)||)$, *where* $\log||R(u_0)||$ *bounds the bit length of the coefficients of the polynomials obtained by specializing the* $u_1, \ldots, u_n$ *in* $R(u)$. *This yields a description of all common roots of an algebraic system. The same results are achieved with a Las Vegas algorithm with bit complexity* $O^*(m^2n^2V_0^2D\ell)$.

Proof and algorithm. The primitive-element algorithm of Canny [Can88] (or the black-box method of Kaltofen and Villard [KT88]) reduces factoring to the computation and manipulation of $2n+1$ univariate polynomials in u_0, denoted by R^0, R_i^+, R_i^-, $i = 1, \ldots, n$. For details, see the proof of Lemma 2.2 in [Can88]. Each new polynomial is defined by specializing the variables $u_1, \ldots, u_n$ to randomly selected constants, hence yielding univariate polynomials in u_0. These polynomials have degree V_0. Their output-sensitive coefficient length is in $O(\log||R(u_0)||)$. Another bound on this length is V_0 times the length of the coefficients in $R(u)$, hence $O(V_0D\ell)$.

We may compute any univariate polynomial via V_0+1 evaluations, each by applying the Monte Carlo version of Theorem 5.3. The total bit complexity for expanding the $2n+1$ polynomials is $O^*(m^2n^2V_0q)$. Here q expresses the sum of the input and the output bit sizes which is asymptotically equal to the actual bit size of the coefficients of the univariate polynomials, expressed by $\log||R(u_0)||$.

The primitive-element algorithm computes the first subresultant of the square-free parts of $R_i^+(u_0)$ and $R_i^-(u_0)$, for $i = 1, \ldots, n$, and reduces the computed polynomials modulo $R^0(u_0)$. These steps have complexity dominated by the $O(n)$ univariate GCD computations, each with bit complexity in $O^*(V_0^2D\ell)$ if we use FFT. The total complexity of these steps is dominated by the cost of computing the $2n+1$ univariate polynomials. This procedure yields n univariate rational functions, whose specializations at the roots of $R^0(u_0)$ yield the original system's common zeros.

The Las Vegas version of this algorithm is obtained by applying the Las Vegas version of Theorem 5.3. Then, in the above bounds, we let $q = V_0D\ell$ and the claim follows. □

Corollary 7.9. *There is a Las Vegas algorithm that solves Problem* 7.6 *by numerically approximating the factors of the u-resultant with bit operation complexity* $O^*(m^2nV_0(m\beta + \log||R(u_0)||))$, *where* $\log||R(u_0)||$ *is as above and* β *bounds the output precision of the factors' coefficients. This yields a Las Vegas algorithm for algebraic system solving with the same complexity, where* β *bounds the output precision of the roots' coordinates.*

Proof and algorithm. We start with the primitive-element algorithm supporting Theorem 7.8 which computes a univariate polynomial $R^0(u_0)$ and n univariate rational expressions in u_0. When the latter are specialized at the roots of $R^0(u_0)$, they yield the factors' coefficients, i.e., the system's common zeros. The most expensive step here is solving $R^0(u_0)$ by the Monte Carlo algorithm of Theorem 7.4, in $O^*(m^2nV_0(q\log\beta + m\beta))$ bit operations. Recall that the additive factor $m\beta$ is due to the scaling of M in order to obtain the output with β bits. Moreover, the $\log\beta$ factor disappears from the final bound, as in the proof of Theorem 7.8.

Now $q = \log||R(u_0)|| + \beta$, and due to the $O^*(\cdot)$ notation, the bound stated at the claim follows. Observe that this complexity dominates that of the Monte Carlo algorithm supporting Theorem 7.8.

The factors' coefficients give the system's roots, which can be checked by substitution into the given equations. The cost of this verification step is dominated, thus yielding a certified randomized method. □

Appendix

A. Computing integer determinants by the Wiedemann–Coppersmith–Kaltofen–Villard algorithm

Let us outline the computation of integer determinants in [KV01]. The algorithm extends the ones in [Cop94,Wie86]. The algorithm in [Wie86] computes the minimum polynomial of a matrix M as the generating polynomial for the sequence of scalars $x^{\mathrm{T}}M^i y$, $i = 0, 1, \ldots$, for two random vectors x, y. Furthermore, in [Wie86] it is proved that for a preconditioned matrix MD, with random diagonal matrix D, this is likely to be equal to the characteristic polynomial $\det(\lambda I - M)$, which turns into $(-1)^m \det M$ for $\lambda = 0$. Wiedemann [Wie86] supplies all details and the probability estimates. An accelerated block version of Wiedemann's algorithm was proposed in [Cop94]. (On the related Lanczos algorithm see [GV96, Chapter 9], on its block version see [GV96, Section 9.2.6].) The algorithm was intensively studied for solving linear systems of equations (see some bibliography in [KV01]). The main result in [KV01] is the elaboration and analysis of this algorithm to yield the following theorem.

Theorem A.1 (Kaltofen and Villard [KV01, Theorem 2]). *The algorithm in* [KV01, Section 3] *computes the determinant of any matrix* $M \in \mathbb{Z}^{m\times m}$ *with* $O^*(m^{10/3}L)$ *bit operations. It utilizes* $O^*(m^{4/3} + m\log L)$ *random bits and either returns the correct determinant or it returns "failure", the latter with probability of no more than* $\frac{1}{2}$.

The main stages of the algorithm in [KV01, Section 3], supporting the above theorem, are

(1) Precondition $M \leftarrow (I + Z)M$, $Z = (z_{i,j})_{i,j}$, $z_{i,j} = 0$ unless $j = i + 1$; $z_{i,i+1}$ are random.

(2) Compute $B^{[i]} := X^{\mathrm{T}} M^i Y$ for $i = 0, 1, \ldots, \Theta(m/t)$, where X, Y are random $m \times t$ matrices, $t \in \mathbb{Z}$ is chosen in $[1, m]$; this is achieved with a baby step/giant step technique. The most costly parts of this stage are substage 2.2, of computing a power M^r of M for $1 \leqslant r \leqslant t$, and substage 2.3, of computing the Krylov sequence $X^{\mathrm{T}} M^{rk}, k = 1, 2, \ldots, s$ for $s = O(m/(tr))$.

(3) Compute the minimal matrix generator for the sequence $(B^{[i]})_{i \geqslant 0}$; this reduces to finding t linearly independent vectors in the kernel of the block Toeplitz matrix $T := (B^{[d+i-j]})_{i,j=0}^{d-1,d}, d = \lceil m/t \rceil$. If $\operatorname{rank}(T) < m$, then output "failure".

(4) Compute $\det M$, which is equal to the ratio of the leading and trailing coefficients of the determinant of this generator.

For computations at stage 3 without fast matrix multiplication, Kaltofen and Villard [KV01] only states the bound $O^*(m^2 t^3 L)$ for $k = 0$, in the notation of Section 2, and derives this bound by using Levinson–Durbin's algorithm. The resulting overall bit cost estimate is $O^*(m^{10/3} L)$. The smaller exponent $\frac{16}{5}$ in [Pan02b] relies on the bound $O^*(m^2 t^2 L)$ at stage 3, which is implied by (2) for $k = 0$ and can be achieved in at least two ways. According to Pan [Pan02c], we may apply the Morf–Bitmead–Anderson (MBA) divide-and-conquer algorithm [Pan01, Chapter 5], complemented by compression of displacement generators by Pan [Pan92, Appendix], and by the randomized preconditioning of Kaltofen and Saunders [KS91]. Alternatively, one may apply the block version of the algorithm of Brent et al. [BGY80] instead of the block MBA algorithm. The latter recipe, in the equivalent form of using Lehmer–Knuth–Schönhage's block half-gcd algorithm [Ber03a] (also supporting the same complexity bound at stage 3 and consequently the overall bound $O^*(m^{16/5} L)$) was cited in [KV01], although only in conjunction with the algorithms using fast matrix multiplication (for $k > 0$ this is the multivariate block half-gcd algorithm). So the bound $O^*(m^{16/5} L)$ is implicit in [KV01].

The block half-gcd algorithm is notorious for being impractical, unlike all other stages of the determinant computation in [KV01] (if we assume the computation modulo sufficiently many smaller primes $s_1, \ldots, s_u$ with the CRA at the end and the application of the classical algorithm at the stages of matrix multiplication). In contrast, the recent works [Pan02a,Pan03a,PMRW03] show practical promise of employing the MBA approach for the integer input matrix M having smaller displacement rank.

B. The generation of random parameters

The determinant algorithm of Appendix A involves $(2t + 1)m - 1$ random entries of the matrices Z, X, Y and u random primes $s_1, \ldots, s_u$. Let us assume a fixed upper bound $\tau > 0$ on the probability of failure in Las Vegas algorithms and of the output errors in Monte Carlo algorithms. According to Kaltofen and Villard [KV01] and Pan [Pan03b], this bound is supported under the random choice of

- the $(2t+1)m-1$ integer entries of the matrices Z, X, Y in a range $(-\beta, \beta)$, $\beta > m^2/\tau$, independently of each other, and under the uniform probability distribution in this range, and
- sufficiently many random primes $s_1, \ldots, s_u$ in a range $(w, 20w]$ for w bounded by (4) such that the product $s_1 \cdots s_u$ exceeds the value $2||\det M||_2$ for the Monte Carlo algorithms or some a priori upper bound on this value, e.g. $||M||_2^m$ for the Las Vegas algorithm. Clearly, m random primes in this range suffice even in the Las Vegas case, and
- $2mt + t^2$ random integers in the range $[0, s_i]$ for every i, $i = 1, \ldots, u$.

We obtain the upper bounds of

$$((2t+1)m-1)\big(1 + \lfloor \log(m^2/\tau) \rfloor\big) = O\big(m^2 \log(m/\tau)\big)$$

random bits for the matrices Z, X, Y, and $(2mt + t^2)m \log w = O((L + \log(m/\tau))m^2 t)$ random bits in all primes $s_1, \ldots, s_u$ and in the random integers in the ranges $[0, s_1,], \ldots, [0, s_u]$. For the Monte Carlo algorithms, the latter upper estimate for primes and the integers decreases to

$$(\log(1 + ||\det M||_2) + L + \log(m/\tau) + O(1))(mt + t^2).$$

To support the claims in our theorems, it remains to show that the bit operation cost of generating the random primes $s_1, \ldots, s_u$ is covered by our estimates for the bit cost of the evaluation of $\det M$.

We recall that $O^*(\log^3 w)$ bit operations are sufficient to generate an integer in a range $(w, 20w]$, which is a prime with a probability of at least $\frac{1}{2}$ [vzGG03, Section 18.4], and that $(\log w)^{O(1)}$ bit operations suffice to test such an integer for primality deterministically [AKS02] (or $O(\log^2 w)$ probabilistically [vzGG03, Section 18.6]). Then, Las Vegas generation of k primes is ensured with probability of at least $1 - \tau$ by using $k \log^{O(1)}(m/\tau)$ bit operations where $k \leqslant m$. Clearly, this bit cost bound is dominated for any fixed constant $\tau > 0$.

References

[ABM99] J. Abbott, M. Bronstein, T. Mulders, Fast deterministic computation of determinants of dense matrices, in: Proceedings of the ACM International Symposium on Symbolic and Algebraic Computation, 1999, pp. 197–203.

[AHU74] A.V. Aho, J.E. Hopcroft, J.D. Ullman, The Design and Analysis of Computer Algorithms, Addison-Wesley, Reading, MA, 1974.

[AKS02] M. Agrawal, N. Kayal, N. Saxena, PRIMES is in P, 2002, Preprint.

[BCS97] P. Bürgisser, M. Clausen, M.A. Shokrollahi, Algebraic Complexity Theory, Springer, Berlin, 1997.

[BEPP99] H. Brönnimann, I.Z. Emiris, V. Pan, S. Pion, Sign determination in residue number systems, Theoret. Comput. Sci. 210 (1) (1999) 173–197 (Special Issue on Real Numbers and Computers).

[Ber03a] D.J. Bernstein, Fast multiplication and its applications, 2003, Preprint, available from http://cr.yp.to/papers.html.

[Ber03b] D.J. Bernstein, Multidigit arithmetic for mathematicians, Adv. Appl. Math. (2003), to appear.

[BGY80] R.P. Brent, F.G. Gustavson, D.Y.Y. Yun, Fast solution of Toeplitz systems of equations and computation of Padé approximations, J. Algorithms 1 (1980) 259–295.

[Bin96] D. Bini, Numerical computation of polynomial zeros by means of Aberth's method, Numer. Algorithms 13 (3–4) (1996) 179–200.

[Can88] J. Canny, Some algebraic and geometric computations in PSPACE, in: Proceedings of the ACM Symposium Theory of Computing, 1988, pp. 460–467.

[Can90] J. Canny, Generalised characteristic polynomials, J. Symbolic Comput. 9 (1990) 241–250.

[CE00] J.F. Canny, I.Z. Emiris, A subdivision-based algorithm for the sparse resultant, J. ACM 47 (3) (2000) 417–451.

[CKL89] J.F. Canny, E. Kaltofen, Y. Lakshman, Solving systems of non-linear polynomial equations faster, in: Proceedings of the ACM International Symposium on Symbolic & Algebraic Computation, 1989, pp. 121–128.

[CL82] G.E. Collins, R. Loos, Real zeros of polynomials, in: B. Buchberger, G.E. Collins, R. Loos (Eds.), Computer Algebra: Symbolic and Algebraic Computation, 2nd Edition, Springer, Wien, 1982, pp. 83–94.

[CLO98] D. Cox, J. Little, D. O'Shea, Using Algebraic Geometry, Graduate Texts in Mathematics, Vol. 185, Springer, New York, 1998.

[Cop94] D. Coppersmith, Solving homogeneous linear equations over GF(2) via block Wiedemann algorithm, Math. Comp. 62 (205) (1994) 333–350.

[DA02] C. D'Andrea, Macaulay-style formulas for the sparse resultant, Trans. AMS 354 (2002) 2595–2629.

[DE01] C. D'Andrea, I.Z. Emiris, Computing sparse projection operators, in: Symbolic Computation: Solving Equations in Algebra, Geometry, and Engineering, Contemporary Mathematics, Vol. 286, AMS, Providence, RI, 2001, pp. 121–139.

[EC95] I.Z. Emiris, J.F. Canny, Efficient incremental algorithms for the sparse resultant and the mixed volume, J. Symbolic Comput. 20 (2) (1995) 117–149.

[EGV00] W. Eberly, M. Giesbrecht, G. Villard, Computing the determinant and Smith form of an integer matrix, in: Proceedings of the Annual IEEE Symposium on Foundations of Computer Science, IEEE Computer Society Press, Los Alamitos, CA, 2000, pp. 675–685.

[Emi98] I.Z. Emiris, A complete implementation for computing general dimensional convex hulls, Internat. J. Comput. Geom. Appl. 8 (2) (1998) 223–253 (Special Issue on Geometric Software).

[EP02a] I.Z. Emiris, V.Y. Pan, Improved algorithms for computing determinants and resultants, Technical Report TR-2002-019, Computer Science Dept, The Graduate Center of the City University of New York, New York, 2002. www.cs.gc.cuny.edu/tr/techreport.php?id=66.

[EP02b] I.Z. Emiris, V.Y. Pan, Symbolic and numeric methods for exploiting structure in constructing resultant matrices, J. Symbolic Comput. 33 (2002) 393–413.

[EP03] I.Z. Emiris, V.Y. Pan, Improved computation of determinants and resultants, in: Proceedings of the International Workshop on Computer Algebra in Scientific Computing (CASC), Passau, Germany, September 2003, pp. 81–94.

[For01] S. Fortune, Polynomial root finding using iterated eigenvalue computation, in : Proceedings of the ACM International Symposium on Symbolic and Algebraic Computation, 2001, pp. 121–128.

[GG74] A.J. Goldstein, R.L. Graham, A Hadamard-type bound on the coefficients of a determinant of polynomials, SIAM Rev. 16 (1974) 394–395.

[GV96] G.H. Golub, C.F. Van Loan, Matrix Computations, 3rd Edition, Johns Hopkins University Press, Baltimore, MD, 1996.

[Kal02] E. Kaltofen, An output-sensitive variant of the baby steps/giant steps determinant algorithm, in: Proceedings of the ACM International Symposium on Symbolic & Algebraic Computation, ACM Press, New York, 2002, pp. 138–144.

[Kim88] M.-H. Kim, On approximate zeros and rootfinding algorithms for complex polynomials, Math. Comp. 51 (1988) 707–719.

[KP91] E. Kaltofen, V.Y. Pan, Processor efficient parallel solution of linear systems over an abstract field, in: Proceedings of the Third Annual ACM Symposium on Parallel Algorithms and Architectures, ACM Press, New York, 1991, pp. 180–191.

[KS91] E. Kaltofen, B.D. Saunders, On Wiedemann's method for solving sparse linear systems, in: H.F. Mattson, T. Mora, T.R.N. Rao (Eds.), Proceedings of the International Symposium on Applied Algebra, Algebraic Algorithms and Error-Correcting Codes, Lecture Notes in Computer Science, Vol. 539, Springer, Berlin, 1991, pp. 29–38.

[KT88] E. Kaltofen, B.M. Trager, Computing with polynomials given black boxes for their evaluations: greatest common divisors, factorization, separation of numerators and denominators, in: Proceedings of the IEEE Symposium on Foundations of Computer Science, IEEE Computer Society Press, Los Alamitos, CA, 1988, pp. 296–305.

[KV01] E. Kaltofen, G. Villard, On the complexity of computing determinants, in: K. Shirayanagi, K. Yokoyama (Eds.), Proceedings of the Fifth Asian Symposium on Computer Mathematics, Lecture Notes in Computing, Vol. 9, World Scientific, Singapore, 2001, pp. 13–27.

[Man93] D. Manocha, Efficient algorithms for multipolynomial resultants, Comput. J. 36 (5) (1993) 485–496.

[MC92] D. Manocha, J. Canny, The implicit representation of rational parametric surfaces, J. Symbolic Comput. 13 (1992) 485–510.

[MC93] D. Manocha, J. Canny, Multipolynomial resultant algorithms, J. Symbolic Comput. 15 (2) (1993) 99–122.

[McN93] J.M. McNamee, A bibliography on roots of polynomials, J. Comput. Appl. Math. 47 (1993) 391–394.

[McN97] J.M. McNamee, A supplementary bibliography on roots of polynomials, J. Comput. Appl. Math. 78 (1) 1997, http://www.elsevier.nl/homepage/sac/cam/mcnamee/index.html.

[Mou98] B. Mourrain, Computing isolated roots by matrix methods, J. Symbolic Comput. 26 (6) (1998) 715–738 (Special Issue on Symbolic-Numeric Algebra for Polynomials).

[Pan92] V.Y. Pan, Parametrization of Newton's iteration for computations with structured matrices and applications, Comput. Math. Appl. 24 (3) (1992) 61–75.

[Pan97] V.Y. Pan, Solving a polynomial equation: some history and recent progress, SIAM Rev. 39 (2) (1997) 187–220.

[Pan01] V.Y. Pan, Structured Matrices and Polynomials: Unified Superfast Algorithms, Birkhäuser, Springer, Boston, New York, 2001.

[Pan02a] V.Y. Pan, Nearly optimal Toeplitz/Hankel computations, Technical Reports 2002-001 and 2002-017, Computer Science Dept, The Graduate Center of the City University of New York, New York, 2002.

[Pan02b] V.Y. Pan, Randomized acceleration of fundamental matrix computations, in: Proceedings of STACS, Lecture Notes in Computer Science, Springer, Berlin, 2002, pp. 215–226.

[Pan02c] V.Y. Pan, Univariate polynomials: nearly optimal algorithms for numerical factorization and rootfinding, J. Symbolic Comput. 33 (5) (2002) 701–733.

[Pan03a] V.Y. Pan, Superfast algorithm for singular integer Toeplitz/Hankel-like matrices, Technical Reports 2002-002 and 2003-004, Computer Science Dept, The Graduate Center of the City University of New York, New York, 2002 and 2003.

[Pan03b] V.Y. Pan, Randomized computation of the determinant and the minimum polynomial of an integer matrix, 2003, Preprint.

[PH01] M. Petkovic, D. Herceg, Point estimation of simultaneous methods for solving polynomial equations: a survey, J. Comput. Appl. Math. 136 (2001) 183–307.

[PHI98] M.S. Petkovic, D. Herceg, S. Illic, Safe convergence of simultaneous method for polynomials zeros, Numer. Algorithms 17 (1998) 313–331.

[PMRW03] V.Y. Pan, B. Murphy, R.E. Rosholt, X. Wang, Toeplitz and Hankel meet Hensel and Newton: nearly optimal algorithms and their safe practical acceleration, 2003, Preprint.

[PPI03] M. Petkovic, L. Petkovic, S. Illic, The guaranteed convergence of Laguerre-like methods, Comput. Math. Appl. 46 (2003) 239–251.

[PY01] V.Y. Pan, Y. Yu, Certification of numerical computation of the sign of the determinant of a matrix, Algorithmica 30 (2001) 708–724.

[Rou99] F. Rouillier, Solving zero-dimensional polynomial systems through the rational univariate representation, Appl. Algebra Engrg. Comm. Comput. 9 (5) (1999) 433–461.

[RZ01] F. Rouillier, P. Zimmermann, Efficient isolation of a polynomial real roots, Technical Report 4113, INRIA–Lorraine, 2001.

[Sma86] S. Smale, Newton's method estimates from data in one point, in: R.E. Ewing, K.I. Gross, C.F. Martin (Eds.), Merging Disciplines: New Directions in Pure, Applied and Computational Mathematics, Springer, New York, 1986, pp. 185–196.

[Sto03] A. Storjohann, High-order lifting and integrality certification, J. Symbolic Comput. 36 (3–4) (2003) 613–648 (Special Issue).

[Too63] A.L. Toom, The complexity of a scheme of functional elements realizing the multiplication of integers, Soviet Math. Dokl. 3 (1963) 714–716.

[vzGG03] J. von zur Gathen, J. Gerhard, Modern Computer Algebra, 2nd Edition, Cambridge University Press, Cambridge, UK, 2003.

[Wie86] D.H. Wiedemann, Solving sparse linear equations over finite fields, IEEE Trans. Inform. Theory 32 (1) (1986) 54–62.

[Yun76] D.Y.Y. Yun, On square-free decomposition algorithms, in: R.D. Jenks (Ed.), Proceedings of the ACM International Symposium on Symbolic and Algebraic Computation, 1976, pp. 26–35.

[Zen03] Z. Zeng, A method computing multiple roots of inexact polynomials, in: Proceedings of the ACM International Symposium on Symbolic and Algebraic Computation, 2003, pp. 266–272.

[Zip93] R. Zippel, Effective Polynomial Computation, Kluwer Academic Publishers, Boston, MA, 1993.

[ZW95] F. Zhao, D. Wang, The theory of Smale's point estimation and its application, J. Comput. Appl. Math. 60 (1995) 253–269.