Deterministic Reduction of Integer Nonsingular Linear System Solving to Matrix Multiplication

Stavros Birmpilis
Cheriton School of Computer Science
University of Waterloo
sbirmpil@uwaterloo.ca

George Labahn
Cheriton School of Computer Science
University of Waterloo
glabahn@uwaterloo.ca

Arne Storjohann
Cheriton School of Computer Science
University of Waterloo
astorjoh@uwaterloo.ca

ABSTRACT

We present a deterministic reduction to matrix multiplication for the problem of linear system solving: given as input a nonsingular $A \in \mathbb{Z}^{n\times n}$ and $b \in \mathbb{Z}^{n\times 1}$, compute $A^{-1}b$. We give an algorithm that computes the minimal integer e such that all denominators of the entries in $2^e A^{-1}$ are relatively prime to 2. Then, for a b that has entries with bitlength $O(n)$ times as large as the bitlength of entries in A, we give an algorithm to produce the 2-adic expansion of $2^e A^{-1}b$ up to a precision high enough such that $A^{-1}b$ over $\mathbb{Q}$ can be recovered using rational number reconstruction. Both e and the 2-adic expansion can be computed in $O(\mathsf{MM}(n, \log n + \log ||A||) \times (\log n)(\log n + \operatorname{loglog} ||A||))$ bit operations. Here, $||A|| = \max_{ij} |A_{ij}|$ and $\mathsf{MM}(n, d)$ is the cost to multiply together, modulo 2^d, two $n\times n$ integer matrices. Our approach is based on the previously known reductions of linear system solving to matrix multiplication which use randomization to find an integer lifting modulus X that is relatively prime to $\det A$. Here, we derandomize by first computing a permutation P, a unit upper triangular M, and a diagonal S with $\det S$ a power of two, and such that $U := APMS^{-1}$ is an integer matrix with $2 \perp \det U$. This allows our modulus X to be chosen a power of 2.

ACM Reference Format:
Stavros Birmpilis, George Labahn, and Arne Storjohann. 2019. Deterministic Reduction of Integer Nonsingular Linear System Solving to Matrix Multiplication. In *2019 Int'l Symposium on Symbolic & Algebraic Computation (ISSAC'19), July 15–18, 2019, Beijing, China.* ACM, New York, NY, USA, 9 pages. https://doi.org/10.1145/XXXXXX.XXXXXX

1 INTRODUCTION

Let $A \in \mathbb{Z}^{n\times n}$ be nonsingular. We are interested in the deterministic reduction to matrix multiplication for the problem of rational system solving, that is, computing $A^{-1}b \in \mathbb{Q}^{n\times 1}$ for a given $b \in \mathbb{Z}^{n\times 1}$. Previously known algorithms [5, 6], which have a running time within the cost reported here, require randomization to find an integer modulus X that is relatively prime to $\det A$.

In this paper we give an algorithm that computes the minimal integer e such that all denominators of the entries in $2^e A^{-1}$ are relatively prime to 2. In this case, for an integer vector b having entries with bitlength $O(n)$ times as large as the bitlength of entries in A, our algorithm also produces the 2-adic expansion of $2^e A^{-1}b$ up to a precision high enough that $A^{-1}b$ over $\mathbb{Q}$ can be recovered using rational number reconstruction. Both e and the 2-adic expansion can be computed in $O(\mathsf{MM}(n, \log n + \log ||A||) \times (\log n)(\log n + \operatorname{loglog} ||A||))$ bit operations where $||A|| = \max_{ij} |A_{ij}|$ and $\mathsf{MM}(n, d)$ is the cost to multiply together, modulo 2^d, two $n\times n$ integer matrices. Our cost analysis makes the following regularity assumptions on $\mathsf{MM}(n, d)$: super-quadricity in n ($\mathcal{H}_{\mathsf{MM}}^{2\le n}$), super-linearity in d ($\mathcal{H}_{\mathsf{MM}}^{1\le d}$), and at most quadratic in d ($\mathcal{H}_{\mathsf{MM}}^{d\le 2}$). Under these assumptions, our cost analysis is valid for any MM that satisfies $\mathsf{MM}(n, d) \in \Omega(n^2 d)$ and $\mathsf{MM}(n, d) \in O(n^3 d^2)$.

Our approach to derandomize integer linear solving is based on ideas for polynomial matrices [3]. Corresponding to A there exists a *2-decomposition*: a tuple (P, H) of matrices where P is a permutation, H is a Hermite form with powers of 2 on the diagonal, and the matrix $U := APH^{-1}$ is nonsingular with odd determinant. For example

$$\overset{AP}{\begin{bmatrix} 27 & 99 & 92 \\ 32 & 116 & -124 \\ 195 & -121 & -8 \end{bmatrix}} = \overset{U}{\begin{bmatrix} 27 & 18 & -19 \\ 32 & 21 & -37 \\ 195 & -79 & -127 \end{bmatrix}} \overset{H}{\begin{bmatrix} 1 & 1 & 12 \\ & 4 & 4 \\ & & 16 \end{bmatrix}}. \quad (1)$$

Unlike the case when working with polynomial entries, in the integer setting, we do not have good *a priori* bounds on the magnitude of entries in $U := APH^{-1}$.

In this paper we introduce the notion of a *2-massager*. This is a tuple of matrices (P, S, M) such that P is a permutation, S is in Smith form with powers of 2 on the diagonal, and M is unit upper triangular with offdiagonal entries in each column of magnitude strictly less than the corresponding diagonal entry of S. In addition, the matrix $APMS^{-1}$ is nonsingular with odd determinant and satisfies $||APMS^{-1}|| \le n||A||$. For example, for the matrix in (1) we have

$$\overset{AP}{\begin{bmatrix} 27 & 99 & 92 \\ 32 & 116 & -124 \\ 195 & -121 & -8 \end{bmatrix}} \overset{M}{\begin{bmatrix} 1 & 3 & 5 \\ & 1 & 15 \\ & & 1 \end{bmatrix}} \overset{S^{-1}}{\begin{bmatrix} 1 & & \\ & 1/4 & \\ & & 1/16 \end{bmatrix}} = \overset{APMS^{-1}}{\begin{bmatrix} 27 & 45 & 107 \\ 32 & 53 & 111 \\ 195 & 116 & -53 \end{bmatrix}}.$$

The remainder of this paper is organised as follows. Cost estimates for basic matrix operations are given in the next section. Section 3 shows how to use a sparse inverse expansion [5] to compute $A^{-1}b$ in the special case when $2 \perp \det A$. Section 4 shows how the derandomization approach for polynomial matrices [3] can be adapted to the integer case in order to compute a 2-decomposition of A. Section 5 defines the 2-massager and present a duality between 2-massagers and 2-decompositions. This section also gives

ISSAC '19, July 15–18, 2019, Beijing, China
© 2019 Association for Computing Machinery.
ACM ISBN 978-1-4503-6084-5/19/07…$15.00
https://doi.org/10.1145/XXXXXX.XXXXXX

an overview of our algorithm to compute a 2-massager. The subroutines to compute a 2-massager, along with a complexity analysis, are given in Sections 6–10. Finally, Section 11 gives the deterministic algorithm for linear system solving.

Cost model

We assume that integers are stored using their binary representation. Thus, for any power of two X, we can deduce the X-adic representation of a positive integer without any computation. The algorithms we propose in this paper are designed to not require any radix conversions.

One of our goals is to design our algorithms so that they behave well under a very general cost model. To this end we will give cost estimates using a function

$$\mathsf{MM}(n,d) : (\mathbb{R}_{\geq 0}, \mathbb{R}_{\geq 0}) \to \mathbb{R}_{\geq 0}$$

that for any nonnegative integers $n' \leq n$ and $d' \leq d$ bounds the number of bit operations required to multiply modulo $2^{d'}$ two square matrices over $\mathbb{Z}/(2^{d'})$ of dimension n'. Here, $\mathbb{Z}/(2^d) := \{0, 1, \ldots, 2^d - 1\}$ is the ring of integers modulo 2^d. A lower bound is $\mathsf{MM}(n,d) \in \Omega(n^2 d)$ and, using an obvious block decomposition, we have $\mathsf{MM}(cn,d) \in O(\mathsf{MM}(n,d))$ for any positive constant c.

In this paper we assume that MM satisfies the following regularity assumptions:

$$\mathcal{H}^{1\leq d}_{\mathsf{MM}} \quad : \quad k\,\mathsf{MM}(n, d/k) \leq \mathsf{MM}(n,d) \ \text{ for all } 1 \leq k \leq d$$

$$\mathcal{H}^{2\leq n}_{\mathsf{MM}} \quad : \quad k^2\,\mathsf{MM}(n/k, d) \leq \mathsf{MM}(n,d) \ \text{ for all } 1 \leq k \leq n$$

$$\mathcal{H}^{d\leq 2}_{\mathsf{MM}} \quad : \quad \mathsf{MM}(n, kd) \leq k^2\,\mathsf{MM}(n,d) \text{ for all } 1 \leq k$$

The first two of these assumptions allow us to simplify the cost estimates of algorithms that recurse on the precision d and dimension n, respectively. Many of the cost estimates derived in susbsequent sections are of the form $\mathsf{MM}(n,d)$ for a d that satisfies $d \in O(\log(n||A||))$ where A is the input matrix to the overall problem. The third assumption gives $\mathsf{MM}(n,d) \in O(\mathsf{MM}(n, \log(n||A||)))$.

2 PRELIMINARIES

For an integer a and precision $d \in \mathbb{Z}_{\geq 0}$ we denote by $\mathrm{Rem}(a, 2^d)$ and $\mathrm{Quo}(a, 2^d)$ the unique integers such that $a = \mathrm{Rem}(a, 2^d) + \mathrm{Quo}(a, 2^d)\, 2^d$ with $\mathrm{Rem}(a, 2^d) \in \mathbb{Z}/(2^d)$. When the first argument of Rem or Quo is a matrix, the intention is to apply the function elementwise to the entries.

In order to compute the product of two matrices A and B over $\mathbb{Z}$ (rather than modulo some power of 2) we can multiply $\mathrm{Rem}(A, 2^d)$ and $\mathrm{Rem}(B, 2^d)$ over $\mathbb{Z}/(2^d)$ for large enough d and then reduce entries in the result in the usual symmetric range $\{-2^{d-1}+1, -2^{d-1}+2, \ldots, 2^{d-1}\}$.

Lemma 1. *Let $A, B \in \mathbb{Z}^{n\times n}$. Then the product $AB \in \mathbb{Z}^{n\times n}$ can be computed in time $O(\mathsf{MM}(n,d))$ for any d that satisfies $n||A||||B|| \leq 2^{d-1} - 1$.*

In some of our algorithms we will need to compute AB for an $A \in \mathbb{Z}^{n\times n}$ and a $B \in \mathbb{Z}^{n\times m}$ where the bitlength of entries in B are approximately p times the bitlength of entries in A. The next lemma shows how to do this efficiently if the dimension × precision compromise $m \times p \in O(n)$ holds.

Lemma 2. *Let $A \in \mathbb{Z}^{n\times n}$ and $X \in \mathbb{Z}_{>0}$ be a power of 2 such that $\log X \in O(\log(n||A||))$. If $B \in \mathbb{Z}/(X^p)^{n\times m}$ with $mp \in O(n)$, then $\mathrm{Rem}(AB, X^p)$ can be computed in time $O(\mathsf{MM}(n, \log(n||A||)))$.*

Proof. Let B have X-adic expansion $B = B_0 + B_1 X + \cdots$ and set

$$B' = \left[\begin{array}{c|c|c|c} B_0 & B_1 & \cdots & B_{p-1} \end{array}\right] \in \mathbb{Z}^{n\times mp}.$$

Let $A^{(1)} \in \mathbb{Z}^{n\times n}_{\geq 0}$ be the matrix obtained from A by replacing all negative entries with zero, and let $A^{(2)} = A^{(1)} - A \in \mathbb{Z}^{n\times n}_{\geq 0}$. Then $A = A^{(1)} - A^{(2)}$ where both $A^{(1)}$ and $A^{(2)}$ are over $\mathbb{Z}_{\geq 0}$. Let d be minimal such that $n||A||X \leq 2^{d-1} - 1$. Then by Lemma 1 we can compute the products

$$A^{(i)}B' = \left[\begin{array}{c|c|c|c} A^{(i)}B_0 & A^{(i)}B_1 & \cdots & A^{(i)}B_{p-1} \end{array}\right] \in \mathbb{Z}^{n\times mp}_{\geq 0}$$

for $i = 1, 2$ within the target complexity. Now compute $A^{(i)}B = \sum_{i=0}^{p-1} X^i A^{(i)} B_i$ for $i = 1, 2$. These sums can be computed in time $O(nmpd) \in O(n^2 d)$ since augmenting $\sum_{j=0}^{k} X^j A^{(i)} B_j$ for some $k < p-1$ by adding $X^{k+1} A^{(i)} B_{k+1}$ to it only needs to work on the leading $O(d)$ bits. Return $\mathrm{Rem}(A^{(1)}B - A^{(2)}B, X^p)$. □

We will make use of some well known algorithms which reduce computations to matrix multiplication. We summarize what we need here. The first two results are needed only for matrices over $\mathbb{Z}/(2)$. The cost of the triangular matrix inversion algorithm [1] follows the recurrence $I(n) \leq 2I(n/2) + O(\mathsf{MM}(n/2, 1))$. Assuming $\mathcal{H}^{2\leq n}_{\mathsf{MM}}$ gives the following.

Lemma 3. *If $A \in \mathbb{Z}/(2)^{n\times n}$ is unit upper triangular, then its inverse can be computed in time $O(\mathsf{MM}(n, 1))$.*

For the next result we can use the `LQUP`-decomposition [4]. For an $m \times n$ matrix with $m \leq n$, the algorithm recurses on m and has complexity following $T(m) \leq 2T(m/2) + I(m) + O((n/m)\mathsf{MM}(m, 1))$. Assuming $\mathcal{H}^{2\leq n}_{\mathsf{MM}}$ gives the following when $m = n$.

Lemma 4. *Given $A \in \mathbb{Z}/(2)^{n\times n}$, the rank r of A together with a nonsingular $U \in \mathbb{Z}/(2)^{n\times n}$ and an $n \times n$ permutation matrix P such that UAP has its first r columns those of I_n and its last $n - r$ rows zero can be computed in time $O(\mathsf{MM}(n, 1)\log n)$.*

A matrix $A \in \mathbb{Z}/(2^d)^{n\times n}$ is unimodular if its determinant is odd. In this case the inverse of A, denoted by A^{-1}, is the unique matrix from $\mathbb{Z}/(2^d)^{n\times n}$ such that $\mathrm{Rem}(A^{-1}A, 2^d) = I_n$. Algebraic Newton iteration [2, Algorithm 9.3] gives the following, assuming $\mathcal{H}^{1\leq d}_{\mathsf{MM}}$.

Lemma 5. *Let $A \in \mathbb{Z}/(2^d)^{n\times n}$ be unimodular. If $\mathrm{Rem}(A^{-1}, 2)$ is known, then $\mathrm{Rem}(A^{-1}, 2^d)$ can be computed in time $O(\mathsf{MM}(n,d))$.*

3 SPECIAL SYSTEM SOLVING

Let $A \in \mathbb{Z}^{n\times n}$. This section shows how to compute $A^{-1}B \bmod 2^d$ for a given $B \in \mathbb{Z}^{n\times m}$ and precision $d \in \mathbb{Z}_{>0}$. Suppose X is a power of 2 that satisfies $\log X \in \Theta(\log(n||A||))$. Then $A^{-1} \bmod 2^d$ is explictly given by its X-adic expansion

$$A^{-1} \equiv A_0 + A_1 X + A_2 X^2 + \cdots + A_{p-1}^{p-1} \bmod X^p, \tag{2}$$

where $p \in \Theta(d/\log X)$. Instead, we rely on an algorithm [5] that combines linear and quadratic lifting to compute a *sparse inverse* *expansion* for A^{-1}: a straight line formula that computes $A^{-1} \bmod X^{2^{k+1}-1}$ for a given $k \in \mathbb{Z}_{\geq 0}$. For example, for $k = 2$ we have

$$A^{-1} \equiv (A_0(I + R_0X) + M_0X^2)(I + R_1X^3) + M_1X^6 \bmod X^{2^{2+1}-1}.$$

Compared to the explicit expansion (2) which requires p coefficient matrices A_* over $\mathbb{Z}/(X)$, the strait line formula encodes $A^{-1} \bmod 2^d$ using only $2(k+1)+1 \in \Theta(\log p)$ integer matrices A_0, R_*, M_* that have entries bounded in bitlength by $O(\log X)$. Once the formula has been computed, the system solution $A^{-1}B \bmod 2^d$ can be computed by premultiplying B by these $2(k+1)+1$ matrices in the correct order.

```
DoublePlusOneLift($A, n, k$)
Input: $A \in \mathbb{Z}^{n\times n}$ with $2 \perp \det A$ and $k \in \mathbb{Z}_{>0}$.
Output: $A_0, R_0, \ldots, R_{k-1}, M_0, \ldots, M_{k-1} \in \mathbb{Z}^{n\times n}$ such that the
        following straight line formula computes
                 $\mathrm{Rem}(A^{-1}, X^{2^{k+1}-1})$,
        where $X = 2^{\lceil \log_2(\max(10^4, 3.61n^2||A||)) \rceil}$.
Note: $||A_0|| < X$, $||M_*|| < X$ and $||R_*|| < 0.6001n||A||$.
    $E, F := 0_{n\times n}, I_n$
    $p := 2^{k+1} - 1$
    $X := 2^{\lceil \log_2(\max(10^4, 3.61n^2||A||)) \rceil}$
    for $i = k-1$ down to $0$ do
        $E := \mathrm{Rem}(E + (X^{2^{i+1}-1})^2 M_i F, X^p)$
        $F := \mathrm{Rem}(F + X^{2^{i+1}-1} R_i F, X^p)$
    od
    $E := \mathrm{Rem}(E + A_0 F, X^p)$
    return $E$
```

Figure 1: Problem `DoublePlusOneLift`

The following result is derived in [5, Section 3].

Lemma 6. *Assuming $k \in O(\log n)$, Problem* `DoublePlusOneLift` *in Figure 1 can be solved in time $O(\mathsf{MM}(n, \log(n||A||)) \log n)$.*

```
SpecialSolve($A, B, d, n, m$)
Input: $A \in \mathbb{Z}^{n\times n}$ with $2 \perp \det A$, $B \in \mathbb{Z}^{n\times m}$ and $d \in \mathbb{Z}_{>0}$.
Output: $\mathrm{Rem}(A^{-1}B, 2^d)$.
```

Figure 2: Problem `SpecialSolve`

Corollary 7. *If the dimension × precision invariant $m \times d \in O(n \log(n||A||))$ holds, then Problem* `SpecialSolve` *in Figure 2 can be solved in time $O(\mathsf{MM}(n, \log(n||A||)) \log n)$.*

Proof. Let $k \in \mathbb{Z}_{\geq 0}$ be minimal such that $(2^{k+1} - 1)\log_2 X \geq d$, X the smallest power of 2 such that $X \geq \max(10000, 3.61n^2||A||)$. Then $k \in O(\log(d/\log X))$. Call `DoublePlusOneLift`(A, n, k) to compute the straight line formula for $A^{-1} \bmod X^{2^{k+1}-1}$.

Next, use the straight line formula shown in Figure 1 but with the first line replaced by "$E, F := 0_{n\times m}, B$." The algorithm will then return $\mathrm{Rem}(A^{-1}B, X^{2^{k+1}-1})$. The cost of applying the straight line formula is determined by the matrix multiplications inside the loop. By Lemma 2, the scheme runs in time $O(\mathsf{MM}(n, \log(n||A||)\, k)$. Finally, the assumption $md \in O(n \log(n||A||))$ implies that $k \in O(\log(n/m))$, which is $O(\log n)$. □

4 2-DECOMPOSITIONS

Two matrices over $\mathbb{Z}/(2^d)$ of equal dimensions are said to be *left equivalent* or *row equivalent* if one can be obtained from the other by premultiplying by a unimodular matrix. The unimodular matrix represents a set of row operations converting one matrix into the other. Corresponding to every $A \in \mathbb{Z}/(2^d)^{n\times n}$ there is a permutation matrix P such that $\mathrm{Rem}(AP, d)$ is left equivalent to a matrix

$$\begin{bmatrix} 2^{e_1} & * & \cdots & * & * & \cdots & * \\ & 2^{e_2} & \cdots & * & * & \cdots & * \\ & & \ddots & \vdots & \vdots & \cdots & * \\ & & & 2^{e_r} & * & \cdots & * \\ & & & & & & \\ & & & & & & \\ & & & & & & \end{bmatrix} \in \mathbb{Z}/(2^d)^{n\times n} \tag{3}$$

that is in *triangular Smith form*: $e_1 \leq e_2 \leq \cdots \leq e_r$ and all entries in row i are divisible by 2^{e_i}, $1 \leq i \leq r$. The e_i are unique.

The above discussion is for a matrix A over $\mathbb{Z}/(2^d)$. Now let $A \in \mathbb{Z}^{n\times n}$ be a nonsingular integer matrix. Then the *2-Smith form* of A is the matrix $\mathrm{diag}(2^{e_1}, 2^{e_2}, \ldots, 2^{e_n})$ with 2^{e_i} the largest power of two which divides the i'th invariant factor of the Smith form of A over $\mathbb{Z}$, $1 \leq i \leq n$. Since 2^{e_n} divides $\det A$, and $|\det A| \leq n^{n/2}||A||^n$, the 2-Smith form of A over $\mathbb{Z}$ can be recovered by computing a triangular Smith form of $\mathrm{Rem}(A, X^{n+1})$ over $\mathbb{Z}/(X^{n+1})$, where X is the smallest power of two such that $X \geq n^{1/2}||A||$. For the remainder of this section we will refer to the exponent of the modulus X^p as the "precision" p.

In [3][1] one finds an algorithm to compute a permutation matrix P such that $\mathrm{Rem}(AP, X^{n+1})$ is left equivalent (over $\mathbb{Z}/(X^{n+1})$) to a triangular Smith form H. As mentioned previously, the precision $n+1$ is large enough to ensure that H will be as in (3) with $r = n$. The matrix $U := APH^{-1}$ is integral with $2 \perp \det A$. We refer to the pair of matrices (P, H) as a *2-decomposition of A* and note that $A = UHP^{-1}$.

Working with precision $n+1$ to compute (P, H) in one fell swoop is too expensive. Instead, the authors in [3] use the observation, summarized in the next lemma, that many of the initial invariant factors can be recovered using a considerably lower precision. For any $0 \leq k \leq n$ we can partition the invariant factors in the 2-Smith form of A as follows:

$$\mathrm{diag}(\overbrace{2^{e_1}, 2^{e_2}, \ldots, 2^{e_k}}^{\text{first } k},\ \overbrace{2^{e_{k+1}}, 2^{e_{k+2}}, \ldots, 2^{e_*}}^{\text{next } \lceil (n-k)/2 \rceil},\ \overbrace{2^{e_*}, 2^{e_*}, \ldots, 2^{e_n}}^{\text{last } \lfloor (n-k)/2 \rfloor})$$

Lemma 8. *Let $A \in \mathbb{Z}^{n\times n}$ be nonsingular and $0 \leq k < n$. If $X \in \mathbb{Z}$ satisfies $X \geq n^{1/2}||A||$ then the 2-Smith form of A has at most $\lfloor (n-k)/2 \rfloor$ invariant factors $\geq X^{2n/(n-k)}$.*

An application of Lemma 8 with $k = 0$ states that a precision 2 is sufficient to compute a permutation P such that $\mathrm{Rem}(AP, X^2)$ is

[1] The algorithms were developed for matrices over $\mathsf{K}[x]$ but the approach carries over directly to integer matrices.

left equivalent to a triangular Smith form as in (3) with $r \geq \lceil n/2 \rceil$. Next, the algorithm works at precision $\lceil 2n/(n-r) \rceil$ to recover at least $\lceil (n-r)/2 \rceil \geq n/4$ of the remaining $n-r$ invariant factors. At each step, the algorithm exploits a natural dimension × precision invariant: the number of remaining invariant factors times the precision is $\Theta(n)$.

At the beginning of an iteration, the algorithm has already computed a permutation P_1 such that for the previous working precision p, the matrix $\mathrm{Rem}(AP_1, X^p)$ is left equivalent (over $\mathbb{Z}/(X^p)$) to a triangular Smith form as in (3). (At the beginning of the first iteration $r = 0$.) Let

$$H_1 = \left[\begin{array}{c|c} E_1 & \\ \hline & I_{n-r} \end{array}\right]$$

be the matrix in (3) but with last $n-r$ columns replaced by those of I_n. We refer to the pair of matrices (P_1, H_1) as an index $(0, r)$ 2-decomposition of A: the matrix $B := AP_1H_1^{-1}$ will be integral with first r columns having full rank modulo 2. If $r = n$ we are done so assume that $r < n$. Next the algorithm updates the precision to $p := \lceil 2n/(n-r) \rceil$ and computes a permutation $P_2 = \mathrm{diag}(I_r, *)$ such that $\mathrm{Rem}(BP_2, X^p)$ is left equivalent to a triangular Smith form having the shape

$$\left[\begin{array}{c|c|c} I_r & V_2 & * \\ \hline & E_2 & * \\ \hline & & \end{array}\right].$$

By Lemma 8, the column dimension m of E_2 satisfies $m \geq \lceil (n-r)/2 \rceil$. Set

$$H_2 = \left[\begin{array}{c|c|c} I_r & V_2 & \\ \hline & E_2 & \\ \hline & & I_{n-r-m} \end{array}\right].$$

We call the pair (P_2, H_2) an index (r, m) 2-decomposition of B. Because of the structure of the matrices, we have $(P_1H_1^{-1})(P_2H_2^{-1}) = (P_1P_2)(H_2H_1)^{-1}$. Then P_1P_2 is a permutation with $\mathrm{Rem}(AP_1P_2, X^p)$ being left equivalent to a triangular Smith form that has first $r+m$ columns equal to those of

$$H_2H_1 = \left[\begin{array}{c|c|c} E_1 & V_2 & \\ \hline & E_2 & \\ \hline & & I_{n-r-m} \end{array}\right].$$

Since each iteration computes at least half of the remaining invariant factors of the 2-Smith form, the number of iterations is bounded by $O(\log n)$. The following theorem captures the approach of [3] described above to compute a 2-decomposition.

Theorem 9. *Let $A \in \mathbb{Z}^{n \times n}$ be nonsingular, $0 \leq r \leq n$, and $0 \leq m \leq n-r$. If (P_1, H_1) is an index $(0, r)$ 2-decomposition for A, and (P_2, H_2) is an index (r, m) 2-decomposition for $B := AP_1H_1^{-1}$, then (P_1P_2, H_2H_1) is an index $(0, r+m)$ 2-decomposition for A.*

To avoid expression swell, the H_i computed at each step can be transformed into *Hermite canonical form*: offdiagonal entries are reduced modulo the diagonal entry in the same column. The H_i are then unique up to the choice of the permutations P_i. In the polynomial setting, the inverse H^{-1} of a matrix H in Hermite form will be a proper matrix fraction, ensuring that $U := AH^{-1}$ will have degree bounded by the degree of A. However, over the integers, H^{-1} may not be proper, frustrating attempts to obtain a good bound for the bitlength of entries in U. We end this section with an example of a class of ill-conditioned Hermite forms.

Example 10. *For $n = 2, 3, 4, \ldots$ consider the family of Hermite forms $H \in \mathbb{Z}^{n \times n}$ that are Toeplitz, with diagonal entry 2 and offdiagonal entries alternating between 1 and 0. For example, for $n = 6$,*

$$H = \begin{bmatrix} 2 & 1 & 0 & 1 & 0 & 1 \\ & 2 & 1 & 0 & 1 & 0 \\ & & 2 & 1 & 0 & 1 \\ & & & 2 & 1 & 0 \\ & & & & 2 & 1 \\ & & & & & 2 \end{bmatrix} \in \mathbb{Z}^{6 \times 6}.$$

Offdiagonal entries in the first row of H^{-1} satisfy

$$H^{-1}_{1,j} = \begin{cases} -1/4 & \text{if } j = 2 \\ 1/8 & \text{if } j = 3 \\ (-1/2)H^{-1}_{1,j-1} + H^{-1}_{1,j-2} & \text{if } j \geq 4 \end{cases}$$

The closed form for this recurrence shows that $\log |H^{-1}_{1,j}| \in \Theta(j)$. For $n \geq 500$ the largest entry in $(\det H)H^{-1}$ has bitlength $\approx 1.35n$ compared to $\det H$ which has bitlength n.

5 2-MASSAGERS

Rather than computing a 2-decomposition for our input matrix, we introduce the notion of a 2-massager.

Definition 11. *Let $A \in \mathbb{Z}^{n \times n}$ be nonsingular. A* 2-massager *for A is a triple of matrices (P, S, M) from $\mathbb{Z}^{n \times n}$ such that:*

- *P is a permutation.*
- *$S = \mathrm{diag}(s_1, \ldots, s_n)$ is the 2-Smith form of A.*
- *M is unit upper triangular.*
- *$APMS^{-1}$ is integral with $2 \perp \det APMS^{-1}$.*

(P, S, M) is a reduced 2-massager *if entries in column i of M are from $\mathbb{Z}/(s_i)$, $1 \leq i \leq n$.*

There is a one to one correspondence between 2-massagers and 2-decompositions. Note that it follows from the uniqueness of the 2-Smith form S of A that the triangular Smith form H from any 2-decomposition of A will have the same diagonal entries as S.

Theorem 12. *Duality between 2-decompositions and -massagers:*

- *If (P, H) is a 2-decomposition of A,*
  then $(P, S, (S^{-1}H)^{-1})$ is a 2-massager for A.
- *If (P, S, M) is a 2-massager for A,*
  then (P, SM^{-1}) is a 2-decomposition of A.

Proof. Since H is in triangular Smith form, $S^{-1}H$ will be unit upper triangular and integral. It suffices to note that $APH^{-1} = AP(SS^{-1}H)^{-1} = AP(S^{-1}H)^{-1}S^{-1}$. □

Suppose (P, S, M) is a 2-massager for A. Since $APMS^{-1}$ is integral, column i of APM must be congruent to zero modulo s_i, $1 \leq i \leq n$. This gives the following.

Lemma 13. *If (P, S, M) is a 2-massager for A, then a reduced 2-massager for A can be obtained by reducing offdiagonal entries in column i of M by s_i, $1 \leq i \leq n$.*

Our algorithm to compute a 2-massager for A will proceed in a similar manner to the algorithm for 2-decomposition sketched in the previous section. In order to simplify the discussion it will be useful to introduce the following definition.

Definition 14. *Let $B \in \mathbb{Z}^{n\times n}$ be nonsingular with first r columns full rank modulo 2. An index (r, m) 2-massager for B is a triple of matrices (P, S, M) from $\mathbb{Z}^{n\times n}$ such that*

- *$P = \operatorname{diag}(I_r, *)$ is a permutation.*
- *$S = \operatorname{diag}(I_r, s_r, \ldots, s_{r+m}, I_{n-r-m})$ with principal $(r+m) \times (r+m)$ submatrix equal to that of the 2-Smith form of B.*
- *M is unit upper triangular with first r and last $n-r-m$ columns those of I_n.*
- *$BPMS^{-1}$ is integral with first $r+m$ columns having full rank modulo 2.*

(P, S, M) is a reduced index (r, m) 2-massager if entries in column i of M are from $\mathbb{Z}/(s_i)$, $r \le i \le m$.

Suppose we have already computed an index $(0, r)$ 2-massager (P_1, S_1, M_1) for A. At the start of the first iteration $r = 0$ and we have the trivial index $(0, 0)$ 2-massager (I_n, I_n, I_n). Then $B := AP_1M_1S_1^{-1}$ will be integral with first r columns of full rank modulo 2. If $r = n$ we are done so assume $r < n$. Our goal now is to compute an index (r, m) 2-massager for B. By Lemma 8 we can guarantee to achieve $m \ge \lceil (n-r)/2 \rceil$ by working modulo X^p where X is the smallest power of two $\ge n^{1/2}||A||$ and $p = \lceil 2n/(n-r) \rceil$. Compute an index (r, m) 2-decomposition (P_2, H_2) for B. Let

$$H_2 := \left[\begin{array}{c|c|c} I_r & V_2 & \\ \hline & E_2 & \\ \hline & & I \end{array}\right].$$

Let S_2 be the diagonal matrix with same diagonal entries as H_2. Define D_2 by writing $S_2 = \operatorname{diag}(I_r, D_2, I_{n-r-m})$, that is, D_2 is the diagonal matrix with same diagonals as E_2. As a corollary of Theorem 12, (P_2, S_2, M_2) will be an index (r, m) 2-massager for B, where

$$M_2 := (S_2^{-1}H_2)^{-1} = \left[\begin{array}{c|c|c} I_r & -V_2(D_2^{-1}E_2)^{-1} & \\ \hline & (D_2^{-1}E_2)^{-1} & \\ \hline & & I_{n-r-m} \end{array}\right]. \tag{4}$$

Then $A(P_1M_1S_1^{-1})(P_2M_2S_2^{-1}) = BP_2M_2S_2^{-1}$ will be integral with first $m + r$ columns having full rank modulo 2. By exploiting the duality of Theorem 12 we can combine (P_1, S_1, M_1) and (P_2, S_2, M_2) to obtain an index $(0, r+m)$ 2-massager. Define D_1 by writing $S_1 = \operatorname{diag}(D_1, I_{n-r-m})$, that is, D_1 is comprised of the first r entries of the 2-Smith form of A. By duality, $(P_1, S_1M_1^{-1})$ is an index $(0, r)$ 2-decomposition of A. Let $P = P_1P_2$ and $S = S_1S_2$. By Theorem 9 $(P, H_2S_1M_1^{-1})$ is then an index $(0, r+m)$ 2-decomposition of A. Using duality in the opposite direction shows that an index $(0, r+m)$ 2-massager for A is given by $(P, S, (S^{-1}H_2S_1M_1^{-1})^{-1})$. Note that $(S^{-1}H_2S_1M_1^{-1})^{-1} = (S_1^{-1}S_2^{-1}H_2S_1M_1^{-1})^{-1} = M_1S_1^{-1}M_2S_1$. We then obtain the following result.

Theorem 15. *Let $A \in \mathbb{Z}^{n\times n}$ be nonsingular, $0 \le r \le n$, and $0 \le m \le n-r$. If (P_1, S_1, M_1) is an index $(0, r)$ 2-massager for A, and (P_2, S_2, M_2) is an index (r, m) 2-massager for $B := AP_1M_1S_1^{-1}$, then $(P_1P_2, S_1S_2, M_1S_1^{-1}M_2S_1)$ is an index $(0, r+m)$ 2-massager for A.*

If we write

$$M_1 = \left[\begin{array}{c|c} F_1 & \\ \hline & I_{n-r} \end{array}\right] \tag{5}$$

and

$$M_2 = \left[\begin{array}{c|c|c} I_r & W_2 & \\ \hline & F_2 & \\ \hline & & I_{n-r-m} \end{array}\right], \tag{6}$$

then by Theorem 15 an index $(0, r+m)$ 2-massager for A is given by (P, S, M) where

$$M = M_1S_1^{-1}M_2S_1 = \left[\begin{array}{c|c|c} F_1 & F_1D_1^{-1}W_2 & \\ \hline & F_2 & \\ \hline & & I_{n-r-m} \end{array}\right]. \tag{7}$$

6 TRIANGULAR SMITH FORM

In this section we give an algorithm for computing a triangular Smith form of a matrix A over $\mathbb{Z}/(d)$.

```
TriangularSmithForm($A, n, d$)
Input:  $A \in \mathbb{Z}/(2^d)^{n\times n}$ for $d \in \mathbb{Z}_{>0}$.
Output: $U, P$ such that $U \in \mathbb{Z}/(2^d)^{n\times n}$ is unimodular, $P$ is an
        $n \times n$ permutation matrix, and Rem($UAP, d$) is in triangular Smith form over $\mathbb{Z}/(2^d)$.
```

Figure 3: Problem `TriangularSmithForm`

Theorem 16. *Problem* `TriangularSmithForm` *in Figure 3 can be solved in time $O(\mathsf{MM}(n, d)(\log n + \log d))$.*

Proof. We describe a divide and conquer algorithm that recurses on the precision parameter d and has running time bounded by the recurrence

$$T(d) \le \left[\begin{array}{ll} T(\lceil d/2 \rceil) + T(\lfloor d/2 \rfloor) + O(\mathsf{MM}(n, d)) & \text{if } d > 1 \\ \mathsf{MM}(n, 1)\log n & \text{if } d = 1. \end{array}\right.$$

Assuming $\mathcal{H}_{\mathsf{MM}}^{1\le d}$ we have the solution $T(d) \in O(d\,\mathsf{MM}(n, 1)\log n + \mathsf{MM}(n, d)\log d)$, which simplifies to the target complexity since $d\,\mathsf{MM}(n, 1) \le \mathsf{MM}(n, d)$ using $\mathcal{H}_{\mathsf{MM}}^{1\le d}$.

For the base case $d = 1$ use Lemma 4. Assume now that $d > 1$. Then set $d_1 = \lceil d/2 \rceil$ and $d_2 = \lfloor d/2 \rfloor$ and recursively compute

$$U_1, P_1 := \texttt{TriangularSmithForm}(\operatorname{Rem}(A, 2^{d_1}), n, d_1).$$

Let r be the number of nonzero rows of $\operatorname{Rem}(U_1AP_1, 2^{d_1})$ and let S be the $r \times r$ diagonal matrix with $S_{ii} = \operatorname{Rem}(U_1AP_1, 2^{d_1})_{ii}$, $1 \le i \le r$. Then

$$\left[\begin{array}{c|c} S^{-1} & \\ \hline & I_{n-r} \end{array}\right] \operatorname{Rem}(U_1AP_1, 2^d)$$

is an integral matrix. We can thus split $\operatorname{Rem}(U_1AP_1, 2^d)$ into two parts using Rem and Quo. Let

$$\left[\begin{array}{c|c} T & T' \\ \hline & \end{array}\right] = \left[\begin{array}{c|c} S & \\ \hline & I_{n-r} \end{array}\right] \operatorname{Rem}\left(\left[\begin{array}{c|c} S^{-1} & \\ \hline & I_{n-r} \end{array}\right] \operatorname{Rem}(U_1AP_1, 2^d), 2^{d_1}\right)$$

and

$$\left[\begin{array}{c|c} B & B' \end{array}\right] = \left[\begin{array}{c|c} S & \\ \hline & I_{n-r} \end{array}\right] \operatorname{Quo}\left(\left[\begin{array}{c|c} S^{-1} & \\ \hline & I_{n-r} \end{array}\right] \operatorname{Rem}(U_1AP_1, 2^d), 2^{d_1}\right).$$

Then

$$\operatorname{Rem}(U_1AP_1, 2^d) = \left[\begin{array}{c|c} T & T' \\ \hline & \end{array}\right] + \left[\begin{array}{c|c} B & B' \end{array}\right] 2^{d_1}.$$

where T is $r \times r$ and B is $n \times r$. Note that by construction both

$$\left[\begin{array}{c|c} S^{-1} & \\ \hline & I_{n-r} \end{array}\right] \left[\begin{array}{c|c} T & T' \end{array}\right]$$

and

$$\left[\begin{array}{c|c} S^{-1} & \\ \hline & I_{n-r} \end{array}\right]\left[\begin{array}{c|c} B & B' \end{array}\right]$$

will be integral.

Let $V = \text{Rem}((S^{-1}T)^{-1}, 2^{d_1})$. Since the diagonal entries of S are powers of 2 of degree at most $d_1 - 1$, the matrix $2^{d_1}S^{-1}$ will be integral and, moreover, $\text{Rem}(2^{d_1}S^{-1}, 2) = 0_{r\times r}$. The matrix

$$U_1' = \left[\begin{array}{c|c} I_r & \\ \hline & I_{n-r} \end{array}\right] - \left[\begin{array}{c|c} 2^{d_1}BVS^{-1} & \end{array}\right]$$

thus satisfies $\text{Rem}(U_1', 2) = I_n$ and hence is unimodular.

Next we show that

$$\text{Rem}(U_1'U_1AP_1, 2^d) = \left[\begin{array}{c|c} T & * \\ \hline & C2^{d_1} \end{array}\right]$$

for an $(n-r)\times(n-r)$ matrix C. Considering the structure of U_1', it suffices to note, on the one hand, that

$$\begin{aligned} &\left[\begin{array}{c|c} 2^{d_1}BVS^{-1} & \end{array}\right]\left[\begin{array}{c|c} T & T' \\ \hline & \end{array}\right] \\ &\quad\equiv \left[\begin{array}{c|c} 2^{d_1}B(VS^{-1}T) & 2^{d_1}BV(S^{-1}T') \end{array}\right] \pmod{2^d} \\ &\quad\equiv \left[\begin{array}{c|c} 2^{d_1}B(I_r + *2^{d_1}) & 2^{d_1}BV* \end{array}\right] \pmod{2^d} \\ &\quad\equiv \left[\begin{array}{c|c} B2^{d_1} & *2^{d_1} \end{array}\right] \pmod{2^d} \end{aligned}$$

where the $*$ are integral matrices. On the other hand we have that

$$\begin{aligned} &\left[\begin{array}{c|c} 2^{d_1}BVS^{-1} & \end{array}\right]\left[\begin{array}{c|c} B & B' \end{array}\right]2^{d_1} \\ &\quad\equiv \left[\begin{array}{c|c} BV2^{d_1} & \end{array}\right]\left[\begin{array}{c|c} S^{-1} & \\ \hline & I \end{array}\right]\left[\begin{array}{c|c} B & B' \end{array}\right]2^{d_1} \pmod{2^d} \\ &\quad\equiv \left[\begin{array}{c|c} BV2^{d_1} & \end{array}\right]\left[\begin{array}{c|c} * & * \end{array}\right]2^{d_1} \pmod{2^d} \\ &\quad\equiv 0_{n\times n} \pmod{2^d} \end{aligned}$$

We can then compute

$$U_2, P_2 := \texttt{TriangularSmithForm}(C, n-r, d_2)$$

and return

$$U, P = \text{Rem}\left(\left[\begin{array}{c|c} I_r & \\ \hline & U_2 \end{array}\right]U_1'U_1, 2^d\right), P_1\left[\begin{array}{c|c} I_r & \\ \hline & P_2 \end{array}\right].$$

It remains to bound the cost of the nonrecursive work. Other than some multiplications of matrices bounded in dimension by n and precision d, the only other computation is that of the inverse V. Lemmas 3 and 5 show that V can be computed in time $O(\mathsf{MM}(n, d))$. □

7 APPLYING AN INDEX MASSAGER

In this section we show how to apply a reduced index massager to an input matrix A in order to produce the massaged matrix $U := APMS^{-1}$ that has $2 \perp \det U$. In the next lemma, by *length* k of a nonzero finite X-adic expansion $a_0 + a_1X + \cdots$ we mean the maximal $k \in \mathbb{Z}_{\geq 1}$ such that a_{k-1} is non-zero.

Lemma 17. *Let (P, S, M) be a reduced index (r, m) 2-massager for a nonsingular $A \in \mathbb{Z}^{n\times n}$. If $X = 2^d$ is the smallest power of 2 such that $X \geq n^{1/2}||A||$, then the sum of the lengths of the X-adic expansions of the columns of M is bounded by $2n$.*

`ApplyMassager`(A, n, P, S, M, r)

Input: Nonsingular $A \in \mathbb{Z}^{n\times n}$ and a reduced index (r, m) 2-massager (P, S, M) for A.

Output: $APMS^{-1}$

Figure 4: Problem `ApplyMassager`

Proof. It will suffice to prove the result for $r = 0$ and $m = n$. Let $S = \text{diag}(2^{e_1}, 2^{e_2}, \ldots, 2^{e_n})$ be the 2-Smith form of A. Since $\det S \mid \det A$, Hadamard's bound gives $\sum_{i=1}^{n} e_i \leq nd$. Since the maximal magnitude entry in column i of M is 2^{e_i}, the length of the X-adic expansion of column i is equal to $\lfloor e_i/d + 1\rfloor$, $1 \leq i \leq n$. Finally, note that $\sum_{i=1}^{n}\lfloor e_i/d + 1\rfloor \leq 2n$ using $\sum_{i=1}^{n} e_i/d \leq n$. □

Theorem 18. *Problem* `ApplyMassager` *in Figure 4 can be solved in time $O(\mathsf{MM}(n, \log(n||A||)))$.*

Proof. Let C be the matrix obtained from M by replacing each column of M with the $n \times \lfloor e_i/d + 1\rfloor$ matrix comprised of the coefficients of the X-adic expansion of the column. By Lemma 17 the number of columns of C will be bounded by $2n$. We now proceed similarly as in the proof of Lemma 2 to recover APM in the allotted time. As in Lemma 2, write $A = A^{(1)} - A^{(2)}$ where $A^{(i)}$ is over $\mathbb{Z}_{\geq 0}$, $i = 1, 2$. Compute $A^{(i)}PC$ and recover $A^{(i)}PM$ for $i = 1, 2$. Finally, compute $APMS^{-1}$ as $(A^{(1)}PM - A^{(2)}PM)S^{-1}$. □

8 COMPUTING AN INDEX MASSAGER

In this section we show how to use the algorithms `SpecialSolve` and `TriangularSmithForm` to compute an index (r, m) 2-massager for an input matrix $B \in \mathbb{Z}^{n\times n}$. The algorithm assume that B is the matrix obtained by applying an index $(0, r)$ 2-massager to the original input matrix A to the overall problem.

`IndexMassager`(B, n, r, p, X)

Input: Nonsingular $B \in \mathbb{Z}^{n\times n}$, $r \in \mathbb{Z}$ with $0 \leq r < n$, $p \in \mathbb{Z}_{\geq 1}$, and X the smallest power of 2 such that $X \geq n^{1/2}||A||$ for a nonsingular matrix $A \in \mathbb{Z}^{n\times n}$.

Output: P, S, M, m such that (P, S, M) is an index (r, m) 2-massager for B, with m maximal such that invariant factor $r + m$ of the 2-Smith form of A is $< X^p$. The matrix M will satisfy $M = \text{Rem}(M, X^p)$.

Condition: $B = AP_1M_1{S_1}^{-1}$ where (P_1, M_1, S_1) is a reduced index $(0, r)$ 2-massager for A.

Figure 5: Problem `IndexMassager`

Theorem 19. *If $p = \lceil 2n/(n-r)\rceil$, then Problem* `IndexMassager` *in Figure 5 can be solved in time*

$$O(\mathsf{MM}(n, \log(n||A||))\log(n\log||A||)).$$

Proof. We describe a 7 step algorithm.

Step 1: Let Q be the permutation P from the LQUP-decomposition of $\text{Rem}(B, 2)^T$. Then Q is such that QB can be written in a block decomposition as

$$QB = \left[\begin{array}{c|c} B_{11} & B_{12} \\ \hline B_{21} & B_{22} \end{array}\right],$$

where $B_{11} \in \mathbb{Z}^{r\times r}$ is nonsingular modulo 2.

Cost 1: $O(\mathsf{MM}(n, 1) \log n)$.

Step 2: Compute

$$\left[\begin{array}{c} C_1 \\ \hline C_2 \end{array}\right] = \mathrm{Rem}\left(\left[\begin{array}{c|c} B_{11} & \\ \hline B_{21} & I_{n-r} \end{array}\right]^{-1} \left[\begin{array}{c} B_{12} \\ \hline B_{22} \end{array}\right], X^p\right)$$

using the algorithm supporting Corollary 7. We now have the partial triangularization

$$\mathrm{Rem}\left(\left[\begin{array}{c|c} B_{11} & \\ \hline B_{21} & I_{n-r} \end{array}\right]^{-1} \left[\begin{array}{c|c} B_{11} & B_{12} \\ \hline B_{21} & B_{22} \end{array}\right], X^p\right) = \left[\begin{array}{c|c} I_r & C_1 \\ \hline & C_2 \end{array}\right] \tag{8}$$

of QB over $\mathbb{Z}/(X^p)$.

Cost 2: $O(\mathsf{MM}(n, \log(n||A||)) \log n)$, by Corollary 7

Step 3: Compute

$$U', P' := \mathtt{TriangularSmithForm}(C_2, n-r, p \log_2 X)$$

using the algorithm supporting Theorem 16. Set $P := \mathrm{diag}(I_r, P')$.

Cost 3: By Theorem 16,

$$O(\overbrace{\mathsf{MM}(n-r, p \lg X)}^{T_1} \overbrace{(\log(n-r) + \log(p \log X))}^{T_2}).$$

First note that

$$\begin{array}{rcl} T_1 & = & \mathsf{MM}\left(\dfrac{n}{n/(n-r)}, \lceil 2n/(n-r) \rceil \log X\right) \\ & \in & O((n/(n-r))^2\, \mathsf{MM}\left(\dfrac{n}{n/(n-r)}, \log X\right) \\ & \in & O(\mathsf{MM}(n, \log X)) \end{array}$$

using $\mathcal{H}^{d\leq 2}_{\mathsf{MM}}$ and $\mathcal{H}^{2\leq n}_{\mathsf{MM}}$ in succession. Next, the logarithmic factor T_2 is simplified using $\log(n-r) \leq \log n$ and $\log(p \lg X) \in O(\log(n \log ||A||))$.

Step 4: We can now complete the triangularization of (8):

$$\mathrm{Rem}\left(\left[\begin{array}{c|c} I_r & \\ \hline & U' \end{array}\right] \left[\begin{array}{c|c} I_r & C_1 \\ \hline & C_2 \end{array}\right] P, X^p\right) = \left[\begin{array}{c|cc} I_r & V & * \\ \hline & E & * \\ \hline & & \end{array}\right].$$

Here, E is an $m \times m$ triangular Smith form and V is the first m columns of $C_1 P'$. Let D be the $m \times m$ diagonal matrix with the same diagonal entries as E, and set $S := \mathrm{diag}(I_r, D, I_{n-r-m})$.

Cost 4: $O(\mathsf{MM}(n-r, p \log X))$.

It remains to compute M. By Lemma 13 we may compute M modulo X^p. As shown in (4), we can take

$$M = \mathrm{Rem}\left(\left[\begin{array}{c|c|c} I_r & -V(D^{-1}E)^{-1} & \\ \hline & (D^{-1}E)^{-1} & \\ \hline & & I_{n-r-m} \end{array}\right], X^p\right).$$

We will compute this M in the final three steps.

Step 5: First compute $T := \mathrm{Rem}((D^{-1}E)^{-1}, X^p)$.

Cost 5: $O(\mathsf{MM}(n-r, p \log X))$, by Lemmas 3 and 5.

Step 6: Instead of computing the product $\mathrm{Rem}(-VT, X^p)$ we will proceed as follows. Let B'_{12} be the first m columns of $B_{12}P'$. Compute the product $\mathrm{Rem}(B'_{12}T, X^p)$.

Cost 6: $O(\mathsf{MM}(n, \log(n||A||)))$, by Lemma 2.

Step 7: Compute $\mathrm{Rem}(VT, X^p) = \mathrm{Rem}(B_{11}^{-1}(B'_{12}T), X^p)$ using the algorithm supporting Corollary 7.

Cost 7: Same as the cost of step 2. □

9 COMBINING INDEX MASSAGERS

In this section we show how to combine an index $(0, r)$ and index (r, m) 2-massager to obtain an index $(0, m)$ 2-massager.

$\mathtt{CombineMassagers}(P_1, S_1, M_1, n, r, P_2, S_2, M_2, m, X)$

Input: A reduced index $(0, r)$ 2-Smith massager (P_1, S_1, M_1) for a nonsingular $A \in \mathbb{Z}^{n\times n}$, an index (r, m) 2-Smith massager (P_2, S_2, M_2) for $\bar{A} := AP_1M_1S_1^{-1}$, and X the smallest power of 2 such that $X \geq n^{1/2}||A||$.

Output: $(P, S, M, r+m)$, with (P, S, M) a reduced index $(0, r+m)$ 2-massager for A.

Condition: $M_2 = \mathrm{Rem}(M_2, X^p)$ with $p = \lceil 2n/(n-r) \rceil$.

Figure 6: Problem `CombineMassagers`

Theorem 20. *Problem* `CombineMassagers` *in Fig. 6 can be solved in time* $O(\mathsf{MM}(n, \log(n||A||)))$.

Proof. Write M_1 and M_2 using a block decomposition as shown in (5) and (6). By Theorem 15, the only computation required to produce (P, S, M) is to compute $F_1 D_1^{-1} W_2$ as shown in (7), where $F_1 \in \mathbb{Z}^{r\times r}$ and $V := D_1^{-1} W_2 \in \mathbb{Z}^{r\times m}$. By Lemma 13, it will suffice to compute $\mathrm{Rem}(F_1 V, X^p)$. For simplicity, and without loss of generality, we will assume that $r = n$ so that F_1 has dimension $n \times n$ and V has dimension $n \times m$.

Let $F_1 = C_0 + C_1 X + C_2 X^2 + \cdots$ and $V = V_0 + V_1 X + V_2 X^2 + \cdots$ be the X-adic expansions of F_1 and V, respectively. Our approach is to compute the integer matrix product

$$\left[\begin{array}{cccc} C_0 & C_1 & \cdots & C_{p-1} \end{array}\right] \left[\begin{array}{cccc} V_0 & V_1 & \cdots & V_{p-1} \\ & V_0 & \cdots & V_{p-2} \\ & & \ddots & \vdots \\ & & & V_0 \end{array}\right],$$

from which $F_1 V$ modulo X^p is easily recovered. To perform the multiplication efficiently we must take into account that the coefficients $C_i \in \mathbb{Z}/(X)^{n\times n}$ may have many zero columns. For $0 \leq i \leq p-1$, let $k_i \in \mathbb{Z}_{\geq 0}$ be minimal such that $C_i = \left[\begin{array}{c|c} 0 & C'_i \end{array}\right]$ with k_i the column dimension of C'_i. By Lemma 17, the column dimension c of $\left[\begin{array}{cccc} C'_0 & C'_1 & \cdots & C'_{p-1} \end{array}\right]$ satisfies $c \leq 2n$.

Compute the $n \times c$ times $c \times mp$ integer matrix product

$$\left[\begin{array}{cccc} C'_0 & C'_1 & \cdots & C'_{p-1} \end{array}\right] \left[\begin{array}{cccc} V_0^{(k_0)} & V_1^{(k_0)} & \cdots & V_{p-1}^{(k_0)} \\ & V_0^{(k_1)} & \cdots & V_{p-2}^{(k_1)} \\ & & \ddots & \vdots \\ & & & V_0^{(k_{p-1})} \end{array}\right], \tag{9}$$

where $V_i^{(k)}$ is the $k \times m$ submatrix of V_i comprised of the last k rows. Since c and mp are $O(n)$, the matrix multiplication in (9) has cost $O(\mathsf{MM}(n, \log(n||A||)))$. Let the result of the multiplication be

$$[\ E_0 \mid E_1 \mid \cdots \mid E_{p-1}\ ] \in \mathbb{Z}^{n \times mp}$$

Then $F_1 D_1^{-1} W_2$ can be recovered by compute the sum $\sum_{i=0}^{p-1} X^i E_i$ as described in the proof of Lemma 2. □

10 COMPUTING A MASSAGER

In this section we show how to use the algorithms presented in the previous three sections to compute an index massager.

```
Massager($A, n$)
Input: Nonsingular $A \in \mathbb{Z}^{n \times n}$.
Output: $(P, S, M)$, a reduced 2-massager for $A$.
$X$ := the smallest power of 2 such that $X \geq n^{1/2}||A||$
$P, S, M, r := I_n, I_n, I_n, 0$
while $r < n$ do
    $B$ := ApplyMassager($A, n, P, S, M, r$)
    $p := \lceil 2n/(n-r) \rceil$
    $P', S', M', m$ := IndexMassager($B, n, r, p, X$)
    $P, S, M, r$ := CombineMassager($P, S, M, n, r, P', S', M', m, X$)
od
return $(P, S, M)$
```

Figure 7: Algorithm Massager

Theorem 21. *Algorithm* Massager *in Figure 7 is correct. The running time is $O(\mathsf{MM}(n, \log(n||A||)) \log(n \log ||A||) \log n)$.*

Proof. Correctness of the algorithm follows from the input and output specifications of the subroutines presented in Sections 7–9. By Lemma 8, the number of loop iterations is bounded by $\log_2 n$ with the cost of each loop iteration being dominated by the call to IndexMassager. The running time estimate is obtained by multiplying the cost estimate of Theorem 19 by $\log n$. □

11 SYSTEM SOLVING

Suppose $a/b \in \mathbb{Q}$ is a signed fraction with $a \in \mathbb{Z}$, $b \in \mathbb{Z}_{>0}$, $a \perp b$ and $b \perp 2$. Then rational number reconstruction [2] can be use to reconstruct a/b from its image $\mathrm{Rem}(a/b, 2^d)$ for large enough d. More precisely, given upper bounds N and D such that $|a| \leq N$ and $b \leq D$, then

$$\mathtt{RatRecon}(\mathrm{Rem}(a/b, 2^d), 2^d, N, D)$$

will reconstruct a/b for any d that satisfies $2^d \geq 2ND$. If the first argument to RatRecon is a vector then the intent is to apply rational reconstruction elementwise to the entries.

Our algorithm for system solving is based on the following observation. If (P, M, S) is a 2-massager for A and $U := APMS^{-1}$, then $2^e A^{-1} = PM(2^e S^{-1})U^{-1}$, where $e = \log_2 S_{nn}$.

Theorem 22. *Algorithm* Solve *in Figure 8 is correct. If $\log ||b|| \in O(n \log(n||A||))$ the running time is*

$$O(\mathsf{MM}(n, \log(n||A||)) \log(n \log ||A||) \log n).$$

```
Solve($A, b, n$)
Input: Nonsingular $A \in \mathbb{Z}^{n \times n}$ and $b \in \mathbb{Z}^{n \times 1}$.
Output: $x, e \in \mathbb{Z}^{n \times 1}, \mathbb{Z}_{\geq 0}$ such that $e$ is minimal such that all
        denominators of the entries in $2^e A^{-1}$ are relatively prime
        to 2, and $x = \mathrm{Rem}(2^e A^{-1} b, 2^d)$ where $d$ is as defined in
        step 3.
Note: $2^e A^{-1} b$ = RatRecon($x, 2^d, N, D$).
1.  $(P, S, M)$ := Massager($A, n$)
    $e := \log_2 S_{nn}$
2.  $U$ := ApplyMassager($A, n, P, S, M, n$)
3.  $N := \lfloor n^{n/2} ||A||^{n-1} ||b|| \rfloor$
    $D := \lfloor n^{n/2} ||A||^n / 2^e \rfloor$
    $d := \lceil \log(2ND) \rceil$
    $y$ := SpecialSolve($U, b, d, n, 1$)
4.  $x := \mathrm{Rem}(PM(2^e S^{-1})y, 2^d)$
    return $x, e$
```

Figure 8: Algorithm Solve

Proof. The correctness of the algorithm follows from the input and output specifications of Massager, ApplyMassager and SpecialSolve. Using Hadamard's bound and Cramer's rule, the denominators and numerators of entries of $2^e A^{-1} b$ are bounded by D and N as computed in the algorithm. This shows that the note added to the algorithm header also holds.

Now consider the running time. By Theorems 21, 18 and 7, the cost of steps 1, 2 and 3, respectively, are within the target cost.

Finally, consider step 4. Let $z = 2^e S^{-1} y$. The computation of Mz is very similar, in terms of structure and magnitudes of entries in the M and z, to the operation of combining a reduced index $(0, n-1)$ 2-Smith massager with an index $(n-1, 1)$ 2-Smith massager. By following the same strategy as in the proof of Theorem 20, step 4 can be done in time $O(\mathsf{MM}(n, \log(n||A||)))$. □

Acknowledgments

The authors are grateful to the referees for their careful reading, helpful comments, and suggestions for improving the exposition.

REFERENCES

[1] J. Bunch and J. Hopcroft. Triangular factorization and inversion by fast matrix multiplication. *Mathematics of Computation*, 28:231–236, 1974.

[2] J. von zur Gathen and J. Gerhard. *Modern Computer Algebra*. Cambridge University Press, 3rd edition, 2013.

[3] S. Gupta, S. Sarkar, A. Storjohann, and J. Valeriote. Triangular x-basis decompositions and derandomization of linear algebra algorithms over $\mathsf{K}[x]$. *Journal of Symbolic Computation*, 47(4), 2012. Festschrift for the 60th Birthday of Joachim von zur Gathen.

[4] O. Ibarra, S. Moran, and R. Hui. A generalization of the fast LUP matrix decomposition algorithm and applications. *Journal of Algorithms*, 3:45–56, 1982.

[5] C. Pauderis and A. Storjohann. Deterministic unimodularity certification. In J. van der Hoeven and M. van Hoeij, editors, *Proc. Int'l. Symp. on Symbolic and Algebraic Computation: ISSAC'12*, pages 281–288. ACM Press, New York, 2012.

[6] A. Storjohann. The shifted number system for fast linear algebra on integer matrices. *Journal of Complexity*, 21(4):609–650, 2005. Festschrift for the 70th Birthday of Arnold Schönhage.