



HAL
open science

Compressed Modular Matrix Multiplication

Jean-Guillaume Dumas, Laurent Fousse, Bruno Salvy

► **To cite this version:**

Jean-Guillaume Dumas, Laurent Fousse, Bruno Salvy. Compressed Modular Matrix Multiplication. 2008. hal-00259950v1

HAL Id: hal-00259950

<https://hal.science/hal-00259950v1>

Preprint submitted on 29 Feb 2008 (v1), last revised 13 Mar 2008 (v2)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Compressed Modular Matrix Multiplication

Jean-Guillaume Dumas* Laurent Fousse* Bruno Salvy†

February 29, 2008

Abstract

We propose to store several integers modulo a small prime into a single machine word. Modular addition is performed by addition and possibly subtraction of a word containing several times the modulo. Modular Multiplication is not directly accessible but modular dot product can be performed by an integer multiplication by the reverse integer. Modular multiplication by a word containing a single residue is also possible. Therefore matrix multiplication can be performed on such a compressed storage. We here give bounds on the sizes of primes and matrices for which such a compression is possible. We also explicit the details of the required compressed arithmetic routines.

1 Introduction

Compression of matrices over fields of characteristic 2 is naturally made via the binary representation of machine integers [1, 9].

The FFLAS/FFPACK project has demonstrated the need of a wrapping of cache-aware routines for efficient small finite field linear algebra [4, 5].

Therefore, a conversion between a modular representation of prime fields of any (small) characteristic and e.g. floating points can be performed via the homomorphism to the integers [2]. In [3] it is proposed to transform polynomial over a prime field into a q -adic representation where q is a prime power larger than the field characteristic. We call this transformation DQT for Discrete Q -adic Transform. With some care, in particular on the size of q , it is possible to map the polynomial operations into the floating point arithmetic realization of this q -adic representation and convert back using an inverse DQT.

In this note we propose to use this fast polynomial arithmetic within machine words to compute dot products. We show in section 2 how to recover a dot product of size $d + 1$ can be recovered from the single coefficient of degree d of a polynomial product. Whenever the prime modulus is small enough this enables to compute several accumulations of binary products in a single machine

*Laboratoire J. Kuntzmann, Université de Grenoble, umr CNRS 5224. BP 53X, 51, rue des Mathématiques, F38041 Grenoble, France. {Jean-Guillaume.Dumas,Laurent.Fousse}@imag.fr

†Projet ALGO, INRIA Rocquencourt, 78153 Le Chesnay. France. Bruno.Salvy@inria.fr

operation. Then we propose in section 3 an alternative matrix multiplication using multiplication of a compressed word by a single residue. The latter requires also a simultaneous modular reduction, called REDQ in [3]. In general, the prime field, the size of matrices and the available mantissa are given. This gives some constraints on the possible choices of q and d . In both cases anyway, we show that these compression techniques represent a speed-up factor of up to the number $d + 1$ of residues stored in the compressed format.

2 Q-adic compression or Dot product via polynomial multiplication

Suppose that $a(X) = \sum_{i=0}^d \alpha_i X^i$ and $b(X) = \sum_{i=0}^d \beta_i X^i$ are two polynomials in $\mathbb{Z}/p\mathbb{Z}[X]$. One can perform the dot product $\sum_{i=0}^d \alpha_i \beta_{d-i}$ by extracting the coefficient of degree d of $a(X)b(X)$.

2.1 Modular dot product via machine word multiplication

The idea here, as in [3], is to replace X by an integer q , usually a power of 2 in order to speed up conversions. Thus the vectors of residues $a = [\alpha_0 \dots \alpha_d]$ and $b = [\beta_0 \dots \beta_d]$ are stored respectively as $\bar{b} = \sum_{i=0}^d \beta_i q^i$ and the *reverse* $\bar{a} = \sum_{i=0}^d \alpha_{d-i} q^i$.

This is done e.g. over floating points via the following compressions:

```
double& init3( double& r,
               const double u, const double v, const double w) {
r=u; r*=_dBase; r+=v; r*=_dBase; return r+=w;
}
```

2.2 Gain

Now for matrix multiplication $A \times B$ one wishes to convert a whole row of the left $m \times k$ matrix A and a whole column of the right $k \times n$ matrix B .

Thus A is transformed into a $m \times \left\lceil \frac{k}{d+1} \right\rceil$ **CompressedRowMatrix**, CA and B is transformed into a $\left\lceil \frac{k}{d+1} \right\rceil \times n$ **CompressedRowMatrix**, CB .

Therefore the matrix multiply $CA \times CB$ can gain a *factor of $d + 1$* over the multiplication of $A \times B$, for classical multiplication. The precise gain will be given in table 2.

The result matrix $C = CA \times CB$ is $m \times n$. Thus in order to compare similar computations one has either to consider multiplication of compressed matrices which is then the procedure

$$C = CA \times CB; \text{ReduceAndCompress}(C) \tag{1}$$

or to consider multiplication of normal matrices via compression and thus the procedure

$$CA = \text{CompressRows}(A); CB = \text{CompressColumns}(B); C = CA \times CB \quad (2)$$

2.3 Partial compression

Note that the last column of CA and the last row of B might not have $d + 1$ elements if $\frac{k}{d+1} \notin \mathbb{Z}$. Thus one has to artificially append some zeroes to the converted values. On \bar{b} this means just do nothing. On the reversed \bar{a} this means multiplying by q several times.

2.4 Delayed reduction and lower bound on q

For the results to be correct the inner dot product must not exceed q . With a positive modular representation mod p (i.e. integers from 0 to $p - 1$), this means that $(d + 1)(p - 1)^2 < q$. Moreover, we would like to use delayed reductions on the intermediate results and thus accumulate the $\bar{a}\bar{b}$ before any modular reduction. It is thus possible to perform matrix multiplications of with common dimension k as long as:

$$\frac{k}{d+1}(d+1)(p-1)^2 = k(p-1)^2 < q. \quad (3)$$

2.5 Available mantissa and upper bound on q

If the product $\bar{a}\bar{b}$ is performed with floating point arithmetic we just need that the coefficient of degree d remains fully in the mantissa m . Write $\bar{a}\bar{b} = c_H q^d + c_L$, the latter means that c_H , and c_H only, must remain lower than 2^m . It could then be exactly recovered by multiplication of $\bar{a}\bar{b}$ by the correctly precomputed and rounded inverse of q^d and floored, as shown e.g. in [3, Lemma 2].

With delayed reduction this means that $\sum_{i=0}^d \frac{k}{d+1}(i+1)(p-1)^2 q^{d-i} < 2^m$. We can use equation 3 in order to show that $\sum_{i=0}^d \frac{k}{d+1}(i+1)(p-1)^2 q^{d-i} \leq q^{d+1}$. With this we just have to enforce that

$$q^{d+1} < 2^m. \quad (4)$$

Thus a single reduction has to be made at the end of the dot product as follows:

```
Element& init( Element& rem, const double dp) const {
double r = dp;
    // Multiply by the inverse of q^d with correct rounding
    r *= _inverseQto_d;
    // Now we just need the part less than q=2^t
    unsigned long rl( static_cast<unsigned long>(r) );
    rl &= _QMINUSONE;
```

```

// And we finally perform a single modular reduction
r1 %= _modulus;
return rem = static_cast<Element>(r1);
}

```

Note that to accelerate this reduction, one can avoid the multiplication by the inverse of q when $q = 2^t$: by adding $q^{2^{d+1}}$ to the final result one is guaranteed that the $t(d+1)$ high bits represent exactly the $d+1$ high coefficients. On the one hand, the multiplication can be advantageously replaced by a bit mask [7]. On the other hand, this doubles the size of the dot product and thus reduces by a factor of $\sqrt[2]{2^{d+1}}$ the largest possible dot product size k .

2.6 Results

One can see on figure 1 that the compression $(d+1)$ is very useful for small primes since the gain over the double floating point routine is quite close to d .




Figure 1: Compressed matrices multiplication of equation 1 compared with dgemm (the floating point double precision routine of GotoBlas) and fgemm (the exact routine of FFLAS) with double or single precision.

Indeed choosing a power of 2 for q simplifies and speeds up conversions and thus gives the following compression factors modulo 3:

Before $n = 256$ the compression is at a factor of five and the time to perform a matrix multiplication is less than a hundredth of a second. Then from 257 to 2048 one has a factor of 4 and the times are roughly 16 times the time of the

Compression	2	3..4	5..8	8	7	6	5	4	3
Degree d	16	12	9	7	6	5	4	3	2
q-adic	2^3	2^4	2^5	2^6	2^7	2^8	2^{10}	2^{13}	2^{17}
Dimensions	2	3..4	5..8	9..16	17..32	33..64	65..256	257..2048	2049..32768

Table 1: Compression factors for different common matrix dimensions modulo 3, with 53 bits of mantissa and q a power of 2.

four times smaller matrix, as visually shown on figure 2, left. The same is true afterwards with the respective factor 3 of compression.

Remark that the curve of fgemm with underlying arithmetic on single floats oscillates and drops. This is because the matrix begins to be too large and that modular reductions are required between the recursive matrix multiplication steps. Then the floating point BLAS¹ routines are used only when the sub-matrices are small enough. One can see the subsequent increase in the number of classical arithmetic steps on the drops around 2048, 4096 and 8192.

3 Right Compressed matrix multiplication

Another way of performing compressed matrix multiplication is to multiply an uncompressed matrix $m \times k$ to the right by a row-compressed $k \times \frac{n}{d+1}$ matrix. A dot product with this algorithm will be of the form $a = [a_0, \dots, a_n] \times [\sum_{j=0}^d b_{0j}q^j, \dots, \sum_{j=0}^d b_{nj}q^j]$. Therefore, a single entry of the resulting matrix will be $\sum_{i=0}^k a_i(\sum_{j=0}^d b_{ij}q^j) = \sum_{j=0}^d (\sum_{i=0}^k a_i b_{ij})q^j$. Here also q and d must satisfy equations (3) and (4).

The major difference is in the reductions. Indeed now one needs to reduce simultaneously the $d + 1$ coefficients of the polynomial in q in order to get the results. This simultaneous reduction can be made by the REDQ algorithm of [3, Algorithm 2].

Thus the whole right compressed matrix multiplication over two compressed matrices CA and CB , is the following algorithm as shown also on figure 2, right:

$$A = \text{Uncompress}(CA); CC = A \times CB; \text{REDQ}(CC) \quad (5)$$

4 Comparison

We summarize the differences of both algorithms in figure 2.

We see that the first algorithm compresses the common dimension whereas the Right (or also Left) compressions compress an external matrix dimension. Thus in the case of rectangular matrices one can choose between those routines the fastest one. This will be the routine compressing the largest dimension as shown on table 2. The table shows that the gain in terms of arithmetic

¹<http://www.tacc.utexas.edu/resources/software/>

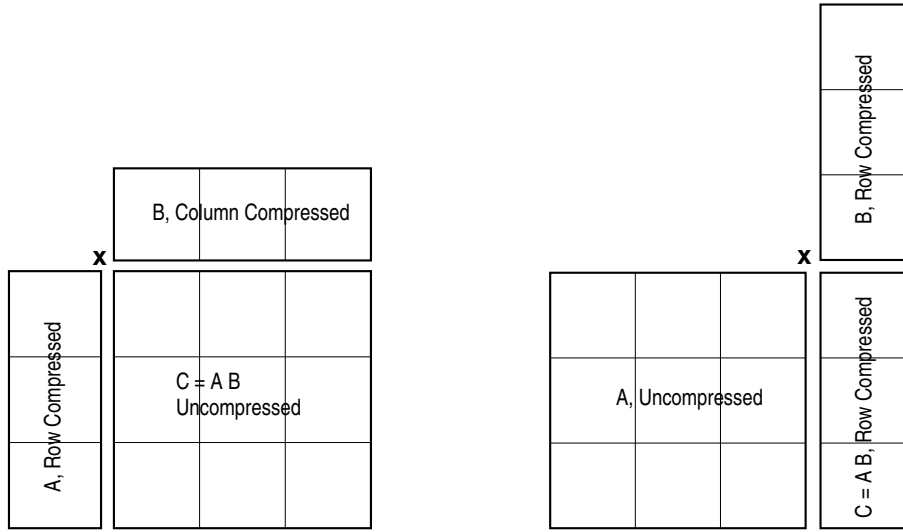


Figure 2: Algorithms of equations (1), left, and (5), right.

operations is $d + 1$ for classical multiplication and $(d + 1)^{\omega-2}$ when fast matrix multiplication is used as e.g. in [6, §3.2].

Algorithm	Operations	Reductions	Conversions
(1)	$\mathcal{O}\left(mn\left(\frac{k}{d+1}\right)^{\omega-2}\right)$	$m \times n \text{ REDC}$	$m \times \frac{n}{d+1} \text{ INIT}_d$
(5)	$\mathcal{O}\left(mk\left(\frac{n}{d+1}\right)^{\omega-2}\right)$	$m \times \frac{n}{d+1} \text{ REDQ}_d$	$m \times \frac{n}{d+1} \text{ EXTRACT}_d$
Left Comp.	$\mathcal{O}\left(nk\left(\frac{m}{d+1}\right)^{\omega-2}\right)$	$\frac{m}{d+1} \times n \text{ REDQ}_d$	$\frac{m}{d+1} \times n \text{ EXTRACT}_d$

Table 2: Number of operations for the different algorithms

For square matrices, the dimensions are equal and thus the number of arithmetic operations is the same for each variant. The difference there will mainly be on the number and on the kind of modular reductions. Since the REDQ_d reduction is faster than $d + 1$ classical reductions, see [3], and since INIT_d and EXTRACT_d are roughly the same operations, the best algorithm would then be one of the Left or Right compression.

For example, with algorithm (1) on matrices of sizes 10000×10000 it took 92.75 seconds to perform the matrix multiplication modulo 3 and 0.25 seconds to convert the resulting C matrix. This is less than 0.3%. For 250×250 matrices it takes less than 0.0028 seconds to perform the multiplication and roughly 0.00008 seconds for the conversions. There, the conversions count for 3%. Further work would thus include implementing the Right or Left compression and comparing

the conversions overhead.

References

- [1] Don Coppersmith. Solving linear equations over $gf(2)$: block Lanczos algorithm. *Linear Algebra and its Applications*, 192:33–60, October 1993.
- [2] Jean-Guillaume Dumas. Efficient dot product over finite fields. In Victor G. Ganzha, Ernst W. Mayr, and Evgenii V. Vorozhtsov, editors, *Proceedings of the seventh International Workshop on Computer Algebra in Scientific Computing, Yalta, Ukraine*, pages 139–154. Technische Universität München, Germany, July 2004.
- [3] Jean-Guillaume Dumas. Q-adic transform revisited. Technical Report 0710.0510 [cs.SC], ArXiv, October 2007. <http://hal.archives-ouvertes.fr/hal-00173894>.
- [4] Jean-Guillaume Dumas, Thierry Gautier, and Clément Pernet. Finite field linear algebra subroutines. In Teo Mora, editor, *Proceedings of the 2002 International Symposium on Symbolic and Algebraic Computation, Lille, France*, pages 63–74. ACM Press, New York, July 2002.
- [5] Jean-Guillaume Dumas, Pascal Giorgi, and Clément Pernet. FFPACK: Finite field linear algebra package. In Jaime Gutierrez, editor, *Proceedings of the 2004 International Symposium on Symbolic and Algebraic Computation, Santander, Spain*, pages 119–126. ACM Press, New York, July 2004.
- [6] Jean-Guillaume Dumas, Pascal Giorgi, and Clément Pernet. Dense linear algebra over prime fields. *ACM Transactions on Mathematical Software*, 2008. to appear.
- [7] Pascal Giorgi. Personal communication, 2008.
- [8] Chao H. Huang and Fred J. Taylor. A memory compression scheme for modular arithmetic. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 27(6):608–611, December 1979.
- [9] Erich Kaltofen and Austin Lobo. Distributed matrix-free solution of large sparse linear systems over finite fields. In A.M. Tentner, editor, *Proceedings of High Performance Computing 1996, San Diego, California*. Society for Computer Simulation, Simulation Councils, Inc., April 1996.