# The M4RI \& M4RIE libraries for linear algebra over $\mathbb{F}_{2}$ and small extensions 

Martin R. Albrecht


Sage/FLINT Days, 19.12.2011, Warwick (UK)

## Outline

M4RI<br>Multiplication<br>Elimination<br>Projects<br>M4RIE<br>Introduction<br>Newton-John Tables<br>Karatsuba Multiplication<br>Results



## Outline

M4RI<br>Multiplication<br>Elimination<br>Projects<br>M4RIE<br>Introduction<br>Newton-John Tables<br>Karatsuba Multiplication<br>Results




## M4RM [ADKF70] I

Consider $C=A \cdot B(A$ is $m \times \ell$ and $B$ is $\ell \times n)$.
A can be divided into $\ell / k$ vertical "stripes"

$$
A_{0} \ldots A_{(\ell-1) / k}
$$

of $k$ columns each. $B$ can be divided into $\ell / k$ horizontal "stripes"

$$
B_{0} \ldots B_{(\ell-1) / k}
$$

of $k$ rows each. We have:

$$
C=A \cdot B=\sum_{0}^{(\ell-1) / k} A_{i} \cdot B_{i} .
$$

## M4RM [ADKF70] II

$$
\begin{gathered}
A=\left(\begin{array}{llll}
1 & 1 & 0 & 1 \\
0 & 0 & 0 & 0 \\
1 & 1 & 1 & 1 \\
0 & 1 & 1 & 1
\end{array}\right), B=\left(\begin{array}{llll}
1 & 0 & 1 & 1 \\
0 & 1 & 1 & 0 \\
0 & 1 & 1 & 0 \\
0 & 1 & 0 & 1
\end{array}\right), A_{0}=\left(\begin{array}{ll}
1 & 1 \\
0 & 0 \\
1 & 1 \\
0 & 1
\end{array}\right) \\
A_{1}=\left(\begin{array}{ll}
0 & 1 \\
0 & 0 \\
1 & 1 \\
1 & 1
\end{array}\right), B_{0}=\left(\begin{array}{llll}
1 & 0 & 1 & 1 \\
0 & 1 & 1 & 0
\end{array}\right), B_{1}=\left(\begin{array}{llll}
0 & 1 & 1 & 0 \\
0 & 1 & 0 & 1
\end{array}\right) \\
A_{0} \cdot B_{0}=\left(\begin{array}{llll}
1 & 1 & 0 & 1 \\
0 & 0 & 0 & 0 \\
1 & 1 & 0 & 1 \\
0 & 1 & 1 & 0
\end{array}\right), A_{1} \cdot B_{1}=\left(\begin{array}{llll}
0 & 1 & 0 & 1 \\
0 & 0 & 0 & 0 \\
0 & 0 & 1 & 1 \\
0 & 0 & 1 & 1
\end{array}\right)
\end{gathered}
$$

M4RM: Algorithm $\mathcal{O}\left(n^{3} / \log n\right)$

1 begin


Algorithm 1: M4RM

## Strassen-Winograd [Str69] Multiplication

- fastest known pratical algorithm
- complexity: $\mathcal{O}\left(n^{\log _{2} 7}\right)$
- linear algebra constant: $\omega=\log _{2} 7$
- M4RM can be used as base case for small dimensions
$\rightarrow$ optimisation of this base case


## Cache Friendly M4RM I



## Cache Friendly M4RM II

```
1 begin
2 C & create an m\timesn matrix with all entries 0;
for 0}\leqstart<m/\mp@subsup{b}{s}{}\mathrm{ do
4 for 0 \leqi< \ell/k) do
                                    we regenerate T for each block
    T}\leftarrow\operatorname{MakeTABLE}(B,i\timesk,0,k)
                for 0}\leqs<\mp@subsup{b}{s}{}\mathrm{ do
                j\longleftarrowstart }\times\mp@subsup{b}{s}{}+s
                id \longleftarrow READBITS}(A,j,i\timesk,k)
                add row id from T to row j of C;
10 return C;
1 1 \text { end}
```


## $t>1$ Gray Code Tables I

- actual arithmetic is quite cheap compared to memory reads and writes
- the cost of memory accesses greatly depends on where in memory data is located
- try to fill all of L1 with Gray code tables.
- Example: $k=10$ and 1 Gray code table $\rightarrow 10$ bits at a time. $k=9$ and 2 Gray code tables, still the same memory for the tables but deal with 18 bits at once.
- The price is one extra row addition, which is cheap if the operands are all in cache.


## $t>1$ Gray Code Tables II

## 1 begin

$2 \quad C \longleftarrow$ create an $m \times n$ matrix with all entries 0 ;
3 for $0 \leq i<(\ell /(2 k))$ do
$T_{0} \leftarrow \operatorname{MakETABLE}(B, i \times 2 k, 0, k)$;
$T_{1} \leftarrow \operatorname{MAKETABLE}(B, i \times 2 k+k, 0, k)$;
for $0 \leq j<m$ do
$i d_{0} \longleftarrow \operatorname{ReadBits}(A, j, i \times 2 k, k) ;$
$i d_{1} \longleftarrow \operatorname{REAdBits}(A, j, i \times 2 k+k, k)$;
add row $i d_{0}$ from $T_{0}$ and row $i d_{1}$ from $T_{1}$ to row $j$ of $C$;

## Results: Multiplication



Figure: 2.66 Ghz Intel i7, 4GB RAM

## Small Matrices

M4RI is efficient for large matrices, but not necessarily for small matrices.

|  | Thomé | M4RI |
| :--- | :--- | :--- |
| transpose | $4.5097 \mu \mathrm{~s}$ | $0.6352 \mu \mathrm{~s}$ |
| copy | $0.2019 \mu \mathrm{~s}$ | $0.2674 \mu \mathrm{~s}$ |
| add | $0.2533 \mu \mathrm{~s}$ | $0.2921 \mu \mathrm{~s}$ |
| mul | $0.2535 \mu \mathrm{~s}$ | $0.4472 \mu \mathrm{~s}$ |

Table: $64 \times 64$ matrices (matops.c)

## Note

One performance bottleneck is that our matrix structure is much more complicated than Emmanuel's.

## Results: Multiplication Revisited



Figure: 2.66 Ghz Intel i7, 4GB RAM

## Outline

M4RI<br>Multiplication<br>Elimination<br>Projects<br>\section*{M4RIE}<br>Introduction<br>Newton-John Tables<br>Karatsuba Multiplication<br>Results




## PLE Decomposition I

## Definition (PLE)

> Let $A$ be a $m \times n$ matrix over a field $K$. A PLE decomposition of $A$ is a triple of matrices $P, L$ and $E$ such that $P$ is a $m \times m$ permutation matrix, $L$ is a unit lower triangular matrix, and $E$ is
> a $m \times n$ matrix in row-echelon form, and

$$
A=P L E .
$$

PLE decomposition can be in-place, that is $L$ and $E$ are stored in $A$ and $P$ is stored as an $m$-vector.

## PLE Decomposition II

From the PLE decomposition we can

- read the rank $r$,
- read the row rank profile (pivots),
- compute the null space,
- solve $y=A x$ for $x$ and
- compute the (reduced) row echelon form.
E. C.-P. Jeannerod, C. Pernet, and A. Storjohann. Fast gaussian elimination and the PLE decomposition. in preparation, 30 pages, 2011.


## Block Recursive PLE Decomposition $\mathcal{O}\left(n^{\omega}\right)$ I

Write

$$
A=\left(\begin{array}{ll}
A_{W} & A_{E}
\end{array}\right)=\left(\begin{array}{ll}
A_{N W} & A_{N E} \\
A_{S W} & A_{S E}
\end{array}\right)
$$

Main steps:

1. Call PLE on $A_{W}$
2. Apply row permutation to $A_{E}$
3. $L_{N W} \leftarrow$ the lower left triangular matrix in $A_{N W}$
4. $A_{N E} \leftarrow L_{N W}^{-1} \times A_{N E}$
5. $A_{S E} \leftarrow A_{S E}+A_{S W} \times A_{N E}$
6. Call PLE on $A_{S E}$
7. Apply row permutation to $A_{S W}$
8. Compress $L$

## Block Recursive PLE Decomposition $\mathcal{O}\left(n^{\omega}\right)$ II




## Block Recursive PLE Decomposition $\mathcal{O}\left(n^{\omega}\right)$ III



## Block Recursive PLE Decomposition $\mathcal{O}\left(n^{\omega}\right)$ IV



## Block Recursive PLE Decomposition $\mathcal{O}\left(n^{\omega}\right)$ V



## Block Recursive PLE Decomposition $\mathcal{O}\left(n^{\omega}\right)$ VI



## Block Recursive PLE Decomposition $\mathcal{O}\left(n^{\omega}\right)$ VII



## Block Iterative PLE Decomposition I

We need an efficient base case for PLE Decomposition

- block recursive PLE decomposition gives rise to a block iterative PLE decomposition
- choose blocks of size $k=\log n$ and use M4RM for the "update" multiplications
- this gives a complexity $\mathcal{O}\left(n^{3} / \log n\right)$


## Block Iterative PLE Decomposition II



## Block Iterative PLE Decomposition III



## Block Iterative PLE Decomposition IV



## Block Iterative PLE Decomposition V




## Block Iterative PLE Decomposition VI



## Block Iterative PLE Decomposition VII




## Block Iterative PLE Decomposition VIII



## Block Iterative PLE Decomposition IX



## Block Iterative PLE Decomposition X



[^0]
## Block Iterative PLE Decomposition XI



## Results: Reduced Row Echelon Form



Figure: 2.66 Ghz Intel i7, 4GB RAM

## Results: Row Echelon Form

Using one core - on sage.math - we can compute the echelon form of a $500,000 \times 500,000$ dense random matrix over $\mathbb{F}_{2}$ in

$$
9711 \text { seconds }=2.7 \text { hours }\left(c \approx 10^{-12}\right) .
$$

Using four cores decomposition we can compute the echelon form of a random dense $500,000 \times 500,000$ matrix in

$$
3806 \text { seconds }=1.05 \text { hours. }
$$

> Anybody got a 256GB RAM machine idlying around so that we can try $1,000,000 \times 1,000,000$ which should take about 20 hours on a single CPU? You know, for science!

## Outline

M4RI<br>Multiplication<br>Elimination<br>Projects<br>M4RIE<br>Introduction<br>Newton-John Tables<br>Karatsuba Multiplication<br>Results




## Sensitivity to Sparsity



Figure: Gaussian elimination of $10,000 \times 10,000$ matrices on Intel 2.33GHz Xeon E5345 comparing Magma 2.17-12 and M4RI 20111004.

## Linear Algebra for Gröbner Basis



| Problem | matrix dimensions | density | PLE | M4RI | GB |
| :---: | :---: | :---: | ---: | ---: | ---: |
| HFE 25 matrix 5 (5.1M) | $12307 \times 13508$ | 0.07600 | 1.03 | 0.59 | 0.81 |
| HFE 30 matrix 5 (16M) | $19907 \times 29323$ | 0.06731 | 4.79 | 2.70 | 4.76 |
| HFE 35 matrix 5 (37M) | $29969 \times 55800$ | 0.05949 | 19.33 | 9.28 | 19.51 |
| Mutant matrix (39M) | $26075 \times 26407$ | 0.18497 | 5.71 | 3.98 | 2.10 |
| random n=24, m=26 matrix 3 (30M) | $37587 \times 38483$ | 0.03832 | 20.69 | 21.08 | 19.36 |
| random n=24, m=26 matrix 4 (24M) | $37576 \times 32288$ | 0.04073 | 18.65 | 28.44 | 17.05 |
| SR(2,2,2,4) compressed, matrix 2 (328K) | $5640 \times 14297$ | 0.00333 | 0.40 | 0.29 | 0.18 |
| SR(2,2,2,4) compressed, matrix 4 (2.4M) | $13665 \times 17394$ | 0.01376 | 2.18 | 3.04 | 2.04 |
| SR(2,2,2,4) compressed, matrix 5 (2.8M) | $11606 \times 16282$ | 0.03532 | 1.94 | 4.46 | 1.59 |
| SR(2,2,2,4) matrix 6 (1.4M) | $13067 \times 17511$ | 0.00892 | 1.90 | 2.09 | 1.38 |
| SR(2,2,2,4) matrix 7 (1.7M) | $12058 \times 16662$ | 0.01536 | 1.53 | 1.93 | 1.66 |
| SR(2,2,2,4) matrix 9 $(36 M)$ | $115834 \times 118589$ | 0.00376 | 528.21 | 578.54 | 522.98 |

## Multi-core Support



Parallel speed-up


M4RI BOpS \& Speed-up



PLE BOpS \& Speed-up

## GF(2) on GFX

Tabelle 3.12: Zeiten auf der GeForce GTX 295 und GeForce GTX 480.

| Matrixgröße | GeForce GTX 295 | GeForce GTX 480 |
| ---: | :---: | :---: |
| $9.984 \times 10.240$ | 0,9 Sek. | 1,2 Sek. |
| $16.384 \times 16.384$ | 2,47 Sek. | 2,9 Sek. |
| $20.000 \times 20.480$ | 4,63 Sek. | 4,63 Sek. |
| $32.000 \times 32.768$ | 13,3 Sek. | 12,2 Sek. |
| $64.000 \times 65.536$ | - | 70,74 Sek. |

Tabelle 3.13: Zeiten auf der CPU [6].

| Matrix Dimension | M4RI/M4RI <br> $20090105^{7}$ | M4RI/M4RI <br> $20100817^{2}$ |
| :---: | :---: | :---: |
| $10.000 \times 10.000$ | 1,532 | 1,050 |
| $16.384 \times 16.384$ | 6,597 | 3,890 |
| $20.000 \times 20.000$ | 12,031 | 7,250 |
| $32.000 \times 32.000$ | 40,768 | 22,560 |
| $64.000 \times 64.000$ | 241,017 | 124,480 |
| [1] 64- bit Debian/GNU Linux, 2.33 Ghz Core2Duo (Macbook Pro, 2nd. Gen.) |  |  |
| [2] 64- bit Debian/GNU Linux, 2.6 Ghz Intel i7 (Macbook Pro 6,2) |  |  |

E. Denise Demirel

Effizientes Lösen linearer Gleichungssysteme über GF(2) mit GPUs
Diplomarbeit, TU Darmstadt, September 2010

## Outline

## M4RI

Multiplication
Elimination
Projects
M4RIE
Introduction
Newton-John Tables
Karatsuba Multiplication
Results



## Motivation I

Your NTL patch worked perfectly for me first try. I tried more benchmarks (on Pentium-M 1.8Ghz):

```
[...] //these are for GF(2^8), malb
sage: n=1000; m=ntl.mat_GF2E(n,n,[ ntl.GF2E_random() for i in xrange(n^2) ])
sage: time m.echelon_form()
1000
Time: CPU 29.72 s, Wall: 43.79 s
```

This is pretty good; vastly better than what's was in SAGE by default, and way better than PARI. Note that MAGMA is much faster though (nearly 8 times faster):
[...]
$>\mathrm{n}:=1000 ; \mathrm{A}:=\operatorname{MatrixAlgebra}\left(\mathrm{GF}\left(2^{\wedge} 8\right), \mathrm{n}\right)!\left[\operatorname{Random}\left(\operatorname{GF}\left(2^{\wedge} 8\right)\right):\right.$ i in $\left.\left[1 . . \mathrm{n}^{\wedge} 2\right]\right]$;
> time E := EchelonForm(A);
Time: 3.440

MAGMA uses (1) [...] and (2) a totally different algorithm for computing the echelon form. [...] As far as I know, the MAGMA method is not implemented anywhere in the open source world. But l'd love to be wrong about that... or even remedy that.

- W. Stein in 01/2006 replying to my 1st non-trivial patch to Sage


## Motivation II

The situation has not improved much in 2011:

| System | Time in ms |
| :--- | ---: |
| Sage 4.7.2 | 97,000 |
| NTL 5.4.2 | 85,000 |
| LinBox SVN + patches | 460 |
| GAP 4.412 | 210 |
| Magma 2.15 | 13 |
| this work | 5.5 |

Table: Product of two dense $1,000 \times 1,000$ matrix over $\mathbb{F}_{2^{2}}$.

## Representation of Elements I

Elements in $\mathbb{F}_{2^{e}} \cong \mathbb{F}_{2}[x] / f$ can be written as

$$
a_{0} \alpha^{0}+a_{1} \alpha^{1}+\cdots+a_{e-1} \alpha^{e-1} .
$$

We identify the bitstring $a_{0}, \ldots, a_{e-1}$ with

- the element $\sum_{i=0}^{e-1} a_{i} \alpha^{i} \in \mathbb{F}_{2^{e}}$ and
- the integer $\sum_{i=0}^{e-1} a_{i} 2^{i}$.

In the datatype mzed_t we pack several of those bitstrings into one machine word:
$a_{0,0,0}, \ldots, a_{0,0, e-1}, a_{0,1,0}, \ldots, a_{0,1, e-1}, \ldots, a_{0, n-1,0}, \ldots, a_{0, n-1, e-1}$.

Additions are cheap, scalar multiplications are expensive.

## Representation of Elements II

- Instead of representing matrices over $\mathbb{F}_{2^{e}}$ as matrices over polynomials we may represent them as polynomials with matrix coefficients.
- For each degree we store matrices over $\mathbb{F}_{2}$ which hold the coefficients for this degree.
- The data type mzd_slice_t for matrices over $\mathbb{F}_{2^{e}}$ internally stores e-tuples of M4RI matrices, i.e., matrices over $\mathbb{F}_{2}$.


## Additions are cheap, scalar multiplications are expensive.

## Representation of Elements III

$$
\begin{aligned}
A & =\left(\begin{array}{cc}
\alpha^{2}+1 & \alpha \\
\alpha+1 & 1
\end{array}\right) \\
& =\left[\begin{array}{ll}
\square 101 & \square 010 \\
\square 011 & \square 001
\end{array}\right] \\
& =\left(\left[\begin{array}{ll}
1 & 0 \\
0 & 0
\end{array}\right],\left[\begin{array}{ll}
0 & 1 \\
1 & 0
\end{array}\right],\left[\begin{array}{ll}
1 & 0 \\
1 & 1
\end{array}\right]\right)
\end{aligned}
$$

Figure: $2 \times 2$ matrix over $\mathbb{F}_{8}$

## Outline

M4RI<br>Multiplication<br>Elimination<br>Projects<br>M4RIE<br>Introduction<br>Newton-John Tables<br>Karatsuba Multiplication<br>Results




## The idea I

Input: $A-m \times n$ matrix
Input: $B-n \times k$ matrix
1 begin
2 for $0 \leq i<m$ do
$3 \quad$ for $0 \leq j<n$ do
4
$L C_{j} \longleftarrow C_{j}+A_{j, i} \times B_{i} ;$
5 return C;
6 end

## The idea II

Input: $A-m \times n$ matrix
Input: $B-n \times k$ matrix
1 begin
2 for $0 \leq i<m$ do
for $0 \leq j<n$ do
$L C_{j} \longleftarrow C_{j}+A_{j, i} \times B_{i} ; / /$ cheap
5 return C;
6 end

## The idea III

Input: $A-m \times n$ matrix
Input: $B-n \times k$ matrix
1 begin
2 for $0 \leq i<m$ do
for $0 \leq j<n$ do
$\left\lfloor C_{j} \longleftarrow C_{j}+A_{j, i} \times B_{i} ; / /\right.$ expensive
5 return C;
6 end

## The idea IV

Input: $A-m \times n$ matrix
Input: $B-n \times k$ matrix
1 begin
2 for $0 \leq i<m$ do
$3 \quad$ for $0 \leq j<n$ do
$L C_{j} \longleftarrow C_{j}+A_{j, i} \times B_{i} ; / /$ expensive
5 return C;
6 end

But there are only $2^{e}$ possible multiples of $B_{i}$.


## The idea V

|  |  |
| :---: | :---: |
|  | Input: $A-m \times n$ matrix |
|  | Input: $B-n \times k$ matrix |
| 2 | for $0 \leq i<m$ do |
| 3 | for $0 \leq j<2^{e}$ do |
| 4 | $T_{j} \longleftarrow j \times B_{i} ;$ |
| 5 | for $0 \leq j<n$ do |
| 6 | $x \longleftarrow A_{j, i}$ |
| 7 | $C_{j} \longleftarrow C_{j}+T_{\chi} ;$ |
| 8 | return $C$; |
|  |  |

$m \cdot n \cdot k$ additions, $m \cdot 2^{e} \cdot k$ multiplications.

## Gaussian elimination \& PLE decomposition

Input: $A-m \times n$ matrix
1 begin

| 2 | $r \longleftarrow 0 ;$ |
| :--- | :--- |
| 3 | for $0 \leq j<n$ do |

$$
\text { for } r \leq i<m \text { do }
$$

$$
\text { if } A_{i, j}=0 \text { then continue; }
$$

$$
\text { rescale row } i \text { of } A \text { such that } A_{i, j}=1 \text {; }
$$

swap the rows $i$ and $r$ in $A$;
$T \longleftarrow$ multiplication table for row $r$ of $A$;

```
                for r +1 \leqk<m do
```

                    \(x \longleftarrow A_{k, j} ;\)
    $A_{k} \longleftarrow A_{k}+T_{x} ;$

$$
r \longleftarrow r+1 ;
$$

return $r$;
14 end

## Outline

## M4RI

Multiplication
Elimination
Projects
M4RIE
Introduction
Newton-John Tables
Karatsuba Multiplication
Results



## The idea

- Consider $\mathbb{F}_{2^{2}}$ with the primitive polynomial $f=x^{2}+x+1$.
- We want to compute $C=A B$.
- Rewrite $A$ as $A_{0} x+A_{1}$ and $B$ as $B_{0} x+B_{1}$.
- The product is

$$
C=A_{0} B_{0} x^{2}+\left(A_{0} B_{1}+A_{1} B_{0}\right) x+A_{1} B_{1} .
$$

- Reduction modulo $f$ gives

$$
C=\left(A_{0} B_{0}+A_{0} B_{1}+A_{1} B_{0}\right) x+A_{1} B_{1}+A_{0} B_{0} .
$$

- This last expression can be rewritten as

$$
C=\left(\left(A_{0}+A_{1}\right)\left(B_{0}+B_{1}\right)+A_{1} B_{1}\right) x+A_{1} B_{1}+A_{0} B_{0} .
$$

Thus this multiplication costs 3 multiplications and 4 adds over $\mathbb{F}_{2}$.

## Outline

M4RI<br>Multiplication<br>Elimination<br>Projects<br>\section*{M4RIE}<br>Introduction<br>Newton-John Tables<br>Karatsuba Multiplication<br>Results




## Results: Multiplication I

| $e$ | Magma <br> $2.15-10$ | GAP <br> 4.4 .12 | SW-NJ | SW-NJ/ <br> M4RI | [Mon05] | Bitslice | Bitslice/ |
| :---: | ---: | ---: | ---: | ---: | ---: | ---: | ---: |
| M4RI |  |  |  |  |  |  |  |
| 1 | 0.100 s | 0.244 s | - | 1 | 1 | 0.071 s | 1.0 |
| 2 | 1.220 s | 12.501 s | 0.630 s | 8.8 | 3 | 0.224 s | 3.1 |
| 3 | 2.020 s | 35.986 s | 1.480 s | 20.8 | 6 | 0.448 s | 6.3 |
| 4 | 5.630 s | 39.330 s | 1.644 s | 23.1 | 9 | 0.693 s | 9.7 |
| 5 | 94.740 s | 86.517 s | 3.766 s | 53.0 | 13 | 1.005 s | 14.2 |
| 6 | 89.800 s | 85.525 s | 4.339 s | 61.1 | 17 | 1.336 s | 18.8 |
| 7 | 82.770 s | 83.597 s | 6.627 s | 93.3 | 22 | 1.639 s | 23.1 |
| 8 | 104.680 s | 83.802 s | 10.170 s | 143.2 | 27 | 2.140 s | 30.1 |

Table: Multiplication of $4,000 \times 4,000$ matrices over $\mathbb{F}_{2^{e}}$

## Results: Multiplication II

Multiplication: Magma vs. Sage


Figure: 2.66 Ghz Intel i7, 4GB RAM

## Results: Reduced Row Echelon Forms I

| e | Magma | GAP | M4RIE |
| :---: | :---: | :---: | :---: |
| 2 | 6.040s | 162.658s | 3.310s |
| 3 | 14.470 s | 442.522s | 5.332s |
| 4 | 60.370s | 502.672s | 6.330s |
| 5 | 659.030s | N/A | 10.511s |
| 6 | 685.460s | N/A | 13.078s |
| 7 | 671.880s | N/A | 17.285s |
| 8 | 840.220s | N/A | 20.247s |
| 9 | 1630.380s | N/A | 260.774s |
| 10 | 1631.350s | N/A | 291.298s |

Table: Elimination of $10,000 \times 10,000$ matrices

## Results: Reduced Row Echelon Forms II

Multiplication: Magma vs. Sage


Figure: 2.66 Ghz Intel i7, 4GB RAM

Fin
E. Vrlazarov, E. Dinic, M. Kronrod, and I. Faradzev. On economical construction of the transitive closure of a directed graph.
Dokl. Akad. Nauk., 194(11), 1970.
(in Russian), English Translation in Soviet Math Dokl.
E. Peter L. Montgomery.

Five, six, and seven-term Karatsuba-like formulae.
IEEE Trans. on Computers, 53(3):362-369, 2005.
Volker Strassen.
Gaussian elimination is not optimal.
Nummerische Mathematik, 13:354-256, 1969.


[^0]:    

