

Error Correction in Fast Matrix Multiplication and Inverse

Daniel S. Roche



Computer Science Department
United States Naval Academy
Annapolis, Maryland, U.S.A.

ISSAC 2018
19 July 2018
New York City

Correcting Errors

General Framework

- Some problem is slow to solve locally.
- So we outsource the computation and get back the result.
- If the result is wrong, we should **detect** that.
- If the result is “close”, we should **correct** it.

Sources of error

- Noisy communication channel
- Faults in computation
- Malicious servers?

Example: Sudoku

Sent to server

	7		8		2			
								7
1	4	8	3	5				
		5			8	3		
	6							9
3						7		
5						1	7	
		4		9				
		7	4					2

Response

9	7	6	8	1	2	4	5	3
2	5	3	9	6	4	8	1	7
1	4	5	3	8	7	9	2	6
7	9	8	6	2	5	3	4	1
4	6	1	5	7	3	2	8	9
3	8	2	1	4	9	7	6	5
5	3	9	2	9	6	1	7	4
6	2	4	7	9	1	5	3	8
8	1	7	4	3	8	6	9	2

Source: <https://www.nytimes.com/crosswords/game/sudoku/hard>

Example: Sudoku

Sent to server

	7		8	2				
								7
1	4	8	3	5				
		5			8	3		
	6							9
3						7		
5						1	7	
		4		9				
		7	4					2

Response

9	7	6	8	1	2	4	5	3
2	5	3	9	6	4	8	1	7
1	4	5	3	8	7	9	2	6
7	9	8	6	2	5	3	4	1
4	6	1	5	7	3	2	8	9
3	8	2	1	4	9	7	6	5
5	3	9	2	9	6	1	7	4
6	2	4	7	9	1	5	3	8
8	1	7	4	3	8	6	9	2

Source: <https://www.nytimes.com/crosswords/game/sudoku/hard>

First Problem: Matrix Multiplication

Matrix Multiplication with Errors

Input: $A, B, C \in \mathbb{F}^{n \times n}$, where $C \approx AB$

Output: $E \in \mathbb{F}^{n \times n}$ s.t. $AB = C - E$

Parameters:

- Dimension n
- Input size $t = \#A + \#B + \#C$
- Error count $k = \#E$

First Problem: Matrix Multiplication

Matrix Multiplication with Errors

Input: $A, B, C \in \mathbb{F}^{n \times n}$, where $C \approx AB$

Output: $E \in \mathbb{F}^{n \times n}$ s.t. $AB = C - E$

Parameters:

- Dimension n
 - Input size $t = \#A + \#B + \#C$
 - Error count $k = \#E$
-
- Naïve recomputation: $O(n^\omega)$
 - Gasieniec, Levcopoulos, Lingas, Pagh, Tokuyama, *Algorithmica* 2017:
 $\tilde{O}(n^2 + kn)$

First Problem: Matrix Multiplication

Matrix Multiplication with Errors

Input: $A, B, C \in \mathbb{F}^{n \times n}$, where $C \approx AB$

Output: $E \in \mathbb{F}^{n \times n}$ s.t. $AB = C - E$

Parameters:

- Dimension n
 - Input size $t = \#A + \#B + \#C$
 - Error count $k = \#E$
-
- Naïve recomputation: $O(n^\omega)$
 - Gasieniec, Levcopoulos, Lingas, Pagh, Tokuyama, *Algorithmica* 2017:
 $\tilde{O}(n^2 + kn)$
 - Ours: $\tilde{O}(t + kn)$ worst-case; $\tilde{O}(t + k^{0.686}n)$ best case

Related Work

Monte Carlo **verification** in Linear Algebra

- Matrix multiplication: Freivalds '79
- Positive semidefiniteness: Kaltofen, Nehring, Saunders '11
- Rank, characteristic polynomial, . . . : Dumas & Kaltofen '14
- Rank profile, triangular forms: Dumas, Lucas, Pernet '17

Error correction in symbolic computation

Complexity of sparse matrix multiplication

Related Work

Monte Carlo **verification** in Linear Algebra

Error correction in symbolic computation

- Chinese remaindering: Goldreich, Ron, Sudan '99
- Rational reconstruction: Khonji, Pernet, Roch, Roche, Stalinski '10
- Sparse interpolation: Comer, Kaltofen, Pernet '12

Complexity of sparse matrix multiplication

Related Work

Monte Carlo **verification** in Linear Algebra

Error correction in symbolic computation

Complexity of sparse matrix multiplication

k = number of nonzeros

- $\tilde{O}(n^2 + k^{0.7}n^{1.2})$: Yuster & Zwick '04
- $\tilde{O}(k^{1.34})$ output sensitive: Amossen & Pagh '09
- $\tilde{O}(k^{0.188}n^2)$: Lingas '10
- $\tilde{O}(n^2 + kn)$ output sensitive: Pagh '13

Outline

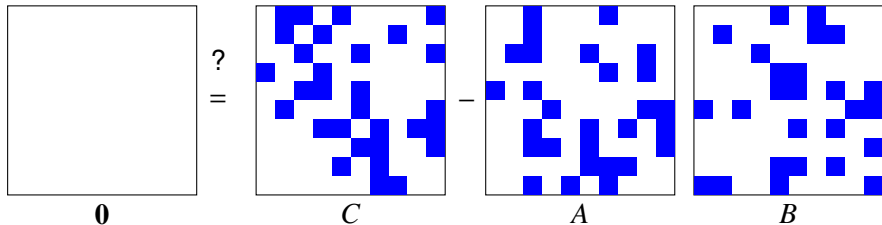
- Introduction
- Matrix Multiplication with Errors
Crucial tools: Nonzero row identification, Sparse interpolation
- Matrix Inverse with Errors
Additional tools: Commutativity, Sparse low-rank linear algebra
- Preliminary implementation timings

Detecting errors

R. Freivalds: Fast Probabilistic Algorithms (1979)

Given A, B, C , does $AB = C$?

- 1 Choose random vector v
- 2 Compute $(C - AB)v$
- 3 Check whether $(C - AB)v = \mathbf{0}$

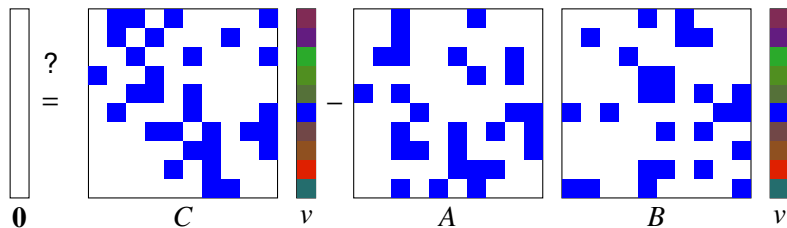


Detecting errors

R. Freivalds: Fast Probabilistic Algorithms (1979)

Given A, B, C , does $AB = C$?

- 1 Choose random vector v
- 2 Compute $(C - AB)v$
- 3 Check whether $(C - AB)v = \mathbf{0}$

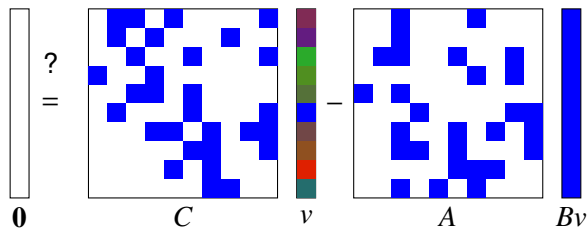


Detecting errors

R. Freivalds: Fast Probabilistic Algorithms (1979)

Given A, B, C , does $AB = C$?

- 1 Choose random vector v
- 2 Compute $(C - AB)v$
- 3 Check whether $(C - AB)v = \mathbf{0}$

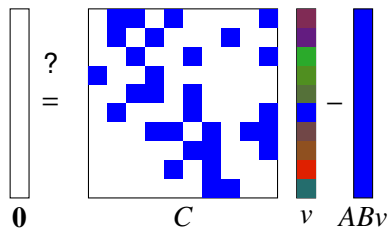


Detecting errors

R. Freivalds: Fast Probabilistic Algorithms (1979)

Given A, B, C , does $AB = C$?

- 1 Choose random vector v
- 2 Compute $(C - AB)v$
- 3 Check whether $(C - AB)v = \mathbf{0}$



Detecting errors

R. Freivalds: Fast Probabilistic Algorithms (1979)

Given A, B, C , does $AB = C$?

- 1 Choose random vector v
- 2 Compute $(C - AB)v$
- 3 Check whether $(C - AB)v = \mathbf{0}$

$$\mathbf{0} = Cv - ABv$$

Row Polynomials

Treat each row of $M \in \mathbb{F}^{n \times n}$ as a degree- $(n - 1)$ polynomial

Example

The diagram illustrates the process of treating the rows of a matrix M as polynomials. On the left, a square matrix is shown with blue squares representing non-zero entries. The matrix is sparse, with non-zero entries scattered across its rows and columns. This matrix is equated to a vertical vector of polynomials. The first row of the matrix corresponds to the polynomial 1 . The second row corresponds to x . The third row corresponds to x^2 . The fourth row corresponds to x^3 . The fifth row corresponds to $x^5 + x^7$. The sixth row corresponds to 0 . The seventh row corresponds to x^9 . The eighth row corresponds to a vertical ellipsis. This vector of polynomials is equated to a vertical vector of functions $f_1(x)$, $f_2(x)$, $f_3(x)$, and a vertical ellipsis.

$$\begin{array}{|c|} \hline 1 \\ \hline x \\ \hline x^2 \\ \hline x^3 \\ \hline \vdots \\ \hline \end{array} = \begin{array}{|c|} \hline \blacksquare x^5 + \blacksquare x^7 \\ \hline 0 \\ \hline \blacksquare x^9 \\ \hline \vdots \\ \hline \end{array} = \begin{array}{|c|} \hline f_1(x) \\ \hline f_2(x) \\ \hline f_3(x) \\ \hline \vdots \\ \hline \end{array}$$

Row Polynomial Evaluation

- Need to evaluate each row polynomial $f_i(x)$ at $1, \theta, \theta^2, \dots$
- This is equivalent to multiplying by a (truncated) DFT matrix.

Example

$$\begin{array}{c}
 \boxed{\begin{array}{cccc} & & \blacksquare & \blacksquare \\ & \blacksquare & & \\ & & \blacksquare & \blacksquare \\ \blacksquare & & & \\ & \blacksquare & & \blacksquare \\ & & \blacksquare & \blacksquare \\ \blacksquare & & & \\ & \blacksquare & & \blacksquare \\ & & \blacksquare & \blacksquare \\ \blacksquare & & & \\ & \blacksquare & & \blacksquare \\ & & \blacksquare & \blacksquare \\ \blacksquare & & & \\ & \blacksquare & & \blacksquare \\ & & \blacksquare & \blacksquare \end{array}} \\
 M \\
 n \times n
 \end{array}
 \quad
 \begin{array}{c}
 \boxed{\begin{array}{c} \text{red gradient bar} \\ \theta^{i+j} \end{array}} \\
 V \\
 n \times s
 \end{array}
 =
 \begin{array}{c}
 \boxed{\begin{array}{cccc}
 f_1(1) & f_1(\theta) & f_1(\theta^2) & \cdots \\
 f_2(1) & f_2(\theta) & f_2(\theta^2) & \cdots \\
 \vdots & \vdots & \vdots & \vdots
 \end{array}} \\
 MV \\
 n \times s
 \end{array}$$

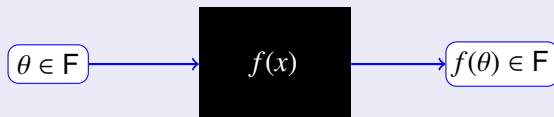
Important: Takes $\tilde{O}(sn + t)$ field ops, where $t = \#M$

Sparse polynomial interpolation

Overview

Given a way to evaluate an unknown polynomial $f \in \mathbb{F}[x]$
with at most s nonzero terms,
how to recover the nonzero coefficients and exponents of f ?

Polynomial black box



Sparse polynomial interpolation

Prony's Method

Prony's (1795) method from exponential analysis can be used here.
(Ben-Or & Tiwari '88, Kaltofen Lakshman & Wiley '91, . . .)

Algorithm Overview

1. Find a high-order element $\omega \in \mathbb{F}$
2. Compute $f(\omega^i)$ for $i = 0, 1, \dots, 2s - 1$
3. Berlekamp-Massey to find annihilator $\Lambda(z)$
4. Find roots ω^{e_i} of Λ
5. Compute discrete logs e_i of Λ
6. Transp. Vandermonde solve to get coefficients c_i

Sparse polynomial interpolation

Prony's Method

Prony's (1795) method from exponential analysis can be used here.
(Ben-Or & Tiwari '88, Kaltofen Lakshman & Wiley '91, ...)

Algorithm Overview

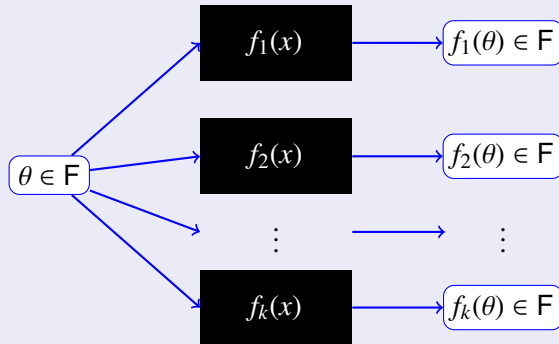
1. Find a high-order element $\omega \in F$
2. Compute $f(\omega^i)$ for $i = 0, 1, \dots, 2s - 1$
3. Berlekamp-Massey to find annihilator $\Lambda(z)$
4. Find roots ω^{e_i} of Λ
5. Compute discrete logs e_i of Λ
6. Transp. Vandermonde solve to get coefficients c_i

- These steps depend on the field F
- Root finding and discrete logarithms are generally the bottleneck (especially over finite fields!)

Batched Sparse Polynomial Interpolation

What if we **simultaneously** interpolate k sparse polynomials?

Simultaneous polynomial black box



Batched Sparse Polynomial Interpolation

Algorithm

Challenge: Simultaneously interpolate $f_1, \dots, f_k \in F[x]$, each having $\leq s$ nonzero terms and degree $\leq n$.

Idea: Precomputation of

- High-order element $\omega \in F$
- First n powers of ω

Result: Recovery in $\tilde{O}(n + sk)$ time for **any field F**

Correcting errors in matrix multiplication

Input: $A, B, C \in \mathbb{F}^{n \times n}$

Output: $E \in \mathbb{F}^{n \times n}$ s.t. $AB = C - E$

Algorithm overview:

- 1 Formulate as $E = C - AB$
- 2 Find nonzero rows
Choose random $v \in \mathbb{F}^n$ and compute Ev
- 3 Evaluate row polynomials
For each evaluation θ , compute $E[1, \theta, \dots, \theta^{n-1}]^T$
- 4 Interpolate row polynomials
Sparse interpolation to recover at least half of the nonzero rows
- 5 Repeat $O(\log k)$ times until all rows are found

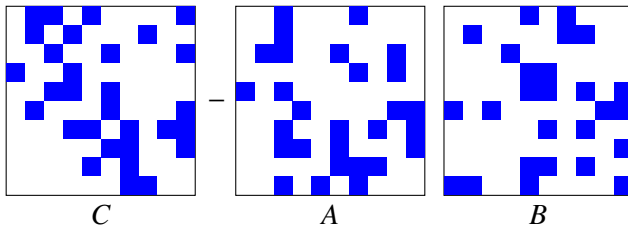
Multiplication Algorithm Overview

Given A, B, C , we want to compute $E = C - AB$



=

E



Multiplication Algorithm Overview

1 Find zero rows of E

$$E = Cv - ABv$$

The diagram shows a 5x3 matrix E with two rows of question marks. To its right is a 5x1 column vector v with colored blocks. Below this is the expression $Cv - ABv$, where Cv is a 5x3 grid of colored blocks and ABv is a 5x3 grid of colored blocks.

$$C - AB = E$$

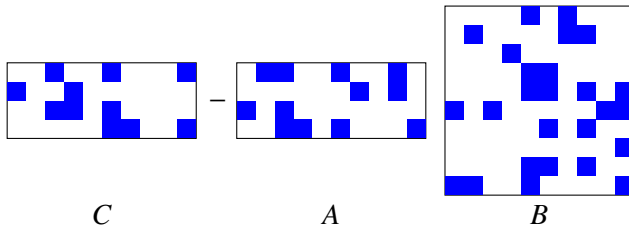
The diagram shows three 5x5 matrices: C , A , and B , each with a 5x1 column vector v to its right. C is a sparse matrix of blue blocks. A and B are also sparse matrices of blue blocks. The equation $C - AB = E$ is shown above the matrices. E is a 5x3 grid of question marks.

Multiplication Algorithm Overview

2 Remove corresponding rows from C and A

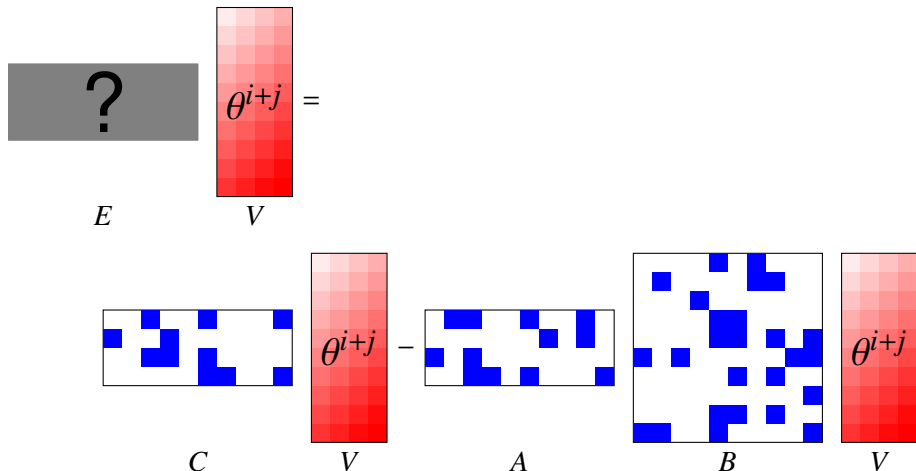
? =

E



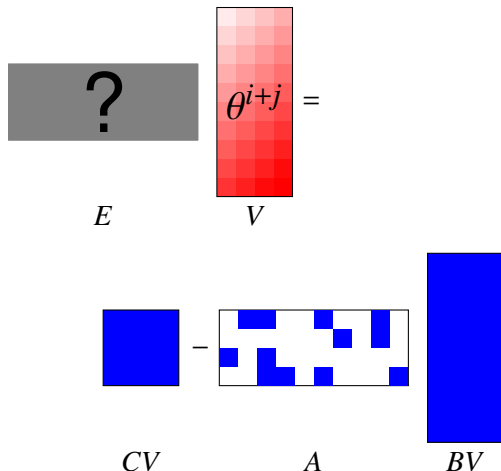
Multiplication Algorithm Overview

- 3 Choose $s \leq n$ and evaluate row polynomials at 2^s powers of θ



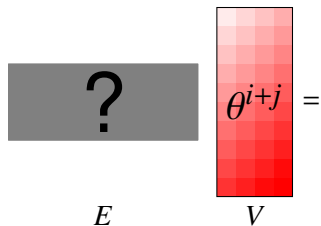
Multiplication Algorithm Overview

- Choose $s \leq n$ and evaluate row polynomials at 2^s powers of θ



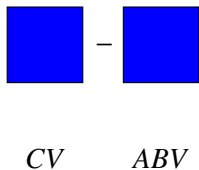
Multiplication Algorithm Overview

- 3 Choose $s \leq n$ and evaluate row polynomials at $2s$ powers of θ



E

V



CV

ABV

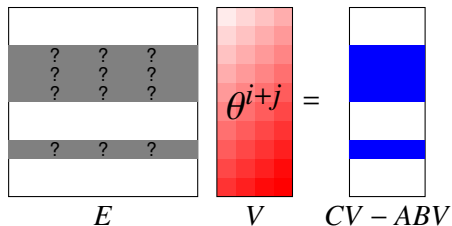
Multiplication Algorithm Overview

- 3 Choose $s \leq n$ and evaluate row polynomials at $2s$ powers of θ

$$E \quad V \quad = \quad CV - ABV$$

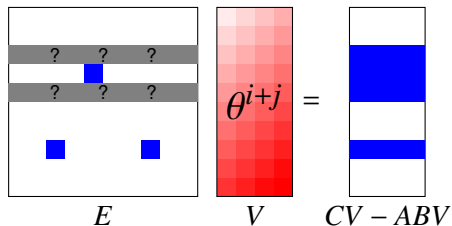
Multiplication Algorithm Overview

- 3 Choose $s \leq n$ and evaluate row polynomials at 2^s powers of θ



Multiplication Algorithm Overview

- 4 Perform sparse interpolation to recover s -sparse rows of E



Choosing the number of evaluations

Each row polynomial f_1, \dots, f_r is evaluated at $\omega^0, \dots, \omega^{2^s-1}$

Question: How large should s be?

- To recover every row, need $s \geq \max_i \#f_i$
Therefore s could be as large as $k - r = O(k)$

Choosing the number of evaluations

Each row polynomial f_1, \dots, f_r is evaluated at $\omega^0, \dots, \omega^{2s-1}$

Question: How large should s be?

- To recover every row, need $s \geq \max_i \#f_i$
Therefore s could be as large as $k - r = O(k)$
- To recover **at least half** of the rows, only need

$$s \geq \frac{2k}{r}$$

Then sparse interpolation costs $\tilde{O}(rs + n) = \tilde{O}(k + n)$

Iterative updates

The problem can be transposed easily:

- $E = C - AB$
- $E^T = C^T - B^T A^T$

After each step, **at least half** of the rows are completely recovered.

We **transpose** at each step to minimize r .

Multiplication Complexity

Total cost in field operations is

$$\tilde{O}\left(n + t + k \cdot \min\left(\frac{t}{r} + 1, \frac{n}{\min(r, \frac{k}{r})^{3-\omega}}\right)\right)$$

- n = dimension
- t = number of nonzeros in A, B, C
- k = number of errors (nonzeros in E)
- r = number of nonzero rows or columns in E , whichever is smaller
- $\omega < 2.373$, exponent of matrix multiplication

Multiplication Complexity

Total cost in field operations is $\tilde{O}\left(n + t + k \min\left(\frac{t}{r} + 1, n / \min\left(r, \frac{k}{r}\right)^{3-\omega}\right)\right)$

- n = dimension
- t = number of nonzeros in A, B, C
- k = number of errors (nonzeros in E)
- r = number of nonzero rows or columns in E , whichever is smaller
- $\omega < 2.373$, exponent of matrix multiplication

Very sparse $t \ll nk^{\omega/2}$	“Compact” errors $r \approx \sqrt{k}$	Worst case	
		$k \leq n$	$k \geq n$
$tk^{0.5}$	$t + k^{0.69}n$ $t + k^{(\omega-1)/2}n$	$t + kn$	$t + k^{0.38}n^{1.63}$ $t + k^{\omega-2}n^{4-\omega}$

Second Problem: Matrix Inverse

Matrix Inverse with Errors

Input: $A, B \in \mathbb{F}^{n \times n}$

Output: $E \in \mathbb{F}^{n \times n}$ s.t. $A^{-1} = B + E$

Parameters:

- Dimension n
 - Input size $t = \#A + \#B$
 - Error count $k = \#E$
-
- Naïve recomputation: $O(n^\omega)$
 - Ours: $\tilde{O}(t + kn + k^\omega)$ worst-case; $\tilde{O}(t + k^{0.69}n)$ best case

Problem Formulation

Rewriting the problem definition:

$$EA = I - BA$$
$$AE = I - AB$$

Crucial steps in multiplication algorithm:

- Computing nonzero rows of E
Evaluate $EA v = v - A(Bv)$ for random vector v
- Evaluating row polynomials at $2s$ powers of θ
(Coming up next. . .)

Low-Rank Linear Algebra

Algorithm (Chung, Kwok, Lau '13)

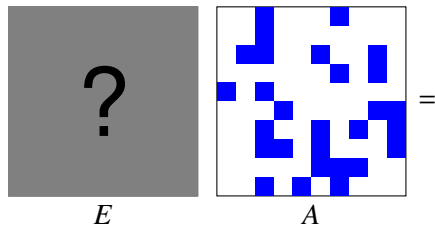
Input: Matrix M with rank r and t nonzeros

Output: Indices of r linearly independent rows

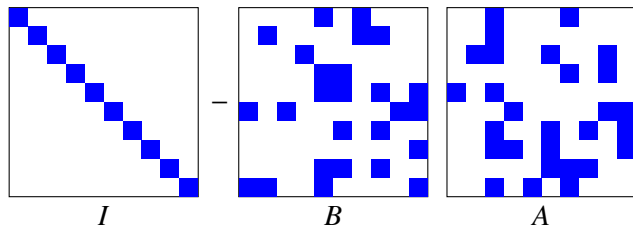
Cost: $\tilde{O}(t + r^\omega)$

Inverse Algorithm Overview

Write $EA = I - BA$



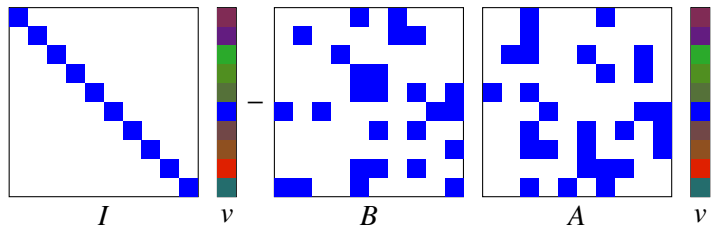
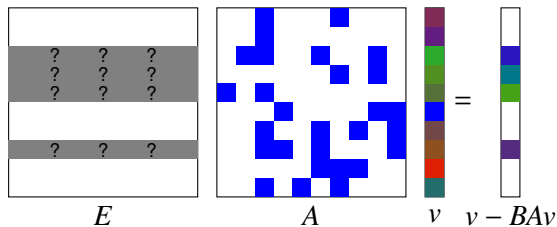
E A =



I B A

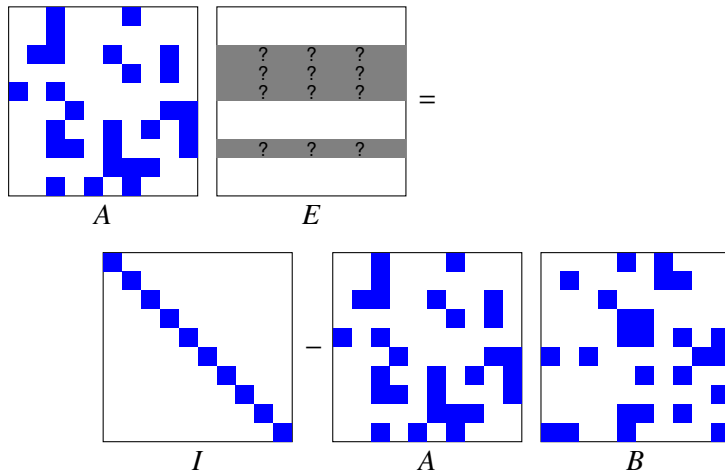
Inverse Algorithm Overview

Find nonzero rows of E



Inverse Algorithm Overview

Write $AE = I - AB$



Inverse Algorithm Overview

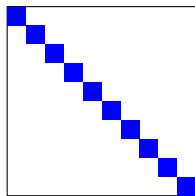
Remove zero rows of E and corresp. columns from A



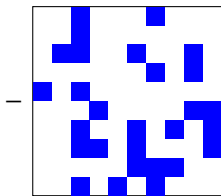
X



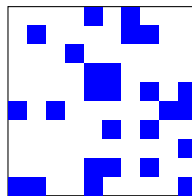
E



I



A



B

Inverse Algorithm Overview

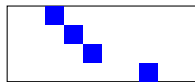
Find linear independent rows of X , remove others from both sides



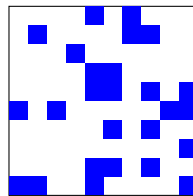
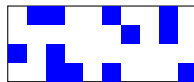
=

X

E



-



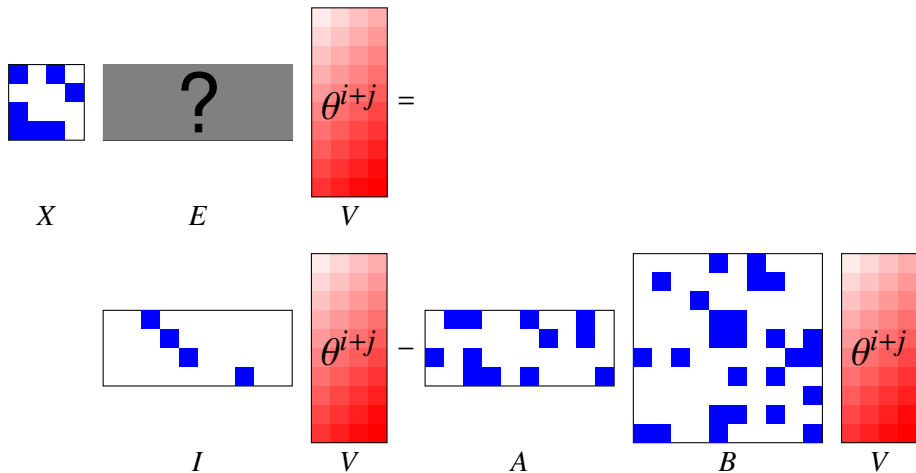
I

A

B

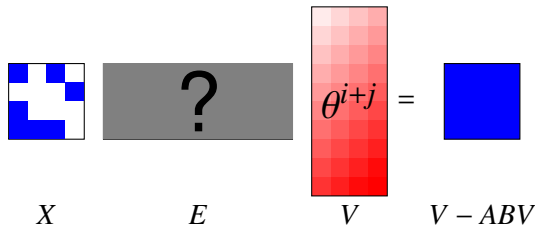
Inverse Algorithm Overview

Evaluate row polynomials at first $2s$ powers of θ



Inverse Algorithm Overview

Evaluate row polynomials at first $2s$ powers of θ



Inverse Algorithm Overview

Compute X^{-1} and apply to both sides

$$E \quad V \quad = \quad X^{-1} \quad V - ABV$$

Inverse Complexity

Total cost in field operations is $\tilde{O}\left(t + r^\omega + kn / \min(r, \frac{k}{r})^{3-\omega}\right)$

- n = dimension
- t = number of nonzeros in A, B, C
- k = number of errors (nonzeros in E)
- r = number of nonzero rows or columns in E , whichever is smaller
- $\omega < 2.373$, exponent of matrix multiplication

"Compact" errors $r \approx \sqrt{k}$	Worst case		
	$k \leq n^{0.73}$	$n^{0.74} \leq k \leq n$	$k \geq n$
$t + k^{0.69}n$	$t + kn$	$t + k^{2.38}$	$n^{2.38}$
$t + k^{(\omega-1)/2}n$		$t + k^\omega$	n^ω

Summary

New algorithms for:

- 1 **Multiplication with errors:** Computing $E = C - AB$ given A, B, C
- 2 **Inverse with errors:** Computing $E = A^{-1} - B$ given A, B

Both algorithms **account for sparsity** in the input and outputs, and can be (almost) **the same cost as reading the input and output** when there are few errors.

Possible application: Chinese remaindering

Error correction algorithms can be used when the output has **varying-sized entries**

Example

$$\begin{pmatrix} 11 & 18 & 20 & -3 \\ 1 & 14 & 9 & -23 \\ 7 & -28 & -18 & -14 \\ 28 & 22 & -16 & -30 \end{pmatrix} * \begin{pmatrix} 15 & 30 & -6 & 30 \\ -5 & 29 & 28 & 21 \\ 15 & -19 & -1 & -21 \\ -15 & -22 & -11 & -29 \end{pmatrix} \\ = \begin{pmatrix} 420 & 538 & 451 & 375 \\ 425 & 771 & 630 & 802 \\ 185 & 48 & -654 & 406 \\ 520 & 2442 & 794 & 2508 \end{pmatrix}$$

Possible application: Chinese remaindering

Error correction algorithms can be used when the output has **varying-sized entries**

Example

$$\begin{pmatrix} 11 & 18 & 20 & -3 \\ 1 & 14 & 9 & -23 \\ 7 & -28 & -18 & -14 \\ 28 & 22 & -16 & -30 \end{pmatrix} * \begin{pmatrix} 15 & 30 & -6 & 30 \\ -5 & 29 & 28 & 21 \\ 15 & -19 & -1 & -21 \\ -15 & -22 & -11 & -29 \end{pmatrix} \pmod{1607}$$

$$= \begin{pmatrix} 420 & 538 & 451 & 375 \\ 425 & 771 & 630 & 802 \\ 185 & 48 & -654 & 406 \\ 520 & -772 & 794 & -706 \end{pmatrix}$$

Updating with **subsequent primes** is an error correction problem!

Implementation

How well does this work in practice?

Preliminary implementation:

- Written in Sage
- Includes many tweaks in sparse interpolation
- Able to count [field operations](#) on actual instances

Experiments so far test the [crossover point](#) where naïve recomputation becomes better than error correction.

Experimental crossover points

