

Solving Sparse Systems with the Block Wiedemann Algorithm

Efficient Implementation over $\text{GF}(2)$

Sonia Belaid

Sylvain Lachartre

Séminaire SALSA

2011, 8 July

Outline

Outline

Introduction

- ◆ Wiedemann's algorithm was **introduced** by Wiedemann in 1986 [?]
- ◆ It was **extended** by Coppersmith in 1994 [?] to perform parallel computations
- ◆ It was **improved and used** by Emmanuel Thomé [?] and an international team of researchers in 2009 to break a 768 bits RSA key [?]

- ◆ Input : **large sparse matrix** B
- ◆ Computation of the **linearly recurrent sequence**

$$x^T \cdot B^k \cdot z$$

with x and z random vectors

- ◆ Computation of the **minimal polynomial** $\mu(X) = X \cdot \mu'(X)$
(because B is singular) using Berlekamp-Massey
- ◆ Exhibition of the **kernel vector** v such that

$$v = \mu'(B)$$

Outline

Outline

Block Wiedemann Algorithm :

- ◆ is used for the resolution of linear systems on finite fields
- ◆ takes a **sparse matrix** and exhibits a **kernel vector**
- ◆ is organized into **three** steps :
 1. BW1 : computes a matrices series $A = (a_k)_{1 \leq k \leq L}$
 2. BW2 : computes a linear generator F_g of A
 3. BW3 : exhibits a kernel vector w
- ◆ is **probabilistic** (depending on the matrix characteristics)

Main Algorithm

Algorithm 1 Wiedemann

Input: matrix $B \in \text{GF}(2)^{N \times N}$

Output: vector $w \in \text{GF}(2)^N$

1. $(A, z) \leftarrow \text{BW1}(B)$ $// z \in \text{GF}(2)^{N \times n}$ (random)
 2. $F_g \leftarrow \text{BW2}(A)$ $// F_g \in \text{GF}(2)^N$ (linear generator)
 3. $w \leftarrow \text{BW3}(F_g, B, A, z)$
 4. **return** w
-

- ◆ **Input** : large sparse matrix B
- ◆ **Output** : kernel vector w
- ◆ **N** : input matrix dimension
- ◆ **m, n** : columns number of random blocks x and z

BW1 : Sequence A

- ◆ The first step consists in computing the series

$$\mathbf{A}(\mathbf{X}) \leftarrow \sum_k a_k \cdot \mathbf{X}^k \quad \text{with} \quad a_k \leftarrow \mathbf{x}^T \cdot \mathbf{B}^k \cdot \mathbf{z} \in \text{GF}(2)^{m \times n}$$

- ◆ Only the first $\mathbf{L} = \frac{\mathbf{N}}{\mathbf{m}} + \frac{\mathbf{N}}{\mathbf{n}} + 1$ coefficients are needed

Algorithm 2 BW1

Input: matrix B

Output: polynomial $A \in \text{GF}(2)[[X]]^{m \times n}$, matrix $z \in \text{GF}(2)^{N \times n}$

1. $(\mathbf{x}, \mathbf{z}) \leftarrow$ random matrices $\in \text{GF}(2)^{N \times m} \times \text{GF}(2)^{N \times n}$
 2. $\mathbf{v} \leftarrow B \cdot \mathbf{z}$
 3. **for** $k \leftarrow 0$ **to** L **do**
 4. $A[k]^a \leftarrow \mathbf{x}^T \cdot \mathbf{v}$
 5. $\mathbf{v} \leftarrow B \cdot \mathbf{v}$
 6. **endfor**
 7. **return** A
-

a. $A[k]$ represents the degree k coefficient of A

Coppersmith's generalization of **Berlekamp-Massey** algorithm :

- ◆ **Initialization** : $m + n$ candidates F_j for the vectorial generator

$$A(X) \cdot F(X) = G(X) + X^t \cdot E(X)$$

with t the step number depending on A at the beginning.

- ◆ **Iteration** : error $E(X)$ reduced by multiplying the previous equality by P s.t. :

$$E[0] \cdot P = 0$$

- ◆ **Termination** : generator discovered when its error is zero

$$F_g(X) \text{ s.t. } A(X) \cdot F_g(X) = G_g(X) \quad (E_g(X) = 0)$$

BW2 : Initialization

- d** **constant** $d \leftarrow \lceil \frac{N}{m} \rceil$
- s** $F^{(t_0)}$ **degree** s.t. columns of $A[0], \dots, A[s-1]$ form a basis of $\text{GF}(2)^m$
- t** **step number** starting with $t \leftarrow s$
- end** **stop condition** $\text{end} \leftarrow 0$
- δ **quantity** s.t. $\delta(f, g) = \max(\deg f, 1 + \deg g)$
- Δ **degrees bounds**
- $\forall j, \quad \delta(F_j, G_j) \leq \Delta_j$
- F** **generator candidates**
- $(\quad I_n \mid X^{s-i_1} \cdot r_1 \quad \dots \quad X^{s-i_m} \cdot r_m \quad)$

BW2 : Iteration (1/2)

- ◆ computation of polynomial \mathbf{P} of degree 1

$$E[0] \cdot \mathbf{P} = 0$$

- ◆ the condition becomes

$$\begin{aligned} A \cdot F^{(t)} \cdot \mathbf{P} &= G^{(t)} \cdot \mathbf{P} + X^t \cdot E^{(t)} \cdot \mathbf{P} \\ A \cdot F^{(t+1)} &= G^{(t+1)} + X^{(t+1)} \cdot E^{(t+1)} \end{aligned}$$

- ◆ update of F and E

$$\begin{aligned} \mathbf{F}^{(t+1)} &= F^{(t)} \cdot \mathbf{P} \\ \mathbf{E}^{(t+1)} &= E^{(t)} \cdot \mathbf{P} \cdot \frac{1}{X} \end{aligned}$$

BW2 : Iteration (2/2)

Algorithm 3 PMatrix

Input: matrix $E0 \in \text{GF}(2)^{m \times (m+n)}$ **tab** $\Delta \in \mathbb{N}^{m+n}$

Output: polynomial $P \in \text{GF}(2)[X]^{(m+n) \times (m+n)}$

```
/* Sort */
1.  $(P[0], \Delta) \leftarrow \text{Sort}(\Delta)$ 
2.  $E0 \leftarrow E0 \cdot P[0]$ 
/* Gaussian Elimination */
3.  $(E0, P[0]) \leftarrow \text{GaussElimE0}(E0, P[0])$ 
/* Elimination of non zero columns */
4. for  $i \leftarrow 1$  to  $m + n$  do
5.   if ( $E0_i^a$  is not null) then
6.      $P \leftarrow \text{MultByX}(P, i)$ 
7.      $\Delta_i \leftarrow \Delta_i + 1$ 
8.   endif
9. endfor
10. return  $P$ 
```

- mean value $\bar{\Delta}$ of Δ_j coefficients increases by $\frac{m}{m+n} = \frac{1}{2}$

$$t - \bar{\Delta} = (t - s) \cdot \frac{n}{m+n}$$

- for $t = s + \lceil \frac{m}{m+n} d \rceil$, $t - \bar{\Delta} \geq d \Rightarrow \exists j, t - \Delta_j \geq d$
- according to theorem 8.6 in [?]

$$\exists j, t - \Delta_j \geq d \Rightarrow E_j(X) = 0$$

Stop Condition

F_j generator if $t - \Delta_j \geq d$ (before $s + \lceil \frac{m}{m+n} d \rceil$ steps)

Algorithm 4 BW2

Input: polynomial A

Output: polynomial $F \in \text{GF}(2)[X]^n$

/ Initialization */*

1. $(d, s, t, \text{end}, \Delta, F) \leftarrow \text{Init}(A)$

2. $E \leftarrow \text{Error}(A, F, t)$

/ Iteration */*

3. **while** ($\text{end} = \text{FALSE}$) **do**

4. $(P, \text{end}, g) \leftarrow \text{PMatrix}(E[0], \Delta)$

5. $F \leftarrow F \cdot P$

6. $E \leftarrow E \cdot P \cdot \frac{1}{X}$

7. $t \leftarrow t + 1$

8. **endwhile**

/ Termination */*

9. **return** F_g

BW3 : Kernel Vector

Kernel vector exhibition :

- ◆ coefficient j of $A(X) \cdot F_g(X)$

$$(A \cdot F_g)[j] = x^T \cdot B^{j - \deg F_g + 1} \cdot v \quad \text{with} \quad v = \sum_{i=0}^{\deg F_g} B^{\deg F_g - i} \cdot z \cdot F_g[i]$$

- ◆ by construction

$$(A \cdot F_g)[j] = 0 \quad \text{for } j \geq \delta(F_g)$$

- ◆ $B^{\delta(F_g) - \deg F_g + 1} \cdot v$ orthogonal to all vectors $(B^T)^i \cdot x_k$
- ◆ if these vectors form a basis of K^N

$$B \cdot \underbrace{(B^{\delta(F_g) - \deg F_g} \cdot v)}_{\text{Kernel Vector}} = 0$$

BW3 Algorithm

Algorithm 5 BW3

Input: polynomial F_g , matrix B , polynomial A , matrix z

Output: vector w

1. $w \leftarrow 0_N$
 2. **for** $i \leftarrow 0$ **to** $\deg F_g$ **do** $w \leftarrow B \cdot w + z \cdot F_g[i]$ **endfor**
 3. **if** ($w \neq 0_N$) **then**
 4. **for** $k \leftarrow 0$ **to** $\delta(F_g)^a$ **do**
 5. $u \leftarrow B \cdot w$
 6. **if** $u = 0$ **then return** w **endif**
 7. $w \leftarrow u$
 8. **endfor**
 9. **endif**
 10. **return** FAILED
-

a. $\delta(F_g) = \max(\deg F_g, \deg(A \cdot F_g) + 1)$

Outline

2009

Researchers broke a **768 bits** RSA key using **NFS** [?]

- ◆ NFS (Number Field Sieve) : factorization of large numbers
 - using Wiedemann's algorithm
 - input binary matrix :
 - * 200 millions of rows
 - * 150 non zeros elements by rows
 - 98 days of computations on a cluster of 576 cores

Outline

Outline

- ◆ Linear algebra library in C language focused on dense matrices over $\text{GF}(2)$
- ◆ Created by Gregory Bard
- ◆ Now maintained by Martin Albrecht
- ◆ We use it to compute operations on dense matrices
 - BW1 : to compute the products of $x \cdot (B^k \cdot z)$
 - BW2 : to perform all the operations involving blocks
 - BW3 : to compute the products $(B \cdot z) \cdot F$

Sparse Matrix

- ◆ Structure for the input sparse matrix :
 - dimensions m, n
 - number of non zeros nb
 - number of non zeros by rows sz
 - positions of non zeros pos
 - rows structures with the same characteristics l

Matrix Example

$$\begin{pmatrix} 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 1 \\ 0 & 0 & 1 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 \end{pmatrix}$$

```
m,n = 5,    nb=6
sz = [1,0,2,2,1]
pos = [0,1,4,2,3,4]
l[0] : nb=1, pos=[0]
l[1] : nb=0, pos=[] ...
```


Sparse Dense Operations

Algorithm 6 Sparse Dense Product

Input: matrix B , matrix v

Output: matrix $res \leftarrow B \cdot v$

1. $res \leftarrow 0$
 2. $p \leftarrow B_{pos}$
 3. **for** $i \leftarrow 0$ **to** N **do**
 4. **for** $j \leftarrow 0$ **to** B_{inb} **do**
 5. $res_i \leftarrow res_i \oplus v_p$
 6. $p \leftarrow p + 1$
 7. **endfor**
 8. **endfor**
 9. **return** res
-

◆ Complexity : **linear** with the number of non zeros.

Outline

Counting Sort

array to sort

Δ	10	9	9	8	12	8	9	8	12	8
----------	----	---	---	---	----	---	---	---	----	---

counting array

values	8	9	10	11	12
occurrences	4	3	1	0	2

sorted array

CountingSort(Δ)	8	8	8	8	9	9	9	10	12	12
--------------------------	---	---	---	---	---	---	---	----	----	----

Complexity

Linear in $O(m + n)$ with $m + n$ the size of Δ

Outline

Practical Optimization : Operations Save

Loop

1. **for** $i \leftarrow 0$ **to** $a \times b$ **do**



Loop

1. *constant* $tmp \leftarrow a \times b$
2. **for** $i \leftarrow 0$ **to** tmp **do**

Product

BW1 : $\forall k \in [1..L],$
 $v \leftarrow B^k \cdot z$
BW3 : $\forall k \in [1..\deg F_g],$
 $v \leftarrow B^k \cdot z$



Product

BW1 : $\forall k \in [1..L'],$
 $v \leftarrow Bz_{save}[k] \leftarrow B^k \cdot z$
BW3 : $\forall k \in [1..\deg F_g],$
 $v \leftarrow Bz_{save}[k]$

Practical Optimization : Reduction of the Number of Tests

Test

```
1. for  $i \leftarrow 1$  to  $n$  do  
2.   if  $(i < \frac{n}{2})$  then Action1  
3.     else Action2 endif  
4. endfor
```



Test

```
1. for  $i \leftarrow 1$  to  $\frac{n}{2}$  do  
2.   Action1  
3. endfor  
4. for  $i \leftarrow \frac{n}{2}$  to  $n$  do  
5.   Action2  
6. endfor
```

Practical Optimization : Adaptation of M4RI¹ Functions

◆ removing **initial tests**

- $A \cdot B$: number of A 's columns = number of B 's rows
- equality to zero

◆ adapting to my matrices **constant dimensions**

- BW1 : products of matrices $(64 \times N) \times (N \times 64)$
- BW2 : products of matrices $(64 \times 128) \times (128 \times 128)$
- global : most matrices whose dimensions are multiples of the size a machine word

◆ **improving** some functions

- `mzd_is_zero` : stopping at the first non zero word
- `mzd_transpose` : clearing before transposing

Practical Optimization : Parallelization

- ◆ **BW1** : parallel
- ◆ **BW2** : sequential \Rightarrow parallelization
- ◆ **BW3** : sequential \Rightarrow operations save

Practical Optimization : Parallelization BW1

- ◆ sparse dense product : $B \cdot (B^k \cdot z)$, $k \in [1..L]$
 - each thread dedicated to a number of rows
 - depending on number of non zeros by rows




Figure: Representation of a 2688×2688 matrix


- ◆ computations of $a_k : a_k \leftarrow x \cdot (B^k \cdot z)$, $k \in [1..L]$
 - each thread dedicated to a number of a_k
 - same number of coefficients for each thread

Practical Optimization : Parallelization BW2


- ◆ polynomials products $F \cdot P$ and $E \cdot P$
 - each thread dedicated to a number of coefficients of F or E
 - same number of coefficients for each thread
- ◆ product of P by X
 - each thread dedicated to a number of columns of P
 - same number of columns for each thread

Practical Optimization : Parallelization : CPU Use


1 thread



3 threads




6 threads



Practical Optimization : Parallelization : Execution Time

Execution times according to the number threads




Dim	$\delta(\%)$	1th	2th	3th	4th	5th	6th	7th	8th
2688	1.3	0.3	0.2	0.2	0.2	0.2	0.2	0.2	0.2
15360	1.5	11.6	7.7	6.2	5.2	4.6	4.1	3.9	3.6
15360	5.2	22.3	12.5	9.6	7.7	7.1	6.9	6.4	5.9
15360	8.2	32.0	18.5	14.2	11.1	10.4	9.5	9.2	8.7

Theoretical Optimizations : Termination Tests

Stop Condition

$$\exists j \in [1..m+n], \quad t \geq \Delta_j + d$$



New Stop Condition

$$\exists j \in [1..\frac{m+n}{2} + 1] \text{ s.t. } \Delta_j^{(t)} = \Delta_j^{(t-1)}, \quad t \geq \Delta_j + d$$

Theoretical Optimizations : Error Degree Update (1/2)

- ◆ for large matrices
 - high error degree in BW2
 - product $E \cdot P$ **expensive**
- ◆ but **only $E[0]$** used in each step
- ◆ knowing the number of reminding steps, we can :
 - determine the **useful degree** of E
 - decrease the number of coefficients to update

Theoretical Optimizations : Error Degree Update (2/2)

- ◆ Stop condition : $t_f = \min(\Delta)^{(t_f)} + d$
- ◆ Worst case
 - $\min(\Delta)$ increases by 0.5 at each step
 - from step t , we still have $\delta_e^{(t)}$ iterations :

$$\begin{aligned}t + \delta_e^{(t)} &= \min(\Delta)^{(t)} + \frac{\delta_e^{(t)}}{2} + d \\ \Rightarrow \delta_e^{(t)} &= 2 \cdot (\min(\Delta)^{(t)} + d - t)\end{aligned}$$

$$\bullet \Rightarrow \begin{cases} \delta_e^{(t+1)} = \delta_e^{(t)} & \text{if } \min(\Delta)^{(t+1)} = \min(\Delta)^{(t)} + 1 \\ \delta_e^{(t+1)} = \delta_e^{(t)} - 2 & \text{otherwise} \end{cases}$$

Error Degree Bound

$$\forall t, \quad \deg E^{(t)} \leq \delta_e^{(t)}$$

Theoretical Optimizations : Candidates Degree Update (1/2)

- ◆ for large matrices
 - candidates degree increases by 1 at each step (becoming high)
 - product $F \cdot P$ becomes **expensive**
- ◆ but **only** F_g (linear generator) will be kept
- ◆ knowing the number of reminding steps, we can :
 - determine the **useful degree** of F (that is the one of F_g)
 - limit the number of coefficients to update

Theoretical Optimizations : Candidates Degree Update (2/2)

- ◆ F degree is limited to δ_f such that

$$\delta_f^{(t)} = \min(\Delta)^{(t)} + r^{(t)}$$

with $r^{(t)} = 2 \cdot (\min(\Delta)^{(t)} + d - t)$ (maximum reminding steps)

$$\Rightarrow \delta_f^{(t)} = \min(\Delta)^{(t)} + 2 \cdot (\min(\Delta)^{(t)} + d - t)$$

- ◆ $\Rightarrow \begin{cases} \delta_f^{(t+1)} = \delta_f^{(t)} + 1 & \text{if } \min(\Delta)^{(t+1)} = \min(\Delta)^{(t)} + 1 \\ \delta_f^{(t+1)} = \delta_f^{(t)} - 2 & \text{otherwise} \end{cases}$

Candidates Degree Bound

$$\forall t, \quad \deg F^{(t)} \leq \delta_f^{(t)}$$

Outline

Outline

Execution Times

Dim	$\delta(\%)$	BW1	BW2	BW3	Global
738	7.00	10%	90%	0%	0''100
3422	7.57	20%	73%	7%	0''450
10000 ²	5.97	46%	43%	11%	1''960
27000 ¹	2.20	60%	28%	12%	14''060
73674	0.61	65%	30%	5%	1'12''760
93913	0.20	51%	41%	8.3%	1'38''710

Figure: Performances on different matrix sizes

Computer Details : PC Xeon (64 bits)

- ◆ 8 Go RAM
- ◆ 2.40 GHz
- ◆ 8 processors Intel(R) Xeon(R) CPU

Outline

Comparisons with Magma and Sage

Matrices		Our Implementation		Speed up (1th)	
Dim	δ	8 th	1 th	Magma	Sage
2688	1.27%	0"200	0"360	$\times 5.8$	$\times 150.0$
15360	1.54%	3"580	11"300	$\times 85.8$	— ³
15360	5.20%	5"870	21"670	$\times 62.4$	— ²
15360	8.21%	8"690	31"480	$\times 62.3$	— ²

Figure: Temporal comparisons with Magma⁴ and Sage⁵

3. not enough memory (> 8 Go)

4. Magma version 2.17-1, released 2010-12-02

5. Sage version 4.6.2, released 2011-02-25

Outline

◆ Summary

- Efficient implementation of the Block Wiedemann algorithm over $\text{GF}(2)$ in C language
- Practical and theoretical optimizations
- Encouraging results compared to existing methods

◆ Further Work

- *Algebraic Cryptanalysis of HFE Cryptosystems Using Gröbner Bases* [?] using Block Wiedemann algorithm
- Use of Gröbner bases algorithm [?]
- Comparisons with LinBox
- Work on applications
- Open source?

Bibliography