Shortest Vector from Lattice Sieving: a Few Dimensions for Free

Léo Ducas 1

Cryptology Group, CWI, Amsterdam, The Netherlands



EUROCRYPT 2018 Tel Aviv, April 30th

¹Supported by a Veni Innovational Research Grant from NWO (639,021,645).

The Shortest Vector Problem

- I: The basis **B** of an *n*-dimensional lattice \mathcal{L}
- **O:** A shortest non-zero vector $\mathbf{v} \in \mathcal{L}$

Algorithm	Running time	Memory
Enumeration	$n^{n/2e} \cdot 2^{O(n)}$	poly(<i>n</i>)
Sieving ²	$[2^{.292n+o(n)}, 2^{.415n+o(n)}]$	$[2^{.2075n+o(n)}, 2^{.292n+o(n)}]$

The paradox

In theory, Sieving is faster. In pratice it is quite slower.

²Given complexity are heuristic, heavily supported by experiments. 💿 🗸 😨 🔊 🛇

The Shortest Vector Problem

- I: The basis **B** of an *n*-dimensional lattice \mathcal{L}
- **O:** A shortest non-zero vector $\mathbf{v} \in \mathcal{L}$

Algorithm	Running time	Memory
Enumeration	$n^{n/2e} \cdot 2^{O(n)}$	poly(<i>n</i>)
$Sieving^2$	$[2^{.292n+o(n)}, 2^{.415n+o(n)}]$	$[2^{.2075n+o(n)}, 2^{.292n+o(n)}]$

The paradox

In theory, Sieving is faster. In pratice it is quite slower.

²Given complexity are heuristic, heavily supported by experiments. $\Xi \mapsto A \equiv 0 \propto 0$ Léo Ducas (CWI, Amsterdam) SVP from Sieving: a Few Dims for Free 30 April 2018 2 / 23

The Shortest Vector Problem

- I: The basis **B** of an *n*-dimensional lattice \mathcal{L}
- **O:** A shortest non-zero vector $\mathbf{v} \in \mathcal{L}$

Algorithm	Running time	Memory
Enumeration	$n^{n/2e} \cdot 2^{O(n)}$	poly(<i>n</i>)
$Sieving^2$	$[2^{.292n+o(n)}, 2^{.415n+o(n)}]$	$[2^{.2075n+o(n)}, 2^{.292n+o(n)}]$

The paradox

In theory, Sieving is faster. In pratice it is quite slower.

²Given complexity are heuristic, heavily supported by experiments. $\blacksquare \land \blacksquare \land \blacksquare$

Many trade-offs



- Our main contribution can also be applied to other sieving algorithms.
- Implementation limited to the version of [Micciancio Voulgaris 2010].

Many trade-offs



- Our main contribution can also be applied to other sieving algorithms.
- Implementation limited to the version of [Micciancio Voulgaris 2010].

Results

Heuristic claim, asymptotic

One can solve SVP in dimension n with a call to SIEVE in dimension n - d

where $d = \Theta(n/\log n)$.

Heuristic claim, concrete

One can solve SVP in dimension *n* making a call to SIEVE in dimension *i* for each $i = 1 \dots n - d$ for

$$d \approx \frac{n \cdot \ln(4/3)}{\ln(n/2\pi e)}$$
 $(d \approx 15 \text{ for } n = 80)$

Experimental claim: A bogey

A SIEVE implem. almost on par with enumeration (within a factor 4 in dims 70–80), with still much room for improvements.

Léo Ducas (CWI, Amsterdam) SVP from Sieving: a Few Dims for Free

30 April 2018 4 / 23

Results

Heuristic claim, asymptotic

One can solve SVP in dimension n with a call to SIEVE in dimension n - d

where $d = \Theta(n/\log n)$.

Heuristic claim, concrete

One can solve SVP in dimension *n* making a call to SIEVE in dimension *i* for each $i = 1 \dots n - d$ for

$$d \approx rac{n \cdot \ln(4/3)}{\ln(n/2\pi e)}$$
 ($d \approx 15$ for $n = 80$)

Experimental claim: A bogey

A SIEVE implem. almost on par with enumeration (within a factor 4 in dims 70–80), with still much room for improvements.

Results

Heuristic claim, asymptotic

One can solve SVP in dimension n with a call to SIEVE in dimension n - d

where $d = \Theta(n/\log n)$.

Heuristic claim, concrete

One can solve SVP in dimension *n* making a call to SIEVE in dimension *i* for each $i = 1 \dots n - d$ for

$$d \approx rac{n \cdot \ln(4/3)}{\ln(n/2\pi e)}$$
 ($d \approx 15$ for $n = 80$)

Experimental claim: A bogey

A SIEVE implem. almost on par with enumeration (within a factor 4 in dims 70–80), with still much room for improvements.

30 April 2018 4 / 23

1 Dimensions for free

2 Implementation and performances

1 Dimensions for free

2 Implementation and performances

Algorithm 1 SIEVE(\mathcal{L})

 $L \leftarrow$ a set of N random vectors from \mathcal{L} where $N \approx (4/3)^{n/2}$. while $\exists (\mathbf{v}, \mathbf{w}) \in L^2$ such that $\|\mathbf{v} - \mathbf{w}\| < \|\mathbf{v}\|$ do $\mathbf{v} \leftarrow \mathbf{v} - \mathbf{w}$ end while return L

The above runs in heuristic time $(4/3)^{n+o(n)}$.

Many concrete and asymptotic improvements: [Nguyen Vidick 2008, Micciancio Voulgaris 2010, Laarhoven 2015, Becker Gamma Joux 2015, Becker D. Gamma Laarhoven 2015, ...].

Algorithm 2 SIEVE(\mathcal{L})

 $L \leftarrow$ a set of N random vectors from \mathcal{L} where $N \approx (4/3)^{n/2}$. while $\exists (\mathbf{v}, \mathbf{w}) \in L^2$ such that $\|\mathbf{v} - \mathbf{w}\| < \|\mathbf{v}\|$ do $\mathbf{v} \leftarrow \mathbf{v} - \mathbf{w}$ end while return L

The above runs in heuristic time $(4/3)^{n+o(n)}$.

Many concrete and asymptotic improvements:

[Nguyen Vidick 2008, Micciancio Voulgaris 2010, Laarhoven 2015, Becker Gamma Joux 2015, Becker D. Gamma Laarhoven 2015, ...].

< /₽ > < E > <

Note that SIEVE returns $N \approx (4/3)^n$ short vectors, not just a shortest vector.

Definition (Gaussian Heuristic: Expected length of the shortest vector)

$$\operatorname{gh}(\mathcal{L}) = \sqrt{n/2\pi e} \cdot \operatorname{vol}(\mathcal{L})^{1/n}$$

Observation (heuristic & experimental)

The output of SIEVE contains almost all vectors of length $\leq \sqrt{4/3} \cdot \text{gh}(\mathcal{L})$:

$$\mathcal{L} := \operatorname{Sieve}(\mathcal{L}) = \left\{ \mathbf{x} \in \mathcal{L} \text{ s.t. } \|\mathbf{x}\| \leq \sqrt{4/3} \cdot \operatorname{gh}(\mathcal{L})
ight\}.$$

Main idea: Sieve in a projected sub-lattice, and lift all candidate solutions.

SUBSIEVE(\mathcal{L}, d)

- Set $\mathcal{L}' = \mathcal{L}(\mathbf{b}_1, \dots, \mathbf{b}_d)$
- Set $\mathcal{L}'' = \pi_{\mathcal{L}'}^{\perp}(\mathcal{L})$
- Compute $L = \text{SIEVE}(\mathcal{L}'')$
- Hope that $\pi_{\mathcal{L}'}^{\perp}(\mathbf{s}) \in L$
- Lift all $\mathbf{v} \in L$ from \mathcal{L}'' to \mathcal{L} and take the shortest (Babai alg.)

Pessimistic prediction for (1)

$$\operatorname{gh}(\mathcal{L}) \leq \sqrt{4/3} \cdot \operatorname{gh}(\mathcal{L}_d).$$

Optimistic prediction for (1)

"left part of \mathcal{L} ", dim=d

"right part of \mathcal{L} ", dim=n-d

$$\sqrt{\frac{n-d}{n}} \cdot \operatorname{gh}(\mathcal{L}) \leq \sqrt{4/3} \cdot \operatorname{gh}(\mathcal{L}_d).$$

Similar to linear pruning for enum.

30 April 2018

9 / 23

Léo Ducas (CWI, Amsterdam) SVP from Sieving: a Few Dims for Free

Main idea: Sieve in a projected sub-lattice, and lift all candidate solutions.

SUBSIEVE (\mathcal{L}, d)

- ▶ Set $\mathcal{L}' = \mathcal{L}(\mathbf{b}_1, \dots, \mathbf{b}_d)$ "left part of \mathcal{L} ", dim=d
- ▶ Set $\mathcal{L}'' = \pi_{\mathcal{L}'}^{\perp}(\mathcal{L})$ "right part of \mathcal{L} ", dim=n d
- Compute $L = \text{SIEVE}(\mathcal{L}'')$
- ▶ Hope that $\pi_{\mathcal{L}'}^{\perp}(\mathbf{s}) \in L$
- Lift all $\mathbf{v} \in L$ from \mathcal{L}'' to \mathcal{L} and take the shortest (Babai alg.)

Pessimistic prediction for (1)

$$\operatorname{gh}(\mathcal{L}) \leq \sqrt{4/3} \cdot \operatorname{gh}(\mathcal{L}_d).$$

Optimistic prediction for (1)

$$\sqrt{rac{n-d}{n}} \cdot \mathrm{gh}(\mathcal{L}) \leq \sqrt{4/3} \cdot \mathrm{gh}(\mathcal{L}_d).$$

(1)

Similar to linear pruning for enum.

Main idea: Sieve in a projected sub-lattice, and lift all candidate solutions.

SUBSIEVE (\mathcal{L}, d)

- ► Set $\mathcal{L}' = \mathcal{L}(\mathbf{b}_1, \dots, \mathbf{b}_d)$ "left part of \mathcal{L} ", dim=d
- ► Set $\mathcal{L}'' = \pi_{\mathcal{L}'}^{\perp}(\mathcal{L})$ "right part of \mathcal{L} ", dim=n d
- Compute $L = \text{SIEVE}(\mathcal{L}'')$
- ▶ Hope that $\pi_{\mathcal{L}'}^{\perp}(\mathbf{s}) \in L$
 - Lift all $\mathbf{v} \in L$ from \mathcal{L}'' to \mathcal{L} and take the shortest (Babai alg.)

Pessimistic prediction for (1)

$$\operatorname{gh}(\mathcal{L}) \leq \sqrt{4/3} \cdot \operatorname{gh}(\mathcal{L}_d).$$

Optimistic prediction for (1)

$$\sqrt{rac{n-d}{n}} \cdot \mathrm{gh}(\mathcal{L}) \leq \sqrt{4/3} \cdot \mathrm{gh}(\mathcal{L}_d).$$

(1)

Similar to linear pruning for enum.

Main idea: Sieve in a projected sub-lattice, and lift all candidate solutions.

SUBSIEVE(\mathcal{L}, d)

- ▶ Set $\mathcal{L}' = \mathcal{L}(\mathbf{b}_1, \dots, \mathbf{b}_d)$ "left part of \mathcal{L} ", dim=d
- ▶ Set $\mathcal{L}'' = \pi_{\mathcal{L}'}^{\perp}(\mathcal{L})$ "right part of \mathcal{L} ", dim=n d
- Compute $L = \text{SIEVE}(\mathcal{L}'')$
- ▶ Hope that $\pi_{\mathcal{L}'}^{\perp}(\mathbf{s}) \in L$
- Lift all $\mathbf{v} \in L$ from \mathcal{L}'' to \mathcal{L} and take the shortest (Babai alg.)

Pessimistic prediction for (1)

$$\mathsf{gh}(\mathcal{L}) \leq \sqrt{4/3} \cdot \mathsf{gh}(\mathcal{L}_d).$$

Optimistic prediction for (1)

$$\sqrt{rac{n-d}{n}} \cdot \mathrm{gh}(\mathcal{L}) \leq \sqrt{4/3} \cdot \mathrm{gh}(\mathcal{L}_d).$$

(1)

Similar to linear pruning for enum.

9 / 23

- ▶ To ensure (1), we need the basis to be as reduced as possible
- We can easily afford BKZ preprocessing with block-size b = n/2
- ▶ Using simple BKZ models³ we can predict $gh(\mathcal{L})$ and $gh(\mathcal{L}')$

Heuristic claim

SUBSIEVE(\mathcal{L}, d) algorithm will successfully find the shortest vector of \mathcal{L} for some $d = \Theta(n/\ln n)$.

 \Rightarrow Improve time & memory by a sub-exponential factor $2^{\Theta(n/\log n)}$

³The Geometric Series Assumption

- ▶ To ensure (1), we need the basis to be as reduced as possible
- We can easily afford BKZ preprocessing with block-size b = n/2
- ▶ Using simple BKZ models³ we can predict $gh(\mathcal{L})$ and $gh(\mathcal{L}')$

Heuristic claim

SUBSIEVE(\mathcal{L}, d) algorithm will successfully find the shortest vector of \mathcal{L} for some $d = \Theta(n/\ln n)$.

 \Rightarrow Improve time & memory by a sub-exponential factor $2^{\Theta(n/\log n)}$

³The Geometric Series Assumption

Idea: Attempt stronger pre-processing.

Algorithm 3 SUBSIEVE⁺(\mathcal{L} , d)

$$\begin{split} & L \leftarrow \text{SIEVE}(\mathcal{L}'') \\ & L = \{ \text{LIFT}_{\mathcal{L}'' \to \mathcal{L}}(v) \text{ for } v \in L \} \\ & \text{for } j = 0 \dots n/2 - 1 \text{ do} \\ & \mathbf{v}_j = \arg\min_{\mathbf{s} \in L} \| \pi_{(\mathbf{v}_0 \dots \mathbf{v}_{j-1})^{\perp}}(\mathbf{s}) \| \\ & \text{end for} \\ & \text{return } (\mathbf{v}_0 \dots \mathbf{v}_{n/2-1}) \end{split}$$

- ▶ Insert $(\mathbf{v}_0 \dots \mathbf{v}_{n/2-1})$ as the new $\mathbf{b}_1 \dots \mathbf{b}_{n/2}$
- ▶ Repeat SUBSIEVE⁺(\mathcal{L} , d) for $d = n 1, n 2, ..., d_{min}$
- ▶ Hope that iteration *d*_{min} + 1 provided a quasi-HKZ basis.

Concrete prediction with quasi-HKZ preprocessing



1 Dimensions for free

2 Implementation and performances

- No gaussian sampling
 - Initial sphericity of L doesn't seem to matter
 - Initial vectors can be made much shorter \Rightarrow speed-up
- Prevent collisions using a hash table
- Terminate when the ball $\sqrt{4}/3 \cdot gh(\mathcal{L})$ is half-saturated
- Sort only periodically
 - Can use faster data-structures
- Vectors represented in bases B and GRAMSCHMIDT(B)
 - Required to work in projected-sublattices
- Kernel in c++, control in python
 - Calls to fpylll to maintain B and GRAMSCHMIDT(B)

- No gaussian sampling
 - Initial sphericity of L doesn't seem to matter
 - Initial vectors can be made much shorter \Rightarrow speed-up
- Prevent collisions using a hash table
- Terminate when the ball $\sqrt{4}/3 \cdot gh(\mathcal{L})$ is half-saturated
- Sort only periodically
 - Can use faster data-structures
- Vectors represented in bases B and GRAMSCHMIDT(B)
 - Required to work in projected-sublattices
- Kernel in c++, control in python
 - Calls to fpylll to maintain B and GRAMSCHMIDT(B)

- No gaussian sampling
 - Initial sphericity of L doesn't seem to matter
 - Initial vectors can be made much shorter \Rightarrow speed-up
- Prevent collisions using a hash table
- Terminate when the ball $\sqrt{4}/3 \cdot \text{gh}(\mathcal{L})$ is half-saturated
- Sort only periodically
 - Can use faster data-structures
- Vectors represented in bases B and GRAMSCHMIDT(B)
 - Required to work in projected-sublattices
- Kernel in c++, control in python
 - Calls to fpylll to maintain B and GRAMSCHMIDT(B)

- No gaussian sampling
 - Initial sphericity of L doesn't seem to matter
 - Initial vectors can be made much shorter \Rightarrow speed-up
- Prevent collisions using a hash table
- Terminate when the ball $\sqrt{4}/3 \cdot \text{gh}(\mathcal{L})$ is half-saturated
- Sort only periodically
 - Can use faster data-structures
- ▶ Vectors represented in bases **B** and GRAMSCHMIDT(**B**)
 - Required to work in projected-sublattices
- Kernel in c++, control in python
 - Calls to fpylll to maintain B and GRAMSCHMIDT(B)

- No gaussian sampling
 - Initial sphericity of L doesn't seem to matter
 - Initial vectors can be made much shorter \Rightarrow speed-up
- Prevent collisions using a hash table
- Terminate when the ball $\sqrt{4}/3 \cdot \text{gh}(\mathcal{L})$ is half-saturated
- Sort only periodically
 - Can use faster data-structures
- Vectors represented in bases B and GRAMSCHMIDT(B)
 - Required to work in projected-sublattices
- Kernel in c++, control in python
 - Calls to fpylll to maintain B and GRAMSCHMIDT(B)

Already used in Sieving [Fitzpatrick et al. 2015]. More generally know as SIMHASH [Charikar 2002].

Idea: Pre-filter pairs $(\mathbf{v}, \mathbf{w}) \in L$ with a fast compressed test.

- ▶ Choose a spherical code $C = {\bf c}_1 \dots {\bf c}_k \} \subset S^n$ and a threshold $t \le k/2$
- Precompute compressions $\tilde{\mathbf{v}} = \text{SIGN}(\mathbf{v}) \in \{0,1\}^k$
- Only test $\|\mathbf{v} \pm \mathbf{w}\| \le \|\mathbf{v}\|$ if

 $|\text{HAMMINGWEIGHT}(\mathbf{v} \oplus \mathbf{w}) - k/2| \ge t.$

- Asymptotic speed-up $\Theta(n/\log n)$?
- ▶ In practice, k = 128 (2 words), t = 18: about 10 cycles per pairs.

通 ト イヨト イヨト

Concurrently and independetly invented in [Mariano Laarhoven 2018].

Idea: Increase the dimension progressively.

- Recursively, Sieve in the lattice $\mathcal{L}(\mathbf{b}_1, \dots, \mathbf{b}_{n-1})$
- Start the sieve in dimension *n* with many short-ish vectors
- Fresh vector get reduced much faster thanks to this initial pool.

Refer to [Mariano Laarhoven 2018] for a full analysis of this trick.

Dimensions for Free (V2 \rightarrow V3)

- Apply the quasi-HKZ preprocessing strategy
- Do not force the choice of d_{\min}
- Simply increase *d* until the shortest vector is found.



Figure: Predictions experiments for d_{\min}

Léo Ducas (CWI, Amsterdam)

SVP from Sieving: a Few Dims for Free

Performances

T(sec.)



Léo Ducas (CWI, Amsterdam)

SVP from Sieving: a Few Dims for Free

30 April 2018 18 / 23

	Algorithms							
	V0	V1	V2	V3	[MV10]	[FBB ⁺ 14]	[ML17]	[HK17]
F .								
Features								
XOR-POPCNT trick		х	х	х		х		
pogressive sieving			х	х				
SubSieve				х				
LSH (more mem.)							x	
tuple (less mem.)								х
Dimension				Runn	ning times			
<i>n</i> = 60	227s	49s	8s	0.9s	464s	79s	13s	1080s
<i>n</i> = 70	-	-	276s	10s	23933s	4500s	250s	33000s
<i>n</i> = 80	-	-	-	234s	-	-	4320s	94700s
CPU freq. (GHz)	3.6	3.6	3.6	3.6	4.0	4.0	2.3	2.3

Summary

Sieving vs. Sieving

- Exploit all outputs of Sieve \Rightarrow Dimensions for Free
- Our implementation is 10x faster than all previous Sieving
- It does not use LSH techniques: further speed-up expected

Sieving vs. Enumeration

- Only a factor 4x slower than Enum for dimensions 70–80
- ► Guesstimates a cross-over at dim ≈ 90 with further improvements (LSH/LSF, fine-tuning, vectorization, ...)

Summary

Sieving vs. Sieving

- Exploit all outputs of Sieve \Rightarrow Dimensions for Free
- Our implementation is 10x faster than all previous Sieving
- It does not use LSH techniques: further speed-up expected

Sieving vs. Enumeration

- Only a factor 4x slower than Enum for dimensions 70–80
- ► Guesstimates a cross-over at dim ≈ 90 with further improvements (LSH/LSF, fine-tuning, vectorization, ...)

Summary

Sieving vs. Sieving

- Exploit all outputs of Sieve \Rightarrow Dimensions for Free
- Our implementation is 10x faster than all previous Sieving
- It does not use LSH techniques: further speed-up expected

Sieving vs. Enumeration

- Only a factor 4x slower than Enum for dimensions 70–80
- ► Guesstimates a cross-over at dim ≈ 90 with further improvements (LSH/LSF, fine-tuning, vectorization, ...)



Léo Ducas (CWI, Amsterdam)

To be continued.

Question ?

æ 30 April 2018 21 / 23

∃ →

Image: A match a ma

A claim of P. Kirchner⁴

A simple Sieving algorithm (e.g. NV) can be implemented with

$$A = 2^{.2075n+o(n)}$$
 and $T = 2^{.2075n+o(n)}$

The circuit is easy to re-invent, using shift registers:

- No long wires: no speed-of-light delays !
- Essentially 1-dimensional



Conjecture / Open Question

There exist a sieving circuit with:

$$A = 2^{.2075n+o(n)}$$
 and $T \le 2^{.142n+o(n)}$

Hint

- [Becker Gama Joux 2015] with only on level of filtration
- ▶ 3 or 4 layers of 2-dimensions should suffice.
- ► Keep shift-registers not fully saturated, for easier on-the-fly insertion.

Conjecture / Open Question

There exist a sieving circuit with:

$$A = 2^{.2075n+o(n)}$$
 and $T \le 2^{.142n+o(n)}$

Hint

- [Becker Gama Joux 2015] with only on level of filtration
- ► 3 or 4 layers of 2-dimensions should suffice.
- ► Keep shift-registers not fully saturated, for easier on-the-fly insertion.