Lower Bounds on Lattice Enumeration with Extreme Pruning

Yoshinori Aono

NICT 国立研究開発法人 情報通信研究機構

Phong Nguyễn

Inria inventors for the digital world cnrs 東京大学 THE UNIVERSITY OF TOKYO

Takenobu Seito

日本銀行 BANK OF JAPAN

Junji Shikata

YNU 横浜国立大学 YOKOHAMA National University

@Crypto2018, Santa Barbara, 20, Aug.

*The views expressed in this talk do not necessarily reflect the official views of BoJ

Background

Motivation is Long-term security for lattice-based crypto.

- NIST will publish PQ standard draft around 2025 and standardized scheme(s) will be used for several decades

- Need to assess performance of core attacking algorithms for setting parameters

- Majority of candidates are lattice-based.

Two-sided estimation for attacks cost

$$\frac{\text{Limit of algorithm efficiency}}{\text{Limit of computing power}} \leq \text{Attack Cost} \leq \frac{\text{Algorithm efficiency at now}}{\text{Computing power at now}}$$

- Lots of efforts have been made to find upper bounds
- How about lower bounds?

Algorithms since 70-80's:
ENUM, BKZ, Sieve, hybrids, etc.

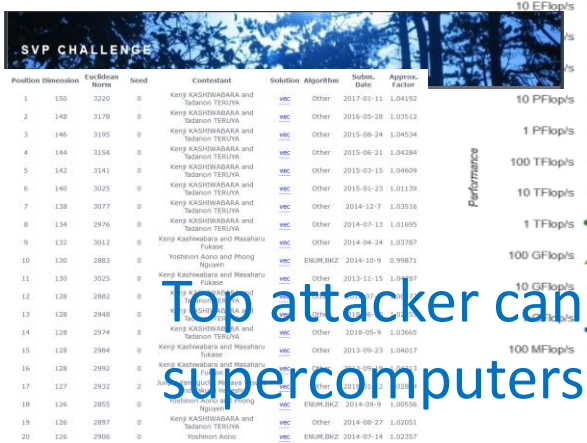

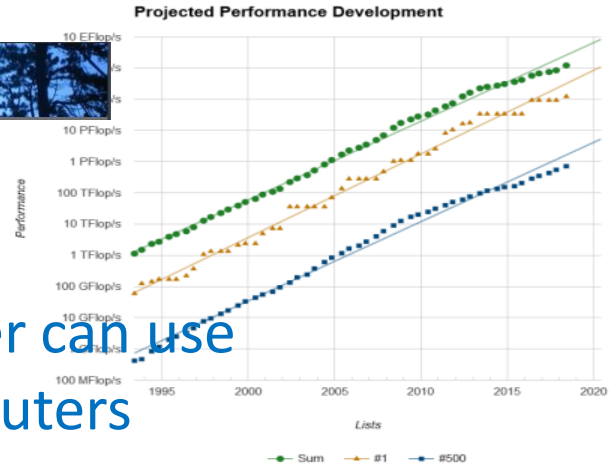

Top attacker can use supercomputers

Lower bounds

$$\frac{\text{Limit of algorithm efficiency}}{\text{Limit of computing power}} \leq \text{Attack Cost}$$

- Proving limit of efficiency of any attacking algorithm is very useful for crypto, though it is extremely hard problem (e.g. P≠NP)

- Efforts to find lower bounds for major algorithms
  - Sieve: $O(2^{0.292n})$ in classical and $O(2^{0.265n})$ in quantum (heuristic)
  - Pruned ENUM: non-trivial lower bound open

We have solved this problem

Technical result

- Lower bounds for cost of pruned lattice enumeration[GNR@EC10] used to solve SVP/BDD and related hard lattice problems

Pros

- Easy to compute ($\leq$10 ms in practice)
- Meaningful: close to upper bounds
- Can also be applied to quantum enumeration [**A**.-Nguyen-Shen@AC18 and ePrint 2018/546]

Cons and Future work

- Non trivial to adapt to other algorithms such as discrete pruning ENUM, Sieve, etc.

Applications

- Comparing our lower bound vs sieve lower bound to solve SVP-β
- State-of-the-art: current algorithms
- Conservative setting: anticipating progress in lattice reduction

Classical hardness

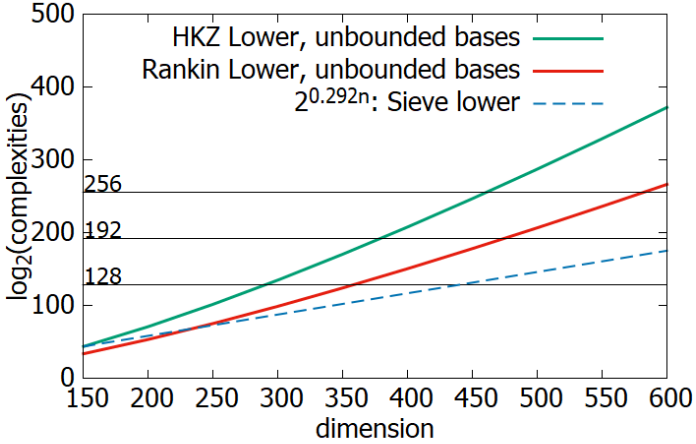

Quantum hardness

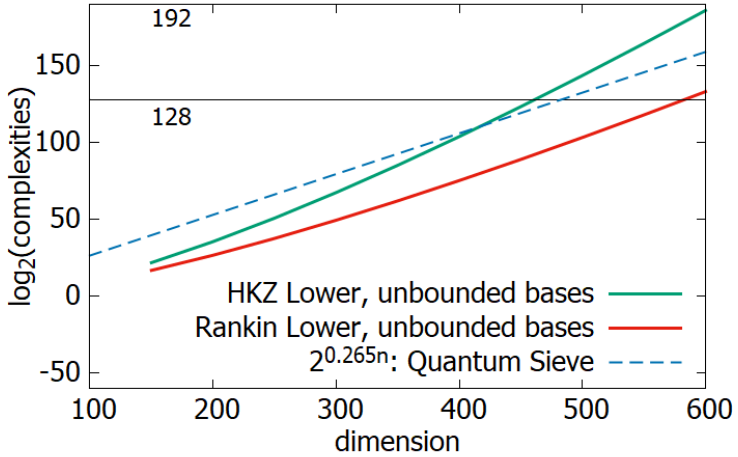

- In quantum setting, the lower bound used in several NIST submissions is not as conservative as previously believed

Agenda

- ~~Background and overview of our results~~
- Pruned ENUM and cost estimation in [GNR@EC10]
- Lower bound via isoperimetry
- Linear lower bound of randomized ENUM and application to SVP-β

ENUM: Lattice vector enumeration

- A core subroutine of BKZ-type lattice algorithms
- Given a basis $B=(\boldsymbol{b}_1,\ldots,\boldsymbol{b}_n)$ of lattice L, enumerate short lattice points

- Depth-first search of a tree depending on the input basis

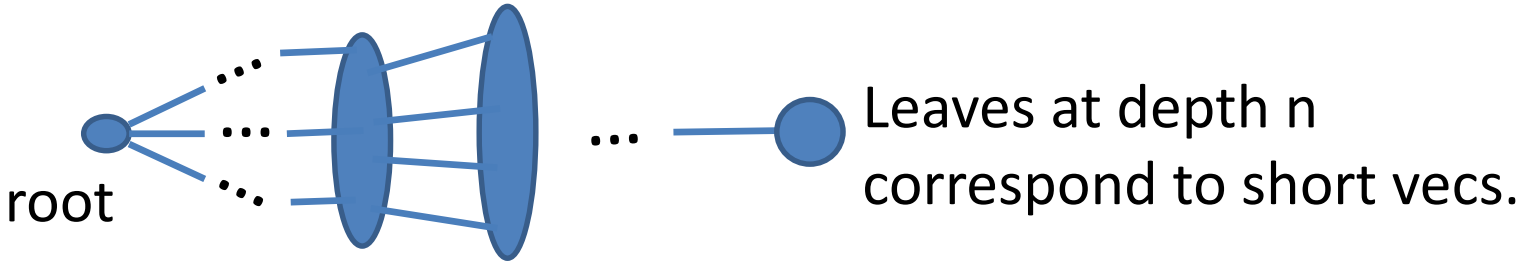

- Huge speed-up with pruned ENUM [SH@EC95,GNR@EC10]: tradeoff with success probability.

Gaussian heuristic assumption

- For a lattice L and a “normal” shaped $S \subseteq \mathbf{R}^n$, we have

$$\#(L \cap S) \approx \frac{\mathrm{vol}(S)}{\mathrm{vol}(L)}$$

- This approximates # nodes by the volume of searching area at each depth

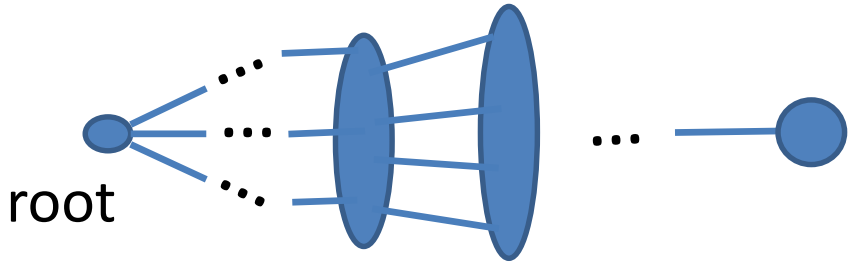

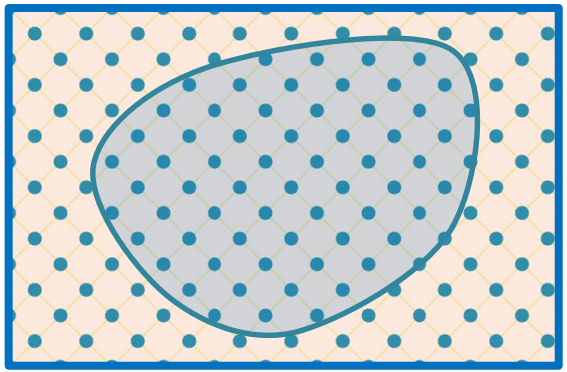

- Under GH, cost of tree enumeration

$$= \sum_{k=1}^{n} (\#\text{nodes at depth } k) \approx \sum_{k=1}^{n} \frac{\text{vol}(C_k)}{\text{vol}(\pi_{n-k}(B))}$$

- C_k is the cylinder-intersection defined by enumeration parameters $0 \leq R_1 \leq R_2 \leq \ldots \leq R_n$ [GNR@EC10]

$$C_k = \left\{ (x_1, \ldots, x_k) : \sum_{i=1}^{\ell} x_i^2 \leq R_\ell^2 \text{ for } \forall i = 1, \ldots, k \right\}$$

Example for k=3: $C_3 = \left\{ (x_1, x_2, x_3) : \begin{array}{ll} x_1^2 & \leq R_1^2 \\ x_1^2 + x_2^2 & \leq R_2^2 \\ x_1^2 + x_2^2 + x_3^2 & \leq R_3^2 \end{array} \right\}$

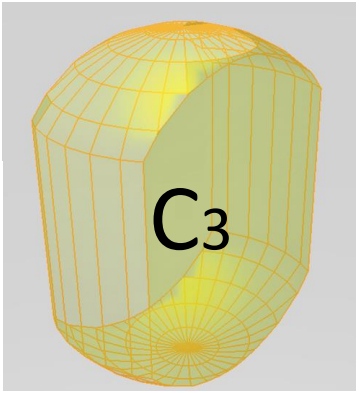

- The cost of pruned ENUM is the minimum of optimization problem

> Given: basis $B=(\boldsymbol{b}_1,\ldots,\boldsymbol{b}_n)$; target probability p_0; radius R_n
> Find: minimum Cost($R_1,\ldots,R_n$)
> Subject to: Prob($R_1,\ldots,R_n$)$\geq p_0$

where

$$\text{Cost}(R_1,\ldots,R_n) \triangleq \sum_{i=1}^{n} \frac{\text{vol}(C_k)}{\text{vol}(\pi_{n-k}(B))}$$

$$\text{Prob}(R_1,\ldots,R_n) \triangleq \frac{\text{vol}(C_n)}{\text{vol}(L)}$$

Note: we have to optimize n-variables $R_1,\ldots,R_n$

Pros of GNR pruned ENUM: speedups

Cost of pruned enumeration with success probability p is much smaller than p·(Cost of enumeration without pruning)

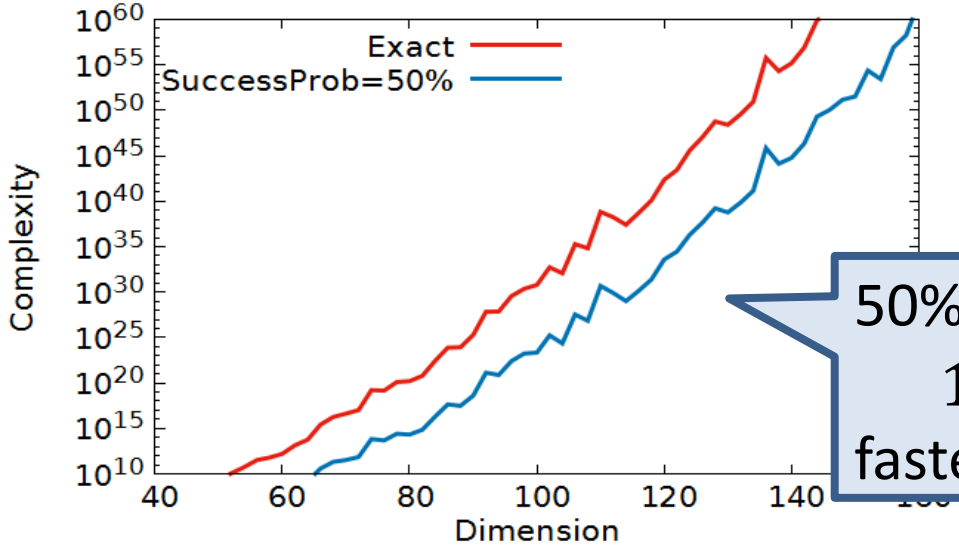

50% algorithm is about
$10^{10} \approx 33$bits
faster than exact alg.

Experiments on LLL-reduced bases

Cons of GNR pruned ENUM

1: No efficient method to find optimal radii: many parameters to opt.

- We propose a variant of the cross-entropy method
- Graph of $(R_1,\ldots,R_n)$ looks good, but no theoretical guarantee of optimality

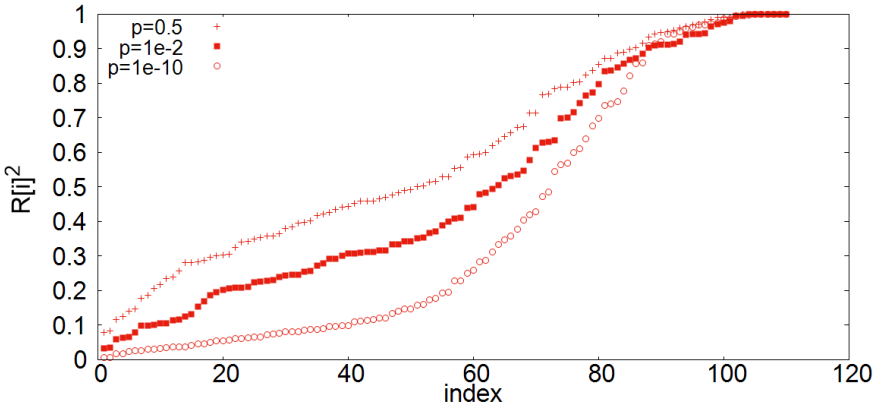

2: Non-trivial cost bounds for arbitrary p_0 unknown

- Naïve lower bound is useless
- We prove the first lower bound result for Cost$(R_1,\ldots,R_n)$

Agenda

- ~~Background and overview of our results~~
- ~~Pruned ENUM and cost estimation in [GNR@EC10]~~
- Lower bound via isoperimetry
- Linear lower bound of randomized ENUM and application to SVP-β

Isoperimetry: our key tool from math.

[Isoperimetry] If an n-dim. object C ⊆ Ball$_n$(1) has an orthogonal projection onto $\mathbf{R}^k$ whose volume is bounded by M,

Then, for the ball-cylinder intersection C':=$\left\{(x_1,\ldots,x_n)\in B_n(1):\sum_{i=1}^{k}x_i^2\le r^2\right\}$

vol(C)≤vol(C')

where r is taken so that the projection volume =M.

Example: k=2 and n=3

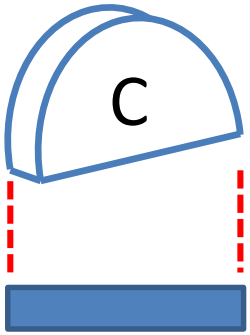

Projection is a bar

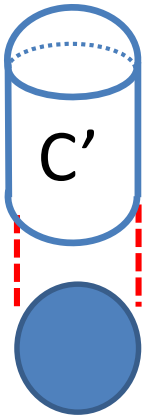

vol(C)≤vol(C')

Circle of equivalent area to bar

Observation on pruned ENUM

- Under GH,

Cost($R_1,\ldots,R_n$) $\triangleq \sum_{i=1}^{n} \frac{\mathrm{vol}(C_k)}{\mathrm{vol}(\pi_{n-k}(B))}$ and Prob($R_1,\ldots,R_n$) $\triangleq \frac{\mathrm{vol}(C_n)}{\mathrm{vol}(L)}$

- Observation:

Each C_k is the orthogonal projection of $C_n \subset$ Ball(R_n)

- Isoperimetry implies that

$\mathrm{vol}(C_n) \leq \mathrm{vol}(C_n')$

where C_n' is the intersection of ball and cylinder

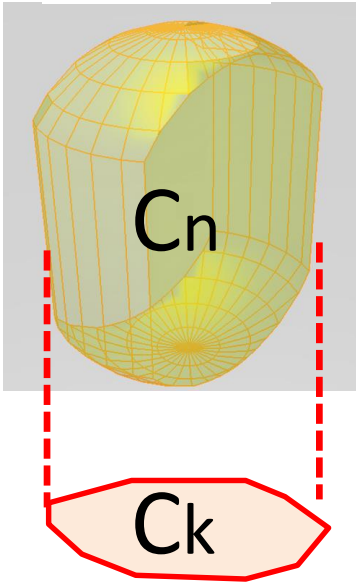

Analytic formula of the maximum volume

- *Isoperimetry* connects vol(C_n) with vol(C_k):

$$\mathrm{vol}(C_n) \le \mathrm{vol}(\text{Ball-cylinder intersection}) = V_n(R_n) \cdot \underbrace{I_{(R'_k/R_n)^2}\left(\frac{k}{2}, 1+\frac{n-k}{2}\right)}_{\text{Incomplete beta function}}$$

where R'_k is the radius satisfying $V_k(R'_k)=\mathrm{vol}(C_k)$

- This formula gives a lower bound for vol(C_k) if p=vol(C_n)/vol(L) is bounded
- The inverse incomplete beta function is implemented by the **boost library**

Advantages in implementation

- About 10 lines in C++ with the boost library
- Less than 10 ms on a standard desktop computer
- Deterministic algorithm

```
bkzfloat cost,lcost;
for (int i=iend;i>=istart+2;i--) {
    bkzfloat lf;
    lf = ibeta_inv_wrapper<bkzfloat>(0.5*(iend-i+1), 0.5*(i-istart),prob);
    lf = pow(lf,0.5*(iend-i+1));
    lcost *= radius / c[i];
    bkzfloat localcost = lcost * bkzconstants::vol_unit_ball(iend-i+1) * lf;
    cost += localcost;
}
return 0.5*cost;        //halved by symmetry
```

In contrast: our optimizing subroutine to find upper bounds is

- About 900 lines in C++, $\geqq$1-10 seconds to compute
- Output is not stable because it uses randomness

Experiment 1: Tightness of radii

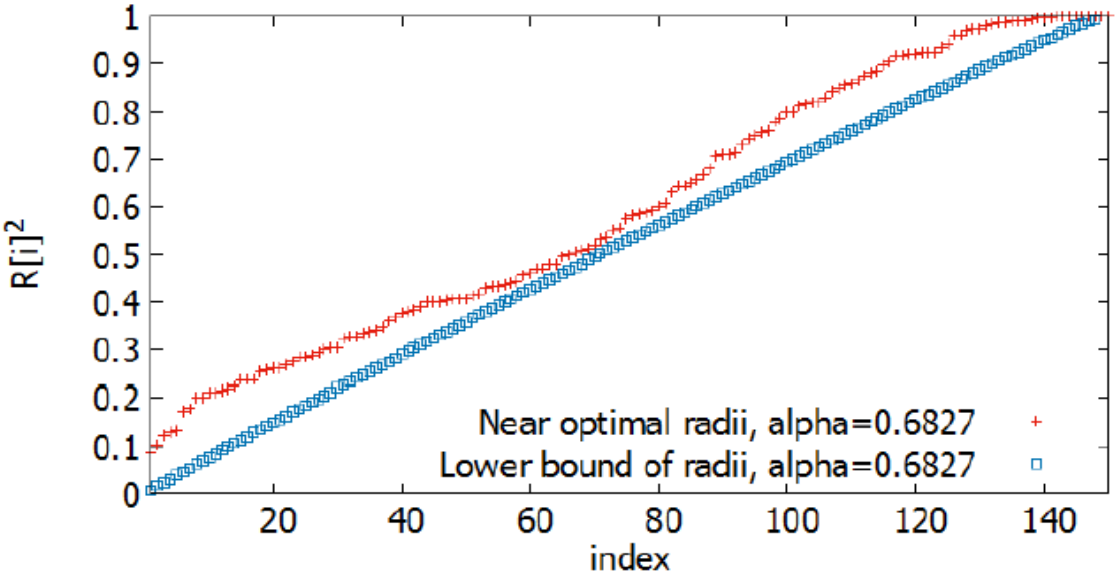

- Numerical experiments to compare upper vs lower bound $(R'_k)^2$

Experiment 2: Tightness of # nodes at depth k

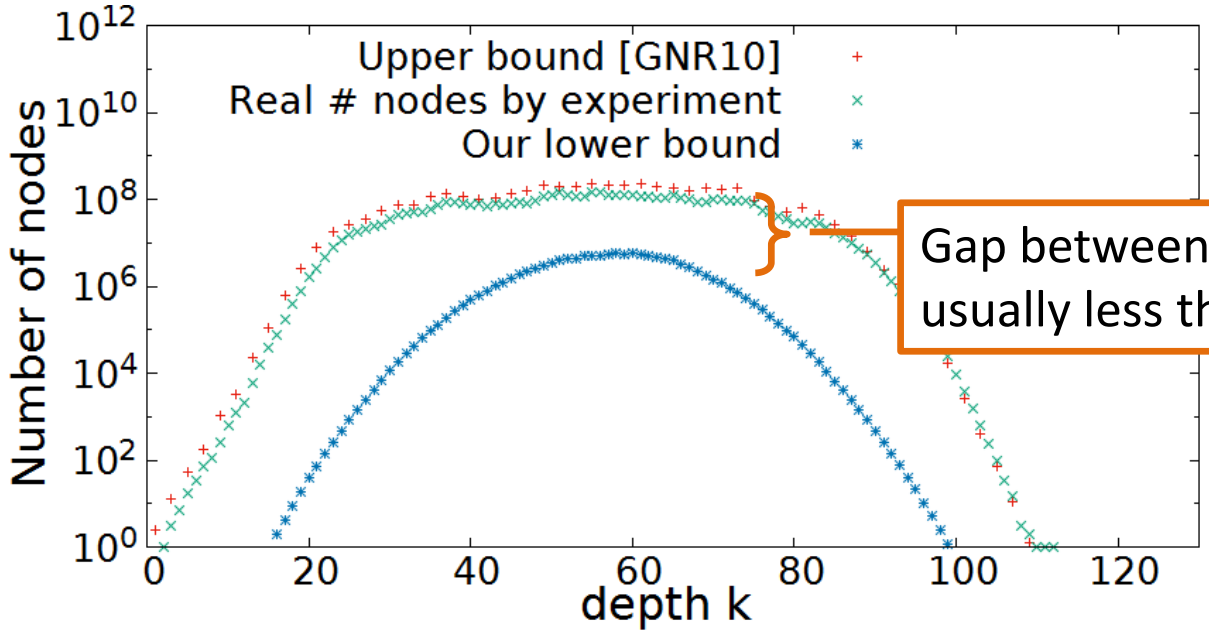

Gap between Upper and Lower is usually less than 20% in log-scale

- Numerical experiments to compare upper vs lower bound
- ENUM with (R=1.1GH, Dim=120, $p=10^{-6}$) for a BKZ reduced basis

Agenda

- ~~Background and overview of our results~~
- ~~Pruned ENUM and cost estimation in [GNR@EC10]~~
- ~~Lower bound via isoperimetry~~
- Linear lower bound of randomized ENUM and application to SVP-β

Lower bounds on randomizing strategy

- [Extreme pruning of GNR10] If we have many random bases $B_1,\ldots,B_M$, do ENUM with tiny probabilities $p_1,\ldots,p_M$

- The total cost $\sum_{i=1}^{M} \mathrm{Cost}(\text{Basis } B_i,\ \text{success prob.} = p_i) + \mathrm{Time}(\text{Randomization})$

  is much smaller than single ENUM with probability $p = \sum_{i=1}^{M} p_i$

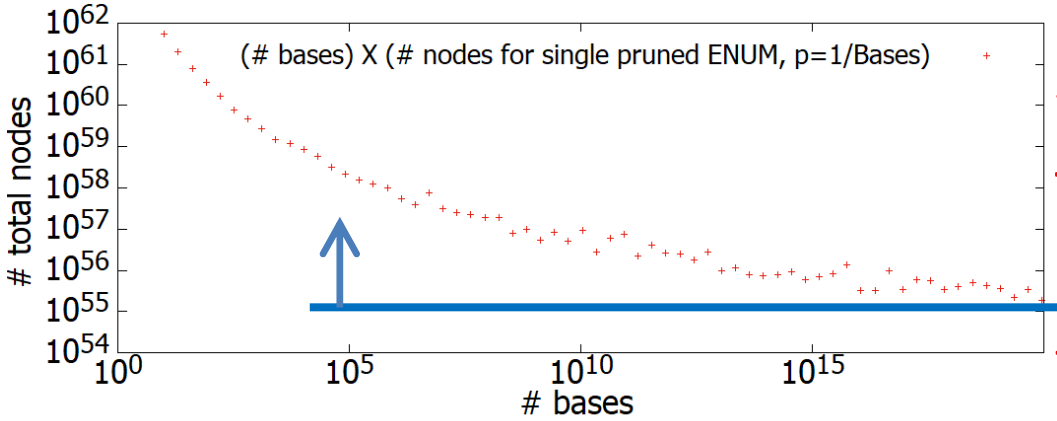

We proved that:
Total cost is lower bounded by a constant independent of #bases

Linear lower bound on randomizing strategy

- We proved that for a basis B and radius R, there is a constant C(B,R)
  (Cost of ENUM with probability p) ≧ p･C(B,R)
- Also, we have showed

$$\frac{(\text{LHS})}{p} \to C(B,R) \text{ if } p \to 0$$

- Gives limitations of randomization even with infinitely many bases:
  Cost(Extreme pruning with *global* probability 1)

$$\geq \sum_{i=1}^{M} p_i \cdot C(B_i, R) \geq C(B_{min}, R)$$

where B_{min} is the basis achieving best lower bound

Two scenarios for $C(B_{min},R)$

- A basis achieving $C(B_{min},R)$ gives us *the limitation* of extreme pruning and useful for security estimation of lattice crypto
- We give two scenarios for the type of bases that attackers in the future can efficiently generate
- State-of-the-art scenario:
  - HKZ is the best basis in practice
  - Strong BKZ-type algorithms try to approximate HKZ
- Conservative scenario:
  - Approximating Rankin problems can be done efficiently
  - Out of reach today

Application to hardness of SVP-β

- Comparing our lower bound vs. sieve lower bound to solve SVP-β
- State-of-the-art scenario: HKZ will be the practical best basis
- Conservative scenario: Rankin basis will be efficiently computable

Classical hardness

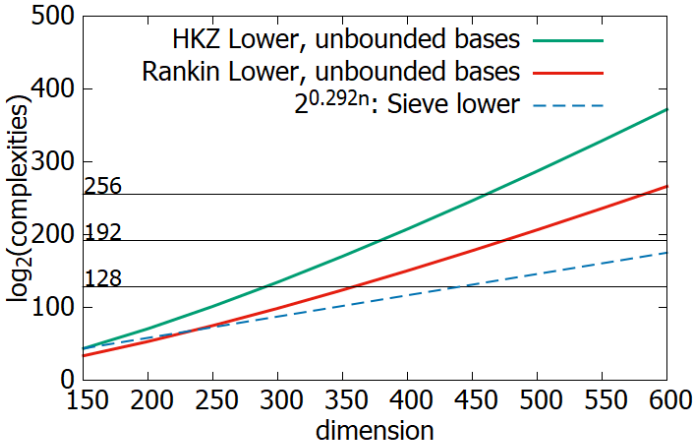

Quantum hardness

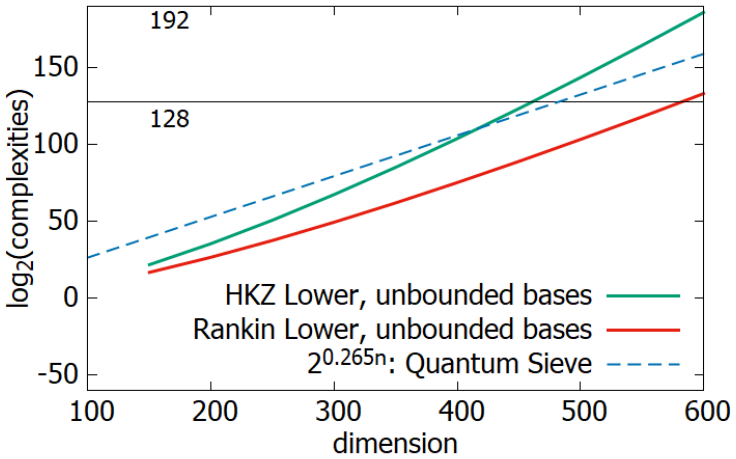

- From the graphs for Quantum, a conservative designer needs to change their parameters

Conclusion

1. Proving lower-bound costs for Gama-Nguyen-Regev's extreme pruning
2. First use of isoperimetry to (lattice) cryptography
3. Impact on parameters of lattice crypto
    - Provides lower bound costs on solving SVP-β by using extreme pruning
    - For typical dimensions,
        - Classical setting: ENUM is slower than Sieve
        - Quantum setting: ENUM is faster than Sieve
    - Thus, conservative designers need to update parameters

Open problems

- On [GNR10]'s extreme pruning ENUM
  - Tighter upper/lower bounds

- Adapt to other algorithms such as Discrete pruning ENUM, Sieve: unified lower bounds ?
  - Only trivial bound is known for discrete pruning ENUM [AN17]

Thank you for your attention

Full-version: https://eprint.iacr.org/2018/586