Revisiting the Expected Cost of Solving uSVP and Applications to LWE

Martin R. Albrecht[1]
Florian Göpfert[2,3]
Fernando Virdia[1]
Thomas Wunderer[3]

[1]Information Security Group, Royal Holloway, University of London,
[2]rockenstein AG,
[3]TU Darmstadt

4 December 2017
Asiacrypt 2017

Overview

Lattice reduction cost models

Our experiments

Experimental and theoretical results

New security estimates

Conclusions

Lattice reduction cost models

>>> Lattice reduction algorithms are a fundamental tool for cryptanalysis of lattice-based cryptographic schemes

>>> A common strategy is to use them to solve the Unique Shortest Vector Problem as part of 'primal lattice attacks'

>>> Costing such algorithms is therefore a fundamental step for choosing secure parameters

>>> Heads-up: cost models disagree on the asymptotic complexity

Lattice reduction cost models

>>> Lattice reduction algorithms are a fundamental tool for cryptanalysis of lattice-based cryptographic schemes

>>> A common strategy is to use them to solve the Unique Shortest Vector Problem as part of 'primal lattice attacks'

>>> Costing such algorithms is therefore a fundamental step for choosing secure parameters

>>> Heads-up: cost models disagree on the asymptotic complexity

- >>> [GN08] is the first systematic study of lattice reduction strategies
- >>> The work looks at using BKZ for solving Unique-SVP, using a statistical approach for estimating its effectiveness
- >>> A necessary condition for successful recovery is obtained
- >>> This approach is later applied to LWE embedding lattices in [AFG14]

>>> [GN08] is the first systematic study of lattice reduction strategies

>>> The work looks at using BKZ for solving Unique-SVP, using a statistical approach for estimating its effectiveness

>>> A necessary condition for successful recovery is obtained

>>> This approach is later applied to LWE embedding lattices in [AFG14]

>>> Let Λ be our lattice of dimension d with a unique shortest vector $\mathbf{v}$ (up to $\pm$ sign), and let λ_i be the i^{th} minima

>>> Let δ be the Hermite factor $\iff$ BKZ recovers vectors long $\approx \delta^d Vol(\Lambda)^{1/d}$

>>> If

$$\lambda_2(\Lambda)/\lambda_1(\Lambda) > \tau\delta^d, \quad \text{for } \tau \in (0,1) \tag{1}$$

the shortest vector is recovered

>>> τ is estimated experimentally

>>> (1) $\Rightarrow$ optimal number of LWE samples m_{2008} and BKZ block size β_{2008} to run the primal attack

>>> We refer to this work as *the 2008 model*

- >>> Let Λ be our lattice of dimension d with a unique shortest vector $\mathbf{v}$ (up to $\pm$ sign), and let λ_i be the i^{th} minima
- >>> Let δ be the Hermite factor $\iff$ BKZ recovers vectors long $\approx \delta^d Vol(\Lambda)^{1/d}$
- >>> If

$$\lambda_2(\Lambda)/\lambda_1(\Lambda) > \tau\delta^d, \quad \text{for } \tau \in (0,1) \tag{1}$$

  the shortest vector is recovered
- >>> τ is estimated experimentally
- >>> (1) $\Rightarrow$ optimal number of LWE samples m_{2008} and BKZ block size β_{2008} to run the primal attack
- >>> We refer to this work as *the 2008 model*

>>> Let Λ be our lattice of dimension d with a unique shortest vector $\mathbf{v}$ (up to $\pm$ sign), and let λ_i be the i^{th} minima

>>> Let δ be the Hermite factor $\iff$ BKZ recovers vectors long $\approx \delta^d Vol(\Lambda)^{1/d}$

>>> If

$$\lambda_2(\Lambda)/\lambda_1(\Lambda) > \tau\delta^d, \quad \texttt{for } \tau \in (0,1) \tag{1}$$

the shortest vector is recovered

>>> τ is estimated experimentally

>>> (1) $\Rightarrow$ optimal number of LWE samples m_{2008} and BKZ block size β_{2008} to run the primal attack

>>> We refer to this work as *the 2008 model*

>>> Let Λ be our lattice of dimension d with a unique shortest vector $\mathbf{v}$ (up to $\pm$ sign), and let λ_i be the i^{th} minima

>>> Let δ be the Hermite factor $\iff$ BKZ recovers vectors long $\approx \delta^d Vol(\Lambda)^{1/d}$

>>> If

$$\lambda_2(\Lambda)/\lambda_1(\Lambda) > \tau\delta^d, \quad \text{for } \tau \in (0,1) \qquad (1)$$

the shortest vector is recovered

>>> τ is estimated experimentally

>>> (1) $\Rightarrow$ optimal number of LWE samples m_{2008} and BKZ block size β_{2008} to run the primal attack

>>> We refer to this work as *the 2008 model*

>>> [ADPS16] introduces a new success condition for solving Unique-SVP with BKZ when $\|\mathbf{v}\|$ is known

>>> The strategy is based on the Geometric Series Assumption, and on the structure of the BKZ algorithm

>>> We refer to this work as *the 2016 model*

>>> To explain the condition we will first review how BKZ works

>>> [ADPS16] introduces a new success condition for solving Unique-SVP with BKZ when $\|\mathbf{v}\|$ is known

>>> The strategy is based on the Geometric Series Assumption, and on the structure of the BKZ algorithm

>>> We refer to this work as *the 2016 model*

>>> To explain the condition we will first review how BKZ works

>>> Let $\{\mathbf{b}_i\}_i$ be Λ's basis, $\{\mathbf{b}_i^*\}_i$ their Gram-Schmidt vectors, and $\mathbf{v}_i^*$ the projection of $\mathbf{v} \perp \{\mathbf{b}_j^*\}_{j=1}^{i-1}$

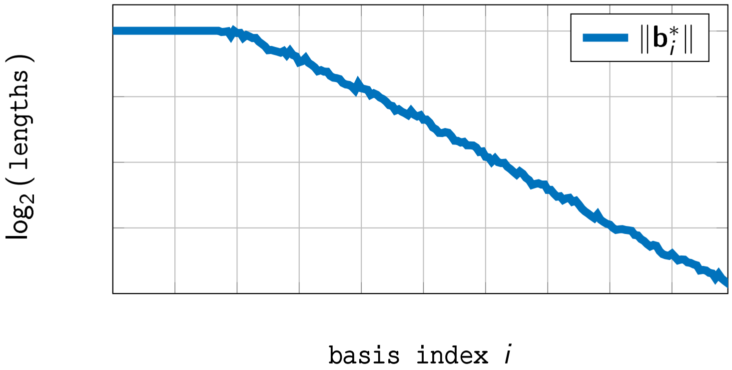

>>> Let $\{\mathbf{b}_i\}_i$ be Λ's basis, $\{\mathbf{b}_i^*\}_i$ their Gram-Schmidt vectors, and $\mathbf{v}_i^*$ the projection of $\mathbf{v} \perp \{\mathbf{b}_j^*\}_{j=1}^{i-1}$

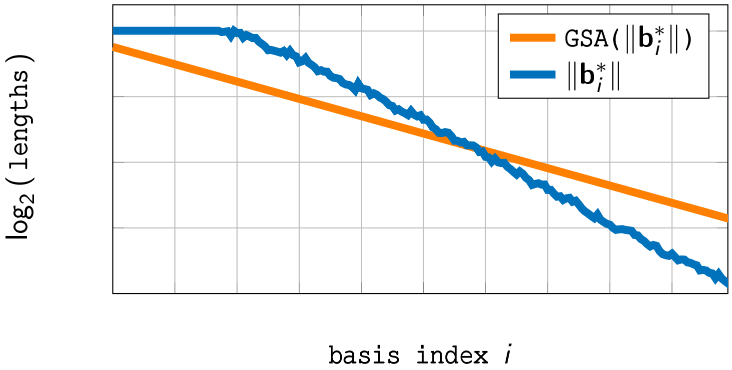

>>> Let $\{\mathbf{b}_i\}_i$ be Λ's basis, $\{\mathbf{b}_i^*\}_i$ their Gram-Schmidt vectors, and $\mathbf{v}_i^*$ the projection of $\mathbf{v} \perp \{\mathbf{b}_j^*\}_{j=1}^{i-1}$

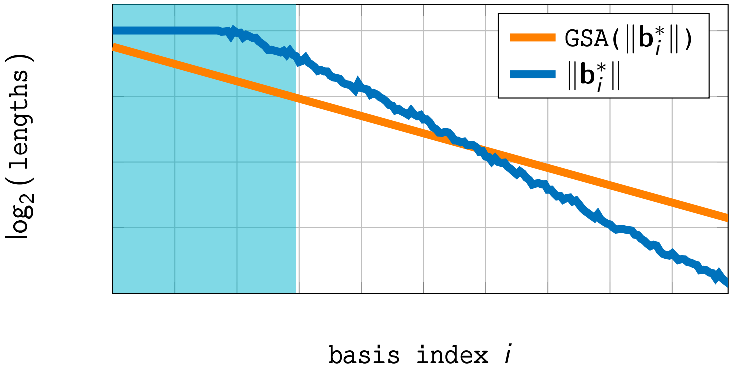

>>> Let $\{\mathbf{b}_i\}_i$ be Λ's basis, $\{\mathbf{b}_i^*\}_i$ their Gram-Schmidt vectors, and $\mathbf{v}_i^*$ the projection of $\mathbf{v} \perp \{\mathbf{b}_j^*\}_{j=1}^{i-1}$

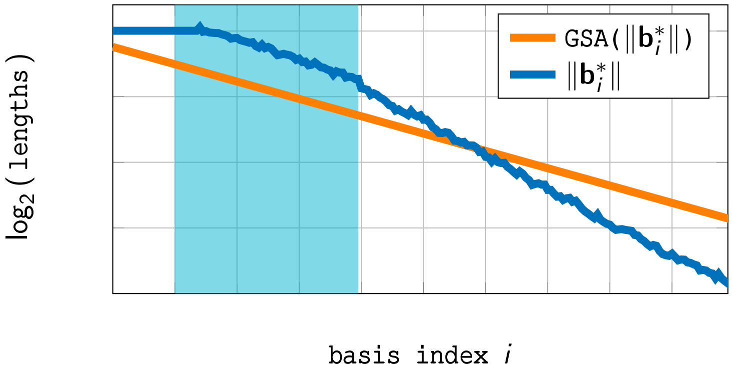

>>> Let $\{\mathbf{b}_i\}_i$ be Λ's basis, $\{\mathbf{b}_i^*\}_i$ their Gram-Schmidt vectors, and $\mathbf{v}_i^*$ the projection of $\mathbf{v} \perp \{\mathbf{b}_j^*\}_{j=1}^{i-1}$

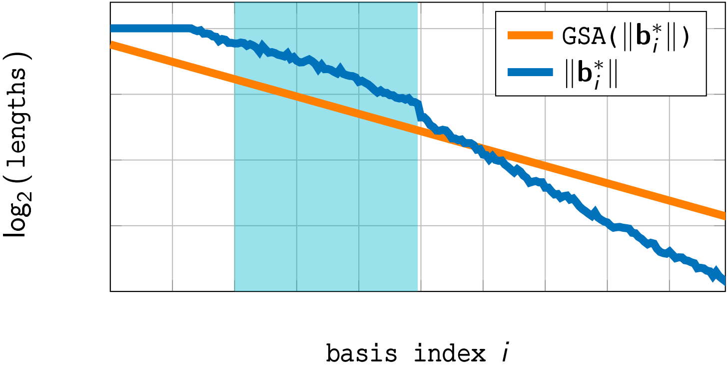

>>> Let $\{\mathbf{b}_i\}_i$ be Λ's basis, $\{\mathbf{b}_i^*\}_i$ their Gram-Schmidt vectors, and $\mathbf{v}_i^*$ the projection of $\mathbf{v} \perp \{\mathbf{b}_j^*\}_{j=1}^{i-1}$

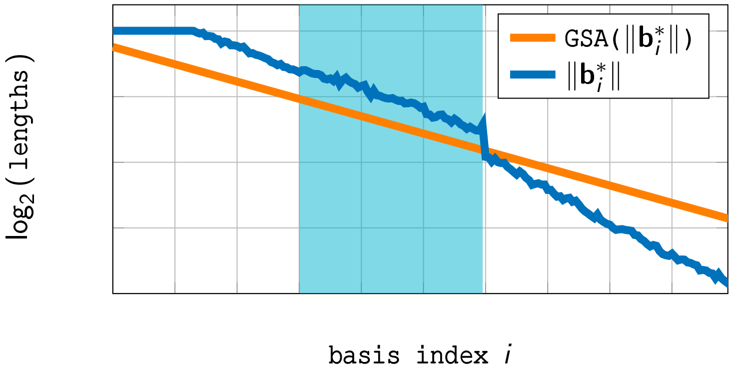

>>> Let $\{\mathbf{b}_i\}_i$ be Λ's basis, $\{\mathbf{b}_i^*\}_i$ their Gram-Schmidt vectors, and $\mathbf{v}_i^*$ the projection of $\mathbf{v} \perp \{\mathbf{b}_j^*\}_{j=1}^{i-1}$

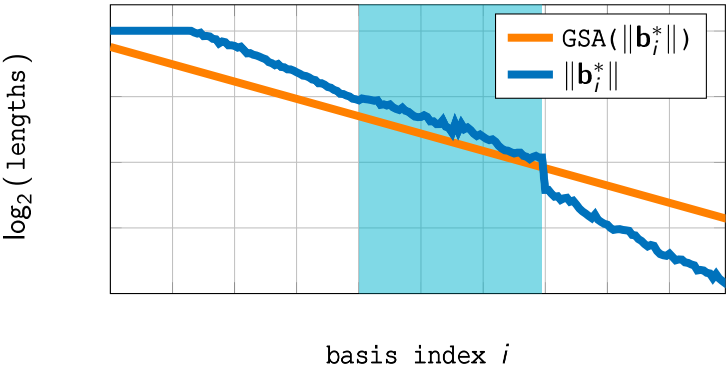

>>> Let $\{\mathbf{b}_i\}_i$ be Λ's basis, $\{\mathbf{b}_i^*\}_i$ their Gram-Schmidt vectors, and $\mathbf{v}_i^*$ the projection of $\mathbf{v} \perp \{\mathbf{b}_j^*\}_{j=1}^{i-1}$

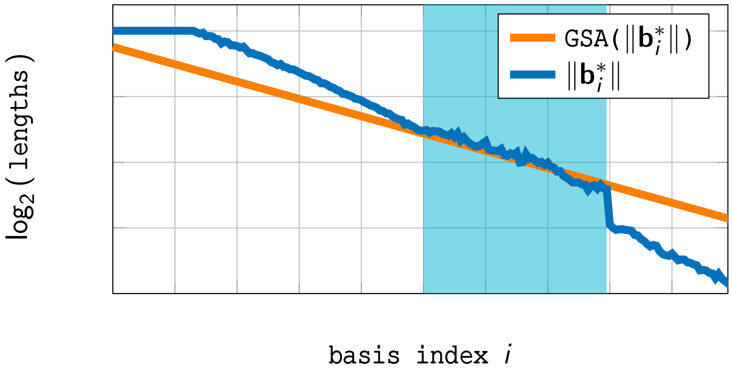

>>> Let $\{\mathbf{b}_i\}_i$ be Λ's basis, $\{\mathbf{b}_i^*\}_i$ their Gram-Schmidt vectors, and $\mathbf{v}_i^*$ the projection of $\mathbf{v} \perp \{\mathbf{b}_j^*\}_{j=1}^{i-1}$

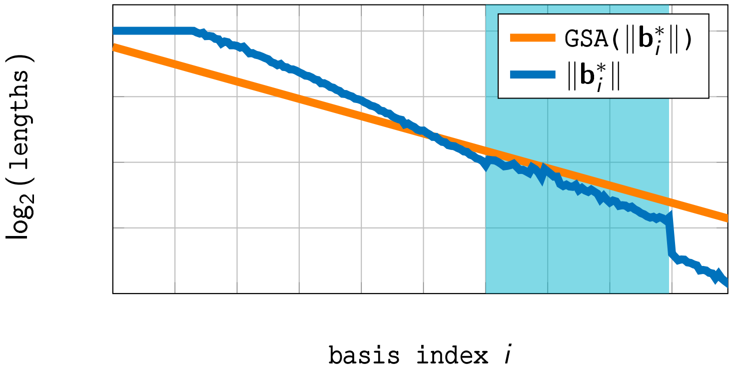

>>> Let $\{\mathbf{b}_i\}_i$ be Λ's basis, $\{\mathbf{b}_i^*\}_i$ their Gram-Schmidt vectors, and $\mathbf{v}_i^*$ the projection of $\mathbf{v} \perp \{\mathbf{b}_j^*\}_{j=1}^{i-1}$

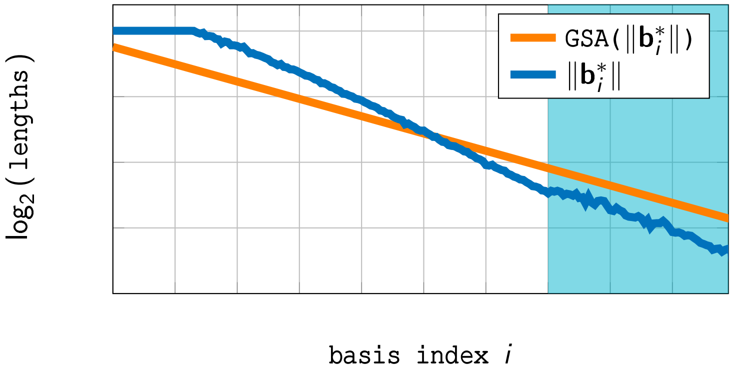

>>> Let $\{\mathbf{b}_i\}_i$ be Λ's basis, $\{\mathbf{b}_i^*\}_i$ their Gram-Schmidt vectors, and $\mathbf{v}_i^*$ the projection of $\mathbf{v} \perp \{\mathbf{b}_j^*\}_{j=1}^{i-1}$

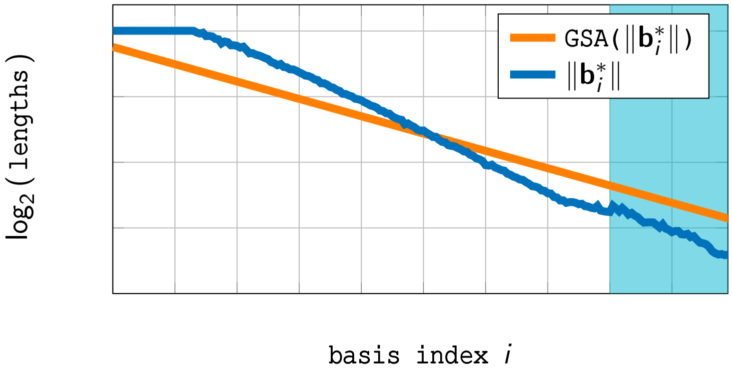

\>>> Let $\{\mathbf{b}_i\}_i$ be Λ's basis, $\{\mathbf{b}_i^*\}_i$ their Gram-Schmidt vectors, and $\mathbf{v}_i^*$ the projection of $\mathbf{v} \perp \{\mathbf{b}_j^*\}_{j=1}^{i-1}$

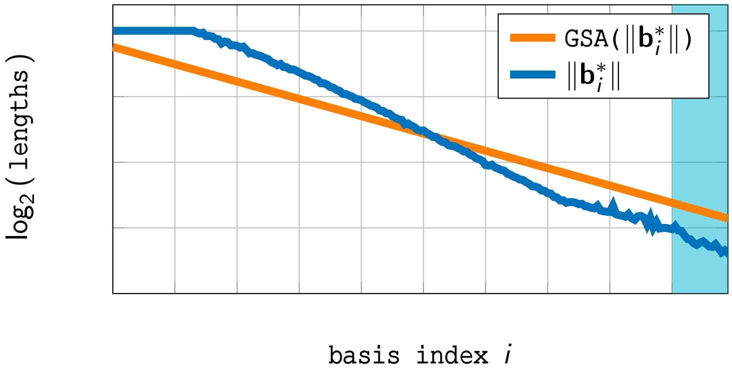

>>> Let $\{\mathbf{b}_i\}_i$ be Λ's basis, $\{\mathbf{b}_i^*\}_i$ their Gram-Schmidt vectors, and $\mathbf{v}_i^*$ the projection of $\mathbf{v} \perp \{\mathbf{b}_j^*\}_{j=1}^{i-1}$

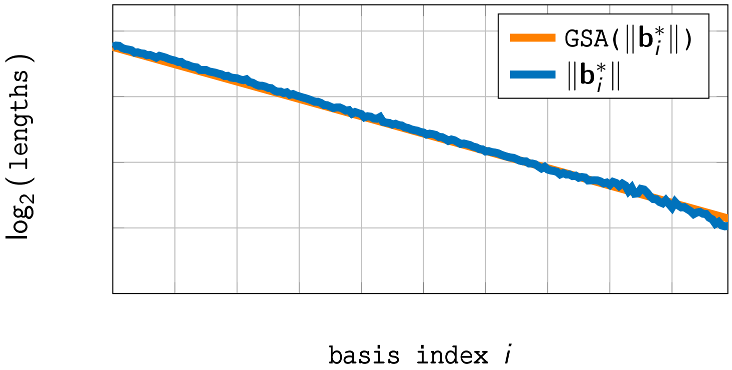

>>> Let $\{\mathbf{b}_i\}_i$ be Λ's basis, $\{\mathbf{b}_i^*\}_i$ their Gram-Schmidt vectors, and $\mathbf{v}_i^*$ the projection of $\mathbf{v} \perp \{\mathbf{b}_j^*\}_{j=1}^{i-1}$

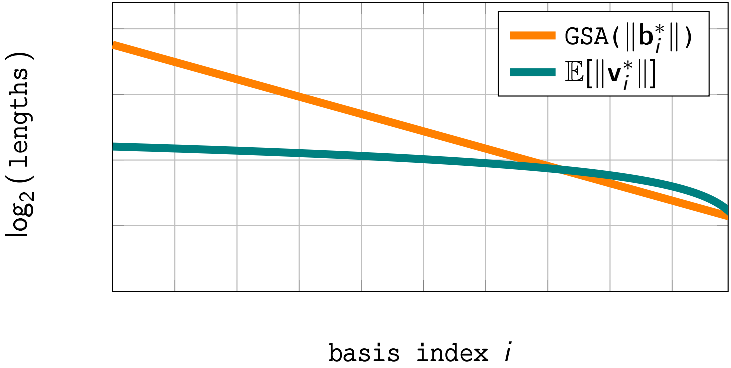

>>> Let $\{\mathbf{b}_i\}_i$ be Λ's basis, $\{\mathbf{b}_i^*\}_i$ their Gram-Schmidt vectors, and $\mathbf{v}_i^*$ the projection of $\mathbf{v} \perp \{\mathbf{b}_j^*\}_{j=1}^{i-1}$

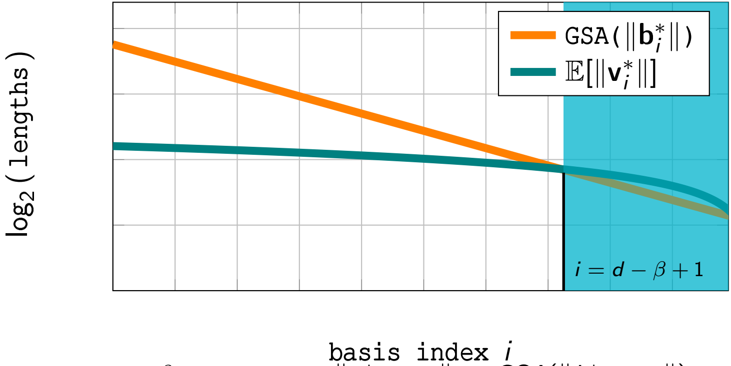

>>> Choose β such that $\|\mathbf{v}_{d-\beta+1}^*\| < GSA(\|b_{d-\beta+1}^*\|)$

>>> Instantly solves Decision-LWE

>>> Let $\{\mathbf{b}_i\}_i$ be Λ's basis, $\{\mathbf{b}_i^*\}_i$ their Gram-Schmidt vectors, and $\mathbf{v}_i^*$ the projection of $\mathbf{v} \perp \{\mathbf{b}_j^*\}_{j=1}^{i-1}$

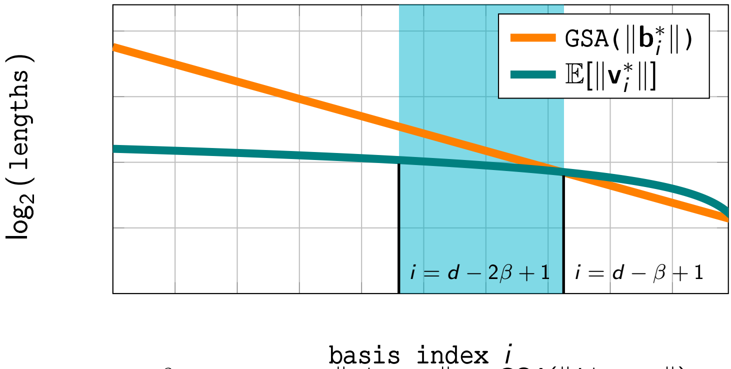

>>> Choose β such that $\|\mathbf{v}_{d-\beta+1}^*\| < GSA(\|b_{d-\beta+1}^*\|)$

>>> Instantly solves Decision-LWE

>>> Let $\{\mathbf{b}_i\}_i$ be Λ's basis, $\{\mathbf{b}_i^*\}_i$ their Gram-Schmidt vectors, and $\mathbf{v}_i^*$ the projection of $\mathbf{v} \perp \{\mathbf{b}_j^*\}_{j=1}^{i-1}$

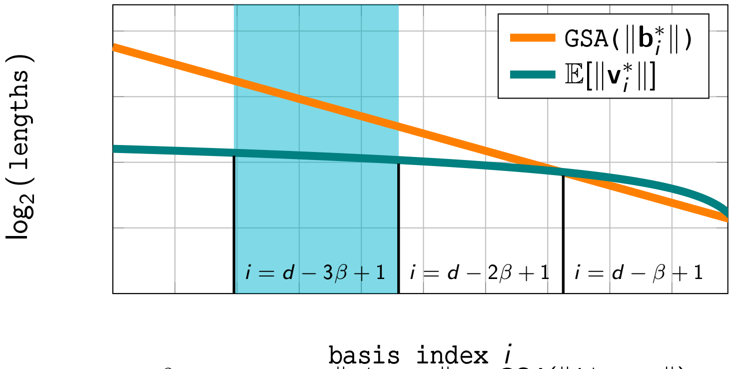

>>> Choose β such that $\|\mathbf{v}_{d-\beta+1}^*\| < GSA(\|b_{d-\beta+1}^*\|)$

>>> Instantly solves Decision-LWE

>>> Let $\{\mathbf{b}_i\}_i$ be Λ's basis, $\{\mathbf{b}_i^*\}_i$ their Gram-Schmidt vectors, and $\mathbf{v}_i^*$ the projection of $\mathbf{v} \perp \{\mathbf{b}_j^*\}_{j=1}^{i-1}$

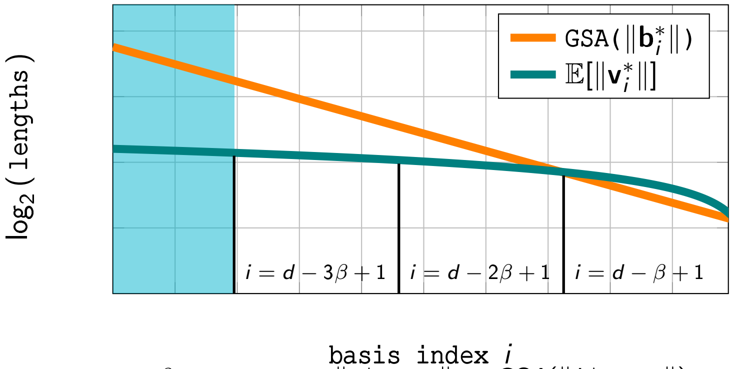

>>> Choose β such that $\|\mathbf{v}_{d-\beta+1}^*\| < GSA(\|b_{d-\beta+1}^*\|)$

>>> Instantly solves Decision-LWE

>>> Should solve Search-LWE with at most $\lceil d/\beta \rceil - 1$ more SVP oracle calls

>>> The two models disagree on the primal attack's asymptotic complexity

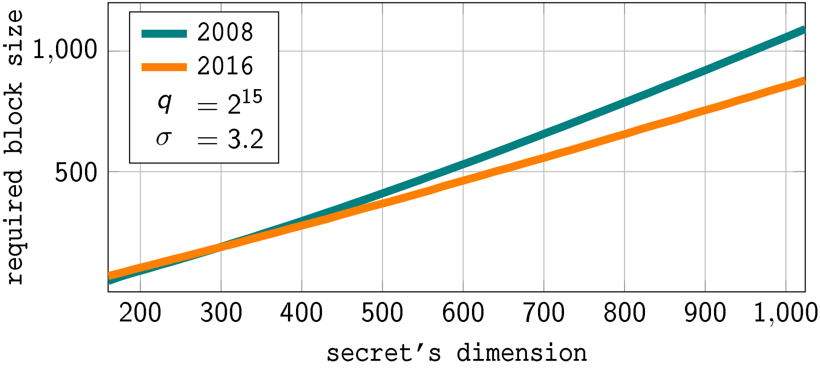

>>> We decided to experimentally investigate the accuracy of the 2016 model's predictions

>>> The two models disagree on the primal attack's asymptotic complexity

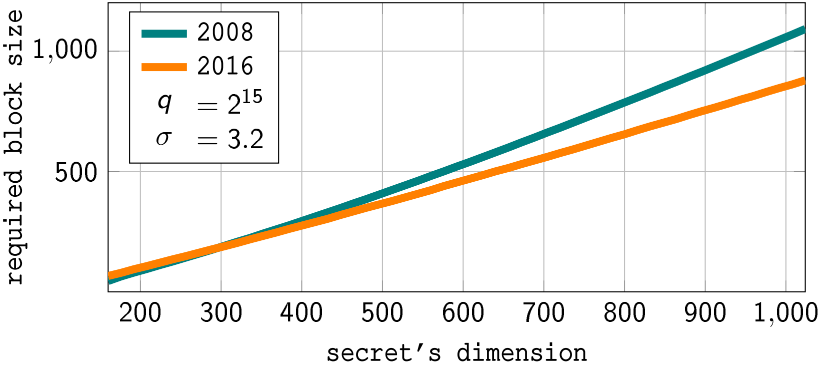

>>> We decided to experimentally investigate the accuracy of the 2016 model's predictions

Our experiments

>>> Given (n, q, σ), the 2008 model provides parameters (m_{2008}, β_{2008}) for 10% recovery probability [AFG14]

>>> We pick (m_{2016}, β_{2016}) according to [ADPS16], run BKZ2 and measure the recovery rate

>>> We instrument BKZ to take detailed statistics about the $\mathbf{v}_i^*$ length and moment of recovery

>>> To simplify analysis we make some changes to subroutine calls to LLL

>>> All our experiments were run using the FpyLLL lattice reduction library [FPL17, FPY17]

Our experiments

>>> Given (n, q, σ), the 2008 model provides parameters (m_{2008}, β_{2008}) for 10% recovery probability [AFG14]

>>> We pick (m_{2016}, β_{2016}) according to [ADPS16], run BKZ2 and measure the recovery rate

>>> We instrument BKZ to take detailed statistics about the $\mathbf{v}_i^*$ length and moment of recovery

>>> To simplify analysis we make some changes to subroutine calls to LLL

>>> All our experiments were run using the FpyLLL lattice reduction library [FPL17, FPY17]

Our experiments

>>> Given (n, q, σ), the 2008 model provides parameters (m_{2008}, β_{2008}) for 10% recovery probability [AFG14]

>>> We pick (m_{2016}, β_{2016}) according to [ADPS16], run BKZ2 and measure the recovery rate

>>> We instrument BKZ to take detailed statistics about the $\mathbf{v}_i^*$ length and moment of recovery

>>> To simplify analysis we make some changes to subroutine calls to LLL

>>> All our experiments were run using the FpyLLL lattice reduction library 😍 [FPL17, FPY17]

Results

LWE parameters			[ADPS16]		Experiments		
n	q	σ	β_{2016}	m_{2016}	β	#	recovery rate
65	521	3.2	56	182	56	10000	93.3%
					51		52.8%
					46		4.8%
80	1031	3.2	60	204	60	1000	94.2%
					55		60.6%
					50		8.9%
					45		0.2%
100	2053	3.2	67	243	67	500	88.8%
					62		39.6%
					57		5.8%
					52		0.2%
108	2053	3.2	77	261	77	5	100.0%
110	2053	3.2	78	272	78	5	100.0%

>>> Experiments agree with the 2016 model, but we noticed two unexpected behaviours

>>> First, while expecting BKZ to recover $\mathbf{v}^*_{d-\beta+1}$, for small experiments we observed

>>> Experiments agree with the 2016 model, but we noticed two unexpected behaviours

>>> First, while expecting BKZ to recover $\mathbf{v}^*_{d-\beta+1}$, for small experiments we observed

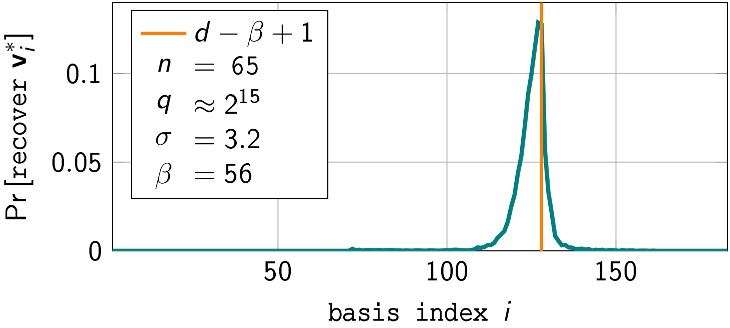

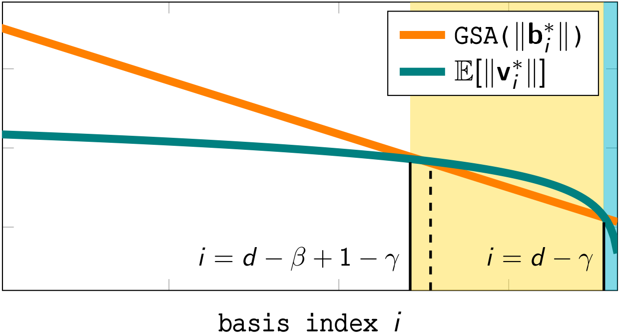

- >>> $\mathbf{v}_i^*$ is first recovered at the rightmost intersection at $i = d - \gamma$
- >>> In the next tour this projection is extended at $i = d - \beta + 1 - \gamma$
- >>> The double intersection is not common for cryptographically chosen parameters, and can be easily avoided

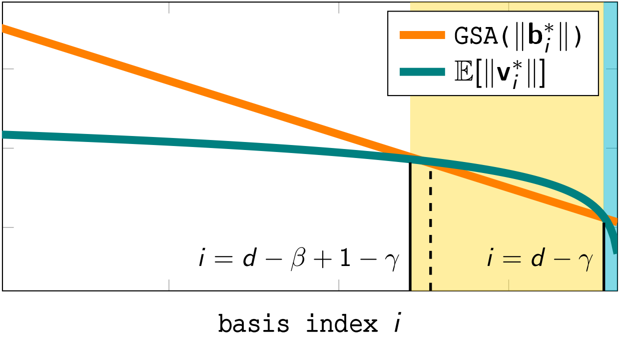

>>> $\mathbf{v}_i^*$ is first recovered at the rightmost intersection at $i = d - \gamma$

>>> In the next tour this projection is extended at $i = d - \beta + 1 - \gamma$

>>> The double intersection is not common for cryptographically chosen parameters, and can be easily avoided

>>> Second, 99.7%+ of the time $\mathbf{v}$ is recovered immediately after the SVP oracle finds its projection

>>> We model the state of the bases after first finding $\mathbf{v}^*_{d-\beta+1}$

>>> Second, 99.7%+ of the time $\mathbf{v}$ is recovered immediately after the SVP oracle finds its projection

>>> We model the state of the bases after first finding $\mathbf{v}^*_{d-\beta+1}$

>>> Lemma For $\beta > 40$, Size Reduction recovers $\mathbf{v}$ from $\mathbf{v}^*_{d-\beta+1}$ with overwhelming probability $\square$

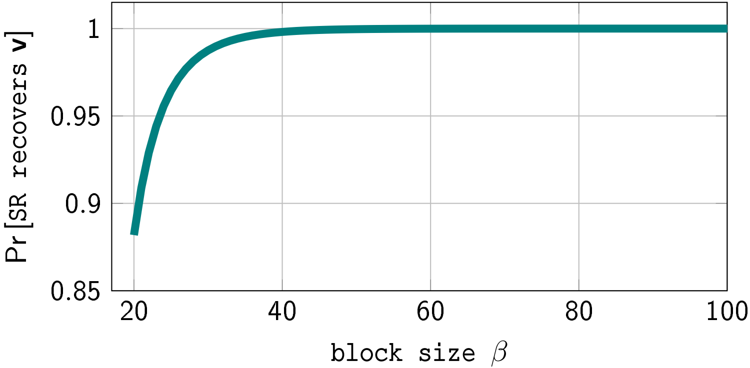

New security estimates

>>> We added the 2016 model to the LWE estimator from [APS15], and used it to recost the primal attack against proposed schemes (as of May 2017)

>>> For each scheme we used their proposed cost strategy

New security estimates

>>> We added the 2016 model to the LWE estimator from [APS15], and used it to recost the primal attack against proposed schemes (as of May 2017)

>>> For each scheme we used their proposed cost strategy

Scheme	Estimate as of May 17	Our estimate
Lizard [CKLS16a, CKLS16b]	129.7--131.6	85.9--88.7
TESLA [BG14, ABBD15]	71.0--142.0	61.5--122.4
SEAL v2.1 [CLP17]	97.6--130.5	99.6--129.5

>>> Security estimates for Lizard (PKE), TESLA (Signatures) and SEAL 2.1 (FHE) under the 2016 model, as of May 2017; more in the paper

>>> Some schemes were parametrised against the dual attack from [Alb17], which is still (often) cheaper against sparse and small secrets. Nontheless, in those cases the gap between primal and dual attack narrows

Conclusions

>>> We confirmed the validity of the 2016 model [ADPS16]

>>> Some existing lattice based schemes may need reparametrisation to resist cryptanalysis via lattice reduction

>>> The double intersection observation and the difference in success probability between models tell a cautionary tale about extrapolating asymptotics from small dimensional experiments

Conclusions

>>> We confirmed the validity of the 2016 model [ADPS16]

>>> Some existing lattice based schemes may need reparametrisation to resist cryptanalysis via lattice reduction

>>> The double intersection observation and the difference in success probability between models tell a cautionary tale about extrapolating asymptotics from small dimensional experiments

Conclusions

>>> We confirmed the validity of the 2016 model [ADPS16]

>>> Some existing lattice based schemes may need reparametrisation to resist cryptanalysis via lattice reduction

>>> The double intersection observation and the difference in success probability between models tell a cautionary tale about extrapolating asymptotics from small dimensional experiments

Thank you

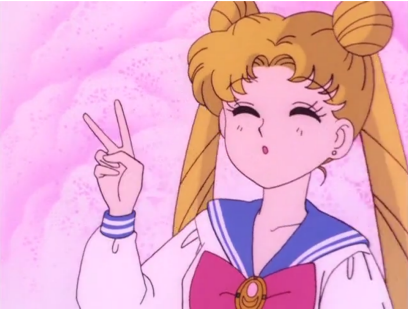

\>>> Paper @ https://ia.cr/2017/815

\>>> Experiments (code && data) @
https://github.com/fvirdia/agvw17-code-data

\>>> Estimator [APS15] @
https://bitbucket.org/malb/lwe-estimator

[ABBD15] Erdem Alkim, Nina Bindel, Johannes Buchmann, and Özgür Dagdelen.
TESLA: Tightly-secure efficient signatures from standard lattices.

Cryptology ePrint Archive, Report 2015/755, 2015.
http://eprint.iacr.org/2015/755.

[ADPS16] Erdem Alkim, Léo Ducas, Thomas Pöppelmann, and Peter Schwabe.
Post-quantum key exchange - A new hope.
In Thorsten Holz and Stefan Savage, editors, *25th USENIX Security Symposium, USENIX Security 16*, pages 327--343. USENIX Association, 2016.

[AFG14] Martin R. Albrecht, Robert Fitzpatrick, and Florian Göpfert.
On the efficacy of solving LWE by reduction to unique-SVP.
In Hyang-Sook Lee and Dong-Guk Han, editors, *ICISC 13*, volume 8565 of *LNCS*, pages 293--310. Springer, Heidelberg, November 2014.

[Alb17] Martin R. Albrecht.
On dual lattice attacks against small-secret LWE and parameter choices in HElib and SEAL.
In Jean-Sébastien Coron and Jesper Buus Nielsen, editors, *EUROCRYPT 2017, Part II*, volume 10211 of *LNCS*, pages 103--129. Springer, Heidelberg, April / May 2017.

[APS15] Martin R. Albrecht, Rachel Player, and Sam Scott.
On the concrete hardness of Learning with Errors.

Journal of Mathematical Cryptology, 9(3):169--203, 2015.

[BG14] Shi Bai and Steven D. Galbraith.
An improved compression technique for signatures based on learning with errors.
In Josh Benaloh, editor, *CT-RSA 2014*, volume 8366 of *LNCS*, pages 28--47. Springer, Heidelberg, February 2014.

[CKLS16a] Jung Hee Cheon, Duhyeong Kim, Joohee Lee, and Yongsoo Song.
Lizard: Cut off the tail! Practical post-quantum public-key encryption from LWE and LWR.
Cryptology ePrint Archive, Report 2016/1126 (20161222:071525), 2016.
http://eprint.iacr.org/2016/1126/20161222:071525.

[CKLS16b] Jung Hee Cheon, Duhyeong Kim, Joohee Lee, and Yongsoo Song.
Lizard: Cut off the tail! Practical post-quantum public-key encryption from LWE and LWR.
Cryptology ePrint Archive, Report 2016/1126, 2016.
http://eprint.iacr.org/2016/1126.

[CLP17] Hao Chen, Kim Laine, and Rachel Player.
Simple encrypted arithmetic library - SEAL v2.1.
Cryptology ePrint Archive, Report 2017/224, 2017.
http://eprint.iacr.org/2017/224.

[FPL17] The FPLLL development team.

fplll, a lattice reduction library.
Available at https://github.com/fplll/fplll, 2017.

[FPY17] The FPYLLL development team.
fypll1, a Python (2 and 3) wrapper for fplll.
Available at https://github.com/fplll/fpylll, 2017.

[GN08] Nicolas Gama and Phong Q. Nguyen.
Predicting lattice reduction.
In Nigel P. Smart, editor, *EUROCRYPT 2008*, volume 4965 of *LNCS*, pages 31--51. Springer, Heidelberg, April 2008.