Challenges for lattice cryptography

Steven Galbraith



University of Auckland, New Zealand

Apology

- I am not an expert on lattice cryptography or PQ crypto.
- I am not going to talk about fully homomorphic encryption or multilinear maps or indistinguishibility obfuscation.
- I am only talking about classical cryptosystems and classical attacks.

Plan

- LWE/(I)SIS and lattice attacks
- A confusion of parameters and a lack of understanding
- Approximate common divisors
- Search-to-decision reductions and adaptive attacks
- Signatures

Please ask questions at any time.

LWE and (I)SIS

- Let $m, n, q \in \mathbb{N}$ with $m > n$.
- Fix a distribution $\mathcal{D} \subseteq \mathbb{Z}^m$.
  Maybe uniform distribution on $\{0,1\}^n$, or discrete Gaussian distribution on $\mathbb{Z}$ with standard deviation $\sigma \ll q$.
- LWE (Regev): Given $(\mathbf{A}, \mathbf{b})$ where $\mathbf{A}$ is an $m \times n$ matrix to find $(\mathbf{s}, \mathbf{e}) \in \mathbb{Z}^n \times \mathbb{Z}^m$, if they exist, such that $\mathbf{e}$ is a likely sample from $\mathcal{D}$ and $\mathbf{b} = \mathbf{A}\mathbf{s} + \mathbf{e} \pmod{q}$.
- (I)SIS (Ajtai): Given $(\mathbf{A}, \mathbf{b})$ where $\mathbf{A}$ is an $n \times m$ matrix to find $\mathbf{s}$ (if it exists) that is a likely sample from $\mathcal{D}$ such that $\mathbf{b} = \mathbf{A}\mathbf{s} \pmod{q}$.
  This is the "inhomogeneous SIS problem" ISIS.
  SIS is the case $\mathbf{b} = \mathbf{0}, \mathbf{s} \neq \mathbf{0}$.
- Regev and Ajtai give very strong evidence that these are hard problems.
- Bewildering parameters m, n, q, σ.

LWE and (I)SIS

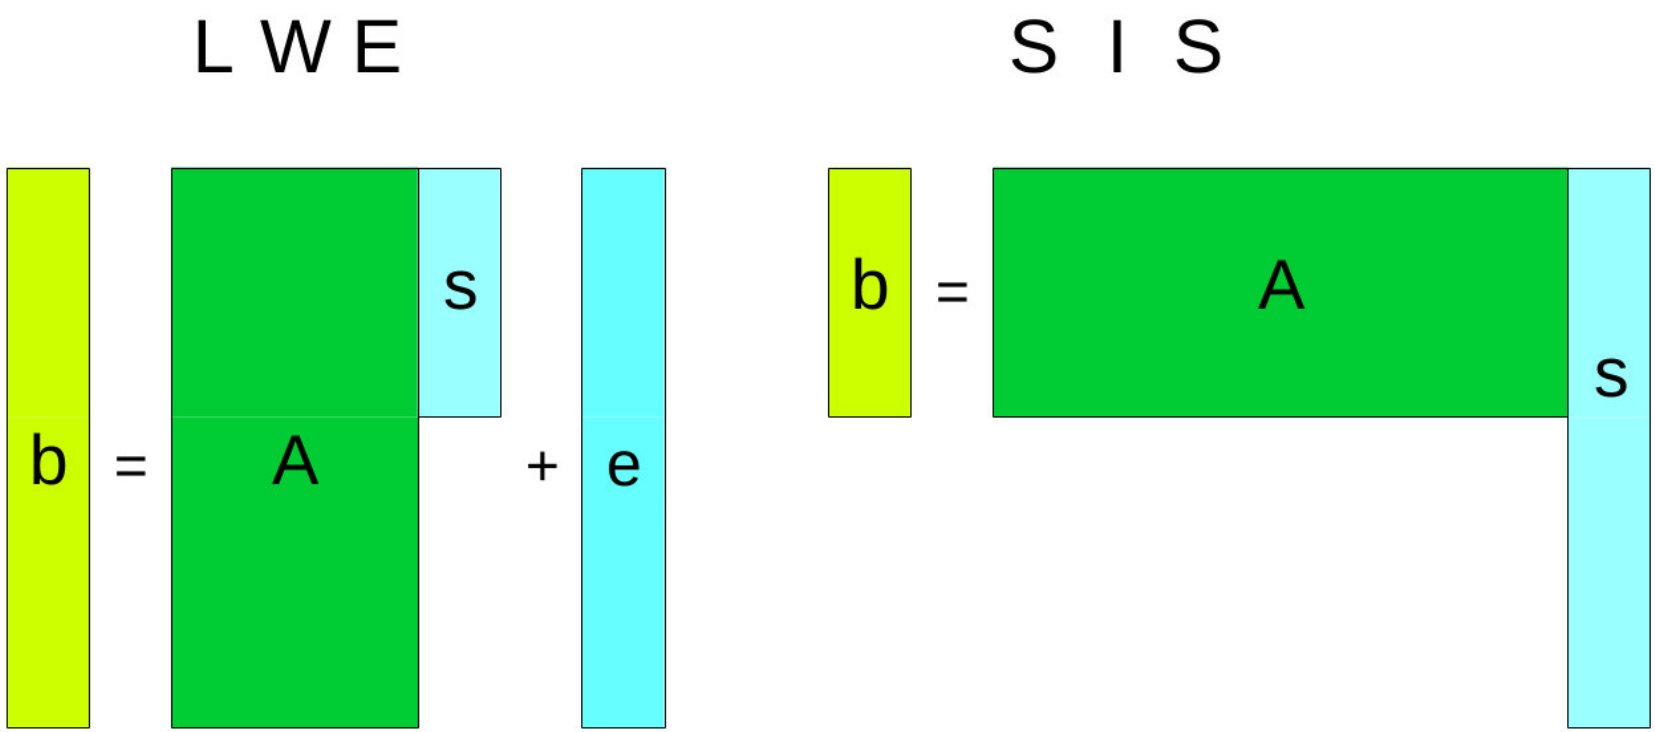

LWE and (I)SIS

- LWE is usually considered in the "low density" case when there is a unique solution $(\mathbf{s}, \mathbf{e})$.
- (I)SIS is usually considered in the "high density" case, when there is more than one solution.
- LWE can be converted to the case where the vector $\mathbf{s}$ is also chosen from $\mathcal{D}$.
- Once $\mathbf{s}$ is a "small vector" one can re-write LWE as ISIS by writing

$$\mathbf{b} = \mathbf{A}\mathbf{s} + \mathbf{e} = (\mathbf{A}|\mathbf{I}_m) \begin{pmatrix} \mathbf{s} \\ \mathbf{e} \end{pmatrix}.$$

- Similarly, ISIS can be converted to LWE.
- Learning with rounding (LWR) is a "compressed" version of LWE.

LWE and (I)SIS

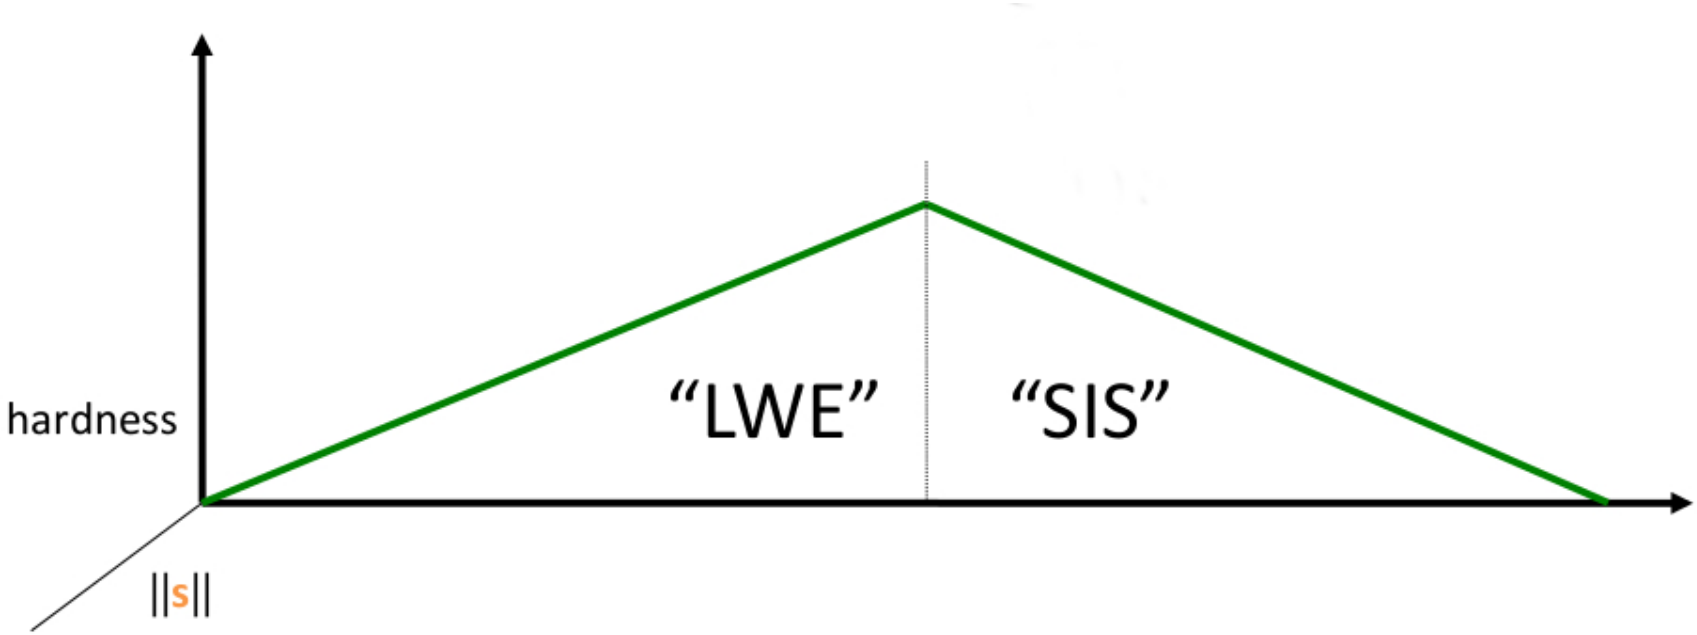

Image by Vadim Lyubashevsky

Lattice algorithms to solve LWE and ISIS

- Both LWE and ISIS can be re-phrased as instances of the closest vector problem (CVP) in the lattice

$$L = \{\mathbf{x} \in \mathbb{Z}^m : \mathbf{x} \equiv A\mathbf{s} \pmod{q}, \mathbf{s} \in \mathbb{Z}^n\}.$$

- Given $\mathbf{b}$ the LWE problem is to compute the lattice point in L closest to $\mathbf{b}$.

Lattice algorithms to solve LWE and ISIS

- LWE and ISIS are cases of CVP, and hence are solved using algorithms for lattice basis reduction.
- A **fundamental challenge** is to predict the running time of lattice attacks for large parameters.
- **Question:** How many people in this room consider themself an expert on floating-point LLL, enumeration algorithms, choice of block size in BKZ, Hermite factors, lattice sieving?
- Reading:
  - N. Gama and P. Q. Nguyen, "Predicting lattice reduction", EUROCRYPT 2008.
  - Y. Chen and P. Q. Nguyen, "BKZ 2.0: Better lattice security estimates", ASIACRYPT 2011.
  - M. Liu and P. Q. Nguyen, "Solving BDD by enumeration: an update", CT-RSA 2013.
  - M. R. Albrecht, R. Player, and S. Scott, "On the concrete hardness of learning with errors", J. Math. Crypt. 2015.

Other algorithms to solve LWE and ISIS

- Arora-Ge: Converts LWE into solving system of multivariate polynomials.
  Errors can't be too small if number of samples is large enough.
- Blum-Kalai-Wasserman: Low weight Gaussian elimination.
  It can be viewed as a variant of the Goldreich-Levin / Kushilevitz-Mansour Fourier learning algorithm.
  In principle it needs a subexponentially large number of samples (i.e., m very large). Lyubashevsky shows (for LPN) the number of samples can be $O(n^{1+\epsilon})$.
- Combinatorial methods (especially for ISIS):
  Schroeppel-Shamir, Camion-Patarin, Wagner, Minder-Sinclair, Howgrave-Graham-Joux, Becker-Coron-Joux.
  See Bai-Galbraith-Li-Sheffield (eprint 2014/593).

What are the important parameters?

- LWE parameters (m, n, q, σ): $\mathbf{A}$ is $m \times n$, $\mathbf{b} = \mathbf{A}\mathbf{s} + \mathbf{e} \pmod{q}$, $\|\mathbf{s}\| \approx \sqrt{n}\sigma$, $\|\mathbf{e}\| \approx \sqrt{m}\sigma$.
- m is not very important.
- The ratio σ/q is very important.
  σ/q very small implies the CVP instance has a target vector very close to a lattice point.
  σ/q too large means the cryptosystem may not function well and there may be many close enough lattice points.
- n is very important as it sets a lower bound on the lattice dimension and the complexity of BKW-type attacks.

Choosing parameters

- I recommend you to read Chris Peikert's blog post "What does GCHQ's cautionary tale mean for lattice cryptography?" `http://web.eecs.umich.edu/~cpeikert/soliloquy.html`

  > *My final conclusion is that worst-case security reductions are really important in lattice cryptography, where there is so much rope to hang oneself with (i.e., flexibility in generating random instances).*

- **Moral:** Having a worst-case reduction gives more security.
- **Question:** How many people in this room consider themself an expert on Regev's reduction to GapSVP, Fourier transform of Gaussian measures, smoothing parameter etc?
- **Challenge**: We need a "beginner's guide" to worst-case reductions and general tools to help determine concrete parameters based on those asymptotic reductions.

Modulus switching (optional)

- Suppose $\mathbf{b} = \mathbf{A}\mathbf{s} + \mathbf{e} + q\mathbf{k}$ with $\mathbf{s}$ short.
- Let q' be another modulus and multiply the equation by q'/q to get

$$\frac{q'}{q}\mathbf{b} = \frac{q'}{q}\mathbf{A}\mathbf{s} + \frac{q'}{q}\mathbf{e} + q'\mathbf{k}.$$

  Writing $\mathbf{b}' = [\frac{q'}{q}\mathbf{b}]$ and $\mathbf{A}' = [\frac{q'}{q}\mathbf{A}]$ then

$$\mathbf{b}' = \mathbf{A}'\mathbf{s} + \mathbf{e}' + q'\mathbf{k}$$

  for some "short" vector $\mathbf{e}' \approx q'\mathbf{e}/q$.
- Hence modulus switching turns LWE modulo q into LWE modulo q'.
- One can do a similar thing for ISIS by converting ISIS to LWE, doing modulus switching, and converting back.
- Note: $|\mathbf{e}'|/q' \approx |\mathbf{e}|/q$, so modulus switching does not change the fundamental parameter σ/q.

Binary LWE

- Let q be an odd prime and $n, m \in \mathbb{N}$.
- Let $\mathbf{s} \in \{0,1\}^n$ or $\{-1,0,1\}^n$ be secret (**column** vector).
- Hardness results: Goldwasser-Kalai-Peikert-Vaikuntanathan, Brakerski-Langlois-Peikert-Regev-Stehlé, Micciancio-Peikert. Need larger n than traditional LWE.
- Challenge: Match theoretical bounds with practical attacks.
- Improved lattice attack (Bai-Galbraith, ACISP 2014).
- Transform LWE instance $\mathbf{b} = \mathbf{A}\mathbf{s} + \mathbf{e} \pmod q$ to $(n+m) \times m$ ISIS instance

$$(\mathbf{A}|\mathbf{I}_m)\begin{pmatrix}\mathbf{s}\\ \mathbf{e}\end{pmatrix} \equiv \mathbf{b} \pmod q.$$

- Re-balance the problem.
- Construct basis $\mathbf{B}$ for the lattice
  $L = \{\mathbf{v} \in \mathbb{Z}^{n+m} : (\mathbf{A} \mid \mathbf{I}_m)\mathbf{v} \equiv 0 \pmod q\}$.
  Rescale $\mathbf{B}$ by multiplying first n rows by σ.

Ring-LWE

- To get more efficient schemes (in terms of key size and speed) we use Ring-LWE.
- The ideas go back to NTRU.
- The difficulty of special cases of these problems is still ongoing research. For instance these papers on eprint from the last year or so contain ideas (not all of them good) worth studying:
  - 2015/106 "Provably weak instances of Ring-LWE"
  - 2015/676 "Quantum Cryptanalysis of NTRU"
  - 2015/971 "Attacks on Search RLWE"
  - 2016/089 "On the Hardness of LWE with Binary Error: Revisiting the Hybrid Lattice-Reduction and Meet-in-the-Middle Attack"
  - 2016/127 "A subfield lattice attack on overstretched NTRU assumptions"
  - 2016/139 "An Algorithm for NTRU Problems and Cryptanalysis of the GGH Multilinear Map without an encoding of zero"
  - 2016/177 "Reduced Memory Meet-in-the-Middle Attack against the NTRU Private Key"

Symmetric encryption from approximate common divisors

(van Dijk, Gentry, Halevi and Vaikuntanathan, 2010)

- Let p be large prime, known to Alice and Bob.
- To encrypt $\mathfrak{m} \in \{0, 1\}$ to Bob, Alice does:
  - Choose $q, e \in \mathbb{Z}$ with $|e| \ll p$ and q large.
  - Compute $\mathsf{c} = pq + 2e + \mathfrak{m}$, and send to Bob.
- To decrypt c Bob does
  - $\mathfrak{m} = [[\mathsf{c}]_p]_2$.
- Here $[\mathsf{c}]_p$ denotes the integer in $(-p/2, p/2]$ congruent modulo p to c.

The approximate common divisor problem (ACD)

- Adversary sees communications of Alice and Bob.
- She sees $c_i = pq_i + (2e_i + \mathfrak{m}_i)$ for $1 \leq i \leq k$.
- Goal: compute p, and hence read all messages.
- Approx Common Divisor problem: Given many samples $pq_i + e_i$ where e_i is "small" relative to p, to compute p.

The approximate common divisor problem (ACD)

- ACD: Find p given $c_i = pq_i + e_i$.
- There are lattice attacks by van Dijk, Gentry, Halevi and Vaikuntanathan. (Also see Ding and Tao.)
- These attacks are powerful (ACD is "easy"), so the integers p and q_i are unbelievably large.
  [Essentially, the lattice attack needs a dimension $> \gamma/\eta$ where $\gamma = \log_2(c_i)$ and $\eta = \log_2(p)$.]
- Also multivariate attacks by Howgrave-Graham, Cohn-Heninger.
- Galbraith-Gebregiyorgis-Murphy "Algorithms for the Approximate Common Divisor Problem" (soon to be on eprint) shows the multivariate attacks are not better than the orthogonal lattice attack.
  We also consider sample amplification and BKW-like approaches to ACD.

Cheon and Stehlé variant

- I recommend you read the paper "Fully Homomorphic Encryption over the Integers Revisited" (EUROCRYPT 2015).
- Cheon and Stehlé give a new variant of ACD that is reduced to LWE.
  (They also give a scale-invariant homomorphic encryption scheme.)
- The new ACD variant is more resistant to lattice attacks than the original version, and so the integers can be a lot smaller.
- **Moral:** Existence of a worst-case security theorem leads to more security.
- A good project is to investigate these parameters, even for simple encryption (rather than FHE).

Adaptive attacks

- It is standard (and realistic) in crypto to consider the setting where an attacker has access to a decryption oracle.
- Recall that decryption of an ACD ciphertext c computes $\mathfrak{m} = [[c]_p]_2$.
- Assume p is odd.
- Attack: Query decryption oracle on even integers $c \approx p/2$. If $c < p/2$ then $[[c]_p]_2 = 0$, while if $p/2 < c < p$ then $[[c]_p]_2 = 1$.
  Hence determine secret key p by binary search.
- Problem: Design an IND-CCA1 variant of this scheme.

Adaptive attacks

- CCA1 attacks are known for almost all homomorphic encryption schemes except Loftus, May, Smart and Vercauteren IND-CCA1 scheme.
- There are similar attacks on many "basic" lattice encryption schemes.
- A related issue is side-channel attacks on lattice crypto.
- Micciancio and Peikert (EUROCRYPT 2012) have IND-CCA1 encryption from LWE, but not homomorphic.

Search-to-decision reductions

- Lattice cryptosystems (such as ACD and Regev) enjoy strong security guarantees: IND-CPA security is reduced to solving worst-case computational problems in lattices.
- A crucial step in Regev's proof is a search-to-decision reduction for learning with errors (LWE).
  See Micciancio-Mol for more details about search-to-decision reductions for LWE.
- This transforms a decisional adversary (as in the IND security game) into a computational adversary.

Dark side of search-to-decision reductions

- Idea: A decryption algorithm provides a decision oracle.
- The search-to-decision reduction allows an adversary to computes a user's private key if they have access to a decryption oracle.
- **Moral:** Existence of a worst-case theorem leads to **less** security.
- For public key encryption: Employ a padding scheme. Decryption returns $\perp$ on most inputs, rendering the decryption oracle useless.
- For homomorphic encryption the problem is more serious. Apart from the proposal by Loftus, May, Smart and Vercauteren, no lattice-based homomorphic encryption scheme has IND-CCA1 security.
- For a solution to this problem in the side-channel world see Fan and Vercauteren, eprint 2012/144.

Search-to-decision reductions for general rings

- For Ring-LWE in general rings there is no general search-to-decision reduction.
- This case is the worst of all possible worlds: we do not have a sufficiently general result to be able to prove a strong IND-CPA theorem, and yet there are enough partial results to ensure that the schemes do not have CCA1 security!
- Paper on this topic is in preparation.

Separating computational assumptions

- It therefore may be necessary to have schemes where the computational assumption to protect the private key is not the same as the computational assumption underlying the semantic security of the ciphertexts.
- Recall the difference between RSA and Rabin:
- The RSA private key is protected by factoring, while the security of ciphertexts is protected using the RSA problem (computing e-th roots).
- The Rabin private key and the security of ciphertexts are both reduced to factoring.
  Hence, a decryption oracle for Rabin (no padding scheme) leads to a total break of the system.
  [Added after feedback: There are many non-mathematical reasons why Rabin was not adopted in practice.]

Lattice Signatures

- There are two paradigms for signatures: full-domain hash ("hash-and-sign") and Fiat-Shamir ("zero knowledge proof").
- Example of full-domain hash: RSA signatures.
- Example of Fiat-Shamir: Schnorr, DSA, ECDSA.
- For lattices:
  - Full-domain hash was done by Gentry-Peikert-Vaikuntanathan.
  - Fiat-Shamir done by Lyubashevsky.
- Both systems are good, but Lyubashevsky's approach seems to have been the most successful for getting efficient schemes.

Lattice Signatures

- Some early lattice signature schemes leaked the private key (see Nguyen-Regev "learning a parallelepiped").
- Very nice approach (Lyubashevsky) is to use rejection sampling to ensure statistical distribution of signatures is independent of the secret.
  This prevents "learning a parallelepiped" attacks.
- **Moral:** Existence of a security theorem leads to more security.
- Open Question 1: Can one can get shorter signatures that are still secure by relaxing the use of rejection sampling?
- Open Question 2: Current Fiat-Shamir signatures are not a ZK proof of knowledge of a solution to LWE or SIS. Can such proofs be done efficiently?

Conclusions

- We need a community of people who are experts in lattice reduction and worst-case reductions.
- We need to understand Ring-LWE.
- Adaptive attacks should be considered (especially for homomorphic encryption).
- Lattice signatures should be made more compact.
- Final comment: post-quantum crypto should be about greater security, not greater efficiency.