From: Dang, Thinh H. (Fed)
To: (b) (6)
Subject: Fw: Rene asked for a summary of the IdealSVP attack
Date: Tuesday, August 24, 2021 3:48:21 PM

From: Apon, Daniel C. (Fed) <daniel.apon@nist.gov>
Sent: Tuesday, August 24, 2021 1:07 PM
To: Moody, Dustin (Fed) <dustin.moody@nist.gov>; Peralta, Rene C. (Fed) <rene.peralta@nist.gov>
Cc: internal-pqc <internal-pqc@nist.gov>; Bassham, Lawrence E. (Fed) <lawrence.bassham@nist.gov>; Liu, Yi-Kai (Fed) <yi-kai.liu@nist.gov>; Cooper, David A. (Fed) <david.cooper@nist.gov>
Subject: Re: Rene asked for a summary of the IdealSVP attack

Gotcha, thanks

From: Moody, Dustin (Fed) <dustin.moody@nist.gov>
Sent: Tuesday, August 24, 2021 1:06 PM
To: Apon, Daniel C. (Fed) <daniel.apon@nist.gov>; Peralta, Rene C. (Fed) <rene.peralta@nist.gov>
Cc: internal-pqc <internal-pqc@nist.gov>; Bassham, Lawrence E. (Fed) <lawrence.bassham@nist.gov>; Liu, Yi-Kai (Fed) <yi-kai.liu@nist.gov>; Cooper, David A. (Fed) <david.cooper@nist.gov>
Subject: Re: Rene asked for a summary of the IdealSVP attack

Daniel,

Thanks for such a great write-up to help us get up to speed.

Small note - the Alice that Leo mentions is not Alice Silverberg. If you look at the bottom of his email, it is actually Alice Pellet-Mary (co-author of ref. 12 you listed).

Dustin

From: Apon, Daniel C. (Fed) <daniel.apon@nist.gov>
Sent: Tuesday, August 24, 2021 12:56 PM
To: Peralta, Rene C. (Fed) <rene.peralta@nist.gov>
Cc: internal-pqc <internal-pqc@nist.gov>; Bassham, Lawrence E. (Fed) <lawrence.bassham@nist.gov>; Liu, Yi-Kai (Fed) <yi-kai.liu@nist.gov>; Cooper, David A. (Fed) <david.cooper@nist.gov>; Moody, Dustin (Fed) <dustin.moody@nist.gov>
Subject: Rene asked for a summary of the IdealSVP attack

Daniel,

For those of us that are unlikely to find a few hours to properly review this,
it would be great if you could give us a summary. Vadim seems to disagree
about the results being directly applicable to PQC candidates. What's up
with that?

Regards, René.

Sure... I'll try to give a brief (lol..) summary of this sub-area, at least so you get a flavor for the technical matter/issues..

Some Historical Background

This line of work (arguably) began in 2014 with a paper written by Campbell/Groves/Shepherd from GCHQ (the "U.K.'s NSA") at an ETSI workshop: https://docbox.etsi.org/workshop/2014/201410_CRYPTO/S07_Systems_and_Attacks/S07_Groves_Annex.pdf
When I say 'arguably,' what I mean is that DJB has argued with that ordering of historical events, and claimed he came up with the idea first in his blog post at https://blog.cr.yp.to/20140213-ideal.html (expressing his idea of NTRU Prime's ring for the first time).
Regardless, if you can also read DJB's contemporary summary in this Google Groups thread https://groups.google.com/g/cryptanalytic-algorithms/c/GdVfp5Kbdb8, and read the later 'blog' post by Chris Peikert here: https://web.eecs.umich.edu/~cpeikert/soliloquy.html

Before moving to more modern developments, it's worth pointing out that cryptographic work using "ideal lattices" was mostly inspired by Craig Gentry's original breakthrough paper in Fully Homomorphic Encryption at STOC 2009: https://dl.acm.org/doi/pdf/10.1145/1536414.1536440
(Or, you can read Craig Gentry's PhD thesis here: https://crypto.stanford.edu/craig/craig-thesis.pdf)

At the point of this first breakthrough of a candidate FHE scheme, many people started exploring cryptography based on ideal lattices (such as the Soliloquy scheme by GCHQ).
While it's worth noting that FHE systems soon moved on be based on different assumptions, like the Learning With Errors problem; e.g. https://eprint.iacr.org/2011/344.pdf from 2011, the original scheme by Gentry was essentially broken (by a quantum key-recovery attack) given the techniques presented in the Campbell/Groves/Shepherd paper and subsequent work. (To be completely accurate, I'd have to think about whether the complete attack has been totally worked out against Gentry's original FHE system, but it certainly should work against a simplified variant of Gentry's original FHE scheme, and it definitely works as-is against Smart/Vercauteren's 2010 follow-up FHE scheme that tried to speed up Gentry's original one: https://www.iacr.org/archive/pkc2010/60560424/60560424.pdf)

It's also worth noting that the polynomial-time quantum step can be done in subexponential time (see http://biasse.myweb.usf.edu/papers/subexp_rel.pdf and https://www.cambridge.org/core/journals/lms-journal-of-computation-and-mathematics/article/subexponential-class-group-and-unit-group-computation-in-large-degree-number-fields/4387ACB036E3358143A563F196E386CB) -- something like exp(n^(1/2)) or exp(n^(2/3)) classically, which is a significantly faster attack classically than we'd typically expect against modern lattice cryptosystems.

Also, worth noting that the original Garg/Gentry/Halevi 2013 multilinear map candidate and its uses in indistinguishability obfuscation can be broken in the same ways..

The Ongoing Line of Work (Key Technical References + a few semi-relevant asides too)

Let me cite a bunch of the important technical references in this line of work, leading up to today. The most important ones to glance at (in my opinion) have a **(*)**.

1. Might as well start at the start.. (Introduction to Cyclotomic Fields, Washington 1997): https://link.springer.com/book/10.1007/978-1-4612-1934-7
2. Quantum Unit/Class Group stuff (STOC 2005): http://www.cse.psu.edu/~sjh26/unitgroup.pdf
3. Quantum Unit/Class Group stuff (ANTS X 2013): https://msp.org/obs/2013/1-1/obs-v1-n1-p17-s.pdf
4. **(*)** As before, Soliloquy (2014): https://docbox.etsi.org/workshop/2014/201410_CRYPTO/S07_Systems_and_Attacks/S07_Groves_Annex.pdf
5. the one most cited-- Quantum Unit Group stuff (STOC 2014): http://personal.psu.edu/kxe8/unitgroup.pdf
6. a heuristic approach to short-PIP by Biasse Song (which I can't find online-- we'd have to ask University of Waterloo for a copy if needed): A note on the quantum attacks against schemes relying on the hardness of finding a short

generator of an ideal in Q(ζ2n). Technical Report 2015-12, The University of Waterloo, 2015
7. the full 2016 paper based on the 2015 technical report above: https://fangsong.info/files/pubs/BS_SODA16.pdf
8. **(*)** a rigorous analysis of Biasse-Song's approach (Cramer/Ducas/Peikert/Regev 2016): https://web.eecs.umich.edu/~cpeikert/pubs/logunit.pdf
9. **(*)** Short Stickelberger Class Relations (Cramer/Ducas/Wesolowski 2016/2017): https://eprint.iacr.org/2016/885
10. **(**)** A very nice, 5-page survey in 2017 by Leo Ducas: http://www.nieuwarchief.nl/serie5/pdf/naw5-2017-18-3-184.pdf
11. The case of Multiquadratics (Bauch/Bernstein/de Balence/Lange/van Vredendaal 2017): https://eprint.iacr.org/2017/404
12. Approx-SVP in Ideal Lattices with Pre-processing (Pellet-Mary/Hanrot/Stehle 2019): https://eprint.iacr.org/2019/215
13. **(*)** analyzing the shortness of vectors by the Ideal-SVP Quantum Alg -- Ducas/Plancon/Wesolowski (DPW 2019): https://eprint.iacr.org/2019/234

The State-of-the-Art Attack

Prior to DJB's talk, the state of the art was the attack by CDW 2016/2017 (whether we get a lot more, a little more, or no more from the S-unit idea is an outstanding question, *smile*).
Let's see if I can describe this at a high-level, without using LaTeX. =)

The attack can be mostly broken down into four steps.

You're given as input an ideal **a** of a cyclotomic number field.

First, you perform a (*quantum*) Class-Group Discrete Logarithm.
Namely, you express the class [**a**] of ideal **a** in a basis **B** = {**p**^sigma | sigma in G} for some prime ideal **p** where G is the Galois group of the number field over the rationals.
To do so, you are running the algorithm of BS16 (see bullet-point 7 or 8 above), which is based on the quantum algorithm for the Hidden Subgroup Problem over R^n from EHKS14 (see bullet-point 4 above).
The output is an element e of \Z[G] such that [**p**^e] = [**a**], where the Galois group ring R = \Z[G] is just the integer linear combinations of elements of the Galois group g.

Second, you (*classically*) solve the Close Principal Multiple problem.
You'll need a *principal* ideal to go to Step 3 next, and **a** is almost certainly not principal. So, you search for a principal ideal **b** = **ac** such that **c** is a "small" ideal.
This is basically done as an application of Stickelberger's Theorem (a structural theorem which is kind of the 'conclusion' of Washington's 1997 textbook on cyclotomic fields).
Assuming a hypothesis about the 'plus-part' of the class group being trivial -- i.e. that so-called (h_m)^+ = 1, which is related to the Generalized Riemann Hypothesis (GRH) -- we obtain a principal ideal **b** of subexponentially bounded norm.

Third, you (*quantumly*) solve the Principal Ideal Problem.
Given **b**, you find some generator h of it. This is from BS16.

Fourth, you (*classically*) solve the Short Generator Problem.
This involves finding a unit (a-ha!) u, such that g = uh (which also generates **b**) has small norm.
In particular, you find a unit u by decoding the log-unit lattice (which is a Closest Vector Problem type of situation, but

where this log-unit lattice is particularly special now, so it's fast).

Finally, you (hopefully!!) have in your hand a short generator g of the ideal **a** of the ideal lattice cryptosystem you were attacking, which is either the secret key of the cryptosystem or something functionally equivalent, so now Bob's your uncle.

Some Initial Discussion Points

- Haven't we just broken cyclotomic lattice cryptosystems? (Un)fortunately, no.
  There are two major obstacles for true cryptanalytic application from the point of view of CDW17:

  - The generator g that you get out from the above attack is only "short" in a very relative sense. In particular, the norm of g is (highly) expected to be around ~exp(n^(1/2)) times the length of the shortest vector.
    That is, you're not getting the shortest vector or even a close approximation, you're getting something that's still subexponentially long.
    Theoretically-secure LWE-type cryptosystems rely on very small approximation factors, involving polynomially-long secret key vectors.
    Practically-secure NIST candidates rely on even shorter (linearly-long) secret key vectors.

  - Even breaking Ideal-SVP with a very small epsilon-approximation factor is not enough to break systems based on Ring-LWE, Ring-SIS, Module-LWE, or Module-SIS.
    While it's known that these Ring/Module problems are at least as hard as Ideal-SVP, the opposite direction is not known.
    Ring-LWE, for example, is better phrased (in terms of SVP problems) as actually an approx-SVP problem on *module* lattices of rank 2.
    "Rank 2 Module-LWE" is, similarly, actually an approx-SVP problem on *module* lattices of rank 3. (In general, rank-k Module-LWE is more accurately related to a rank-(k+1) Approx-ModuleSVP problem for k >= 1, where Ring-LWE is k = 1.)

- How far apart are IdealSVP (the rank-1 "ModuleSVP" case) and rank-2 ModuleSVP? It's still mostly unclear, but it's been a fundamental, known barrier for a long time.
  For recent work along these lines, you can see a CRYPTO 2020 paper https://eprint.iacr.org/2019/1142.pdf
  The salient point is that their reduction from higher rank-k only goes down to rank-2 (exactly), and completely falls apart if you try to push to rank-1.
  (I remember asking Noah S-D a question about this during his talk on the publication, and his response as to why they couldn't get to rank-1 or what technique might help out was essentially "Pfft, I have no idea.")

- Is it the case that DJB's new S-unit method resolves the first barrier above for the case of IdealSVP / rank-1 ModuleSVP? Maybe. He certainly is claiming so. However, it's worth noting that there seems to be contention on the forum as to whether the claim is in fact accurate/true or not. For instance, Leo Ducas points out that he doesn't make that concrete claim until the final slide, and it seems essentially unsupported by some kind of evidence in the rest of the preceding talk. (Apparently Leo says that Alice Silverberg and he were both confused about this claim, which is also odd, since DJB claimed Alice Silverberg as a co-author on the work for the sake of the talk..)

  Some of my questions on the pqc-forum were designed to try to nail down the "how and why" of where this claim came from; I'm hoping Dan responds with some more information about this..

Alright, that's enough typing for one morning. Hope this is helpful, Rene,
--Daniel