From: Smith-Tone, Daniel (Fed)
To: Moody, Dustin (Fed)
Subject: RE: Isogeny Paper Review
Date: Thursday, December 21, 2017 10:42:11 AM

I agree completely, and it means that the paper requires a LOT of work to get in a presentable form. Still, this area suffers from problems that nobody cares about and proofs that are of the form "see problem A." This is an advance from that state, I think.

From: Moody, Dustin (Fed)
Sent: Thursday, December 21, 2017 10:40 AM
To: Smith-Tone, Daniel (Fed) <daniel.smith@nist.gov>
Subject: Re: Isogeny Paper Review

Oh, I've thought about the first main comment where he says there is a one paragraph proof instead of a 3 page proof.

Yes, the one-paragraph proof is correct, and not that hard. The authors should incorporate that into their paper and make the suggested corrections.

From: Smith-Tone, Daniel (Fed)
Sent: Thursday, December 21, 2017 10:27:26 AM
To: Moody, Dustin (Fed)
Subject: RE: Isogeny Paper Review

Thanks, Dustin,

He wasn't the only one to have comments like that. There were others as well that said that it is not too impressive in light of the fact that Lemma 5 is actually easy and that the stronger version highlighted makes all the other results trivial.

I've looked at it, too, and I think that I completely agree with you (which I already said, essentially, in the discussion). I think that the paper probably needs some work but I'm happy with the problems presented. I'll try to champion the paper to get it through if it's not too big a deal to the other reviewers.

Thanks, again.

Cheers,
Daniel

From: Moody, Dustin (Fed)
Sent: Thursday, December 21, 2017 6:38 AM
To: Smith-Tone, Daniel (Fed) <daniel.smith@nist.gov>
Subject: Re: Isogeny Paper Review

My first take on reading it once is that I think I know who the reviewer is. And if it is that person, I'd trust what he says. I didn't dig into the proof details to check if they were right, but it looks like this reviewer did. Similar to how he reviewed this paper, he reviewed the paper published on my dissertation, basically saying (like here: "not too impressive, i already knew this" or "this is elementary".

I'd still say it's a good paper in that it advances the security models/notions for SIDH and will attract wider study. I'll look more into the details over the break.

Hope you and your family have a Merry Christmas!

Dustin

P.S. When's the next time we'll see you?

From: Smith-Tone, Daniel (Fed)
Sent: Wednesday, December 20, 2017 4:23:59 PM
To: Moody, Dustin (Fed)
Subject: Isogeny Paper Review

Hi, Dustin,

Thanks again for reviewing that paper I sent you. I wanted to let you know about some of the discussion raised by another reviewer. I'm posting it below. I'm curious what you think about it.

Cheers!

This paper defines a set of SIDH-related computational problems
(all intractable as of today) and shows randomized polynomial-
time equivalences between them.

In a nutshell, the main results of the paper are:

(Thm 3) Certain computational and decisional problems that
naturally occur in SIDH are equivalent.

(Thm 2) One of those problems is public key validation
(an algorithm that would detect an active attack like
the one described by Galbraith-Petit-Shani-Ti).

(Thm 4) An oracle that computes Alice's secret isogeny from
her public key is equivalent to oracles that
(a) compute the shared secret from the public data and
(b) decide whether a given instance of Alice's public
key and an alleged shared key is consistent with some
fixed instance of Bob's public key and shared secret.

One of the main contributions of the paper is a coherent and
precise formulation of SIDH-related computational problems.

The paper also contains a gentle introduction to basic isogeny-related algorithms for elliptic curves (find a torsion basis, decompose a point w.r.t. a basis, invert action on coprime torsion, from kernels to maps and back, dual isogenies).

Unfortunately, several of the non-obvious reductions contain mistakes; see the comments below.

Lemma 5 including its rather complicated three-page intersection-theoretical proof is beaten by the following simple argument:

Claim: Let f, g: E1 -> E2 be isogenies of degrees <= d.
If f and g agree on > 4d points, then f = g.

Proof: Consider the isogeny f-g. By Silverman, The Arithmetic of Elliptic Curves, Lemma V.1.2 and Corollary III.6.3,
|deg(f-g) - deg f - deg g| <= 2 sqrt(deg f * deg g).
This implies
deg(f-g) <= deg f + deg g + 2 sqrt(deg f * deg g).
<= 4d by assumption.
Since the separable degree of f-g is never greater than its degree, we get
#ker(f-g) = deg_s(f-g) <= deg(f-g) <= 4d
unless f-g = 0.
Therefore, if f and g agree on more than 4d points, they must be equal.

With the better bound (4d instead of 3d^2), Lemma 6 is trivial except for unreasonably large \kappa >= log(n/4l^e): it implies that there is only one isogeny of degree dividing l^e that has the given action on the torsion subgroup, hence finding all of them is exactly the same task as finding (the only) one.
Note we also got rid of the complication that Lemma 5 only applied to isogenies of exact degree d before, whereas we can argue about all (potentially distinct) degrees bounded by d at the same time. (But it looks like a similar generalization should be possible for the original proof in Appendix A.)

Note that the better bound also impacts the discussion about decisional problems on page 13.

Moreover, there are some issues with the original proof of Lemma 6:

- The fact that oracles (1) and (3) compute isogenies of degree *dividing* l^e seems to have been forgotten. In the easy case where you can simply call the oracle, what if there exist isogenies of degrees l^{e-1} and l^e that satisfy the conditions? Oracle (1) will only find one of them, but (3) should return both.

- The claim that two isogenies must be the same if they pass through the same nodes on the isogeny graph is not true: There exist curves with two distinct isogenies of the same degree between them. (Easy example: Consider a prime l > p, then each node has out-degree l+1 which is larger than the number of nodes, thus some codomains must collide.)

- The reduction in Lemma 6 also fails to capture isogenies of degree smaller than l^k, but this can be fixed easily.

Luckily, the improved bound above avoids all these issues, but

the way it does this is by rendering the distinction between
oracles (1) and (3) nonexistent (except for obscure parameters).
That they are essentially the same problem is of course still
a somewhat interesting result, but the discussion above shows
it is not very hard to see.

In the second part of the proof of Theorem 2, the queries to
oracle (2)^ are nonsensical: The query
(E', [deg phi] P, [deg phi] Q, R, S, id)
contains some information (deg phi) dependent on an object the
oracle is supposed to reason about. I assume it was intended
to iterate over all r between 1 and k-e and query
(E', [l^r]P, [l^r]Q, R, S, id)
for each r, but then oracle (2)^ (and thus "we" falsely) might
respond "yes" because there may be an isogeny whose degree is
different from l^r (in fact, it could even be larger than
l^{k-e}!) with the required properties.
(Moreover, the requirement "of degree dividing l^{e-k}" was
silently dropped in "The question is the same as asking [...]".)

The "dual" problems that follow Theorem 1 and their equivalence
to problems (1)-(3) are immediate consequences to being capable
of computing dual isogenies. I do not think they add any value
to the paper; they are literally obtained by adding one \hat on
top of everything and flipping some arrows. The results that
make use of those dual problems could easily be reformulated to
use (1)-(3) instead.

Footnote 5 in the proof of Theorem 4 seems a bit clumsy. Why
not just *declare* that the CSIDH oracle will return a random
value if it is given invalid input? This is clearly the worst
failure mode for someone using the oracle: in case it actually
"crashes" in some sense instead, one could still simply assume
some random return value and proceed as in the "invalid return
value" case. That way the given practical argument could be
transformed into a precise theoretical statement.

In the introduction, the authors claim to show an unconditional
equivalence between their six problems and "the decisional
variant considered by Thormaker [...]". However, Thormaker's
version of the decisional oracle does not verify the isogeny's
action on a torsion subgroup, so the oracle considered in this
paper is actually more specific. Therefore what was proved is
not really the same result unconditionally; it is a similar
reduction under different assumptions (that permit applying
Lemma 5, which Thormaker could not have used in his setting).

I like Section 2 ("preliminaries on isogeny problems") as a
well-written and clear introduction to basic helper algorithms
used in computations with elliptic curve torsion subgroups.
Introduction and conclusion are also well-written.

==========

Conclusion:

The paper introduces some new (and some known) computational problems in isogeny-based cryptography and reductions between them. While the authors succeeded in providing sufficient motivation and precise statements of the discussed problems, several of the nontrivial reductions unfortunately contain mistakes or become obsolete by using the improved bound for Lemma 5. The remaining results are easy to prove and little surprising. Overall, although the paper is a nice read, it unfortunately lacks deeper insights.

==========

Remarks and typos:

many places: "Name[42]" lacks a space in many citations.

p. 3: Maybe define things for "an isogeny" instead of "isogenies" (plural -> singular) for clearer language.

p. 3: "seperable" -> "separable"

p. 3: The sentence on VÃ©lu would benefit from the addition of "in time polynomial in the degree".

p. 4: Not every supersingular elliptic curve defined over F_{p^2} for has that group of F_{p^2}-rational points: for instance, if E *is* a curve as claimed, then its quadratic twist has 4p more or fewer F_{p^2}-rational points. What is true is that every *isomorphism class* of supersingular elliptic curves over F_{p^2} contains representants with (p+1)^2 and (p-1)^2 points over F_{p^2}. Of course, there is a long-standing tradition of sweeping this detail under the rug in the published SIDH literature.

p. 4: I would omit "for i=1 and i=2" in the second paragraph.

p. 4: Lemma 1 talks about a Z-basis, but E[l^e] is not a free Z-module. All occurrences of this should be replaced by "Z/l^e-basis".

p. 5: In remark 1, insert "optimized variants of" before "the above computations". I doubt that the algorithm described in Lemma 1 is that fast.

p. 5: Again, "Z-basis" -> "Z/l^e-basis".

p. 6: To what base are the logarithms? Of course it doesn't matter asymptotically, but concrete values of \kappa are considered, so we need that. If I take Costello-Longa-Naehrig's parameters and try to get \kappa < 3 from that, I would have to assume a base of about 5, which is not common at all, so I suspect some arith-metic mistake there.

p. 6: Who is "one" in "one expects there to be enough primes of the right form [...]"?

p. 8: Switch the footnote 4 and the period at the end of line 2.

p.10: In the proof of Theorem 2, "k = deg psi'" -> "k = log_l(deg psi')".

p.11: Maybe add brackets to E/A in E/A[l_2^e_2] (same for B).

p.12: "A and B Isogeny Problems" -> "A- and B-Isogeny Problems"

p.12: At the end of the third paragraph, it is claimed that
"the theorems of Section 3 could not have been proven
in [Petit's] setting". I doubt this; couldn't you
just treat each prime power in N_1 individually?

p.13: I don't really get what is meant by "as l increases the
number of cyclic subgroups of E[l^e] [...] must tend to
0 [...]". This is not true if you also decrease e, and
more importantly I don't see how the number of private
keys relates to the requirement log(l^e) ~ log(n/l^e).
Rather, I feel like the chosen formalization of these
problems simply does not permit talking about any other
l^e since n, l, e were "hardcoded" globally instead of
given as inputs to each argument.

p.14: "observatoin" -> "observation"

p.16: It is claimed that "improving the natural attack strate-
gies [...] is just as difficult as solving the isogeny
problems themselves": It should be made clearer that
this only applies to "perfect" improvements (which is
precisely why they immediately give solutions to the
underlying conjecturally hard problems!). These
reductions do not at all preclude the existence of
improvements to isogeny path search algorithms which
do not completely break SIDH (in polynomial time).

References:
3. Capitalize "Diffie-Hellman" (and maybe some more words).
8. Capitalize "protocols".
12. The editor is the only name of the form "Last, First".
17. has a random space in the word "isogen ies".
18. has a random space in the word "computati on".
19. Capitalize "Diffie-Hellman" (and maybe some more words).

Appendix A:
p.19: It is claimed that "\eta(p_1,p_2) /= 0" since p is a point
on E. However, \eta is the denominator of the rational
maps defining an isogeny with nontrivial kernel, so it
actually has just the points in the kernel as its roots!
Although the isogeny was defined everywhere, the images
of the kernel points were "lost" by passing to an affine
patch of the codomain.