

The Case for SIKE

A Decade of the Supersingular Isogeny Problem

<https://eprint.iacr.org/2021/543.pdf>

"SIKE is a fantastic scheme, but its computation is by far the most expensive, and the problem is relatively new. Who knows here?"

- Daniel Apon (NIST)

Craig Costello*

Microsoft®

Research

*Disclaimer: opinions are my own and indicate a strong partisan bias.

10 topics: each has 1 quote + 1 slide + 1 sentence summary


0. Quick prelims
1. A decade unscathed
2. The rise and rise of classical hardness
3. Quantum computers don't really help
4. Concrete cryptanalytic clarity
5. Side-channel security
6. The efficiency drawback
7. Happy hybrids
8. Other avenues of attack
9. Elegance
10. The \$IKE challenges

most references won't be in talk
(see <https://eprint.iacr.org/2021/543.pdf>)

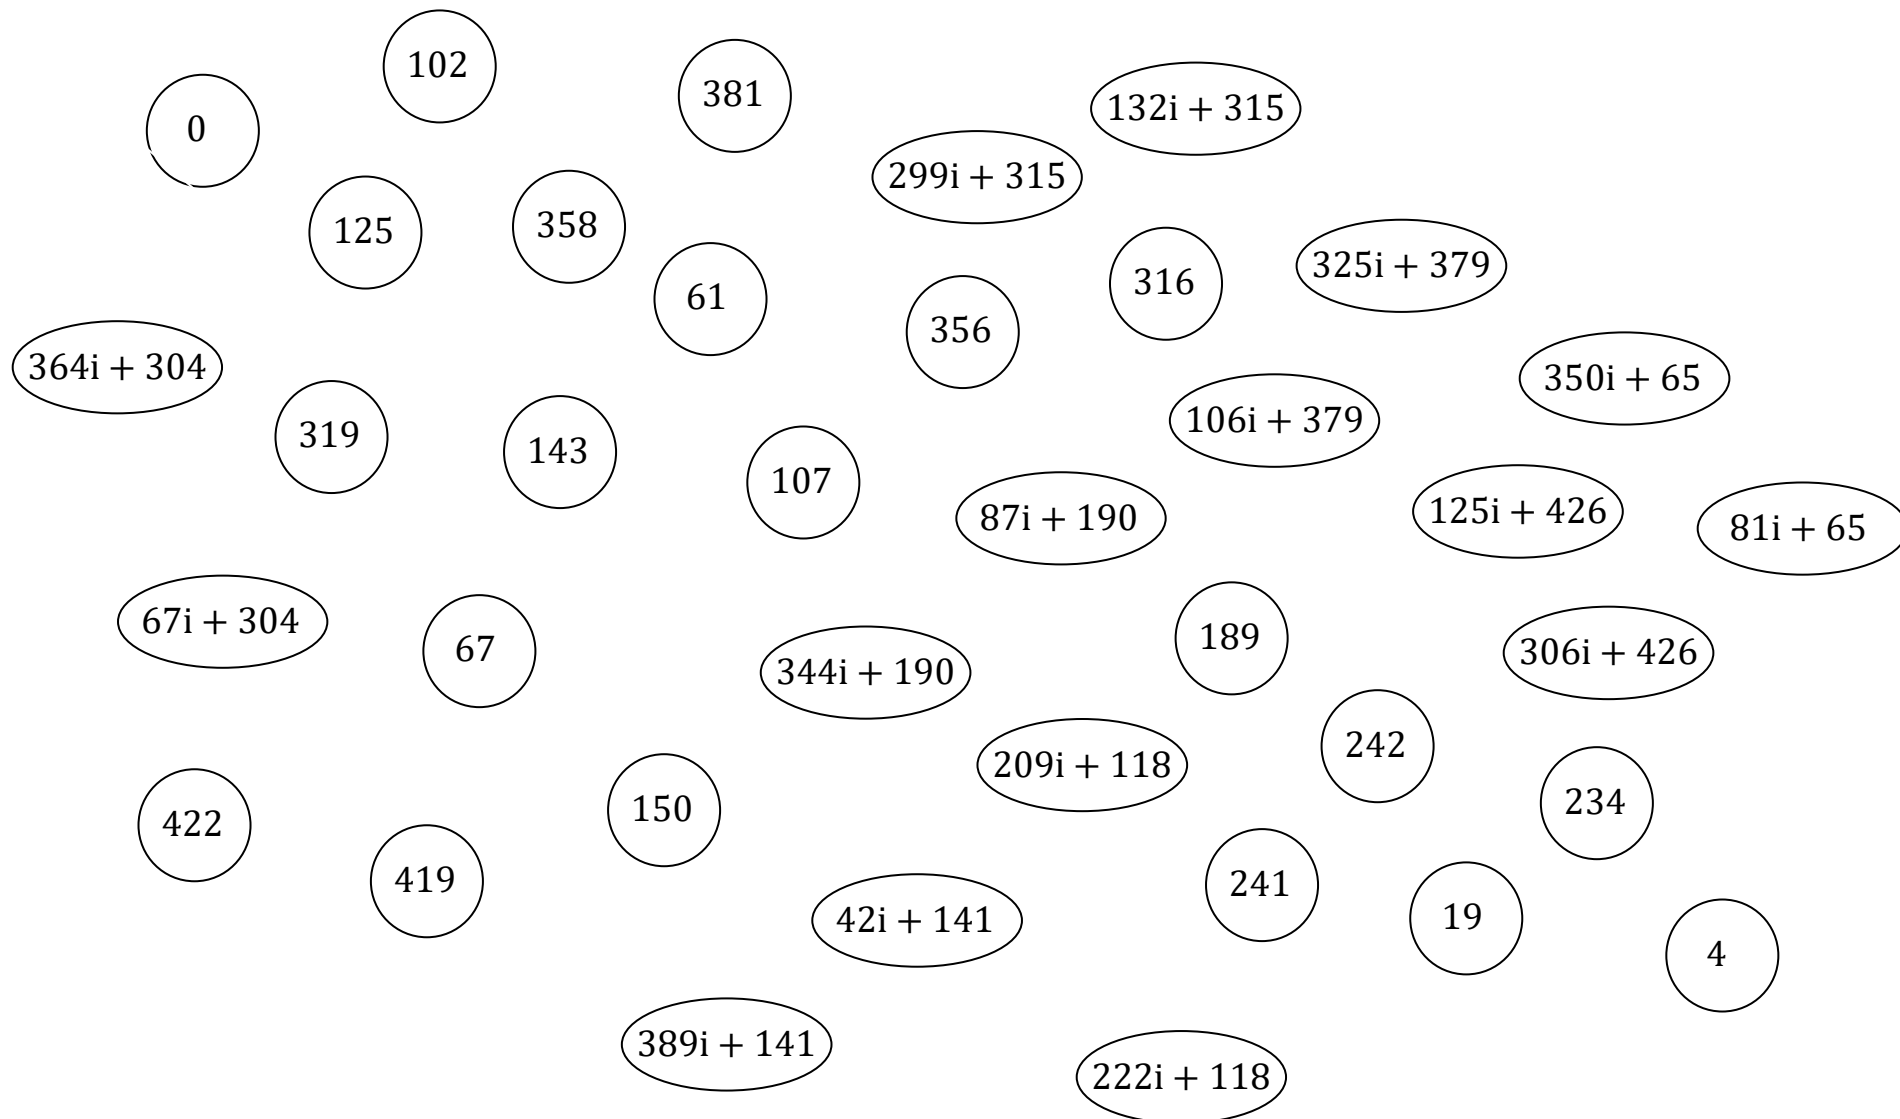
0

Quick prelims

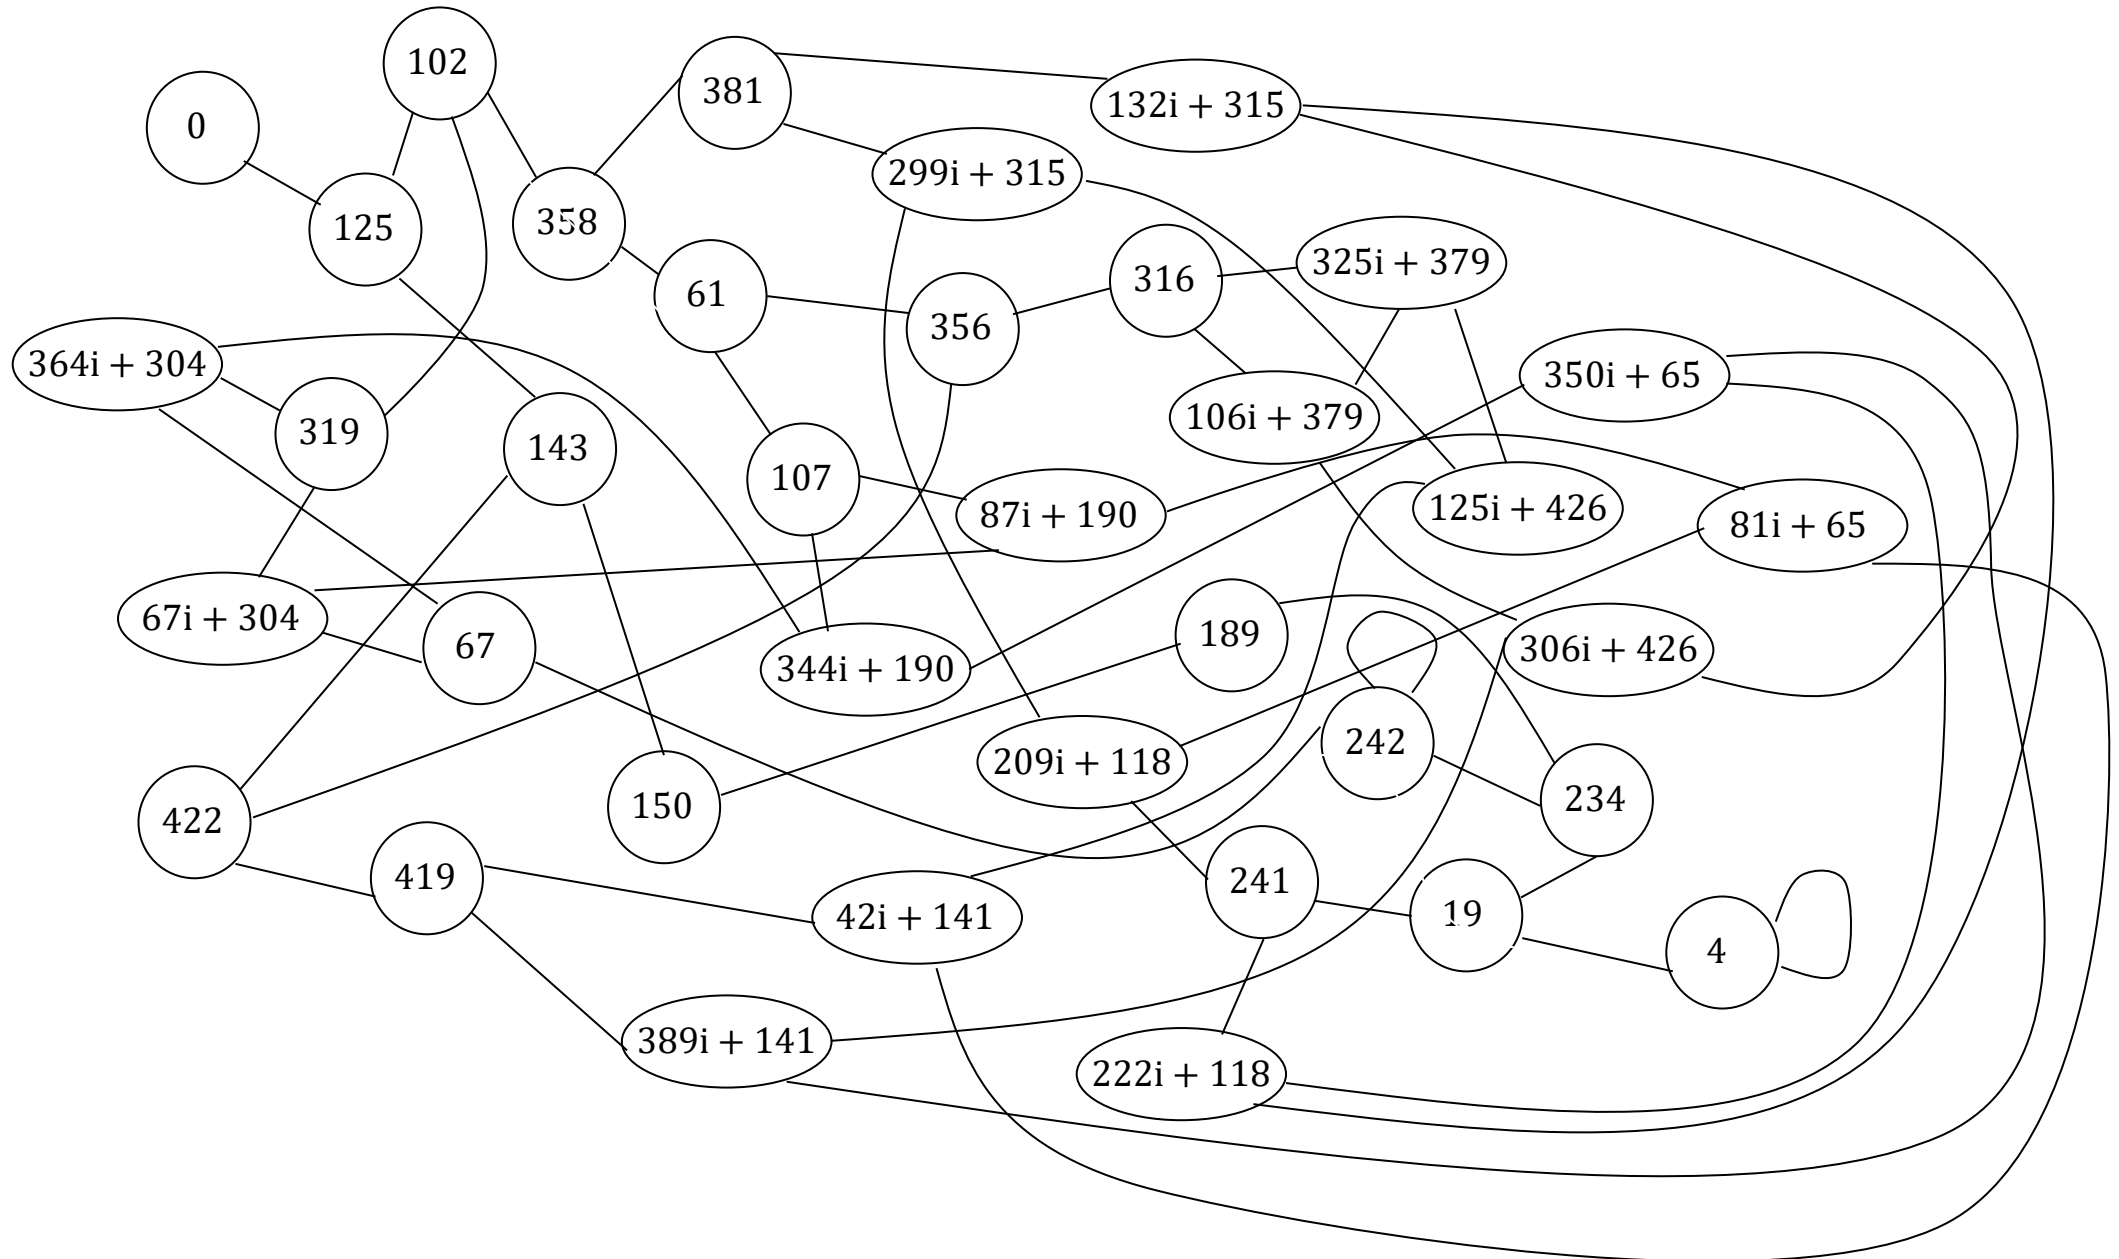
SIDH/SIKE






e.g. supersingular isogeny graph – the nodes



$p := 431$: there are 37 supersingular j 's (all over $\mathbb{F}_{p^2} := \mathbb{F}_p(i), i^2 + 1 = 0$)







1

A decade unscathed

"Which post-quantum submissions (1) haven't suffered security losses since the #NISTPQC competition began and (2) are among the 26 submissions in round 2 (which is ending soon)? I think there are exactly 3: SIKE (which scares me for being too new), Classic McEliece, and SPHINCS+."

- Daniel J. Bernstein

A lot *can* happen in 10 years...

- DLP

- Diffie-Hellman'76: "taking logs mod $q \approx 2^{200}$ requires approximately 2^{100} operations"
- Adleman'79: index calculus runs in subexponential time, completely breaks $q \approx 2^{200}$

- Factoring

- RSA'78: " $N \approx 2^{266}$ moderate security [...], $N \approx 2^{664}$ margin of security against future developments"
- 1978-1988: Pomerance quadratic sieve, Lenstra's ECM, Pollard's NFS (used to factor RSA-768 in 2010)

- McEliece

- McEliece'78: proposes $n = 1024$ Goppa codes targeting 2^{64} security
- Lee-Brickell'88, Leon'88, van Tilburg'88, Stern'88. E.g., Lee-Brickell 2^{11} improvement for $n = 1024$...

- NTRU

- Hoffstein-Pipher-Silverman'96: NTRU presented at CRYPTO rump session
- Coppersmith-Shamir'97: improved lattice attacks at CRYPTO rump session, forces params to increase

- ECDLP

- Miller'85: specifies $E: y^2 = x^3 - ax$ with $p \equiv 3 \pmod{4}$, "may be prudent to avoid curves with CM"
- MOV'93/Frey-Rück'94: attacks use E being supersingular (CM curves are fine)



A decade unscathed

CSSI problem

Jao-De Feo '11: propose CSSI, specify a range of parameters with known (generic) attack complexities

2011 – present:



Proposed metric:

not time that problem has been around...



but time that attack complexities have not decreased



Summary: The SSI problem may only be 10 years old, but it's off to a better start than any other public key problem I can think of.

2

The rise and rise of classical hardness

"So, on average, the 2^{80} storage device will be accessed 2^{48} times during each unit of time. The cost of these accesses will certainly dominate the computational costs. Thus, our security estimates, which ignore communication costs, should be regarded as being conservative."

- Gora Adj, Daniel Cervantes-Vázquez, Jesús-Javier Chi-Domínguez,
Alfred Menezes and Francisco Rodríguez-Henríquez

The rise and rise of classical hardness...

- **2011:** Jao-De Feo
 - Generic, meet-in-the-middle, claw-finding algorithms
 - Classical: $O(p^{1/4})$ time and space
- **2017:** Adj, Cervantes-Vázquez, Chi-Domínguez, Menezes, Rodríguez-Henríquez
 - $O(p^{1/4})$ memory does not make sense / exist
 - Fix $w = 2^{80}$ as upper bound on memory, then analyze runtime
 - van Oorschot – Wiener (vOW) is now the best attack
 - Runs in time $\frac{2.5}{m} \cdot \left(\frac{p^{3/8}}{\sqrt{w}}\right) \cdot t$
 - SIKE parameters decrease in Round 2 to match NIST security levels more closely
- **2021:** Longa, Wang, Szefer
 - Real-world budget-based cost model shows current parameters still offer a wide security margin
 - SIKE parameters could be safely decreased again



Summary: The SIKE parameters have recently decreased because the SSI problem is harder than was initially thought, and this could well happen again.

3

Quantum computers don't really help

"An adversary with enough quantum memory to run Tani's algorithm with the query-optimal parameters could break SIKE faster by using the classical control hardware to run van Oorschot–Wiener."

- Samuel Jaques and John M. Schanck

instance	best quantum			best classical		
	2^{96}	2^{64}	2^{40}	2^{96}	2^{64}	2^{40}
SIKEp434	124	147	178	117	133	135
SIKEp503	134	179	234	142	158	160
SIKEp610	181	189	307	183	199	201
SIKEp751	219	274	345	235	251	253

Best known classical and quantum attack complexities (base 2 logarithms) for the four SIKE instances

Classical: 2^{96} is far greater than the current memory of the planet

Quantum: 2^{96} approximate number of gates that atomic scale qubits with speed of light propagation times could perform in a millennium

Summary: I think it would be poetic to have a post-quantum standard where quantum computers don't really help to break it.

4

Concrete cryptanalytic clarity

"NIST believes it is important to understand how exactly this CoreSVP security translates into "true" bit security strength for KYBER."

- NIST

Concrete cryptanalytic clarity

- Constructively, lattice-based schemes much easier to understand than SIDH/SIKE
- Cryptanalytically, the opposite is true...


$\frac{2.5}{m} \cdot \left(\frac{p^{3/8}}{\sqrt{w}} \right) \cdot t$	
---	---

Figure: the concrete complexity of the best-known isogeny (left) and lattice (right) attacks

Summary: While understanding lattice-based cryptosystems is typically much easier than understanding the SIKE cryptosystem, understanding the state-of-the-art in attacking the SSI problem is much easier than understanding the state-of-the-art in attacking the LWE problem.

5

Side-channel security

"It is a very well-understood operation. We know how to attack it; we also know how to defend against it."

- David Jao

Two phases of isogeny computation

- **Phase 1:** elliptic curve scalar multiplication ($\leq 20\%$ total runtime)
 - Compute secret kernel element $S = P + [k]Q$, $P, Q \in E$ (both public) and $k \in \mathbb{Z}$ (secret)
 - At a high level, operations depend on the secret
 - Constant-time techniques needed to avoid timing attacks
 - Fortunately, we have **20+ years of post-Kocher research** into this
- **Phase 2:** isogeny evaluation ($\geq 80\%$ total runtime)
 - Compute secret isogeny with kernel $\phi : E \rightarrow E/\langle S \rangle$, where E and $E/\langle S \rangle$ (both public) and S secret
 - Operations are now public, it's the data (field elts.) that are secret
 - Need to protect field elements from timing/power analysis
 - Fortunately, we have **20+ years of post-Kocher research** into this

Summary: Decades of ECC side-channel analysis have given SIKE a good head start in the knowledge and implementation of side-channel protection and protecting it in most scenarios would be relatively cheap.

6

The efficiency drawback

"The main drawback to SIKE is that its performance is roughly an order of magnitude worse than many of its competitors. Much work has been done to optimize implementations, including the compressed-key version, and it is hoped that such optimizations continue."

- NIST

Performance evolution of curve-based cryptos

- Elliptic curve cryptography

- **2000**: NIST **standardizes** in FIPS 186-2
- **2001**: \approx **2 million cycles** (Pentium II) by Brown, Hankerson, Lopez-Hernandez and Menezes
- **2006**: \approx **800k cycles** (Pentium III) by Bernstein
- **2016**: \approx **60k cycles** (Intel Core i7) by Longa
-

- Pairing-based cryptography

- **2000**: Joux reports **one second**
- **2010**: **less than a millisecond** (Intel Core i7) by Beuchat, Gonzalez-Diaz, Mitsunari, Okamoto, Rodriguez-Henriquez, Teruya
- **2010**: **less than half a millisecond** (Intel Core i7) by Aranha, Karabina, Longa, Gebotys, Lopez-Hernandez
-

- Isogeny-based cryptography

- **2011**: Jao – De Feo report **758 milliseconds** at 128-bit security level
- **2021**: Longa (a few hours ago) reports **5.9 milliseconds** for SIKEp434 (encaps+decaps)

Summary: When it comes to curve-based cryptography versus its counterparts, performance disparity tends only to be temporary...

... but keysize disparity lasts forever!!!

7

Happy hybrids

"The submission package shall include a statement that lists and describes the advantages and limitations of the cryptosystem. [...] This could include, for example, the suitability of the algorithm for use in hybrid schemes...."

- NIST

A happy hybrid

- Modern ECC (e.g. curve25519 and Goldilocks) uses curves obtained deterministically over \mathbb{F}_p - see RFC7748 "Elliptic curves for security"
- SIKE already has large prime fields and modern ECC arithmetic, e.g. SIKEp434 uses 434-bit \mathbb{F}_p and Montgomery arithmetic. Use this for ECC hybrid!
- Including secure ECC alongside SIKE adds small overhead in (1) runtime, (2) public keys, and (3) code complexity.
- All the "fast" proposals slow down significantly in hybrid mode...

Summary: SIKE is the only NIST candidate that has a nice hybrid.

8

Other avenues of attack

"We can therefore conclude that at least heuristically, it seems extremely unlikely that Petit's attack can possibly apply to the actual, balanced SIDH parameters."

- Chloe Martindale and Lorenz Panny

Other avenues of attack

- SIDH public keys encode an elliptic curve E and two torsion points $P, Q \in E$
- 2016: Galbraith-Petit-Shani-Ti give active attack that exploits torsion points
 - Adversary can modify a valid public key (E, P, Q) into a malicious (E, P, \tilde{Q}) and send to Alice
 - Alice has no way of knowing this is malicious, proceeds, and adversary learns a bit of her secret s_A
 - Adversary does this $\lceil \log s_A \rceil$ times and learns Alice's full secret
- GPST attack is the reason SIDH became SIKE, the actively secure incarnation
- **NOTE:** every other NIST candidate scheme has the same problem, they all need to be modified for active security!
- Petit gives passive torsion point attacks to variant isogeny-based schemes, but they do not apply to SIDH/SIKE

Summary: A lot of great cryptanalytic work has improved our knowledge of the isogeny problem landscape in the last decade, but the best attack against the SSI problem and against SIKE remain the generic collision-finding attacks.


9

Elegance

"It is my intent to show that elliptic curves have a rich enough arithmetic structure so that they will provide a fertile ground for planting the seeds of cryptography."

- Victor S. Miller

Elegance



=



noise, rounding, errors,
decryption failure,
decoding, etc

=



Summary: In terms of the remaining key encapsulation candidates, a search over each of the Round 3 specification documents for the number of instances of the string "failure" finds (in alphabetical order) BIKE = 27, FrodoKEM = 29, HQC = 16, Kyber = 61, McEliece = 18, NTRU = 6, NTRU Prime = 42, Saber = 24, and finally, SIKE = 0.

10

The \$IKE challenges

"When it comes to betting on yourself [...] you're a chicken-livered coward if you hesitate."

- Bertie Charles Forbes

$$\frac{2.5}{m} \cdot \left(\frac{p^{3/8}}{\sqrt{w}} \right) \cdot t$$

best classical

instance	2^{64}	2^{40}
\$IKEp182	45	50
\$IKEp217	54	63
SIKEp434	133	135
SIKEp503	158	160
SIKEp610	199	201
SIKEp751	251	253

~~\$5,000 USD~~

\$50,000 USD

Congratulations
Aleksei Udovenko and Guiseppe Vitto!
Solved August 28, 2021

Your name could be here!

Summary: If you are not yet convinced that breaking SIKE is hard, then perhaps some prize money will encourage you to get *cracking*.

<https://github.com/microsoft/SIKE-challenges/>