National Institute of Standards and Technology
U.S. Department of Commerce

NIST's PQC Standardization

Suggested Avenues for Lattice-Based Research

Jacob Alperin-Sheriff

NIST

May 22, 2018

CSRC

KEMs–Key and Ciphertext Sizes

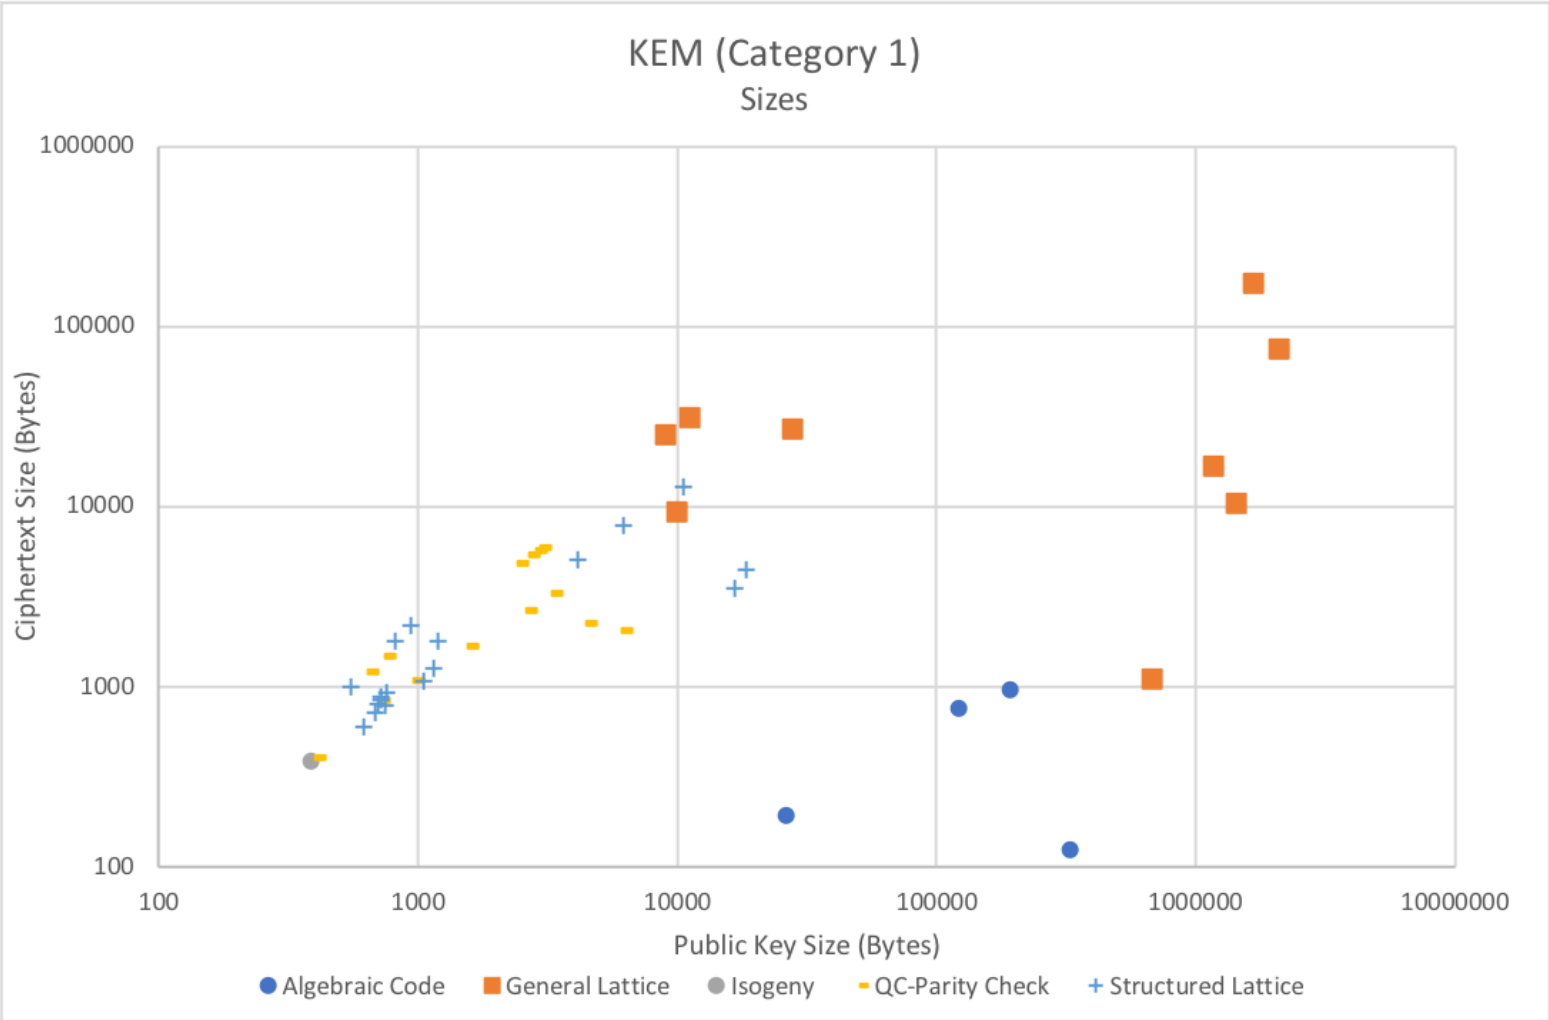

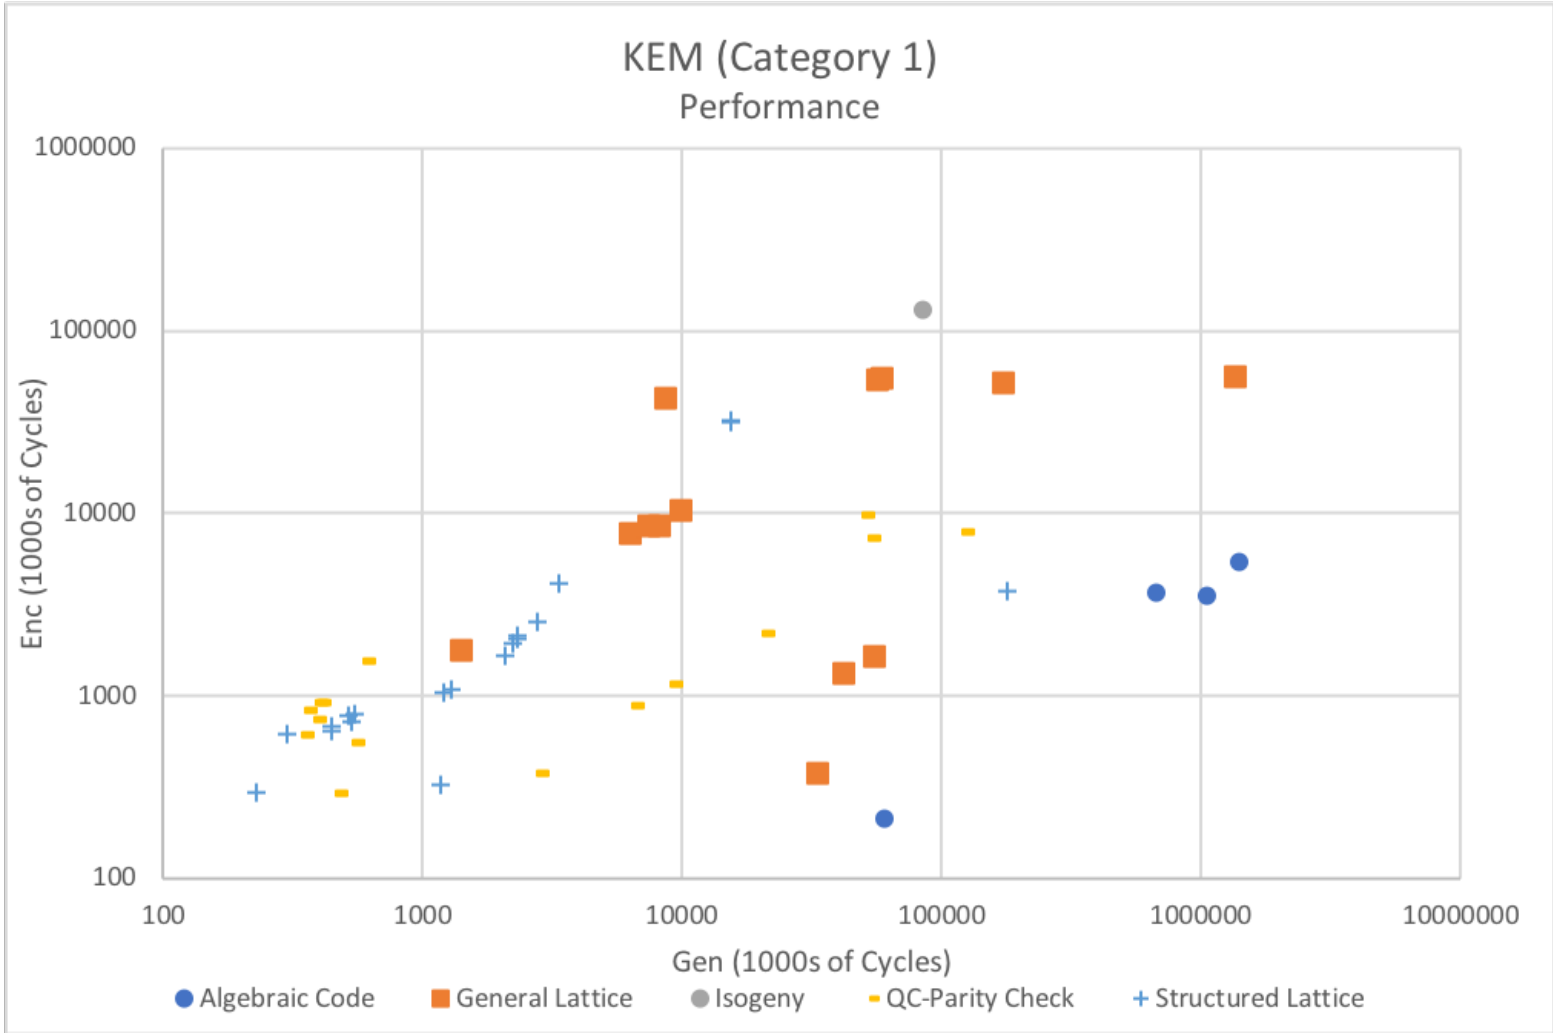
KEM (Category 1)
Performance
1000000
100000
10000
1000
100
Enc (1000s of Cycles)
100
1000
10000
100000
1000000
10000000
Gen (1000s of Cycles)
Algebraic Code
General Lattice
Isogeny
QC-Parity Check
Structured Lattice

KEMs–Performance by Sizes

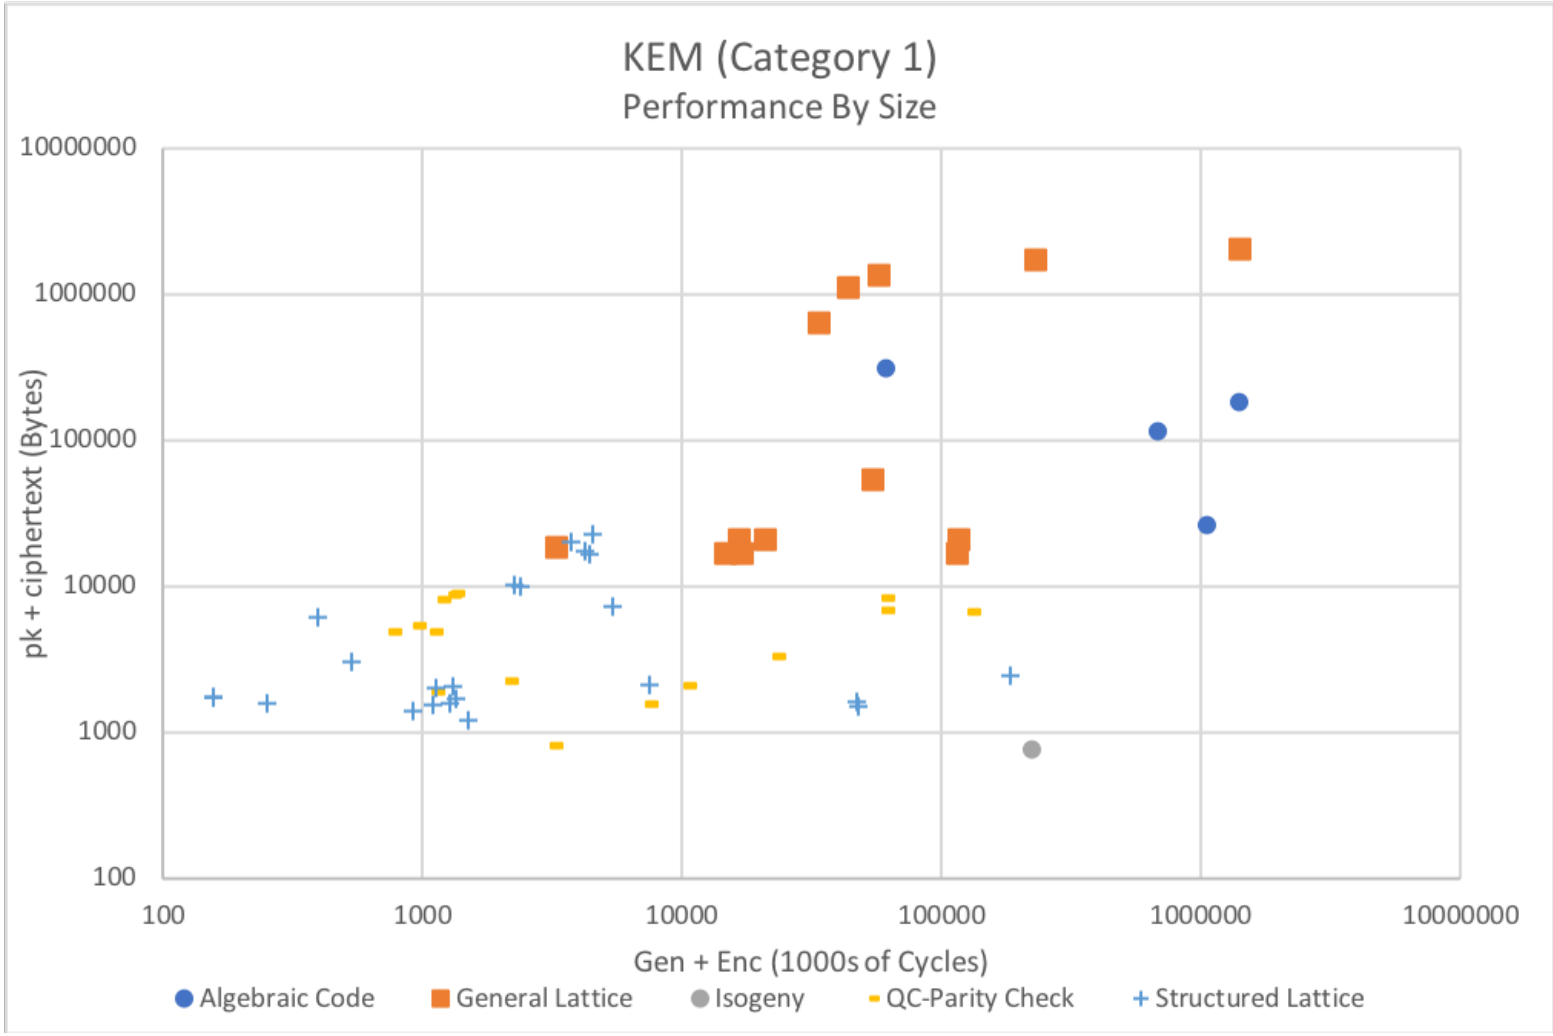

Signature–Sizes

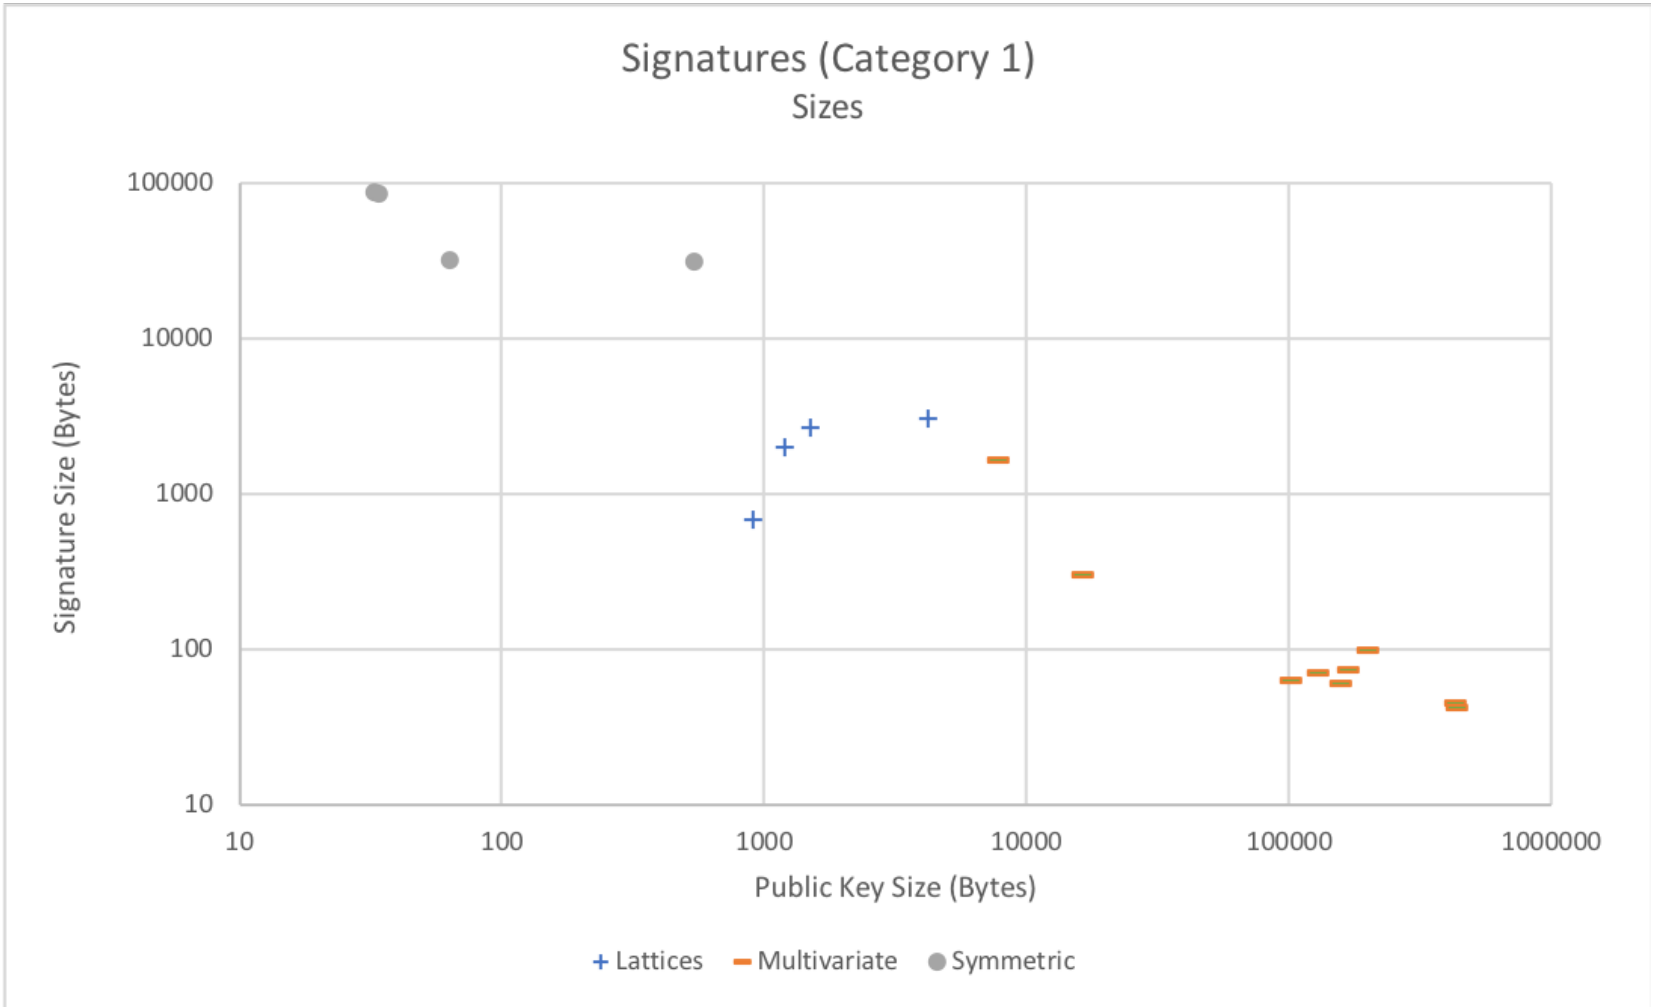

Signature–Performance

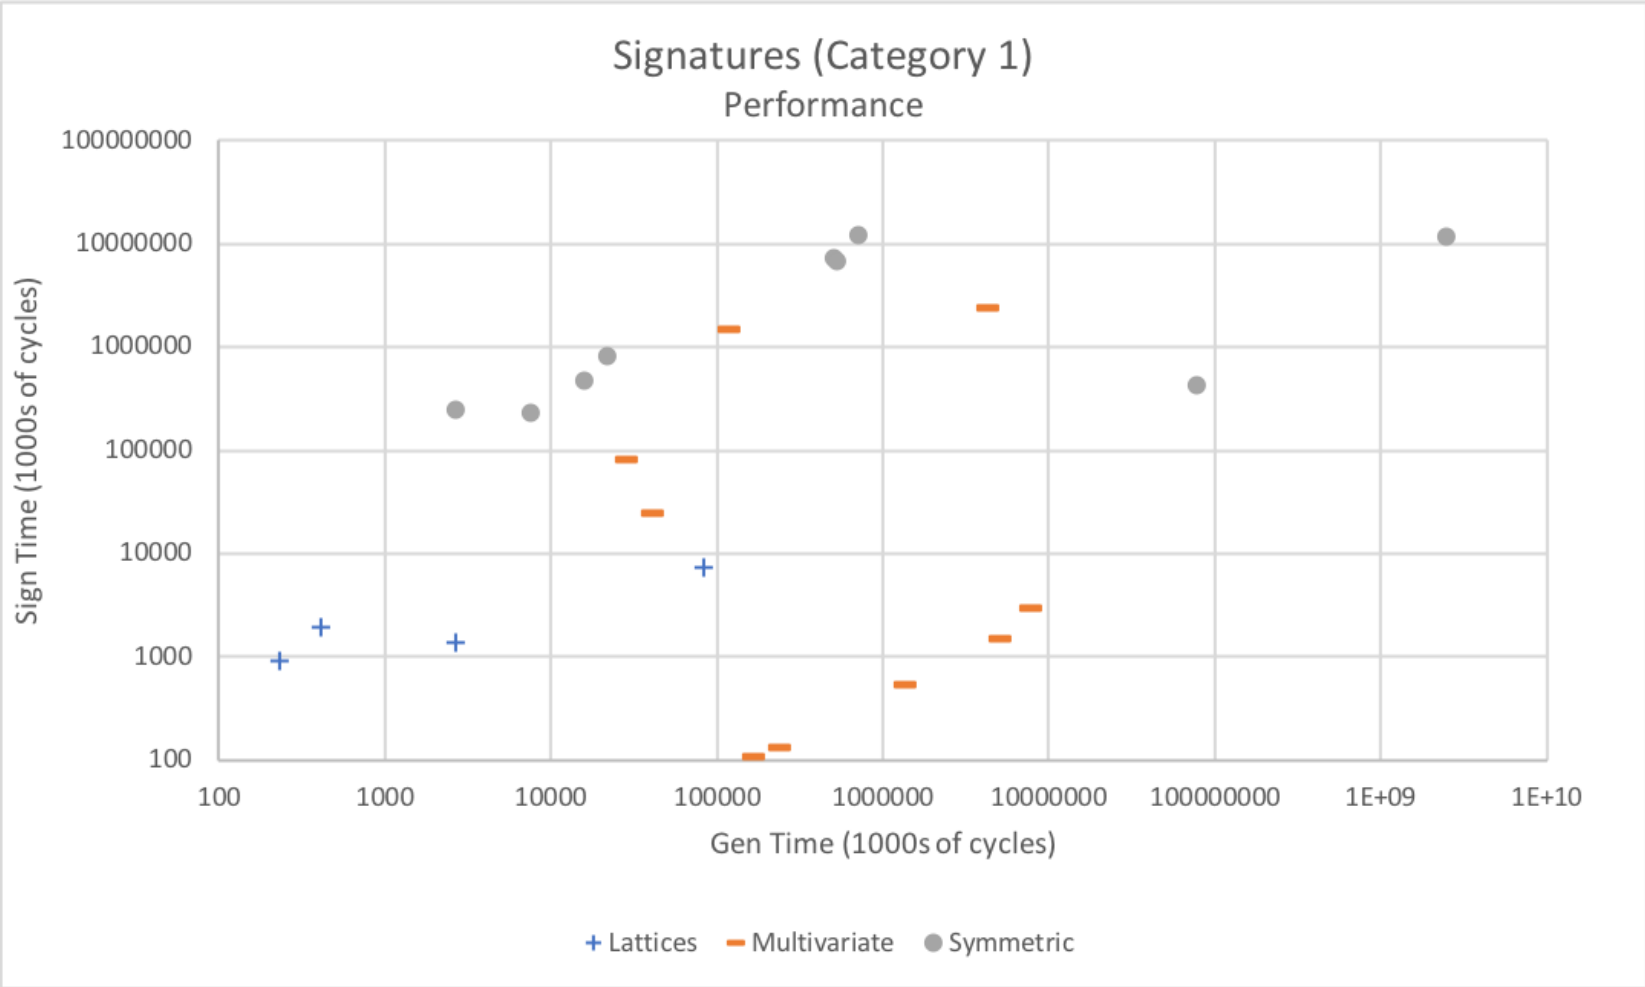

Signature–Performance by Sizes

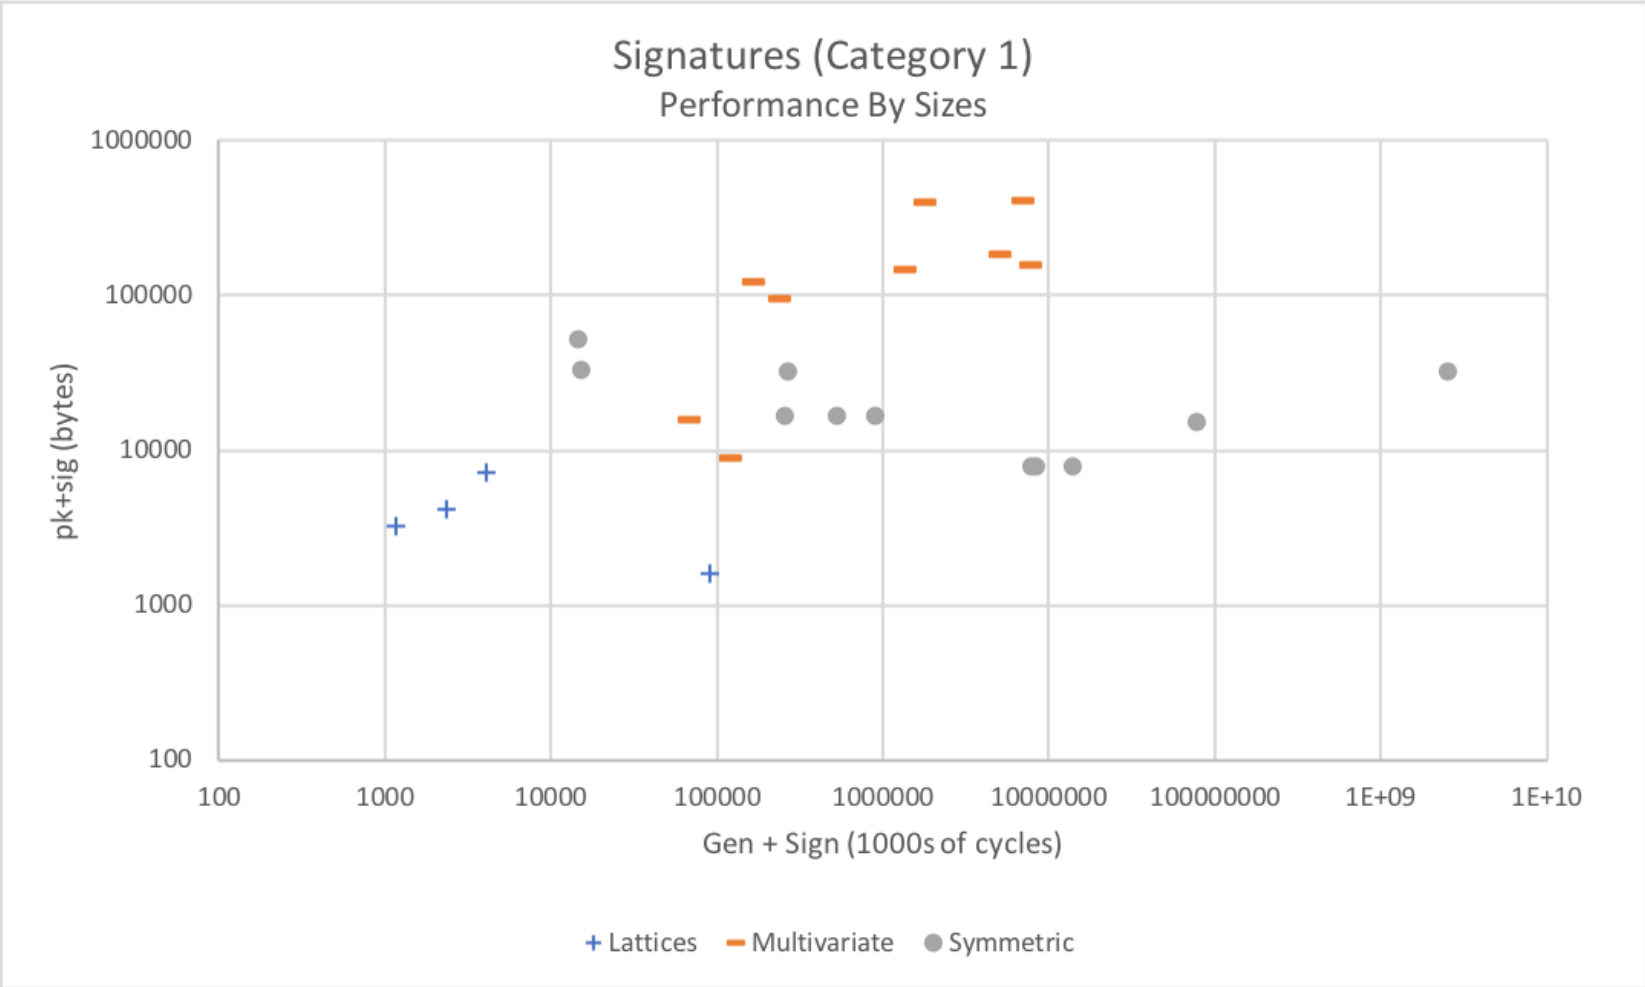

Breaks by Problem Category

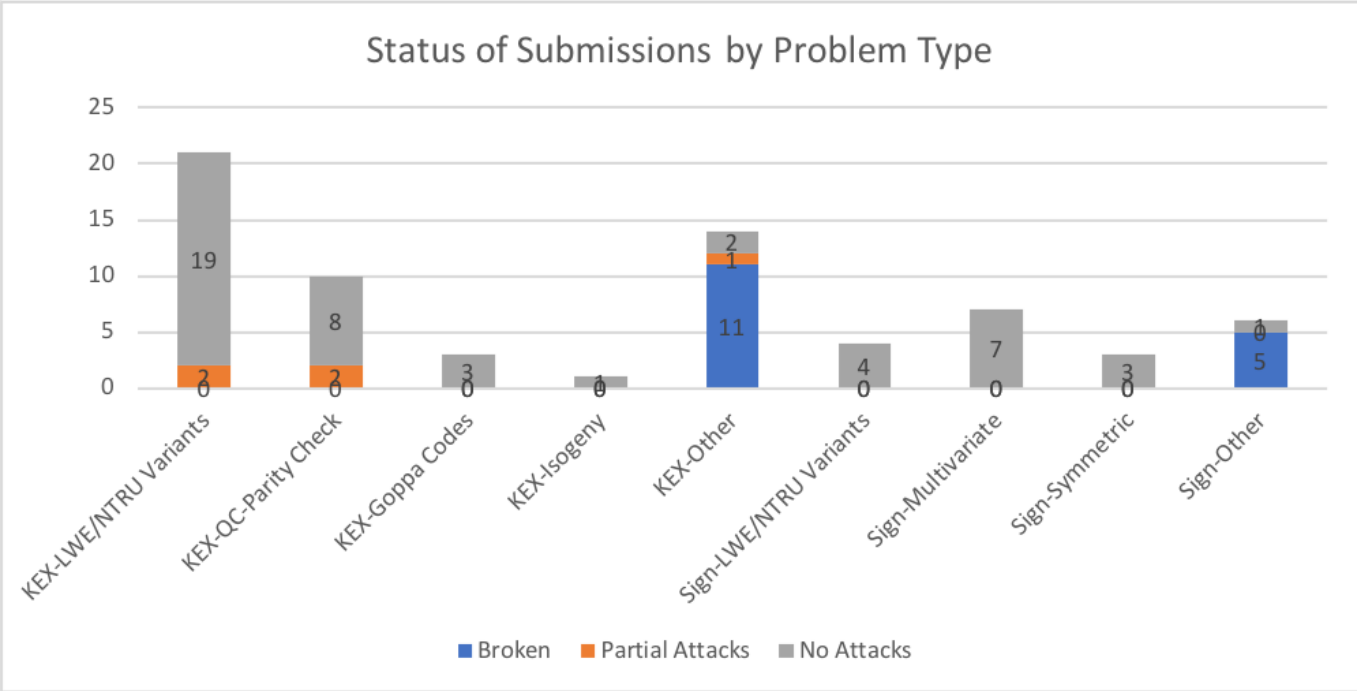

Breaks by Problem Category

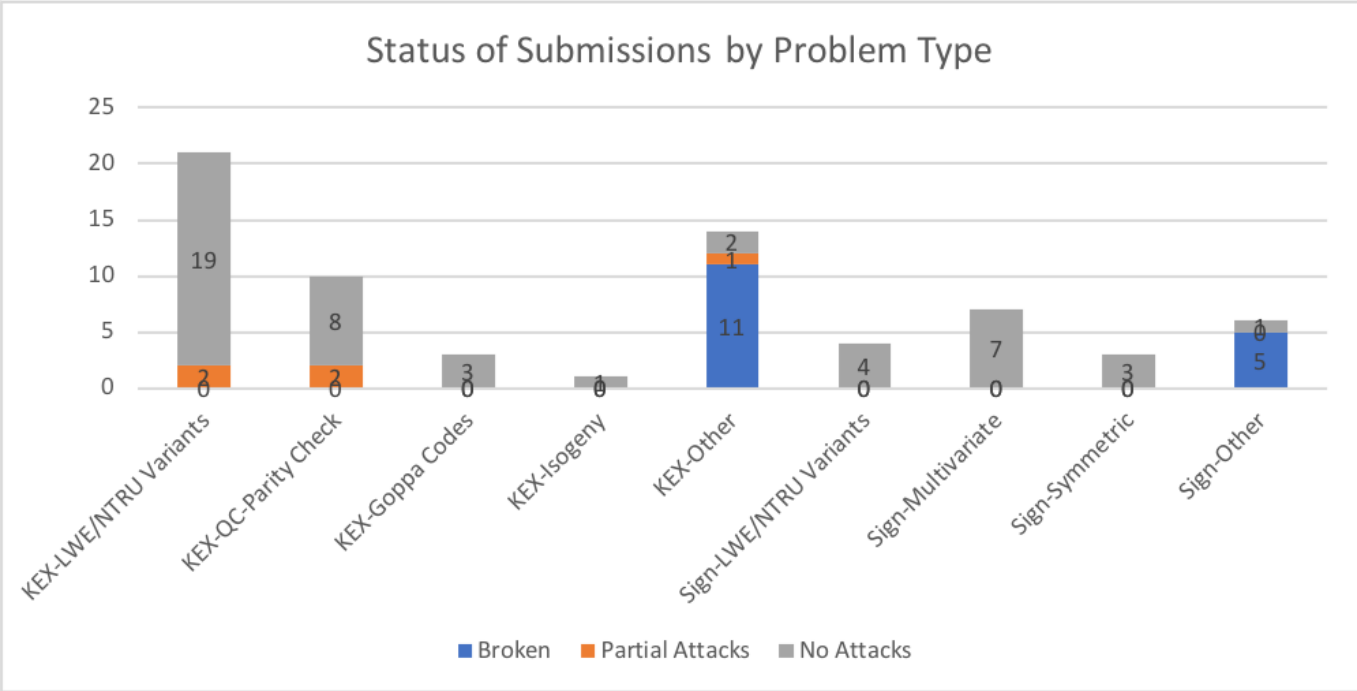

- Problem: Have people only looked at the "low hanging fruit?"

Breaks by Problem Category

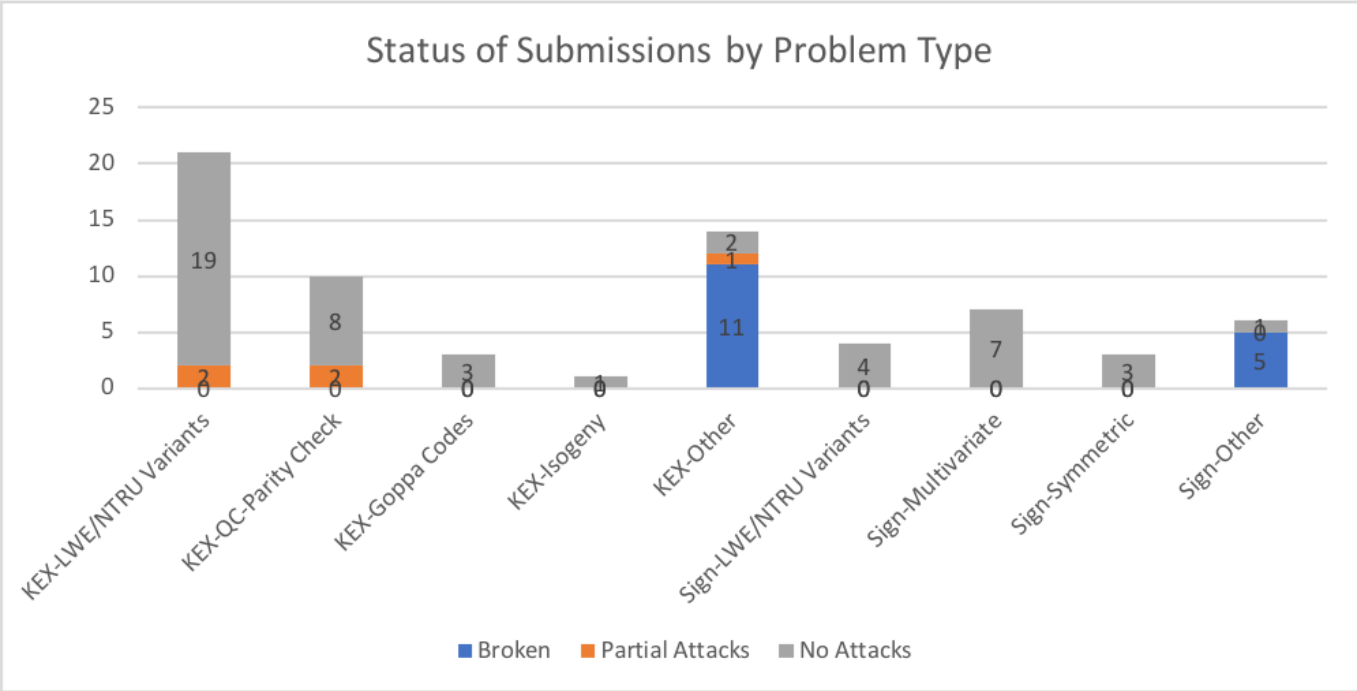

- Problem: Have people only looked at the "low hanging fruit?"
- Remainder of Talk: 10 questions re: security of lattice-based schemes
  - Please jump in with answers!

Question 1

What is the "right" cost model for BKZ?

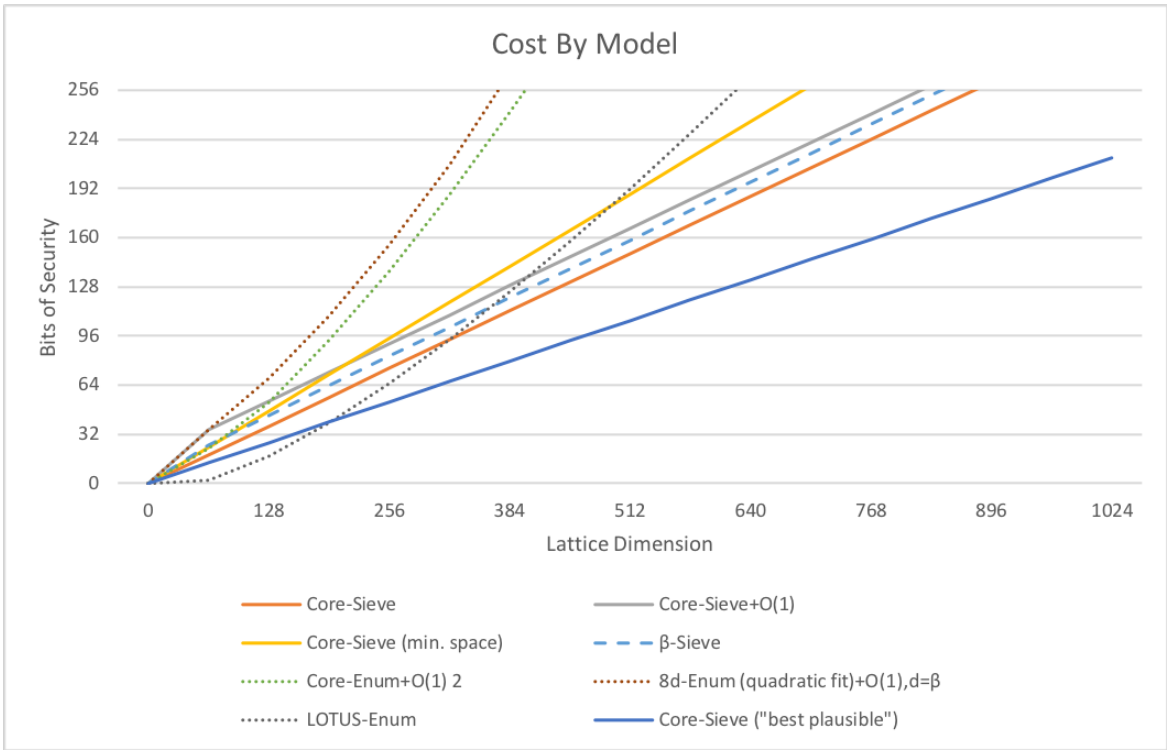

Figure: Cost Models (Lattice Dimension vs bits of security)

- Core or No Core?
  - Core models consider cost of 1 oracle call in BKZ
  - Is the # of calls important for parameter setting?

Main Disputes on Cost Models

- Core or No Core?
  - Core models consider cost of 1 oracle call in BKZ
  - Is the # of calls important for parameter setting?
- Can Sieving Be Considered Feasible?
  - What are reasonable physical MAXSPACE requirements?
  - Is there any hope of reducing sieving's space requirements?

Main Disputes on Cost Models

- Core or No Core?
  - Core models consider cost of 1 oracle call in BKZ
  - Is the # of calls important for parameter setting?
- Can Sieving Be Considered Feasible?
  - What are reasonable physical MAXSPACE requirements?
  - Is there any hope of reducing sieving's space requirements?

- Can sieving algorithms beat/meet the "kissing constant" bound?

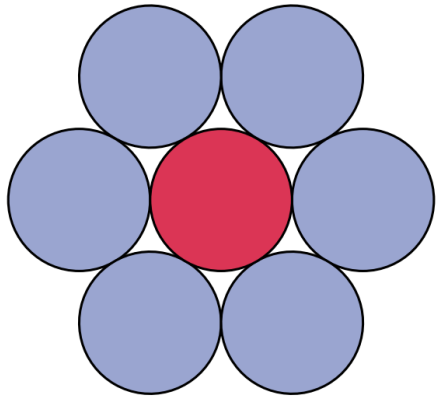

Figure 1. The perfect kissing arrangement for $n = 2$.

- "Interested parties" have contested meaningfulness of reductions to worst-case problems

- "Interested parties" have contested meaningfulness of reductions to worst-case problems

Analysis of Chatterjee et al. (2016)

Let $\alpha = \tilde{O}(1/n)$ (so that $\gamma = \tilde{O}(n^2)$. If there exists an algorithm W, that can distinguish $m = n^c$ LWE_α samples from m uniform samples for a $1/n^{d_1}$ fraction of all $\mathbf{s} \in \mathbb{Z}_q^n$ with advantage at least $1/n^{d_2}$, then there is a polynomial-time algorithm W' for solving SIVP_γ that calls the W oracle a total of $6n^{10+c+d_1+2d_2}q$ times.

- "Interested parties" have contested meaningfulness of reductions to worst-case problems

Analysis of Chatterjee et al. (2016)

Let $\alpha = \tilde{O}(1/n)$ (so that $\gamma = \tilde{O}(n^2)$. If there exists an algorithm W, that can distinguish $m = n^c$ LWE_α samples from m uniform samples for a $1/n^{d_1}$ fraction of all $\mathbf{s} \in \mathbb{Z}_q^n$ with advantage at least $1/n^{d_2}$, then there is a polynomial-time algorithm W' for solving SIVP_γ that calls the W oracle a total of $6n^{10+c+d_1+2d_2}q$ times.

- Other reductions also fail to yield useful concrete (non-asymptotic) bounds

Question 2

For a non-"weird" adversary, how non-tight are the existing reductions?

Lemma 3.7 in Regev '05

If there exists an oracle W that solves LWE_α in time t with probability ϵ, then there exists an oracle W' that, given samples from $A_{\mathbf{s},\beta}$, $\beta \leq \alpha$, solves LWE_β in time $tn + n^c$ with probability $\epsilon - \mathrm{negl}(n)$.

Question 2

For a non-"weird" adversary, how non-tight are the existing reductions?

Lemma 3.7 in Regev '05

If there exists an oracle W that solves LWE_α in time t with probability ϵ, then there exists an oracle W' that, given samples from $A_{\mathbf{s},\beta}$, $\beta \leq \alpha$, solves LWE_β in time $tn + n^c$ with probability $\epsilon - \mathrm{negl}(n)$.

- Necessary if $|\mathcal{A}(\mathsf{LWE}_\beta), \mathcal{A}(U)| \geq 1/n^c$, $|\mathcal{A}(\mathsf{LWE}_{\beta/2}), \mathcal{A}(U)| \leq \mathrm{negl}(n)$
  - Seems very "weird" for an adversary to "naturally" have this property

Question 2

For a non-"weird" adversary, how non-tight are the existing reductions?

Lemma 3.7 in Regev '05

If there exists an oracle W that solves LWE_α in time t with probability ϵ, then there exists an oracle W' that, given samples from $A_{\mathbf{s},\beta}$, $\beta \leq \alpha$, solves LWE_β in time $tn + n^c$ with probability $\epsilon - \mathrm{negl}(n)$.

- Necessary if $|\mathcal{A}(\mathsf{LWE}_\beta), \mathcal{A}(U)| \geq 1/n^c$, $|\mathcal{A}(\mathsf{LWE}_{\beta/2}), \mathcal{A}(U)| \leq \mathrm{negl}(n)$
  - Seems very "weird" for an adversary to "naturally" have this property
- May just be an artifact of black box proofs rather than real problem

Question 2

For a non-"weird" adversary, how non-tight are the existing reductions?

Lemma 3.7 in Regev '05

If there exists an oracle W that solves LWE_α in time t with probability ϵ, then there exists an oracle W' that, given samples from $A_{\mathbf{s},\beta}$, $\beta \leq \alpha$, solves LWE_β in time $tn + n^c$ with probability $\epsilon - \mathrm{negl}(n)$.

- Necessary if $|\mathcal{A}(\mathsf{LWE}_\beta), \mathcal{A}(U)| \geq 1/n^c$, $|\mathcal{A}(\mathsf{LWE}_{\beta/2}), \mathcal{A}(U)| \leq \mathrm{negl}(n)$
  - Seems very "weird" for an adversary to "naturally" have this property
- May just be an artifact of black box proofs rather than real problem
- Full analysis of "necessity" of each step would be nice

Question 3

Can we get better reductions for LWR?

- State of the art for regular LWR
  - From standard LWE:
    samples bounded in terms of amount of rounding
  - From LWE with bounded uniform error:
    q/p factor loss in number of samples

Question 3

Can we get better reductions for LWR?

- State of the art for regular LWR
  - From standard LWE:
    samples bounded in terms of amount of rounding
  - From LWE with bounded uniform error:
    q/p factor loss in number of samples
- Ring/Module-LWR:
  - Reduction to Extended LWE over Rings ... but is that helpful?
  - Otherwise, nothing except the trivial reduction

Question 4

Are there attacks performing better on LWR than "equivalent" LWE?

- LWE with (small) ($\ll \sqrt{n}$) bounded uniform errors:
  - Insecure with sufficient samples [Arora-Ge]
  - Secure for small numbers of samples [Micciancio-Peikert, etc.]

Question 4

Are there attacks performing better on LWR than "equivalent" LWE?

- LWE with (small) ($\ll \sqrt{n}$) bounded uniform errors:
  - Insecure with sufficient samples [Arora-Ge]
  - Secure for small numbers of samples [Micciancio-Peikert, etc.]
- For LWR:
  - Attack works
  - Reductions don't
  - I have some ideas here, could use help ...

Question 5

Is it reasonable to suspect that worst-case ideal lattice problems are not quantum-hard for reasonable parameters?

- hardness gap for ideal vs. general at $\exp(\tilde{O}(\sqrt{n}))$ approximation factors [Cramer-Ducas-Wesolowski]

Question 5

Is it reasonable to suspect that worst-case ideal lattice problems are not quantum-hard for reasonable parameters?

- hardness gap for ideal vs. general at $\exp(\tilde{O}(\sqrt{n}))$ approximation factors [Cramer-Ducas-Wesolowski]
- Are subexponential algorithms possible for polynomial approx. factors?

Question 6

If worst-case ideal lattice problems aren't quantum hard, should we worry about ring and module LWE?

- Key Question: Is there a "phase shift" between dimensions 1 and 2 for module lattices?

Question 6

If worst-case ideal lattice problems aren't quantum hard, should we worry about ring and module LWE?

- Key Question: Is there a "phase shift" between dimensions 1 and 2 for module lattices?
- How hard is GapSVP for dimension 2 module lattices (over a large ring R)?
  - Like dimension 2 general lattices, doesn't follow from determinant (and Minkowski) alone.
  - If no harder than for dimension 2 general lattices, might have reason to worry

Question 7

Are there other subfield attacks possible?

- Known ones (Lauter et al) don't seem to directly apply to NIST submissions

Question 7

Are there other subfield attacks possible?

- Known ones (Lauter et al) don't seem to directly apply to NIST submissions
- Bounds from [Peikert'16] don't apply to many either

Question 8

Are algebraic attacks possible that beat basis reduction?

Arora-Ge Attack on (discretized) LWE with errors in $[-\beta, \beta]$

$f_{(\mathbf{a},b)}(\mathbf{s}) = \prod_{i\in[-\beta,\beta]}(\langle \mathbf{a}, \mathbf{s}\rangle - b - 1 - i) = 0 \bmod q$

where f has $\approx n^{2\beta+1}$ monomials of distinct products of the s_i.

Question 8

Are algebraic attacks possible that beat basis reduction?

Arora-Ge Attack on (discretized) LWE with errors in $[-\beta, \beta]$

$f_{(\mathbf{a},b)}(\mathbf{s}) = \prod_{i\in[-\beta,\beta]}(\langle \mathbf{a}, \mathbf{s}\rangle - b - 1 - i) = 0 \bmod q$

where f has $\approx n^{2\beta+1}$ monomials of distinct products of the s_i.

- Requires $\approx n^{2\beta+1}$ samples to linearize, not enough for any submissions

Question 8

Are algebraic attacks possible that beat basis reduction?

Arora-Ge Attack on (discretized) LWE with errors in $[-\beta, \beta]$

$f_{(\mathbf{a},b)}(\mathbf{s}) = \prod_{i\in[-\beta,\beta]}(\langle \mathbf{a}, \mathbf{s}\rangle - b - 1 - i) = 0 \bmod q$

where f has $\approx n^{2\beta+1}$ monomials of distinct products of the s_i.

- Requires $\approx n^{2\beta+1}$ samples to linearize, not enough for any submissions
- Direct use of Gröbner bases appears worse than basis reduction
  - A cleverer approach using sparseness of secrets might ...

Question 8

Are algebraic attacks possible that beat basis reduction?

Arora-Ge Attack on (discretized) LWE with errors in $[-\beta, \beta]$

$f_{(\mathbf{a},b)}(\mathbf{s}) = \prod_{i\in[-\beta,\beta]}(\langle \mathbf{a}, \mathbf{s}\rangle - b - 1 - i) = 0 \bmod q$

where f has $\approx n^{2\beta+1}$ monomials of distinct products of the s_i.

- Requires $\approx n^{2\beta+1}$ samples to linearize, not enough for any submissions
- Direct use of Gröbner bases appears worse than basis reduction
  - A cleverer approach using sparseness of secrets might ...
- Could some hybrid approach work?

Question 9

What amount of decryption error is small enough to prevent CCA attacks?

- (Nearly) all submissions use Fujisaki-Okamoto to get CCA KEMs.
  - Uses cryptographic hash of the (random) message as source of all encryption randomness
  - Re-encrypts to check validity

Question 9

What amount of decryption error is small enough to prevent CCA attacks?

- (Nearly) all submissions use Fujisaki-Okamoto to get CCA KEMs.
  - Uses cryptographic hash of the (random) message as source of all encryption randomness
  - Re-encrypts to check validity
- Precomputation of "good" set of attack messages meaningful when Hamming weights of $\mathbf{s}, \mathbf{e}$ not fixed
  - Search offline for input messages yielding high Hamming weight $\mathbf{s}, \mathbf{e}$
  - Allows attack on scheme with "honest" decryption failure rate below 2^{-128} with fewer than 2^{64} queries

Question 9

What amount of decryption error is small enough to prevent CCA attacks?

- (Nearly) all submissions use Fujisaki-Okamoto to get CCA KEMs.
  - Uses cryptographic hash of the (random) message as source of all encryption randomness
  - Re-encrypts to check validity
- Precomputation of "good" set of attack messages meaningful when Hamming weights of $\mathbf{s}, \mathbf{e}$ not fixed
  - Search offline for input messages yielding high Hamming weight $\mathbf{s}, \mathbf{e}$
  - Allows attack on scheme with "honest" decryption failure rate below 2^{-128} with fewer than 2^{64} queries
- With fixed Hamming weights, can we find "good" set?

Question 10

Are there variants of or alternatives to reconciliation and error correction for key exchange not covered by patents?

- Round 2 tweaks or mergers that clarify the patent situation could be very useful!