From: מצו״ב <matzov@idf.il> via pqc-forum@list.nist.gov
To: pqc-forum@list.nist.gov
Subject: [pqc-forum] Improved Dual Lattice Attack
Date: Monday, April 04, 2022 12:39:02 PM ET

- בלמ"ס -

Dear PQC researchers,

The Center of Encryption and Information Security (MATZOV) of the IDF has conducted an internal audit of leading Post-Quantum cryptographic (PQC) schemes, focusing on the Learning With Errors and Rounding problems.

After consultations with NIST over the last few months – we have decided to release the audit as a Technical Report available for public review.

https://doi.org/10.5281/zenodo.6412487

Our report presents several improvements to the dual lattice attack, which induce a noticeable reduction in the security estimation for Kyber, Saber and Dilithium, bringing them below the required threshold.

The report does not intend to provide a complete analysis of all post-quantum candidates, nor to recommend usage of specific algorithms. Rather, this publication is meant to share advances in the cryptanalysis of lattices which we believe to be relevant to the academic research in the field.

We acknowledge the remarkable work done by NIST in the process and its impact – creating interest in the post-quantum field and promoting new cryptographic schemes.

A prudent approach for these schemes is recommended, as research in the field is constantly evolving and much remains unstudied. Therefore, as a contribution to the community, the report includes further research ideas which we deem interesting.

MATZOV, IDF

From: Daniel Apon <dapon.crypto@gmail.com> via pqc-forum@list.nist.gov
To: מצו״ב <matzov@idf.il>
CC: pqc-forum@list.nist.gov
Subject: Re: [pqc-forum] Improved Dual Lattice Attack
Date: Monday, April 04, 2022 01:05:14 PM ET

Good afternoon (evening) MATZOV,

Thank you very kindly for releasing this paper! I'm sure it must have been a challenge. I'm looking forward to reading it in much detail.

In the paper, you highlight in multiple places that your analysis occurs in the RAM model. There are several variables involved in assessing the cost of an actual attack. Have you considered "more realistic" memory-costing in your analyses?

There is a long history of discussion on this forum this past summer and fall about what the proper way to model attacker memory costs are. Broadly, there is the RAM model as compared to a variety of so-called "local" models. Some prominent examples of these alternative/local models include the 2D nearest neighbor model (often called the Square-Root model) and the 3D nearest neighbor model (often called the Cube-Root model). As a further example, the NIST PQC call for proposals defined a version of the (quantum) circuit model involving a MAXDEPTH parameter for gate-operations in series.

Note that this question is particularly relevant as the defining cost-metric (the computational hardness vs. AES) involves essentially no memory costs, whereas lattice sieving historically involves high memory costs.

One approach to addressing alternative models of memory-costing would be to define a single 'trade-off value' between max memory and computational properties (width, depth, network topology of the cryptanalytic device, etc.) and simply add some number of bit operations to the bit-complexity of algorithms in the RAM model. Typically, jointly settling on such a value is a difficult task, with many unknowns. See, for example, the discussion in the Kyber Round 3 spec that gives ranges of bit-complexities (either positive or negative) based on various uncertainty-factors. Do you have a preferred view on how to do this generic analysis?

However, more importantly: Do you find that any of your new algorithmic approaches have a concrete cost that would differ from a generic model-to-model analysis? (I hope to read through the work and answer "No!" but perhaps you have an opinion you could share now.)

Thank you for your insightful work and significant contribution to the science.

Best regards,
--Daniel Apon
Cryptography Lead, the MITRE Corporation
dapon@mitre.org

On Mon, Apr 4, 2022 at 12:38 PM מצו״ב <Matzov@idf.il> wrote:

> - בלמ"ס -
>
> Dear PQC researchers,
>
> The Center of Encryption and Information Security (MATZOV) of the IDF has conducted an internal audit of leading Post-Quantum cryptographic (PQC) schemes, focusing on the Learning With Errors and Rounding problems.
>
> After consultations with NIST over the last few months – we have decided to release the audit as a Technical Report available for public review.
>
> https://doi.org/10.5281/zenodo.6412487
>
> Our report presents several improvements to the dual lattice attack, which induce a noticeable reduction in the security estimation for Kyber, Saber and Dilithium, bringing them below the required threshold.
>
> The report does not intend to provide a complete analysis of all post-quantum candidates, nor to recommend usage of specific algorithms. Rather, this publication is meant to share advances in the cryptanalysis of lattices which we believe to be relevant to the academic research in the field.
>
> We acknowledge the remarkable work done by NIST in the process and its impact – creating interest in the post-quantum field and promoting new cryptographic schemes.
>
> A prudent approach for these schemes is recommended, as research in the field is constantly evolving and much remains unstudied. Therefore, as a contribution to the community, the report includes further research ideas which we deem interesting.
>
> MATZOV, IDF
>
> --
>

> pqc-forum/DM5PR14MB140491EED7763525AB6C67E2A1E59%40DM5PR14MB1404.namprd14.prod.outlook.com.

--
You received this message because you are subscribed to the Google Groups "pqc-forum" group.
To unsubscribe from this group and stop receiving emails from it, send an email to pqc-forum+unsubscribe@list.nist.gov.
To view this discussion on the web visit https://groups.google.com/a/list.nist.gov/d/msgid/pqc-forum/CAPxHsSKV5J7pUA6CKttgyeTYqxqd2_J4Z67nuLvOTgf4_nzT1w%40mail.gmail.com.

From: Daniel Apon <dapon.crypto@gmail.com> via pqc-forum@list.nist.gov
To: pqc-forum <pqc-forum@list.nist.gov>
CC: D. J. Bernstein <djb@cr.yp.to>, pqc-...@list.nist.gov <pqc-forum@list.nist.gov>
Subject: Re: [pqc-forum] Improved Dual Lattice Attack
Date: Sunday, April 10, 2022 10:30:27 AM ET

Good morning (or other state-of-day), Dan,

For the sake of focusing this discussion to the open scientific questions, how about I simply concede the points in your commentary ("NIST sucks," "Lattices are frogs," etc.). I think what's most interesting here is getting to a precise (non-napkin-math) calculation of the memory costs in MATZOV's new algorithms. (Although I've only read through so far for a high level understanding, the new algorithms initially appear correct and well-analyzed in the RAM model to me.)

As a starting point for what I'd like to get at, consider the example range of calculations you began with (along with the caveat):
"The numerical examples on page 103 show the extra security ranging from
a 2^40 factor for Core-SVP 2^129 to a 2^90 factor for Core-SVP 2^271.

See Section 6 of the same document for many reasons that these numbers
can be underestimating or overestimating the actual attack costs."

Let me try to replicate that by hand for Kyber-1024 (to show some insufficiencies with a napkin math approach, whether this one I'm cooking up now or any other). The Kyber-1024 Round 3 spec claims classical Core-SVP hardness of 256. Using what I'll informally call "Thijs's May 2019 heuristic" (see https://groups.google.com/a/list.nist.gov/g/pqc-forum/c/VHw5xZOJP5Y/m/nklFHrY4AwAJ), we can set o(1)=0 in the exponent of list size estimates from sieving analyses. Let's arbitrarily pick the 3-sieve algorithm's minimal list size costs from the G6K paper as the memory size estimate, ignoring runtime overhead induced by the smaller list size. This is 2^{0.1788n} from Fig 2 / page 19 of https://eprint.iacr.org/2019/089.pdf. Let's use the NTRU Prime memory-costing heuristic of N^{.5} / (2^5).

Putting all of this together, we get log_2((2^{.1788 * 1024})^(1/2) / (2^5)) ~= 86.55.
(Note that doing the same calculation for Kyber-512 gets you log_2((2^{.1788 * 512})^(1/2) / (2^5)) ~= 40.77 in a straightforward way.)

Now I want to just 'scale this up' to a hypothetical Kyber (ignoring issues like powers of 2 in the dimension) by taking percentages of Core-SVP values as
log_2((2^{.1788 * (271/256) * 1024})^(1/2) / (2^5)) ~= 91.91. //and I'm already off by a couple bits from 2^90

A more rigorous approach would begin by calculating the concrete list size of the new dual lattice attack algorithm (the paper gives a concrete way to calculate the number of samples, D, on page 39 -- even if it's a mess to unravel). Then, one should look at the precise movement of memory required by Algorithms 2 (page 15) and 3 (page 17). It's important here to consider the memory architecture and the actual steps of the algorithm. In order to arrive at a conservative lower bound for this algorithm, it's probably best to model the memory architecture in the most ideal way possible (simply a single, uniform 2D grid, perhaps).

That is, following the NTRU Prime 3rd Round spec, if 2^30 bits of DRAM at 22nm fit in a 5mm x 5mm square, and one needs 2^90 lattice vectors, then we're talking about a 2D grid of bits spanning at least 2^(60/2) * 5mm on each side if arranged in a square, or approximately 42% of the width of Planet Earth. This is approximately the width of 1.5 our moon Luna. Re-arranging as a roughly spherical shape (perhaps in layers of 2D grids) for efficiency of communication, one derives approximately a small-moon-sized cryptanalytic device: https://www.youtube.com/watch?v=8Nho44lGVV8

So it's clearly critical that parameterizations for the new dual lattice attack consider values of D that are sufficiently small to fit on a real-world cryptanalytic device that could be constructed without importing matter from other solar systems. But that issue aside, then it's important to consider the cost of memory movement (as you highlight). A precise and perspicuous analysis of those exact costs is still outstanding.

Taking D^(1/2) / (2^5) as the additional running-time cost-factor is a reasonable first approach (even if I believe that calculation is largely over-estimating the run-time costs associating with large memory..), but a more rigorous analysis should consider the concrete steps performed by the algorithm, as well as any improvements that might be gained by how lattice vectors are laid out in memory during sieving. (This is not an idle intellectual exercise, since insights here will be applicable even when re-tooling the algorithm to achieve smaller values of D that could be effective in the real world.)

Toward that end, "simple" optimizations like in MATZOV's paper, Section 5.4 (Efficient Updating of the FFT Input) will be very strong.

Best regards,
--Daniel Apon

On Saturday, April 9, 2022 at 7:54:17 AM UTC-4 D. J. Bernstein wrote:

> Daniel Apon writes:
> > lattice sieving historically involves high memory costs
>
> The best estimates available here aren't comforting for Kyber.
>
> https://ntruprime.cr.yp.to/nist/ntruprime-20201007.pdf, using energy
> numbers published by Intel, estimates each access to a bit within N bits
> of memory as matching the cost of N^0.5/2^5 bit operations (page 57),
> and uses this to estimate how RAM costs affect concrete sieving costs.
> The numerical examples on page 103 show the extra security ranging from
> a 2^40 factor for Core-SVP 2^129 to a 2^90 factor for Core-SVP 2^271.
>
> See Section 6 of the same document for many reasons that these numbers
> can be underestimating or overestimating the actual attack costs. The
> round-3 Kyber documentation estimates that Kyber-512 attacks cost
> between 2^135.5 and 2^165.5 "gates", where the "floor" for NIST's lowest
> security category is 2^143 "gates".
>
> The new attack paper starts from the Kyber documentation's middle
> estimate, namely 2^151.5, and says it's reducing the attack costs by a
> factor 2^14, to 2^137.5, using better attack algorithms.
>
> It's clear that at least some of the algorithm-analysis uncertainties
> stated in the round-3 Kyber documentation are regarding speedups that
> combine with the speedups in the new paper. The optimistic possibility
> for the attacker is that the new paper is actually giving a 2^14 savings
> from 2^135.5, i.e., 2^121.5 bit operations.
>
> Does accounting for real RAM costs close the gap between 2^121.5 and

> 2^143? One might think that, sure, this is covered by the 2^40 mentioned above: Kyber-512 previously had security 2^40*2^135.5 = 2^175.5, so a 32.5-bit security margin, and the new paper is reducing this to an 18.5-bit security margin: i.e., the new paper is merely cutting out 40% of the Kyber security margin, rather than breaking Kyber outright.
>
> But let's look more closely at the numbers. As a preliminary point, round-3 Kyber-512 is starting from Core-SVP just 2^112 and revised-Core-SVP just 2^118, with exponent 87% and 91% of 129 respectively, so the obvious estimate is about 2^36 instead of 2^40.
>
> Furthermore, this 2^36 is accounting for the energy cost of accesses to a giant RAM array, while it's clear that many of the bits of security beyond Core-SVP claimed in the round-3 Kyber security analysis are coming from accounting for the cost of local bit operations. These effects don't multiply; they add!
>
> Internally, Core-SVP is starting from estimates of the number of "operations" inside sieving. It makes sense to say that the attacker needs to pay for the large-scale memory access inside each "operation". It also makes sense to say that the attacker needs to pay for all the bit operations inside each "operation". But the local bit operations are an asymptotically irrelevant extra cost on top of the memory access, and the best bet is that they don't make much difference for Kyber-512. The real cost of this type of algorithm is, at a large scale, driven primarily by data motion, not by local computation.
>
> (The new paper seems to have some local speedups to the sieving inner loop, which similarly should be presumed to make little difference next to the memory-access bottleneck, but my understanding is that this is under half of the bits of security loss that the paper is reporting.)
>
> So I don't see how current knowledge can justify suggesting that the costs of RAM rescue Kyber-512 from the new attack. It seems entirely possible that the real costs of this Kyber-512 attack are considerably below the costs of a brute-force AES-128 attack. Deciding this one way or the other will require much more serious analysis of attack costs.

Certainly it's disturbing to see Kyber-512 dropping from (1) supposedly "conservative" to (2) bleeding-edge in bit operations and then to (3) apparently broken in bit operations and bleeding-edge in real cost. The Kyber documentation really should stop using the word "conservative".

It's also deeply concerning that the uncertainties in evaluating costs of lattice attacks, such as the 2^30 uncertainty factor (2^135.5 through 2^165.5) in the round-3 Kyber submission, have been weaponized again and again to suggest that we shouldn't worry about a paper speeding up lattice attacks by a factor 2^5 or 2^10 or 2^15. The cumulative effect of years of such speedups has clearly been far more than 2^30. (The mascot for lattice-based cryptography should be a slow-boiled frog.)

As a historical matter, we've seen again and again in cryptography that a series of public attack advances has culminated in a feasible attack. The community recognizes and promotes progress by putting serious effort into _quantifying_ the attack costs. Security evaluation is obviously the most important input to NISTPQC, so the NISTPQC rules should have been designed to prioritize and assist quantification of attack costs.

Back in 2016, NIST proposed NISTPQC evaluation rules that instead prioritized fake confidence in staying above cutoffs such as AES-128 and AES-192. I correctly predicted in

https://blog.cr.yp.to/20161030-pqnist.html

that "Quantitatively comparing post-quantum public-key security levels is going to be a nightmare". I recommended throwing away the security cutoffs and replacing them with the traditional focus on analyzing algorithm costs as accurately as possible. NIST's subsequent arguments for prioritizing cutoffs don't stand up to examination---for details see Appendix B.5 of

https://ntruprime.cr.yp.to/latticerisks-20211031.pdf

---and, even worse, the cutoffs are in cost metrics that NIST _pretends_ to have defined but has never actually defined. (See below.)

Going forward, clearly NIST is going to include some lattice systems in its first standards; supposedly we'll find out which ones any moment now. Maybe NIST is going to recklessly include the smallest proposed parameters---but apparently we won't find this out for a while; NIST indicated, surprisingly, that it _isn't_ planning to name parameters yet. Given how much supposed security lattices have lost over the years and how large the remaining attack surface is, there's a worrisome level of risk even for bigger parameters such as Kyber-1024. So there's an ongoing need for clear quantification of the costs of lattice attacks.

> the NIST PQC call for proposals defined a version of the (quantum)
> circuit model involving a MAXDEPTH parameter for gate-operations in series.

False. NIST described some desiderata for a model, such as MAXDEPTH, but never defined a model to be used for NISTPQC. In particular, NIST never defined the set of allowed "gates", despite requests for clarification. (The algorithms literature includes many different gate sets, often giving wildly different algorithm costs; see, e.g., how Ambainis's distinctness paper uses quantum RAM gates.) Section 5.4 of

https://cr.yp.to/papers/categories-20200918.pdf

gives quotes, references, and numerical examples to illustrate the lack of definition.

We've seen repeatedly how the ambiguities in NIST's pseudo-definitions have been exploited to downplay attacks against some systems---this reached amazing heights with last month's excuses for not withdrawing the claim that the dimension-256 parameters in the preliminary Frodo design from Lindner--Peikert "appear to be at least as secure as AES-128"---and at the same time to hype attacks against other systems. NIST's failure to pick a metric has done far more damage to comparisons than whatever damage would have been done from the selection of a metric that turns out to not be perfectly realistic.

> ---D. J. Bernstein

From: D. J. Bernstein <djb@cr.yp.to> via pqc-forum@list.nist.gov
To: pqc-forum@list.nist.gov
Subject: Re: [pqc-forum] Improved Dual Lattice Attack
Date: Tuesday, April 12, 2022 11:51:33 AM ET
Attachments: smime.p7m

Daniel Apon writes:

> That is, following the NTRU Prime 3rd Round spec, if 2^30 bits of DRAM at 22nm fit in a 5mm x 5mm square, and one needs 2^90 lattice vectors, then we're talking about a 2D grid of bits spanning at least 2^(60/2) * 5mm on each side if arranged in a square, or approximately 42% of the width of Planet Earth.

First, 22nm is very far from the latest chip technology. See, e.g.,

https://en.wikipedia.org/wiki/5_nm_process

and the links from there to upcoming technology nodes. The advance from 22nm to 5nm took just 8 years, and made transistors about 3x smaller in each direction.

Second, dividing the width mentioned above by about 30 gets down to the radius of existing tracts of uninhabited land owned by governments with a history of carrying out attacks.

Third, the NISTPQC call for proposals says "The security provided by a cryptographic scheme is the most important factor in the evaluation", not "The security provided by a cryptographic scheme against 22nm attackers is the most important factor in the evaluation". It would be astonishing if a project trying to protect against the long-term quantum threat were allowing such shortsighted security goals.

Fourth, the "Improved Dual Lattice Attack" that this thread is about appears to considerably reduce the attack costs. The Kyber documentation mentions various other reasons for a 2^30 uncertainty factor regarding the attack costs. Given the context, the above mention of "2^90 lattice vectors" needs to be accompanied by a warning that the costs could be much lower.

Fifth, restricting attention to 22nm is not endorsed by the NTRU Prime documentation. On the contrary, the documentation explicitly points to the trend towards smaller technology---and stays away from selecting any bleeding-edge parameters in the first place.

The reason 22nm shows up in the documentation is that, as a separate question from how expensive computation is on an absolute scale, it's important to understand the _relative_ costs of different operations inside attacks. Intel was nice enough to publish detailed energy figures for 22nm in 2015; the NTRU Prime documentation compares those to readily available data regarding 22nm RAM, obtaining an estimated sqrt(N)/2^5 _ratio_ between the cost of accessing a bit in N bits of RAM and the cost of a bit operation. The documentation then explains why it's reasonable to guess that future technology will have similar ratios:

> Smaller technology than 22nm reduces the cost of bit operations, as noted above, while also packing memory more densely. It is reasonable to guess that these effects will stay approximately balanced: compared to performing an AND or XOR on two bits within a tiny distance, moving a bit over a tiny distance uses the same basic physical phenomena but uses those phenomena in a simpler way, and having it cost a constant factor less is unsurprising. This guess is not meant as a substitute for continuing to monitor technology trends.

None of this is endorsing the idea that the security goal should be security against 22nm attackers.

> consider values of D that are sufficiently small to fit on a
> real-world cryptanalytic device that could be constructed without importing
> matter from other solar systems

Building something in the ballpark of 2^60 grams of chips doesn't require "importing matter from other solar systems": the Earth weighs 2^92 grams, and silicon etc. are very common.

Is an attacker in the foreseeable future going to build 2^60 grams of chips, and have the energy budget to run those chips? No, and no. See my pqc-forum email dated 20 Nov 2016 05:14:07 +0000. But it would be crazy for NISTPQC to set its minimum security level at just barely stopping attacks with current technology.

The NISTPQC call for proposals sets a minimum security level considerably above this. Specifically, it sets brute-force search for a single AES-128 key as a "floor" for security. It's clear that attackers aren't anywhere near carrying out 2^128 operations.

What happens if an attack falls in the gap: easier to break than AES-128 but still not feasible for attackers today? Unfortunately, we've seen that the answer depends on the cryptosystem being attacked:

* For some cryptosystems, the infeasibility is hyped. So much equipment needed to finish in a reasonable time! So much energy! Look at how hard this would be!

* For other cryptosystems, we instead hear that anything below 2^128 operations, even with access to a massive memory array counted as just one "operation", counts as a break.

Specifically, the pqc-forum comparisons of attack costs to Earth resources have been encouraging consideration of larger-scale attacks for cryptosystems in general (Perlner email dated 17 Aug 2020 17:41:27 +0000) and for LAC in particular (Hamburg email dated 12 Apr 2018 14:50:16 -0400). For most other specific attacks (including infeasible attacks), a comparison to Earth resources isn't even mentioned. But, for the latest Kyber-512 security loss, we're seeing an Earth comparison being used to suggest that the attack should be ignored. See also

https://twitter.com/mjos_crypto/status/1511027605033652234

saying that dropping below AES-128 security is ok.

Some of the NIST statements have suggested that AES-128 isn't a hard floor for security. What _is_ the floor, then? If an attack uses only 2^128 _bit_ operations, does that count as a break? What if it also needs 2^70 bits of RAM? NIST keeps dodging concrete questions, and keeps dodging the question of which "gates" it allows. The unclear boundaries are then used to reward some cryptosystems and punish others.

The documented facts are that _some_ attack speedups are big asymptotic changes (e.g., L(1) down to L(1/2)), but most attack speedups (and many breaks) come from people looking more closely at attack costs. For these people, the way that NIST promotes a yes/no cutoff question regarding security, without actually defining the cutoff, is a big disincentive to the necessary research. Instead of saying, wow, this makes an attack 1000x or 1000000x faster, one is faced with people asking whether this speedup crosses NIST's cutoff. How is one supposed to answer this question when NIST doesn't say which cutoff definitions it allows?

So, instead of a scientific process studying clearly defined questions, there's a political process weaponizing a lack of clarity. At some point observers are forced to ask whether the lack of clarity is deliberate.

---D. J. Bernstein

From: Carl Mitchell <carl.mitchell@gomotive.com> via pqc-forum <pqc-forum@list.nist.gov>
To: pqc-forum@list.nist.gov
Subject: Re: [pqc-forum] Improved Dual Lattice Attack
Date: Tuesday, April 12, 2022 12:21:48 PM ET

The "nm" process sizes are almost entirely marketing terms. TSMCs' 5nm SRAM bit cell size is 0.021µm^2 according to TSMC's marketing materials on semiWiki[1], which is closer to 22nm than 5nm.

That said, I agree that clarity of the metrics is important. Is RAM access considered? What is a "gate operation"? A single clearly-defined metric for security evaluation would ease comparisons, even if it's a non-physical one with planet-sized RAM arrays and instantaneous access at a distance. As long as the assumptions of the metric make attacks *easier* than in real-world systems any real attack should also meet (at least) the same target security level.

[1] https://semiwiki.com/semiconductor-manufacturers/tsmc/283487-tsmcs-5nm-0-021um2-sram-cell-using-euv-and-high-mobility-channel-with-write-assist-at-isscc2020/

—Carl Mitchell

My views herein are my own, not those of my employer (Motive).

On Tue, Apr 12, 2022 at 11:51 AM D. J. Bernstein <djb@cr.yp.to> wrote:

> Daniel Apon writes:
> > That is, following the NTRU Prime 3rd Round spec, if 2^30 bits of DRAM at
> > 22nm fit in a 5mm x 5mm square, and one needs 2^90 lattice vectors, then
> > we're talking about a 2D grid of bits spanning at least 2^(60/2) * 5mm on
> > each side if arranged in a square, or approximately 42% of the width of
> > Planet Earth.
>
> First, 22nm is very far from the latest chip technology. See, e.g.,
>
> https://en.wikipedia.org/wiki/5_nm_process
>
> and the links from there to upcoming technology nodes. The advance from 22nm to 5nm took just 8 years, and made transistors about 3x smaller in

> each direction.
>
> Second, dividing the width mentioned above by about 30 gets down to the radius of existing tracts of uninhabited land owned by governments with a history of carrying out attacks.
>
> Third, the NISTPQC call for proposals says "The security provided by a cryptographic scheme is the most important factor in the evaluation", not "The security provided by a cryptographic scheme against 22nm attackers is the most important factor in the evaluation". It would be astonishing if a project trying to protect against the long-term quantum threat were allowing such shortsighted security goals.
>
> Fourth, the "Improved Dual Lattice Attack" that this thread is about appears to considerably reduce the attack costs. The Kyber documentation mentions various other reasons for a 2^30 uncertainty factor regarding the attack costs. Given the context, the above mention of "2^90 lattice vectors" needs to be accompanied by a warning that the costs could be much lower.
>
> Fifth, restricting attention to 22nm is not endorsed by the NTRU Prime documentation. On the contrary, the documentation explicitly points to the trend towards smaller technology---and stays away from selecting any bleeding-edge parameters in the first place.
>
> The reason 22nm shows up in the documentation is that, as a separate question from how expensive computation is on an absolute scale, it's important to understand the _relative_ costs of different operations inside attacks. Intel was nice enough to publish detailed energy figures for 22nm in 2015; the NTRU Prime documentation compares those to readily available data regarding 22nm RAM, obtaining an estimated sqrt(N)/2^5 _ratio_ between the cost of accessing a bit in N bits of RAM and the cost of a bit operation. The documentation then explains why it's reasonable to guess that future technology will have similar ratios:
>
> Smaller technology than 22nm reduces the cost of bit operations, as noted above, while also packing memory more densely. It is reasonable

to guess that these effects will stay approximately balanced: compared to performing an AND or XOR on two bits within a tiny distance, moving a bit over a tiny distance uses the same basic physical phenomena but uses those phenomena in a simpler way, and having it cost a constant factor less is unsurprising. This guess is not meant as a substitute for continuing to monitor technology trends.

None of this is endorsing the idea that the security goal should be security against 22nm attackers.

> consider values of D that are sufficiently small to fit on a
> real-world cryptanalytic device that could be constructed without importing
> matter from other solar systems

Building something in the ballpark of 2^60 grams of chips doesn't require "importing matter from other solar systems": the Earth weighs 2^92 grams, and silicon etc. are very common.

Is an attacker in the foreseeable future going to build 2^60 grams of chips, and have the energy budget to run those chips? No, and no. See my pqc-forum email dated 20 Nov 2016 05:14:07 +0000. But it would be crazy for NISTPQC to set its minimum security level at just barely stopping attacks with current technology.

The NISTPQC call for proposals sets a minimum security level considerably above this. Specifically, it sets brute-force search for a single AES-128 key as a "floor" for security. It's clear that attackers aren't anywhere near carrying out 2^128 operations.

What happens if an attack falls in the gap: easier to break than AES-128 but still not feasible for attackers today? Unfortunately, we've seen that the answer depends on the cryptosystem being attacked:

* For some cryptosystems, the infeasibility is hyped. So much equipment needed to finish in a reasonable time! So much energy! Look at how hard this would be!

> * For other cryptosystems, we instead hear that anything below 2^128 operations, even with access to a massive memory array counted as just one "operation", counts as a break.
>
> Specifically, the pqc-forum comparisons of attack costs to Earth resources have been encouraging consideration of larger-scale attacks for cryptosystems in general (Perlner email dated 17 Aug 2020 17:41:27 +0000) and for LAC in particular (Hamburg email dated 12 Apr 2018 14:50:16 -0400). For most other specific attacks (including infeasible attacks), a comparison to Earth resources isn't even mentioned. But, for the latest Kyber-512 security loss, we're seeing an Earth comparison being used to suggest that the attack should be ignored. See also
>
> https://twitter.com/mjos_crypto/status/1511027605033652234
>
> saying that dropping below AES-128 security is ok.
>
> Some of the NIST statements have suggested that AES-128 isn't a hard floor for security. What _is_ the floor, then? If an attack uses only 2^128 _bit_ operations, does that count as a break? What if it also needs 2^70 bits of RAM? NIST keeps dodging concrete questions, and keeps dodging the question of which "gates" it allows. The unclear boundaries are then used to reward some cryptosystems and punish others.
>
> The documented facts are that _some_ attack speedups are big asymptotic changes (e.g., L(1) down to L(1/2)), but most attack speedups (and many breaks) come from people looking more closely at attack costs. For these people, the way that NIST promotes a yes/no cutoff question regarding security, without actually defining the cutoff, is a big disincentive to the necessary research. Instead of saying, wow, this makes an attack 1000x or 1000000x faster, one is faced with people asking whether this speedup crosses NIST's cutoff. How is one supposed to answer this question when NIST doesn't say which cutoff definitions it allows?
>
> So, instead of a scientific process studying clearly defined questions, there's a political process weaponizing a lack of clarity. At some point

> observers are forced to ask whether the lack of clarity is deliberate.
>
> ---D. J. Bernstein
>
> --
> You received this message because you are subscribed to the Google Groups "pqc-forum" group.
> To unsubscribe from this group and stop receiving emails from it, send an email to pqc-forum+unsubscribe@list.nist.gov.
> To view this discussion on the web visit https://groups.google.com/a/list.nist.gov/d/msgid/pqc-forum/20220412155050.752181.qmail%40cr.yp.to.

--
You received this message because you are subscribed to the Google Groups "pqc-forum" group.
To unsubscribe from this group and stop receiving emails from it, send an email to pqc-forum+unsubscribe@list.nist.gov.
To view this discussion on the web visit https://groups.google.com/a/list.nist.gov/d/msgid/pqc-forum/CAEkod37Oiibp82Babfc0yzXCuhAuswiG2UVx4ehqeS0mymfTtw%40mail.gmail.com.

From: Leo Ducas <leo.ducas1@gmail.com> via pqc-forum@list.nist.gov
To: pqc-forum <pqc-forum@list.nist.gov>
CC: מצו״ב <matzov@idf.il>
Subject: [pqc-forum] Re: Improved Dual Lattice Attack
Date: Tuesday, April 12, 2022 02:30:49 PM ET

Dear MATZOV researchers, dear all,

thank you for sharing your work, and in particular bringing attention to the cost model for BDGL sieve, where, you claim a few bits of security can be shaved. I am looking at the estimates from [AGPS20, page 47], and found the following model for [BDGL16] list decoding:

insert_cost = filters * C (d,T2) * COST_IP(d) * log2(d)
query_cost = filters * C (d,T1) * COST_IP(d) * log2(d)

(the d parameter is not explicit for ip_cost, but I'll need to tweak it below).

I agree with you that this model is not adequate; the original algorithm from [BDGL16] should instead have cost:

Z = filters^(d/m)
insert_cost = Z * COST_IP(d/m) * m + m * Z * log(Z) * COST_COMPARE_SWAP + m * filters * C(d , T2) * COST_TREE_ITER
query_cost = Z * COST_IP(d/m) * m + m * Z * log(Z) * COST_COMPARE_SWAP + m * filters * C(d , T1) * COST_TREE_ITER

(Some discussion on the choice of m is necessary, I am delaying to PS for readability. TLDR: there are some overheads that have not ben accounted by the literature, and they get worse as m increases).

1/ Can you confirm that this is how you modeled its cost ? And if so, how did you choose m ? If not, then what ?
2/ Even better, would you be kind enough to provide you modified scripts for obtaining your conclusions ?

Best regards

- Leo Ducas

PS: On the choice of m

The choice of m we propose in BDGL is m = O(log d), though the explicit constant is hard to choose. Why not take very large m to thwart the first terms ? After all, one could take Z=2 my carefully choosing m, but doing so essentially sends us back to Hyperplane LSH, whose complexity is exponentially worse.

Choosing m = O(log d), our Theorem 5.1 in [BDGL16] shows that the loss compared to the idealized model is at most sub-exponential 2^{~O(\sqrt n)}. Unfortunately, this overhead has never been quantified concretely in the literature, and has essentially been ignored in the NIST estimates.

If forced to guess, I note that in practice [DSvW21] m=3 for d=120 seems to be the optimal trade-off between probaby loss and speed of list decoding, so a reasonable choice for d~400 might be m = 5 or 6.

[BDGL16] https://eprint.iacr.org/2015/1128
[AGPS20] https://eprint.iacr.org/2019/1161
[DSvW21] https://eprint.iacr.org/2021/141

Le lundi 4 avril 2022 à 18:38:57 UTC+2, מצו״ב a écrit :

> - בלמ"ס -
>
> Dear PQC researchers,
>
> The Center of Encryption and Information Security (MATZOV) of the IDF has conducted an internal audit of leading Post-Quantum cryptographic (PQC) schemes, focusing on the Learning With Errors and Rounding problems.
>
> After consultations with NIST over the last few months – we have decided to release the audit as a Technical Report available for public review.
>
> https://doi.org/10.5281/zenodo.6412487
>
> Our report presents several improvements to the dual lattice attack, which induce a noticeable reduction in the security estimation for Kyber, Saber and Dilithium, bringing them below the required threshold.
>
> The report does not intend to provide a complete analysis of all post-quantum candidates, nor to recommend usage of specific algorithms. Rather, this publication is meant to share advances in the cryptanalysis of lattices which we believe to be relevant to the academic research in the field.

> We acknowledge the remarkable work done by NIST in the process and its impact – creating interest in the post-quantum field and promoting new cryptographic schemes.
>
> A prudent approach for these schemes is recommended, as research in the field is constantly evolving and much remains unstudied. Therefore, as a contribution to the community, the report includes further research ideas which we deem interesting.
>
> MATZOV, IDF

--

From: Daniel Apon <dapon.crypto@gmail.com> via pqc-forum@list.nist.gov
To: מצו״ב <matzov@idf.il>
CC: Leo Ducas <leo.ducas1@gmail.com>, pqc-forum <pqc-forum@list.nist.gov>
Subject: Re: [pqc-forum] Re: Improved Dual Lattice Attack
Date: Tuesday, April 26, 2022 03:19:25 PM ET

Dear MATZOV,

I tend to agree with your comments on my questions. Thank you for sharing them.

Best regards,
--Daniel Apon

On Tue, Apr 26, 2022 at 3:06 PM מצו״ב <Matzov@idf.il> wrote:

> - בלמ"ס -
>
> Dear PQC researchers,
>
> Thank you for your comments.
>
> 1. Mr. Apon,
>
> Please note the analysis in the RAM model in the report is largely based on previous works, such as Albrecht et al (https://ia.cr/2019/1161) and more, as described in sections 6 and 7.
>
> The choice of this model enables proper comparison of the attack to other similar methods and to external results, and we expect its relative improvement to be similar in other computation and memory models. In fact, because the sieve is performed on smaller dimensional lattices than comparable attacks (the main source of improvement), then the memory requirements should only be smaller, and the comparative improvement should possibly be even more significant in other models that penalize large memory usage.
>
> Beyond the memory cost of the sieve, there is the addition of the FFT step, similar to the work presented by Guo and Johansson at AsiaCrypt 2021 (https://doi.org/10.1007/978-3-030-92068-5_2). This step also requires accesses to somewhat large memory, but its requirements are smaller than the sieve in most parameters, and even increasing its cost does not affect the model significantly. To give an example, we tried

assuming the FFT costs were larger by a factor of 1,000 and re-optimized the parameters, and the result was a change of less than a factor of 2 for the overall costs in most cases.

> However, more importantly: Do you find that any of your new algorithmic approaches have a concrete cost that would differ from a generic model-to-model analysis? (I hope to read through the work and answer "No!" but perhaps you have an opinion you could share now.)

The short answer is indeed "no", as our work does not present increased memory usage compared to other similar attacks published for lattice candidates.

2. Mr. Ducas,

> I agree with you that this model is not adequate; the original algorithm from [BDGL16] should instead have cost:

> Z = filters^(1/m)

> insert_cost = Z * COST_IP(d/m) * m + m * Z * log(Z) * COST_COMPARE_SWAP + m * filters * C(d , T2) * COST_TREE_ITER

> query_cost = Z * COST_IP(d/m) * m + m * Z * log(Z) * COST_COMPARE_SWAP + m * filters * C(d , T1) * COST_TREE_ITER

These formulas do reflect the cost of algorithm from [BDGL16]. However, in Section 6.2.1 we describe a different decoding algorithm, which utilizes a somewhat different iteration tree. In this algorithm, the subcode lists are relabeled in a way that allows for more efficient pruning. The formulas for the cost of this algorithm are

insert_cost = Z * COST_IP(d/m) * m + m * Z * log(Z) * COST_COMPARE_SWAP + filters * C(d , T2) * COST_TREE_ITER

query_cost = Z * COST_IP(d/m) * m + m * Z * log(Z) * COST_COMPARE_SWAP + filters * C(d , T1) * COST_TREE_ITER

The first two terms, which account for preprocessing, are the same as [BDGL16] (as the preprocessing is very similar). The cost of iterating over the tree, which is usually the significant part, is different. While the algorithm in [BDGL16] costs O(m) operations per obtained filter, the algorithm in our report costs O(1) operations per obtained filter.

> Note that this also means the choice of m has a far less significant effect on the runtime. For 400 < d < 1000, choosing m = 5 or 6 ensures that the cost of preprocessing is negligible relative to the cost of iteration.
>
> 3. We have received another comment from Mr. Ducas regarding the usage of the G6K model for estimations.
>
> We acknowledge that extrapolating the model to higher sieving dimensions was adviced against in (https://doi.org/10.1007/978-3-030-17656-3_25), and was added mainly for comparison to previous results of Guo and Johannson.
>
> Creating an adequate model for higher sieving dimensions other than the asymptotic model is a matter of further research.
>
> Our report has been updated accordingly: https://doi.org/10.5281/zenodo.6493704
>
> We thank Mr. Ducas for the clarification.
>
> We are thankful for all the comments, and will be happy to answer questions from either the community or NIST's team regarding the presented algorithmic improvements.
>
> Best regards,
>
> MATZOV
>
> ---
>
> **מאת:** pqc-forum@list.nist.gov <pqc-forum@list.nist.gov> בשם Leo Ducas <leo.ducas1@gmail.com>
> **נשלח:** יום שלישי 12 אפריל 2022 18:55
> **אל:** pqc-forum <pqc-forum@list.nist.gov>
> **עותק:** Leo Ducas <leo.ducas1@gmail.com>; מצו״ב <Matzov@idf.il>
> **נושא:** Re: Improved Dual Lattice Attack [pqc-forum]
>
> :Typo
> Z = filters^(d/m) --> Z = filters^(1/m)
>
> : Le mardi 12 avril 2022 à 20:29:54 UTC+2, Leo Ducas a écrit
>
>> ,Dear MATZOV researchers, dear all
>>
>> thank you for sharing your work, and in particular bringing attention to the cost model for

> > BDGL sieve, where, you claim a few bits of security can be shaved. I am looking at the estimates from [AGPS20, page 47], and found the following model for [BDGL16] list :decoding
> >
> > insert_cost = filters * C (d,T2) * COST_IP(d) * log2(d)
> > query_cost = filters * C (d,T1) * COST_IP(d) * log2(d)
> >
> > .(the d parameter is not explicit for ip_cost, but I'll need to tweak it below)
> >
> > I agree with you that this model is not adequate; the original algorithm from [BDGL16] :should instead have cost
> >
> > Z = filters^(d/m)
> > * insert_cost = Z * COST_IP(d/m) * m + m * Z * log(Z) * COST_COMPARE_SWAP + m * filters
> > C(d , T2) * COST_TREE_ITER
> > * query_cost = Z * COST_IP(d/m) * m + m * Z * log(Z) * COST_COMPARE_SWAP + m * filters
> > C(d , T1) * COST_TREE_ITER
> >
> > .Some discussion on the choice of m is necessary, I am delaying to PS for readability)
> > TLDR: there are some overheads that have not ben accounted by the literature, and they .(get worse as m increases
> >
> > Can you confirm that this is how you modeled its cost ? And if so, how did you choose /1 ? m ? If not, then what
> > Even better, would you be kind enough to provide you modified scripts for obtaining /2 ? your conclusions
> >
> > Best regards
> > Leo Ducas -
> >
> > PS: On the choice of m
> >
> > The choice of m we propose in BDGL is m = O(log d), though the explicit constant is hard to choose. Why not take very large m to thwart the first terms ? After all, one could take ,Z=2 my carefully choosing m, but doing so essentially sends us back to Hyperplane LSH .whose complexity is exponentially worse

>> Choosing m = O(log d), our Theorem 5.1 in [BDGL16] shows that the loss compared to the idealized model is at most sub-exponential 2^{~O(\sqrt n)}. Unfortunately, this overhead has never been quantified concretely in the literature, and has essentially been ignored in .the NIST estimates
>>
>> If forced to guess, I note that in practice [DSvW21] m=3 for d=120 seems to be the optimal trade-off between probabity loss and speed of list decoding, so a reasonable .choice for d~400 might be m = 5 or 6
>>
>> https://eprint.iacr.org/2015/1128 [BDGL16]
>> https://eprint.iacr.org/2019/1161 [AGPS20]
>> https://eprint.iacr.org/2021/141 [DSvW21]
>> : a écrit מצו״ב ,Le lundi 4 avril 2022 à 18:38:57 UTC+2
>>
>>> - בלמ"ס -
>>>
>>> Dear PQC researchers,
>>>
>>> The Center of Encryption and Information Security (MATZOV) of the IDF has conducted an internal audit of leading Post-Quantum cryptographic (PQC) schemes, focusing on the Learning With Errors and Rounding problems.
>>>
>>> After consultations with NIST over the last few months – we have decided to release the audit as a Technical Report available for public review.
>>>
>>> https://doi.org/10.5281/zenodo.6412487
>>>
>>> Our report presents several improvements to the dual lattice attack, which induce a noticeable reduction in the security estimation for Kyber, Saber and Dilithium, bringing them below the required threshold.
>>>
>>> The report does not intend to provide a complete analysis of all post-quantum candidates, nor to recommend usage of specific algorithms. Rather, this publication is meant to share advances in the cryptanalysis of lattices which we believe to be relevant to the academic research in the field.
>>>
>>> We acknowledge the remarkable work done by NIST in the process and its impact – creating interest in the post-quantum field and promoting new cryptographic schemes.
>>>
>>> A prudent approach for these schemes is recommended, as research in the field is constantly evolving and much remains unstudied. Therefore, as a contribution to the community, the report includes further research ideas which we deem interesting.
>>>
>>> MATZOV, IDF

> --
>
> "You received this message because you are subscribed to the Google Groups "pqc-forum

> .group
> -pqc To unsubscribe from this group and stop receiving emails from it, send an email to
> .forum+unsubscribe@list.nist.gov
> /https://groups.google.com/a/list.nist.gov/d/msgid To view this discussion on the web visit
> .pqc-forum/bceaaee6-50b0-45e7-b4f3-beba7c95043dn%40list.nist.gov
>
> --
> You received this message because you are subscribed to the Google Groups "pqc-forum" group.
> To unsubscribe from this group and stop receiving emails from it, send an email to pqc-forum+unsubscribe@list.nist.gov.
> To view this discussion on the web visit https://groups.google.com/a/list.nist.gov/d/msgid/pqc-forum/DM5PR14MB14045D3A0AB14C4D9831D352A1FB9%40DM5PR14MB1404.namprd14.prod.outlook.com.

--
You received this message because you are subscribed to the Google Groups "pqc-forum" group.
To unsubscribe from this group and stop receiving emails from it, send an email to pqc-forum+unsubscribe@list.nist.gov.
To view this discussion on the web visit https://groups.google.com/a/list.nist.gov/d/msgid/pqc-forum/CAPxHsSLfac0Vx9Np0334zP0aaW43%3DC2p6ERfSA_KA%3DQytBjfyw%40mail.gmail.com.

From: Martin R. Albrecht <martinralbrecht@googlemail.com> via pqc-forum <pqc-forum@list.nist.gov>
To: pqc-forum@list.nist.gov
Subject: Re: [pqc-forum] Improved Dual Lattice Attack
Date: Monday, May 02, 2022 02:46:10 PM ET

Hi all,

We looked at the changes to sieving estimates in the "Report on the Security of LWE: Improved Dual Lattice Attack" by MATZOV available at https://doi.org/10.5281/zenodo.6412487 and agree

1. We made a mistake in costing visiting each node as an inner product. These costs can be amortised by preprocessing as already described in the original BDGL paper.

2. We agree that the enumeration described in Section 6.2.1 of the MATZOV paper reduces the cost per solution from O(m) to constant.

We have updated out scripts and cost estimates here:

https://github.com/jschanck/eprint-2019-1161/pull/3

We stress that our scripts and estimates were developed to assess the effect of quantum computing on sieving. For this reason, i.e. the seemingly prohibitive cost of QRAM, we explicitly ignore all memory access costs. Accounting for such memory access costs would affect other trade-offs.

We also note that our estimates are slightly lower than those reported by MATZOV. We assume this is down to some choice of magic constants somwhere. Thus, we would appreciate if MATZOV could publish their estimation scripts to allow us to reproduce and compare to them.

In addition, the lattice estimator <https://github.com/malb/lattice-estimator/> has been updated with these new costs

https://github.com/malb/lattice-estimator/pull/35

Note that these corrections and improvements to sieving are not restricted to the dual attack but also apply to the primal attack.

Best,
John, Eamonn and Martin

PS: We expect that the lattice estimator will be updated shortly with the dual attack model from the above mentioned report, too.

PPS: We thank Léo Ducas for helpful discussions on list decoding in BDGL.

On Mon, Apr 04 2022, מצו״ב wrote:
> - בלמ"ס -
>
>
> Dear PQC researchers,
>
>
>
> The Center of Encryption and Information Security (MATZOV) of the IDF has
> conducted an internal audit of leading Post-Quantum cryptographic (PQC) schemes,
> focusing on the Learning With Errors and Rounding problems.
>
> After consultations with NIST over the last few months – we have decided to
> release the audit as a Technical Report available for public review.
>
> https://doi.org/10.5281/zenodo.6412487<https://doi.org/10.5281/zenodo.6412487>
>
>
>
> Our report presents several improvements to the dual lattice attack, which
> induce a noticeable reduction in the security estimation for Kyber, Saber and
> Dilithium, bringing them below the required threshold.
>
> The report does not intend to provide a complete analysis of all post-quantum
> candidates, nor to recommend usage of specific algorithms. Rather, this
> publication is meant to share advances in the cryptanalysis of lattices which we
> believe to be relevant to the academic research in the field.
>

> We acknowledge the remarkable work done by NIST in the process and its impact – creating interest in the post-quantum field and promoting new cryptographic schemes.
>
> A prudent approach for these schemes is recommended, as research in the field is constantly evolving and much remains unstudied. Therefore, as a contribution to the community, the report includes further research ideas which we deem interesting.
>
> MATZOV, IDF

--

_pgp: https://gcc02.safelinks.protection.outlook.com/?url=https%3A%2F%2Fkeybase.io%2Fmartinralbrecht&amp;data=05%7C01%7Cyi-kai.liu%40nist.gov%7C58f851ab2fad434af44d08da2c6c0432%7C2ab5d82fd8fa4797a93e054655c61dec%7C1%7C0%7C637871139706116111%7CUnknown%7CTWFpbGZsb3d8eyJWIjoiMC4wLjAwMDAiLCJQIjoiV2luMzIiLCJBTiI6Ik1haWwiLCJXVCI6Mn0%3D%7C3000%7C%7C%7C&amp;sdata=a58cwhIHPGADqL%2BiHf%2BA2r%2BYbPhkYT%2BVKgW8lbLjrnA%3D&amp;reserved=0
_www: https://gcc02.safelinks.protection.outlook.com/?url=https%3A%2F%2Fmalb.io%2F&amp;data=05%7C01%7Cyi-kai.liu%40nist.gov%7C58f851ab2fad434af44d08da2c6c0432%7C2ab5d82fd8fa4797a93e054655c61dec%7C1%7C0%7C637871139706116111%7CUnknown%7CTWFpbGZsb3d8eyJWIjoiMC4wLjAwMDAiLCJQIjoiV2luMzIiLCJBTiI6Ik1haWwiLCJXVCI6Mn0%3D%7C3000%7C%7C%7C&amp;sdata=EjqmifstJcPMxYiPg3IBlEcNiy2NOVxhDVb7CUzReHI%3D&amp;reserved=0
_prn: he/him or they/them

--

You received this message because you are subscribed to the Google Groups "pqc-forum" group.

From: Leo Ducas <leo.ducas1@gmail.com> via pqc-forum@list.nist.gov
To: pqc-forum <pqc-forum@list.nist.gov>
CC: מצו״ב <matzov@idf.il>
Subject: [pqc-forum] Re: Improved Dual Lattice Attack
Date: Monday, May 02, 2022 09:18:53 PM ET

Dear MATZOV researchers, dear all,

Another brief remark. I note that you are using the GSA to model BKZ, but are using progressive-BKZ to cost it. The GSA requires many tours at the same blocksize to reach. Progressive BKZ lags behind a bit. At the relevant dimension this is about 9 blocksize, so about 2.5 bits extra on the cost of BKZ. See data and script at https://github.com/lducas/simu-vs-gsa/blob/main/simu-vs-gsa.sage

Its probably less impactful over the whole attack by re-balancing the various steps.

Its small, but we seem to be at that level of details. I'm pointing it out for completeness, and for comparing apple to apple (our costing of the primal attack uses the simulator).
Best regards
-- Léo

d=1024

beta, b0_simul, b0_gsa

...

378 73.277 68.534

379 72.721 68.010

380 72.173 **67.493**

381 71.631 66.982

382 71.095 66.477

383 70.567 65.977

384 70.045 65.483

385 69.529 64.995

386	69.020	64.513
387	68.517	64.036
388	68.020	63.564
389	**67.530**	63.098
390	67.045	62.637
391	66.566	62.181

Le lundi 4 avril 2022 à 18:38:57 UTC+2, מצו״ב a écrit :

> - בלמ"ס -
>
> Dear PQC researchers,
>
> The Center of Encryption and Information Security (MATZOV) of the IDF has conducted an internal audit of leading Post-Quantum cryptographic (PQC) schemes, focusing on the Learning With Errors and Rounding problems.
>
> After consultations with NIST over the last few months – we have decided to release the audit as a Technical Report available for public review.
>
> https://doi.org/10.5281/zenodo.6412487
>
> Our report presents several improvements to the dual lattice attack, which induce a noticeable reduction in the security estimation for Kyber, Saber and Dilithium, bringing them below the required threshold.
>
> The report does not intend to provide a complete analysis of all post-quantum candidates, nor to recommend usage of specific algorithms. Rather, this publication is meant to share advances in the cryptanalysis of lattices which we believe to be relevant to the academic research in the field.
>
> We acknowledge the remarkable work done by NIST in the process and its impact – creating interest in the post-quantum field and promoting new cryptographic schemes.
>
> A prudent approach for these schemes is recommended, as research in the field is constantly evolving and much remains unstudied. Therefore, as a contribution to the community, the report includes further research ideas which we deem interesting.
>
> MATZOV, IDF

From: Christopher J Peikert <cpeikert@alum.mit.edu> via pqc-forum@list.nist.gov
To: מצו״ב <matzov@idf.il>
CC: pqc-forum@list.nist.gov
Subject: Re: [pqc-forum] Improved Dual Lattice Attack
Date: Friday, May 06, 2022 11:33:57 AM ET

Dear all: on the question of the memory cost of this attack, I'd like to highlight some concrete numbers.
Tables 3--6 of the MATZOV report show that for (say) Kyber-512, the attack uses sieving in dimensions close to 380 (+-3, depending on choice of models).

How much memory does this need? A fairly precise estimate is at least 2^90 bits for pair-sieving (which the MATZOV report uses for its runtime analysis), and at least 2^85 bits for triple-sieving (which is slower than pair-sieving).

(I derived these numbers using the algorithms' models and real-world experiments, which closely align. For example, the data in Table 1 of https://eprint.iacr.org/2021/141.pdf nicely fits the triple-sieving model of 2^{(0.1887+o(1))*d}. The pair-sieving model has 0.2075 in place of 0.1887.)

Sieving algorithms are highly memory-bound, so these large memory requirements would impose a significant real-world cost that is not counted in the RAM-model analysis (and would also affect the overall optimization of parameters). Of course, quantifying this precisely is an important research question.

Sincerely yours in cryptography,

Chris

From: Martin R. Albrecht <martinralbrecht@googlemail.com> via pqc-forum <pqc-forum@list.nist.gov>
To: pqc-forum@list.nist.gov
Subject: Re: [pqc-forum] Improved Dual Lattice Attack
Date: Saturday, May 07, 2022 11:42:01 AM ET

Hi Dan,

I assume you're referring to the difference between "usvp" and "bdd" in the estimator, e.g. here: https://github.com/malb/lattice-estimator/

This corresponds to Q7 of

https://gcc02.safelinks.protection.outlook.com/?url=https%3A%2F%2Fpq-crystals.org%2Fkyber%2Fdata%2Fkyber-specification-round3-20210804.pdf&amp;data=05%7C01%7Candrew.regenscheid%40nist.gov%7Ce8ea6b5cf2e34a485d0b08da30401f97%7C2ab5d82fd8fa4797a93e054655c61dec%7C1%7C0%7C637875349210183605%7CUnknown%7CTWFpbGZsb3d8eyJWIjoiMC4wLjAwMDAiLCJQIjoiV2luMzIiLCJBTiI6Ik1haWwiLCJXVCI6Mn0%3D%7C3000%7C%7C%7C&amp;sdata=2fjeqVaxXWcjwrxLL8vov2x0%2Bdy9aFn0hBESldZ7BaA%3D&amp;reserved=0

It is worth reiterating that this improvement has a time-memory trade-off flavour to it since the final sieving step is over a larger dimension than the BKZ sieving steps. However, the cost of memory access is not costed (in the estimator nor usually the wider literature)

To me this supports a point that you've been making for many years: we should cost memory access, too, to get a better understanding of the true costs of these attacks.

I should also stress that the Kyber spec uses the CN11 simulator to predict the shape after lattice reduction while the estimator uses the GSA by default. The former is more precise, the latter is faster.

Here's the effect of that difference:

```
sage: LWE.primal_usvp(Kyber512, red_shape_model="CN11")
rop: ≈2^146.8, red: ≈2^146.8, δ: 1.003869, β: 417, d: 994, tag: usvp
sage: LWE.primal_bdd(Kyber512, red_shape_model="CN11")
rop: ≈2^142.2, red: ≈2^141.1, svp: ≈2^141.3, β: 396, η: 430, d: 1013, tag: bdd
```

```
sage: LWE.primal_usvp(Kyber512, red_shape_model="GSA")
rop: ≈2^143.8, red: ≈2^143.8, δ: 1.003941, β: 406, d: 998, tag: usvp
sage: LWE.primal_bdd(Kyber512, red_shape_model="GSA")
rop: ≈2^140.3, red: ≈2^139.7, svp: ≈2^138.8, β: 391, η: 421, d: 1013, tag: bdd
```

Skimming through “5.3 Approximations, overheads, and foreseeable improvements” of

https://gcc02.safelinks.protection.outlook.com/?url=https%3A%2F%2Fpq-crystals.org%2Fkyber%2Fdata%2Fkyber-specification-round3-20210804.pdf&amp;data=05%7C01%7Candrew.regenscheid%40nist.gov%7Ce8ea6b5cf2e34a485d0b08da30401f97%7C2ab5d82fd8fa4797a93e054655c61dec%7C1%7C0%7C637875349210183605%7CUnknown%7CTWFpbGZsb3d8eyJWIjoiMC4wLjAwMDAiLCJQIjoiV2luMzIiLCJBTiI6Ik1haWwiLCJXVCI6Mn0%3D%7C3000%7C%7C%7C&amp;sdata=2fjeqVaxXWcjwrxLL8vov2x0%2Bdy9aFn0hBESldZ7BaA%3D&amp;reserved=0

nothing stood out as already covered by the estimator.

Cheers,
Martin

On Sat, May 07 2022, D. J. Bernstein wrote:

> [[PGP Signed Part:Undecided]]
> 'Martin R. Albrecht' via pqc-forum writes:
>> Note that these corrections and improvements to sieving are not
>> restricted to the dual attack but also apply to the primal attack.
>
> Am I correctly understanding that your latest "lattice-estimator" cost
> estimate for breaking Kyber-512 is 2^140.3 bit operations: i.e., several
> times fewer bit operations than AES-128 key search, and 2400x fewer bit
> operations than the 2^151.5 from the round-3 Kyber documentation?
>
> The round-3 Kyber documentation also mentions various other speedups
> that could save "a factor of up to 2^16" without any new attack ideas.
> Are any of these speedups covered by your cost estimate? If so, is there
> a chart making clear which speedups are covered and which aren't? Thanks
> in advance for any clarification you can provide.

>
> ---D. J. Bernstein

--

_pgp: https://gcc02.safelinks.protection.outlook.com/?url=https%3A%2F%2Fkeybase.io%2Fmartinralbrecht&amp;data=05%7C01%7Candrew.regenscheid%40nist.gov%7Ce8ea6b5cf2e34a485d0b08da30401f97%7C2ab5d82fd8fa4797a93e054655c61dec%7C1%7C0%7C637875349210183605%7CUnknown%7CTWFpbGZsb3d8eyJWIjoiMC4wLjAwMDAiLCJQIjoiV2luMzIiLCJBTiI6Ik1haWwiLCJXVCI6Mn0%3D%7C3000%7C%7C%7C&amp;sdata=oZ%2FNK6C%2BrAYlVGfRPtc2lmR%2B%2BM7b28OUhjP1khX7gI8%3D&amp;reserved=0
_www: https://gcc02.safelinks.protection.outlook.com/?url=https%3A%2F%2Fmalb.io%2F&amp;data=05%7C01%7Candrew.regenscheid%40nist.gov%7Ce8ea6b5cf2e34a485d0b08da30401f97%7C2ab5d82fd8fa4797a93e054655c61dec%7C1%7C0%7C637875349210183605%7CUnknown%7CTWFpbGZsb3d8eyJWIjoiMC4wLjAwMDAiLCJQIjoiV2luMzIiLCJBTiI6Ik1haWwiLCJXVCI6Mn0%3D%7C3000%7C%7C%7C&amp;sdata=57TOyzSICM8c5j15Lq6BP64V1JtkrLxPe%2BD1KtKPXHM%3D&amp;reserved=0
_prn: he/him or they/them

--

From: Leo Ducas <leo.ducas1@gmail.com> via pqc-forum@list.nist.gov
To: pqc-forum <pqc-forum@list.nist.gov>
CC: pqc-...@list.nist.gov <pqc-forum@list.nist.gov>
Subject: Re: [pqc-forum] Improved Dual Lattice Attack
Date: Sunday, May 08, 2022 06:15:07 AM ET

Dear Dan, dear all,

> So, just to make sure I'm clear about the conclusion: Your current
> estimate is 2^142.2 bit operations to break Kyber-512, i.e., 600x fewer
> bit operations than what the Kyber-512 round-3 documentation said
> (namely 2^151.5)? And this doesn't account for the known speedups that
> the documentation says could save "a factor of up to 2^16"?

It *does* account for some of them. I am unsure how you misread those
lines of Martin:

> This corresponds to Q7 of
>
> https://pq-crystals.org/kyber/data/kyber-specification-round3-20210804.pdf

Q7 being one of the 8 open questions leading to the potential 2^16 speed-up you are referring too. This Q7 accounted for 2^8 of this potential speed-up. It is in fact a 2^4.4 speed-up. The rest comes from the mis-costing of BDGL that Matzof pointed too and that we have been discussing earlier in this thread and correcting in the estimator.

I'm also pointing to the fact that this list of open questions does not only list potential speed-ups, but also unaccounted potential overheads.

Best regards.

Léo

--
You received this message because you are subscribed to the Google Groups "pqc-forum" group.
To unsubscribe from this group and stop receiving emails from it, send an email to pqc-

forum+unsubscribe@list.nist.gov.
To view this discussion on the web visit https://groups.google.com/a/list.nist.gov/d/msgid/pqc-forum/5bbeabef-0feb-4646-b0d6-b30123ac16e4n%40list.nist.gov.

From: D. J. Bernstein <djb@cr.yp.to> via pqc-forum@list.nist.gov
To: pqc-forum <pqc-forum@list.nist.gov>
Subject: Re: [pqc-forum] Improved Dual Lattice Attack
Date: Sunday, May 08, 2022 09:32:41 AM ET
Attachments: smime.p7m

Leo Ducas writes:

> It is in fact a 2^4.4 speed-up.

So, to make sure I'm clear about your position regarding the overall status of Kyber-512:

* Compared to the round-3 Kyber documentation estimating 2^151.5 "gates" to break Kyber-512, your current estimate after the latest attack paper is 600x fewer "gates", i.e., 2^142.2? Is this also the official Kyber position?

* Furthermore, within the known speedups that the documentation says could save "a factor of up to 2^16", you're saying that 2^8 could apply to this 2^142.2, i.e., that known Kyber-512 attacks could cost just 2^134.2 "gates", well below the AES-128 attack cost?

I understand that you're also pointing to "potential overheads", but is the Kyber team now claiming on this basis that known attacks require "2^143 classical gates"?

The estimate of 2^151.5 "gates" also appears to be the basis for NIST's 2020 claim that "Kyber clearly meets the security categories defined in the CFP". Is the Kyber team continuing to claim these categories?

> It *does* account for some of them. I am unsure how you misread those
> lines of Martin:

Hmmm. For some reason you (1) omit the lines that I actually quoted from Martin right above my question, (2) actively substitute other lines, and (3) on this basis claim a misreading. The lines that I actually quoted fully justify the clarification question that I asked.

If Martin erred in writing "nothing" rather than "nothing except X" for
some specific X, then this error is something to attribute to him,
certainly not to the followup clarification question.

———D. J. Bernstein

From: D. J. Bernstein <djb@cr.yp.to> via pqc-forum@list.nist.gov
To: pqc-forum@list.nist.gov
Subject: Re: [pqc-forum] Improved Dual Lattice Attack
Date: Sunday, May 08, 2022 10:06:22 AM ET
Attachments: smime.p7m

Christopher J Peikert writes:

> How much memory does this need? A fairly precise estimate is at least
> 2^90 bits for pair-sieving (which the MATZOV report uses for its runtime
> analysis), and at least 2^85 bits for triple-sieving (which is slower than
> pair-sieving).

The numbers here are similar to numbers posted in the same thread a month ago, but somehow they seem to have mutated from being rough estimates a month ago into something that sounds reasonably confident ("fairly precise ... at least").

Would you describe 2^90 and 2^85 as "barriers"? Will there be an admission of error if these attacks against Kyber-512 are shown to fit into less memory? Or is the word "fairly" intended to allow subsequent wiggle room, eliminating falsifiability? Thanks in advance for clarifying the status of your claim.

> (I derived these numbers using the algorithms' models and real-world
> experiments, which closely align. For example, the data in Table 1 of
> https://eprint.iacr.org/2021/141.pdf nicely fits the triple-sieving model
> of 2^{(0.1887+o(1))*d}. The pair-sieving model has 0.2075 in place of
> 0.1887.)

This doesn't make sense: "o(1)" by definition says nothing about any concrete size, so the claims of fit and alignment must be based on something else. Can you please spell out your calculations, to support public assessment of the risks of the 90 and 85 being overestimates?

---D. J. Bernstein

--

From: Martin R. Albrecht <martinralbrecht@googlemail.com> via pqc-forum <pqc-forum@list.nist.gov>
To: pqc-forum@list.nist.gov
Subject: Re: [pqc-forum] Improved Dual Lattice Attack
Date: Sunday, May 08, 2022 10:14:16 AM ET

On Sun, May 08 2022, D. J. Bernstein wrote:

> Hmmm. For some reason you (1) omit the lines that I actually quoted from
> Martin right above my question, (2) actively substitute other lines, and
> (3) on this basis claim a misreading. The lines that I actually quoted
> fully justify the clarification question that I asked.
>
> If Martin erred in writing "nothing" rather than "nothing except X" for
> some specific X, then this error is something to attribute to him,
> certainly not to the followup clarification question.

For the avoidance of doubt, I did mean “nothing except X” where “X” is the thing (Q7) I had mentioned as being in that list.

--

_pgp: https://gcc02.safelinks.protection.outlook.com/?url=https%3A%2F%2Fkeybase.io%2Fmartinralbrecht&amp;data=05%7C01%7Candrew.regenscheid%40nist.gov%7C6e0d8ddbb09a43308df408da30fd07de%7C2ab5d82fd8fa4797a93e054655c61dec%7C1%7C0%7C637876160560549344%7CUnknown%7CTWFpbGZsb3d8eyJWIjoiMC4wLjAwMDAiLCJQIjoiV2luMzIiLCJBTiI6Ik1haWwiLCJXVCI6Mn0%3D%7C3000%7C%7C%7C&amp;sdata=5×4Hfy2jirKYGHx8gV2d2dFIN3BB3HGzfABDKCeBZ3k%3D&amp;reserved=0
_www: https://gcc02.safelinks.protection.outlook.com/?url=https%3A%2F%2Fmalb.io%2F&amp;data=05%7C01%7Candrew.regenscheid%40nist.gov%7C6e0d8ddbb09a43308df408da30fd07de%7C2ab5d82fd8fa4797a93e054655c61dec%7C1%7C0%7C637876160560549344%7CUnknown%7CTWFpbGZsb3d8eyJWIjoiMC4wLjAwMDAiLCJQIjoiV2luMzIiLCJBTiI6Ik1haWwiLCJXVCI6Mn0%3D%7C3000%7C%7C%7C&amp;sdata=m%2FZL%2FqkMrGw6QWh9vYxrGCY8pswfSXsOVpO%2B89oy2%2FM%3D&amp;reserved=0
_prn: he/him or they/them

--

From: 赵运磊 <ylzhao@fudan.edu.cn> via pqc-forum <pqc-forum@list.nist.gov>
To: D. J. Bernstein <djb@cr.yp.to>
CC: pqc-forum@list.nist.gov
Subject: On the possibility of achieving NIST security goals with the recent advances of dual attacksRe: Re: [pqc-forum] Improved Dual Lattice Attack
Date: Thursday, May 12, 2022 06:01:48 AM ET

Dear Prof. Bernstein and deal all in PQC community:

The recent advances of dual attacks might bring the worry the possibility of achieving the security goals set by NIST for lattice-based KEM schemes, particularly on dimension of 512. Our recent work shows it may still be possible, but with optimized constructions.

In our recent work: https://gcc02.safelinks.protection.outlook.com/?url=https%3A%2F%2Farxiv.org%2Fabs%2F2205.05413&amp;data=05%7C01%7Candrew.regenscheid%40nist.gov%7C3c65bf19266b45bcaafe08da33fe6b84%7C2ab5d82fd8fa4797a93e054655c61dec%7C1%7C0%7C637879465085560322%7CUnknown%7CTWFpbGZsb3d8eyJWIjoiMC4wLjAwMDAiLCJQIjoiV2luMzIiLCJBTiI6Ik1haWwiLCJXVCI6Mn0%3D%7C3000%7C%7C%7C&amp;sdata=pdjb98yqyZzgysEbJb7ZbFhpwpim3QFKwEEnYclWXuY%3D&amp;reserved=0 CNTR-512 can have 2^{170.4} gate complexity at 2^{107.4} memory complexity with error probability 2^{-94}. It is tested by running the script provided by Kyber. Assuming each secret key will not be used to decrypt for more than 2^{94} times in its lifttime, this parameter set may achieve security level II (2^{143} gates required by NIST) even if with the recent advances on dual attacks. The details are given in Appendix E in the mentioned paper.

It also appears that the technique used by CNTR may also be applied to NTRU-prime. As it is a new work, we sincerely look forward to your kind comments and critiques to further improve it.

All my best
Yours sincerely
Yunlei

> -----原始邮件-----
> 发件人: "D. J. Bernstein" <djb@cr.yp.to>
> 发送时间: 2022-05-09 01:29:56 (星期一)
> 收件人: pqc-forum@list.nist.gov
> 抄送:
> 主题: Re: [pqc-forum] Improved Dual Lattice Attack
>
> 'Martin R. Albrecht' via pqc-forum writes:
> > For the avoidance of doubt, I did mean “nothing except X” where “X” is
> > the thing (Q7) I had mentioned as being in that list.
>
> Um, the message mentioned Q7 as corresponding to a usvp-bdd difference,
> not as being in the 5.3 list (a list that the message cited separately).
> Even the weaker notion that the message _hinted_ at Q7 being in 5.3
> seems impossible to reconcile with the plain meaning of the word
> "nothing" in the self-contained sentence that I quoted before:
>
> > > > Skimming through “5.3 Approximations, overheads, and foreseeable improvements” of
> > > > https://gcc02.safelinks.protection.outlook.com/?url=https%3A%2F%2Fpq-crystals.org%2Fkyber%2Fdata%2Fkyber-specification-round3-20210804.pdf&amp;data=05%7C01%7Candrew.regenscheid%40nist.gov%7C3c65bf19266b45bcaafe08da33fe6b84%7C2ab5d82fd8fa4797a93e054655c61dec%7C1%7C0%7C637879465085560322%7CUnknown%7CTWFpbGZsb3d8eyJWIjoiMC4wLjAwMDAiLCJQIjoiV2luMzIiLCJBTiI6Ik1haWwiLCJXVCI6Mn0%3D%7C3000%7C%7C%7C&amp;sdata=GX6NcnVtR85zD8MgqkEKODhAJ59BRQd0GXeAVGOuFzg%3D&amp;reserved=0
> > > > nothing stood out as already covered by the estimator.
>
> Rewriting "nothing" as "nothing except the Q7 mentioned above, which is
> in 5.3" isn't resolving ambiguity; it's retroactively switching to a
> different statement with different consequences.
>
> It's amazing that, when I quote a questionable sentence and politely ask
> for confirmation of what the sentence is communicating, I'm accused of
> misreading---by someone who omits the quote I gave and substitutes a

> different quote!---and after two further messages there's still no
> admission of error from the actual source of the error.
>
> ---D. J. Bernstein
>
> --
> You received this message because you are subscribed to the Google Groups "pqc-forum" group.
> To unsubscribe from this group and stop receiving emails from it, send an email to pqc-forum+unsubscribe@list.nist.gov.
> To view this discussion on the web visit https://groups.google.com/a/list.nist.gov/d/msgid/pqc-forum/20220508172956.158683.qmail%40cr.yp.to.

--

You received this message because you are subscribed to the Google Groups "pqc-forum" group.
To unsubscribe from this group and stop receiving emails from it, send an email to pqc-forum+unsubscribe@list.nist.gov.
To view this discussion on the web visit https://groups.google.com/a/list.nist.gov/d/msgid/pqc-forum/31828a6c.b11c.180b7b6bd26.Coremail.ylzhao%40fudan.edu.cn.

From: 赵运磊 <ylzhao@fudan.edu.cn> via pqc-forum <pqc-forum@list.nist.gov>
To: D. J. Bernstein <djb@cr.yp.to>
CC: pqc-forum@list.nist.gov
Subject: Re: Re: On the possibility of achieving NIST security goals with the recent advances of dual attacksRe: Re: [pqc-forum] Improved Dual Lattice Attack
Date: Thursday, May 12, 2022 09:42:48 AM ET

Dear Prof. Bernstein:

Thanks for your question.

Indeed, Kyber is covered by our patents (not only the two patents mentioned in the KCL proposal, but also more patent afterforwards). It can be clearly seen from the following two works:

https://gcc02.safelinks.protection.outlook.com/?url=https%3A%2F%2Farxiv.org%2Fabs%2F2109.02893&amp;data=05%7C01%7Candrew.regenscheid%40nist.gov%7C4320b7f38b124e26046508da341d4b70%7C2ab5d82fd8fa4797a93e054655c61dec%7C1%7C0%7C637879597686664904%7CUnknown%7CTWFpbGZsb3d8eyJWIjoiMC4wLjAwMDAiLCJQIjoiV2luMzIiLCJBTiI6Ik1haWwiLCJXVCI6Mn0%3D%7C3000%7C%7C%7C&amp;sdata=dIssyKjk%2BlAfVOXg27Lg4GkkdGi7bVmZjVJPtDBbjBQ%3D&amp;reserved=0

https://gcc02.safelinks.protection.outlook.com/?url=https%3A%2F%2Farxiv.org%2Fabs%2F1611.06150&amp;data=05%7C01%7Candrew.regenscheid%40nist.gov%7C4320b7f38b124e26046508da341d4b70%7C2ab5d82fd8fa4797a93e054655c61dec%7C1%7C0%7C637879597686664904%7CUnknown%7CTWFpbGZsb3d8eyJWIjoiMC4wLjAwMDAiLCJQIjoiV2luMzIiLCJBTiI6Ik1haWwiLCJXVCI6Mn0%3D%7C3000%7C%7C%7C&amp;sdata=WSsiW7BaGGEek8GduTvbP%2B1cnpSqc53DCxHlfaMaOjQ%3D&amp;reserved=0

From these works, it is clear that if we interpret Kyber within our AKCN mechanism proposed in 1611.06150 in 2016 (also in the two patents mentioned in KCL), i.e., if we focus on the con/rec mechanisms of Kyber and AKCN-LWE, the con part of Kyber and AKCN are the same, but the rec part of Kyber is less efficient. To be frank, after we posted 1611.06150, we sent an email to inform some authors of "NewHope without reconciliation", but we didn't receive response until we notice the paper of "NewHope without reconciliation" .

Kyber and Saber face more patent threats than our patents as discussed in the past in the forum. NTRU has no patent issue, but the current version of NTRU and NTRU-prime might not in its best forms. CTRU and CNTR could eliminate most of the existing patent threats against LWE/LWR-based KEM. CTRU and CNTR may combine the advantages of both NTRU and LWE/LWR. Note also that CNTR and CTRU have the same KeyGen and Decryption processes, which means that we can easily switch between NTRU-RLWE/RLWR.

Yours sincerely
Yunlei

> ———原始邮件———
> 发件人: "D. J. Bernstein" <djb@cr.yp.to>
> 发送时间: 2022-05-12 20:55:14 (星期四)
> 收件人: pqc-forum@list.nist.gov
> 抄送:
> 主题: Re: On the possibility of achieving NIST security goals with the recent advances of dual attacksRe: Re: [pqc-forum] Improved Dual Lattice Attack
>
> '赵运磊' via pqc-forum writes:
> > The recent advances of dual attacks might bring the worry the
> > possibility of achieving the security goals set by NIST for
> > lattice-based KEM schemes, particularly on dimension of 512. Our
> > recent work shows it may still be possible, but with optimized
> > constructions.
>
> Can you please comment on what's covered by your patents related to this
> work? I noticed that your patents
>
>   https://gcc02.safelinks.protection.outlook.com/?url=https%3A%2F%2Fpatents.google.com%2Fpatent%2FCN107566121A%2Fen&amp;data=05%7C01%7Candrew.regenscheid%40nist.gov%7C4320b7f38b124e26046508da341d4b70%7C2ab5d82fd8fa4797a93e054655c61dec%7C1%7C0%7C637879597686664904%7CUnknown%7CTWFpbGZsb3d8eyJWIjoiMC4wLjAwMDAiLCJQIjoiV2luMzIiLCJBTiI6Ik1haWwiLCJXVCI6Mn0%3D%7C3000%7C%7C%7C&amp;sdata=ySsuo17S%2BPlxaOsj4mB38kchAr9UruAj5tgg1mRxiao%3D&amp;reserved=0

>   https://gcc02.safelinks.protection.outlook.com/?url=https%3A%2F%2Fpatents.google.com%2Fpatent%2FCN108173643B%2Fen&amp;data=05%7C01%7Candrew.regenscheid%40nist.gov%7C4320b7f38b124e26046508da341d4b70%7C2ab5d82fd8fa4797a93e054655c61dec%7C1%7C0%7C637879597686664904%7CUnknown%7CTWFpbGZsb3d8eyJWIjoiMC4wLjAwMDAiLCJQIjoiV2luMzIiLCJBTiI6Ik1haWwiLCJXVCI6Mn0%3D%7C3000%7C%7C%7C&amp;sdata=ZyLAOBz4kDqXG9WoT2b066IggMyE1NCIpw12hmRBbpM%3D&amp;reserved=0
>
> were reported in the KCL/OKCN/AKCN/CNKE submission, which is very
> similar to "NewHope without reconciliation". The patents were filed a
> month before "NewHope without reconciliation" was published, and I
> haven't seen any analysis of the patent coverage.
>
> It would be useful to see public assurances as to your company's
> position regarding usage of "NewHope without reconciliation" and its
> variants, such as Kyber, SABER, and your latest proposals.
>
> ———D. J. Bernstein
>
> --
> You received this message because you are subscribed to the Google Groups "pqc-forum" group.
> To unsubscribe from this group and stop receiving emails from it, send an email to pqc-forum+unsubscribe@list.nist.gov.
> To view this discussion on the web visit https://groups.google.com/a/list.nist.gov/d/msgid/pqc-forum/20220512125514.219585.qmail%40cr.yp.to.

--

You received this message because you are subscribed to the Google Groups "pqc-forum" group.

To unsubscribe from this group and stop receiving emails from it, send an email to pqc-forum+unsubscribe@list.nist.gov.

To view this discussion on the web visit https://groups.google.com/a/list.nist.gov/d/msgid/pqc-forum/36e11573.891e.180b881bae4.Coremail.ylzhao%40fudan.edu.cn.

From: 赵运磊 <ylzhao@fudan.edu.cn> via pqc-forum <pqc-forum@list.nist.gov>
To: D. J. Bernstein <djb@cr.yp.to>
CC: pqc-forum@list.nist.gov
Subject: Re: Re: On the possibility of achieving NIST security goals with the recent advances of dual attacksRe: Re: [pqc-forum] Improved Dual Lattice Attack
Date: Thursday, May 12, 2022 09:50:35 AM ET

With respect to the patents, we ever mentioned in the KCL submission we would like to give up all the patents for using our proposals. We hold the patents only for protection. This position applies to all of our proposals.

All my best

Yunlei

```
> ———原始邮件———
> 发件人: "D. J. Bernstein" <djb@cr.yp.to>
> 发送时间: 2022-05-12 20:55:14 (星期四)
> 收件人: pqc-forum@list.nist.gov
> 抄送:
> 主题: Re: On the possibility of achieving NIST security goals with the recent
advances of dual attacksRe: Re: [pqc-forum] Improved Dual Lattice Attack
>
> '赵运磊' via pqc-forum writes:
> > The recent advances of dual attacks might bring the worry the
> > possibility of achieving the security goals set by NIST for
> > lattice-based KEM schemes, particularly on dimension of 512. Our
> > recent work shows it may still be possible, but with optimized
> > constructions.
>
> Can you please comment on what's covered by your patents related to this
> work? I noticed that your patents
>
```

>   https://gcc02.safelinks.protection.outlook.com/?url=https%3A%2F%2Fpatents.google.com%2Fpatent%2FCN107566121A%2Fen&amp;data=05%7C01%7Candrew.regenscheid%40nist.gov%7C3d6652e208b94dedf1cc08da341e62d2%7C2ab5d82fd8fa4797a93e054655c61dec%7C1%7C0%7C637879602354032345%7CUnknown%7CTWFpbGZsb3d8eyJWIjoiMC4wLjAwMDAiLCJQIjoiV2luMzIiLCJBTiI6Ik1haWwiLCJXVCI6Mn0%3D%7C3000%7C%7C%7C&amp;sdata=tTgc478xwDvVhCaZfbyLmXwj3h92NRCTIkbBOIPwgDQ%3D&amp;reserved=0
>   https://gcc02.safelinks.protection.outlook.com/?url=https%3A%2F%2Fpatents.google.com%2Fpatent%2FCN108173643B%2Fen&amp;data=05%7C01%7Candrew.regenscheid%40nist.gov%7C3d6652e208b94dedf1cc08da341e62d2%7C2ab5d82fd8fa4797a93e054655c61dec%7C1%7C0%7C637879602354032345%7CUnknown%7CTWFpbGZsb3d8eyJWIjoiMC4wLjAwMDAiLCJQIjoiV2luMzIiLCJBTiI6Ik1haWwiLCJXVCI6Mn0%3D%7C3000%7C%7C%7C&amp;sdata=tvTv7l3Nt5SydIyr5mTbAig%2FFI1uaxcrg6t%2B0SC6Xwk%3D&amp;reserved=0
>
> were reported in the KCL/OKCN/AKCN/CNKE submission, which is very
> similar to "NewHope without reconciliation". The patents were filed a
> month before "NewHope without reconciliation" was published, and I
> haven't seen any analysis of the patent coverage.
>
> It would be useful to see public assurances as to your company's
> position regarding usage of "NewHope without reconciliation" and its
> variants, such as Kyber, SABER, and your latest proposals.
>
> ---D. J. Bernstein
>
> --
> You received this message because you are subscribed to the Google Groups "pqc-forum" group.
> To unsubscribe from this group and stop receiving emails from it, send an email to pqc-forum+unsubscribe@list.nist.gov.
> To view this discussion on the web visit https://groups.google.com/a/list.nist.gov/d/msgid/pqc-forum/20220512125514.219585.qmail%40cr.yp.to.

--

From: 赵运磊 <ylzhao@fudan.edu.cn> via pqc-forum <pqc-forum@list.nist.gov>
To: D. J. Bernstein <djb@cr.yp.to>
CC: pqc-forum@list.nist.gov
Subject: [pqc-forum] More clarifications about patents
Date: Thursday, May 12, 2022 12:03:48 PM ET

Dear Prof. Bernstein and dear all in PQC community:

Here, we would like to make the patent issues clearer.

For all the KEM schemes based on LWE/MLWE/LWR/MLWR, they actually have the same scheme structures. The key differences can be well interpreted w.r.t what are referred to as the Con/Rec mechanism in https://gcc02.safelinks.protection.outlook.com/?url=https%3A%2F%2Farxiv.org%2Fabs%2F1611.06150&amp;data=05%7C01%7Candrew.regenscheid%40nist.gov%7C3cb0946bc8fb4e5257ab08da3430fe1d%7C2ab5d82fd8fa4797a93e054655c61dec%7C1%7C0%7C637879682279510328%7CUnknown%7CTWFpbGZsb3d8eyJWIjoiMC4wLjAwMDAiLCJQIjoiV2luMzIiLCJBTiI6Ik1haWwiLCJXVCI6Mn0%3D%7C3000%7C%7C%7C&amp;sdata=MdmO07Iz1HLyoZnRaKfVQqBvvVvL58WSgnQF6qwc%2FfE%3D&amp;reserved=0 (as well as in our KCL proposal). Every KEM based on LWE/MLWE/LWR/MLWR implies a Con/Rec mechanism. The difference between LWE\MLWE-based KEM and LWR\MLWR-based KEM is that Con/Rec in LWE\MLWE-based is w.r.t. the modulus $q$, but Con/Rec in LWR\MLWR-based is w.r.t the compression parameter $p$. The Con/Rec implied by Frodo is just one previously proposed, but it is not optimal (as a consequence Frodo does not violate our patents). To the best of our knowledge, AKCN in https://gcc02.safelinks.protection.outlook.com/?url=https%3A%2F%2Farxiv.org%2Fabs%2F1611.06150&amp;data=05%7C01%7Candrew.regenscheid%40nist.gov%7C3cb0946bc8fb4e5257ab08da3430fe1d%7C2ab5d82fd8fa4797a93e054655c61dec%7C1%7C0%7C637879682279510328%7CUnknown%7CTWFpbGZsb3d8eyJWIjoiMC4wLjAwMDAiLCJQIjoiV2luMzIiLCJBTiI6Ik1haWwiLCJXVCI6Mn0%3D%7C3000%7C%7C%7C&amp;sdata=MdmO07Iz1HLyoZnRaKfVQqBvvVvL58WSgnQF6qwc%2FfE%3D&amp;reserved=0 (as well as in our KCL proposal) is the first one that is proved to be optimal. The Con/Rec mechanisms in Kyber and Saber are also optimal in correcting errors, but Rec in Kyber involves an unnecessary rounding operation which makes it less efficient and more error-prone (the Con of AKCN and that of Kyber are the same). Con/Rec of AKCN-MLWE and Saber are essentially the same, but w.r.t. the compression parameter $p$ in Saber. These differences can be clearly noted from the mentioned two arXiv reports:

https://gcc02.safelinks.protection.outlook.com/?url=https%3A%2F%2Farxiv.org%2Fabs%2F2109.02893&amp;data=05%7C01%7Candrew.regenscheid%40nist.gov%7C3cb0946bc8fb4e5257ab08da3430fe1d%7C2ab5d82fd8fa4797a93e054655c61dec%7C1%7C0%7C637879682279510328%7CUnknown%7CTWFpbGZsb3d8eyJWIjoiMC4wLjAwMDAiLCJQIjoiV2luMzIiLCJBTiI6Ik1haWwiLCJXVCI6Mn0%3D%7C3000%7C%7C%7C&amp;sdata=NCyPjQEssc6uc0%2B%2BOUQ%2FAE7Ub2%2B01VPePwrSYHosD0w%3D&amp;reserved=0

https://gcc02.safelinks.protection.outlook.com/?url=https%3A%2F%2Farxiv.org%2Fabs%2F1611.06150&amp;data=05%7C01%7Candrew.regenscheid%40nist.gov%7C3cb0946bc8fb4e5257ab08da3430fe1d%7C2ab5d82fd8fa4797a93e054655c61dec%7C1%7C0%7C637879682279510328%7CUnknown%7CTWFpbGZsb3d8eyJWIjoiMC4wLjAwMDAiLCJQIjoiV2luMzIiLCJBTiI6Ik1haWwiLCJXVCI6Mn0%3D%7C3000%7C%7C%7C&amp;sdata=MdmO07Iz1HLyoZnRaKfVQqBvvVvL58WSgnQF6qwc%2FfE%3D&amp;reserved=0

Finally, we would like to stress again we hold all the patents only for protection against credit (not for economic reasons). We hope the above clarifications could make the situation clearer.

All my best
Yunlei

> ———原始邮件———
> 发件人: "D. J. Bernstein" <djb@cr.yp.to>
> 发送时间: 2022-05-12 20:55:14 (星期四)
> 收件人: pqc-forum@list.nist.gov
> 抄送:
> 主题: Re: On the possibility of achieving NIST security goals with the recent advances of dual attacksRe: Re: [pqc-forum] Improved Dual Lattice Attack
>
> '赵运磊' via pqc-forum writes:
> > The recent advances of dual attacks might bring the worry the
> > possibility of achieving the security goals set by NIST for
> > lattice-based KEM schemes, particularly on dimension of 512. Our
> > recent work shows it may still be possible, but with optimized
> > constructions.
>

```
> Can you please comment on what's covered by your patents related to this
> work? I noticed that your patents
>
>    https://gcc02.safelinks.protection.outlook.com/?
url=https%3A%2F%2Fpatents.google.com%2Fpatent%2FCN107566121A%2Fen&amp;data=05%7C01%7C
andrew.regenscheid%40nist.gov%7C3cb0946bc8fb4e5257ab08da3430fe1d%7C2ab5d82fd8fa4797a9
3e054655c61dec%7C1%7C0%7C637879682279510328%7CUnknown%7CTWFpbGZsb3d8eyJWIjoiMC4wLjAwM
DAiLCJQIjoiV2luMzIiLCJBTiI6Ik1haWwiLCJXVCI6Mn0%3D%7C3000%7C%7C%7C&amp;sdata=xta8zJ6d8
fsJwGkmQpM47ExNGVHaT76DTcweGXKe5ck%3D&amp;reserved=0
>    https://gcc02.safelinks.protection.outlook.com/?
url=https%3A%2F%2Fpatents.google.com%2Fpatent%2FCN108173643B%2Fen&amp;data=05%7C01%7C
andrew.regenscheid%40nist.gov%7C3cb0946bc8fb4e5257ab08da3430fe1d%7C2ab5d82fd8fa4797a9
3e054655c61dec%7C1%7C0%7C637879682279510328%7CUnknown%7CTWFpbGZsb3d8eyJWIjoiMC4wLjAwM
DAiLCJQIjoiV2luMzIiLCJBTiI6Ik1haWwiLCJXVCI6Mn0%3D%7C3000%7C%7C%7C&amp;sdata=mIAdFWLaj
n9%2FBHcg%2BG5s0jwdd2fxJRH9KbDJ3r%2F4DeQ%3D&amp;reserved=0
>
> were reported in the KCL/OKCN/AKCN/CNKE submission, which is very
> similar to "NewHope without reconciliation". The patents were filed a
> month before "NewHope without reconciliation" was published, and I
> haven't seen any analysis of the patent coverage.
>
> It would be useful to see public assurances as to your company's
> position regarding usage of "NewHope without reconciliation" and its
> variants, such as Kyber, SABER, and your latest proposals.
>
> ---D. J. Bernstein
>
> --
> You received this message because you are subscribed to the Google Groups "pqc-
forum" group.
> To unsubscribe from this group and stop receiving emails from it, send an email to
pqc-forum+unsubscribe@list.nist.gov.
> To view this discussion on the web visit https://groups.google.com/a/list.nist.gov/
d/msgid/pqc-forum/20220512125514.219585.qmail%40cr.yp.to.
```

From: Blumenthal, Uri - 0553 - MITLL <uri@ll.mit.edu> via pqc-forum@list.nist.gov
To: 赵运磊 <ylzhao@fudan.edu.cn>
CC: pqc-forum@list.nist.gov
Subject: Re: [pqc-forum] More clarifications about patents
Date: Wednesday, May 18, 2022 01:53:41 PM ET
Attachments: smime.p7m

Dear Yunlei,

Thank you for your proposals - they are very interesting. I have a few questions.

1. Are there reference implementations of CTRU, OSKR, and others? Optimized implementations?

2. Would your proposed algorithms, such as OSKR, still be potential subjects to the same patent claims that, e.g., Kyber is dealing with?

3. You (University, Company, etc.) have some patents covering CTRU, OSKR, and other algorithms that you proposed. It is nice that your email stated: "we would like to give up all the patents for using our proposals. We hold the patents only for protection." Not being a lawyer, I cannot evaluate whether that statement is sufficient from legal point of view. Would the patent(s) holders be willing to make a more "official" statement to that extent?

Please feel free to answer on this mailing list, or privately - as you prefer.

Thank you!

--

V/R,

Uri

On 5/12/22, 12:03, "'赵运磊' via pqc-forum" <pqc-forum@list.nist.gov> wrote:

> Dear Prof. Bernstein and dear all in PQC community:
>
> Here, we would like to make the patent issues clearer.

For all the KEM schemes based on LWE/MLWE/LWR/MLWR, they actually have the same scheme structures. The key differences can be well interpreted w.r.t what are referred to as the Con/Rec mechanism in

https://arxiv.org/abs/1611.06150 (as well as in our KCL proposal). Every KEM based on LWE/MLWE/LWR/MLWR implies a Con/Rec mechanism. The difference between LWE\MLWE-based KEM and LWR\MLWR-based KEM is that Con/Rec in LWE\MLWE-based is w.r.t. the modulus $q$, but Con/Rec in LWR\MLWR-based is w.r.t the compression parameter $p$. The Con/Rec implied by Frodo is just one previously proposed, but it is not optimal (as a consequence Frodo does not violate our patents). To the best of our knowledge, AKCN in https://arxiv.org/abs/1611.06150 (as well as in our KCL proposal) is the first one that is proved to be optimal. The Con/Rec mechanisms in Kyber and Saber are also optimal in correcting errors, but Rec in Kyber involves an unnecessary rounding operation which makes it less efficient and more error-prone (the Con of AKCN and that of Kyber are the same). Con/Rec of AKCN-MLWE and Saber are essentially the same, but w.r.t. the compression parameter $p$ in Saber. These differences can be clearly noted from the mentioned two arXiv reports:

https://arxiv.org/abs/2109.02893

https://arxiv.org/abs/1611.06150

Finally, we would like to stress again we hold all the patents only for protection against credit (not for economic reasons). We hope the above clarifications could make the situation clearer.

All my best
Yunlei

> ———原始邮件———
>
> 发件人: "D. J. Bernstein" <djb@cr.yp.to>
>
> 发送时间: 2022-05-12 20:55:14 (星期四)
>
> 收件人: pqc-forum@list.nist.gov
>
> 抄送:
>
> 主题: Re: On the possibility of achieving NIST security goals with the recent advances of dual attacksRe: Re: [pqc-forum] Improved Dual Lattice Attack
>
>

> '赵运磊' via pqc-forum writes:
> > The recent advances of dual attacks might bring the worry the
> > possibility of achieving the security goals set by NIST for
> > lattice-based KEM schemes, particularly on dimension of 512. Our
> > recent work shows it may still be possible, but with optimized
> > constructions.
>
> Can you please comment on what's covered by your patents related to this
> work? I noticed that your patents
>
>   https://patents.google.com/patent/CN107566121A/en
>   https://patents.google.com/patent/CN108173643B/en
>
> were reported in the KCL/OKCN/AKCN/CNKE submission, which is very
> similar to "NewHope without reconciliation". The patents were filed a
> month before "NewHope without reconciliation" was published, and I
> haven't seen any analysis of the patent coverage.
>
> It would be useful to see public assurances as to your company's
> position regarding usage of "NewHope without reconciliation" and its
> variants, such as Kyber, SABER, and your latest proposals.
>
> ———D. J. Bernstein
>
> --
> You received this message because you are subscribed to the Google Groups "pqc-forum" group.
> To unsubscribe from this group and stop receiving emails from it, send an email to pqc-forum+unsubscribe@list.nist.gov.
> To view this discussion on the web visit https://groups.google.com/a/list.nist.gov/d/msgid/pqc-forum/20220512125514.219585.qmail%40cr.yp.to.

--

You received this message because you are subscribed to the Google Groups "pqc-forum" group.

To unsubscribe from this group and stop receiving emails from it, send an email to pqc-forum+unsubscribe@list.nist.gov.

To view this discussion on the web visit https://groups.google.com/a/list.nist.gov/d/msgid/pqc-forum/5a91d094.8ba5.180b902887e.Coremail.ylzhao%40fudan.edu.cn.

--

You received this message because you are subscribed to the Google Groups "pqc-forum" group.

To unsubscribe from this group and stop receiving emails from it, send an email to pqc-forum+unsubscribe@list.nist.gov.

To view this discussion on the web visit https://groups.google.com/a/list.nist.gov/d/msgid/pqc-forum/740E846A-A9A9-4CC3-BD7E-8C5FF3DD4F3E%40ll.mit.edu.

From: 赵运磊 <ylzhao@fudan.edu.cn> via pqc-forum <pqc-forum@list.nist.gov>
To: Blumenthal, Uri - 0553 - MITLL <uri@ll.mit.edu>
CC: pqc-forum@list.nist.gov
Subject: Re: Re: [pqc-forum] More clarifications about patents
Date: Wednesday, May 18, 2022 10:05:05 PM ET

Dear Uri:

Thanks for your interest in our CTRU and OSKR works.

(1)For OSKR, we have implementations in C, AVX2, ARM Cortex-M4. For CTRU, we currently only have reference implementations in C, but we are optimizing the implementations in C and AVX, which should be given in the near future. About the benchmark results of CTRU, we stress that the results are true. But the benchmark is performed at the student’s laptop computer that was bought about 4 years earlier. The benchmark results may show that NTRU-HRSS may be more influenced with a relatively older computer or platform. If needed, we can send the implementation codes in a private mail.

(2)For patents, yes, all our proposals are under patent protection now. As mentioned, we hold patents mainly for protection to be against discredits. If needed, I will do my best to coordinate towards a positive outputs for freely using our proposal.

All my best
Sincerely yours
Yunlei

> ———原始邮件———
>
> 发件人: "Blumenthal, Uri - 0553 - MITLL" <uri@ll.mit.edu>
>
> 发送时间: 2022-05-19 01:53:21 (星期四)
>
> 收件人: "赵运磊" <ylzhao@fudan.edu.cn>
>
> 抄送: "pqc-forum@list.nist.gov" <pqc-forum@list.nist.gov>
>
> 主题: Re: [pqc-forum] More clarifications about patents
>
> Dear Yunlei,
>
> Thank you for your proposals - they are very interesting. I have a few questions.

>

> 1. Are there reference implementations of CTRU, OSKR, and others? Optimized implementations?

>

> 2. Would your proposed algorithms, such as OSKR, still be potential subjects to the same patent claims that, e.g., Kyber is dealing with?

>

> 3. You (University, Company, etc.) have some patents covering CTRU, OSKR, and other algorithms that you proposed. It is nice that your email stated: "we would like to give up all the patents for using our proposals. We hold the patents only for protection." Not being a lawyer, I cannot evaluate whether that statement is sufficient from legal point of view. Would the patent(s) holders be willing to make a more "official" statement to that extent?

>

> Please feel free to answer on this mailing list, or privately - as you prefer.

>

> Thank you!

> --

> V/R,

> Uri

>

>

> On 5/12/22, 12:03, "'赵运磊' via pqc-forum" <pqc-forum@list.nist.gov> wrote:

>

>   Dear Prof. Bernstein and dear all in PQC community:

>

>   Here, we would like to make the patent issues clearer.

>

>   For all the KEM schemes based on LWE/MLWE/LWR/MLWR, they actually have the same scheme structures. The key differences can be well interpreted w.r.t what are referred to as the Con/Rec mechanism in

>   https://gcc02.safelinks.protection.outlook.com/?url=https%3A%2F%2Farxiv.org%2Fabs%2F1611.06150&amp;data=05%7C01%7Candrew.regenscheid%40nist.gov%7C259d37b9af8b4cd98d6708da393bfd0f%7C2ab5d82fd8fa4797a93e054655c61dec%7C1%7C0%7C637885227052437744%7CUnknown%7CTWFpbGZsb3d8eyJWIjoiMC4wLjAwMDAiLCJQIjoiV2luMzIiLCJBTiI6Ik1haWwiLCJXVCI6Mn0%3D%7C3000%7C%7C%7C&amp;sdata=gmyyzUWiPClJPoub9gBLddeIQpBclcIbS%2FSPJsDH4dE%3D&amp;reserved=0 (as well as in our KCL proposal). Every KEM based on LWE/MLWE/LWR/MLWR implies a Con/Rec mechanism. The difference between LWE\MLWE-based KEM and LWR\MLWR-based KEM is that Con/Rec in LWE\MLWE-based is w.r.t. the modulus $q$, but Con/Rec in LWR\MLWR-based is w.r.t the compression parameter $p$. The Con/Rec implied by Frodo is just one previously proposed, but it is not optimal (as a consequence Frodo does not violate our patents). To the best of our knowledge, AKCN in https://gcc02.safelinks.protection.outlook.com/?url=https%3A%2F%2Farxiv.org%2Fabs%2F1611.06150&amp;data=05%7C01%7Candrew.regenscheid%40nist.gov%7C259d37b9af8b4cd98d6708da393bfd0f%7C2ab5d82fd8fa4797a93e054655c61dec%7C1%7C0%7C637885227052593370%7CUnknown%7CTWFpbGZsb3d8eyJWIjoiMC4wLjAwMDAiLCJQIjoiV2luMzIiLCJBTiI6Ik1haWwiLCJXVCI6Mn0%3D%7C3000%7C%7C%7C&amp;sdata=UAZ6o7dF5sjsD3HuGbCv9kDXREsdc8GsXZeLrCP%2BEqY%3D&amp;reserved=0 (as well as in our KCL proposal) is the first one that is proved to be optimal. The Con/Rec mechanisms in Kyber and Saber are also optimal in correcting errors, but Rec in Kyber involves an unnecessary rounding operation which makes it less efficient and more error-prone (the Con of AKCN and that of Kyber are the same). Con/Rec of AKCN-MLWE and Saber are essentially the same, but w.r.t. the compression parameter $p$ in Saber. These differences can be clearly noted from the mentioned two arXiv reports:

>

>   https://gcc02.safelinks.protection.outlook.com/?url=https%3A%2F%2Farxiv.org%2Fabs%2F2109.02893&amp;data=05%7C01%7Candrew.regenscheid%40nist.gov%7C259d37b9af8b4cd98d6708da393bfd0f%7C2ab5d82fd8fa4797a93e054655c61dec%7C1%7C0%7C637885227052593370%7CUnknown%7CTWFpbGZsb3d8eyJWIjoiMC4wLjAwMDAiLCJQIjoiV2luMzIiLCJBTiI6Ik1haWwiLCJXVCI6Mn0%3D%7C3000%7C%7C%7C&amp;sdata=HQunkfpEpmom4Frv59q1pSYVbyzhalAQDSTPlXe27xs%3D&amp;reserved=0

>

>   https://gcc02.safelinks.protection.outlook.com/?url=https%3A%2F%2Farxiv.org%2Fabs%2F1611.06150&amp;data=05%7C01%7Candrew.regenscheid%40nist.gov%7C259d37b9af8b4cd98d6708da393bfd0f%7C2ab5d82fd8fa4797a93e054655c61dec%7C1%7C0%7C637885227052593370%7CUnknown%7CTWFpbGZsb3d8eyJWIjoiMC4wLjAwMDAiLCJQIjoiV2luMzIiLCJBTiI6Ik1haWwiLCJXVCI6Mn0%3D%7C3000%7C%7C%7C&amp;sdata=UAZ6o7dF5sjsD3HuGbCv9kDXREsdc8GsXZeLrCP%2BEqY%3D&amp;reserved=0

>

```
>     Finally, we would like to stress again we hold all the patents only for protection against credit (not for economic reasons). We hope the above clarifications could make the situation clearer.
>
>     All my best
>     Yunlei
>
>
>
>     > ————原始邮件————
>     > 发件人: "D. J. Bernstein" <djb@cr.yp.to>
>     > 发送时间: 2022-05-12 20:55:14 (星期四)
>     > 收件人: pqc-forum@list.nist.gov
>     > 抄送:
>     > 主题: Re: On the possibility of achieving NIST security goals with the recent advances of dual attacksRe: Re: [pqc-forum] Improved Dual Lattice Attack
>     >
>     > '赵运磊' via pqc-forum writes:
>     > > The recent advances of dual attacks might bring the worry the
>     > > possibility of achieving the security goals set by NIST for
>     > > lattice-based KEM schemes, particularly on dimension of 512. Our
>     > > recent work shows it may still be possible, but with optimized
>     > > constructions.
>     >
>     > Can you please comment on what's covered by your patents related to this
>     > work? I noticed that your patents
>     >
>     >    https://gcc02.safelinks.protection.outlook.com/?url=https%3A%2F%2Fpatents.google.com%2Fpatent%2FCN107566121A%2Fen&amp;data=05%7C01%7Candrew.regenscheid%40nist.gov%7C259d37b9af8b4cd98d6708da393bfd0f%7C2ab5d82fd8fa4797a93e054655c61dec%7C1%7C0%7C637885227052593370%7CUnknown%7CTWFpbGZsb3d8eyJWIjoiMC4wLjAwMDAiLCJQIjoiV2luMzIiLCJBTiI6Ik1haWwiLCJXVCI6Mn0%3D%7C3000%7C%7C%7C&amp;sdata=5AGdprKMaX6c%2BauD6aqLzwbUXbvBd%2F7syoEjDwm2Ins%3D&amp;reserved=0
>     >    https://gcc02.safelinks.protection.outlook.com/?url=https%3A%2F%2Fpatents.google.com%2Fpatent%2FCN108173643B%2Fen&amp;data=05%7C01%7Candrew.regenscheid%40nist.gov%7C259d37b9af8b4cd98d6708da393bfd0f%7C2ab5d82fd8fa4797a93e054655c61dec%7C1%7C0%7C637885227052593370%7CUnknown%7CTWFpbGZsb3d8eyJWIjoiMC4wLjAwMDAiLCJQIjoiV2luMzIiLCJBTiI6Ik1haWwiLCJXVCI6Mn0%3D%7C3000%7C%7C%7C&amp;sdata=xhjVHWZY4u7miJWhGta4d%2Brjv1MLUWkg1UmfTmGq6Gg%3D&amp;reserved=0
>     >
>     > were reported in the KCL/OKCN/AKCN/CNKE submission, which is very
>     > similar to "NewHope without reconciliation". The patents were filed a
>     > month before "NewHope without reconciliation" was published, and I
>     > haven't seen any analysis of the patent coverage.
>     >
>     > It would be useful to see public assurances as to your company's
>     > position regarding usage of "NewHope without reconciliation" and its
>     > variants, such as Kyber, SABER, and your latest proposals.
>     >
>     > ---D. J. Bernstein
>     >
>     > --
>     > You received this message because you are subscribed to the Google Groups "pqc-forum" group.
>     > To unsubscribe from this group and stop receiving emails from it, send an email to pqc-forum+unsubscribe@list.nist.gov.
>     > To view this discussion on the web visit https://groups.google.com/a/list.nist.gov/d/msgid/pqc-forum/20220512125514.219585.qmail%40cr.yp.to.
>
>
>
>
>
>     --
>     You received this message because you are subscribed to the Google Groups "pqc-forum" group.
>     To unsubscribe from this group and stop receiving emails from it, send an email to pqc-forum+unsubscribe@list.nist.gov.
>     To view this discussion on the web visit https://groups.google.com/a/list.nist.gov/d/msgid/pqc-forum/5a91d094.8ba5.180b902887e.Coremail.ylzhao%40fudan.edu.cn.
```

>
> --
> You received this message because you are subscribed to the Google Groups "pqc-forum" group.
> To unsubscribe from this group and stop receiving emails from it, send an email to pqc-forum+unsubscribe@list.nist.gov.
> To view this discussion on the web visit https://groups.google.com/a/list.nist.gov/d/msgid/pqc-forum/740E846A-A9A9-4CC3-BD7E-8C5FF3DD4F3E%40ll.mit.edu.

--

You received this message because you are subscribed to the Google Groups "pqc-forum" group.
To unsubscribe from this group and stop receiving emails from it, send an email to pqc-forum+unsubscribe@list.nist.gov.
To view this discussion on the web visit https://groups.google.com/a/list.nist.gov/d/msgid/pqc-forum/2d0f4351.70e6.180da0f8e96.Coremail.ylzhao%40fudan.edu.cn.