Response to Recent Paper by Ward Beullens

The Rainbow Team

December 22, 2020

Abstract

We would like to acknowledge the new attacks by Ward Beullens [1] as being sound and his results as basically correct. We carefully analyzed the new attacks and refined the analysis in these new attacks. In addition, we tried to find all possible enhancements to these attacks.

In terms of security analysis, we present a new practical security model by taking into consideration of the memory access cost. Under this view, we find that our parameter sets I, III, and V as proposed to NIST Round 3 still fit the corresponding NIST security levels in being as hard to cryptanalyze as AES-128, -192, and -256 respectively.

Fundamentally, the Rectangular MinRank attacks are still variants of known exponential attacks. While [1] neatly summarizes and clarifies the design of Rainbow and its results present a new and different perspective, the new modelling of MinRank developed in [2] by Bardet *et al* probably served to take as many bits of security away from Rainbow as the fact that there is a newer and better attack based on MinRank. We may say that as a result of [1], we are actually more confident that Rainbow will stay secure than before.

1 Synopsis

We acknowledge the new attacks by Ward Beullens [1] as being legitimate and his results as basically correct. Indeed, following the announcement, we actually worked with Mr. Beullens and improved his paper somewhat.

In the following we will therefore evaluate Beullens' New (Rectangular) MinRank attack following the Support Minors Modeling methodology of Bardet et al [2], and his new Intersection attack following the Bipartite XL methodology of Perlner and Smith-Tone [3].

For the Intersection Attack, we assume that Beullens' conjectured methodology to evaluate the bi-degrees of operation is valid as opposed to previous evaluations which simply throws away bilinear equations that lead to

Round	Instance	Intersection		New MinRank		Target
		mults	cost	mults	cost	cost
Second	Ia $(32, 32, 32, \mathbb{F}_{16})$	$2^{116.8}$	$2^{144.3}$	$2^{120.2}$	$2^{149.5}$	2^{143}
	IIIc $(68, 36, 36, \mathbb{F}_{256})$	$2^{405.0}$	$2^{441.5}$	$2^{145.8}$	$2^{183.6}$	2^{207}
	Vc $(92, 48, 48, \mathbb{F}_{256})$	$2^{541.0}$	$2^{587.8}$	$2^{184.8}$	$2^{232.6}$	2^{272}
Third	I $(36, 32, 32, \mathbb{F}_{16})$	$2^{134.3}$	$2^{162.1}$	$2^{122.4}$	$2^{152.3}$	2^{143}
	III $(68, 32, 48, \mathbb{F}_{256})$	$2^{205.4}$	$2^{248.3}$	$2^{171.8}$	$2^{216.2}$	2^{207}
	V $(96, 36, 64, \mathbb{F}_{256})$	$2^{254.5}$	$2^{309.9}$	$2^{219.6}$	$2^{276.2}$	2^{272}

Table 1: Intersection and New MinRank Attack Complexities vs. Rainbow

bidegree-(2,1) syzygies. For the New (Rectangular) MinRank attack, we take into account some other possibilities not covered in [1] (see below).

Where we differ from [1] is that in our cost model we are taking into consideration of the memory access. Under this view, we find that our parameter sets I, III, and V as proposed to NIST Round 3 has the security levels as given by Table 1, and still fit the corresponding NIST security levels in being as hard to cryptanalyze as AES-128, -192, and -256 respectively.

1.1 Cost Analysis

Other Round 3 NIST PQC Candidates have submitted instances where the analyzed cost in accepted models are lower than the NIST requirements, yet argued that they meet those security levels. E.g., the Crystals-Kyber team, which has Core-SVP security levels [7] of 118, 182, and 256 bits[1] for the Kyber level 1,3,5 instances (respectively), instead presumes security $2^{151.5}$, $2^{215.1}$, and $2^{287.3}$ (respectively) "gates" [5].

We note that the Crystals-Kyber team doesn't commit to $2^{151.5}$ gates, but instead say that there is an error range of ± 16 bits, and that their range is between 2^{135} and 2^{167} "gates". They further argue that even if it is 2^{135} it is okay because memory access (which they do not analyze) will be large and push it above 2^{143}.

For another example, the Crystals-Dilithium team, which has Core-SVP security [7] of 2^{123}, 2^{182}, and 2^{252} for their level 2,3,5 instances respectively, consider them to take 2^{159}, 2^{217}, and 2^{285} "gates" [6]. Note that the Dilithium briefly surveys the Kyber analysis (citing from [5]), says that the analysis is the basis of [6, Table 4], and briefly mentions that the ± 16 also applies (so that their level 3 instance could be below 2^{207} if not for memory costs).

[1]This number has been disputed as not according to the Core-SVP methodology of Albrecht et al[7], but we follow the Crystals team's numbers.

We concur with the Crystals team that the memory cost can be significant (even though they did not so state outright, and did not quantify it). We will present our cost model below, in which Beullens' attacks notwithstanding, Rainbow-I, III, and V still have the required security levels.

2 Multiplication Costs

Again, where we differ from Beullens is that instead of assuming that multiplication takes $2(k^2+k)$ "gates", where $k = \lg q$ is the bit-length of the field, we take memory access costs into account. Citing Bernstein et al [4]:

> Update for round 3: As in round 2, we report "free" security estimates that disregard the cost of memory (such as "Core-SVP"), and "real" security estimates that account for the cost of memory. For round 2, we did not try to pin down constant factors in the cost of memory. For round 3, we estimate the cost of each access to a bit within N bits of memory as the cost of $\sqrt{N}/2^5$ "bit operations". Here $1/2^5$ arises from comparing Intel's energy figures from [69]:
>
> - Intel reported an energy cost of 6.4 pJ at 22nm for a double-precision floating-point multiplication. This is roughly $2^{-11.3}\times$ #bitops pJ, since a multiplication uses roughly 2^{14} bit operations.
> - As noted above, Intel reported an energy cost of 11.20 pJ "per 5 mm" to move 8 bytes at 22nm. Roughly 2^{30} bits of DRAM at 22nm fit in a 5mm$\times$ 5mm square; if moving 64 bits 5mm takes 11.20 pJ then moving 1 bit the full 5mm horizontally plus the full 5mm vertically takes roughly 0.35 pJ, i.e., $2^{-16.5} \times \sqrt{2^{30}}$ pJ. 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: ...

What we will take as the cost of multiplication is the bit-length of the information that has to be sent to and from memory during that operation, times the square root of the number of bits that needs to be randomly (non-sequentially) accessed divided by 2^5, as a rough equivalent in "gates". In other words, if the vector length in the Wiedemann algorithm that dominates the runtime of both the Intersection Attack and the New (rectangular)

MinRank attack is V, then the cost of multiplication is $\lg V\sqrt{kV}/2^5$ if we have to make an equivalent cost in "gates".

Note that this V is not the matrix size, since we simplistically assume that we can recreate the matrix on the fly from the original equations (this underestimates the complexity). We have also not considered losses that must always result when we parallelize (this also underestimates the complexity).

As a result, we get Table 1. We conclude from this analysis that although our numbers are lower than 2^{128}, 2^{192}, and 2^{256} in multiplications for our level 1,3,5 parameter sets proposed for Round 3, our practical security is greater than what NIST requires.

3 Further Notes on Rainbow Cryptanalysis

3.1 Notes on Security Levels

Again without detracting from the basic correctness of the Beullens attacks, one may note that Beullens' evaluation and ours differ by a bit here and there. For example, he evaluates the intersection attack against Rainbow-Ia $(32, 32, 32, \mathbb{F}_{16})$ as taking 2^{123} "gates". At 40 "gates" per multiplication, that is off by one bit from our $2^{116.8}$. This seems to be a case where [1] assumes an upper bound for the density, while we (as did Smith-Tone and Perlner in [3]) assume that the attacker takes as many of the less dense equations as possible. In the case of the new MinRank attack against Rainbow-I $(36, 32, 32, \mathbb{F}_{16})$, [1] has 2^{127} which is again off by about one bit from our evaluation, which may have been caused by a rounding error.

3.2 A Recap, Prelude to New Modes of Attack

In our analysis, we also consider several possible extensions of the attacks proposed by [1]. To introduce the new attacks, first we recap Beullens' Rectangular MinRank attack. We refer the reader to [1, Fig. 2]. There exists subspaces $W \subset \mathbb{F}_q^m$ and $O_2 \subset O_1 \subset \mathbb{F}_q^n$ such that

$$\mathcal{P}'(\mathbf{x}, \cdot)(O_2) \subset W, \mathcal{P}(O_1) \subset W, \mathcal{P}(O_2) = \{0\}.$$

where $\mathcal{P}'(\mathbf{x}, \mathbf{y}) = \mathcal{P}(\mathbf{x}+\mathbf{y}) - \mathcal{P}(\mathbf{x}) - \mathcal{P}(\mathbf{y})$ is the polar form of $\mathcal{P}$, the rainbow public map. [1] introduce a new MinRank attack that exploits the property that for $\mathbf{y} \in O_2$, we have that $\mathcal{P}'(\mathbf{x}, \mathbf{y}) \in W, \forall \mathbf{x}$. Let $\mathbf{e}_1, \ldots, \mathbf{e}_n$ be the

standard basis for $\mathbb{F}_q^n$. For a vector $\mathbf{x} \in \mathbb{F}_q^n$ Define the matrix

$$L_{\mathbf{x}} = \begin{bmatrix} \mathcal{P}'(\mathbf{e}_1, \mathbf{x}) \\ \vdots \\ \mathcal{P}'(\mathbf{e}_n, \mathbf{x}) \end{bmatrix}.$$

So if $\mathbf{y} \in O_2$, then all the rows of $L_{\mathbf{y}}$ are in W, which implies that the matrix has rank at most $\dim W = o_2$. The matrices, $L_{\mathbf{y}} = \sum_{j=1}^{n} y_j L_{\mathbf{e}_j}$ where $L_{\mathbf{e}_j}$ are known. Furthermore, we can set the last $o_2 - 1$ components of $\mathbf{y}$ to 0 and still expects a result. Thus, there is a nontrivial linear combination of $k = n - o_2 + 1$ matrices $L_{\mathbf{e}_1}, \ldots, L_{\mathbf{e}_{n-o_2+1}}$ of shape $n \times m$ that has rank o_2.

Support Minors Modeling (as introduced in Bardet et al in [2]) of this MinRank problem proceeds as follows: We take a basis of W and diagonalize it, obtaining the $o_2 \times m$ matrix C. Take any row $\mathbf{r}_i$ of the matrix $L_{\mathbf{y}}$ written as linear forms in the y_i. Consider the $(o_2+1) \times (o_2+1)$ minors of the matrix

$$\begin{bmatrix} \mathbf{r}_i \\ C \end{bmatrix}$$

and consider all $o_2 \times o_2$ minors of C plus the y_i as variables. We have obtained a bilinear system with $n\binom{m}{o_2+1}$ equations and can solve using (bipartite) XL, immediately so if $n\binom{m}{o_2+1} \geq (n - o_2 + 1)\binom{m}{o_2} - 1$.

3.3 Columns instead of Rows

A variant is to run the above attack with a transposition (swapping columns and rows). For this MinRank problem there turns out to be no gain.

3.4 Two y's and Elimination

As an extension of the Rectangular MinRank attack, we can set the last o_2 components of $\mathbf{y}$ to either $(1, 0, 0, \ldots, 0)$ or $(0, 1, 0, \ldots, 0)$. Therefore, consider that we have the rows $\mathbf{r}_i$ of $L_{\mathbf{y}}$ and the rows $\mathbf{r}'_i$ of $L_{\mathbf{y}'}$ both of which span W.

Thus, we can repeat the Support Minors Modeling attack using both the minors of $\begin{bmatrix} \mathbf{r}_i \\ C \end{bmatrix}$ and that of $\begin{bmatrix} \mathbf{r}'_i \\ C \end{bmatrix}$. *Assuming that* $n\binom{m}{o_2+1} > \left(n - o_2 + \frac{1}{2}\right)\binom{m}{o_2}$, *we can eliminate the equations with* $\mathbf{y}$ *down to linear forms in just the* C *variables, and do the same with the* $\mathbf{y}'$ *equations, and there will be sufficiently many equations to solve for the* C *variables.*

Note that this attack can be obviously extended to multiple $\mathbf{y}$'s up to $o_2 - 1$. However, none of this seems to matter within the range of parameters that we are discussing, and we did not spot any case where the above applies.

3.5 Tripartite XL with two y's

Let's further consider an extension of the above idea. We simply try to solve the bilinear equations in the $\mathbf{y}$ and the C variables, and those in the $\mathbf{y}'$ and the C variables, all together. We consider raising to the tri-degree $(b, c, 1)$ where b is the degree in the $\mathbf{y}$ variables, c is that of the $\mathbf{y}'$ variables and still linear in the C variables. There are

$$\left(\binom{n-o_2+b-1}{b-1}\binom{n-o_2+c}{c} + \binom{n-o_2+b}{b}\binom{n-o_2+c-1}{c-1}\right) \times n \binom{m}{o_2+1}$$

equations and

$$\binom{n-o_2+b}{b}\binom{n-o_2+c}{c} \times (n-o_2+1)\binom{m}{o_2}$$

variables. However, not all equations are independent. As in the Rectangular Minrank attack, we can derive syzygies from the minors of matrices $\begin{bmatrix} \mathbf{r}_i \\ \mathbf{r}_j \\ C \end{bmatrix}$ and $\begin{bmatrix} \mathbf{r}'_i \\ \mathbf{r}'_j \\ C \end{bmatrix}$ but also the combination $\begin{bmatrix} \mathbf{r}_i \\ \mathbf{r}'_j \\ C \end{bmatrix}$ and so on with even more rows from $\mathbf{y}$ and $\mathbf{y}'$. We wrote scripts to enumerate over b, c to see if we arrive at an improved attack.

Conclusion from the above: We checked over the Rainbow parameter sets and this new variant attack turns out not to improve the complexity over the Rectangular MinRank attack as described in [1].

4 Summary

Our parameter sets I, III, and V as proposed to NIST PQC Round 3 still meet the corresponding NIST security level requirements[2].

While Beullens' paper centers on the new attacks, it neatly summarizes and clarifies the design of Rainbow and identify attack vectors. Critics of Rainbow had voiced the concern that with concentrated attention a new and powerful attack might be found, completely breaking Rainbow. It seems that the new attacks did come and they turned out to be still fundamentally exponential, and in such a way that does not leave Rainbow-I in the range achievable by humans.

[2]Indeed the parameter set Ia as proposed to NIST Round 1 and 2 still does too.

To put things in perspective, had the new insight on the design of Rainbow arrived *prior* to the new and improved MinRank analysis from [2], Rainbow-I would have been assessed to have security greater than 2^{150} multiplications. So the improved modelling by Bardet *et al* deserves some of the credit for this reduction in assessed security levels.

Finally the new perspective from [1] also serves to reassure us that it is unlikely that there will be future attacks heretofore unknown.

References

[1] W. Beullens, *Improved cryptanalysis of UOV and Rainbow*, `eprint.iacr.org/2020/1343`

[2] M. Bardet, M. Bros, D. Carbacas, P. Gaborit, R. Perlner, D. Smith-Tone, J.-P. Tillich, and J. Verbel, *Improvements of Algebraic Attacks for solving the Rank Decoding and MinRank problems*, arXiv 2002.08322v3 (`arxiv.org/pdf/2002.08322.pdf`).

[3] R. Perlner and D. Smith-Tone, *Rainbow Band Separation is Better than we Thought*, `eprint.iacr.org/2020/702`

[4] The NTRU Prime Team, NTRU Prime: NIST Round 3 submission document, available from `csrc.nist.gov/projects/post-quantum-cryptography/round-3-submissions`

[5] The Kyber Team, Crystals-Kyber: NIST Round 3 submission document (see above)

[6] The Dilithium Team, Crystals-Dilithium: NIST Round 3 submission document (see above)

[7] M. Albrecht, B. Curtis, A. Deo, A. Davidson, R. Player, E. Postlethwaite, F. Virdia, and T. Wunderer, *Estimate All The {LWE,NTRU} Schemes*, `estimate-all-the-lwe-ntru-schemes.github.io`.