From: Apon, Daniel C. (Fed)
To: Smith-Tone, Daniel C. (Fed)
Subject: (b) (6)
Date: Monday, April 22, 2019 12:13:38 PM

Yes -- naively, I assume Dist_s is high min-entropy, supported on (say) {-1,0,1} coefficients. (Other options are allowed..)

I don't think there's any issue with assuming every term is invertible as needed. (You and your weird talk about "units".. =))

I'm less sure about restricting the support of Dist_s to a coset of some subgroup of the unit group. It's not AS general of a result, but it's still interesting.
(I would prefer we only restrict Dist_s according to entropy and magnitude of coefficients, for the noise flooding argument.)

Anyway, what does some coset of some subgroup of the unit group of the ring even look like..?

By the way: I really, really, really, really want you to read this paper so that we can discuss =) https://eprint.iacr.org/2018/494.pdf

Order-LWE and the Hardness of Ring-LWE with Entropic Secrets - eprint.iacr.org

eprint.iacr.org

uniform and e iare small integers (say sampled from a discrete Gaussian with parameter ˝q).The adversary’s goal is to distinguish this oracle from one where b iis random. For RLWE, we consider an extremely simpli ed setting for the sake of this high level overview.

From: Smith-Tone, Daniel (Fed)
Sent: Monday, April 22, 2019 11:38:24 AM
To: Apon, Daniel C. (Fed)
Subject: (b) (6)

Hi, Daniel,

I take it that there needs to be a notion of smallness in the weird distribution of s so that the fs+e can be subsumed by the other e. So I’m thinking that we just need to have the c’s be drawn from some subset of the unit group so that the set of products cs is uniform on some set. If the c's are drawn uniformly from some subgroup of the unit group and if the s's are restricted to some coset of this subgroup in R_q, then we get this result for free.

It will take more time to figure out if the other RLWE problems are hard (provably or believably). In any case, the algebra seems to lend itself to this problem.

Cheers,
Daniel

From: Apon, Daniel C. (Fed)
Sent: Monday, April 22, 2019 10:06 AM
To: Smith-Tone, Daniel (Fed) <daniel.smith@nist.gov>
Subject: (b) (6)

Hey Daniel,

(b) (6)
(b) (6)

Anyway, I wanted to try to get concrete about the entropic-Ring-LWE stuff -- writing it down, at least in email, is a first step.

So recall the set-up and the goal.

Initially, we are given a sample (or samples) of the form

a, a*s+e

where a is uniform from ring R_q, s is drawn from a distribution with min-entropy k over R_q, and e is from a discrete Gaussian distribution over R with sufficiently large magnitude.
We also assume that the distribution of s has "high min-entropy," meaning that k is at least omega(log(n)), where n is the ring dimension.
(We've been thinking about R_q = Z[x] / <x^n+1> where the ring dimension n is a power of 2, but variants/extensions R'_q are perfectly fine, of course -- so long as the standard-form RLWE remains plausibly hard over the variant ring R'_q.)

The goal is to 're-randomize' the secret term s -- presumably by changing the ambient space we are working in -- into some s' s.t. s' is 'uniform' or, in particular, so that it looks like a standard-form RLWE instance in the end of things.
This would prove that RLWE with strange distributions is computationally indistinguishable from RLWE with standard distributions (up to the size of the instance, as dictated by the entropy of s).

The argument then proceeds as

a, a*s+e ~= b*c+f, (b*c+f)*s+e //Ring-LWE assumption 'in reverse' by treating a as the uniform part of the RLWE assumption

Then given (b* c+f)*s+e, we re-write as
b*c*s+f*s+e

By noise flooding of Gaussians, this is statistically IND from
b*c*s+e

And we're at the point where we want to hit the b*c*s term with a statistical argument to get IND from
b*s'+e

So, just to write down a simple question:

If we use some R'_q that is 2-good, will the multiplication c*s -- where s is from the 'strange' distribution, but c is from a chosen distribution -- allow s' to land uniformly in some subspace of R'_q appropriately?

Note that, for whatever distribution we choose c from, the problem (b, b*c+f) should be a hard RLWE problem on its own.

So the theorem we would be aiming for is something like...
"RLWE over 2-good rings with a structured secret implies that RLWE with strange distributions over 2-good rings is as hard as RLWE with standard distributions over 2-good rings."

(So, note that we are actually relating **three** distinct forms of RLWE together in order to achieve the result.)

--Daniel