An Investigation into Sparse Secret-LWE and Hybrid Attacks

Ben Curtis and Rachel Player

Information Security Group, Royal Holloway, University of London, UK

Learning with Errors Refresher

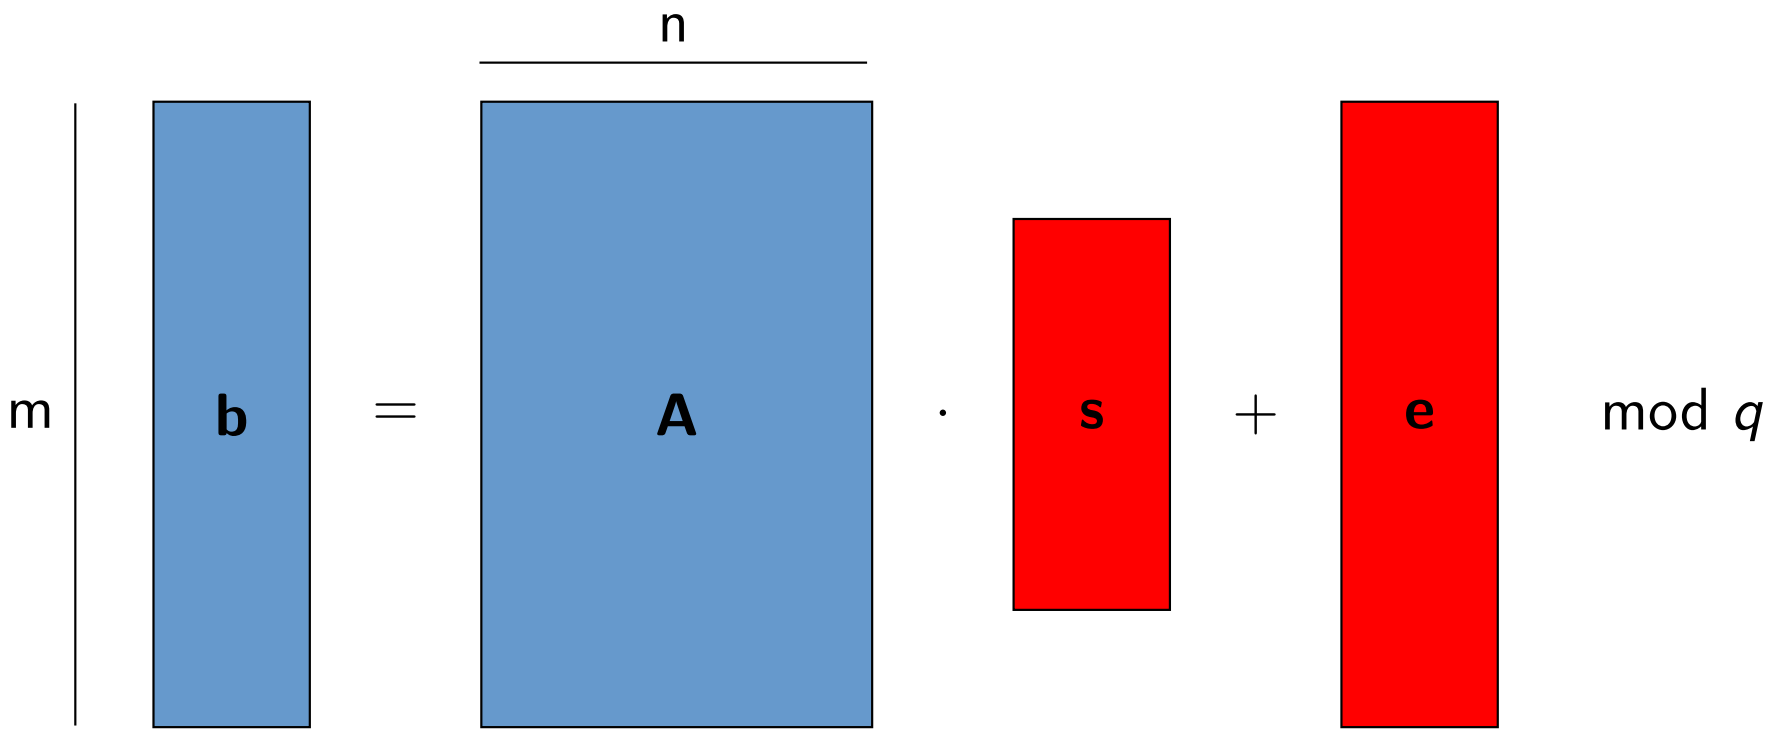

Search: Given $(\mathbf{A}, \mathbf{b})$, aim is to recover $\mathbf{s}$

Decision: Decide if pairs $(\mathbf{A}, \mathbf{b})$ are formed as above or uniformly at random

O. Regev. On lattices, learning with errors, random linear codes, and cryptography. In *STOC*, 2005.

Sparse-secret LWE

Sparse-secret LWE typically considers LWE secrets with small entries and a *low Hamming weight* e.g.

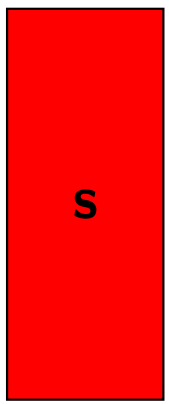

$\in \{-1, 0, 1\}^n$ with at most h nonzero entries

Currently Standardised Parameters: Methodology

The currently standardised parameters are formatted as tables, generated with the following methodology:

- A target security level $\lambda \in \{128, 192, 256\}$, power-of-two ring dimension n and standard deviation $\sigma \approx 3.2$ is fixed
- The LWE Estimator is used to find the maximal modulus q achieving security level λ for the given (n, σ)
- The `usvp`, `dec`, and `dual` attacks (i.e. the default output of the Estimator) are considered

Currently Standardised Parameters: Tables

n	$\log q$	α	usvp	dec	dual	λ_{target}
1024	27	8/q	**131.6**	160.2	138.7	128
2048	54	8/q	**129.7**	144.4	134.2	
4096	109	8/q	**128.1**	134.9	129.9	
8192	218	8/q	**128.5**	131.5	129.2	
16384	438	8/q	**128.1**	129.9	129.0	
32768	881	8/q	**128.5**	129.1	**128.5**	
1024	19	8/q	**193.0**	259.5	207.7	192
2048	37	8/q	**197.5**	233.0	207.8	
4096	75	8/q	**194.7**	212.2	198.5	
8192	152	8/q	**192.2**	200.4	194.6	
16384	305	8/q	**192.1**	196.2	193.2	
32768	611	8/q	**192.7**	194.2	193.7	
1024	14	8/q	**265.6**	406.4	293.8	256
2048	29	8/q	**259.1**	321.7	273.5	
4096	58	8/q	**260.4**	292.6	270.1	
8192	118	8/q	**256.7**	270.4	260.6	
16384	237	8/q	**256.9**	264.2	259.8	
32768	476	8/q	**256.4**	260.2	258.2	

Table: Currently standardised LWE parameters at the 128-, 192- and 256-bit security level for uniform ternary secret and estimates of their security against usvp, dec, and dual attacks under the BKZ cost model $T(\beta, d) = 2^{0.292\beta + 16.4 + \log(8d)}$, where β is the blocksize and d is the dimension. The best performing attack for each parameter set is highlighted in bold.

Hybrid attacks

Combinatorial 'hybrid attacks' apply to small and sparse-secret LWE

Hybrid attacks are not currently supported by the LWE Estimator

We consider the `hybrid-dec` and `hybrid-dual` attacks

N. Howgrave-Graham. A Hybrid Lattice-Reduction and Meet-in-the-Middle Attack Against NTRU. *CRYPTO*, 2007.
J. H. Cheon, *et al.*. A Hybrid of Dual and Meet-in-the-Middle Attack on Sparse and Ternary Secret LWE. *IEEE Access 7*, 8949789506, 2019.

Hybrid Attacks: Intuition

Idea: Splitting $\mathbf{A} = [\mathbf{A}_1|\mathbf{A}_2]$. We have:

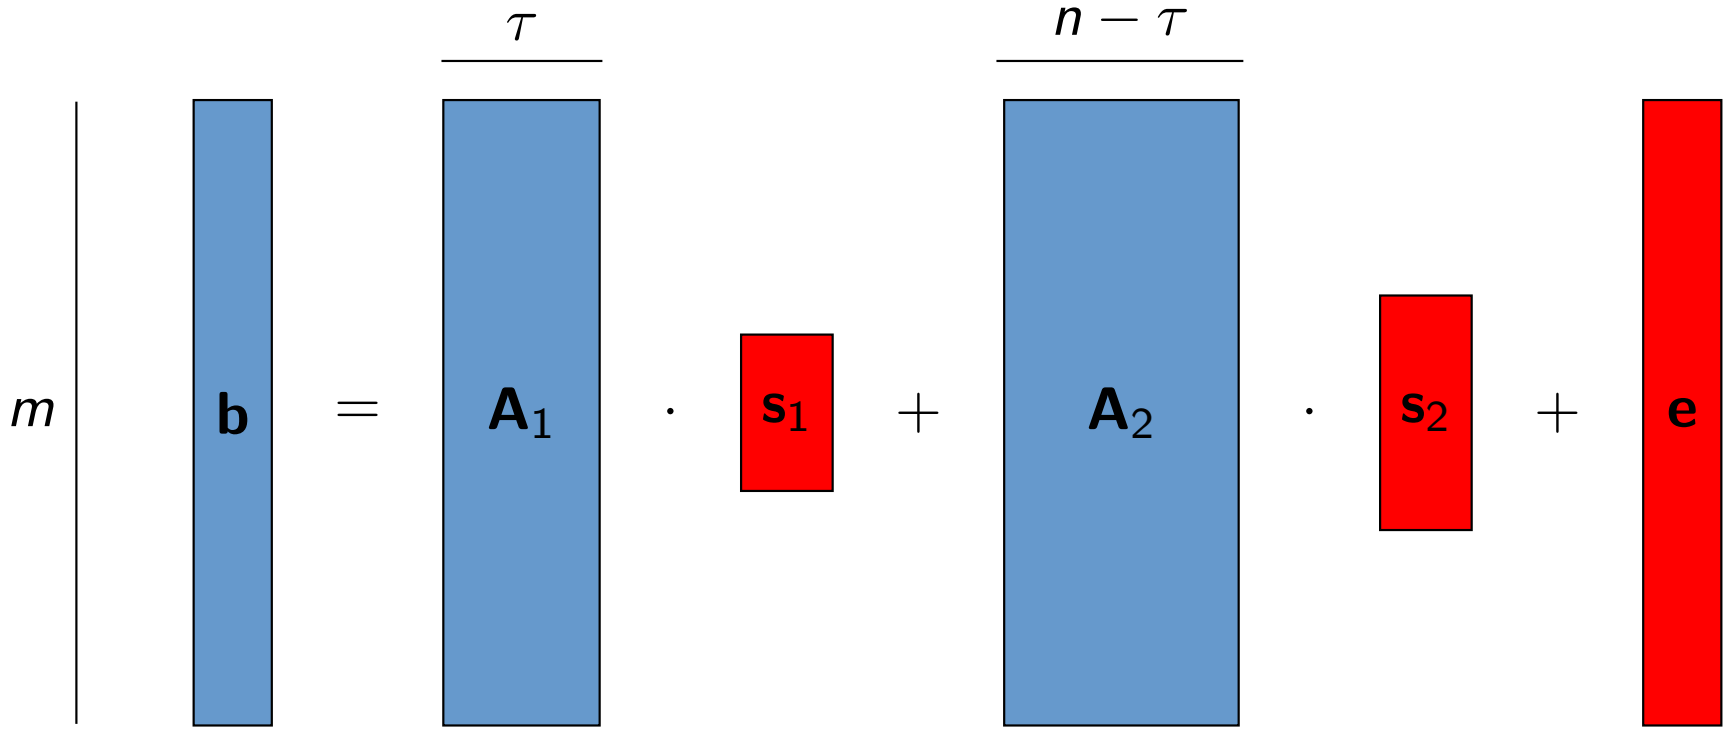

Hybrid Attacks allow for multiple guesses to be made per lattice reduction step - this is *particularly* efficient in the sparse-secret regime.

Security of Current Parameter Sets under Hybrid Attacks

Under a **conservative** analysis of `hybrid-dec` and `hybrid-dual`, current params (with uniform ternary secret) are mostly fine.

n	$\log q$	α	`usvp`	`hybrid-dec`	`hybrid-dual`	λ_{target}
1024	27	8/q	131.6	138.7	**129.6**	128
2048	54	8/q	**129.7**	135.2	131.7	
4096	109	8/q	**128.1**	131.3	128.7	
8192	218	8/q	**128.5**	130.2	128.6	
16384	438	8/q	**128.1**	130.6	**128.1**	
32768	881	8/q	128.5	130.4	**128.3**	
1024	19	8/q	193.0	204.9	**186.4**	192
2048	37	8/q	**197.5**	208.3	197.9	
4096	75	8/q	**194.7**	201.8	193.5	
8192	152	8/q	**192.2**	197.2	192.8	
16384	305	8/q	**192.1**	198.2	192.4	
32768	611	8/q	**192.7**	199.2	193.1	
1024	14	8/q	265.6	287.2	**255.6**	256
2048	29	8/q	259.1	275.0	**249.5**	
4096	58	8/q	**260.4**	272.5	258.0	
8192	118	8/q	**256.7**	263.5	258.2	
16384	237	8/q	256.9	261.0	**256.3**	
32768	476	8/q	**256.4**	258.8	257.0	

Table: Current Parameter Sets with Hybrid Attacks

Sparse Secrets

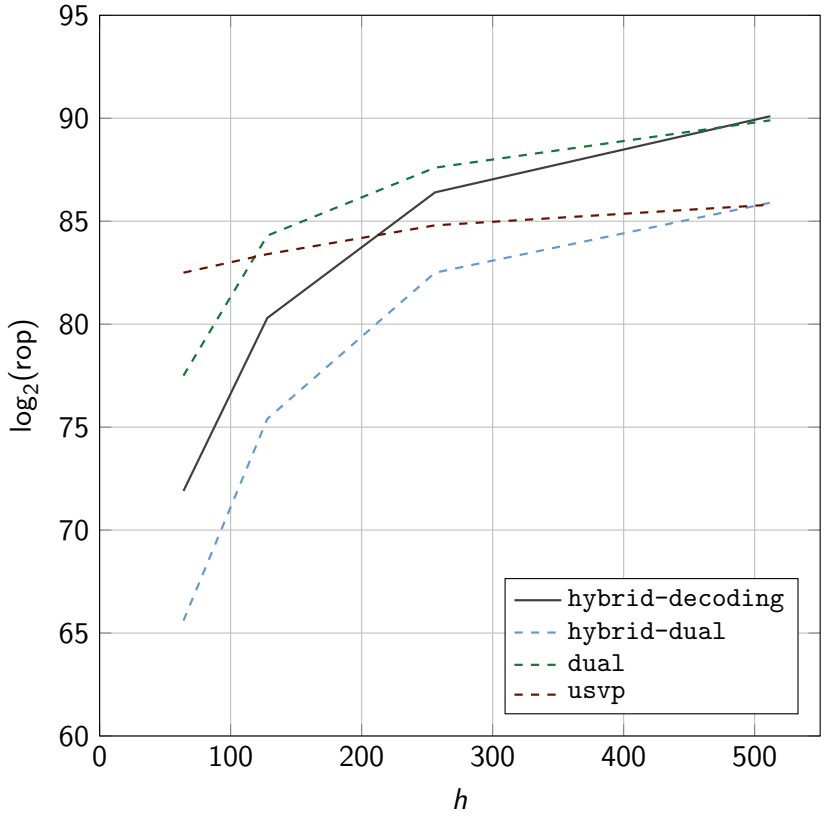

Figure: A comparison of the `usvp`, `dual`, `hybrid-dual` and `hybrid-dec` attacks, for the parameter set $n = 1024, q = 2^{40}$ and $\sigma \approx 3.2$ considering a sparse ternary secret with a variety of Hamming weights $h \in \{64, 128, 256, 512\}$.

Tables of $\log(q)$ for Varying Sparsity

n	λ	$h = \frac{\lambda}{2}$	$h = \frac{3\lambda}{4}$	$h = \lambda$	$h = \frac{3\lambda}{2}$	Current Standard
1024	128	14	19	21	23	27
	192	9	13	14	16	19
	256	7	10	11	12	14
2048	128	27	37	41	46	54
	192	19	26	29	32	37
	256	15	19	22	24	29
4096	128	55	74	83	92	109
	192	37	52	57	64	75
	256	30	39	44	49	58
8192	128	111	148	171	186	218
	192	84	100	114	130	152
	256	60	79	89	98	118
16384	128	223	300	342	377	438
	192	157	201	232	265	305
	256	115	161	176	202	237
32768	128	496	619	699	767	881
	192	350	411	479	523	611
	256	263	313	361	408	476

Table: The reduction in bitsize $\log q$ of the modulus q required to retain the desired level of security against `dual`, `usvp`, `hybrid-dual` and `hybrid-dec`, under our assumptions, when using a sparse ternary secret compared to a uniform ternary secret considering the BKZ cost model $T_{\mathsf{BKZ}}(\beta, d) = 2^{0.292\beta + 16.4 + \log(8d)}$

Discussion (1)

Rationale for standardising secure parameters

Should the recommendation of parameters be lead by implementation choices / requirements?

- Historically this has been the case, e.g. power-of-two n, $\sigma = 3.19$
- Alternatively, the Standard could encourage implementors to select parameters we are confident about

Discussion (2)

Should we standardise sparse-secret parameter sets?

If yes, is there an appropriate sparseness that balances performance (e.g. practicality of bootstrapping) with security?

Discussion (3)

Is the existing methodology for the Standard reasonable?

At present, sets of LWE parameters themselves are standardised

Is it better to standardise the methodology?

Would we have to standardise a BKZ cost model?
Should we include Hybrid Attacks?

Discussion Summary

1. Rationale for standardising secure parameters
2. Should we standardise sparse-secret parameter sets?
3. Is the existing methodology for the Standard reasonable?

   ...Any other comments / questions?

Exploration of Current Parameter Sets

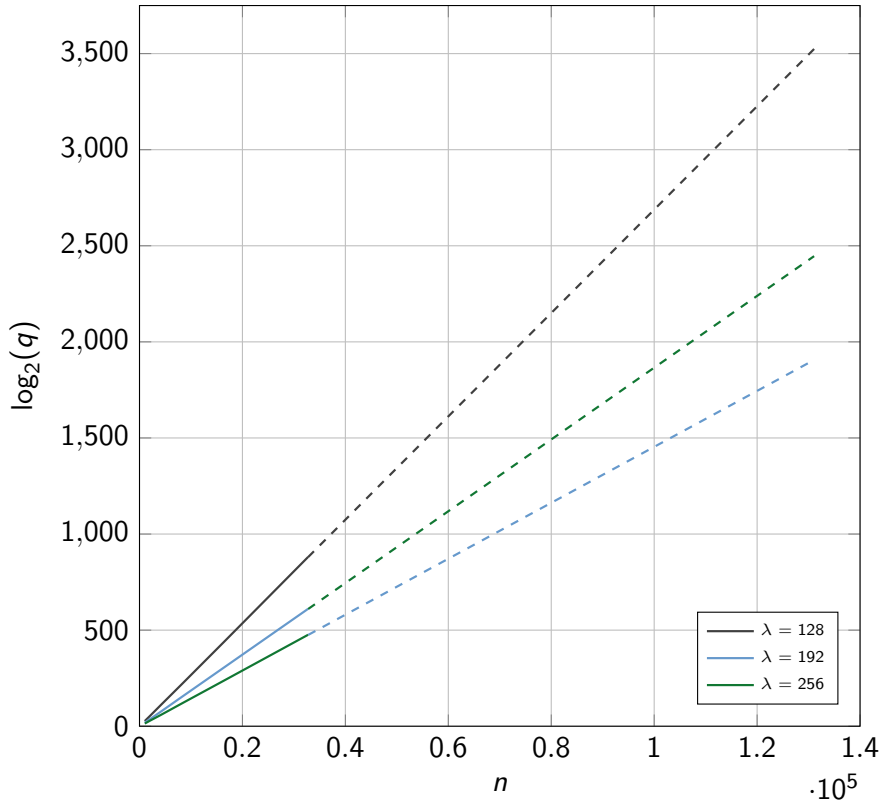

Figure: Extrapolation to $n = 65536$ and $n = 131072$ for uniform ternary secrets. Here, we consider the lattice reduction cost model $T_{\mathsf{BKZ}}(\beta, d) = 2^{0.292\beta + 16.4 + \log(8d)}$ and extrapolate using the Sage function `find_fit`. Note that the solid lines represents data points and the dashed line represents extrapolation.