

Lecture 11

The BGV Scheme

SEC 595: Encrypted Computing

Roadmap

Lecture 10 built BFV. BGV (Brakerski, Gentry and Vaikuntanathan) is best described as BFV with two changes: a chain of moduli instead of one modulus, and the message placed in the low bits with every error scaled by t . Section 1 states the two changes and Section 2 the modulus chain. Sections 3–5 give the scheme and its noise; the new feature is that multiplication *squares* the noise. Section 6 adapts relinearization, and Section 7 introduces modulus switching, which keeps the squared noise in check. Section 8 runs everything on one small instance, and Section 9 places BGV beside BFV and CKKS.

1. BGV Against BFV

1.1 What Stays and What Changes

The ring $R_q = \mathbb{Z}_q[X]/(X^n + 1)$, the distributions, the encoding, the interface (KeyGen, Encrypt, Decrypt, EvalAdd, EvalMult, Relinearize) and the decision-RLWE assumption are those of Lecture 10. Two things change.

1. **A modulus chain.** BFV fixes one q . BGV fixes $q_L > q_{L-1} > \dots > q_0$ and gives every ciphertext a *level*.
2. **Where the message sits.** BFV scales the message by $\Delta = \lfloor q/t \rfloor$ and places it in the high bits. BGV scales every *error* by t and leaves the message in the low bits.

Everything else in this lecture follows from these two changes.

1.2 Where the Message Sits

Think of one ciphertext coefficient as a word of $\log_2 q$ bits. Both schemes must fit a message and a noise term into it without letting them collide.

One ciphertext coefficient, $\log_2 q$ bits wide			
	MSB side	middle	LSB side
BFV	payload ΔM	noise budget	noise v
BGV	noise budget	noise $t v$	payload M

BFV removes the noise by scaling by t/q and *rounding*. BGV removes it by reducing *modulo* t , since $t v \equiv 0 \pmod{t}$. The factor must be exactly t or a multiple of it: any other factor would survive the reduction modulo t .

2. The Modulus Chain

2.1 The Chain and the Level

Definition 1 (Modulus chain). Fix pairwise coprime $p_0, p_1, \dots, p_L$ and set

$$q_\ell = p_0 p_1 \cdots p_\ell, \quad \ell = 0, 1, \dots, L.$$

A ciphertext at **level** ℓ lies in $R_{q_\ell} \times R_{q_\ell}$.

Each $q_{\ell-1}$ divides q_ℓ , with $q_\ell/q_{\ell-1} = p_\ell$. In the RNS representation (Lecture 4), dropping from level ℓ to level $\ell - 1$ deletes one residue. A ciphertext is encrypted at level L , moves down one level with each multiplication, and is decrypted at level 0.

2.2 Consequences of Levels

1. **L is the multiplicative depth**, fixed when the chain is fixed. Deeper circuits need bootstrapping (Lecture 14).
2. **Operands must be aligned**. Two ciphertexts combine only at the same level; the higher one is switched down first.
3. **Ciphertexts shrink**. At level ℓ a ciphertext costs $2n \log_2 q_\ell$ bits, so late operations are cheaper.

A longer chain means a larger q_L , which needs a larger n for the same security (Lecture 6).

3. Keys, Encryption, Decryption

3.1 Spaces and Parameters

$$\mathcal{P} = R_t, \quad \mathcal{C} = R_{q_\ell} \times R_{q_\ell}, \quad q_\ell \gg t,$$

with coefficients in the centred range. The secret key and the encryption randomness u are drawn from R_2 (ternary coefficients), the errors e, e_1, e_2 from $\mathcal{X}$ (discrete Gaussian, tail-cut at β), and the mask a uniformly from R_{q_ℓ} . The chain carries one new constraint,

$$p_i \equiv 1 \pmod{t} \quad \text{for every } i \geq 1, \tag{1}$$

which is needed by ModSwitch (Section 7).

3.2 Key Generation

Definition 2 (KeyGen). Sample $\text{SK} \leftarrow R_2$, $a \leftarrow R_{q_L}$ and $e \leftarrow \mathcal{X}$, and set

$$\text{PK}_0 = \lceil -(a \cdot \text{SK} + t \cdot e) \rceil_{q_L}, \quad \text{PK}_1 = a.$$

Compared with BFV's $\text{PK}_0 = \lceil -(a\text{SK} + e) \rceil_q$, one factor t is inserted. Then

$$\text{PK}_0 + \text{PK}_1 \text{SK} = -te, \tag{2}$$

which is small *and* a multiple of t . The pair is still one RLWE sample, now with error distribution $t\mathcal{X}$; security is unchanged provided $t \ll q$. At a level $\ell < L$, PK is reduced modulo q_ℓ .

3.3 Encryption

Definition 3 ($\text{Encrypt}(\text{PK}, M)$). Sample $u \leftarrow R_2$ and $e_1, e_2 \leftarrow \mathcal{X}$, and output, at level L ,

$$c_0 = [\text{PK}_0 u + t e_1 + M]_{q_L}, \quad c_1 = [\text{PK}_1 u + t e_2]_{q_L}.$$

Side by side with BFV,

$$\text{BFV: } c_0 = \text{PK}_0 u + e_1 + \Delta M, \quad \text{BGV: } c_0 = \text{PK}_0 u + t e_1 + M.$$

The scaling has moved from the message to the errors.

3.4 Decryption and the tv Cancellation

Definition 4 ($\text{Decrypt}(\text{SK}, \text{ct})$ at level ℓ).

$$M = \left[[c_0 + c_1 \text{SK}]_{q_\ell} \right]_t.$$

Reduce modulo q_ℓ into the centred range, then reduce modulo t . There is no division and no rounding; compare BFV's $\left\lfloor \frac{t}{q} [\cdot]_q \right\rfloor$.

Proposition 1 (Decryption identity). *For a fresh ciphertext,*

$$c_0 + c_1 \text{SK} = M + t v \pmod{q_L}, \quad v = -eu + e_1 + e_2 \text{SK}.$$

Proof. Substitute the definitions: $c_0 + c_1 \text{SK} = -au\text{SK} - teu + te_1 + M + au\text{SK} + te_2 \text{SK} = M + t(-eu + e_1 + e_2 \text{SK})$. $\square$

The mask $au\text{SK}$ cancels, as in BFV. Every surviving error carries the factor t , so $[M + tv]_t = M$ for *any* v , once the reduction modulo q_ℓ has returned $M + tv$ exactly. The polynomial v is BFV's fresh noise, so the bound of Lecture 10 carries over: with u, SK ternary and $\|e_i\| \leq \beta$,

$$\|v\| \leq n\beta + \beta + n\beta = 2n\beta + \beta. \quad (3)$$

4. Correctness

Proposition 2 (Correctness). *Let ct at level ℓ satisfy $c_0 + c_1 \text{SK} = M + tv \pmod{q_\ell}$. If $\|M + tv\| < q_\ell/2$, decryption returns M . It suffices that*

$$\|v\| < \frac{q_\ell}{2t}$$

(up to the contribution $\|M\| \leq t/2$).

Proof. Over $\mathbb{Z}$, $c_0 + c_1 \text{SK} = M + tv + q_\ell k$. If $\|M + tv\| < q_\ell/2$, the centred reduction modulo q_ℓ removes $q_\ell k$ and returns $M + tv$ exactly; reduction modulo t then removes tv . $\square$

The threshold is BFV's with q replaced by q_ℓ , and the failure has the same form. Past the threshold the centred reduction subtracts j extra multiples of q_ℓ , with $j = \lfloor (M + tv)/q_\ell \rfloor$, so decryption returns

$$\lfloor M - j q_\ell \rfloor_t = \lfloor M - j r \rfloor_t, \quad r = \lfloor q_\ell \rfloor_t.$$

By (1), $r = \lfloor p_0 \rfloor_t$ at every level. BFV's rounding instead returns $M + \lfloor tv/q \rfloor$, essentially $M + j$. When $r = 1$, as in this lecture's examples, the two failures are mirror images: off by j in opposite directions. Both are silent.

Example 1 (One failure, in BGV and in BFV). *BGV*: $t = 8$, prime $q = 1009$, one coefficient $M = 5$. The bound is $q/(2t) = 1009/16 = 63.06$.

$$v = 30 : \quad M + tv = 245, \quad \lfloor 245 \rfloor < 504.5, \quad 245 \bmod 8 = 5 \quad \checkmark$$

$$v = 70 : \quad M + tv = 565 \equiv 565 - 1009 = -444 \pmod{1009}, \quad (-444) \bmod 8 = 4 \times$$

BFV: $t = 8$, $q = 1024$, $m = 5$, so $\Delta = 128$, $\Delta m = 640$, and the bound is $q/(2t) = 64$.

$$v = 30 : \quad \frac{8}{1024}(640 + 30) = 5.2344 \longrightarrow \lfloor 5.2344 \rfloor = 5 \quad \checkmark$$

$$v = 70 : \quad \frac{8}{1024}(640 + 70) = 5.5469 \longrightarrow \lfloor 5.5469 \rfloor = 6 \times$$

The thresholds agree to within one unit. BGV returns $4 = 5 - 1$: one multiple of q is subtracted, and $1009 \equiv 1 \pmod{8}$. BFV returns $6 = 5 + 1$. Each result is a well-formed plaintext, so neither failure is detectable.

5. EvalAdd and EvalMult

5.1 Addition

Definition 5 (EvalAdd). For two ciphertexts at the same level ℓ ,

$$\text{ct}^{(3)} = \left([c_0^{(1)} + c_0^{(2)}]_{q_\ell}, [c_1^{(1)} + c_1^{(2)}]_{q_\ell} \right).$$

Adding the decryption identities gives $c_0^{(3)} + c_1^{(3)} \text{SK} = (M^{(1)} + M^{(2)}) + t(v_1 + v_2)$. The factor t survives, the noise adds, $\|v_3\| \leq \|v_1\| + \|v_2\|$, and no level is consumed.

5.2 Multiplication

Definition 6 (EvalMult). For two ciphertexts at level ℓ ,

$$\text{ct}^* = \left([c_0^{(1)} c_0^{(2)}]_{q_\ell}, [c_0^{(1)} c_1^{(2)} + c_1^{(1)} c_0^{(2)}]_{q_\ell}, [c_1^{(1)} c_1^{(2)}]_{q_\ell} \right).$$

These are the coefficients of S^0, S^1, S^2 in $(c_0^{(1)} + c_1^{(1)} S)(c_0^{(2)} + c_1^{(2)} S)$ (Lecture 4), and ct^* is decrypted with $c_0^* + c_1^* \text{SK} + c_2^* \text{SK}^2$. There is no rescaling: BFV must remove one Δ from $\Delta^2 M^{(1)} M^{(2)}$, whereas BGV has no Δ , so the tensor product is the whole operation.

5.3 The Product Noise

Proposition 3 (Multiplicative noise). *Let $\text{ct}^{(i)}(\text{SK}) = M^{(i)} + tv_i \pmod{q_\ell}$. Then $\text{ct}^*(\text{SK}) = M^{(1)}M^{(2)} + tv_3 \pmod{q_\ell}$ with*

$$v_3 = M^{(1)}v_2 + M^{(2)}v_1 + tv_1v_2.$$

If $\|v_i\| \leq V$ and $\|M^{(i)}\| \leq t/2$, then $\|v_3\| \leq ntV + tnV^2 = ntV(1 + V)$.

Proof. Over $\mathbb{Z}$, $\text{ct}^{(i)}(\text{SK}) = M^{(i)} + tv_i + q_\ell r_i$; every term of the product containing q_ℓ vanishes modulo q_ℓ , and $(M^{(1)} + tv_1)(M^{(2)} + tv_2) = M^{(1)}M^{(2)} + t(M^{(1)}v_2 + M^{(2)}v_1 + tv_1v_2)$. For the bound, $\|M^{(1)}v_2\| \leq n(t/2)V$, likewise for the second term, and $\|tv_1v_2\| \leq tnV^2$ (Lecture 3). $\square$

The term tv_1v_2 is **quadratic**. BFV multiplies the noise by a constant G ; BGV multiplies it by ntV , which depends on the noise itself.

5.4 Doubly Exponential Growth, and the Repair

Ignoring the linear terms, $V_{\ell+1} \approx tnV_\ell^2$, so after L multiplications

$$V_L \approx \frac{(tnV_0)^{2^L}}{tn},$$

doubly exponential in L . At $n = 4096$, $t = 2$, $\beta = 19$, the fresh bound (3) is $V_0 = 155667 \approx 2^{17.25}$, so $tnV_0 \approx 2^{30.25}$ and

$$\log_2 V_3 \approx 8 \cdot 30.25 - 13 = 229, \quad \log_2 V_6 \approx 64 \cdot 30.25 - 13 = 1923.$$

The repair is to divide by a prime $p_\ell \approx ntV$ after each multiplication:

$$V_{\ell+1} \approx \frac{tnV_\ell^2}{p_\ell} + B_{\text{add}} \approx V_\ell + B_{\text{add}}.$$

The noise stays near a floor and the modulus loses one prime per level. Section 7 carries this out.

6. Relinearization

6.1 The Evaluation Key

The goal is to turn (c_0^*, c_1^*, c_2^*) into $(\hat{c}_0, \hat{c}_1)$ with $\hat{c}_0 + \hat{c}_1 \text{SK} \approx c_0^* + c_1^* \text{SK} + c_2^* \text{SK}^2$.

Definition 7 (EK and Relinearize). Sample $a \leftarrow R_{q_\ell}$ and $e \leftarrow \mathcal{X}$ afresh and publish

$$\text{EK} = \left(\left[-(a\text{SK} + te) + \text{SK}^2 \right]_{q_\ell}, a \right), \quad \text{so that } \text{EK}_0 + \text{EK}_1 \text{SK} = \text{SK}^2 - te.$$

Relinearize by

$$\hat{c}_0 = [c_0^* + \text{EK}_0 c_2^*]_{q_\ell}, \quad \hat{c}_1 = [c_1^* + \text{EK}_1 c_2^*]_{q_\ell}.$$

Proposition 4. $\hat{c}_0 + \hat{c}_1 \text{SK} = c_0^* + c_1^* \text{SK} + c_2^* \text{SK}^2 - t c_2^* e$.

Proof. $\hat{c}_0 + \hat{c}_1 \text{SK} = c_0^* + c_1^* \text{SK} + c_2^* (\text{EK}_0 + \text{EK}_1 \text{SK}) = c_0^* + c_1^* \text{SK} + c_2^* (\text{SK}^2 - te)$. $\square$

The key error carries t , like every other error in BGV, so the added term is invisible modulo t and only adds $-c_2^* e$ to v . But c_2^* is spread over all of $\mathbb{Z}_{q_\ell}$, so $\|c_2^* e\| \approx \frac{1}{2} n q_\ell \beta$, which is far too large.

6.2 Base Decomposition

Fix a base w , let $\ell_w = \lceil \log_w q_\ell \rceil$, and write $c_2^* = \sum_i w^i d_i$ coefficientwise with $\|d_i\| < w$ (Lecture 4). Publish ℓ_w keys and relinearize with the digits:

$$\text{EK}^{(i)} = \left(\left[- (a_i \text{SK} + t e_i) + w^i \text{SK}^2 \right]_{q_\ell}, a_i \right), \quad \hat{c}_j = \left[c_j^* + \sum_i d_i \text{EK}_j^{(i)} \right]_{q_\ell}.$$

Since $\sum_i d_i w^i \text{SK}^2 = c_2^* \text{SK}^2$, the added noise is $-\sum_i d_i e_i$, and

$$\|v_{\text{Relinearize}}\| \leq n \ell_w (w - 1) \beta.$$

A blow-up of order q_ℓ becomes an additive term of order w , at the cost of ℓ_w times more key material.

7. ModSwitch

7.1 The Operation

Definition 8 (ModSwitch). For ct modulo q and $q' \mid q$,

$$\text{ct}' = \left\lfloor \frac{q'}{q} \text{ct} \right\rfloor,$$

coefficientwise, where each coefficient is rounded to the nearest integer *congruent to the original coefficient modulo t* . ct' is read modulo q' .

In BGV, $q = q_\ell$, $q' = q_{\ell-1}$ and $q/q' = p_\ell$. The level-consuming step at level ℓ is

$$\text{EvalMult} \rightarrow \text{Relinearize} \rightarrow \text{ModSwitch (to level } \ell - 1).$$

Switching last divides the relinearization noise by p_ℓ as well.

7.2 Why It Reduces the Noise

Scaling by q'/q divides message, noise and modulus alike, so the ratio $t\|v\|/q$ barely changes. What matters is the *absolute* noise, because v_3 contains $t v_1 v_2$:

$$V \xrightarrow{\text{mult}} ntV^2 \xrightarrow{\text{switch by } p \approx ntV} V + B_{\text{add}}.$$

Without the switch the sequence is $V \rightarrow ntV^2 \rightarrow (nt)^3 V^4 \rightarrow \dots$. The additive term comes from the rounding (Lecture 4):

$$\|v'\| \leq \frac{q'}{q} \|v\| + \frac{1}{2} (1 + n \|\text{SK}\|) + \frac{\|M\|}{t}. \quad (4)$$

7.3 The Condition $q \equiv q' \pmod{t}$

Over $\mathbb{Z}$, $\text{ct}(\text{SK}) = M + tv + qk$, where k is *not* small. Since $\text{ct}' \equiv \text{ct} \pmod{t}$ coefficientwise, $\text{ct}'(\text{SK}) \equiv M + qk \pmod{t}$, and $\text{ct}'(\text{SK})$ reduced modulo q' is $w = \text{ct}'(\text{SK}) - q'k$. Hence

$$w \equiv M + qk - q'k = M + (q - q')k \pmod{t}.$$

Since k cannot be bounded, the term $(q - q')k$ must vanish identically, that is, $q \equiv q' \pmod{t}$.

Proposition 5. Let $q = p_\ell q'$ with $\gcd(q', t) = 1$. Then

$$q \equiv q' \pmod{t} \iff p_\ell \equiv 1 \pmod{t} \iff p_\ell^{-1} \equiv 1 \pmod{t}.$$

Proof. $q - q' = q'(p_\ell - 1)$ and q' is invertible modulo t , so $t \mid q - q'$ iff $t \mid p_\ell - 1$. The inverse of a unit congruent to 1 is congruent to 1. $\square$

Dividing by p_ℓ multiplies the plaintext by $p_\ell^{-1} = 1$ in $\mathbb{Z}_t$: the message is untouched while everything else shrinks. This is condition (1). Combined with the NTT condition $p_i \equiv 1 \pmod{2n}$ (Lecture 4), each prime satisfies $p_i \equiv 1 \pmod{\text{lcm}(t, 2n)}$.

Example 2 (Checking a chain). Take $p_0 = 17$, $p_1 = 41$, $p_2 = 97$.

(a) $t = 8$. $41 = 5 \cdot 8 + 1$ and $97 = 12 \cdot 8 + 1$, so $p_1 \equiv p_2 \equiv 1 \pmod{8}$ and both switches $q_2 \rightarrow q_1 \rightarrow q_0$ are legal.

(b) $t = 16$. $97 = 6 \cdot 16 + 1 \equiv 1$, so $q_2 \rightarrow q_1$ is legal; but $41 = 2 \cdot 16 + 9 \equiv 9 \pmod{16}$, so $q_1 \rightarrow q_0$ is not. The ladder stops at level 1.

(c) $t = 2$. The condition says only that every p_i is odd.

8. A Worked Instance

8.1 The Instance

$$n = 4, \quad t = 8, \quad p_0 = 1009, \quad p_1 = 97, \quad L = 1, \quad q_1 = p_0 p_1 = 97873, \quad q_0 = 1009.$$

Both primes are $\equiv 1 \pmod{8}$, which is both $p_i \equiv 1 \pmod{t}$ and $p_i \equiv 1 \pmod{2n}$. The error bound is $\beta = 1$. The secret key, the error and the encryption randomness are those of the BFV instance of Lecture 10:

$$\begin{aligned} \text{SK} &= 1 - X + X^3, & e &= 1 - X^2 + X^3, \\ a &= 42445 + 19772X + 51750X^2 + 85319X^3, \\ \text{PK}_0 &= 87746 + 74423X + 53349X^2 + 21851X^3, & \text{PK}_1 &= a. \end{aligned}$$

Example 3 (Budgets and keys). (a) *Budgets.* At level 1, $q_1/(2t) = 97873/16 = 6117.06$; at level 0, $q_0/(2t) = 1009/16 = 63.06$.

(b) *Fresh noise.* $2n\beta + \beta = 8 + 1 = 9$.

(c) *One multiplication.* With $V = 9$, Proposition 3 gives $ntV(1 + V) = 4 \cdot 8 \cdot 9 \cdot 10 = 2880$. This fits the level-1 budget 6117.06 but not the level-0 budget 63.06: the product must be computed at level 1 and switched down before decryption.

(d) *Public key.* Computing $a\text{SK}$ in R_{q_1} and adding PK_0 gives, centred,

$$\text{PK}_0 + \text{PK}_1 \text{SK} = -8 + 8X^2 - 8X^3 = -8(1 - X^2 + X^3) = -te,$$

and every coefficient is a multiple of 8, as (2) requires.

Example 4 (Encrypt and decrypt). Encode $m_1 = 5$ and $m_2 = 2$ in base 2: $M^{(1)} = 1 + X^2$, $M^{(2)} = X$. With

$$\begin{aligned} u^{(1)} &= 1 - X^2 + X^3, & e_1^{(1)} &= X - X^2, & e_2^{(1)} &= 1 - X, \\ u^{(2)} &= -1 + X + X^2, & e_1^{(2)} &= 1 - X^3, & e_2^{(2)} &= -X + X^2 + X^3, \end{aligned}$$

encryption gives

$c_0^{(1)}$	$66673 + 42933X + 41618X^2 + 35174X^3$
$c_1^{(1)}$	$74431 + 53333X + 21859X^2 + 10119X^3$
$c_0^{(2)}$	$32808 + 89346X + 10947X^2 + 8040X^3$
$c_1^{(2)}$	$16232 + 35219X + 10475X^2 + 84084X^3$

Computing $[c_0 + c_1\text{SK}]_{q_1}$, subtracting M and dividing by t gives

$$v_1 = 2 - 3X + 3X^2 - X^3, \quad \|v_1\| = 3, \quad v_2 = 4 - 2X - X^2 + X^3, \quad \|v_2\| = 4,$$

and computing $-eu + e_1 + e_2\text{SK}$ directly gives the same two polynomials, confirming Proposition 1. Both are below the fresh bound 9. Reducing $[c_0 + c_1\text{SK}]_{q_1}$ modulo 8 returns $1 + X^2$ and X , which decode at $X = 2$ to 5 and 2. Since $v = -eu + e_1 + e_2\text{SK}$ does not involve q , t or a , these are exactly the noises of the BFV instance of Lecture 10.

Example 5 (Add and multiply). *Add.* $v_1 + v_2 = 6 - 5X + 2X^2$, so $\|v\| = 6$, below $\|v_1\| + \|v_2\| = 7$ because coefficients partly cancel. Decryption returns $1 + X + X^2$, which decodes to $1 + 2 + 4 = 7$.

Multiply. The tensor product is

$$\begin{aligned} c_0^* &= 15380 + 80778X + 16830X^2 + 78638X^3, \\ c_1^* &= 39785 + 69521X + 34001X^2 + 24344X^3, \\ c_2^* &= 35022 + 32246X + 52247X^2 + 73000X^3, \end{aligned}$$

and $[c_0^* + c_1^*\text{SK} + c_2^*\text{SK}^2]_{q_1} = M^{(3)} + t v_3$ with

$$M^{(3)} = X + X^3, \quad v_3 = 102 - 161X + 136X^2 - 38X^3, \quad \|v_3\| = 161.$$

$M^{(3)}$ decodes to $2 + 8 = 10 = 5 \times 2$.

Split. The three parts of v_3 are

$$M^{(1)}v_2 = 5 - 3X + 3X^2 - X^3, \quad M^{(2)}v_1 = 1 + 2X - 3X^2 + 3X^3, \quad t v_1 v_2 = 96 - 160X + 136X^2 - 40X^3,$$

which sum to v_3 . The linear parts contribute at most 5; the quadratic part contributes 160 and dominates. The bound of Proposition 3 with $V = 4$ is $ntV(1 + V) = 4 \cdot 8 \cdot 4 \cdot 5 = 640$.

Example 6 (Relinearize). Use base decomposition with $w = 4$ and errors $t e_i$ in every key.

(a) $4^8 = 65536 < 97873 \leq 262144 = 4^9$, so $\ell_w = \lceil \log_4 97873 \rceil = 9$ keys are published.

(b) The bound is $n\ell_w(w - 1)\beta = 4 \cdot 9 \cdot 3 \cdot 1 = 108$.

(c) The result is

$$\hat{c}_0 = 2831 + 18002X + 61140X^2 + 85983X^3, \quad \hat{c}_1 = 32787 + 81132X + 67514X^2 + 46330X^3.$$

Evaluating $[\hat{c}_0 + \hat{c}_1 \text{SK}]_{q_1}$ gives $M^{(3)} + t v$ with $M^{(3)} = X + X^3$ and

$$v = 102 - 146X + 149X^2 - 36X^3, \quad \|v\| = 149.$$

The increment is $v - v_3 = 15X + 13X^2 + 2X^3$, so $\|v - v_3\| = 15 \leq 108$. Decryption still returns $X + X^3 = 10$.

Example 7 (One modulus switch). Switch $\hat{\text{ct}}$ from $q_1 = 97873$ to $q_0 = 1009$, dropping $p_1 = 97$. In centred form,

$$\hat{c}_0 = 2831 + 18002X - 36733X^2 - 11890X^3, \quad \hat{c}_1 = 32787 - 16741X - 30359X^2 + 46330X^3.$$

(a) For each coefficient c , compute $c \bmod 8$ and $(q_0/q_1)c = c/97$, and take the nearest integer congruent to c modulo 8:

c	$c \bmod 8$	$(q_0/q_1)c$	c'
2831	7	29.186	31
18002	2	185.588	186
-36733	3	-378.691	-381
-11890	6	-122.577	-122
32787	3	338.010	339
-16741	3	-172.588	-173
-30359	1	-312.979	-311
46330	2	477.629	474

The first four rows are $\hat{c}_0$, the last four $\hat{c}_1$. In the third row the nearest integer to -378.691 is $-379 \equiv 5 \pmod{8}$, but 3 is needed, so $c' = -381$. The extra distance is at most $t/2 = 4$.

(b) Modulo 1009 the switched ciphertext is

$$c'_0 = 31 + 186X + 628X^2 + 887X^3, \quad c'_1 = 339 + 836X + 698X^2 + 474X^3,$$

and $[c'_0 + c'_1 \text{SK}]_{q_0} = M^{(3)} + t v'$ with $M^{(3)} = X + X^3$ and

$$v' = 1 - 2X + 2X^2 - X^3, \quad \|v'\| = 2.$$

Decryption returns $X + X^3$, which decodes to 10.


(c) The bound (4) with $\|v\| = 149$, $\|\text{SK}\| = 1$, $\|M\| = 1$ is $149/97 + \frac{1}{2}(1 + 4) + 1/8 = 4.16$; the observed value is 2. The switch was necessary: $149 > 63.06$, so the unswitched noise would not fit the level-0 budget.

8.2 The Instance in Numbers

Stage	$\ v\ $	
fresh $\text{ct}^{(1)}, \text{ct}^{(2)}$	3, 4	same as BFV
after EvalAdd	6	decrypts to 7
after EvalMult	161	160 from $t v_1 v_2$
after Relinearize ($w = 4$)	149	increment 15
after ModSwitch	2	decrypts to 10

Budgets are 6117.06 at level 1 and 63.06 at level 0. The chain had one prime to spend and one multiplication to spend it on.

9. Where the Message Sits



The word is $\log_2 q = 60$ bits wide with $t = 2^{20}$, $\Delta = 2^{40}$; hatching marks the message. BFV puts it in the top twenty bits, BGV in the bottom twenty, CKKS in the top with a dotted boundary, there being no separation at all.

	BFV	BGV	CKKS
Decryption	$c_0 + c_1 SK$ $= \Delta m + v$	$c_0 + c_1 SK$ $= M + tv$	$c_0 + c_1 SK$ $= \Delta \mu + v$
Recover by	scaling by t/q , rounding, then mod t	reducing mod t ; no rounding	dividing by Δ and accepting the error
Correct when	$\ v\ < q/(2t)$	$\ v\ < q_\ell/(2t)$	never exactly; precision about $\log_2 \Delta - \log_2 \ v\ $, in $\mathbb{C}^{n/2}$

Section 1.2 set BFV and BGV side by side. BGV keeps the message at the bottom and the noise tv above it; the same picture extends to CKKS (Lecture 12).

Summary

1. BGV is BFV with two changes: a modulus chain $q_\ell = p_0 \cdots p_\ell$, with one level spent per multiplication, and the message in the low bits with every error scaled by t .
2. $PK_0 = -(aSK + te)$ and $c_0 + c_1 SK = M + tv$; decryption is $[(c_0 + c_1 SK)_{q_\ell}]_t$, correct while $\|v\| < q_\ell/(2t)$. Overflow silently shifts M by $-j [q_\ell]_t$, where j counts the extra wraps.
3. EvalAdd: the noise adds. EvalMult: a bare tensor product with noise $M^{(1)}v_2 + M^{(2)}v_1 + t v_1 v_2$, which is quadratic.
4. Relinearize uses $EK = (-(aSK + te) + SK^2, a)$ with base- w decomposition; the added noise is at most $n\ell_w(w-1)\beta$.
5. ModSwitch divides the absolute noise by p_ℓ and keeps the message, provided $p_\ell \equiv 1 \pmod{t}$.

Exercises

Throughout, “the instance” means the parameters, keys and ciphertexts of Section 8.

1. Encrypt $M = 1 + X$ at the instance parameters with $u = 1$ and $e_1 = e_2 = 0$, and decrypt. Compute v both as $([c_0 + c_1 SK]_{q_1} - M)/t$ and as $-eu + e_1 + e_2 SK$. Why is v nonzero although $e_1 = e_2 = 0$?

2. At level 0 of the instance, take the single coefficient $M = 5$ with $v = 64$ and compute what decryption returns. Show that the result is off by exactly one, and explain this using $q_0 \equiv 1 \pmod{8}$. Repeat with q_0 replaced by the prime $1013 \equiv 5 \pmod{8}$.
3. Assume the pure recurrence $V_{\ell+1} = tnV_{\ell}^2$ with $t = 2$, $n = 4096$, $V_0 = 2^{17.25}$. Compute $\log_2 V_{\ell}$ for $\ell = 1, 2, 3, 4$. Then assume each multiplication is followed by a switch dividing by $p = 2^{40}$ and adding $B_{\text{add}} = 2^{12}$, and recompute.
4. Identify the term of the BFV product expansion (Lecture 10) that forces the division by q/t , and show that its BGV analogue is a multiple of q_{ℓ} . What does BGV pay for not rescaling?
5. Show that if the evaluation-key error were e rather than te , relinearization would add $-c_2^*e$, and decryption would return $M^{(3)} - c_2^*e \pmod{t}$ even if c_2^*e were small.
6. At $t = 16$ and $n = 4$, say which of the primes $17, 41, 73, 89, 97$ satisfy $p \equiv 1 \pmod{\text{lcm}(t, 2n)}$. Explain, in terms of k in Section 7.3, why a bound on $\|k\|$ would not rescue a chain that violates the congruence.
7. Redo Example 7 with ordinary rounding to the nearest integer, ignoring the congruence modulo 8, and decrypt. Which coefficients of the message change?