Functional Encryption

Lecture 23

ABE from LWE

Functional Encryption

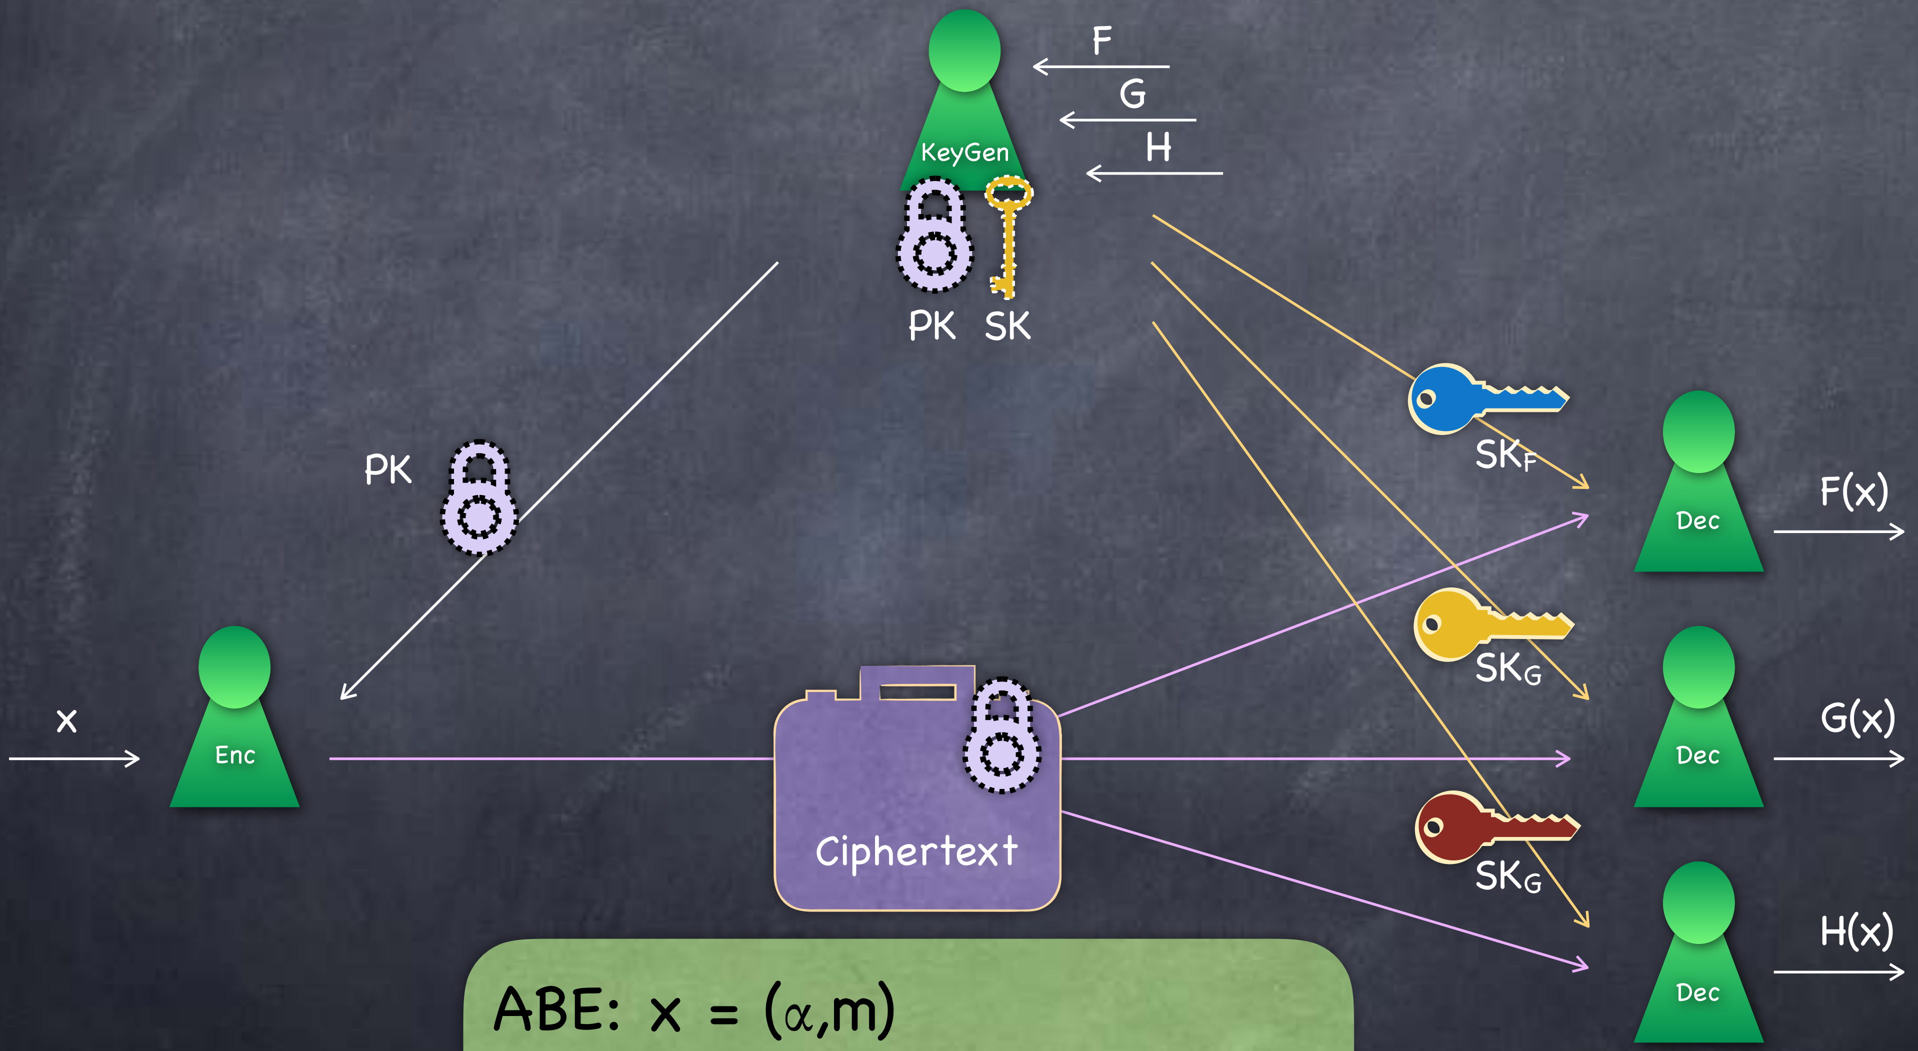

ABE: $x = (\alpha, m)$

$F_{f,z}(x) = (\alpha,\ m \text{ iff } f(\alpha)=z)$

Functional Encryption

Security

F etc. adaptively chosen by adversary. Need $F(x^*_0) = F(x^*_1)$ etc.
ABE: $\alpha^*_0 = \alpha^*_1 = \alpha^*$, and
$F_{f,z}$ s.t. $f(\alpha^*) \neq z$

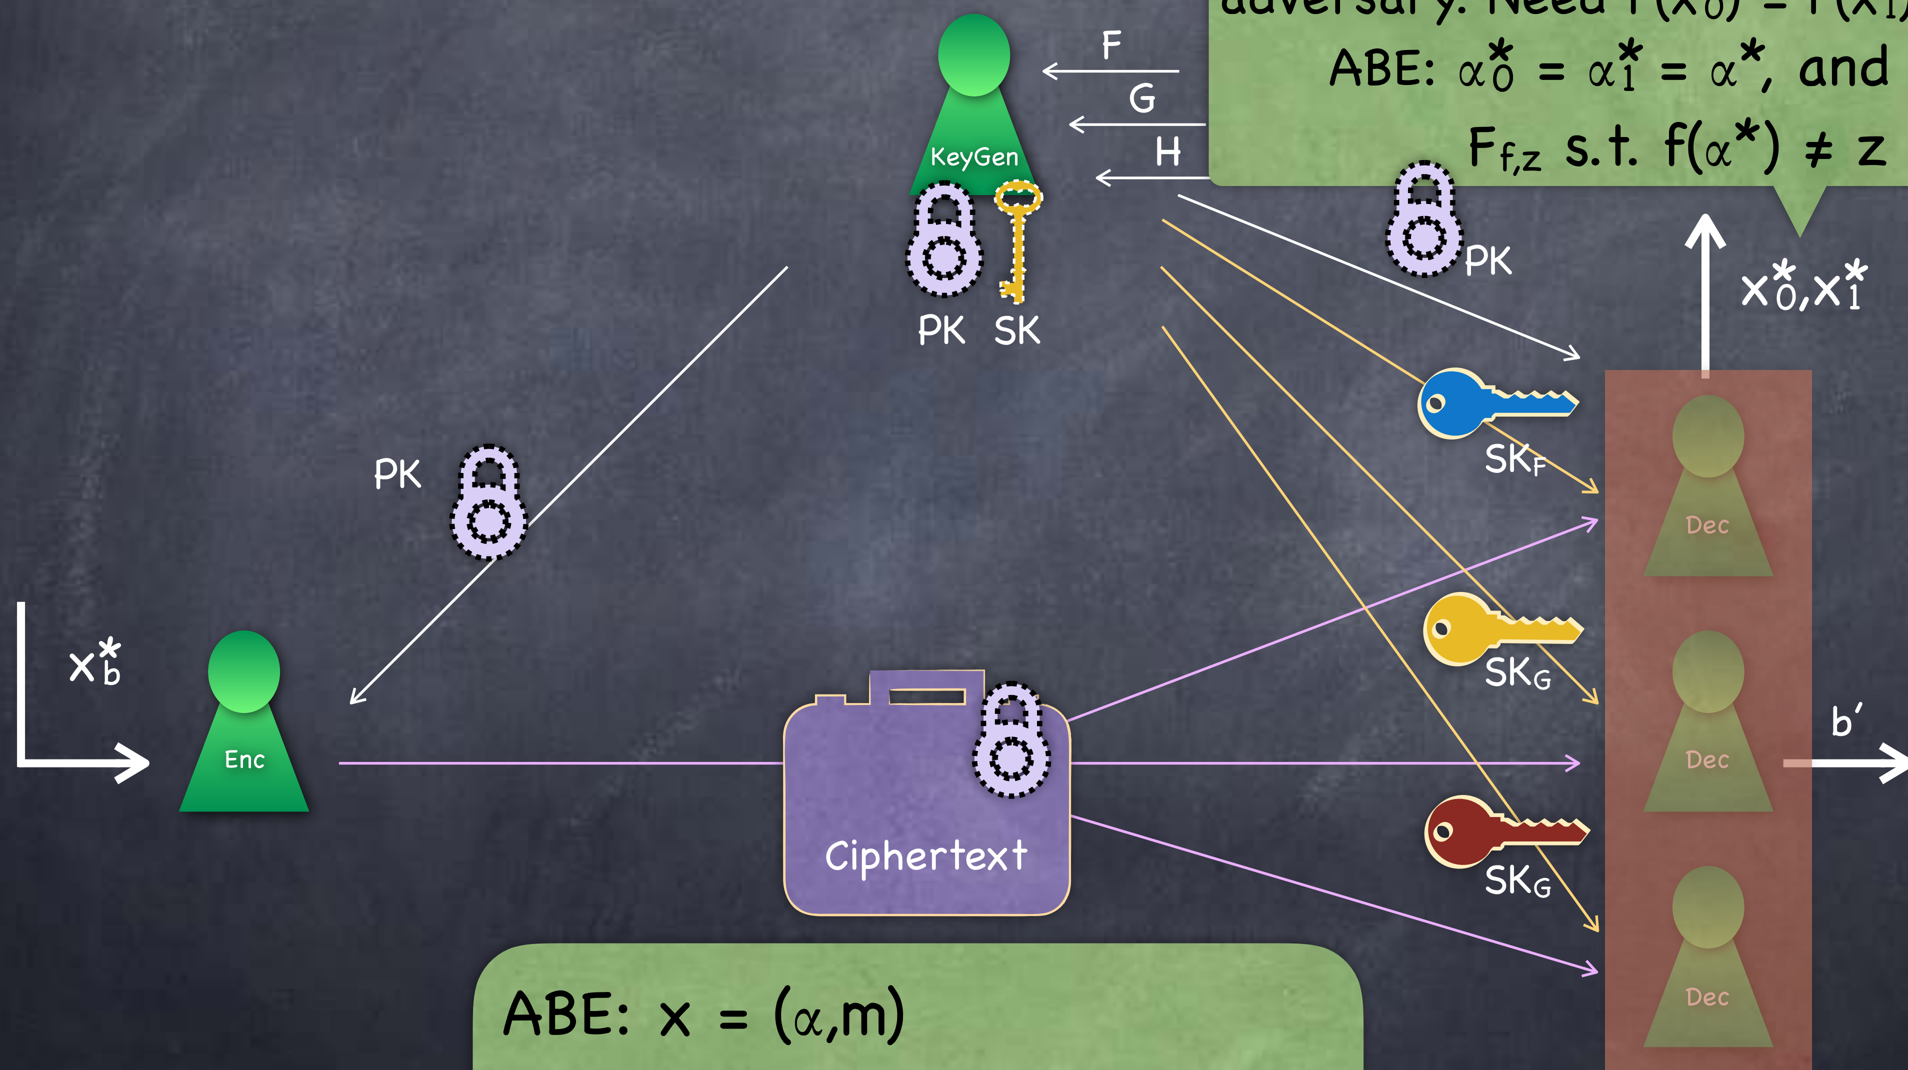

ABE: $x = (\alpha, m)$

$F_{f,z}(x) = (\alpha,\ m \text{ iff } f(\alpha)=z)$

Functional Encryption

Selective Security

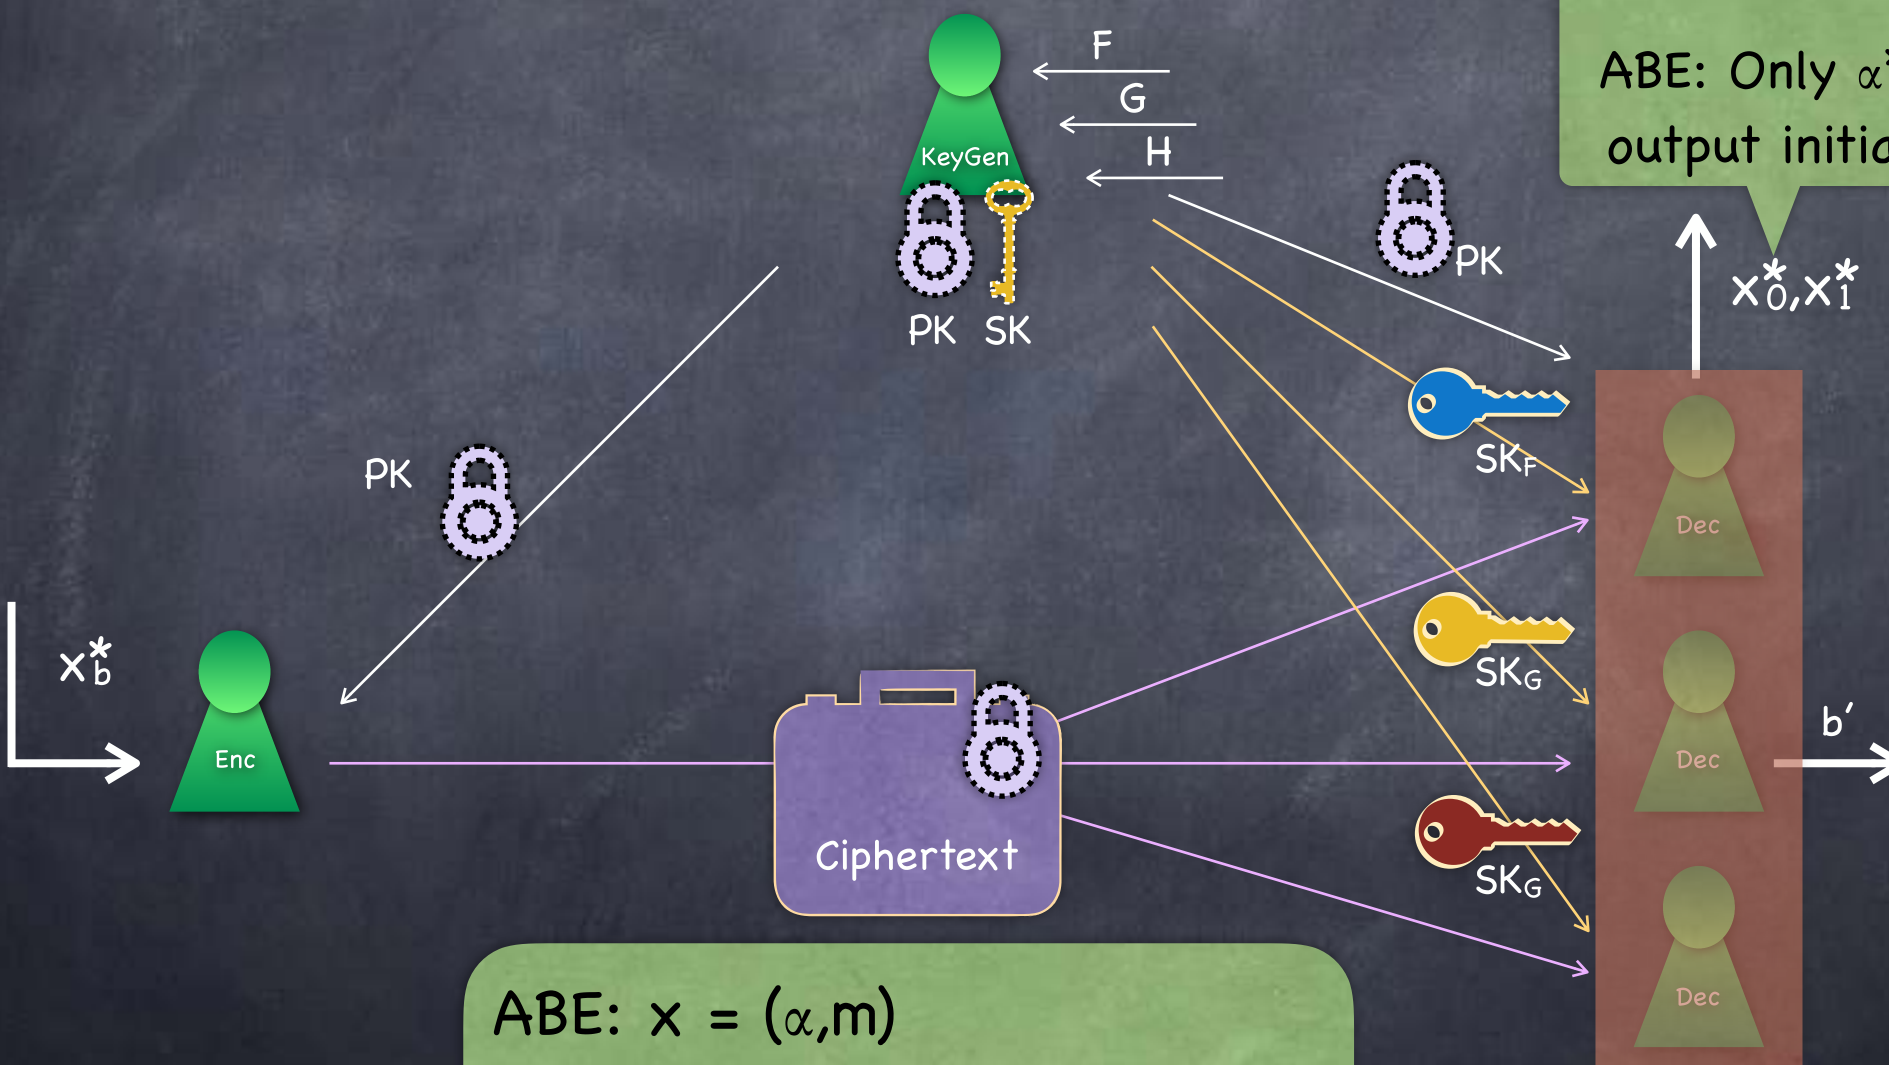

Selective: (x^*_0, x^*_1) output before PK

ABE: Only α^* is output initially

ABE: $x = (\alpha, m)$

$F_{f,z}(x) = (\alpha,\ m \text{ iff } f(\alpha)=z)$

Today: ABE From LWE

- Policy given as an arithmetic circuit f: $\mathbb{Z}_q^t \rightarrow \mathbb{Z}_q$ and a value z. Key $SK_{f,z}$ decrypts ciphertext with attribute α iff $f(\alpha) = z$.
- Very expressive policy $\Rightarrow$ no conceptual distinction between CP-ABE and KP-ABE
  - Can implement CP-ABE also as KP-ABE: α encodes a policy (as bits representing a circuit) and f implements evaluating this policy on attributes hardwired into it

ABE From IBE?

- Policy is (f,z) where f comes from a very large function family
- But instead suppose we had a small number of functions f
- Then enough to have a set of IBE instances one for each f
  - $PK = \{ K_f \}$ one for each f
  - $SK_{f,z}$ = SK for ID z under scheme for f
  - $Enc_{PK}(\alpha,m) = (\alpha, \{ Enc_{K_f}(m;f(\alpha)) \}_f)$
- At a high level, will emulate this idea. But instead of listing K_f and $Enc_{K_f}(m;f(\alpha))$ for each f, will include elements from which any of them can be <u>constructed</u> at the time of decryption
  - Key Homomorphism (BGGHNSVV'14)

Key-Homomorphism

- Overview:
  - Suppose each attribute α has t bits, and f given as a circuit
  - Public key K_f constructed from PK = $\{ K_i \}_{i=1,\ldots,t}$
  - Ciphertext $Enc_{K_f}(m;f(\alpha))$ would be of the form $(Q_{f,f(\alpha)}(s), mask(s)+m)$ where s is randomly chosen
  - $Q_{f,f(\alpha)}(s)$ can be constructed from $\{ Q_{i,\alpha_i}(s) \}_{i=1,\ldots,t}$ (which is included in the actual ciphertext)
  - $SK_{f,z}$ can extract mask(s) from $Q_{f,z}(s)$

ABE From LWE

KeyGen $\leftarrow (f,z)$

$PK = (K_1,\ldots,K_t,K_{mask})$

$SK_{f,z}$ can transform $Q_{f,z}(s)$ into $Mask(s;K_{mask})$

$CT = [\ \alpha,\ Q_{1,\alpha_1}(s),\ldots,\ Q_{t,\alpha_t}(s),$

$m + Mask(s;K_{mask})\ ]$

$(\alpha,m) \rightarrow$ Enc

Dec

If $f(\alpha)=z$, decode $Q_{f,f(\alpha)}$ using $SK_{f,z}$ to get $Mask(s;K_{mask})$

$K_1 \ldots K_t \rightarrow \boxed{PKEval_f} \rightarrow K_f$

$Q_{1,\alpha_1} \ldots Q_{t,\alpha_t} \rightarrow \boxed{CTEval_f} \rightarrow Q_{f,f(\alpha)}$

ABE From LWE

- PK: $K_i = [\ A_0 \mid A_i\ ]$ and $K_{mask} = D$, where $A_0, A_i \leftarrow \mathbb{Z}_q^{n\times m}$, $D \leftarrow \mathbb{Z}_q^{n\times d}$
  - $m \gg n \log q$ so that $A\underline{\mathbf{r}}$ is statistically close to uniform even when $\underline{\mathbf{r}}$ has small entries (e.g., bits)
- Fact: Can pick A along with a trapdoor T_A (a "good" basis for the lattice $L_A{}^{\perp}$) so that, given for any $\underline{\mathbf{u}} \in \mathbb{Z}_q^n$, one can use T_A to sample $\underline{\mathbf{r}}$ with small $\mathbb{Z}_q$ entries (from a discrete Gaussian) that satisfies $A\underline{\mathbf{r}} = \underline{\mathbf{u}}$
  - $\Rightarrow$ sample R with small entries so that AR=D for $D \in \mathbb{Z}_q^{n\times d}$
  - $\Rightarrow$ can sample such an R so that $[\ A \mid B\ ]R = D$, for any B
    - Need $[\ A \mid B\ ]\ [\ R_1 \mid R_2\ ]^T = D$. Sample R_2. Then use T_A to sample $R_1{}^T$ s.t. $AR_1{}^T = D - BR_2{}^T$
- MSK: Trapdoor T_{A_0}

ABE From LWE

Underlying IBE

- PK: $K = [\ A_0 \mid A\ ]$ and $K_{mask} = D$, where $A_0, A \leftarrow \mathbb{Z}_q^{n\times m}$, $D \leftarrow \mathbb{Z}_q^{n\times d}$ and MSK: Trapdoor T_{A_0}

Used for key-homomorphism. Not needed for IBE

- For an identity $z \in \mathbb{Z}_q$ let $K\boxplus z$ denote $[A_0 \mid A + zG]$, where G is the matrix to invert bit decomposition
- Enc(m;z) = $(\ Q_z(\underline{\mathbf{s}}), \text{mask}(\underline{\mathbf{s}}) + \lfloor q/2 \rfloor\, m)$ where $Q_z(\underline{\mathbf{s}}) \approx (K\boxplus z)^T\underline{\mathbf{s}}$ and $\text{mask}(\underline{\mathbf{s}}) \approx D^T\underline{\mathbf{s}}$. Here ≈ stands for adding a small noise (as in LWE)
- SK_z: R_z with small entries s.t. $(K\boxplus z)\, R_z = D$ (computed using T_{A_0})
- Decryption: $R_z^T \cdot Q_z(\underline{\mathbf{s}}) \approx \text{mask}(\underline{\mathbf{s}})$. Recover $m \in \{0,1\}^d$.

ABE From LWE

- PK: $K_i = [\, A_0 \mid A_i \,]$ and $K_{mask} = D$, where $A, A_i \leftarrow \mathbb{Z}_q^{n\times m}$, $D \leftarrow \mathbb{Z}_q^{n\times d}$ and MSK: Trapdoor T_{A_0}
- $K_f = [\, A_0 \mid A_f \,]$ where $A_f = \mathrm{PKEval}(f, A_1, \ldots, A_t)$ (To be described)
- $Q_{i,\alpha_i}(\underline{\mathbf{s}}) \approx (K_i \boxplus \alpha_i)^T \underline{\mathbf{s}}$ where $\underline{\mathbf{s}} \leftarrow \mathbb{Z}_q^n$. (Across all i, same noise used for $A_0^T\underline{\mathbf{s}}$ part.)
- Include $\mathrm{mask}(\underline{\mathbf{s}}) + \lfloor q/2 \rfloor\, m$ in ciphertext, where $\mathrm{mask}(\underline{\mathbf{s}}) \approx D^T\underline{\mathbf{s}}$.
- $Q_{f,f(\alpha)}(\underline{\mathbf{s}}) = \mathrm{CTEval}(f, \alpha, Q_{1,\alpha_1}(\underline{\mathbf{s}}) \ldots, Q_{t,\alpha_t}(\underline{\mathbf{s}})) \approx (K_f \boxplus f(\alpha))^T \underline{\mathbf{s}}$ (To be described)
- $SK_{f,z}$: Compute K_f. Use T_{A_0} to get $R_{f,z}$ s.t. $(K_f \boxplus z)\, R_{f,z} = D$
- Decryption: If $f(\alpha)=z$, then $R_{f,z}^T \cdot Q_{f,f(\alpha)}(\underline{\mathbf{s}}) \approx D^T\underline{\mathbf{s}}$. Recover $m \in \{0,1\}^d$.

ABE From LWE

- $K_f = [\ A_0 \mid A_f\ ]$ where $A_f = \text{PKEval}(f, A_1, \ldots, A_t)$ (To be described)
- $Q_{f,f(\alpha)}(\underline{s}) = \text{CTEval}(f, \alpha, Q_{1,\alpha_1}(\underline{s}) \ldots, Q_{t,\alpha_t}(\underline{s})) \approx (K_f \boxplus f(\alpha))^T \underline{s}$ (To be described)
- CTEval computed gate-by-gate
  - Enough to describe $\text{CTEval}(f_1+f_2, (z_1,z_2), Q_{f_1,z_1}(\underline{s}), Q_{f_2,z_2}(\underline{s}))$ and $\text{CTEval}(f_1 \cdot f_2, (z_1,z_2), Q_{f_1,z_1}(\underline{s}), Q_{f_2,z_2}(\underline{s}))$
  - Recall $Q_{f_1,z_1}(\underline{s}) \approx (K_{f_1} \boxplus z_1)^T \underline{s} = [\ A_0 \mid A_{f_1} + z_1 G\ ]^T \underline{s}$
  - Keep $\approx A_0^T \underline{s}$ aside. To compute $[\ A_{g(f_1,f_2)} + g(z_1,z_2) G\ ]^T \underline{s}$ for $g = +, \cdot$
  - $[\ A_{f_1} + z_1 G\ ]^T \underline{s} + [\ A_{f_2} + z_2 G\ ]^T \underline{s} = [\ A_{f_1+f_2} + (z_1 + z_2)\ G\ ]^T \underline{s}$ with $A_{f_1+f_2} = A_{f_1} + A_{f_2}$ (errors add up)
  - $A_{f_1 \cdot f_2}$
  - $z_2 \cdot [\ A_{f_1} + z_1 G\ ]^T \underline{s} - B(A_{f_1})^T [\ A_{f_2} + z_2 G\ ]^T \underline{s} = [-A_{f_2} B(A_{f_1}) + z_1 z_2 G]^T \underline{s}$
    - $\text{err} = z_2 \cdot \text{err}_1 + B(A_{f_1})^T \text{err}_2$. Need z_2 to be small.

ABE From LWE

- Security?
- Sanity check: Is it secure when no function keys $SK_{f,z}$ are given to the adversary?
- Security from LWE
  - All components in the ciphertext are LWE samples of the form $\langle \underline{\mathbf{a}}, \underline{\mathbf{s}} \rangle$+noise, for the same $\underline{\mathbf{s}}$ and random $\underline{\mathbf{a}}$.
  - Hence all pseudorandom, including the mask $D^T\underline{\mathbf{s}}$ + noise
- Do the secret keys $SK_{f,z}$ make it easier to break security?
- Claim: No!

ABE From LWE

- Scheme is <u>selective-secure</u> (under LWE)
- Recall selective security for ABE:
  Adversary first outputs α^* first, before seeing PK.
  Then obtains keys $SK_{f,z}$ for $F_{f,z}$ s.t. $f(\alpha^*) \neq z$.
  Gives $x_0^* = (\alpha^*, m_0)$ and $x_1^* = (\alpha^*, m_1)$ and gets challenge $Enc(x^*_b)$.
- Simulated execution (indistinguishable from real) where PK* is designed such that without MSK* can generate $SK_{f,z}$ for all f and all $z \neq f(\alpha^*)$
  - Breaking encryption for α^* will still need breaking LWE!

ABE From LWE

- Simulated execution (indistinguishable from real) where PK* is designed such that without MSK* can generate $SK_{f,z}$ for all (f,z) s.t. $z \neq f(\alpha^*)$
  - D, A_0 as before but without trapdoor (i.e., given from outside)
  - Other keys A_i are (differently) trapdoored: $A_i^* = A_0S_i - \alpha^*_iG$ where S_i have small entries
    - A_0S_i close to uniform (like A_i) by extraction argument
  - Consider a query (f,z) where $z \neq f(\alpha^*) =: z^*$
    - Need to give $R_{f,z}$ s.t. $(K_f \boxplus z)\, R_{f,z} = D$
    - Do not have a the trapdoor for $K_f = [\ A_0 \mid A_f - z^*G\ ]$
    - Will use a trapdoor for $A_f - z^*G$ instead!

Two Trapdoors

- Given A_0, $B \in \mathbb{Z}_q^{n\times m}$ of rank n, and D, can find small R s.t. $[\, A_0 \mid B \,]\, R = D$ if we have:

  a "small" basis T_{A_0} for $\Lambda^\perp{}_{A_0}$

  - Either the trapdoor T_{A_0} for sampling small R_0 s.t. $A_0R_0 = U$
  - Or (S, T_{B-A_0S}) s.t. $B - A_0S$ has full rank and S "small"
    - E.g., small S s.t. $B = A_0S + z'G$ for $z' \neq 0$ and G has a known trapdoor T_G (which is also a trapdoor for $z'G$)
- In the actual construction, we used the fact that (A_0, T_{A_0}) can be generated together, to be able to give out function keys $R_{f,z}$. (A_i picked randomly, resulting in random A_f.)
- In the security proof, given an A_0 from outside, will construct $A^*_i = A_0S_i - \alpha_i{}^*G$ and maintain $A^*_f = A_0S_f - f(\alpha^*)G$. Then, if $z \neq f(\alpha^*)$ and so $B = A^*_f + zG = A_0S_f + z'G$ for $z' = z - f(\alpha^*) \neq 0$, can sample $R_{f,z}$.

Simulation of Keys

- Simulated KeyGen (given α^*) produces keys which are statistically close to the original keys
  - Public Key: Accepts A_0 from outside. Picks $A_i^* = A_0S_i - \alpha^*_iG$ where S_i have small entries.
    - For each f, A_f^* defined by EvalPK: $A_f^* = A_0S_f - f(\alpha^*)G$
  - Function Keys: Given (f,z) s.t. $z \neq f(\alpha^*)$, $R_{f,z}$ s.t. $(K_f^* \boxplus z)\, R_{f,z} = D$.
    - $A_f^* \boxplus z = [\, A_0 \mid A_f^* + zG] = [\, A_0 \mid A_0S_f - f(\alpha^*)G + zG]$
      $= [\, A_0 \mid A_0S_f + z'G]$ where $z' \neq 0$
    - S_f remains small (assuming $f_2(\alpha^*)$ is small in products $f_1 \cdot f_2$ in the circuit for computing $f(\alpha^*)$)
    - So can sample small $R_{f,z}$ as required (type 2 trapdoor)
- Simulated keys are statistically indistinguishable from the keys in the real experiment

Simulation of Ciphertext

- Accepts $\approx A_0^T\underline{\mathbf{s}}$ and $\approx D^T\underline{\mathbf{s}}$ from outside, and produces a ciphertext (corresponding to the given $\underline{\mathbf{s}}$, but without knowing $\underline{\mathbf{s}}$)
  - Need $Q_{i,\alpha^*_i}(\underline{\mathbf{s}}) \approx (K^*_i \boxplus \alpha^*_i)^T\underline{\mathbf{s}}$ and $\mathrm{mask}(\underline{\mathbf{s}}) \approx D^T\underline{\mathbf{s}}$
    - For $Q_{i,\alpha^*_i}(\underline{\mathbf{s}})$, need $\approx (A_i^* + \alpha^*_i G)^T\underline{\mathbf{s}} = (A_0 S_i)^T\underline{\mathbf{s}} = S_i^T A_0^T\underline{\mathbf{s}}$.

      Can derive this from $\approx A_0^T\underline{\mathbf{s}}$ and S_i ($S_i^T\cdot$noise is fresh noise)
- Simulated $Q_{i,\alpha^*_i}(\underline{\mathbf{s}})$ and $\mathrm{mask}(\underline{\mathbf{s}})$ are statistically indistinguishable from the real experiment (conditioned on the keys)
- But if $\approx A_0^T\underline{\mathbf{s}}$ and $\approx D^T\underline{\mathbf{s}}$ are replaced by random vectors, then:
  - No information about the message (because random mask)
  - Indistinguishable from the simulation above (by LWE)
    - In turn statistically indistinguishable from the real experiment