Functional Encryption

Lecture 23

ABE from LWE

Functional Encryption

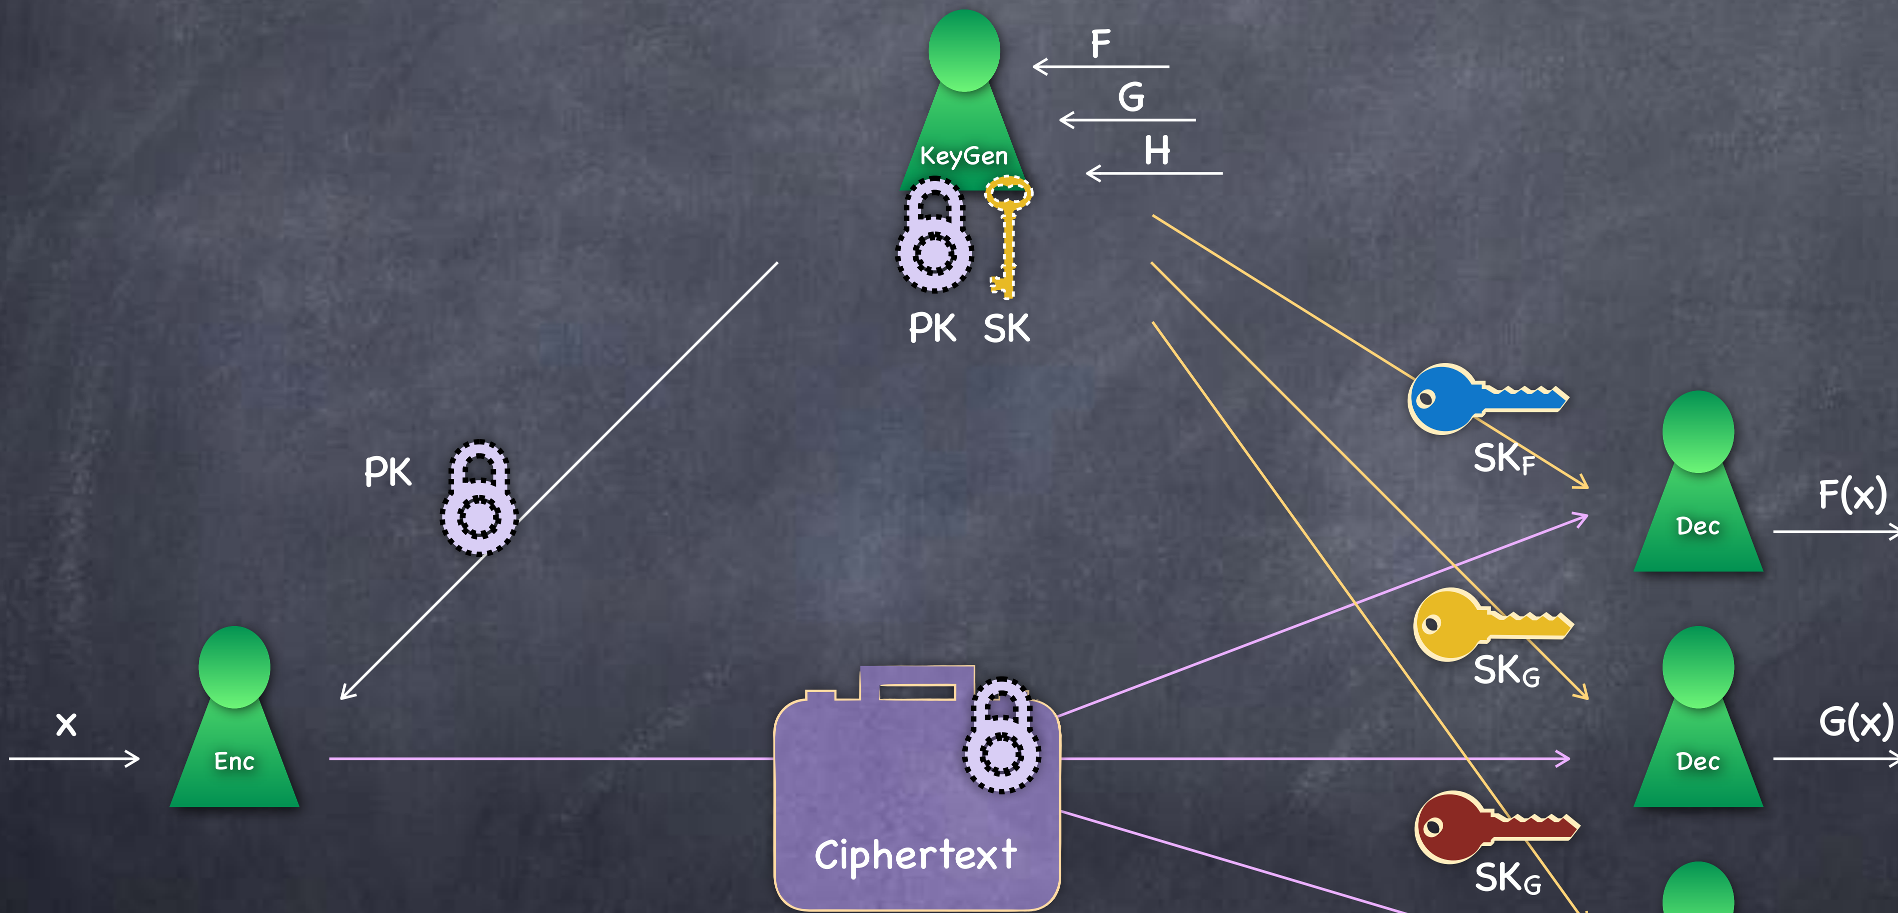

ABE: $x = (a,m)$

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

Functional Encryption

Security

F etc. adaptively chosen by adversary. Need $F(x_0^*) = F(x_1^*)$ etc.

ABE: $a_0^* = a_1^* = a^*$, and $F_{f,z}$ s.t. $f(a^*) \neq z$

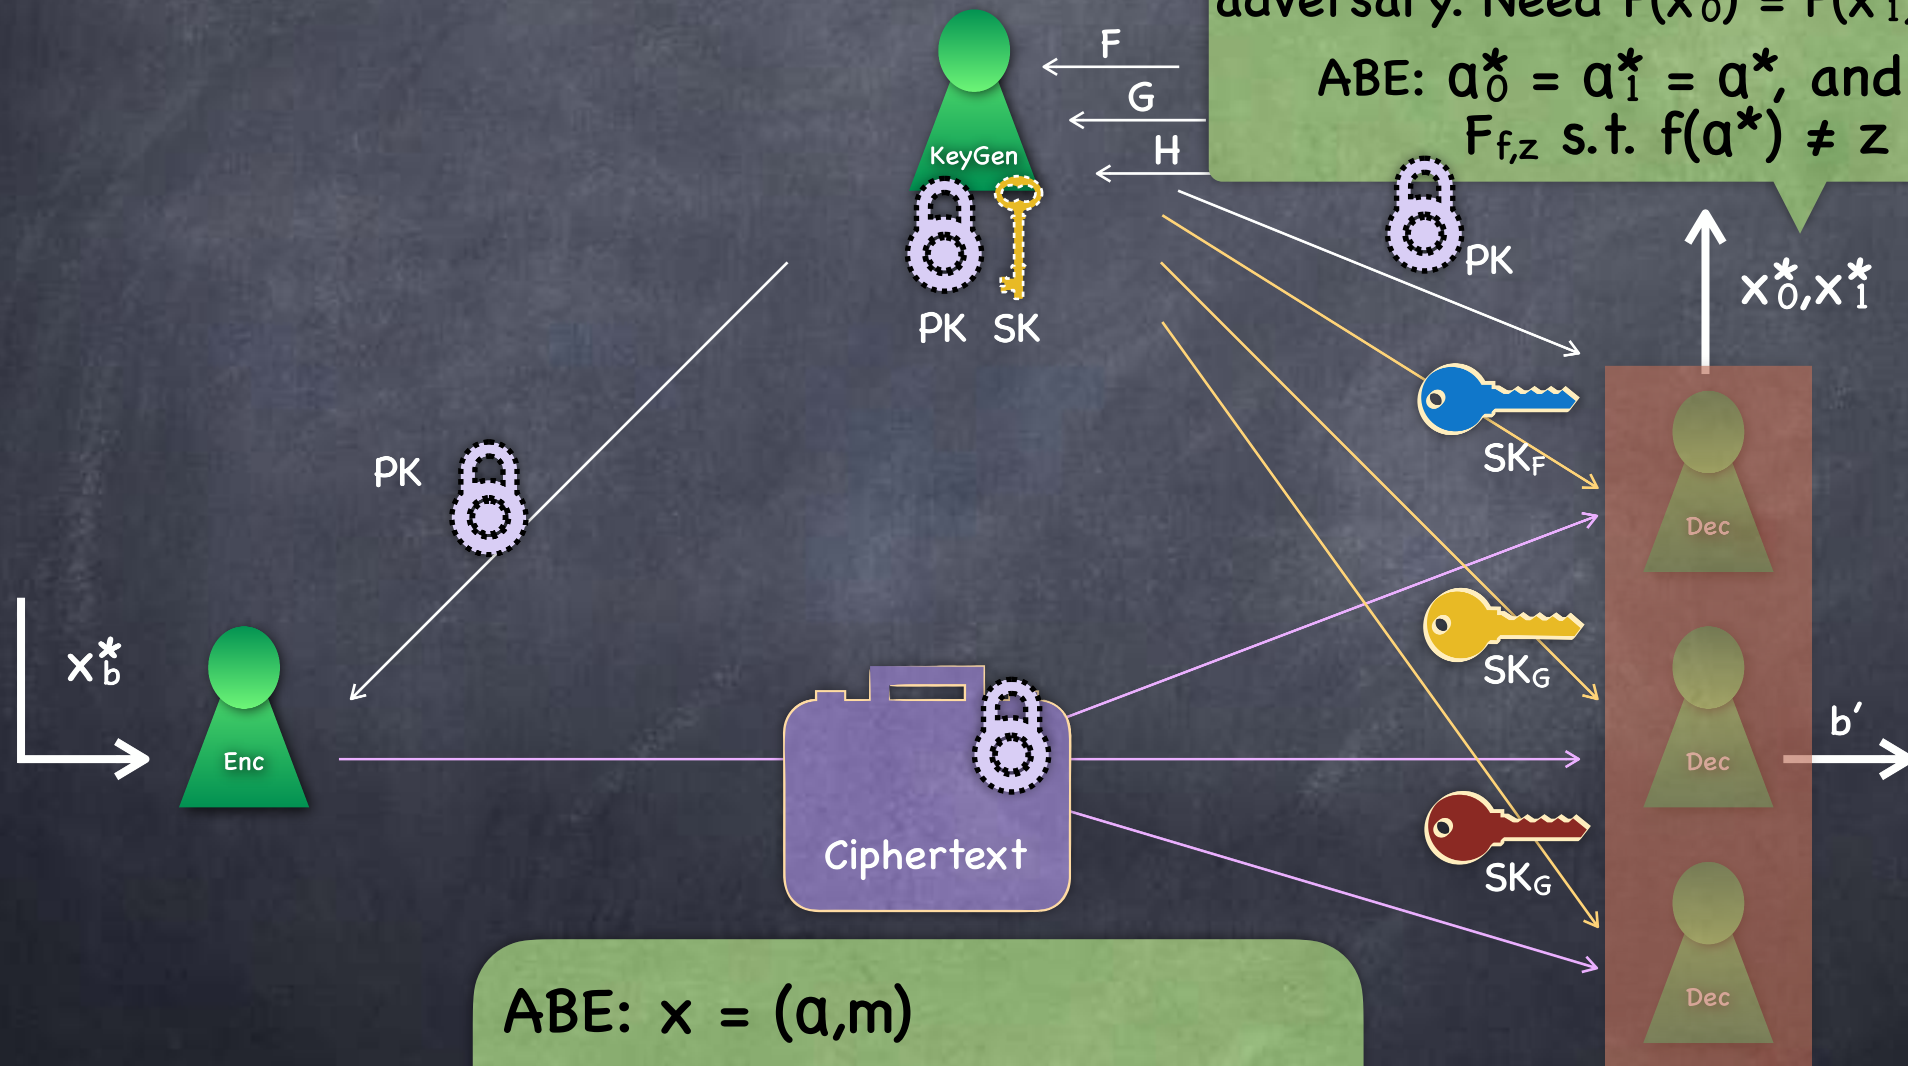

ABE: $x = (a,m)$

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

Functional Encryption

Selective Security

Selective: (x_0^*, x_1^*) output before PK

ABE: Only a^* is output initially

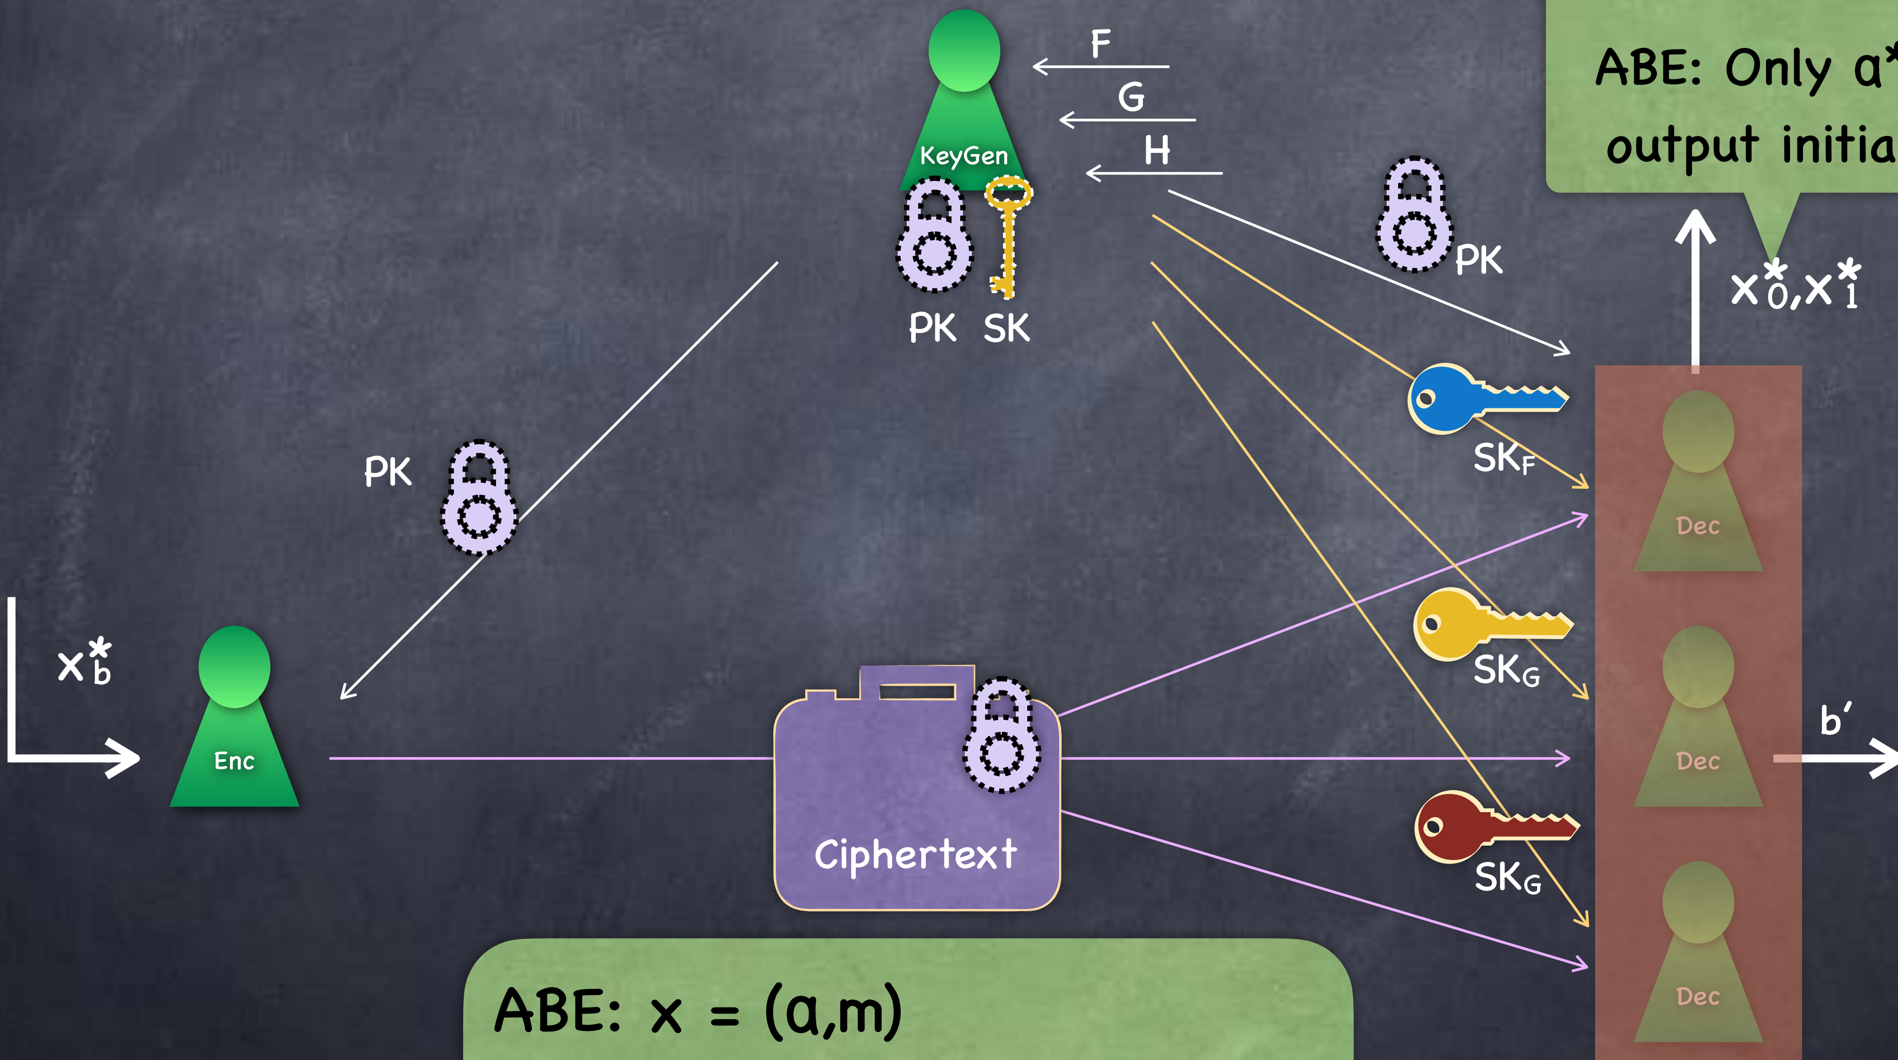

ABE: $x = (a,m)$

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

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?

- Key-policy is (f,z) where f comes from a very large function family
- But instead suppose we had a small number of functions f (but z comes from an exponentially large range)
- 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 $Enc_{K_f}(m;f(\alpha))$ for each f, will include elements from which any of them can be constructed at the time of decryption
  - Key Homomorphism (BGGHNSVV'14)

Key-Homomorphism

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

ABE From LWE

(f,z)

KeyGen

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

K_f

$PKEval_f$

$K_1 \ \ldots \ K_t$

$SK_{f,z}$ can recover

$Mask(s;K_{mask})$ from $Q_{f,z}(s)$

$CT = [\ a,\ Q_{1,a_1}(s),\ldots,\ Q_{t,a_t}(s),$

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

$Q_{f,f(a)}$

$CTEval_f$

$Q_{1,a_1} \ \ldots \ Q_{t,a_t}$

(a,m)

Enc

Dec

If $f(a)=z$, decode $Q_{f,f(a)}$

using $SK_{f,z}$ to get

$Mask(s;K_{mask})$

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 $A\underline{r}$ is statistically close to uniform even when $\underline{r}$ has small entries (e.g., bits)

a "small" basis for $\Lambda_A^{\perp}$

- Fact: Can pick A along with a trapdoor T_A so that, given $\underline{u} \in \mathbb{Z}_q^n$, one can use T_A to sample $\underline{r}$ with small $\mathbb{Z}_q$ entries s.t. $A\underline{r} = \underline{u}$
  - $\Rightarrow$ 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 $[\ A \mid H\ ]R = D$, for any H, D
    - Need $[\ A \mid H\ ]\,[\ R_1 \mid R_2\ ]^T = D$. Sample R_2. Then use T_A to sample R_1^T s.t. $AR_1^T = D - HR_2^T$
- 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
- $\mathrm{Enc}(m;z) = (\, Q_z(\underline{s}),\ \mathrm{mask}(\underline{s}) + \lfloor q/2 \rfloor\, m)$ where $Q_z(\underline{s}) \approx (K \boxplus z)^T \underline{s}$ and $\mathrm{mask}(\underline{s}) \approx D^T \underline{s}$

Using ≈ to denote 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: $R_z^T \cdot Q_z(\underline{s}) \approx \mathrm{mask}(\underline{s})$. Recover $m \in \{0,1\}^d$.

ABE From LWE

- PK: $K_i = [\ A_0 \mid A_i\ ]$ and $K_{mask} = D$, where $A_0, A_i \leftarrow \mathbb{Z}_q^{n\times m}$, $D \leftarrow \mathbb{Z}_q^{n\times d}$ and MSK: Trapdoor T_{A_0}
- $Q_{i,a_i}(\underline{s}) \approx (K_i \boxplus a_i)^T\underline{s}$ where $\underline{s} \leftarrow \mathbb{Z}_q^n$.

  $\uparrow$ Across all i, use same $\approx A_0^T\mathbf{s}$ part.
- $CT = (\{Q_{i,a_i}(\underline{s})\}_i,\ mask(\underline{s}) + \lfloor q/2 \rfloor m)$, where $m \in \{0,1\}^d$, $mask(\underline{s}) \approx D^T\underline{s}$
- $K_f = [\ A_0 \mid A_f\ ]$ where $A_f = PKEval(f,A_1,\ldots,A_t)$ (To be described)
- $Q_{f,f(a)}(\underline{s}) = CTEval(f,a,Q_{1,a_1}(\underline{s})\ldots,Q_{t,a_t}(\underline{s})) \approx (K_f \boxplus f(a))^T\underline{s}$ (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: Compute $Q_{f,f(a)}(\underline{s})$. If $f(a)=z$, then $R_{f,z}^T \cdot Q_{f,f(a)}(\underline{s}) \approx D^T\underline{s}$. Recover $m \in \{0,1\}^d$.

ABE From LWE

- $K_f = [\ A_0 \mid A_f\ ]$ where $A_f = \mathrm{PKEval}(f, A_1, \ldots, A_t)$ (To be described)
- $Q_{f,f(a)}(\underline{s}) = \mathrm{CTEval}(f, a, Q_{1,a_1}(\underline{s}) \ldots, Q_{t,a_t}(\underline{s})) \approx (K_f \boxplus f(a))^T \underline{s}$ (To be described)
- CTEval computed gate-by-gate
  - Enough to describe $\mathrm{CTEval}(f_1+f_2, (z_1,z_2), Q_{f_1,z_1}(\underline{s}), Q_{f_2,z_2}(\underline{s}))$ and $\mathrm{CTEval}(f_1 \cdot f_2, (z_1,z_2), Q_{f_1,z_1}(\underline{s}), Q_{f_2,z_2}(\underline{s}))$
  - 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)
  - $z_2 \cdot [\ A_{f_1} + z_1 G\ ]^T \underline{s} - B(A_{f_1})^T [\ A_{f_2} + z_2 G\ ]^T \underline{s} = [\overbrace{-A_{f_2} B(A_{f_1})}^{A_{f_1 \cdot f_2}} + z_1 z_2 G]^T \underline{s}$
    - $\mathrm{err} = z_2 \cdot \mathrm{err}_1 + B(A_{f_1})^T \mathrm{err}_2$. 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 selective-secure (under LWE)
- Recall selective security for ABE:
  - Adversary first outputs α^*, before seeing PK
  - Then obtains keys $SK_{f,z}$ s.t. $f(\alpha^*) \neq z$
  - Gives $x_0^* = (\alpha^*, m_0)$ and $x_1^* = (\alpha^*, m_1)$ and gets challenge $Enc(x_b^*)$
- Plan: Simulated execution (indistinguishable from real) where PK^* is designed such that, without MSK^*, one 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

- Plan: Simulated execution (indistinguishable from real) where PK* is designed such that, without MSK*, one 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 trapdoor for $K_f = [\, A_0 \mid A_f - z^*G \,]$
    - Will use a trapdoor for $A_f - z^*G$ instead!

Two Trapdoors

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

  a "small" basis 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_{H-A_0S}) s.t. $H - A_0S$ has full rank and S "small"
    - E.g., small S s.t. $H = A_0S + z'G$ 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 $A^*_i = A_0S_i - a_i^*G$ and maintain $A^*_f = A_0S_f - f(a^*)G$. Then, if $z \neq f(a^*)$ and so $H = A^*_f+zG = A_0S_f + z'G$ for $z' = z-f(a^*) \neq 0$, 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.
    - Given f, A_f^* defined by PKEval (& S_f s.t. $A_f^* = A_0S_f - f(\alpha^*)G$)
  - Function Keys: Given (f,z) s.t. $z \neq f(\alpha^*)$, $R_{f,z}$ s.t. $(K_f^* \boxplus z) R_{f,z} = D$.
    - $K_f^* \boxplus z = [\ A_0 \mid A_f^* + zG] = [\ A_0 \mid A_0S_f - f(\alpha^*)G + zG]$
      $= [\ A_0 \mid A_0S_f + z'G]$ 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 $\approx A_0^T\underline{s}$ and $\approx D^T\underline{s}$ from outside, and produces a ciphertext (corresponding to the given $\underline{s}$, but without knowing $\underline{s}$)
  - Need $Q_{i,a^*_i}(\underline{s}) \approx (K^*_i \boxplus a^*_i)^T\underline{s}$ and mask($\underline{s}$) $\approx D^T\underline{s}$
    - For $Q_{i,a^*_i}(\underline{s})$, need $\approx (A_i^* + a^*_iG)^T\underline{s} = (A_0S_i)^T\underline{s} = S_i^TA_0^T\underline{s}$. Can derive this from $\approx A_0^T\underline{s}$ and S_i ($S_i^T\cdot$noise is fresh noise)
- Simulated $Q_{i,a^*_i}(\underline{s})$ and mask($\underline{s}$) are statistically indistinguishable from the real experiment (conditioned on the keys)
- But if $\approx A_0^T\underline{s}$ and $\approx D^T\underline{s}$ 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