Multi-Authority ABE for DNFs from LWE

Pratish Datta, Ilan Komargodski, Brent Waters

NTT Research, Hebrew University, UT Austin

Lattices Reunion, Simons Institute

Traditional Encryption

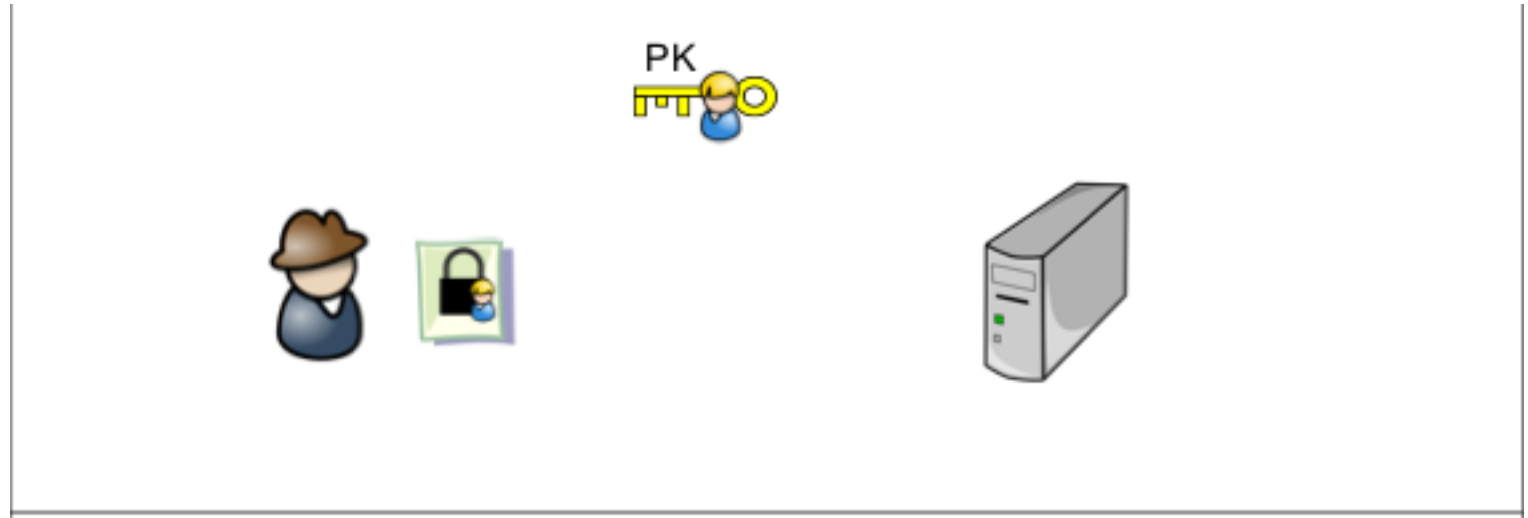

Traditional Encryption

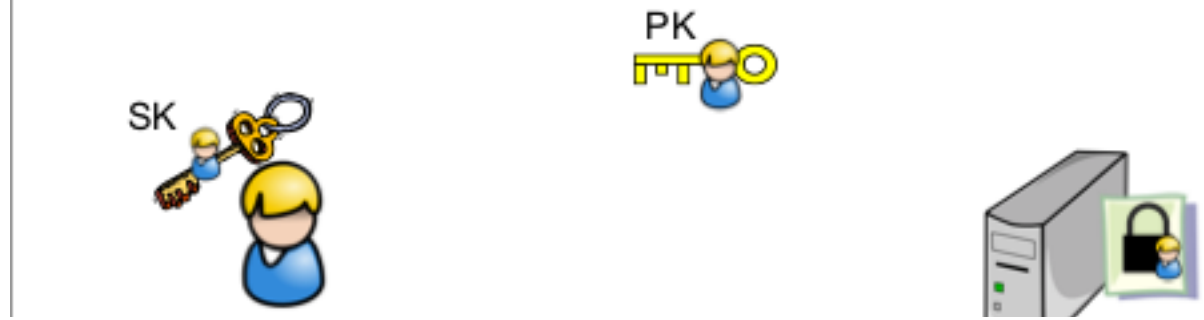

Attribute-Based Encryption

[SW05,GPSW06,...]

- A secret key allows one to decrypt all messages that satisfy

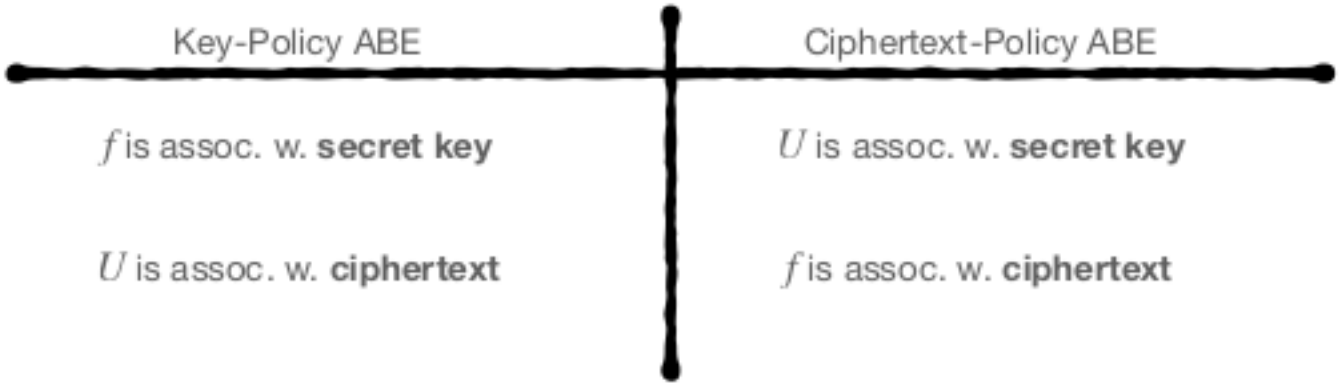

Attribute-Based Encryption

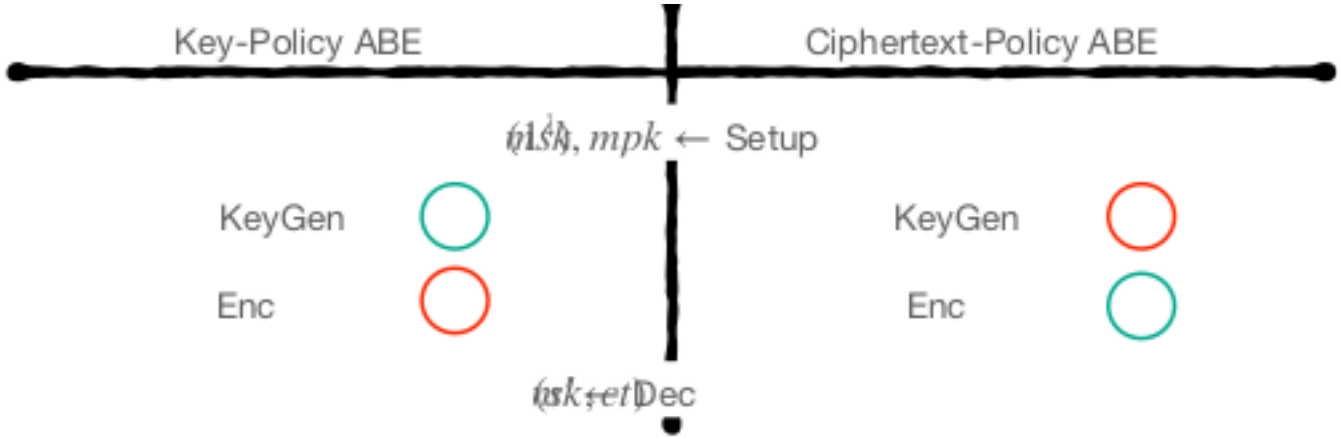

Correctness: If $f(U) = 1$ then $m' = m$

Security: If $f(U) = 0$ then m is "hidden"

Attribute-Based Encryption

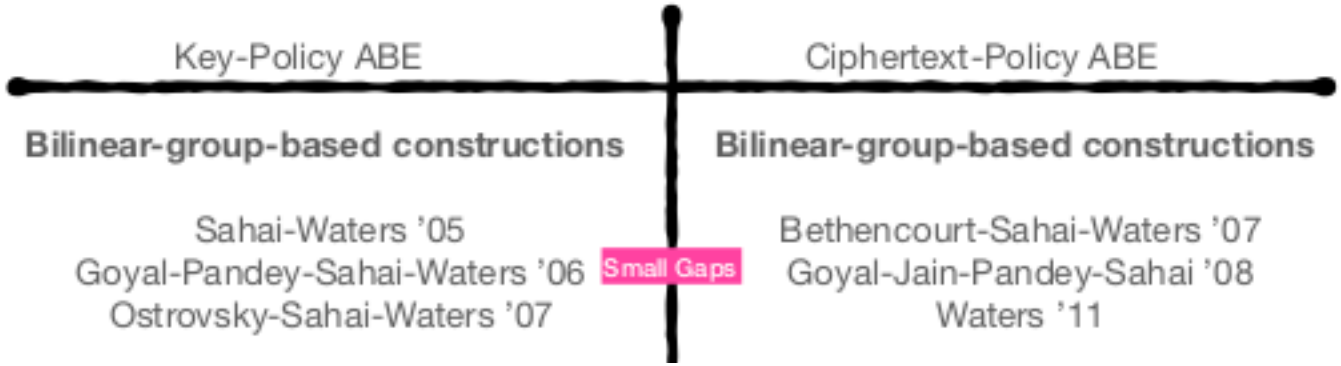

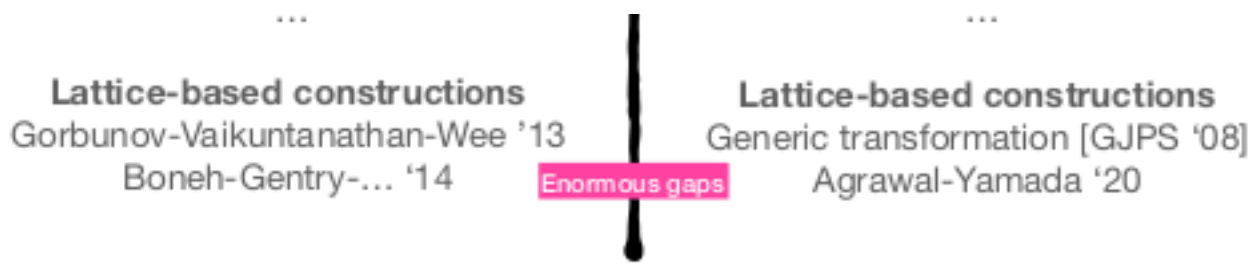

Multi-Authority ABE

- In ABE, *one central authority who verifies attributes and issues secret keys*
- In reality, multiple authorities are in charge of different attributes
  - DMV for "holds a driving license"
  - Military for "veteran"
  - University for "holds a Ph.D"
- "Multi-Authority" ABE
  - Chase '07, ..., Lewko-Waters '11, Okamoto-Takashima '13, Rouselakis-Waters '15

MA-ABE

- Anyone can become an authority
  - No coordination except global PublicParams
- Different authorities control different attributes
- No bound on # of authorities
- Each authority can issue secret keys to users possessing attributes under their control
  - without any interaction with other authorities

MA-ABE

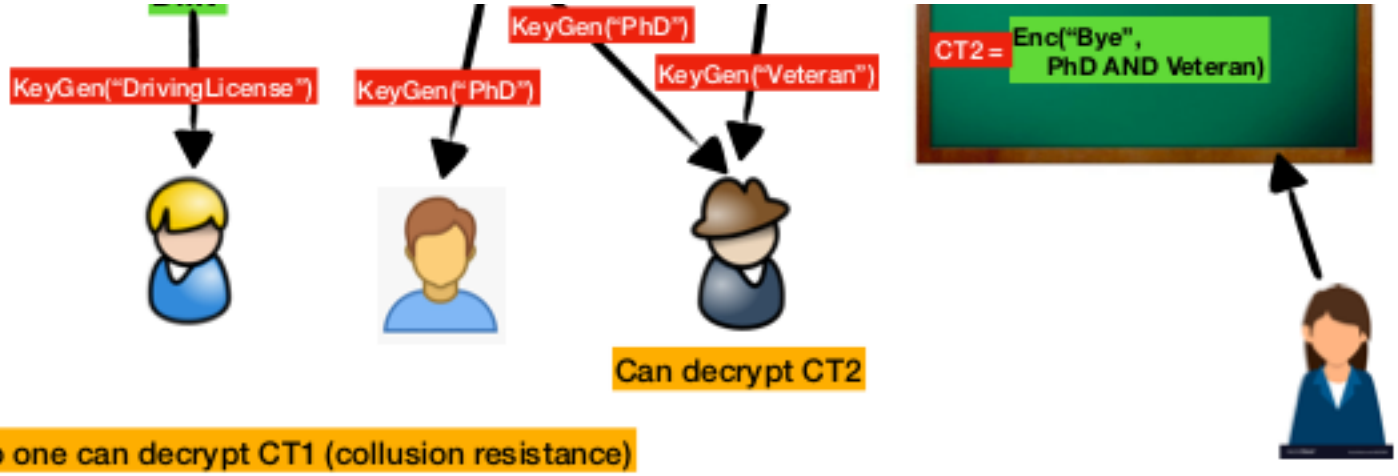

No one can decrypt CT1 (collusion resistance)

The GID Model

Chase '07

- How to uniquely identify a user?
- Associate a unique global *verifiable* identifier (GID)
- The global identity of a user remains fixed for the entire lifetime of the system
- Users have no freedom to choose their global identities

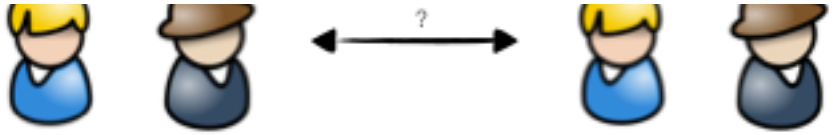

MA-ABE Syntax

(Assume one attribute per authority)

$(pp) \leftarrow$ GlobalSetup

AuthSetup

KeyGen

Enc

$(\{sk_{GID,i}\}, ct)$

(All with same GID)

Correctness: If $f(U) = 1$ then $m' = m$

Security: If $f(U) = 0$ then m is "hidden"

MA-ABE (Main) Known Constructions

& Our Main Result

	Supported Policy Class	Assumption
Lewko-Waters '11	NC1	Subgroup decision (composite order)
Okamoto-Takashima '13	NC1	DLIN (prime order)
Rouselakis-Waters '15	NC1	q-type (prime order)

(All schemes are in the random oracle model)

Our Security Definition

(Static Security)

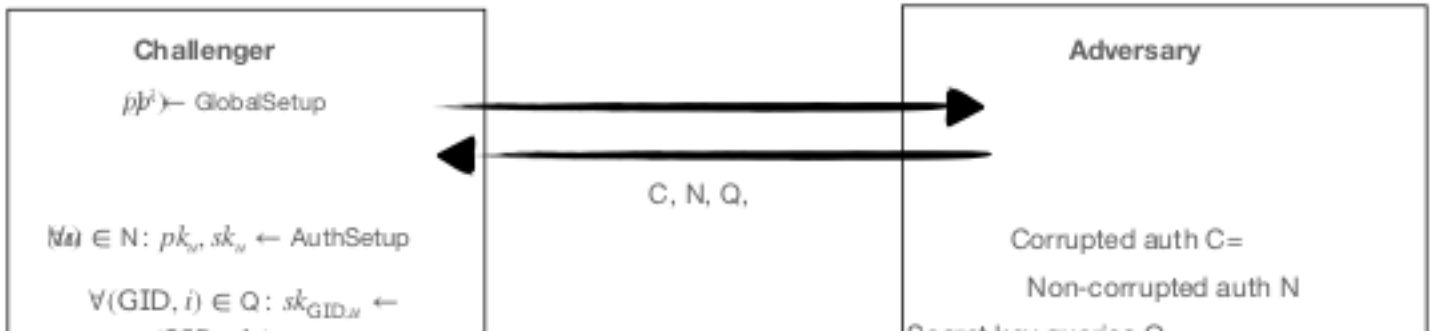

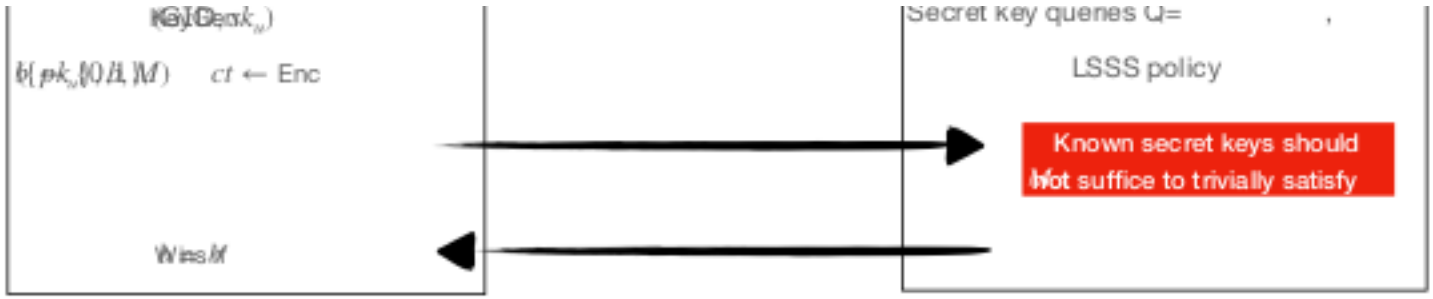

Our Main Theorem

DNF = OR of ANDs

$(x \wedge y \wedge z) \vee (x \wedge \neg y \wedge \neg z) \vee (\neg x \wedge y \wedge \neg z)$

There exist an MA-ABE for access policies captured by DNF formulas.

Our scheme is (statically) secure against an arbitrary collusion of parties in the random oracle model and assuming the LWE assumption.

sub-exp. modulus-to-noise ratio

Open Problems

- "More adaptive" security?
- NC1? P?
- Better efficiency parameters?

Challenges

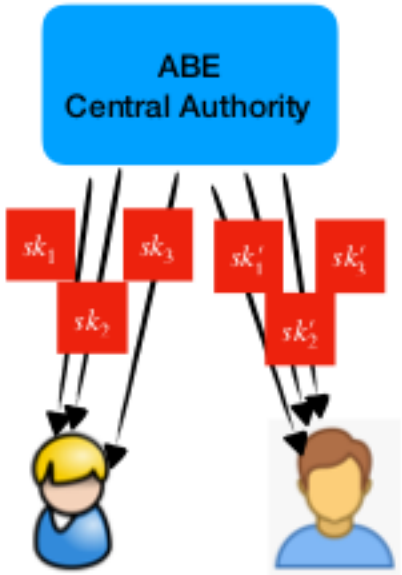

Collusion resistance is obtained by using fresh randomness for every sk specific to the user

sk_is (and sk'_is) are compatible with each other, but not intertwined

MA challenge 1

The randomness used to tie together different key components is obtained from $H(\mathrm{GID})$

Randomness is essentially public

MA challenge 2

Should support arbitrary authorities joining on the fly

Components should be "piecewise". Every authority pk and user sk should be associated with its own attributes

(Non-Monotone) Linear Secret Sharing Schemes

[Shamir, KW93...]

- A secret sharing scheme where sharing & reconstruction are linear functions
- Equivalent to span programs

- Equivalent to *span programs*

$$M \in \mathbb{Z}_q^{\ell \times s}$$

$$\rho : [\ell] \mapsto \text{parties} \cup \neg\text{parties}$$

Monotone / non-monotone

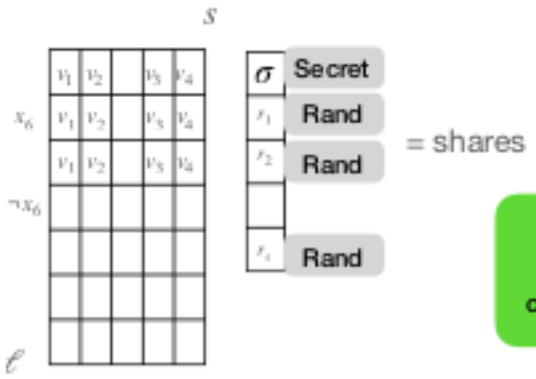

Reconstruction by appropriate linear combination of shares

Our (Non-Monotone) Linear Secret Sharing Schemes

- **Small reconstruction coefficients:** Reconstruction of the secret can be done by small coefficients, i.e., coming from .
- **Linear independence for unauthorized rows:** Shares of an unauthorized set are linearly independent.

Agrawal et al. '20, Lewko-Waters '11

Theorem: There exists such a monotone LSSS* for DNFs

Theorem:
There exists such a non-monotone LSSS* for LOGSPACE (implicit in GVW13)

A different construction

for NC1 in the paper

Open: monotone LSSS for NC1?

Share size ≈ circuit size

The Recipe

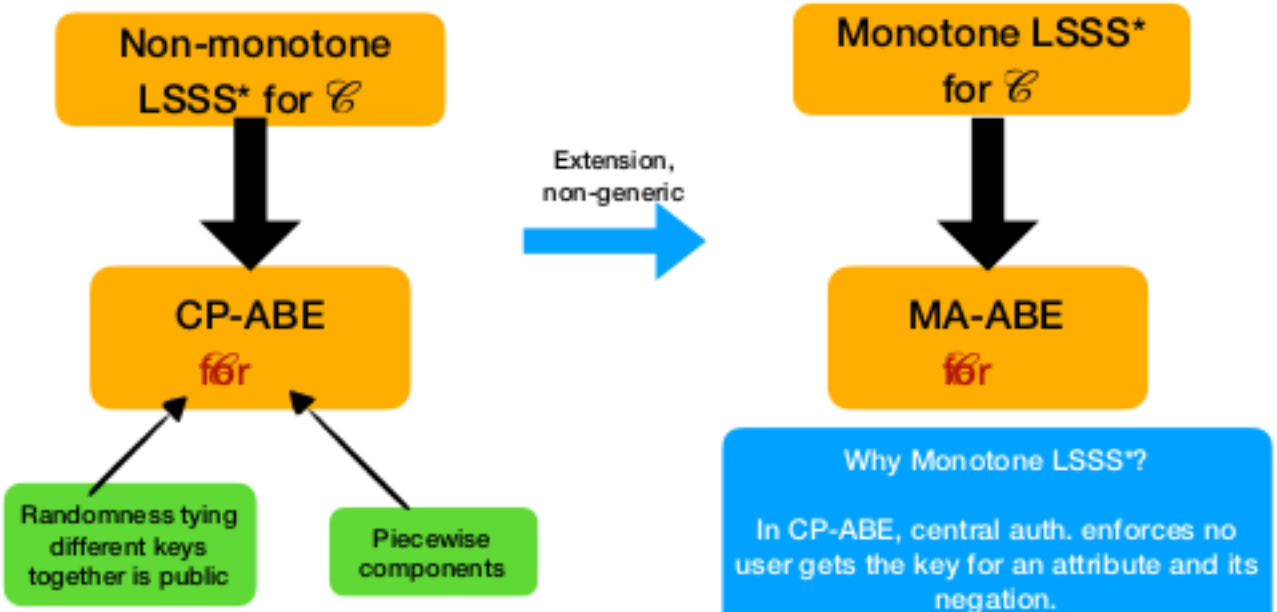

The CP-ABE Scheme

Setup:

For each attribute , sample

- together with a trapdoor
-
-

Output:

,

The CP-ABE Scheme

KeyGen(**)** : // is a set of attributes

Sample and set

For each attribute , sample

Output:

Public nature is important for MA-ABE

t is the (public) randomness that ties together different key components

The CP-ABE Scheme

Enc( ): //

Sample

Compute

Output

The CP-ABE Scheme

Dec( ): //

— Indices of rows of available attributes

— Reconstruction coefficients

Compute:

Output:

Correctness

(Ignoring small noise-like terms)

$$K' = \sum_{i \in I} w_i (\mathbf{c}_i \tilde{\mathbf{k}}_{\rho(i)}^{\top} + \hat{\mathbf{c}}_i \mathbf{t}^{\top})$$

$$msg' = C \oplus \mathrm{MSB}(K')$$

Recall

Reconstruction gives

Security

- Need: challenge independent of
- Get-rid of in and

$$\mathbf{c}_i = \mathbf{s}\mathbf{A}_{\rho(i)} + \text{noise}$$

$$\hat{\mathbf{c}}_i = M_i \cdot \begin{bmatrix} \mathbf{s}\mathbf{y}^{\perp} & 0 & \dots & 0 \\ & \mathbf{v}_2 & & \\ & \mathbf{v}_3 & & \end{bmatrix} - \mathbf{s}\mathbf{H}_{\rho(i)} + \text{noise}$$

- is by programming
- is by LWE
  - Need to get rid of

$$\mathbf{v}_s \end{bmatrix}$$

$\mathrm{MSB}(\mathbf{sy}^\top) \oplus \mathrm{msg}$

Programming

$$\mathbf{H}_{\rho(i)} = M_i \cdot \begin{bmatrix} \mathbf{y}^\perp & \mathbf{0}^\top & \dots & \mathbf{0}^\top \\ & \mathbf{B}_2 & & \\ & \mathbf{B}_3 & & \\ & \dots & & \\ & \mathbf{B}_s & & \end{bmatrix} + \mathbf{A}_{\rho(i)}\mathbf{R}_{\rho(i)}$$

$(\mathbb{Z}_q^{n \times m})^{s \times 1}$

$\{-1,1\}^{m \times m}$

Indistinguishability:
By *Leftover-Hash-Lemma*,
$\mathbf{H}_{\rho(i)}$ is close to uniform

$$\mathbf{c}_i = \mathbf{s}\mathbf{A}_{\rho(i)} + \text{noise}$$

$$\hat{\mathbf{c}}_i = M_i \cdot \begin{bmatrix} \mathbf{s}\mathbf{y}^\perp & 0 & \dots & 0 \\ & \mathbf{v}_2 & & \\ & \mathbf{v}_3 & & \\ & \dots & & \\ & \mathbf{v}_s & & \end{bmatrix} - \mathbf{s}\mathbf{H}_{\rho(i)} + \text{noise}$$

$$\hat{\mathbf{c}}_i = M_i \cdot \begin{bmatrix} \mathbf{0} \\ \hat{\mathbf{v}}_2 \\ \hat{\mathbf{v}}_3 \\ \dots \\ \hat{\mathbf{v}}_s \end{bmatrix} - \mathbf{s}\mathbf{A}_{\rho(i)}\mathbf{R}_{\rho(i)} + \text{noise}$$

$\mathbf{c}_i - \text{noise}$

$$\approx M_i \cdot \begin{bmatrix} \mathbf{0} \\ \hat{\mathbf{v}}_2 \\ \hat{\mathbf{v}}_3 \\ \dots \\ \hat{\mathbf{v}}_s \end{bmatrix} - \mathbf{c}_i\mathbf{R}_{\rho(i)}$$

Getting-Rid of

Remember:

$$sk_U = (\{\tilde{\mathbf{k}}_u\}, \mathbf{t})$$

$$\forall u: \quad \tilde{\mathbf{k}}_u \leftarrow \mathbf{A}_u^{-}(\mathbf{H}_u \cdot \mathbf{t}^\top)$$

$$\mathbf{H}_u = M_{\rho^{-1}(u)} \cdot \begin{bmatrix} \mathbf{y}^\perp & \mathbf{0}^\top & \dots & \mathbf{0}^\top \\ & \mathbf{B}_2 & & \\ & \mathbf{B}_3 & & \\ & \dots & & \\ & \mathbf{B}_s & & \end{bmatrix} + \mathbf{A}_u \mathbf{R}_u$$

$$\mathbf{t} = (1, \mathbf{t})$$

$$\hat{\mathbf{t}} \leftarrow \text{noise}^{m-1}$$

$$\forall u: \quad \tilde{\mathbf{k}}_u \leftarrow \mathbf{A}_u^- \left(M_{\rho^{-1}(u)} \cdot \begin{bmatrix} \mathbf{y}^\perp & \mathbf{0}^\top & \dots & \mathbf{0}^\top \\ & \mathbf{B}_2 & & \\ & \mathbf{B}_3 & & \\ & \dots & & \\ & \mathbf{B}_s & & \end{bmatrix} \mathbf{t}^\top \right) + \mathbf{R}_u \mathbf{t}^\top$$

Getting-Rid of

$$sk_U = (\{\tilde{\mathbf{k}}_u\}, \mathbf{t})$$

$$\forall u: \quad \tilde{\mathbf{k}}_u \leftarrow \mathbf{A}_u^- \left(M_{\rho^{-1}(u)} \cdot \begin{bmatrix} \mathbf{y}^\perp & \mathbf{0}^\top & \dots & \mathbf{0}^\top \\ & \mathbf{B}_2 & & \\ & \mathbf{B}_3 & & \\ & \dots & & \\ & \mathbf{B}_s & & \end{bmatrix} \mathbf{t}^\top \right) + \mathbf{I}$$

$$\mathbf{t} = (1, \hat{\mathbf{t}})$$

$$\mathbf{A}_u \cdot ? = \mathbf{z}_u$$

1) Sample $?$ directly
2) Compute $\mathbf{A}_u \cdot ? = \mathbf{z}_u$
3) Sample B_i's with a trapdoor
4) Choose $\hat{\mathbf{t}}$ that will make it work

This is where linear independence is used

Conclusion

- The first MA-ABE for a non-trivial class from LWE
- A direct LWE-based approach for CP-ABE

Open problems:

- More than DNFs
- Better security (we only get static)
- Better parameters (even for CP-ABE)

New Results

Thank you!

	Supported Policy Class	Assumption
Lewko-Waters '11	NC1	Subgroup decision (composite order)
Okamoto-Takashima '13	NC1	DLIN (prime order)
Rouselakis-Waters '15	NC1	q-type (prime order)
This Work	DNF	LWE