

Encrypted Distributed Systems

Seny Kamara

with Archita Agarwal



14,717,618,286*

4%

Why so Few?



Incompetence?



Lazyness?



Cost?

“...because it would have hurt Yahoo’s ability to index and search message data...”

— J. Bonforte in NY Times

Q: can we search on encrypted data?

Encrypted Search Algorithms



- Major companies
 - MongoDB, Google
 - Meta, Microsoft
 - Amazon, Cisco
 - Hitachi, Fujitsu
 - more...
- Funding agencies
 - NSF
 - IARPA
 - DARPA
- Startups
 - Aroki Systems (acquired)
 - too many to list...

Encrypted Search Algorithms

Property-Preserving
Encryption (PPE)

[[BBO06](#)]

Functional Encryption

[[BSW11](#)]

Structured Encryption
(STE)

[[CGK06](#),[CK10](#)]

Fully-Homomorphic
Encryption (FHE)

[[Gentry09](#)]


Oblivious RAM (ORAM)

[[GO96](#)]

Multi-Party
Computation

[[Yao86](#),[GMW87](#)]

Efficiency

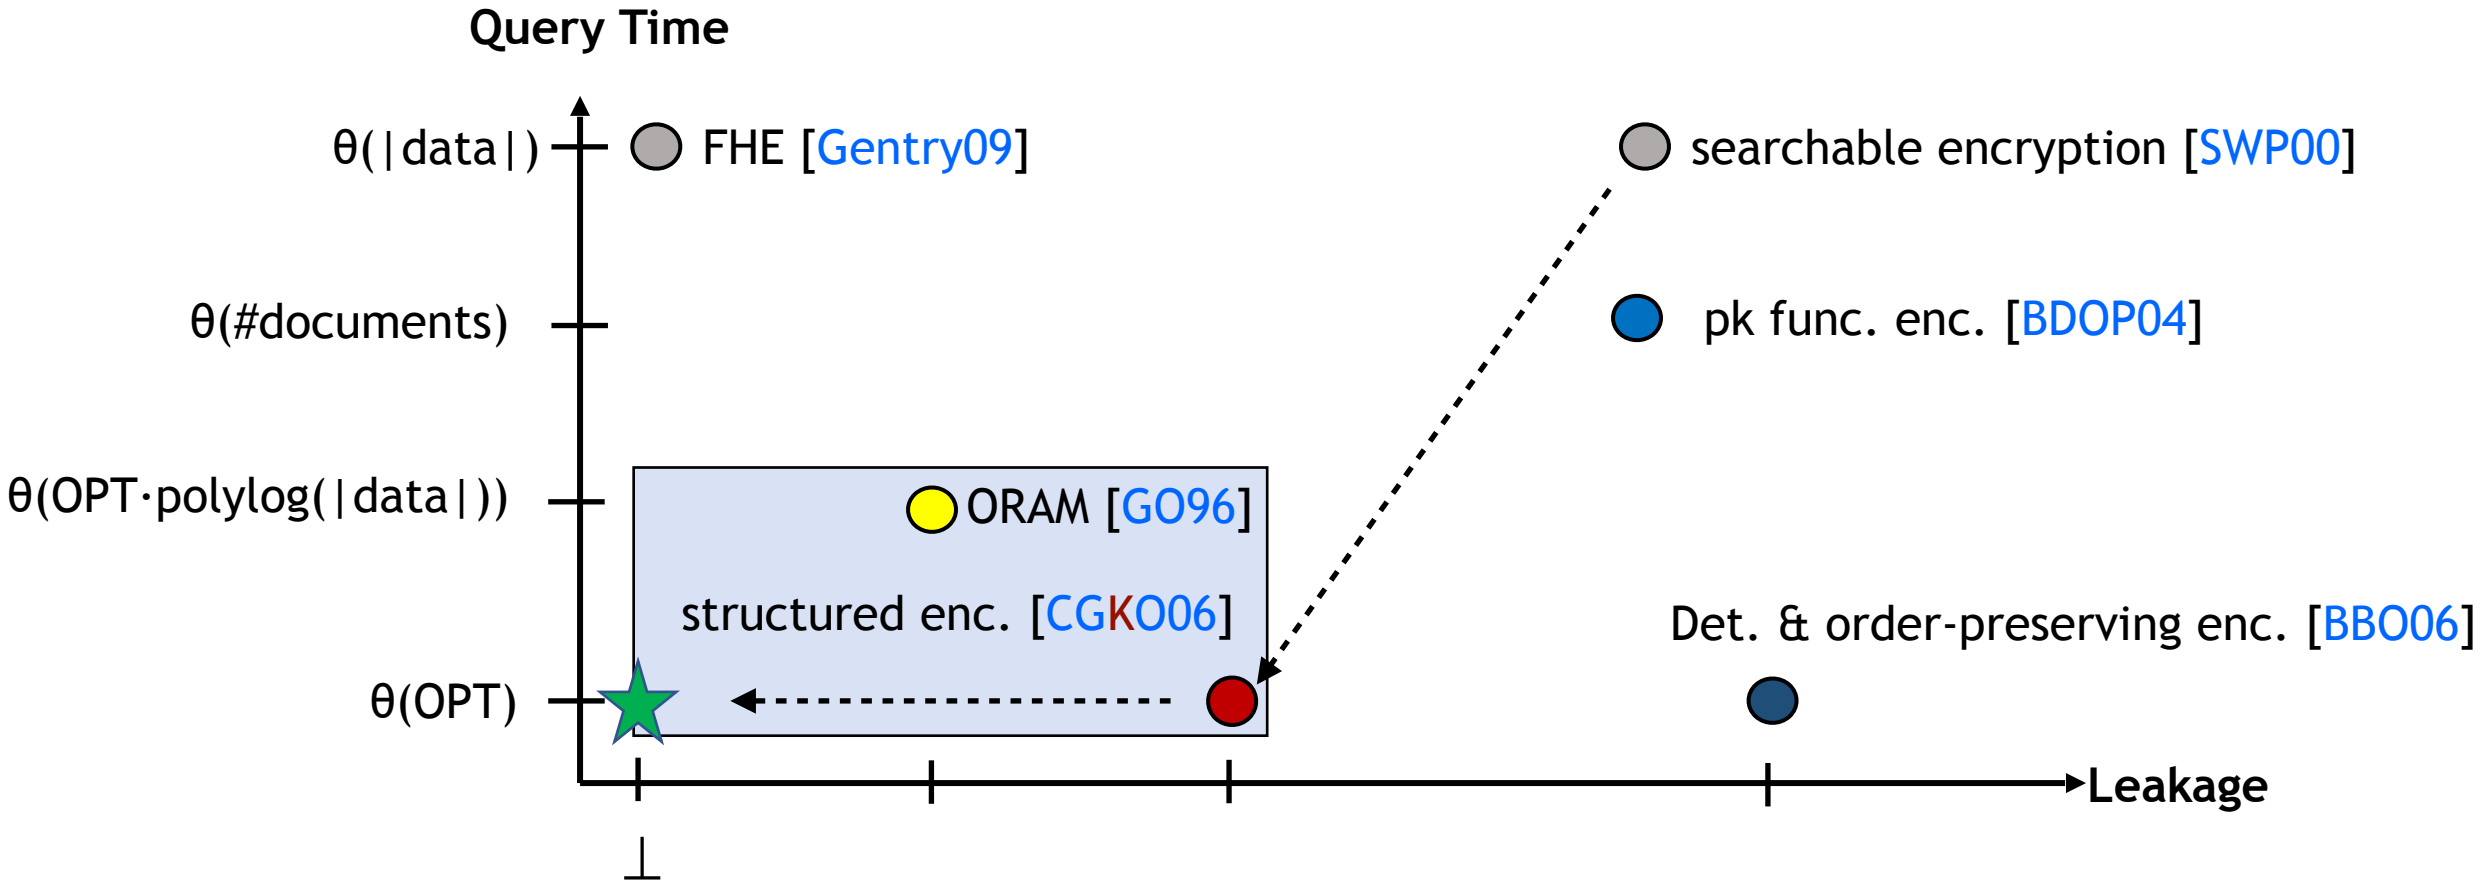


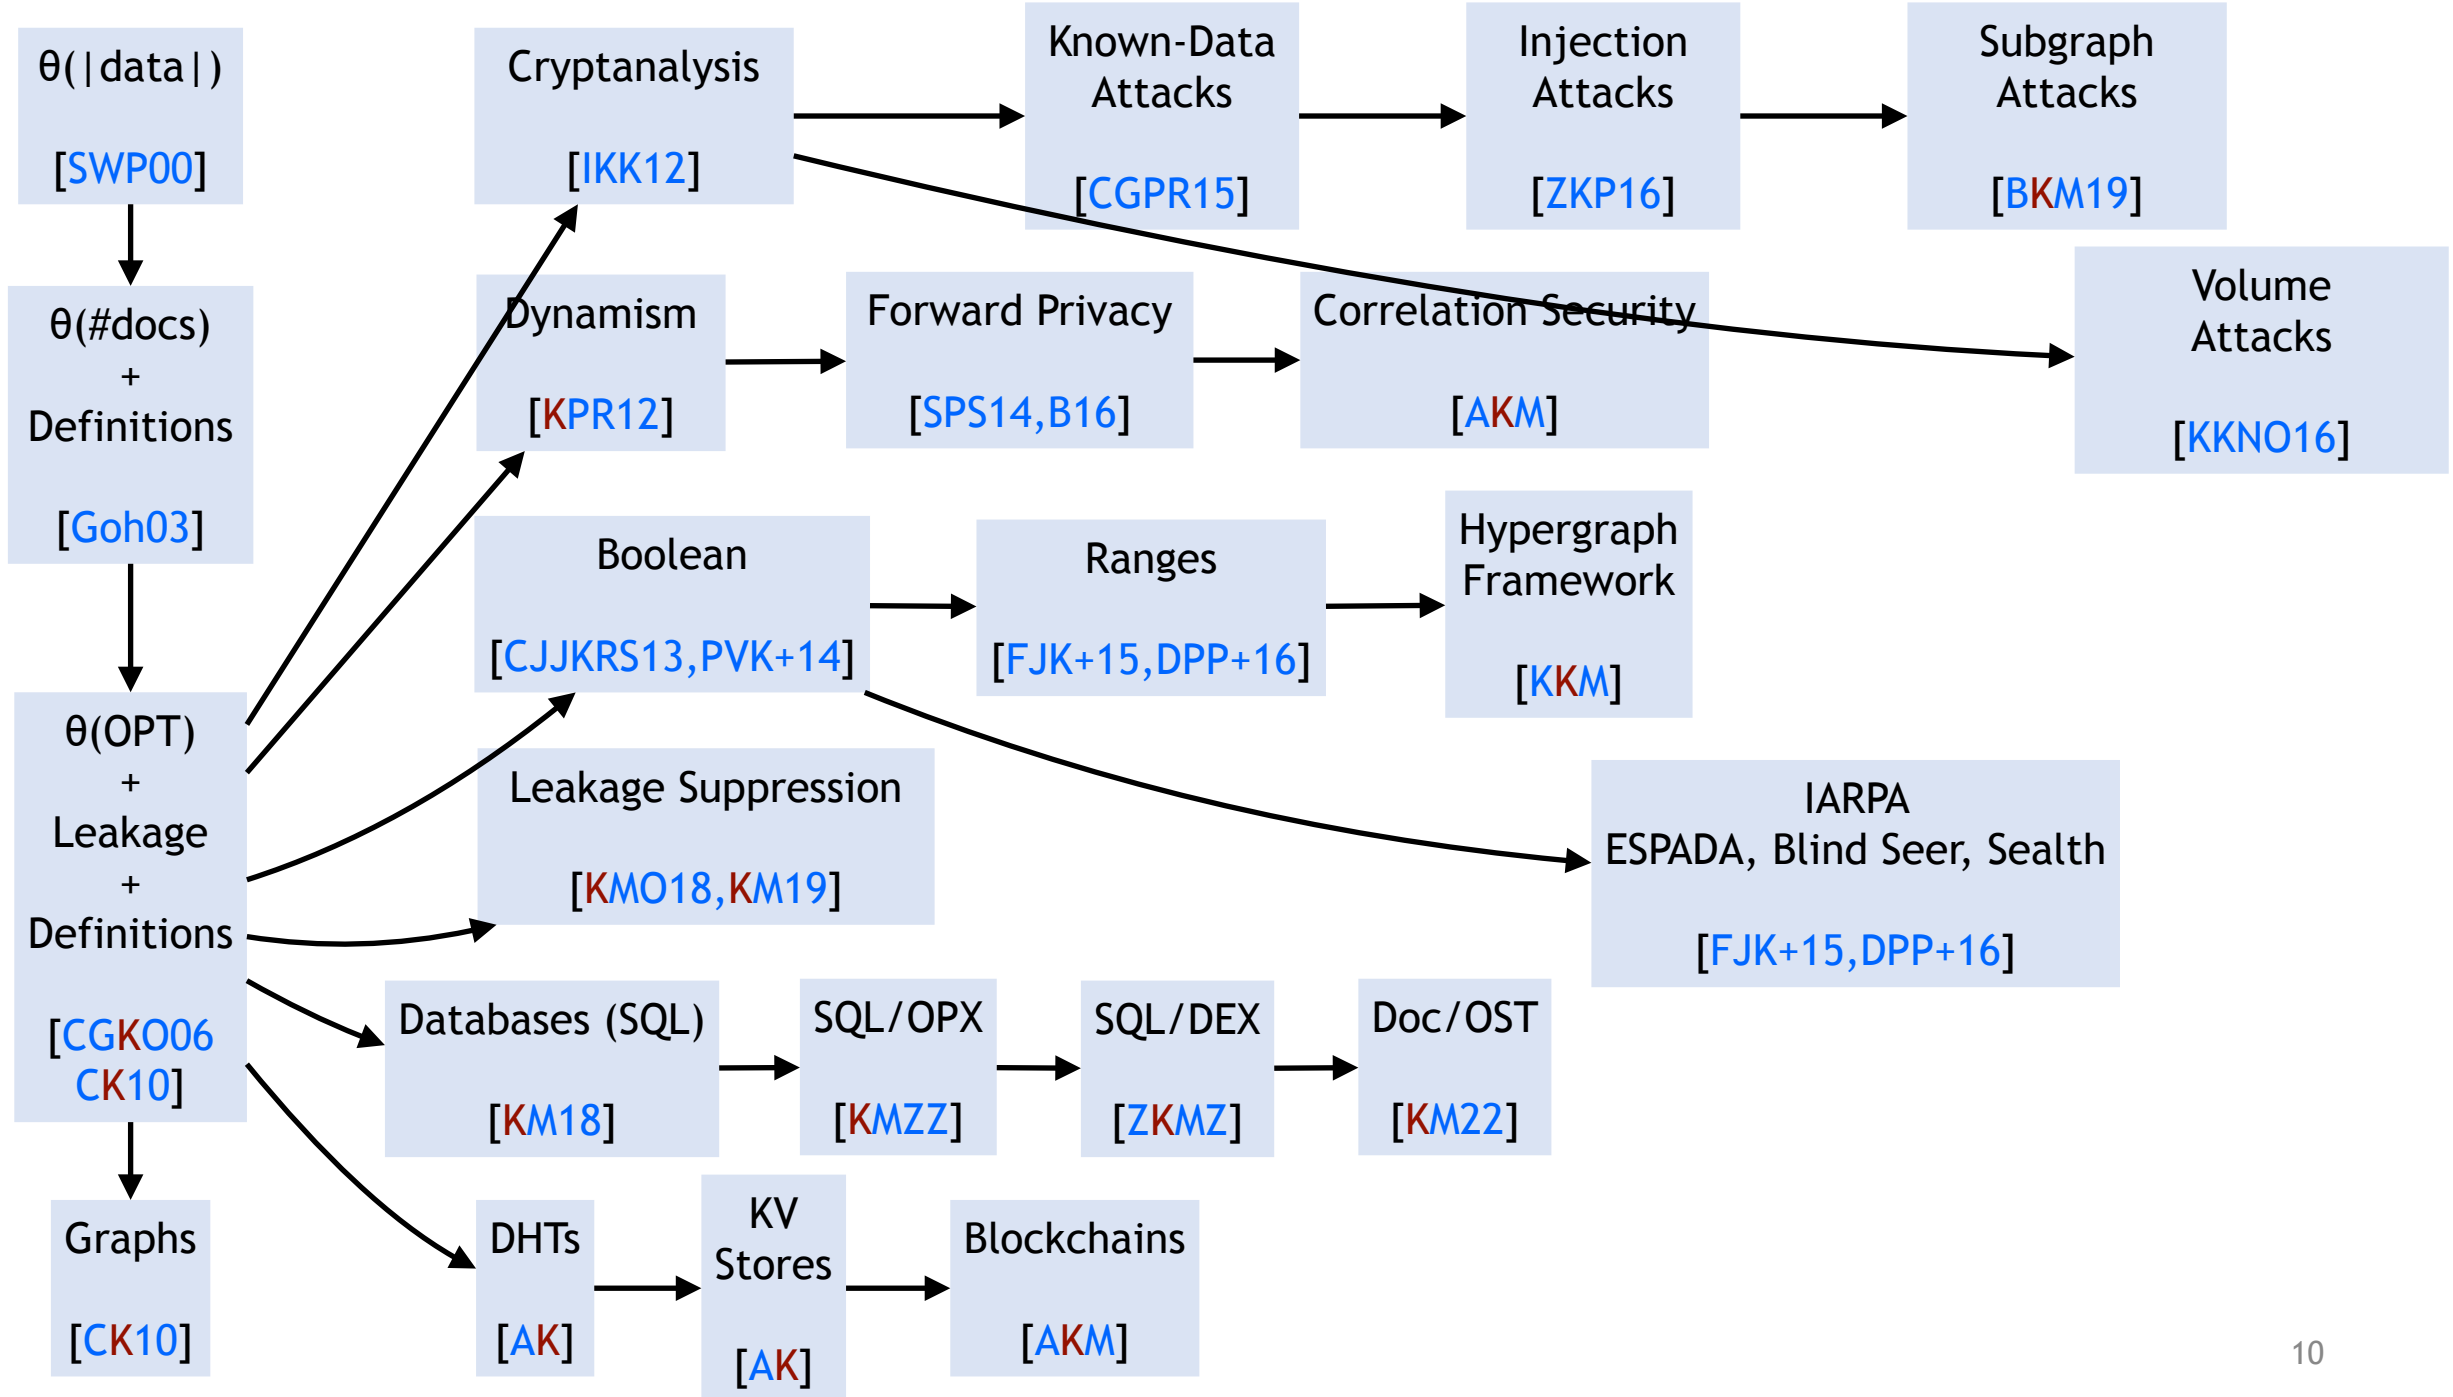
Functionality

Leakage

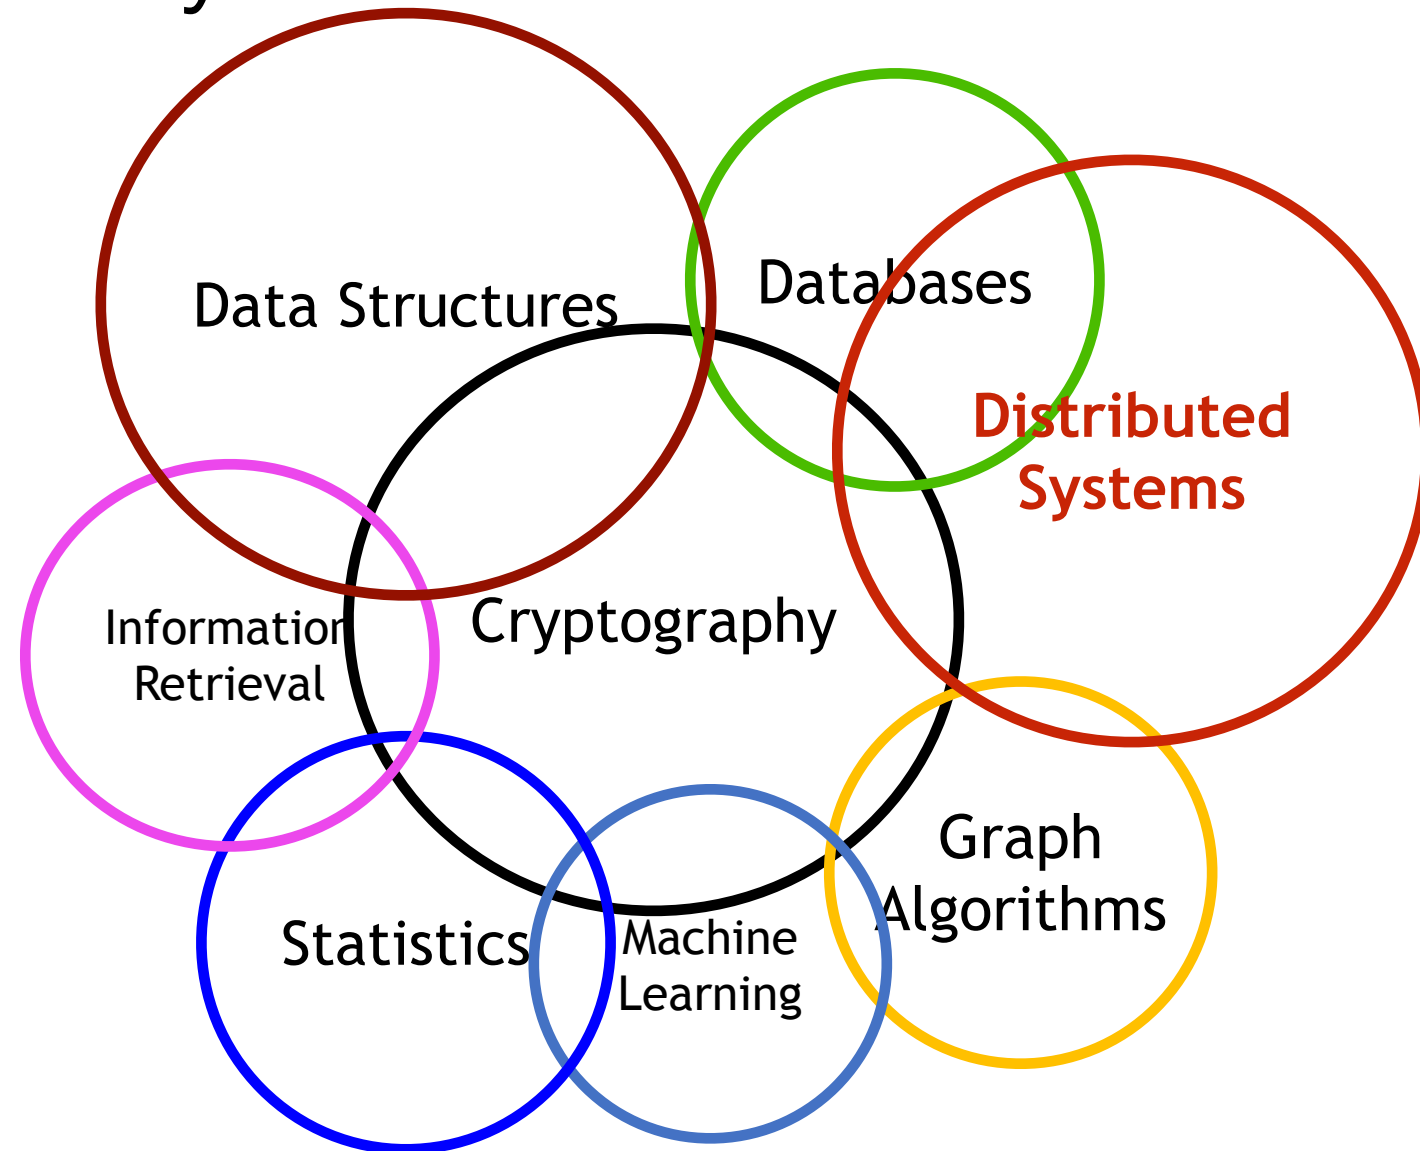
Not Scientific!

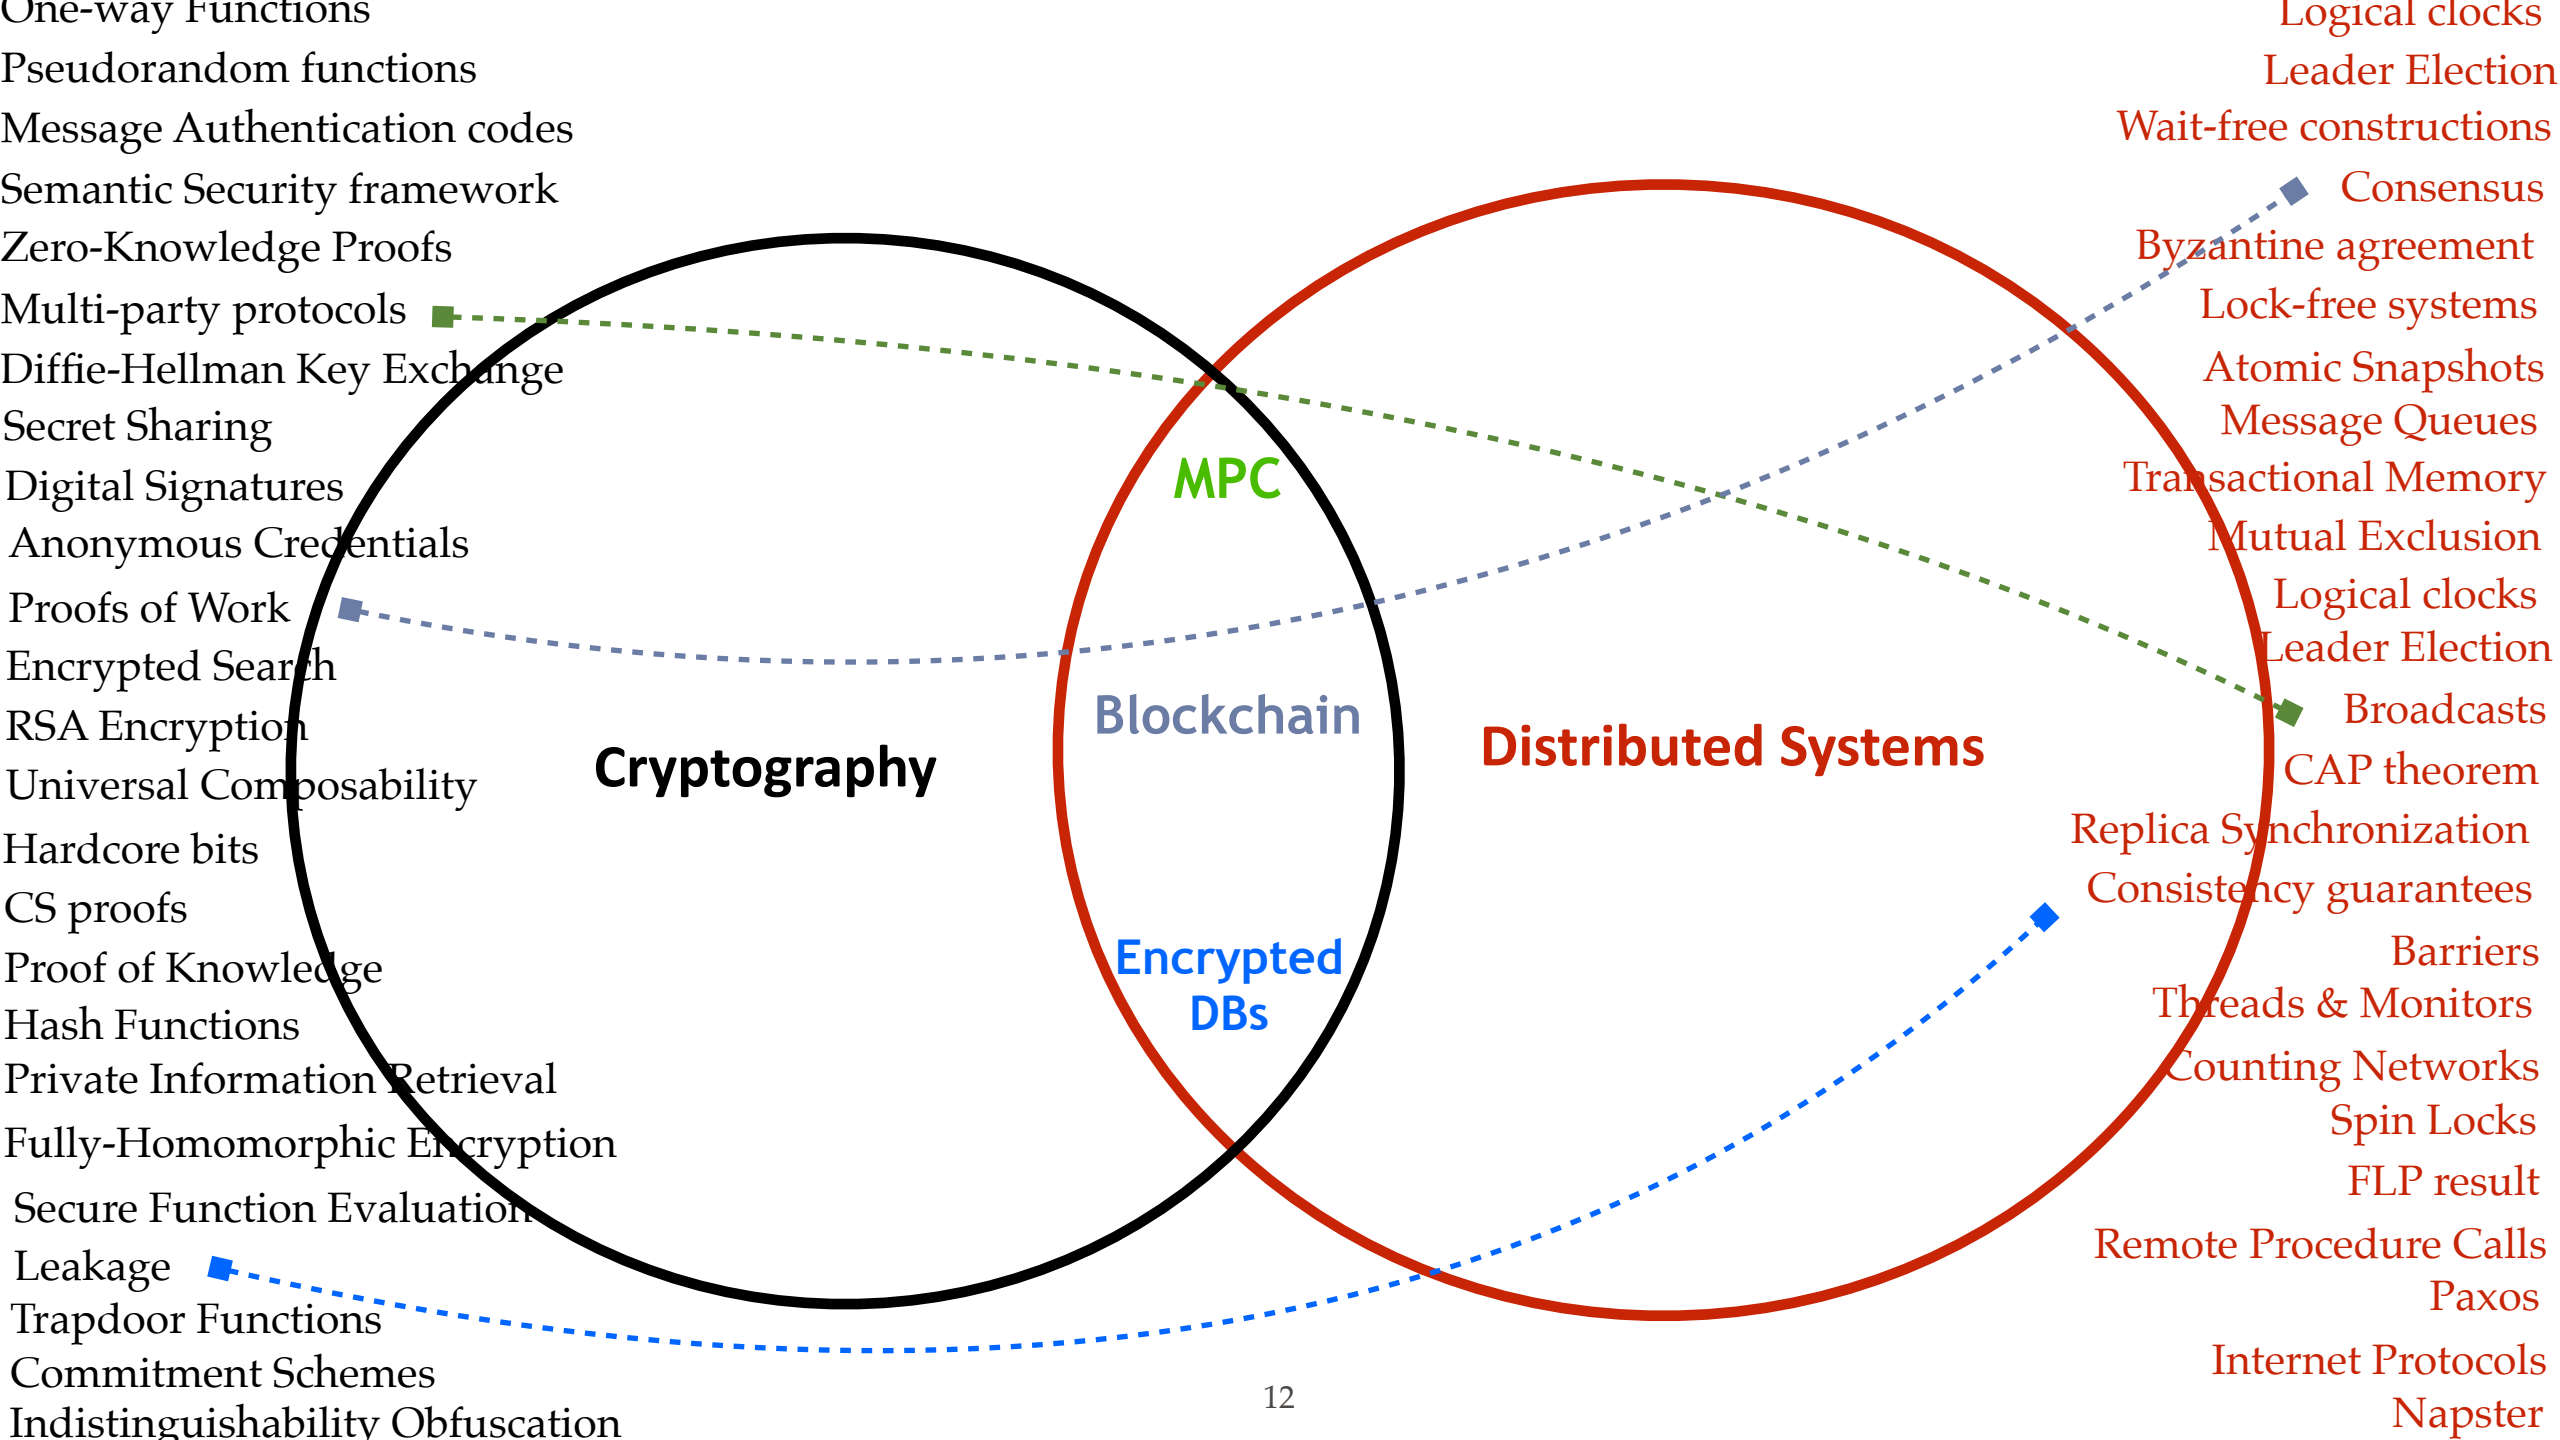
Efficiency vs. Security






Interdisciplinary





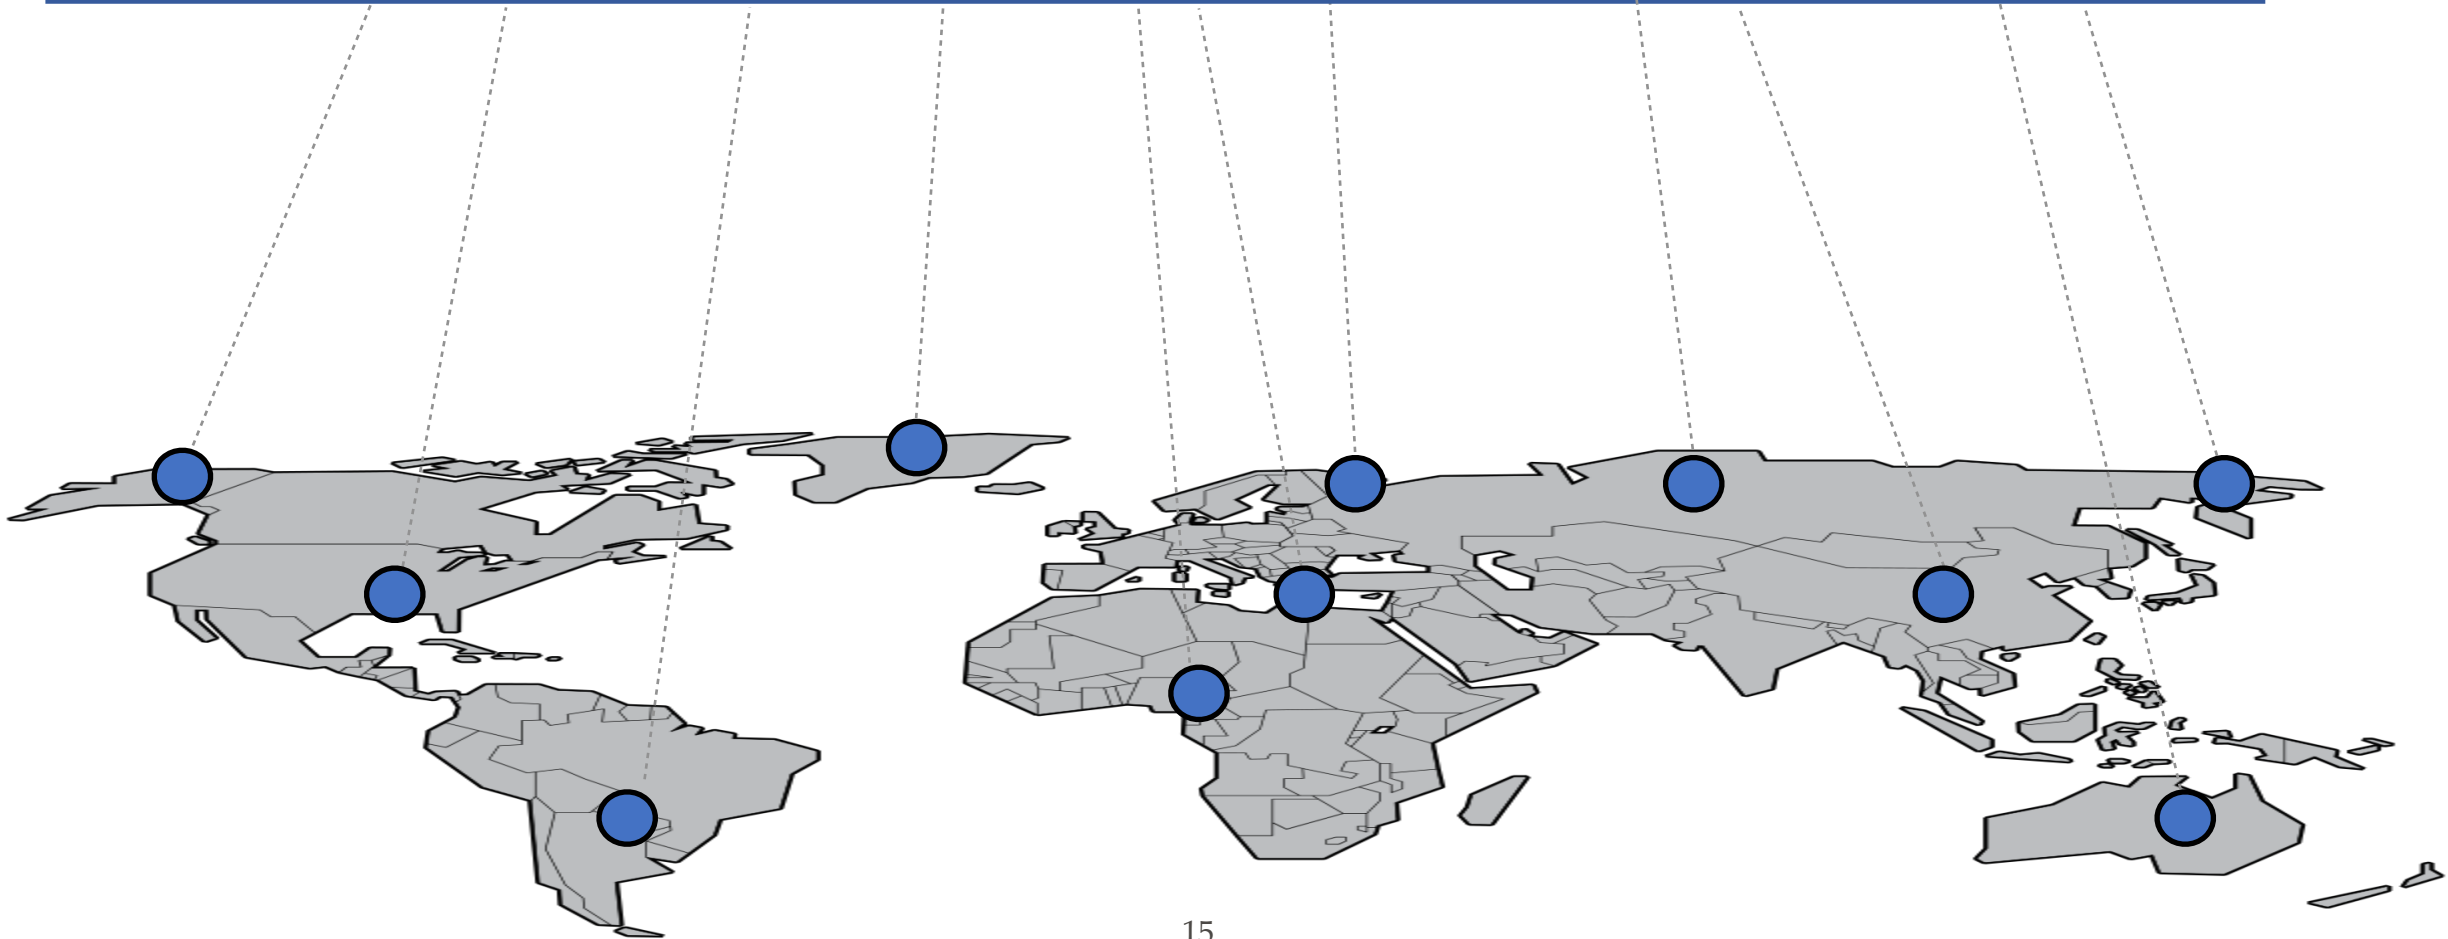
Encrypted Algorithms & Encrypted Systems

- **Q:** can we design algorithms that operate on encrypted data?
- **Q:** can we build systems that run on encrypted data?
 - databases, key-value stores, blockchains, ...



Q: what's the **simplest**
distributed data structure?

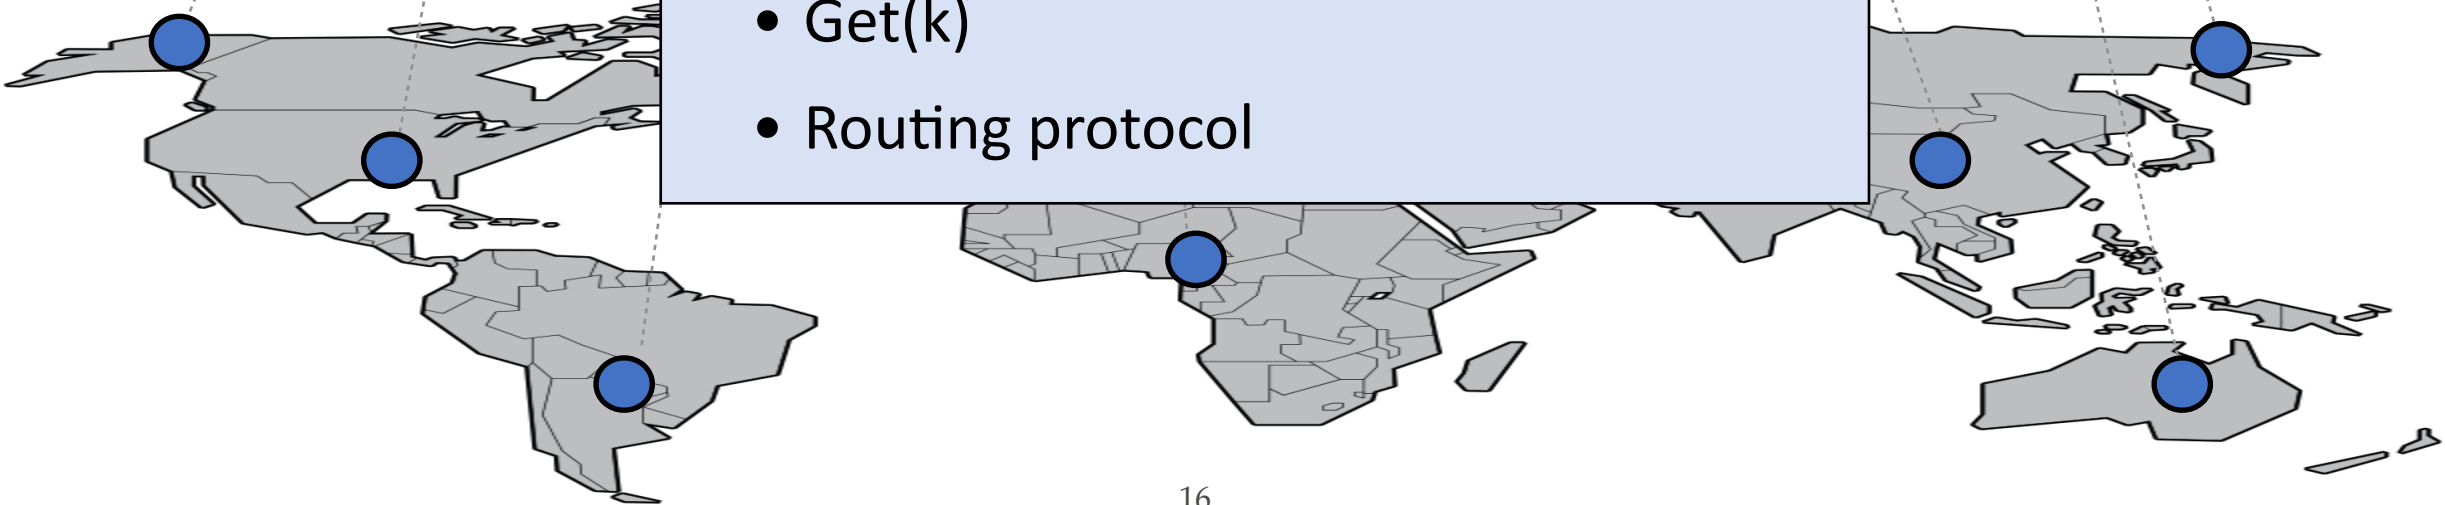
DHT



DHT

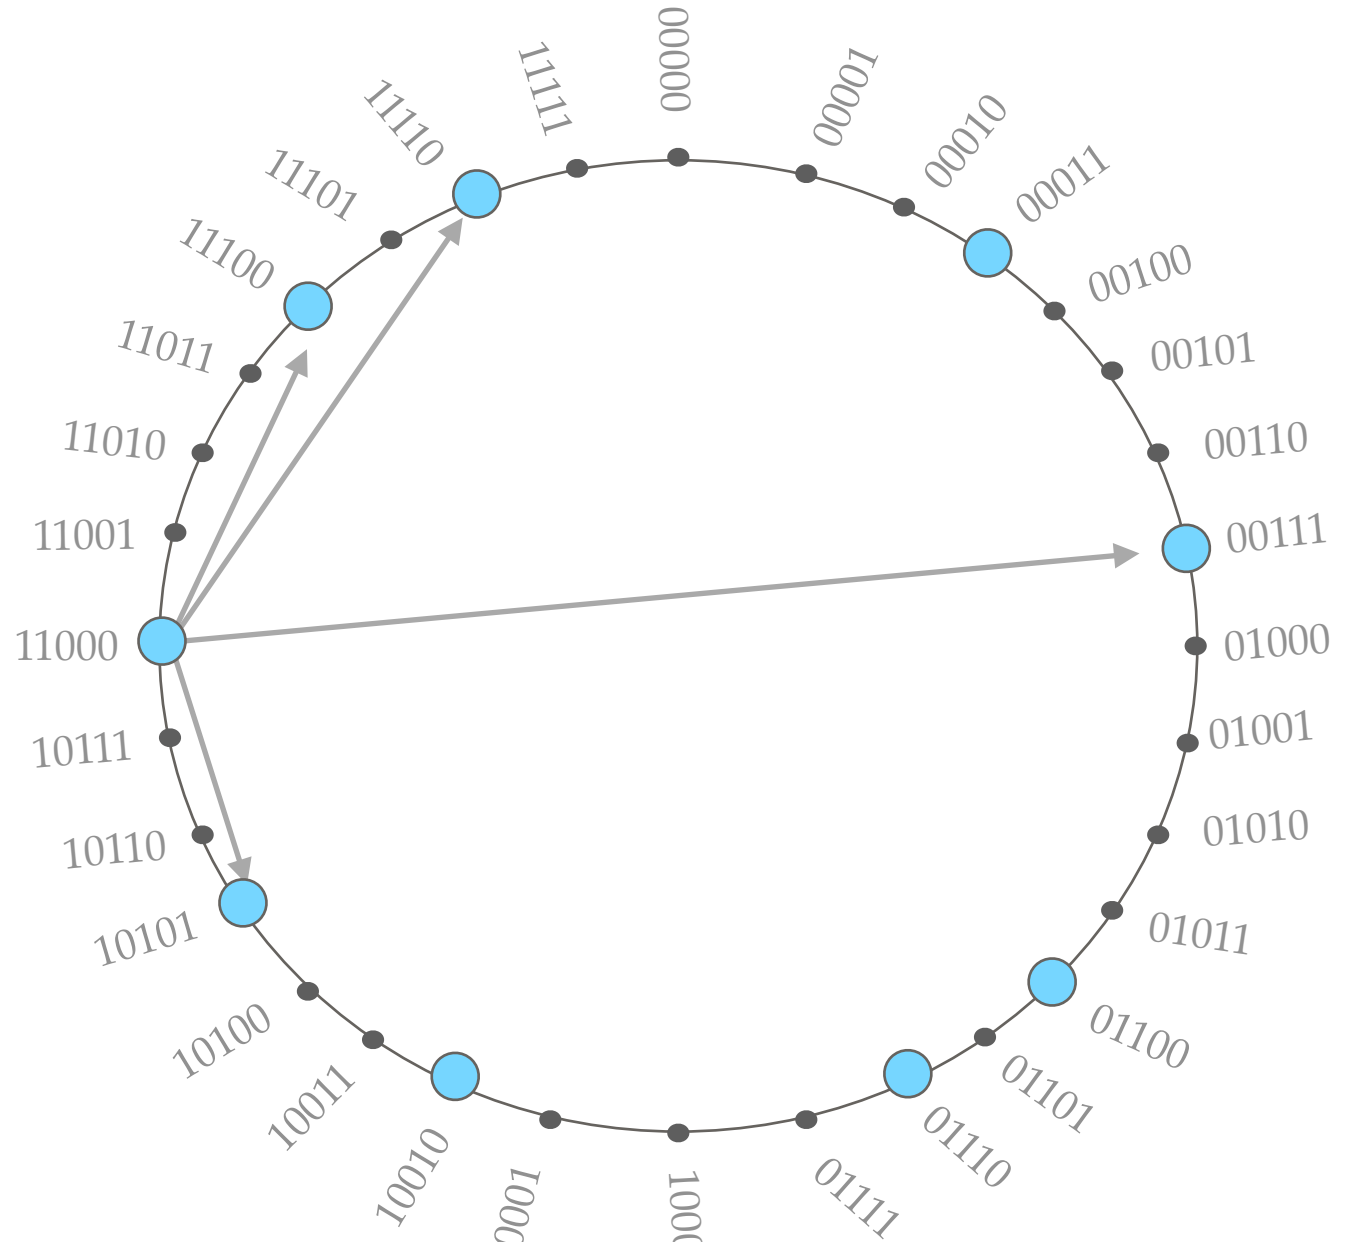
- Protocols:
 - Setup
 - Put(k, v)
 - Get(k)
 - Routing protocol

key $k \Rightarrow$ label ℓ




Chord DHT


- Logical Address Space : **A**
- $(H_1, H_2) \leftarrow \text{Setup}()$
 - H_1 : hashes node ids to addresses
 - H_2 : hashes labels to addresses
- Routing
 - Logarithmic sized routing tables
 - Logarithmic sized paths



Chord DHT : Put()



$\text{Put}(\ell_4, v_4)$



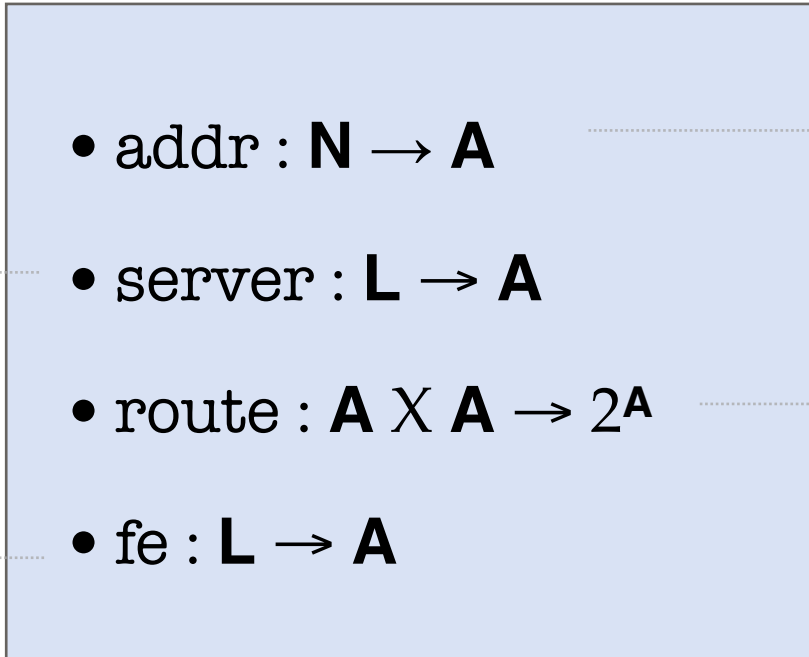
Abstraction of DHTs

succ ◦ **H₂**

Where should labels be stored?

Where should clients send their requests

H₃



All nodes are assigned logical addresses

H₁

How should messages be routed?

Fixed

OUTLINE

(I) Encrypted DHTs

- ❖ What are DHTs
 - ▶ Abstraction of core components
- ❖ **Formalize EDHTs**
 - ▶ Syntax & Security defn
- ❖ Construction
- ❖ Analysis of EDHTs
 - ▶ Main security theorem

(III) Takeaways & Conclusion

Formalizing EDHTs

- Define the **syntax** of EDHTs
- Define the **security** of EDHTs

Formalizing EDHTs

- Define the **syntax** of EDHTs
- Define the **security** of EDHTs

Formalizing EDHTs : Syntax

EDHT = (Gen, Setup, Put, Get)

- Executed by user
- Generates cryptographic keys

- Executed by trusted party
- sets up system

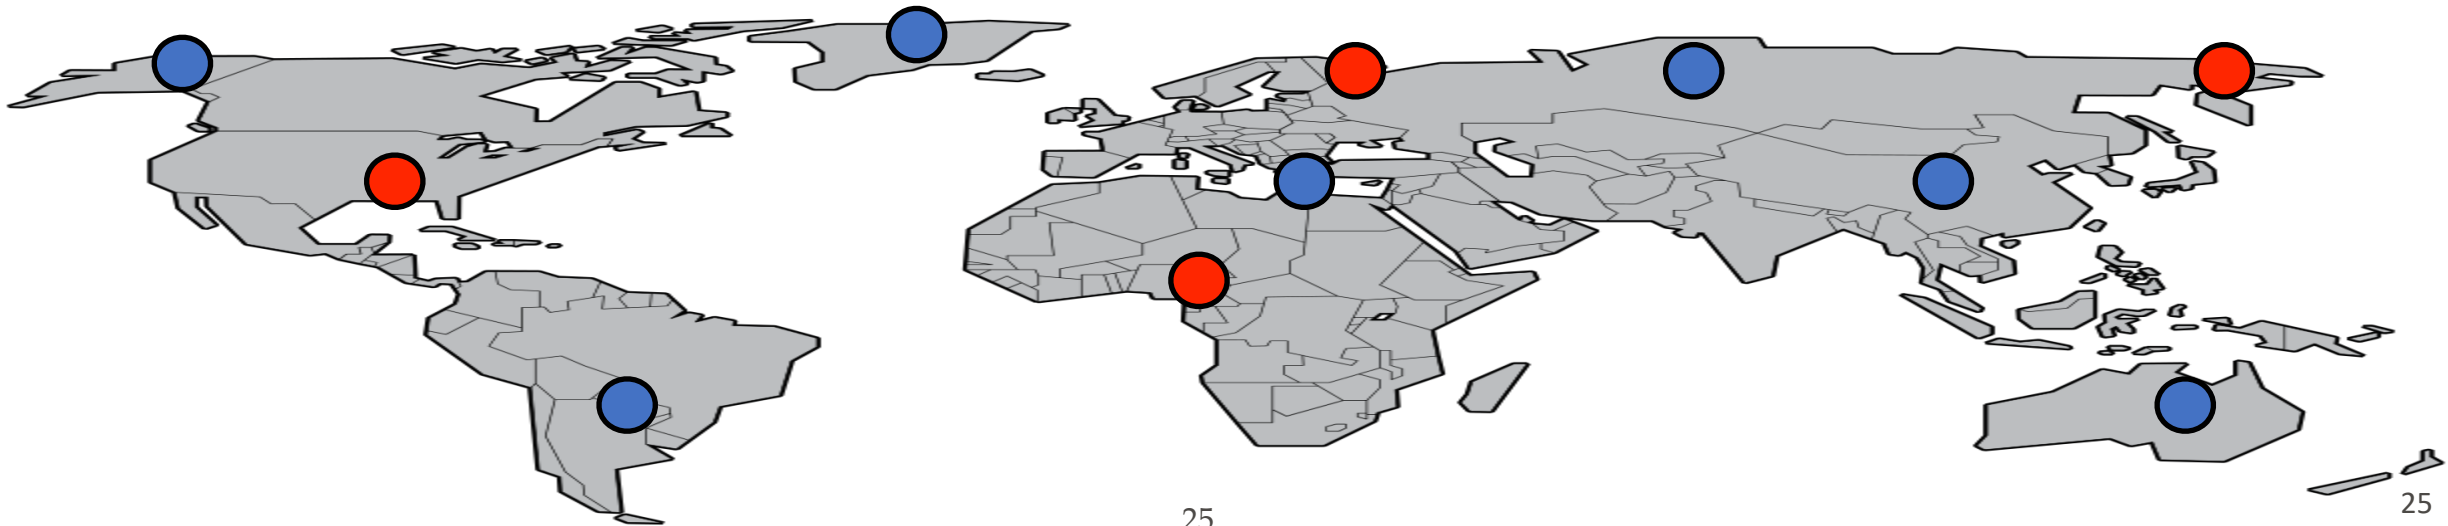
- Executed by user
- Put(K, ℓ , v): stores (ℓ , v)
- Get(K, ℓ): retrieves (ℓ , v)

Formalizing EDHTs

- Define the **syntax** of EDHT
- Define the **security** of EDHTs

Adversarial Model

- Static
- Semi-honest



EDHTs Security

Real

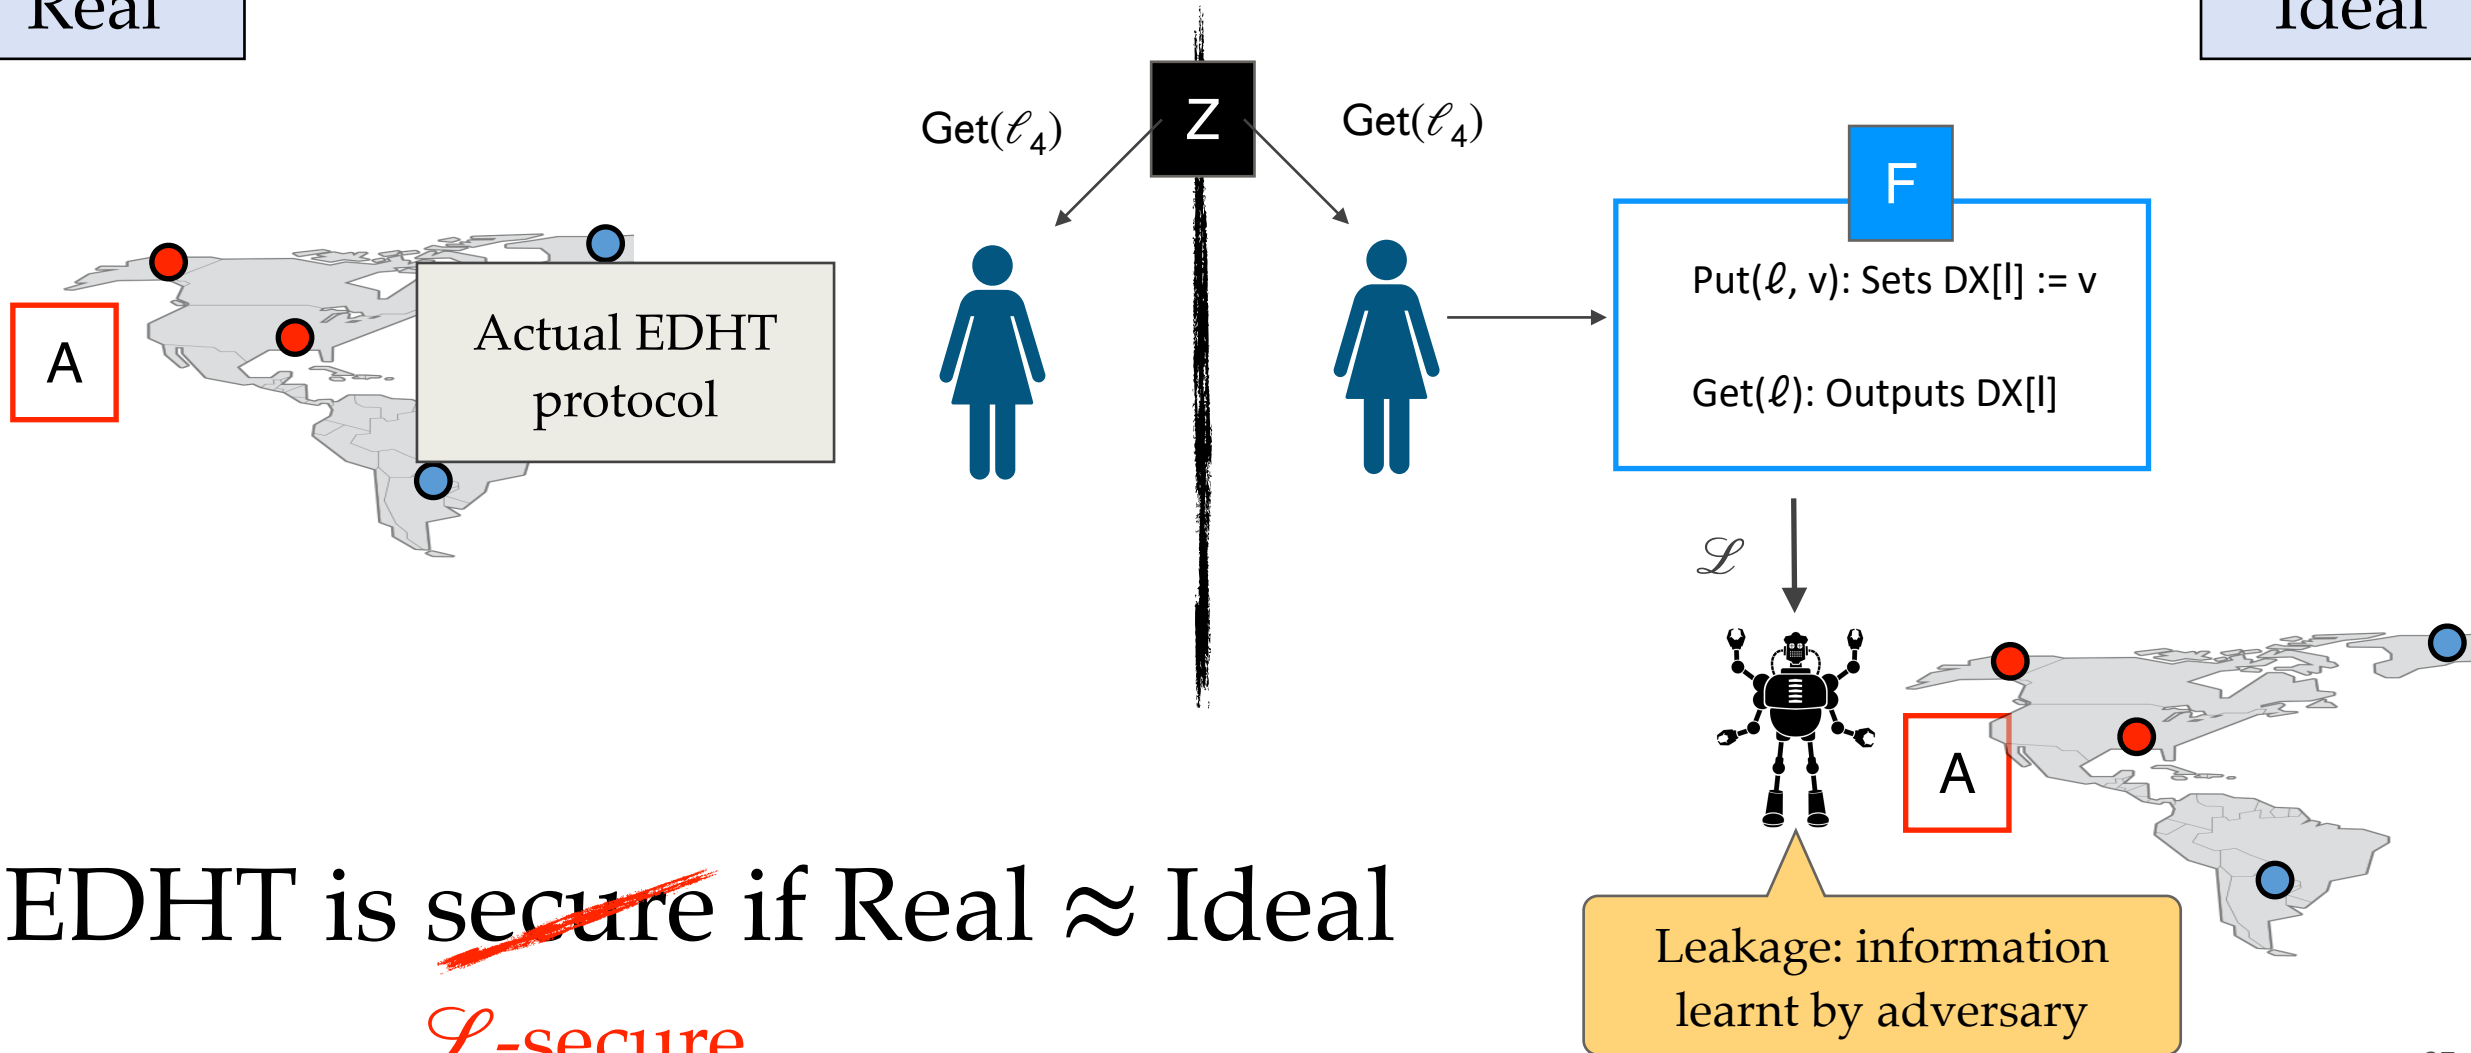
Ideal



EDHT Security

Real

Ideal



EDHT is ~~secure~~ if Real \approx Ideal
 \mathcal{L} -secure

OUTLINE

(I) Encrypted DHTs

- ❖ What are DHTs
 - ▶ Abstraction of core components
- ❖ Formalize EDHTs
 - ▶ Syntax & Security defn
- ❖ **Construction**
- ❖ Analysis of EDHTs
 - ▶ Main security theorem

(III) Takeaways & Conclusion

EDHT Construction

Setup()

- DHT.Setup()

Gen(1^k)

- Sample $K_1 \leftarrow \{0, 1\}^k$
- $K_2 \leftarrow \text{SKE.Gen}(1^k)$
- Output (K_1, K_2)

Put(K, ℓ, v)

- $K = (K_1, K_2)$
- $t = F_{K_1}(\ell)$
- $e = \text{SKE.Enc}_{K_2}(v)$
- DHT.Put(t, e)

Get(K, ℓ)

- $K = (K_1, K_2)$
- $t = F_{K_1}(\ell)$
- $e \leftarrow \text{DHT.Get}(t)$
- $v \leftarrow \text{SKE.Dec}_{K_2}(e)$

OUTLINE

(I) Encrypted DHTs

- ❖ What are DHTs
 - ▶ Abstraction of core components
- ❖ Formalize EDHTs
 - ▶ Syntax & Security defn
- ❖ Construction
- ❖ **Analysis of EDHTs**
 - ▶ **Main security theorem**

(III) Takeaways & Conclusion

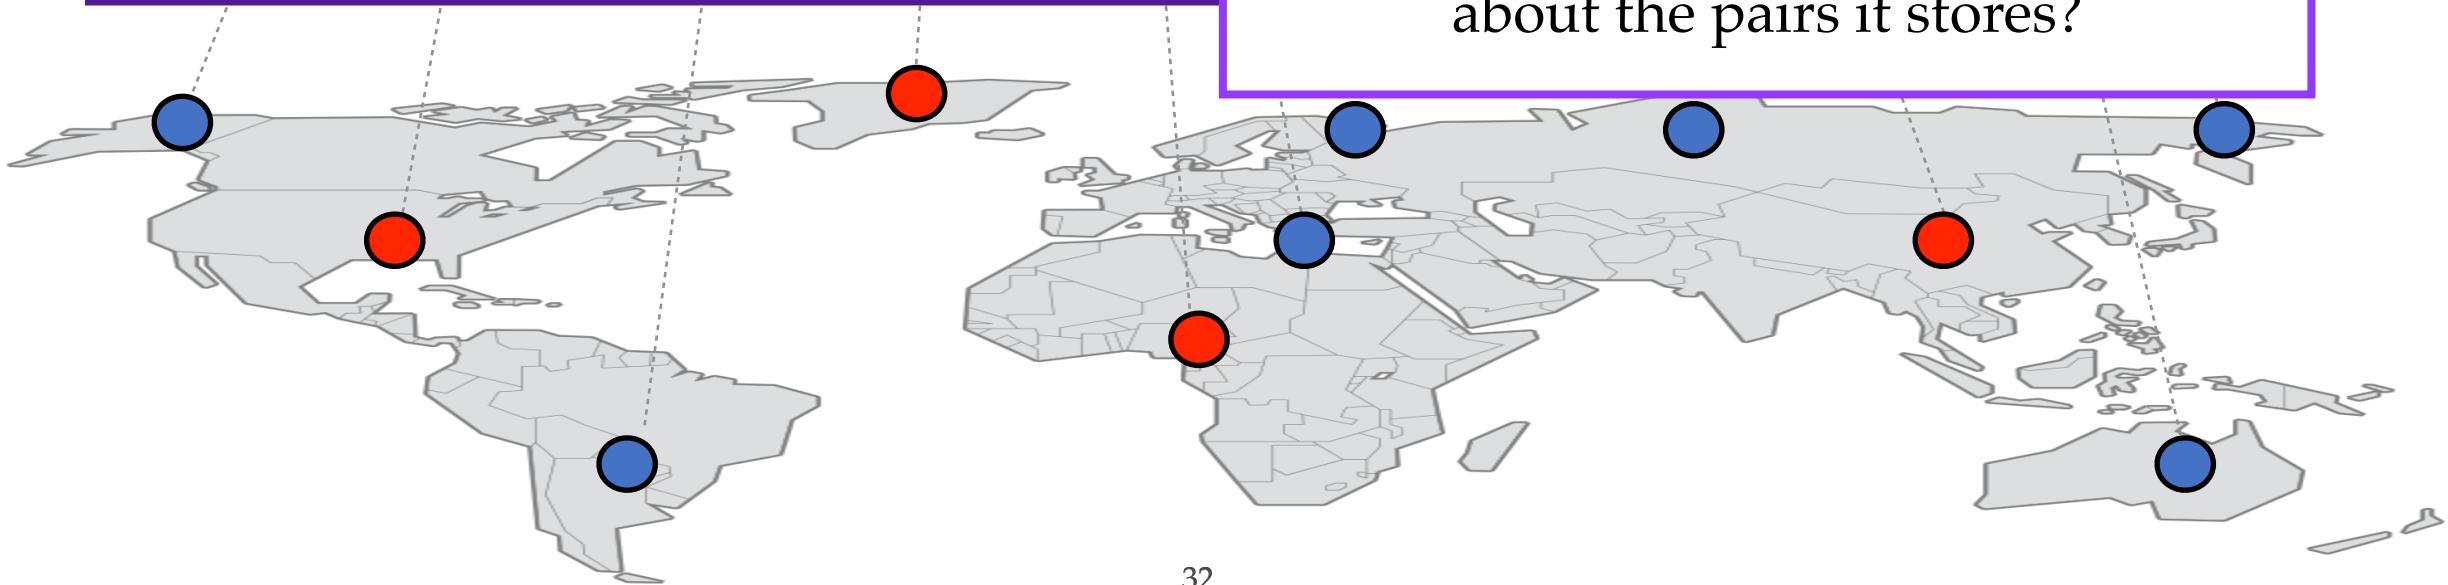
Q: What kind of **security** are we getting?

What does the Adversary learn?

EDHT

Q1: What information does the Adversary learn about pairs stored on corrupted nodes?

Q2: Does it **only** learn information about the pairs it stores?



What does the Adversary learn?

Example:


Infer a good approximation of total number of pairs!

- ❖ Total pairs adv. holds : m
- ❖ Total expected pairs : $\sim mn/t$
 - ❖ if DHTs are load balanced

Q2: Does it **only** learn information about the pairs it stores?

NO

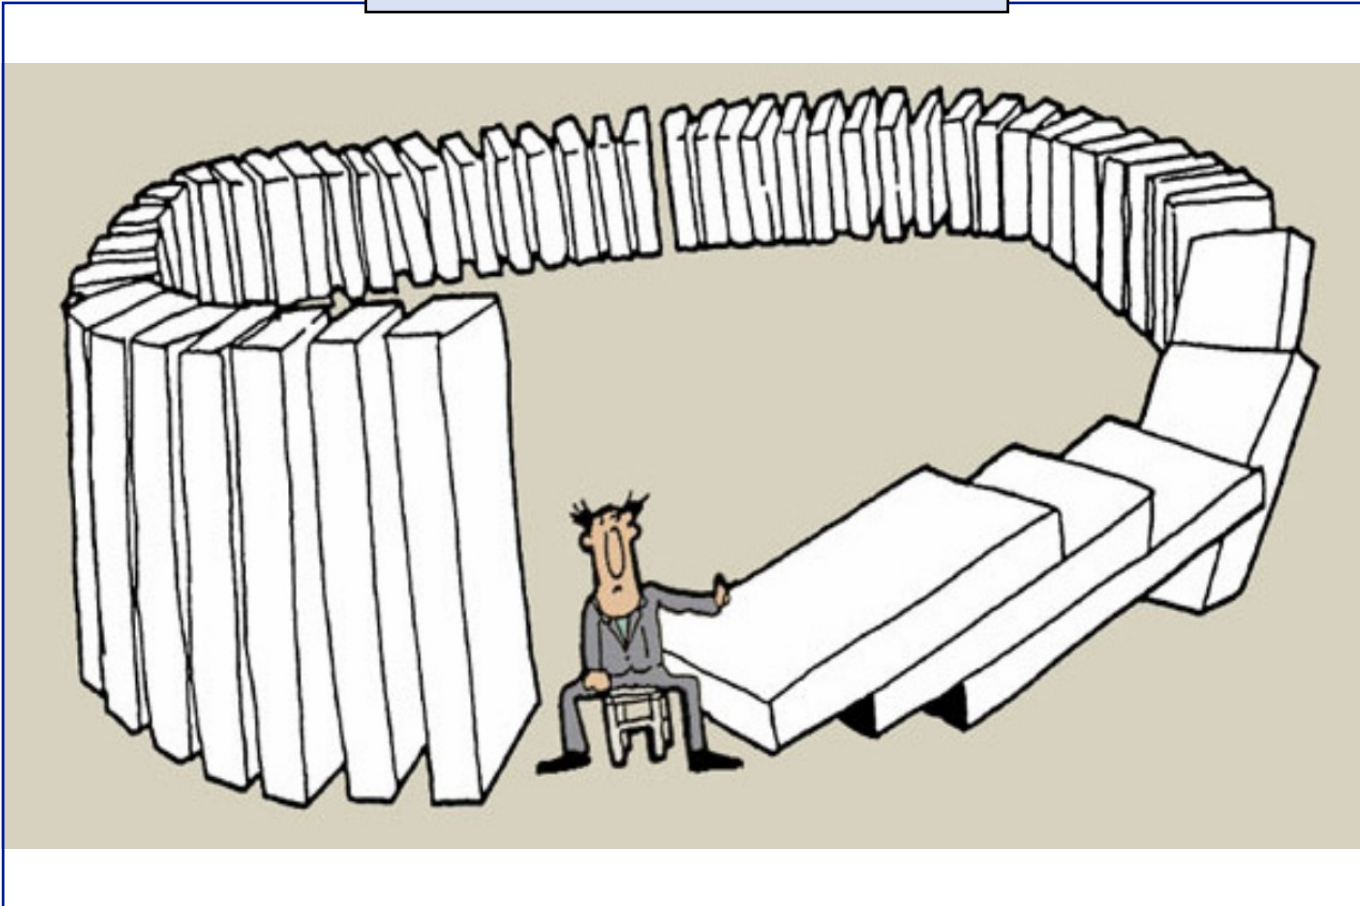
System architecture



Security

Properties of DHTs

P1: Balance



P2: Non-committing allocations



"And if elected, I promise
to keep making promises."

Properties of DHTs

P1: Balance

whp, the probability of
any θ -bounded adversary
seeing a label
should not be more than ϵ

- $\text{addr} : \mathbf{N} \rightarrow \mathbf{A}$
- $\text{server} : \mathbf{L} \rightarrow \mathbf{A}$
- $\text{route} : \mathbf{A} \times \mathbf{A} \rightarrow 2^{\mathbf{A}}$
- $\text{fe} : \mathbf{L} \rightarrow \mathbf{A}$

P2: Non-committing allocations



"And if elected, I promise
to keep making promises."

Properties of DHTs

P1: Balance

whp, the probability of
any θ -bounded adversary
seeing a label
should not be more than ϵ

- $\text{addr} : \mathbf{N} \rightarrow \mathbf{A}$
- $\text{server} : \mathbf{L} \rightarrow \mathbf{A}$
- $\text{route} : \mathbf{A} \times \mathbf{A} \rightarrow 2^{\mathbf{A}}$
- $\text{fe} : \mathbf{L} \rightarrow \mathbf{A}$

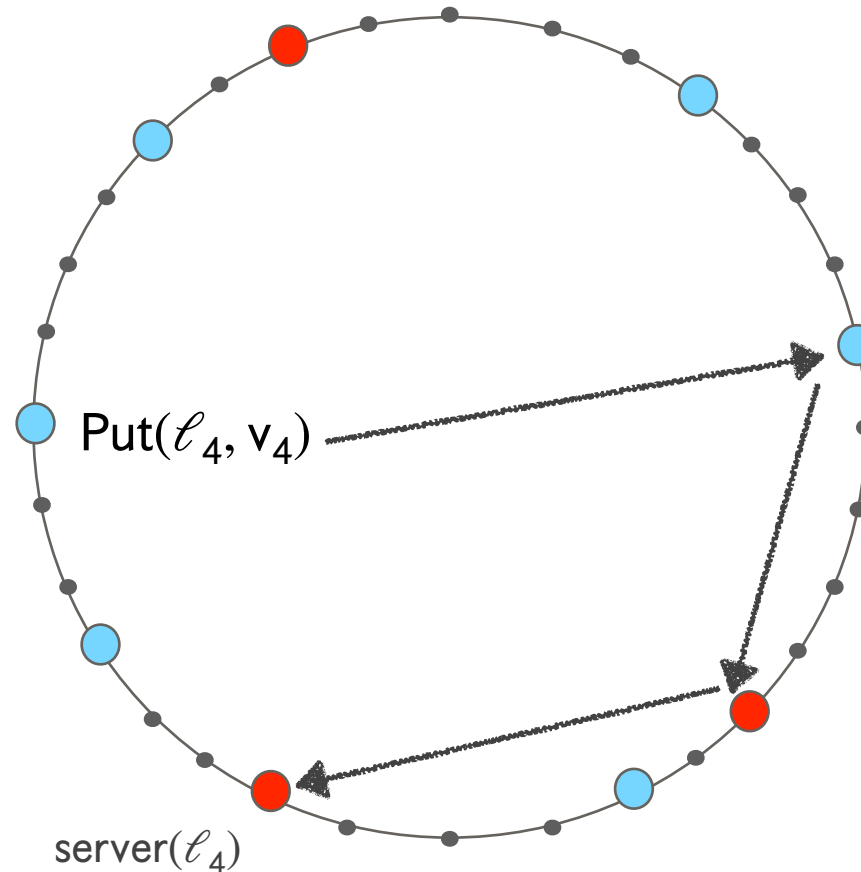
P2: Non-committing allocations



"And if elected, I promise
to keep making promises."

When does an adversary see a label?

- When it **stores** the label or **routes** the label



Properties of DHTs

P1: Balance

whp, the probability of
any θ -bounded adversary
seeing a label
should not be more than ϵ

- $\text{addr} : \mathbf{N} \rightarrow \mathbf{A}$
- $\text{server} : \mathbf{L} \rightarrow \mathbf{A}$
- $\text{route} : \mathbf{A} \times \mathbf{A} \rightarrow 2^{\mathbf{A}}$
- $\text{fe} : \mathbf{L} \rightarrow \mathbf{A}$

P2: Non-committing allocations

much more technical!

Storing or routing a label

Leakage



L_ε

leaks the repetition pattern (when a query for the same label is repeated) for an ε -fraction of queries

affected by balance ε of DHT

Main Security Theorem

Th :

If DHT is $(\epsilon, \theta, \delta)$ -balanced and has non-committing allocations, then EDHT is L_ϵ -secure with prob at least $1 - \delta - \text{negl}(k)$

Balance of Chord

Th :

Chord is $(\varepsilon, \theta, \delta)$ -balanced for

$$\varepsilon = \frac{\theta}{n} \left(\log n + 6 \log \left(\frac{n}{\theta} \right) \right), \quad \delta = \frac{1}{n^2} \quad \text{and} \quad \theta \leq \frac{n}{e \log n}$$

Balance of Chord

Th :

Chord is $(\varepsilon, \theta, \delta)$ -balanced for

$$\varepsilon = \frac{\theta}{n} \left(\log n + 6 \log \left(\frac{n}{\theta} \right) \right), \quad \delta = \frac{1}{n^2} \quad \text{and} \quad \theta \leq \frac{n}{e \log n}$$



$$\varepsilon = O\left(\frac{\theta}{n} \log n\right)$$

VS

$$\varepsilon = O\left(\frac{\theta}{n}\right)$$

optimal

OUTLINE

(I) Encrypted DHTs

Transient DHTs

- ❖ What are DHTs
 - ▶ Abstraction of core components
- ❖ Formalize EDHTs
 - ▶ Syntax & Security defn
- ❖ Construction
- ❖ Analysis of EDHTs
 - ▶ Main security theorem

(III) Takeaways & Conclusion

Outline

✓ (0) Introduction

✓ (I) Encrypted DHTs

(II) Encrypted Key-Value Stores

(III) Future Directions

What are Key-Value Stores?

Same as DHTs

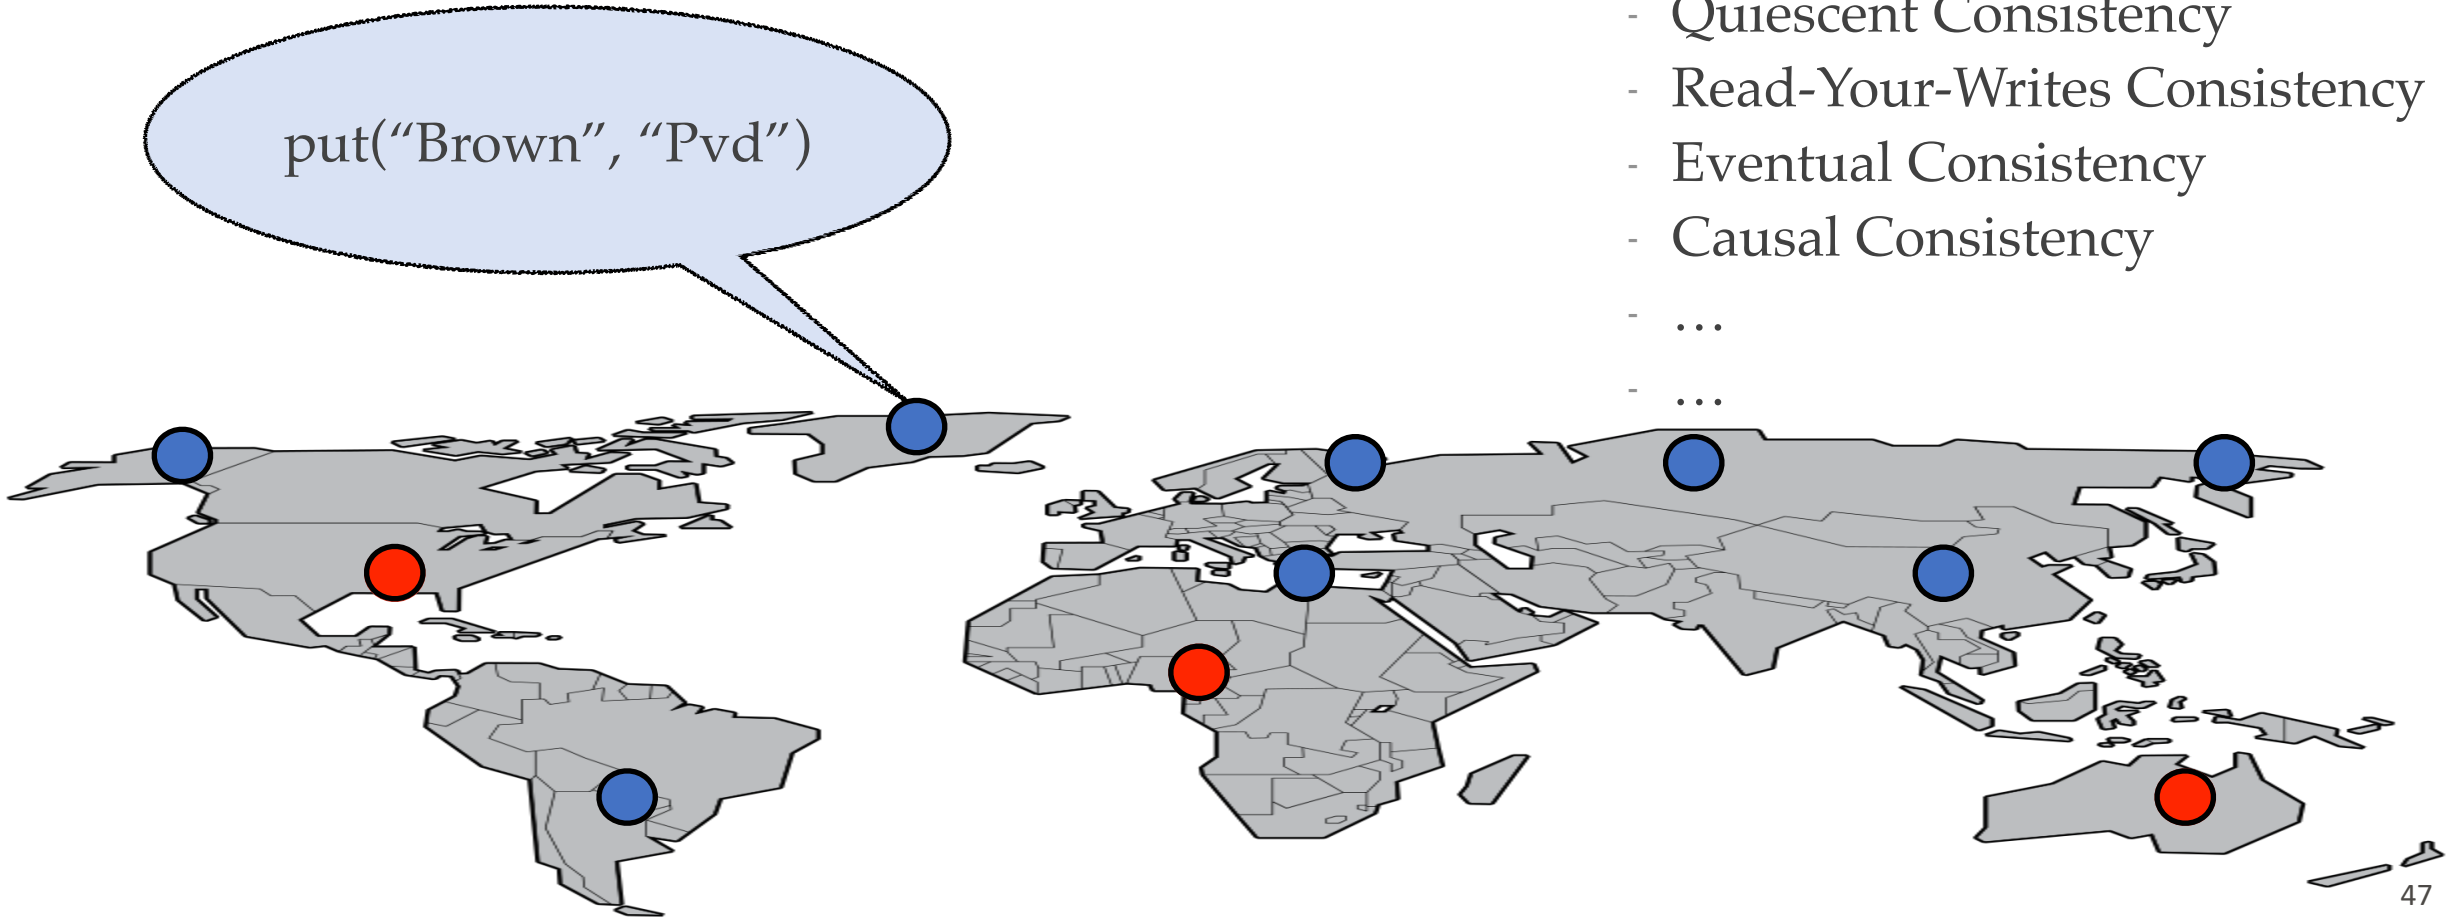
+

Replication

KVS

CONSISTENCY ??

put("Brown", "Pvd")



Abstraction of KVS

- $\text{addr} : \mathbf{N} \rightarrow \mathbf{A}$
- ~~$\text{server} : \mathbf{L} \rightarrow \mathbf{A}$~~ $\text{replicas} : \mathbf{L} \rightarrow 2^{\mathbf{A}}$
- $\text{route} : \mathbf{A} \times \mathbf{A} \rightarrow 2^{\mathbf{A}}$
- $\text{fe} : \mathbf{L} \rightarrow \mathbf{A}$

Construction of EKVS

SAME AS
BEFORE

Put(K, ℓ, v)

- $t = F_{K_1}(\ell)$
- $e = \text{SKE.Enc}_{K_2}(v)$
- ~~DIIT.Put(t, e)~~
KVS.Put(t, e)

Security of EKVS

Single user setting

Clients do not share data

Multi user setting

Clients can share data

concurrent operations on
same piece of data possible

Properties of KVSs

P1: Balance

whp, the probability of any θ -bounded adversary seeing a label should not be more than ϵ

P2: Non-committing

much more technical!

P3: Consistency

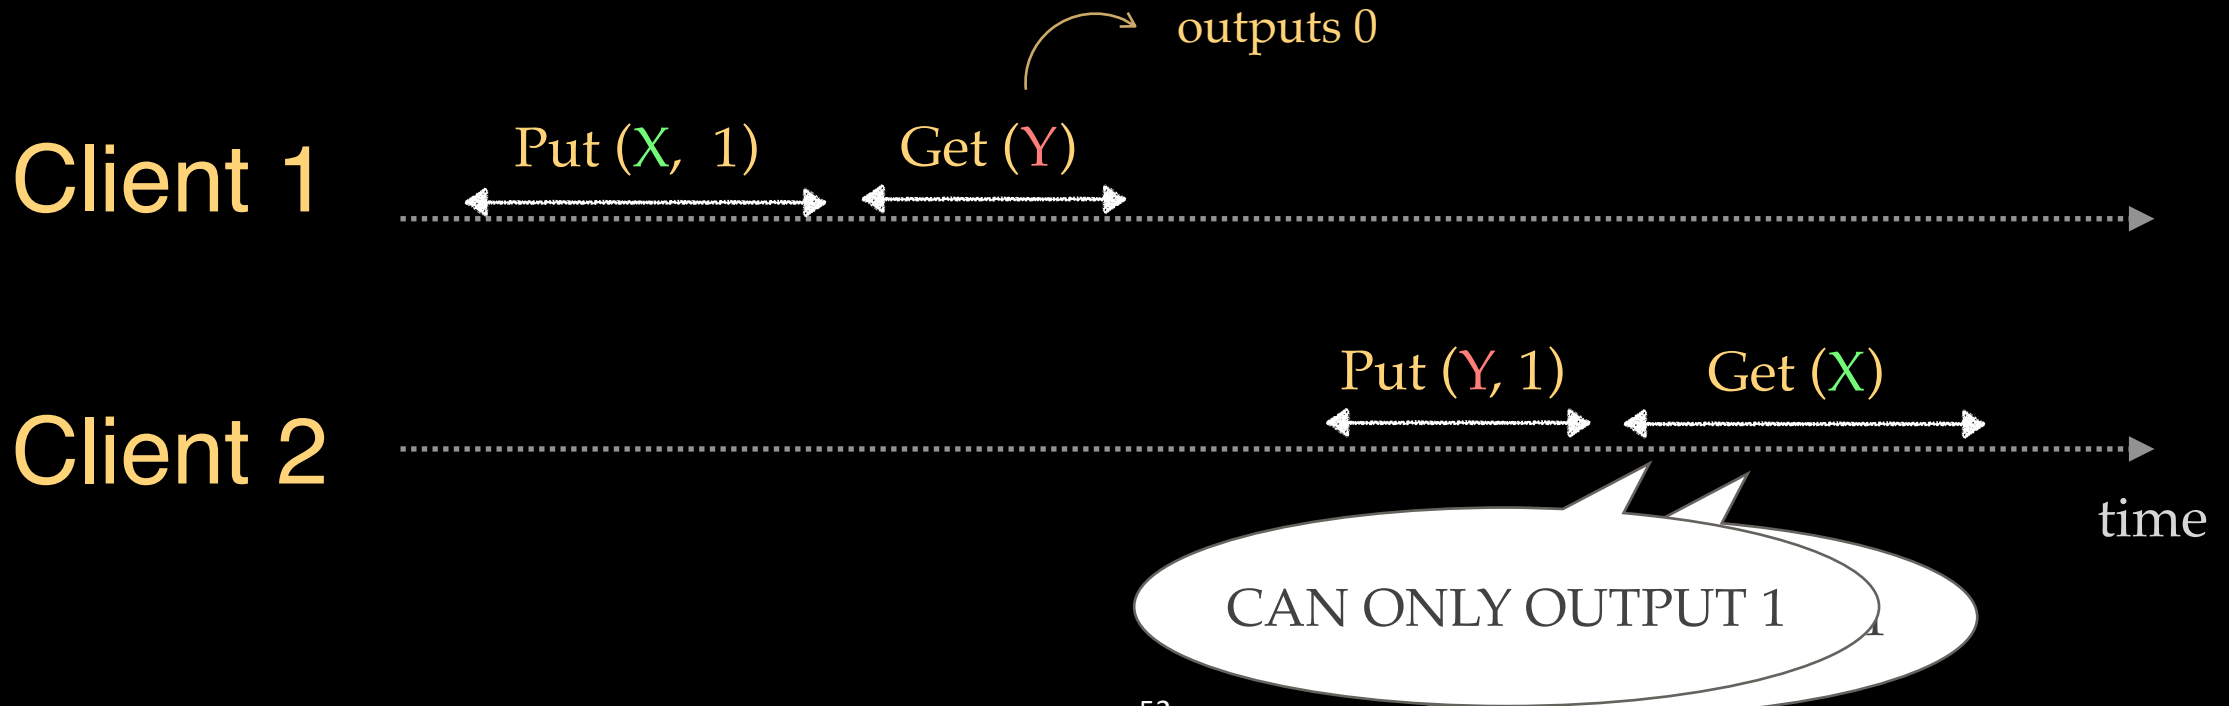


"I know my grades are straight 'Cs,' but don't I get some credit for consistency?"

Security of EKVS

good nodes → **Label X, Label Y** ← bad nodes

KVS is Sequentially Consistent



Security of EKVS

Single user setting

If KVS is $(\epsilon, \theta, \delta)$ -balanced, and **RYW consistent**, then

EKVS is L_ϵ -secure
with prob at least $1 - \delta - \text{negl}(k)$


repetition pattern
on pairs **visible to the
adversary**

Multi user setting

Clients can share data
concurrent operations on
same piece of data possible

Security of EKVS

Single Client



will always output 1
because in *single-user* setting
RYW guarantees
Get(X) reads last Put(X) independently
of operations on Y

Security of EKVS

Single user setting

If KVS is $(\epsilon, \theta, \delta)$ -balanced, and **RYW consistent**, then EKVS is L_ϵ -secure with prob at least $1 - \delta - \text{negl}(k)$

repetition pattern on pairs **visible to the adversary**

Multi user setting

EKVS is L -secure with prob at least $1 - \text{negl}(k)$

repetition pattern on **all the pairs**

Outline

✓ (0) Introduction

✓ (I) Encrypted DHTs

✓ (II) Encrypted Key Value Stores

(III) Future Directions

Security of EKVS

Single user setting

If KVS is $(\epsilon, \theta, \delta)$ -balanced, and
RYW consistent, then
EKVS is L_ϵ -secure
with prob at least $1 - \delta - \text{negl}(k)$

Q1: What happens w/ other consistency guarantees?

Q2: Are stronger notions of consistency better for privacy?

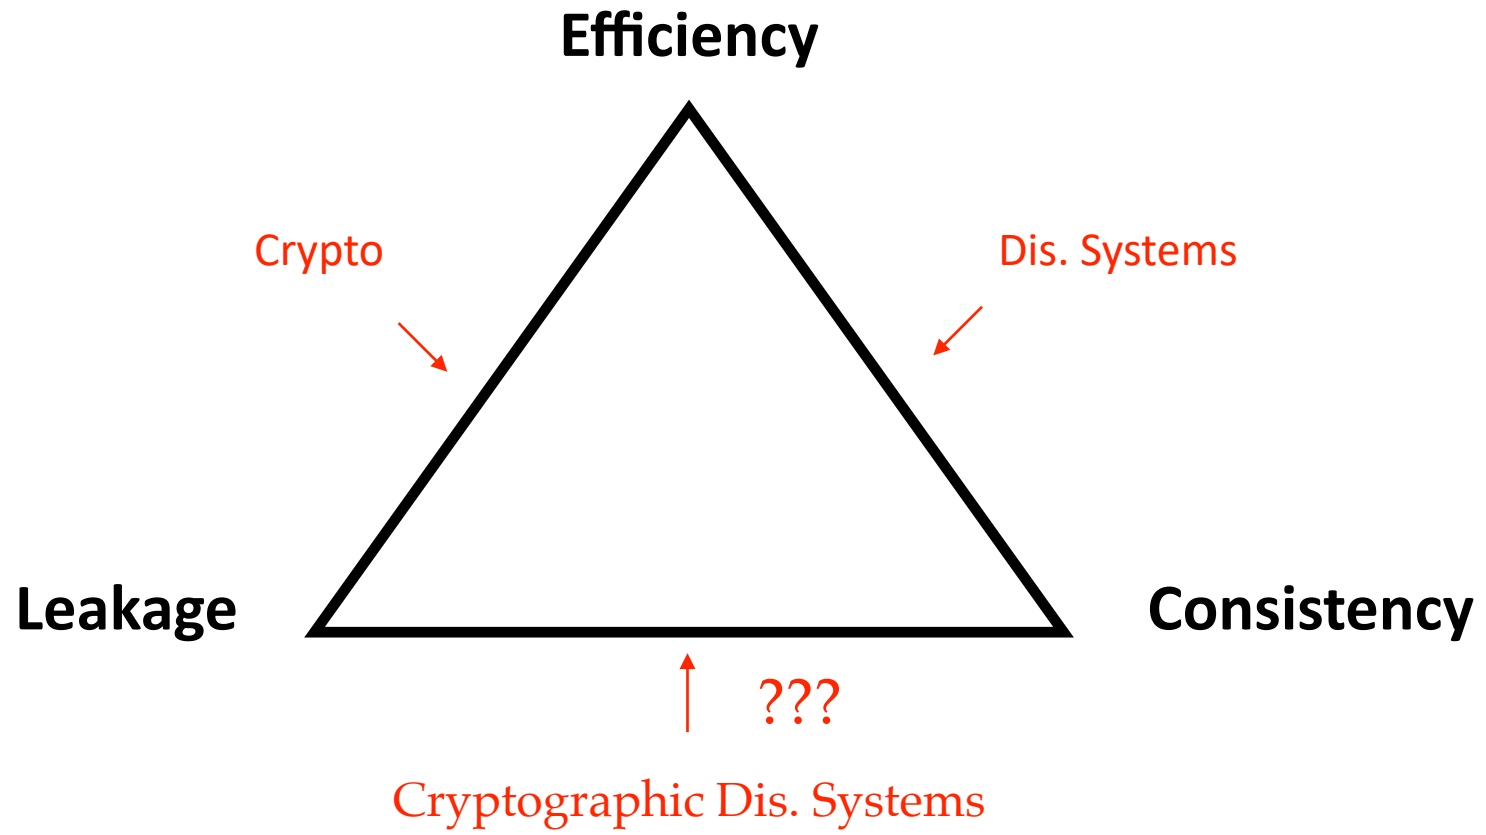
Security of EKVS

Q3: Can we improve security by assuming some consistency guarantees?

Q4: If no, can we show a lower bound on the leakage?

Multi user setting

EKVS is **L-secure**
with prob at least $1 - \text{negl}(k)$





- Acknowledgements

- Archita Agarwal, MongoDB
- Tarik Moataz, MongoDB

- References

- *Encrypted Distributed Storage Systems (thesis)*, A. Agarwal
- *Encrypted Distributed Hash Tables*, A. Agarwal, S. Kamara
- *Encrypted Key Value Stores*, A. Agarwal, S. Kamara

Thank You

Maake Asante Shukria Dhanyavadagalu
Vinaka Kaitos Kam Sah Hammida Manana Dankon
감사합니다 Dank Je Maururu Biyan Matondo
Blagodaram Dankscheen Chokrane Diolch i Chi Terima Kasih Tack
Niyabonga Dziekuje Arigato Grazie Mochchakkeram
Juspaxar Gracias Tingki
நன்றி Bedankt धन्यवाद cam ơn ban Khap Paldies Gratias Tibi
Ua Tsaug Rau Koj Dakujem Grazas Nirringrazzjak Obrigado
Suksama Dėkuji Welalin Di Ou Mési Kia Ora Kop Khun Khap ありがとう
Misaotra Rahmat Matur Nuwun 谢 Di Ou Mési Kia Ora Kop Khun Khap Djijere Dieuf
Najis Tuke Eskerrik Askó