

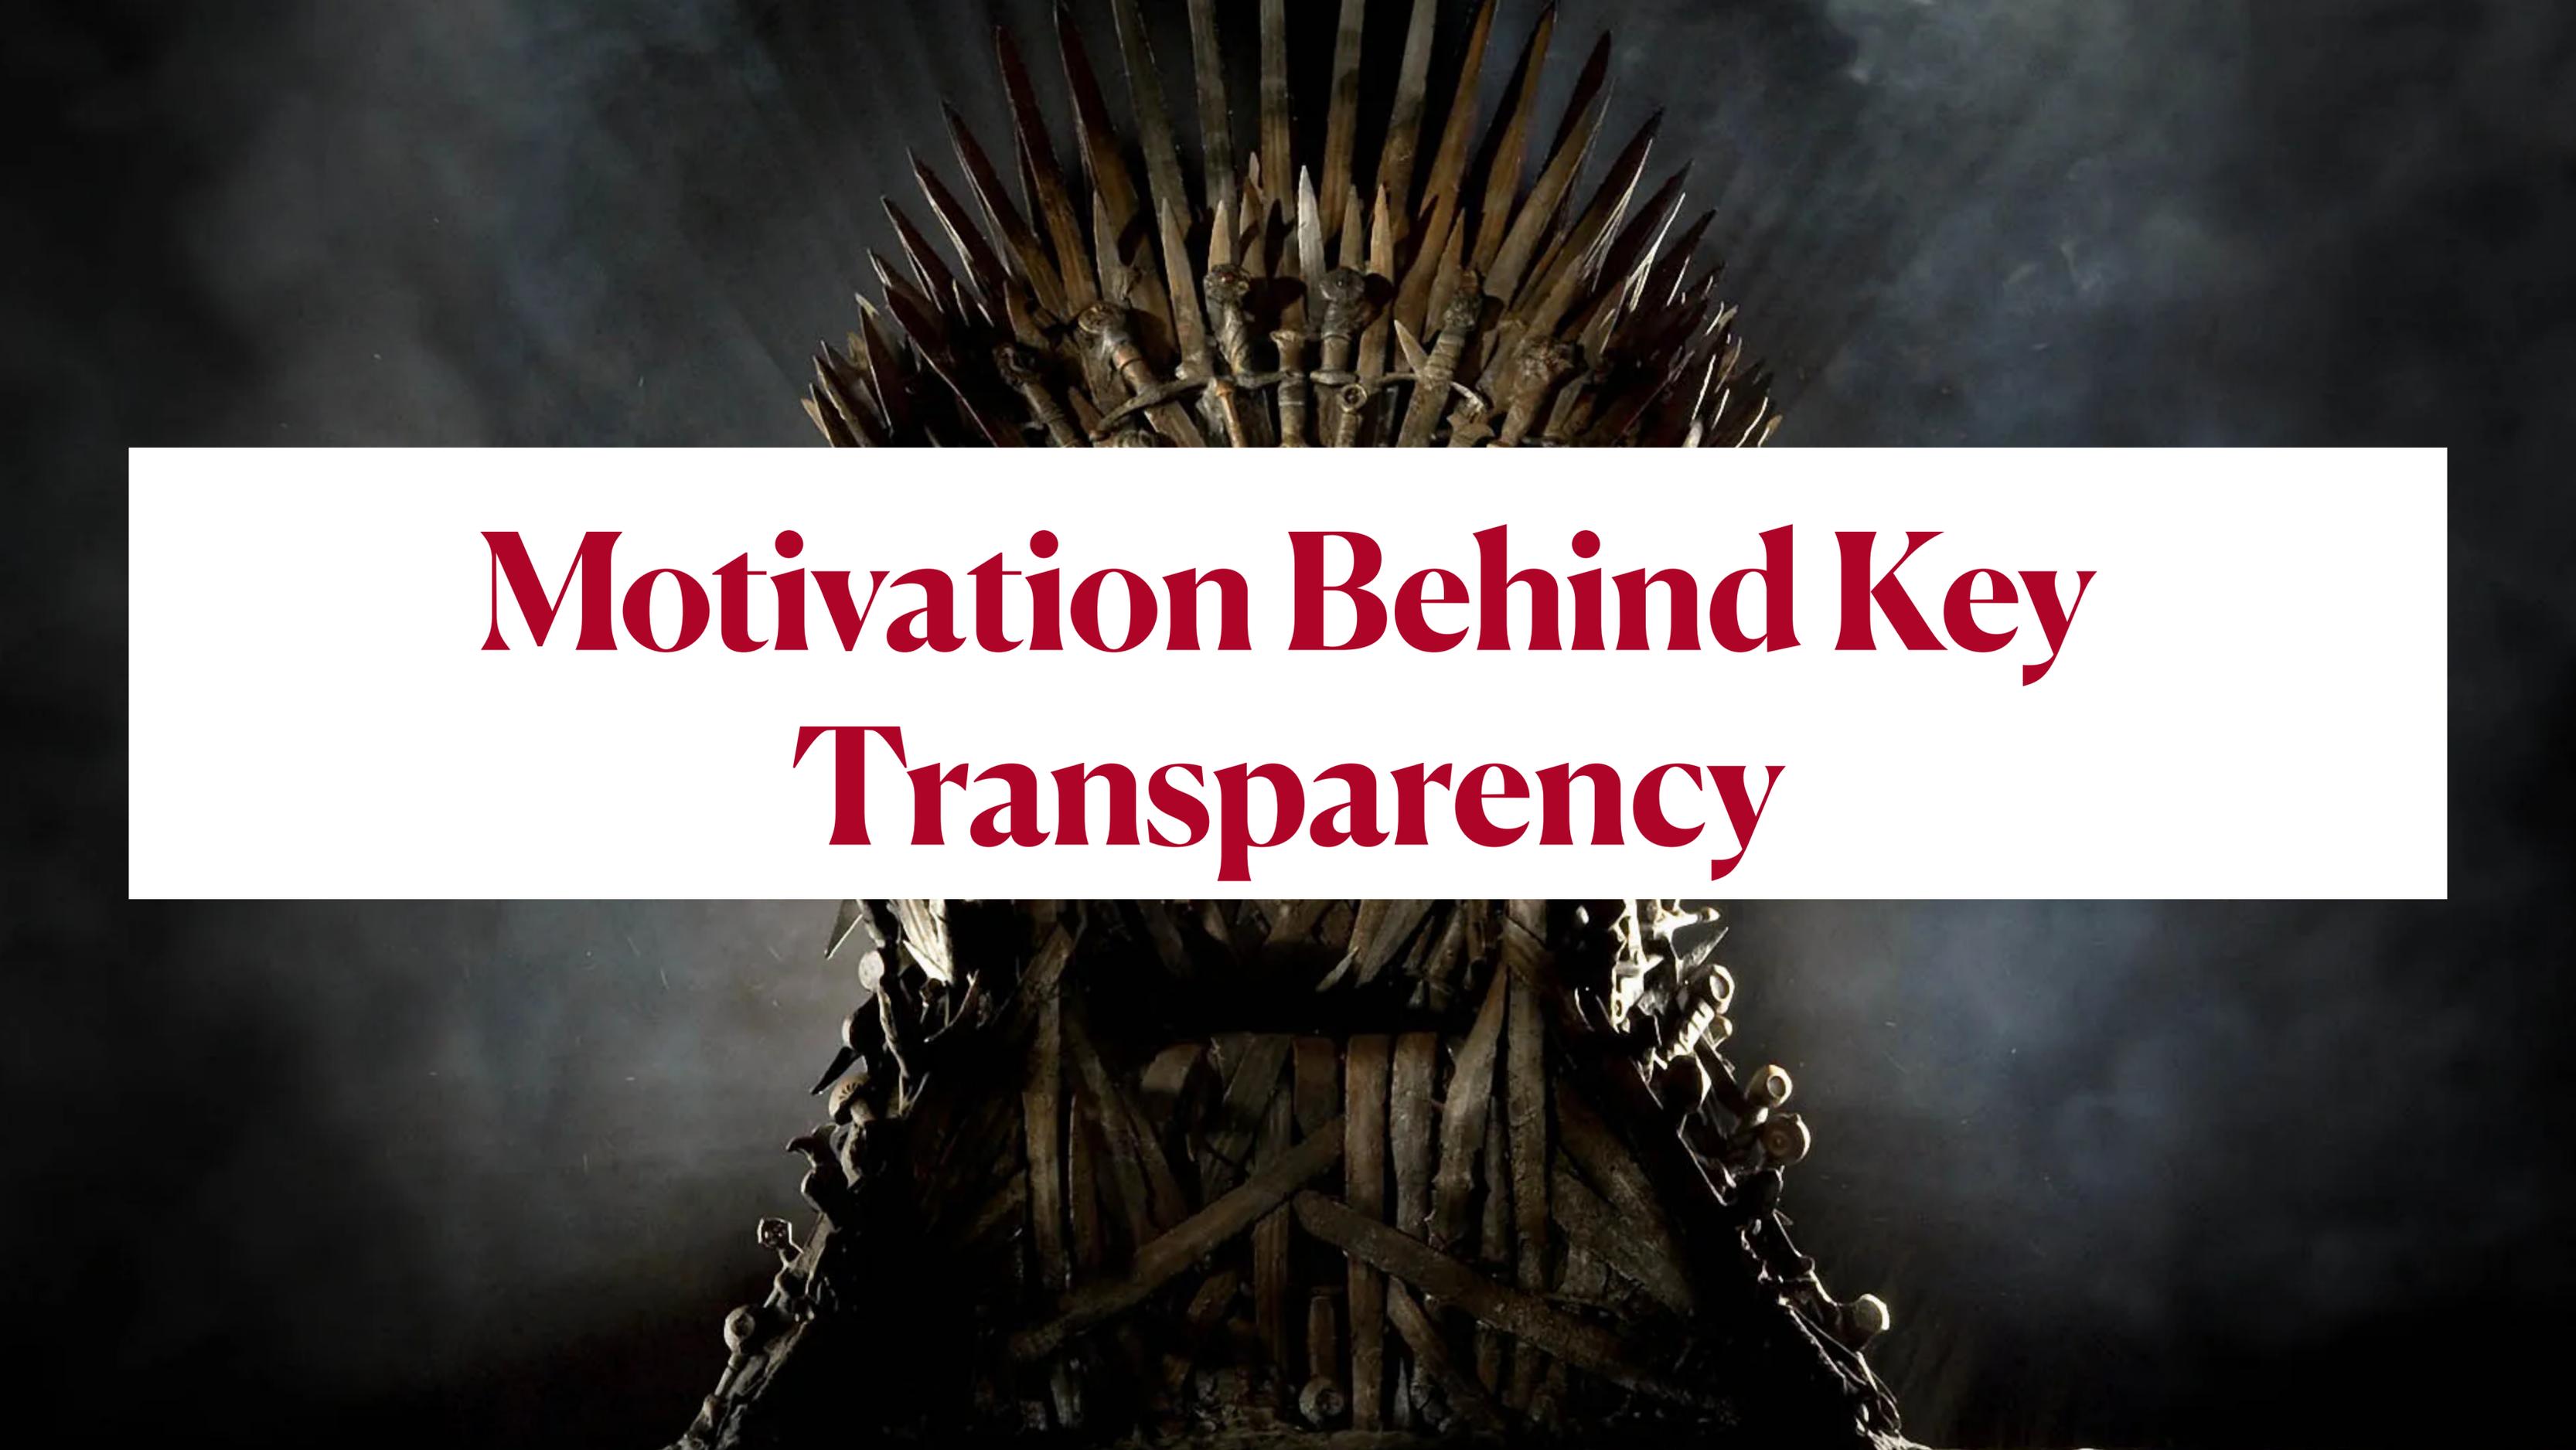
Self-Auditable Key Transparency

Hossein Hafezi



UNIVERSITY OF
CAMBRIDGE





Motivation Behind Key Transparency

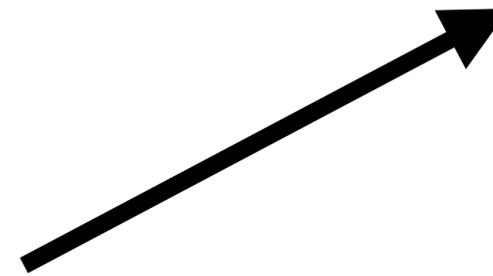
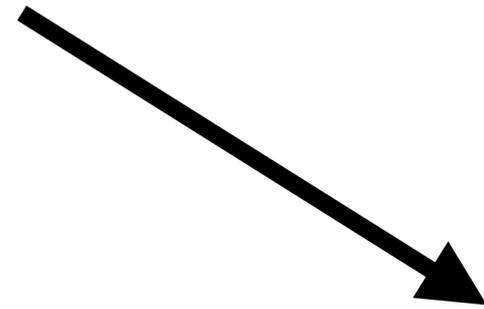
End-to-End Encrypted Messaging



$(pk_{\text{Jon}}, pk_{\text{Ygritte}})$



$(pk_{\text{Jon}}, pk_{\text{Ygritte}})$



End-to-End Encrypted Messaging



$(pk_{\text{Jon}}, pk_{\text{Ygritte}})$



$(pk_{\text{Jon}}, pk_{\text{Ygritte}})$

**The service provider only transmits
ciphertexts and cannot decrypt
them**



End-to-End Encrypted Messaging



$(pk_{\text{Jon}}, pk_{\text{Ygritte}})$



$(pk_{\text{Jon}}, pk_{\text{Ygritte}})$

The service provider only transmits ciphertexts and cannot decrypt them

How do users get each other public keys from example from a phone number?



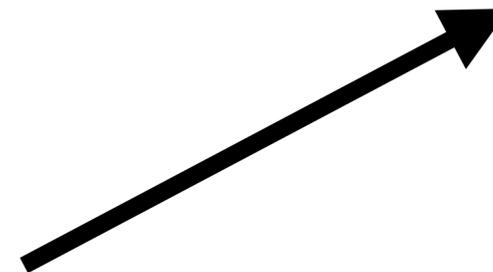
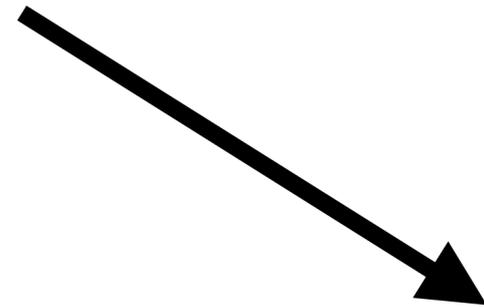
End-to-End Encrypted Messaging



$(pk_{\text{Jon}}, pk_{\text{NK}})$



$(pk_{\text{Jon}}, pk_{\text{NK}})$



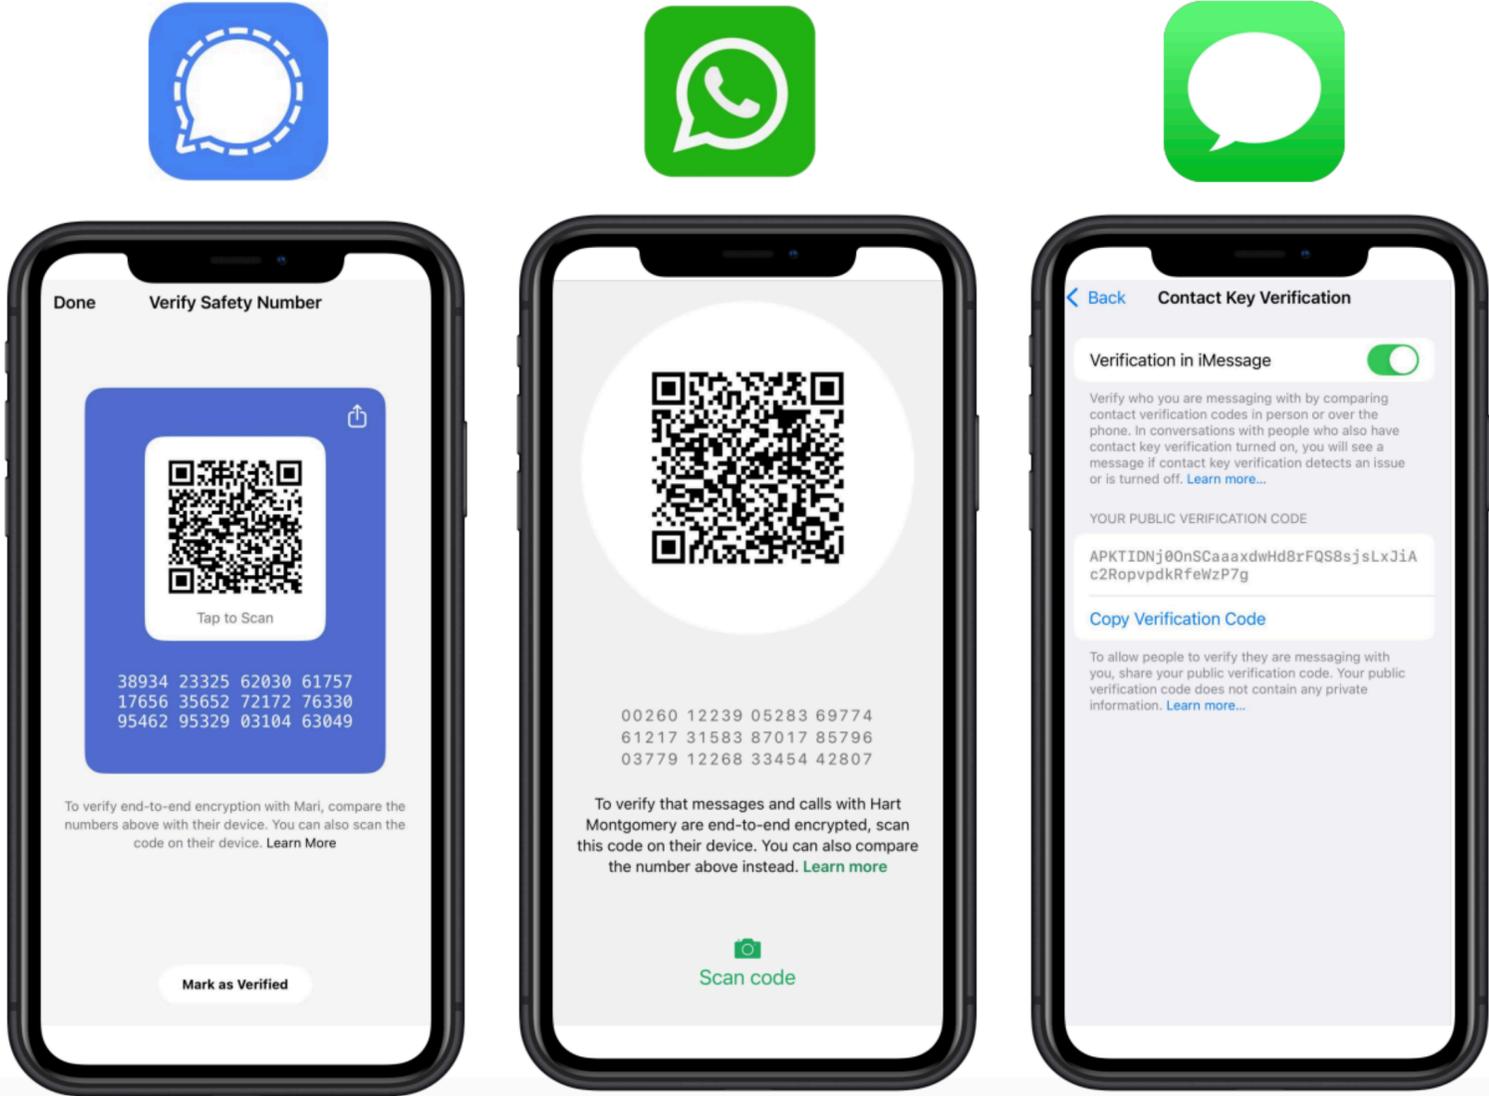
Give Jon a malicious public key



"You know nothing, Jon Snow."



Traditional Solution: Out-of-Band Verification



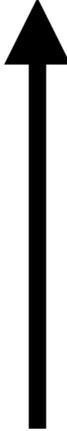
Traditional Solution: Out-of-Band Verification



Key Transparency: CONIKS [MBBFF15]



Key Transparency: CONIKS [MBBFF15]



Periodically post
Signed commitments
to directory



Key Transparency: CONIKS [MBBFF15]



Periodically post
Signed commitments
to directory

$pk_{Ygritte}, \pi$



Key Transparency: CONIKS [MBBFF15]

Verify pk is consistent
with the commitment to
the bulletin board



Periodically post
Signed commitments
to directory

$pk_{Ygritte}, \pi$



Key Transparency: CONIKS [MBBFF15]

Verify pk is consistent
with the commitment to
the bulletin board



Periodically post
Signed commitments
to directory

$pk_{Ygritte}, \pi$



$pk_{Ygritte}, \pi$

Key Transparency: CONIKS [MBBFF15]

Verify pk is consistent with the commitment to the bulletin board



Verify my key is included correctly in the latest epoch



Periodically post Signed commitments to directory

$pk_{Ygritte}, \pi$



$pk_{Ygritte}, \pi$

Key Transparency: CONIKS [MBBFF15]

Verify pk is consistent with the commitment to the bulletin board



Verify my key is included correctly in the latest epoch



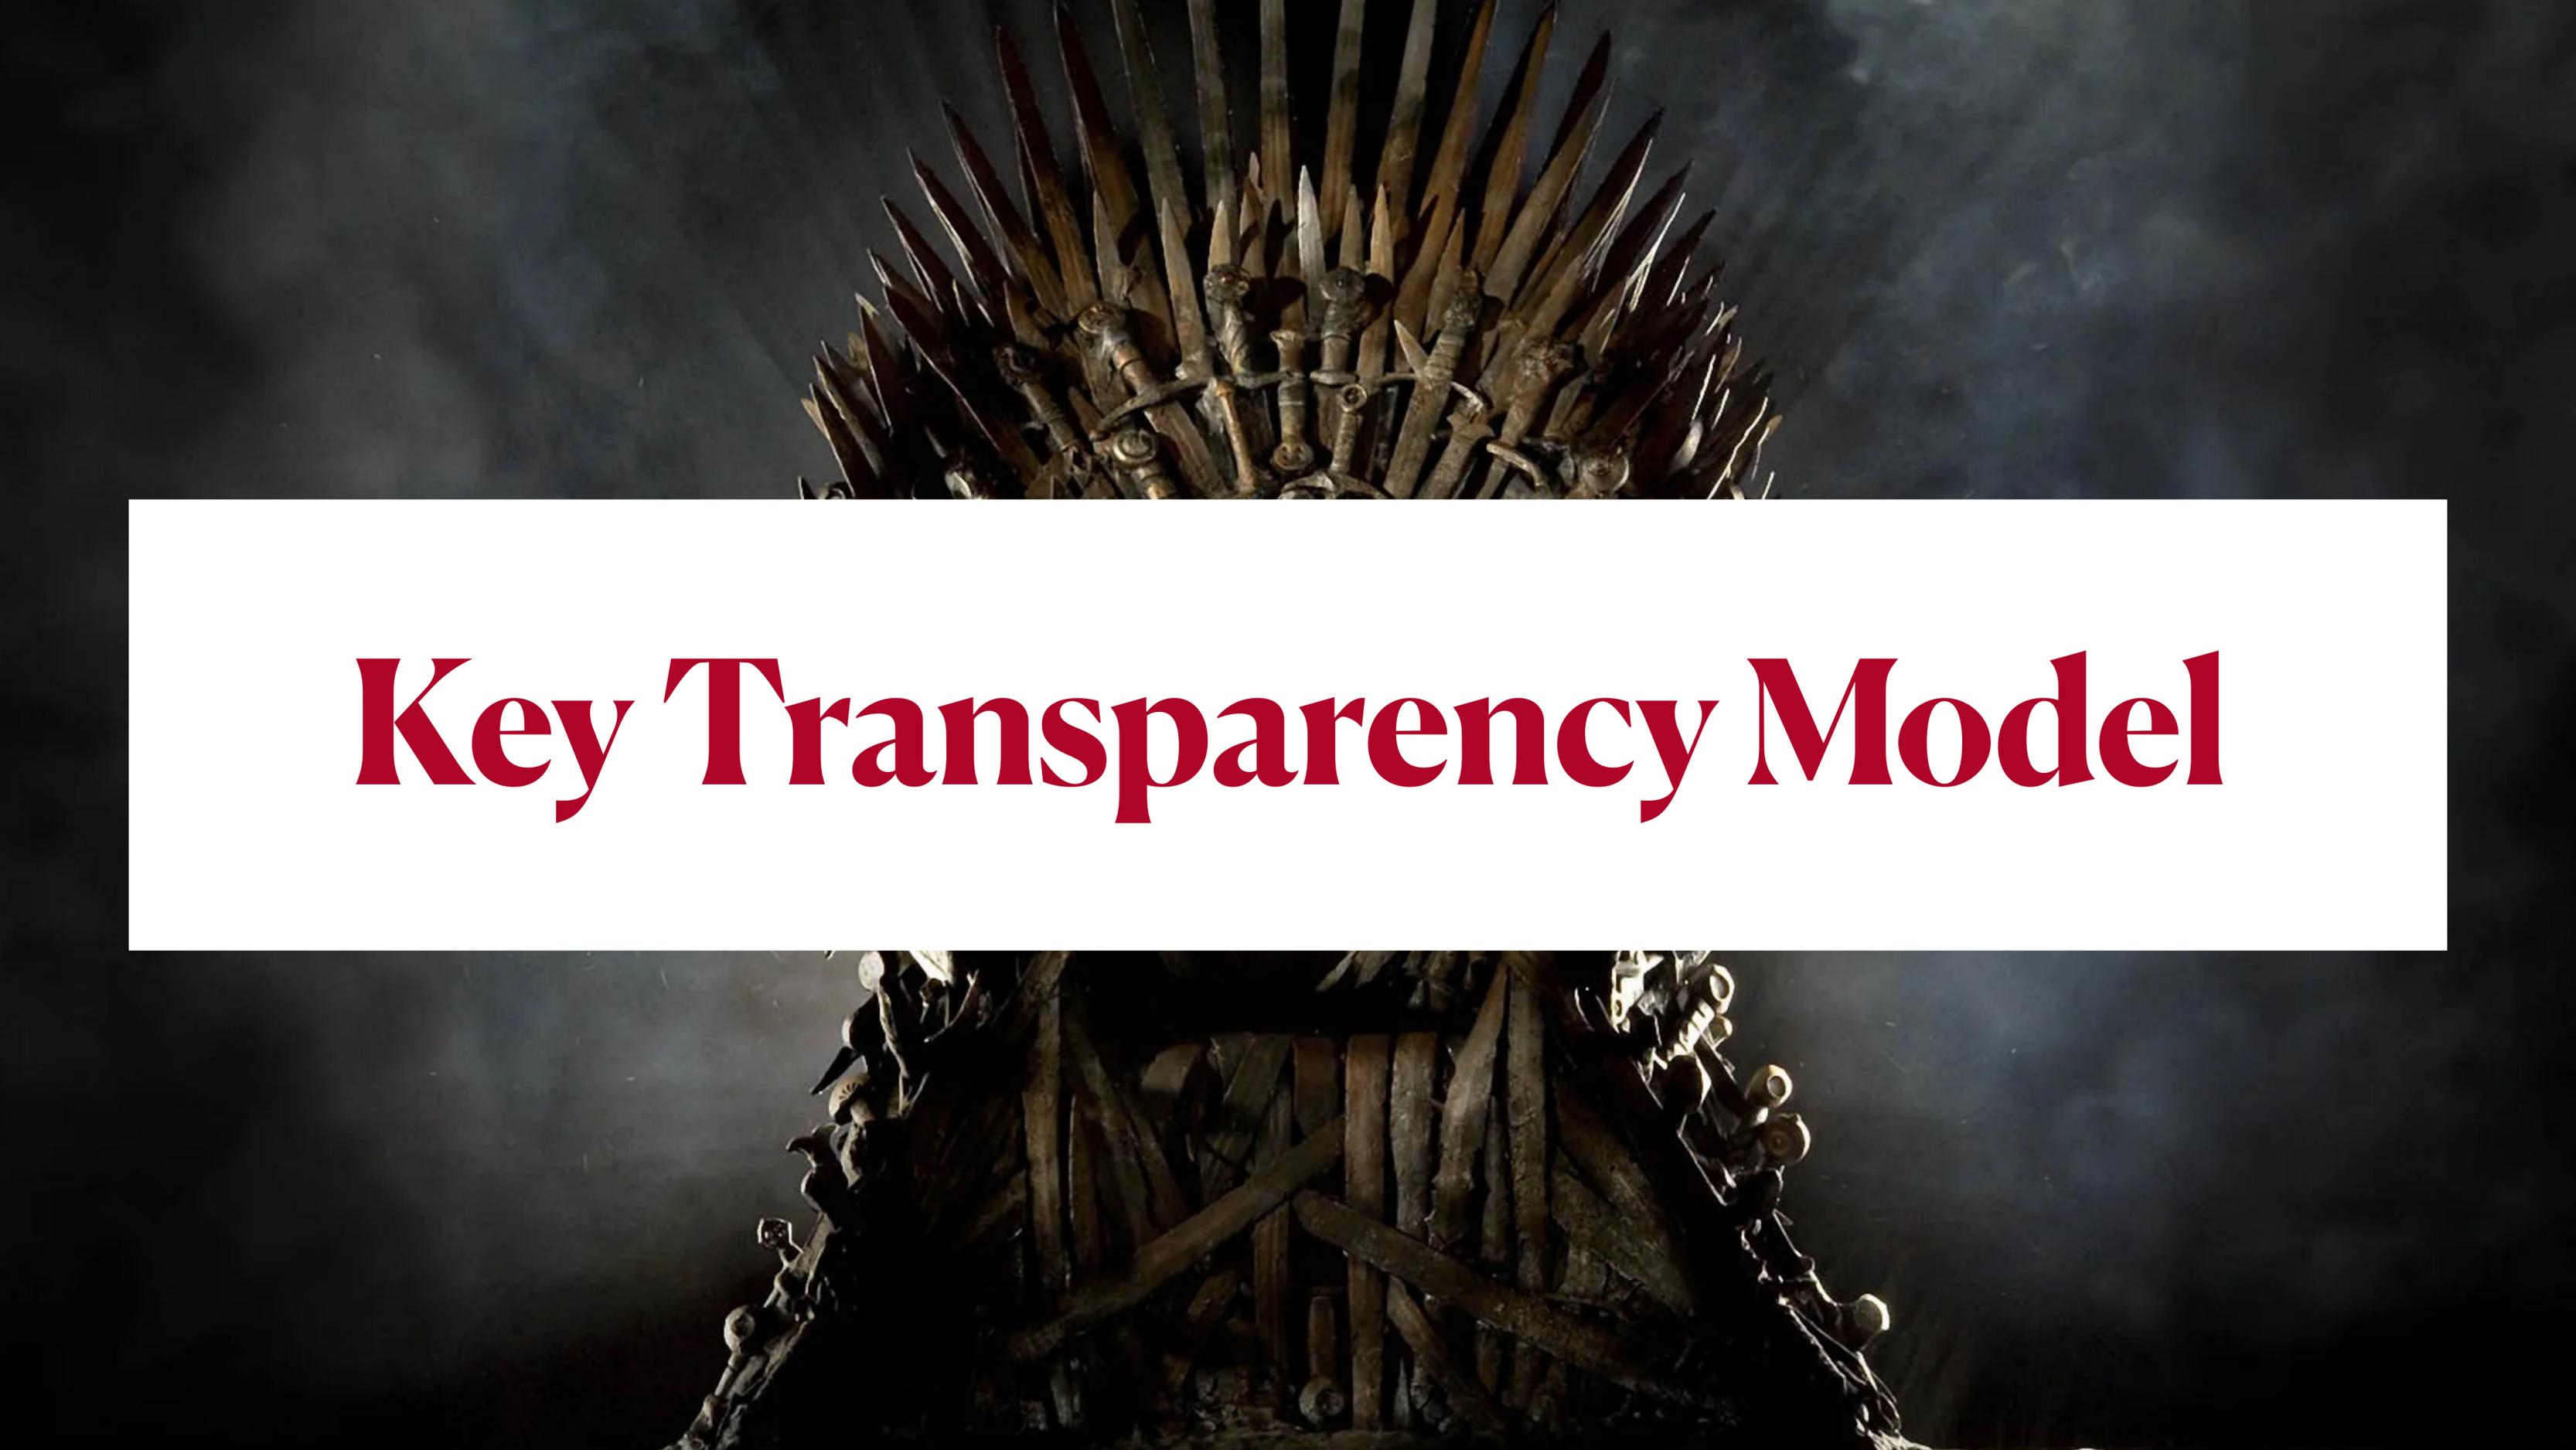
Periodically post Signed commitments to directory

$pk_{Ygritte}, \pi$



$pk_{Ygritte}, \pi$

AUTOMATED SOLUTION: no manual effort for Jon and Ygrittee



Key Transparency Model

Key Transparency Scheme

- Key transparency scheme consists of a server that at every epoch publishes the new commitment to the directory along with auxiliary data (invariance proof)



Commitment 0



Commitment 1

...



Commitment n



Key Transparency Scheme



Jon

(I) what is Ygritte's latest public key?



Lookup

(II) Prove my key hasn't changed since last year?

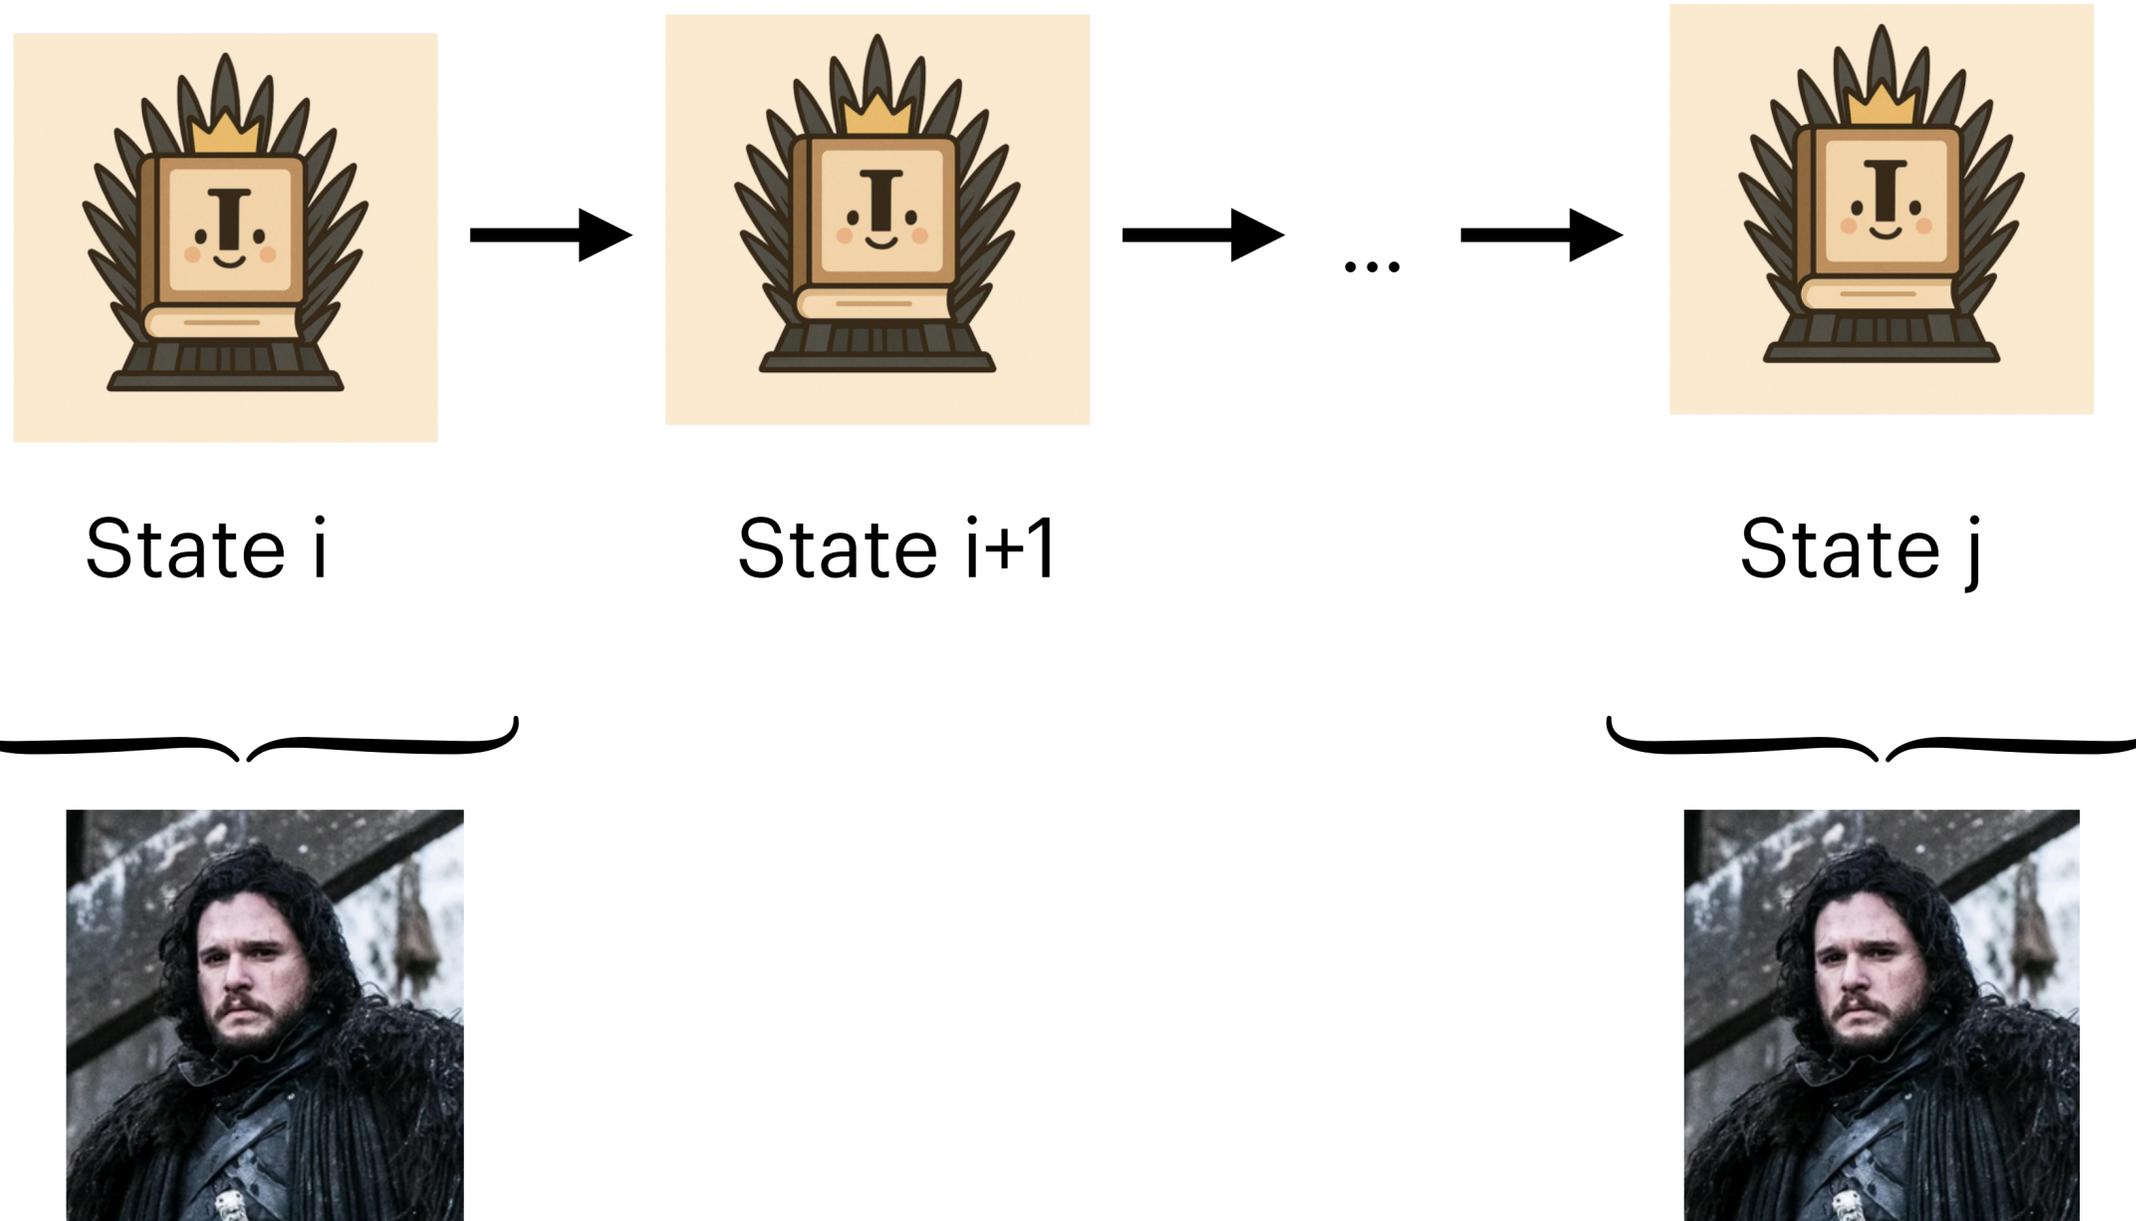


Proof of consistency



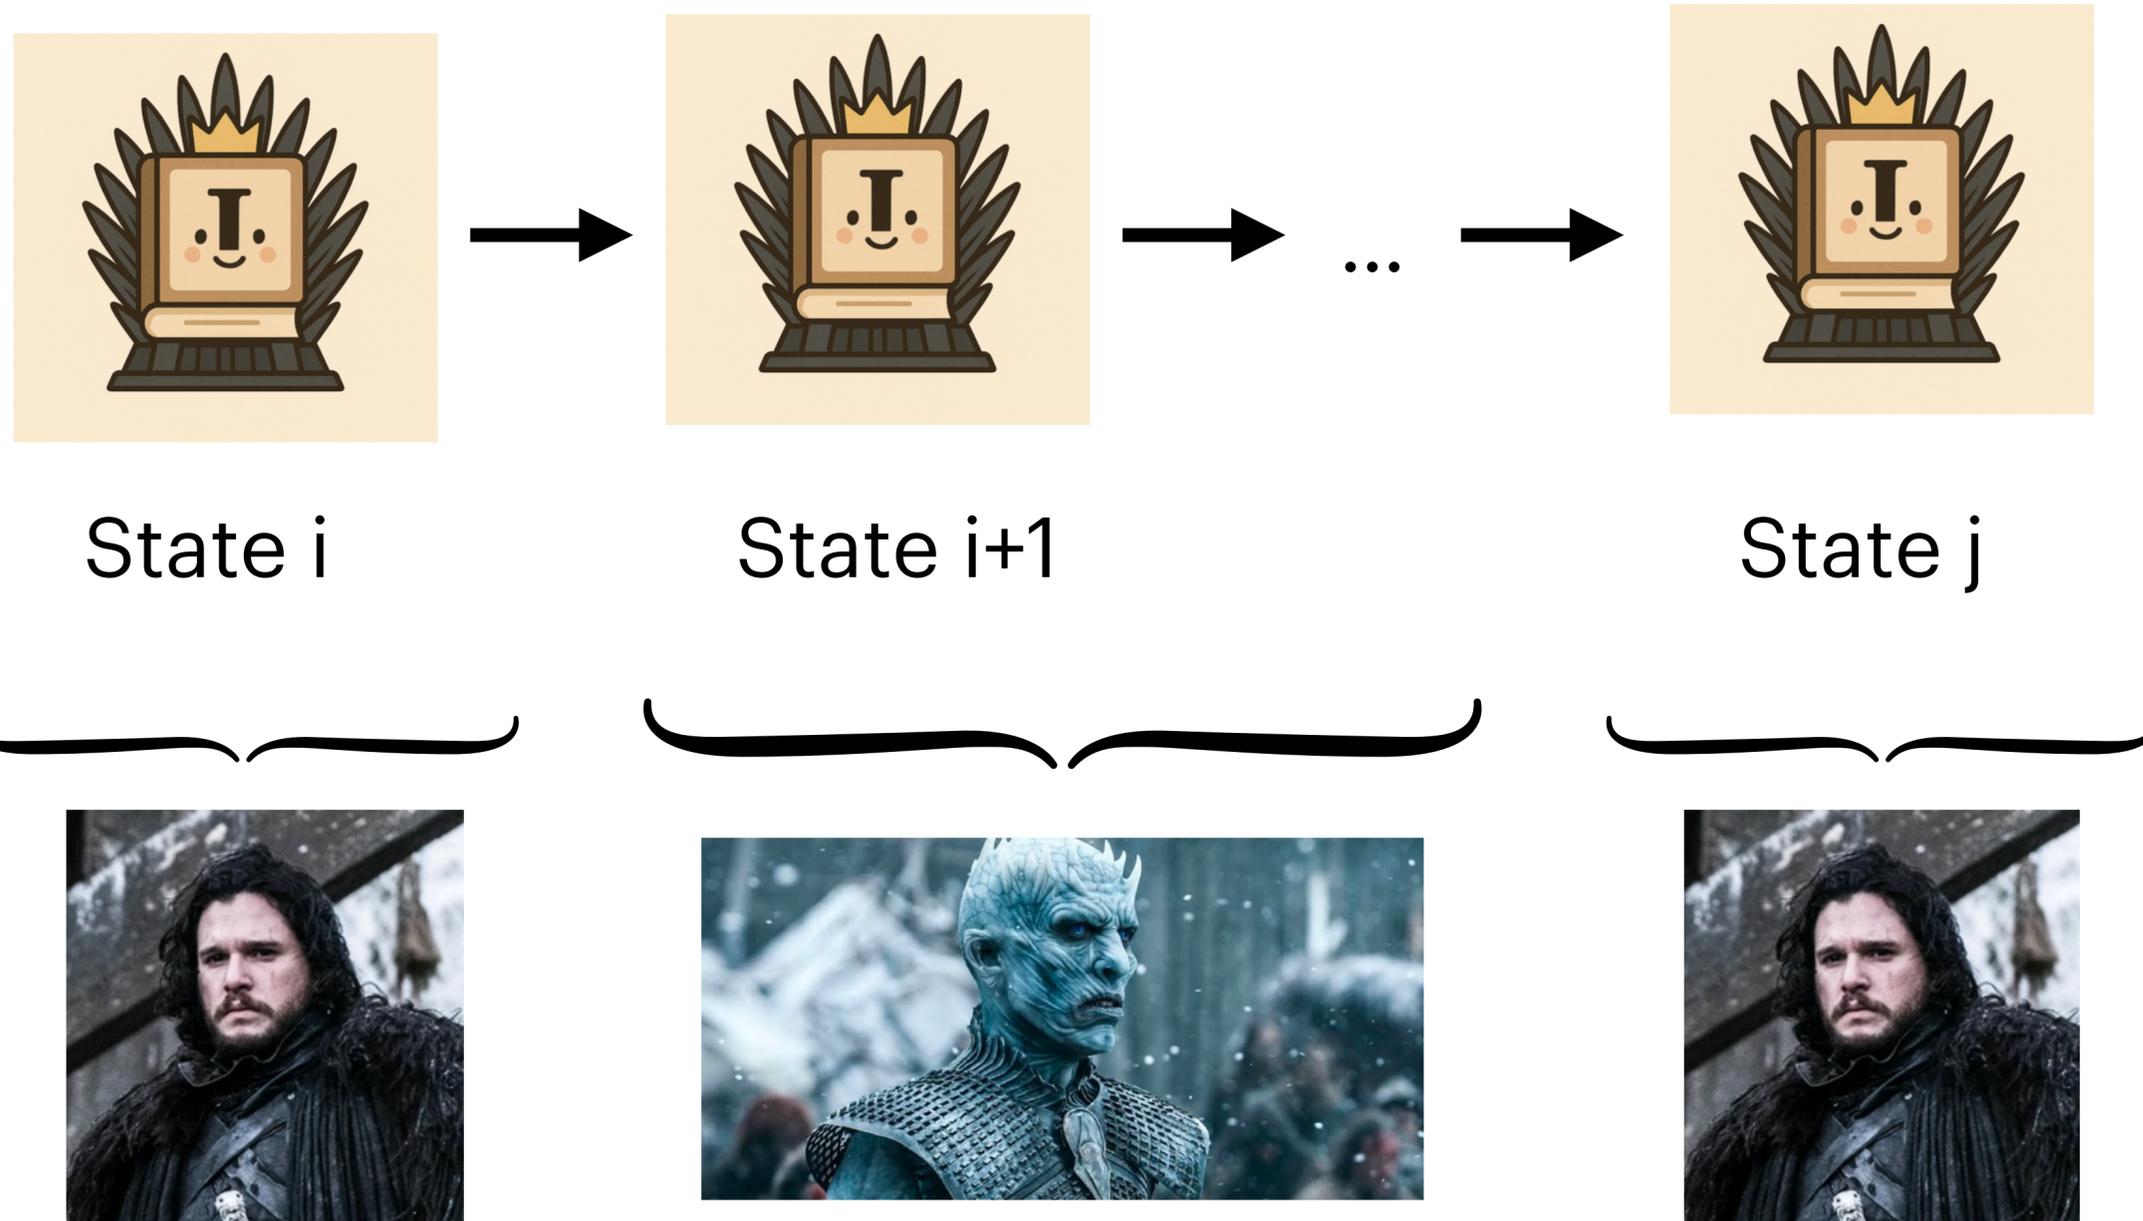
Necessity of Proof of Consistency

Assume Jon is offline from state $i+1$ to state $j-1$



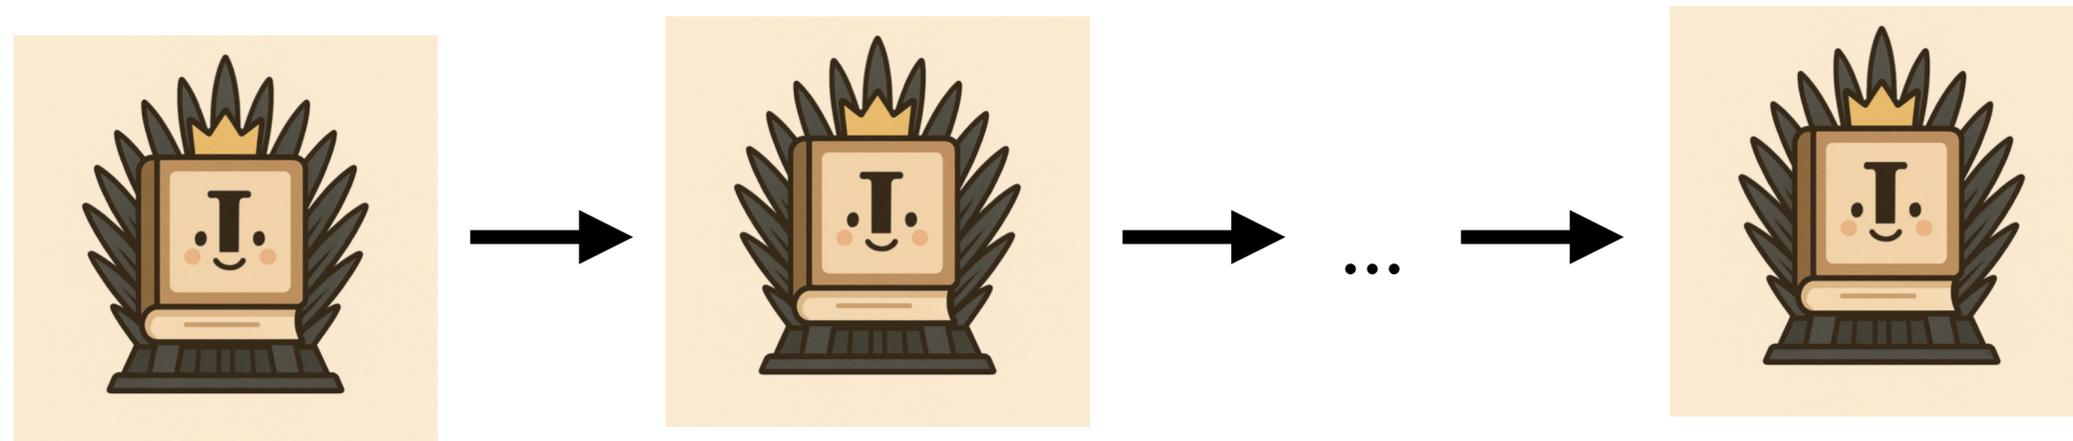
Necessity of Proof of Consistency

Assume Jon is offline from state $i+1$ to state $j-1$



Necessity of Proof of Consistency

Assume Jon is offline from state $i+1$ to state $j-1$



State i

State $i+1$

State j

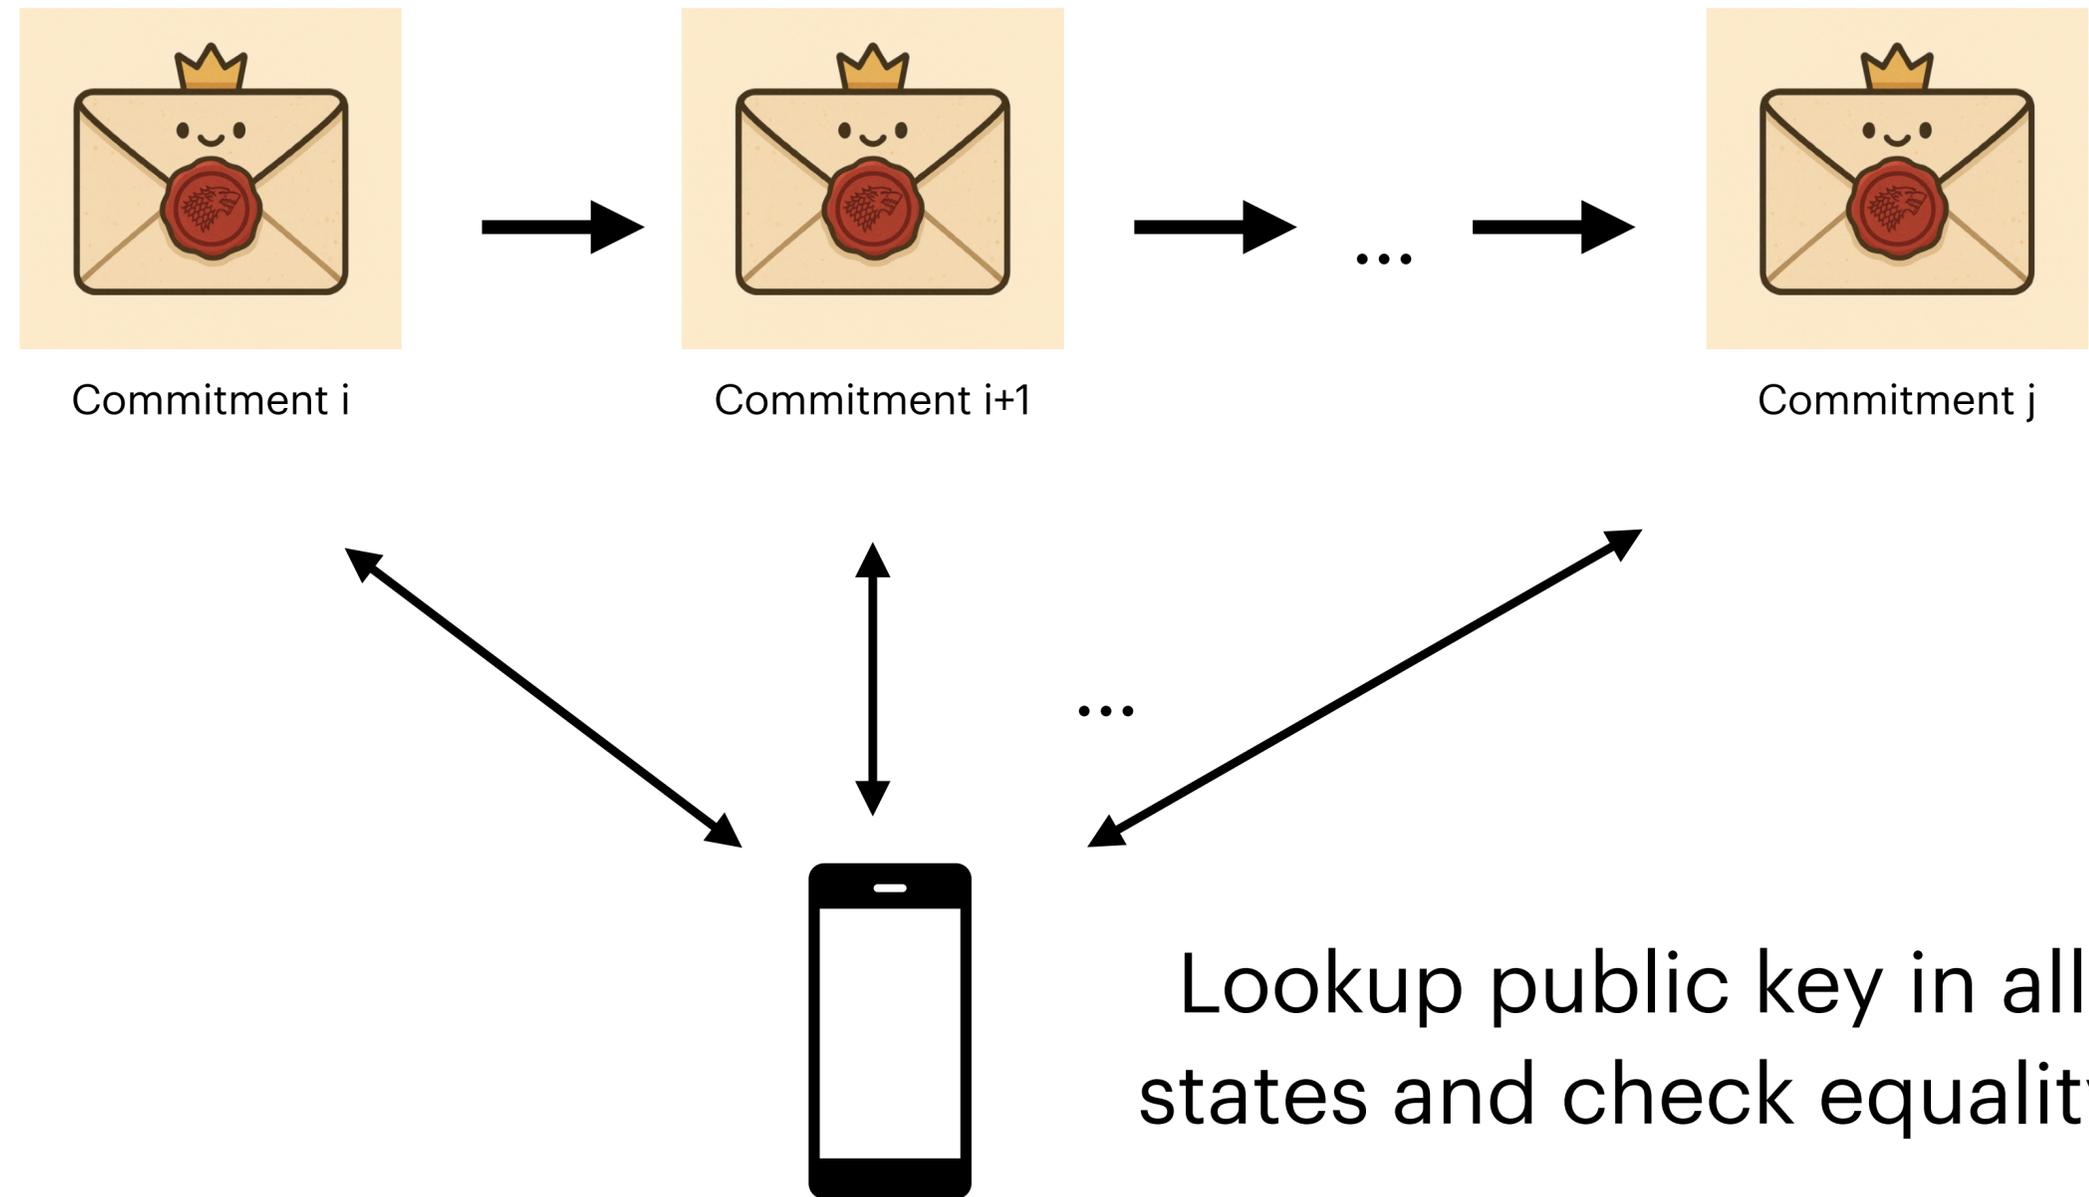


Request a proof
of consistency
from epoch i to
epoch j

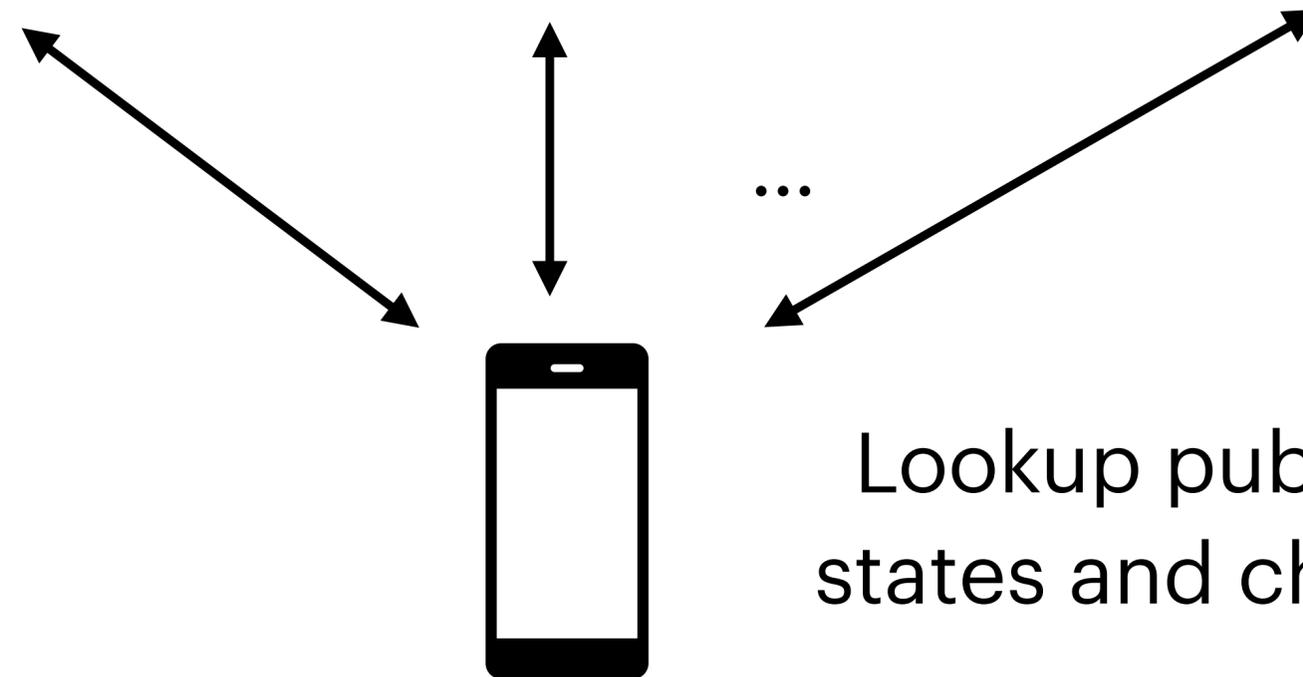
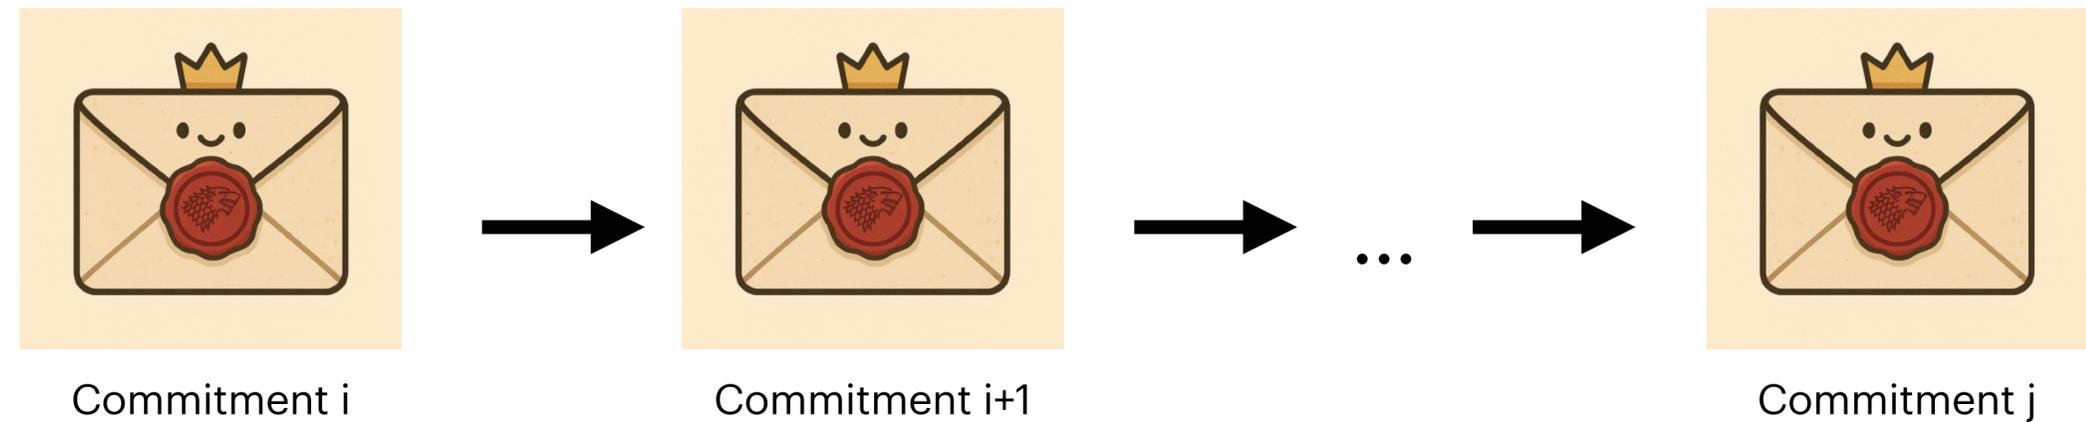
Different Solutions to Proof of Consistency

- Naively performing lookup for every epoch, e.g. CONIKS [Mel+15]
- Version Invariant, e.g. SEEMless [Cha+19], Parakeet [Mal+23], Optiks [Len+24]
- Random Version Invariant, e.g. IronDict [Haf+26] and Aegon [Haf+26]

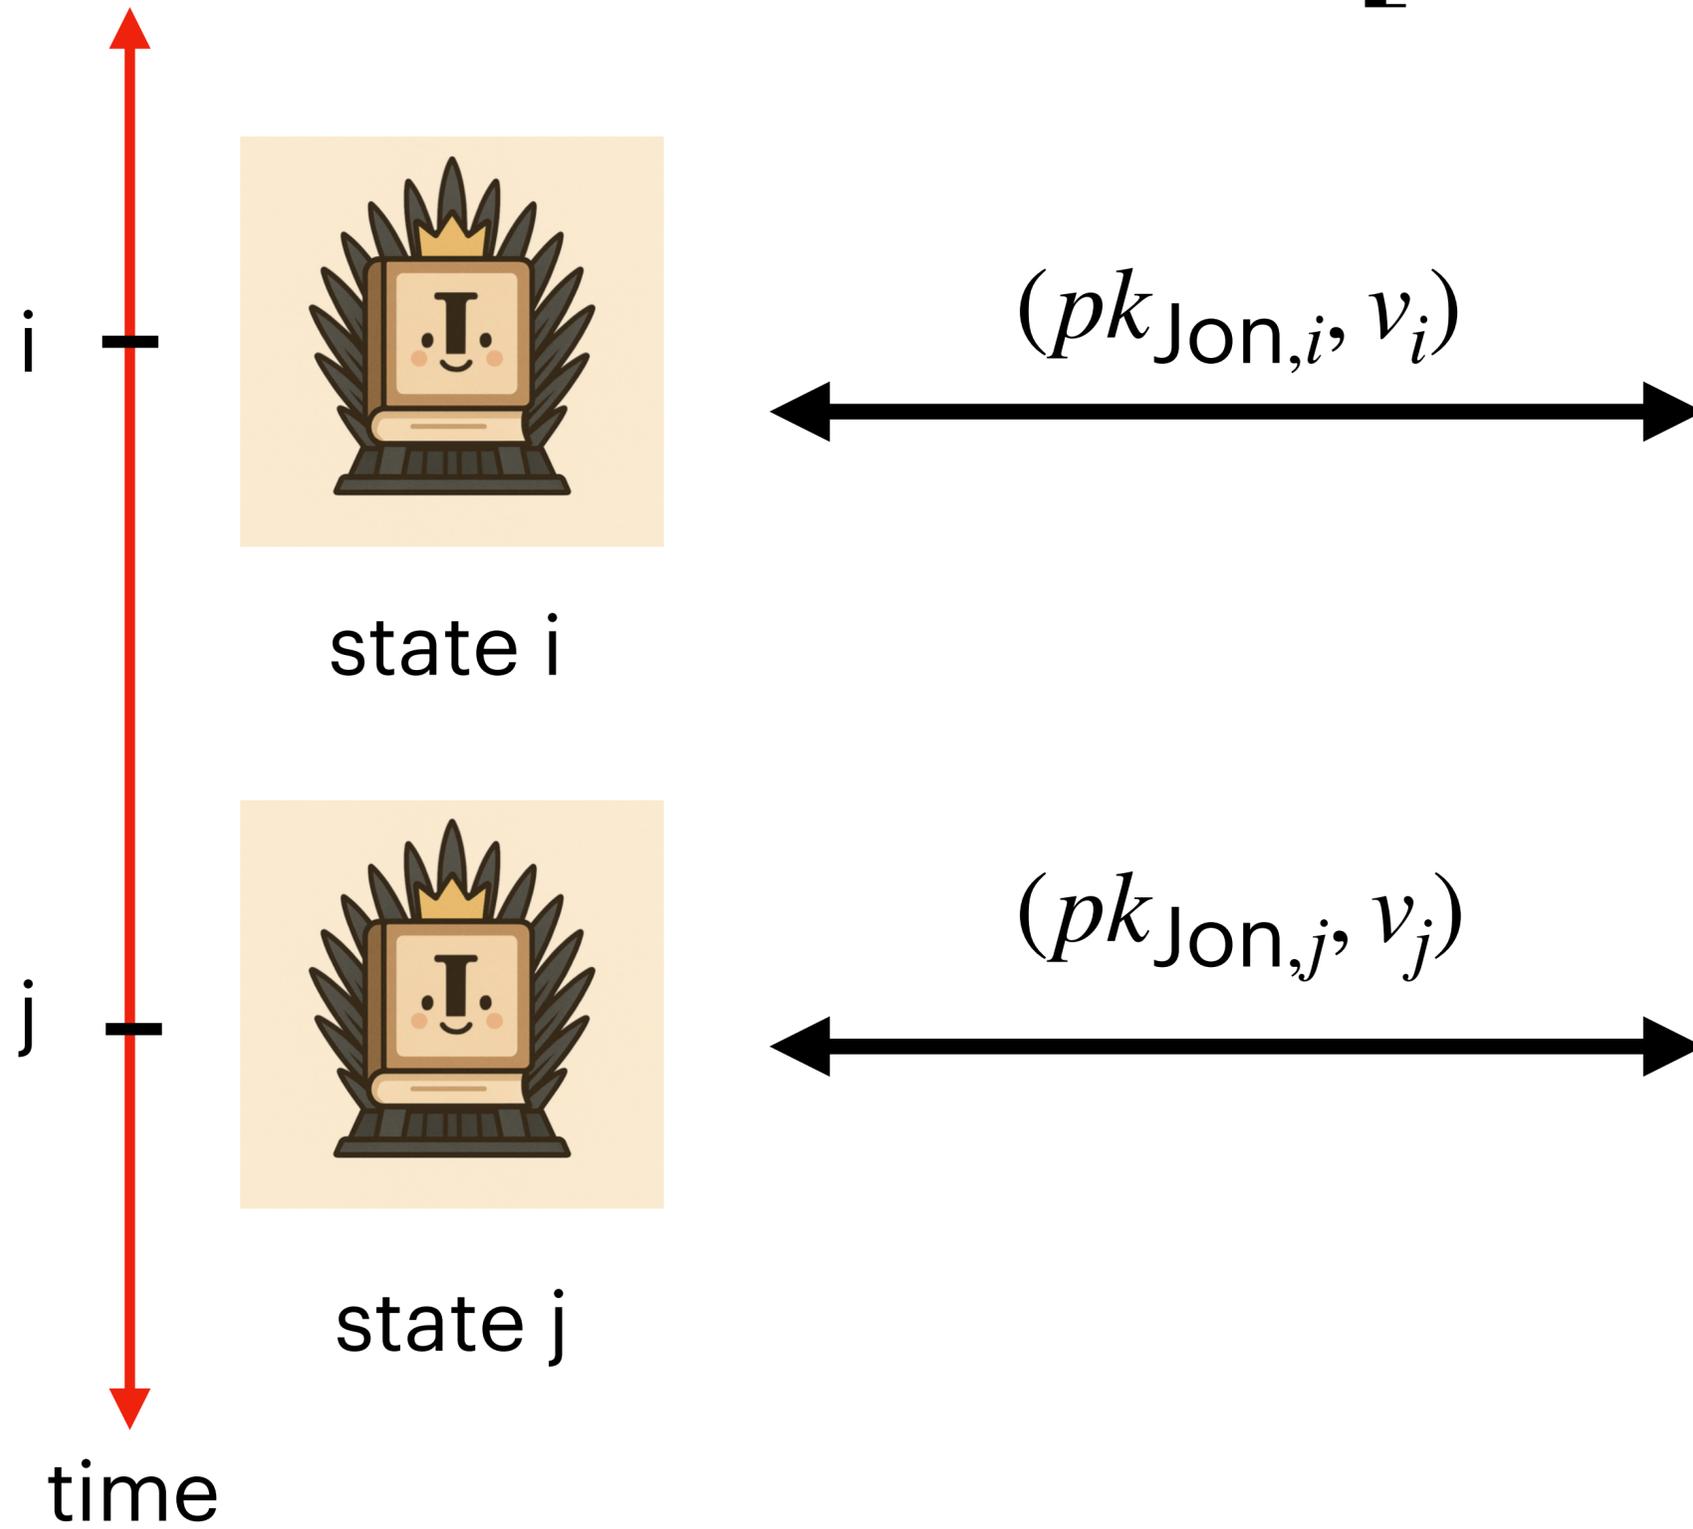
Proof of Consistency: Naive Solution [Mal+15]



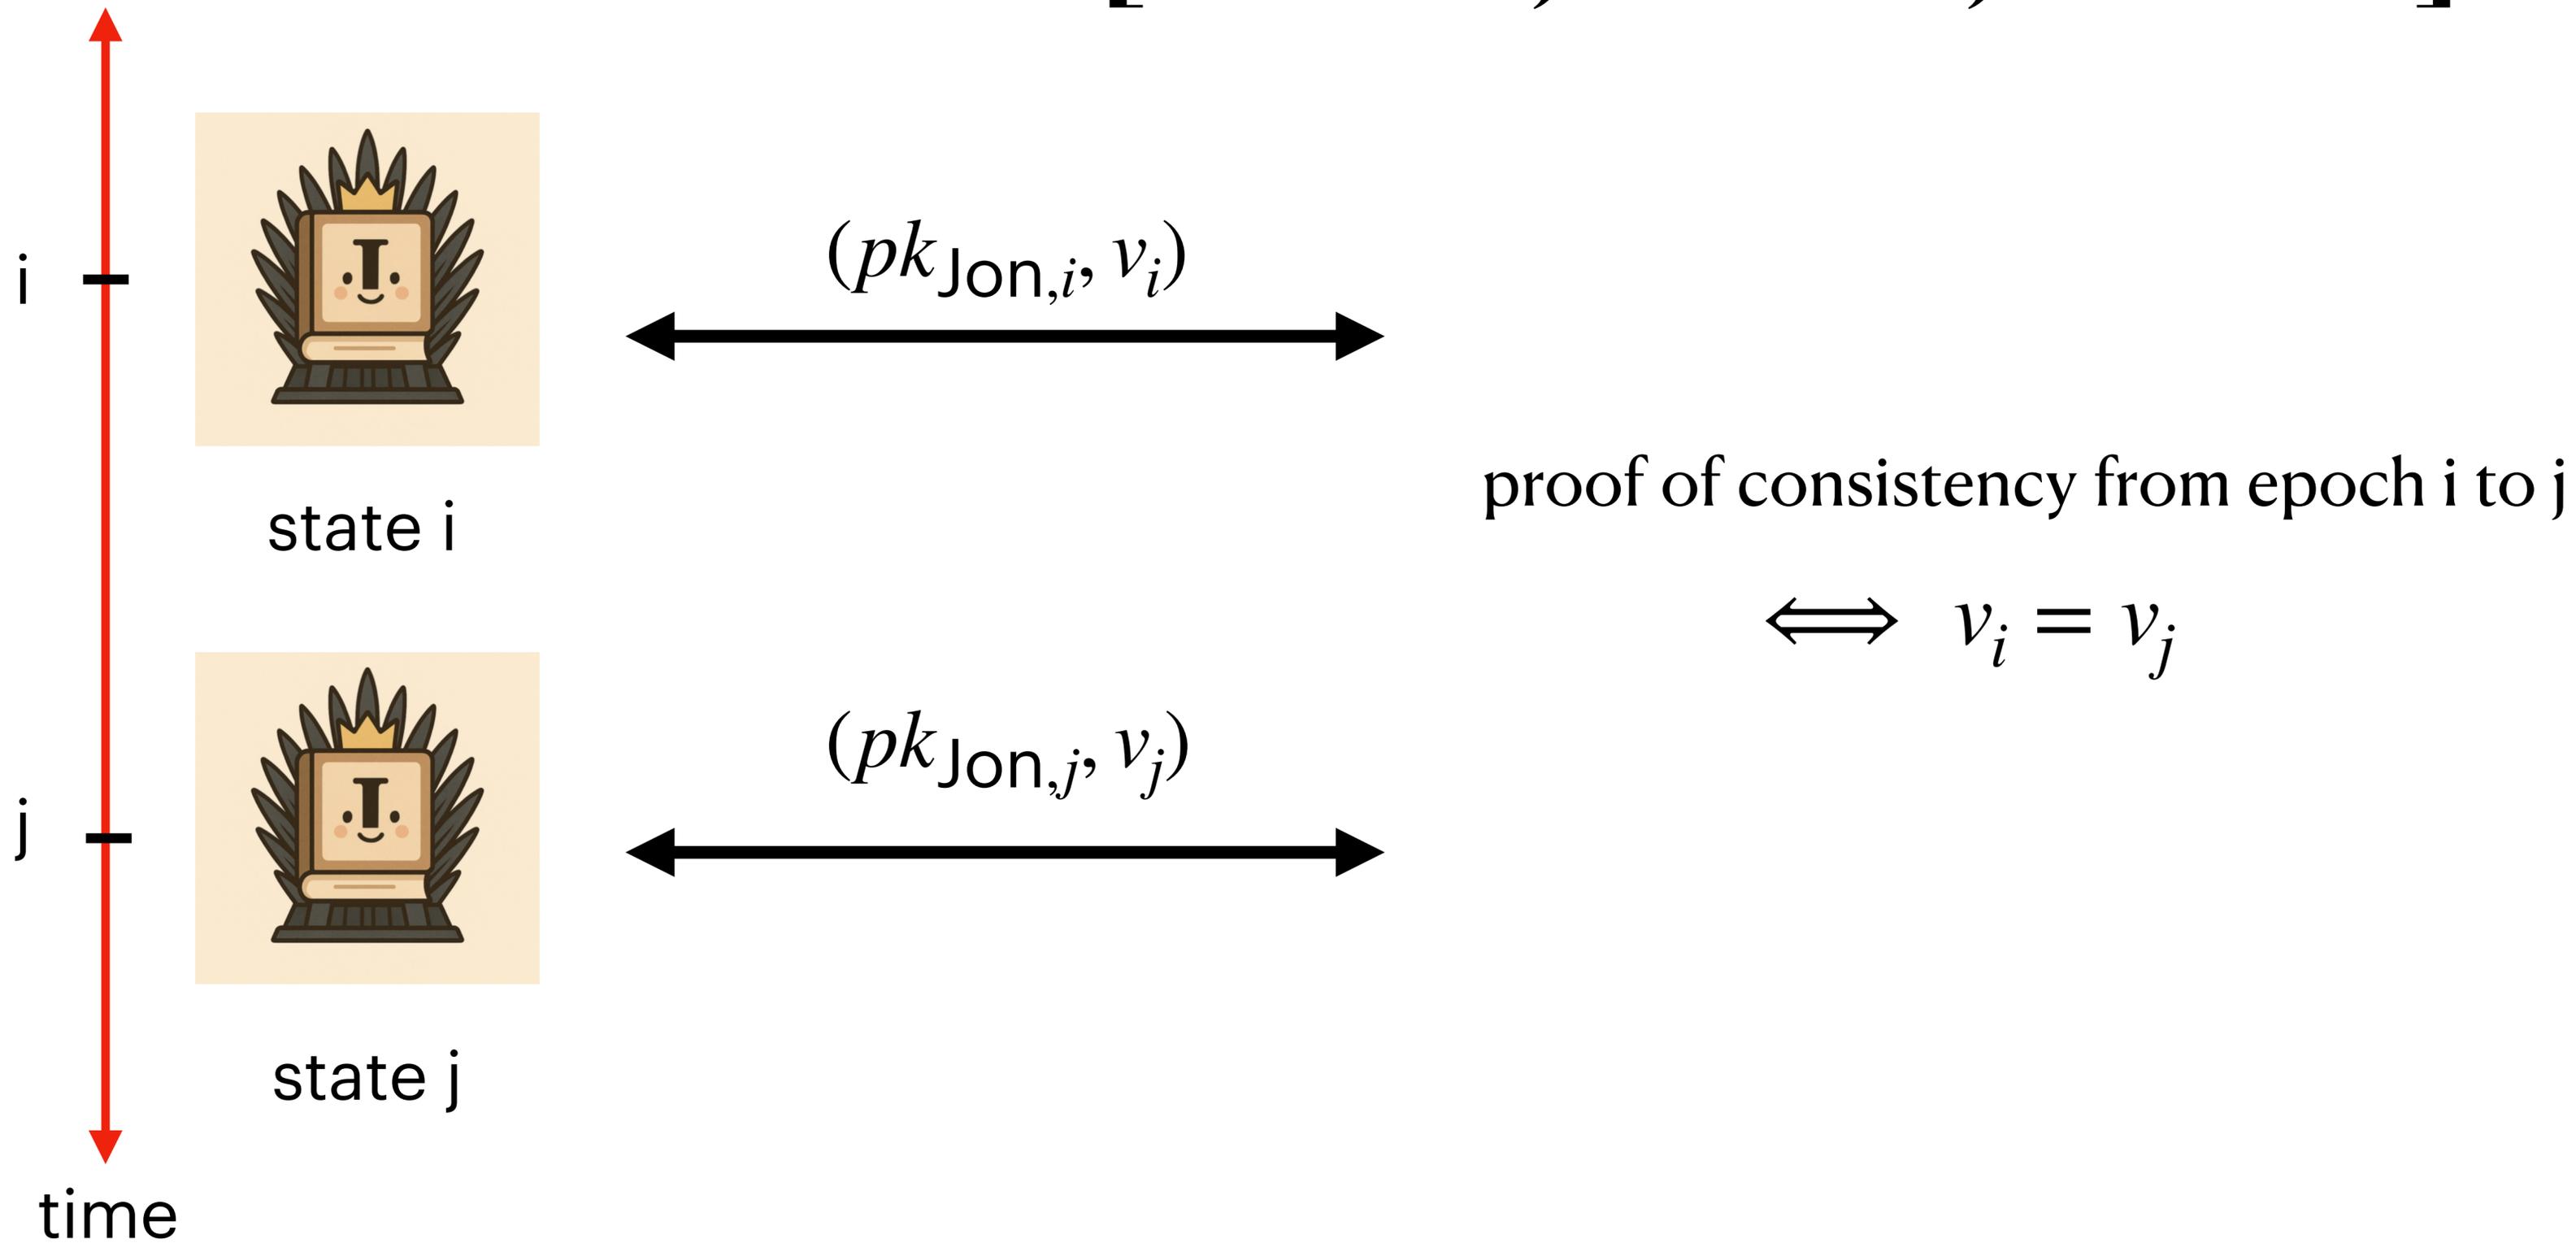
Proof of Consistency: Naive Solution [Mal+15]



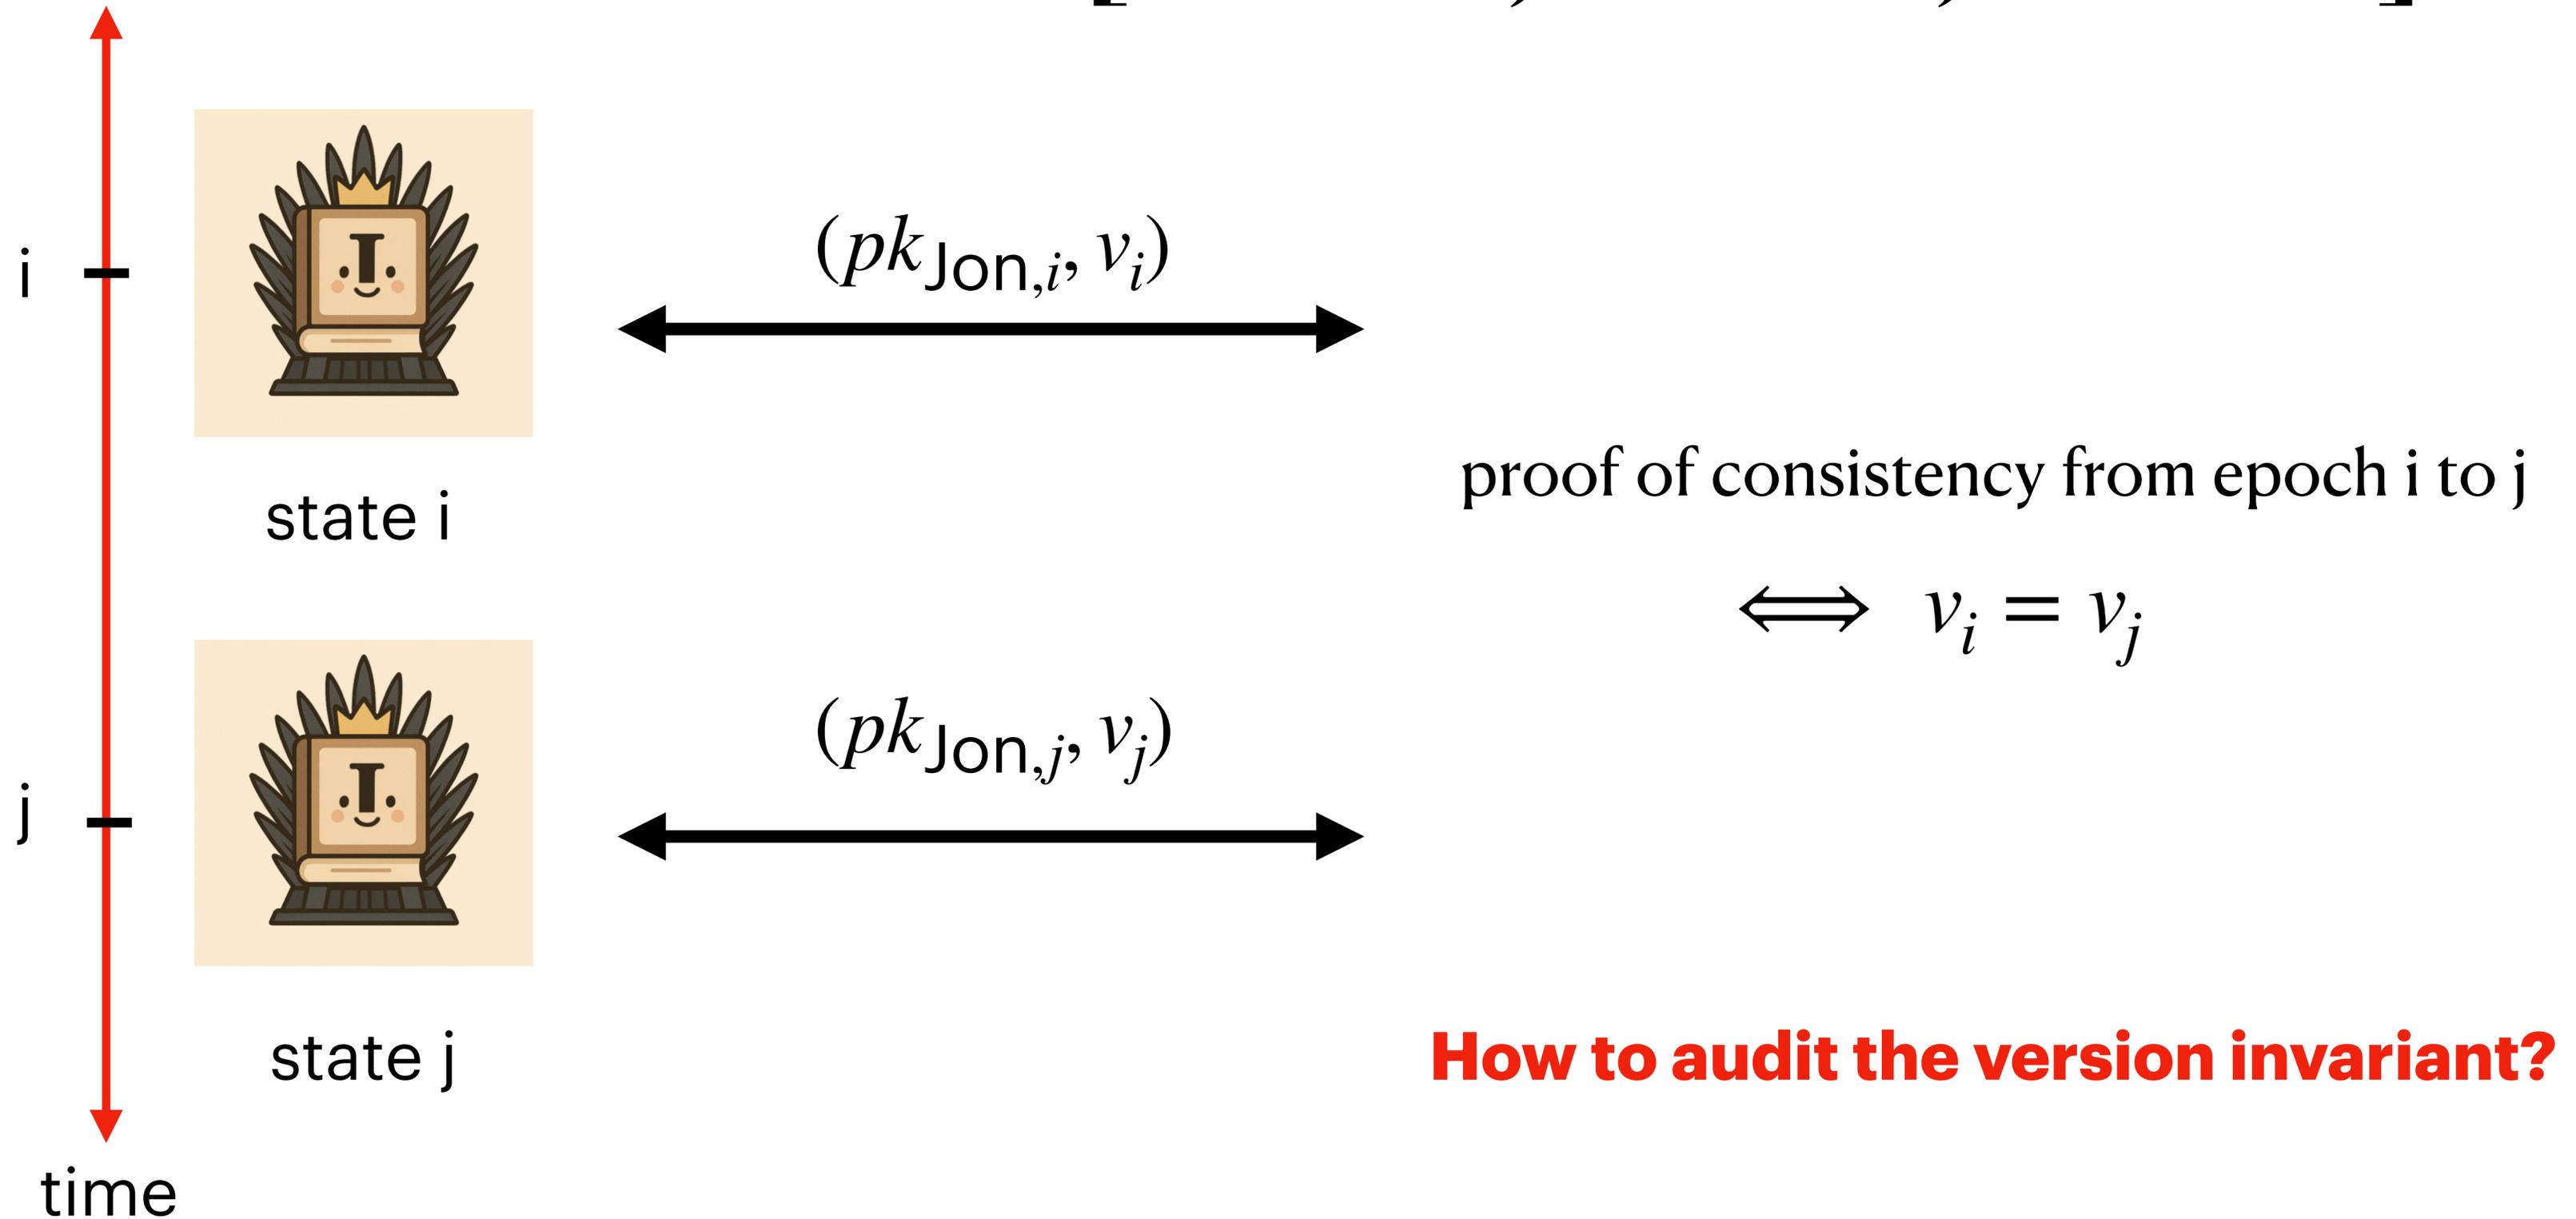
Version Invariant [Cha+19, Mal+23, Len+24]



Version Invariant [Cha+19, Mal+23, Len+24]

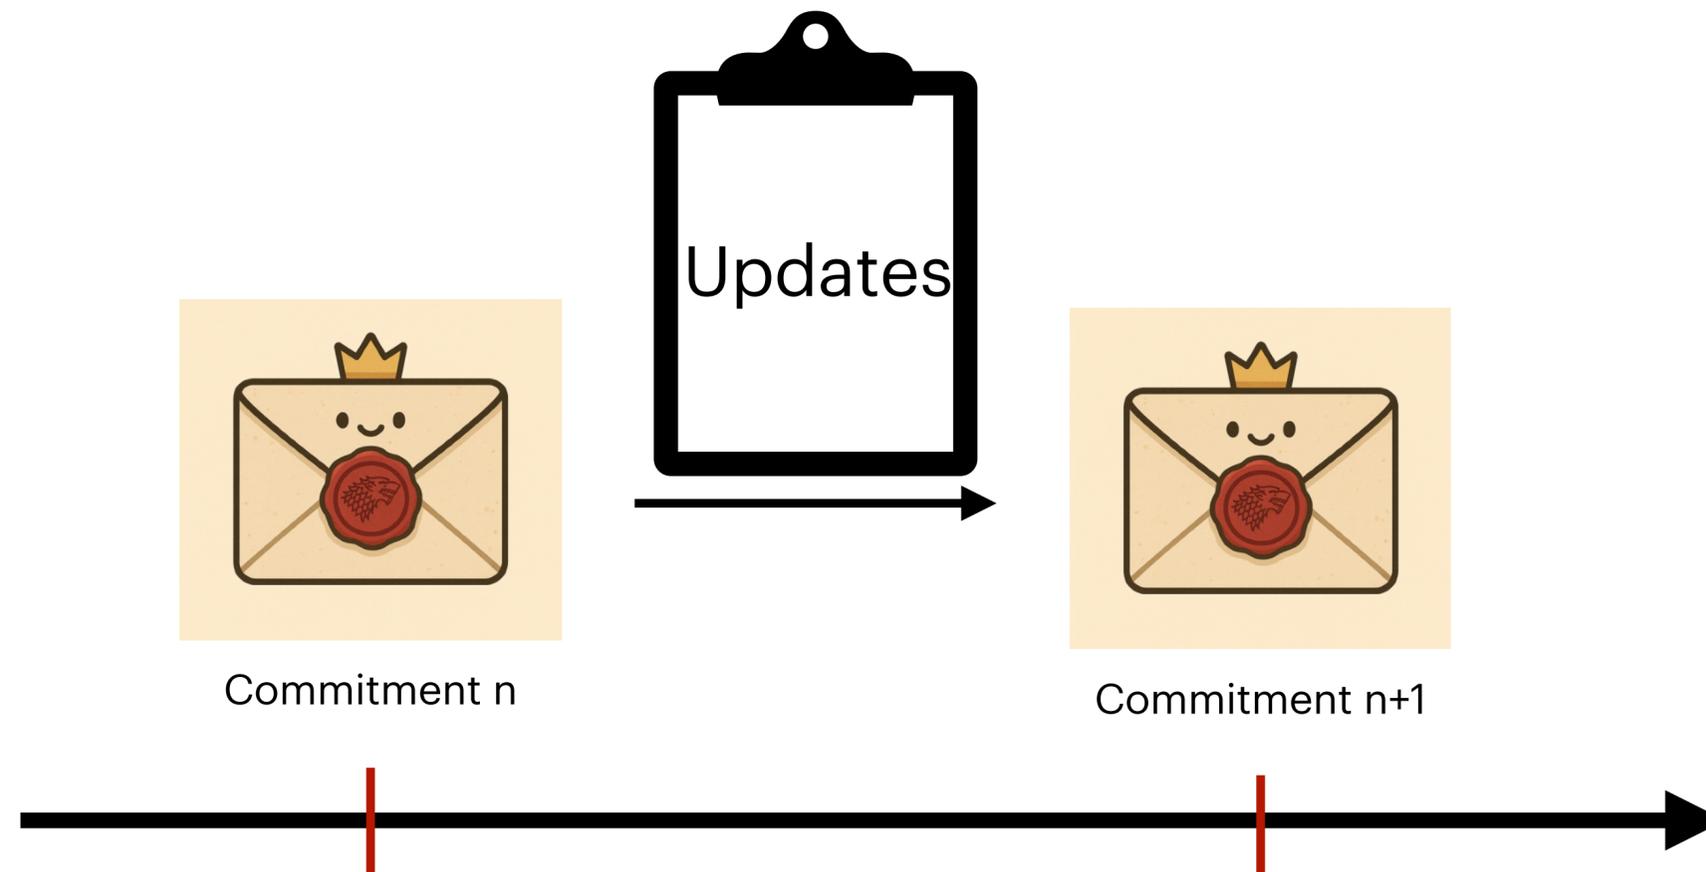


Version Invariant [Cha+19, Mal+23, Len+24]



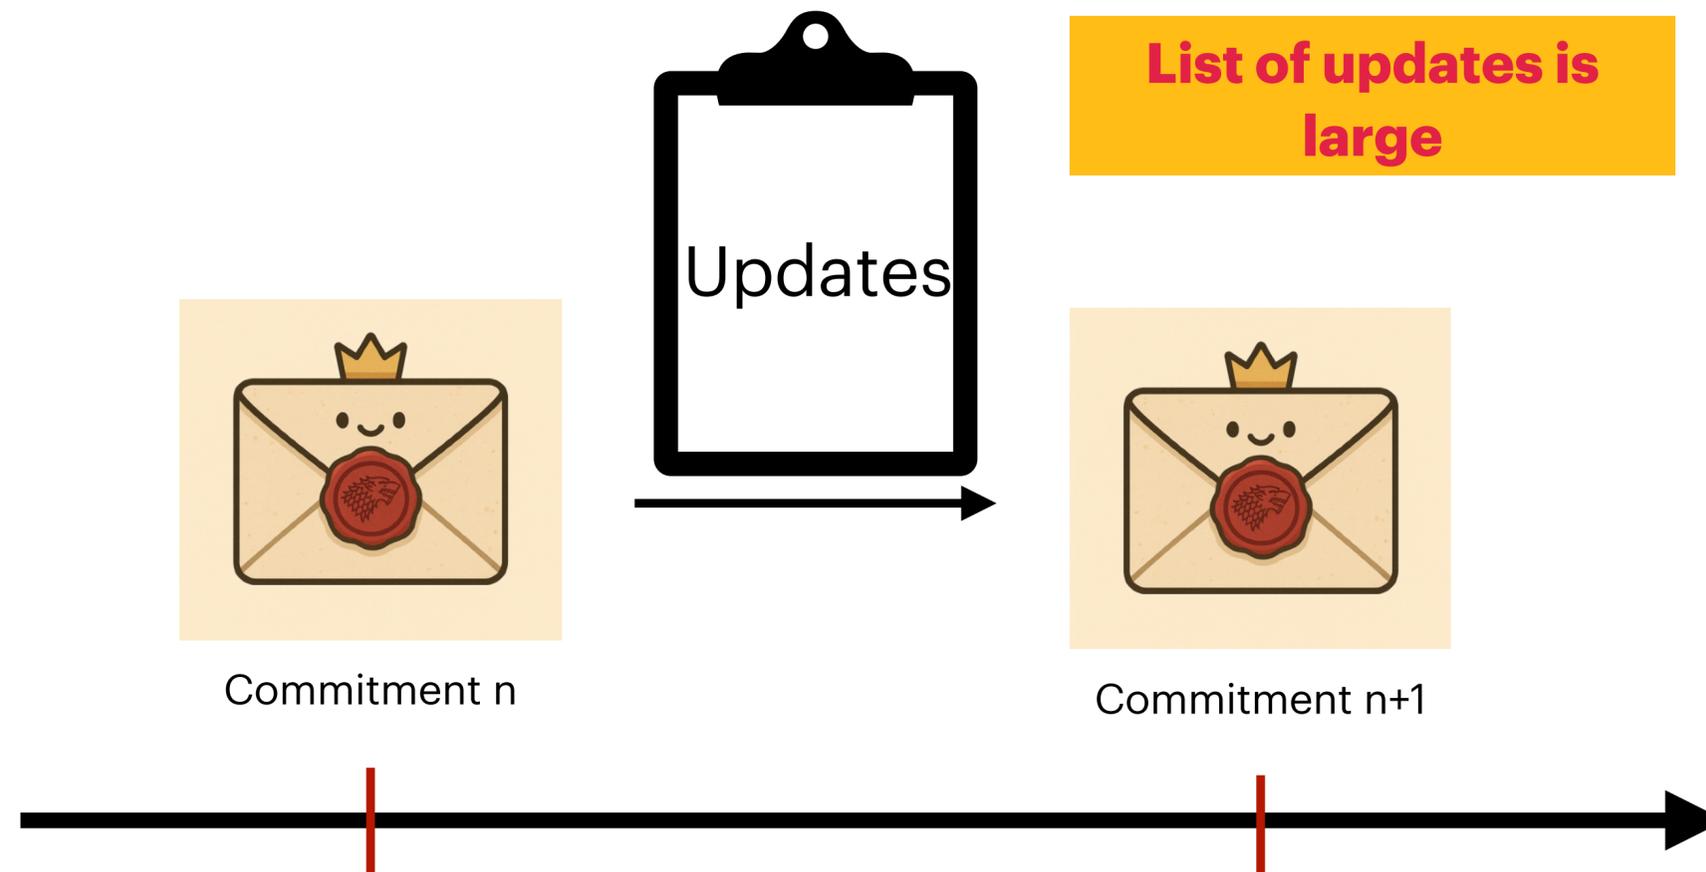
Auditing Version Invariant [Cha+19, Mal+23, Len+24]

- Server publishes invariance proof == the list of updates.
- “Auditors” verify the updates have been appended correctly.



Auditing Version Invariant [Cha+19, Mal+23, Len+24]

- Server publishes invariance proof == the list of updates.
- “Auditors” verify the updates have been appended correctly.





Drawbacks of Existing Schemes:

What Are Our Main Challenges?

SEEMless [Cha+19], Parakeet [Mal+23], Optiks [Len+24]

- Self-Auditability
- Perfect Privacy
- History Pruning

Challenge I: Self-Auditability

- Auditor's proof is large \implies linear in the number of updates in the epoch.

Challenge I: Self-Auditability

- Auditor's proof is large \implies linear in the number of updates in the epoch.
- = hashing 80MB per minutes for WhatsApp KT.

Challenge I: Self-Auditability

- Auditor's proof is large \implies linear in the number of updates in the epoch.
- = hashing 80MB per minutes for WhatsApp KT.
- Auditing delagated to Cloudflare

Challenge I: Self-Auditability

- Auditor's proof is large \implies linear in the number of updates in the epoch.
- = hashing 80MB per minutes for WhatsApp KT.
- Auditing delegated to Cloudflare

Attempts to make the proof succinct using SNARKs (e.g., Verdict [Tzi+22] and Hekaton [Ros+24]) \implies high hash-function costs in circuits limit scalability.

Challenge II: Perfect Privacy

- Privacy is crucial in the context of key transparency:
 - E.g. lookups shouldn't reveal pattern of changing keys of users.
 - Otherwise an adversary observing a user doesn't change their key regularly, makes it a potential target for them.

Challenge II: Perfect Privacy

- Privacy is crucial in the context of key transparency:
 - E.g. lookups shouldn't reveal pattern of changing keys of users.
 - Otherwise an adversary observing a user doesn't change their key regularly, makes it a potential target for them.
- Tree-based solutions do not offer perfect privacy.
 - e.g., SEEMless reveal the last time a key was updated during a lookup.

Challenge III: History Pruning

- The tree-based models (except for CONIKS) are based on append-only trees.

Challenge III: History Pruning

- The tree-based models (except for CONIKS) are based on append-only trees.
- All unnecessary historical data are kept in the tree and its size grows indefinitely.

Challenge III: History Pruning

- The tree-based models (except for CONIKS) are based on append-only trees.
- All unnecessary historical data are kept in the tree and its size grows indefinitely.
- WhatsApp has a user base of 3-4 billion, but its tree size is more than 100 billion.



Our Proposed Schemes: IronDict and Aegon



Optiks

Parakeet

SEEMless

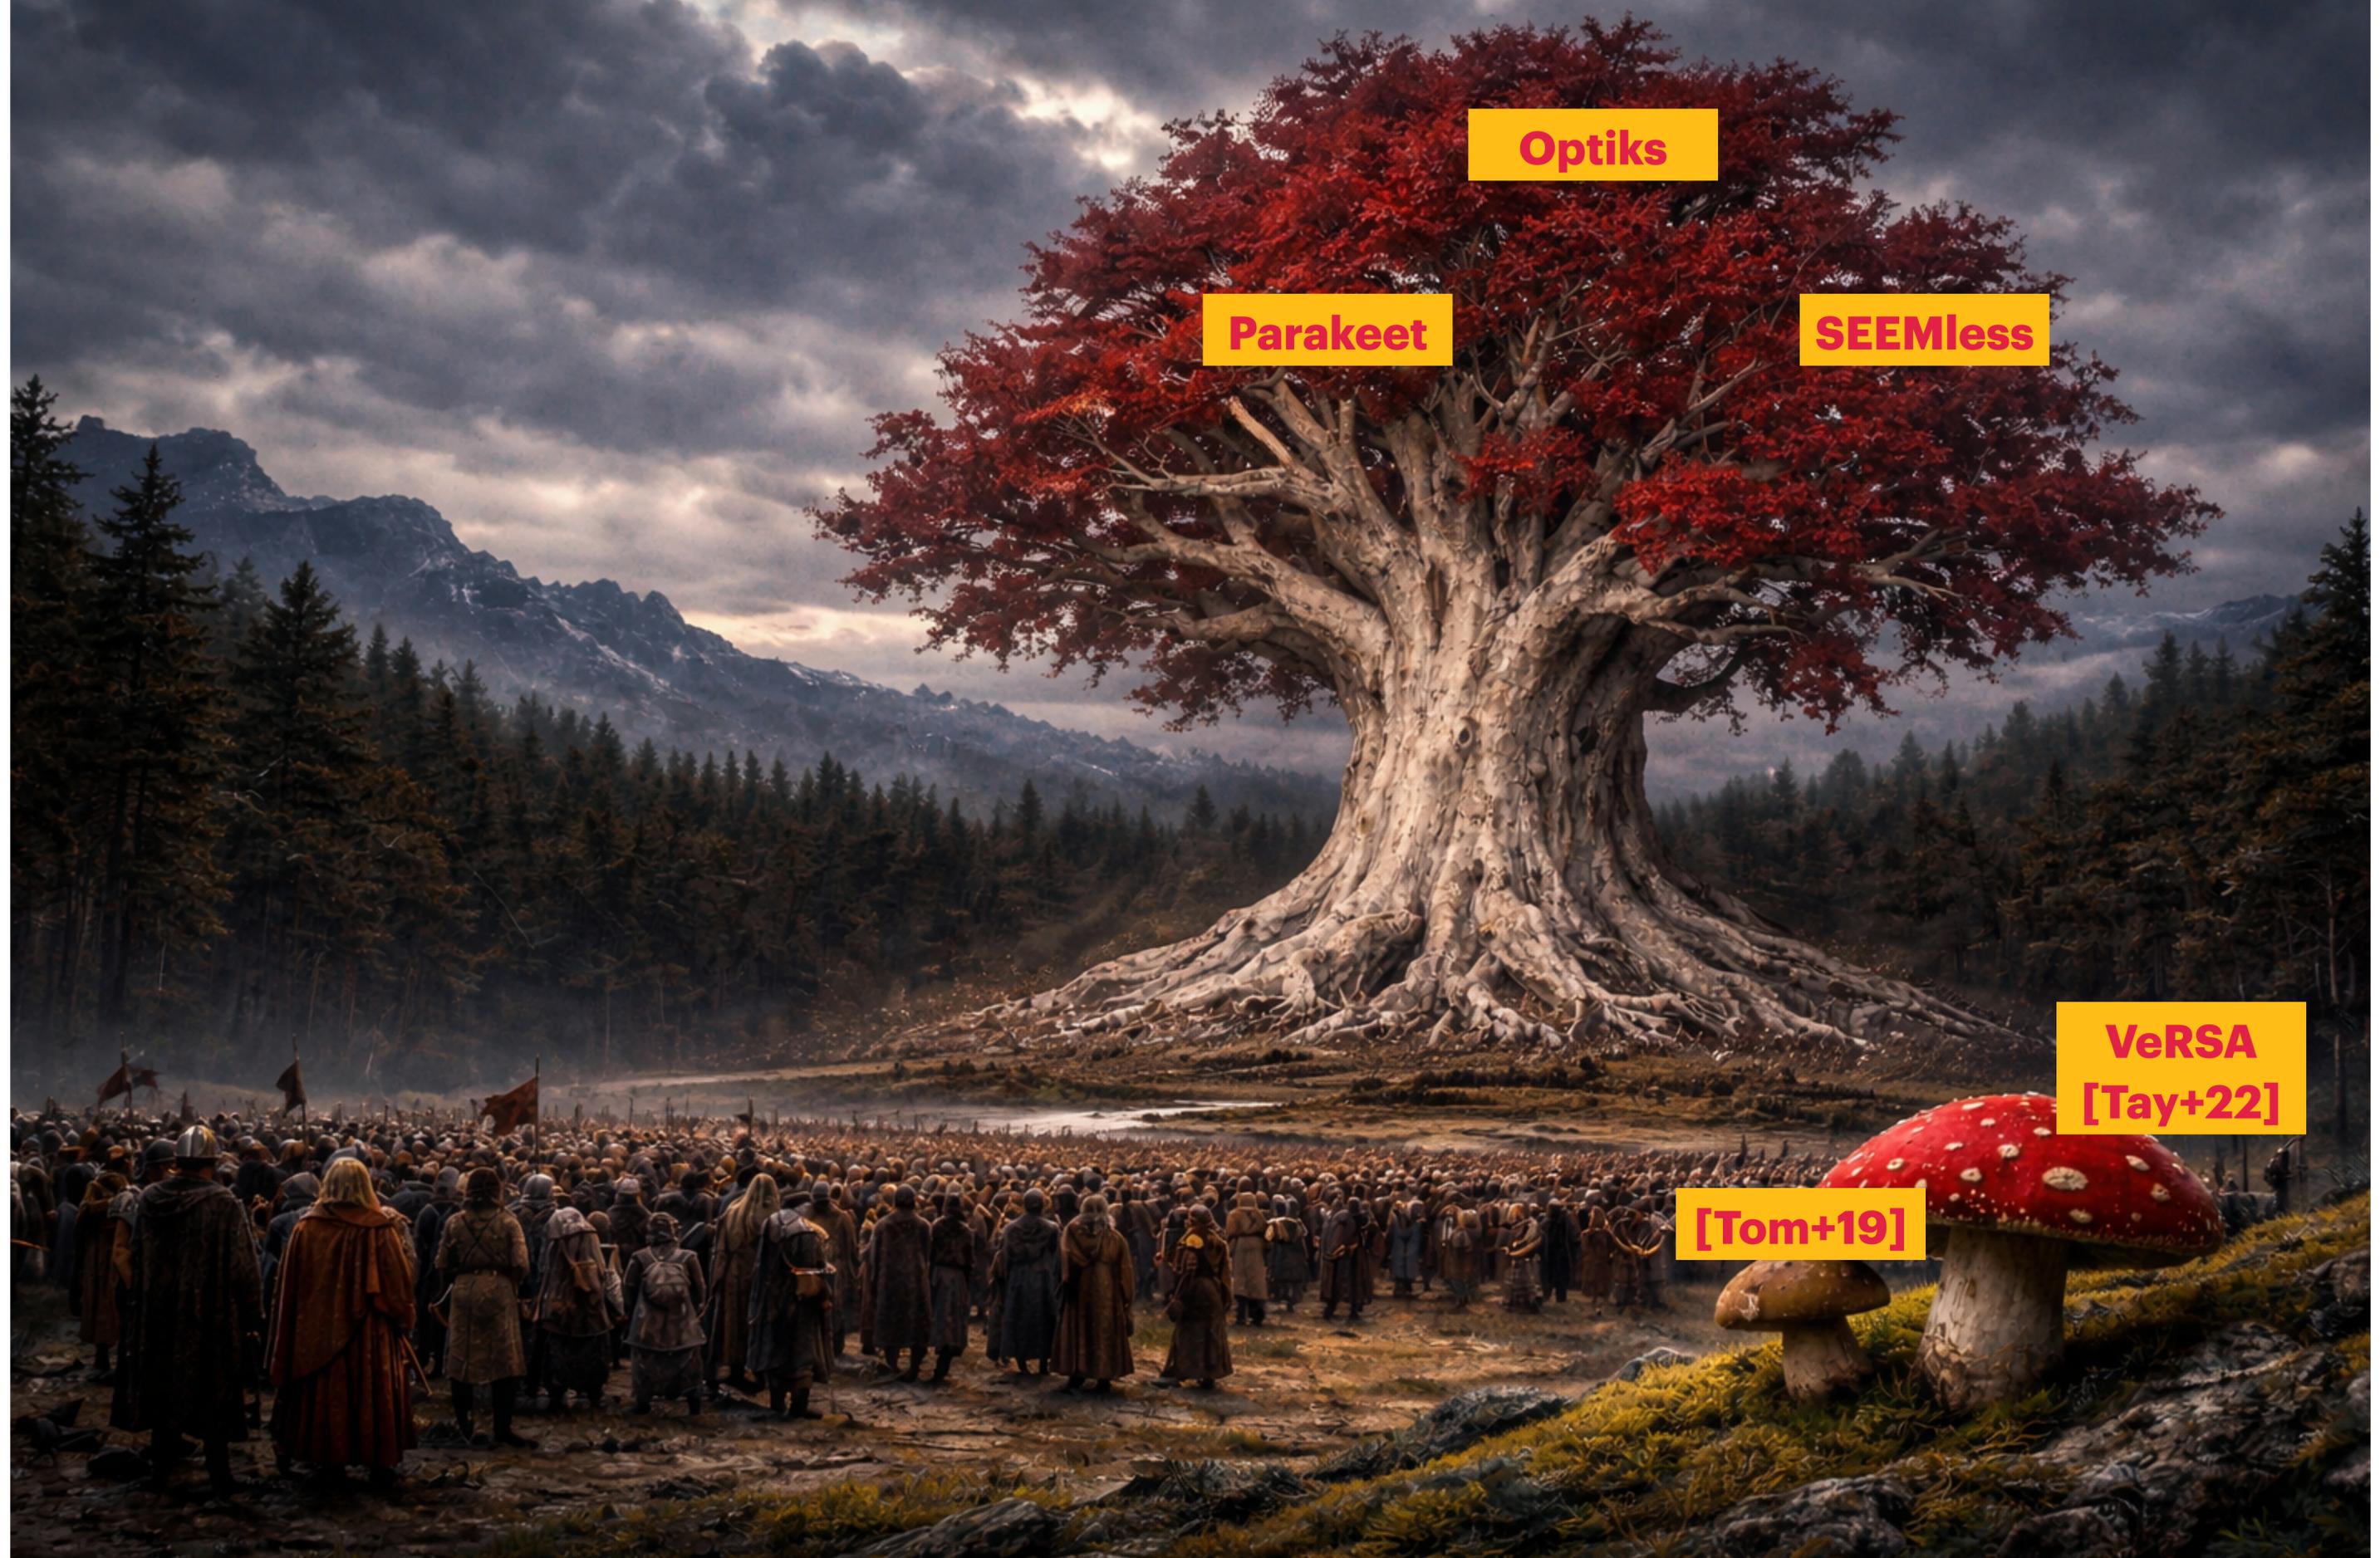
Optiks

Parakeet

SEEMless

VeRSA
[Tay+22]

[Tom+19]



Optiks

Parakeet

SEEMless

**VeRSA
[Tay+22]**

[Tom+19]



**IronDict / Aegon
based on polynomials**

Optiks

Parakeet

SEEMless

**VeRSA
[Tay+22]**

[Tom+19]



IRONDICT:

Transparent Dictionaries from Polynomial Commitments

Hossein Hafezi^{†1}, Alireza Shirzad^{†2}, Benedikt Bünz³, and Joseph Bonneau⁴

^{1,3,4}New York University

²University of Pennsylvania

USENIX Security 2026

IRONDICT:

Transparent Dictionaries from Polynomial Commitments

Hossein Hafezi^{†1}, Alireza Shirzad^{†2}, Benedikt Bünz³, and Joseph Bonneau⁴

^{1,3,4}New York University

²University of Pennsylvania

USENIX Security 2026

Aegon: Scalable and Self-Auditable Key Transparency

Hossein Hafezi*

University of Cambridge

Alireza Shirzad*

University of Pennsylvania

Benedikt Bünz

NYU

Kevin Lewi

Meta

Dillon George

Meta

Joseph Bonneau

NYU

IronDict

- **Self-auditability:** constant proof size, 8 kB for dictionary of size 1 billion.

IronDict

- **Self-auditability:** constant proof size, 8 kB for dictionary of size 1 billion.
 - **fast-forwarding:** Check a single audit proof and verify the whole history
 - Self-auditability without fast-forwarding is **pointless**.

IronDict

- **Self-auditability:** constant proof size, 8 kB for dictionary of size 1 billion.
 - **fast-forwarding:** Check a single audit proof and verify the whole history
 - Self-auditability without fast-forwarding is **pointless**.
- **Perfect-privacy**



AEGON

IRONDICT

High Epoch Latency

Requirement of a Linear-Sized Setup

Not Horizontally Scalable

Storage Inefficiency

Random Version Invariant of IronDict and Aegon

$$\text{Dict}_i = (pk_1^{(i)}, pk_2^{(i)}, \dots, pk_N^{(i)}),$$

Random Version Invariant of IronDict and Aegon

$$\text{Dict}_i = (pk_1^{(i)}, pk_2^{(i)}, \dots, pk_N^{(i)}), \quad \Delta_i = \text{Dict}_{i+1} - \text{Dict}_i$$

Random Version Invariant of IronDict and Aegon

$$\text{Dict}_i = (pk_1^{(i)}, pk_2^{(i)}, \dots, pk_N^{(i)}), \quad \Delta_i = \text{Dict}_{i+1} - \text{Dict}_i$$

Proof of consistency for index t: $\text{Dict}_i[t] = \text{Dict}_i[t] = \dots = \text{Dict}_j[t]$

Random Version Invariant of IronDict and Aegon

$$\text{Dict}_i = (pk_1^{(i)}, pk_2^{(i)}, \dots, pk_N^{(i)}), \quad \Delta_i = \text{Dict}_{i+1} - \text{Dict}_i$$

Proof of consistency for index t: $\text{Dict}_i[t] = \text{Dict}_{i+1}[t] = \dots = \text{Dict}_j[t]$

$$\iff \Delta_{i+1}[t] = \Delta_{i+2}[t] = \dots = \Delta_j[t] = 0$$

Random Version Invariant of IronDict and Aegon

$$\text{Dict}_i = (pk_1^{(i)}, pk_2^{(i)}, \dots, pk_N^{(i)}), \quad \Delta_i = \text{Dict}_{i+1} - \text{Dict}_i$$

Proof of consistency for index t : $\text{Dict}_i[t] = \text{Dict}_i[t] = \dots = \text{Dict}_j[t]$

$$\iff \Delta_{i+1}[t] = \Delta_{i+2}[t] = \dots = \Delta_j[t] = 0$$

Ghost key attack : $\sum_{i < k \leq j} \Delta_k[t] = 0$ and $\exists i < i_0 \leq j : \Delta_{i_0}[t] \neq 0$

Random Version Invariant of IronDict and Aegon

$$\text{Dict}_i = (pk_1^{(i)}, pk_2^{(i)}, \dots, pk_N^{(i)}), \quad \Delta_i = \text{Dict}_{i+1} - \text{Dict}_i$$

Proof of consistency for index t : $\text{Dict}_i[t] = \text{Dict}_i[t] = \dots = \text{Dict}_j[t]$

$$\iff \Delta_{i+1}[t] = \Delta_{i+2}[t] = \dots = \Delta_j[t] = 0$$

Ghost key attack : $\sum_{i < k \leq j} \Delta_k[t] = 0$ and $\exists i < i_0 \leq j : \Delta_{i_0}[t] \neq 0$

Use randomness so the differences do not cancel out each other

Random Version Invariant of IronDict and Aegon

Rand vector : $\text{Rand}_n = \text{Dict}_0 + \sum_{i \leq n} r_i \cdot \Delta_i \equiv \text{Rand}_{n+1} = \text{Rand}_n + r_{n+1} \cdot \Delta_{n+1}$

Random Version Invariant of IronDict and Aegon

$$\text{Rand vector : } \text{Rand}_n = \text{Dict}_0 + \sum_{i \leq n} r_i \cdot \Delta_i \equiv \text{Rand}_{n+1} = \text{Rand}_n + r_{n+1} \cdot \Delta_{n+1}$$

proof of consistency from epoch i to epoch j $\iff \text{Rand}_i[t] = \text{Rand}_j[t]$

Random Version Invariant of IronDict and Aegon

Rand vector : $\text{Rand}_n = \text{Dict}_0 + \sum_{i \leq n} r_i \cdot \Delta_i \equiv \text{Rand}_{n+1} = \text{Rand}_n + r_{n+1} \cdot \Delta_{n+1}$

proof of consistency from epoch i to epoch $j \iff \text{Rand}_i[t] = \text{Rand}_j[t]$

auditor's task \implies ensure Rand_n is computed correctly

See Numbers!

	Aegon	WhatsApp KT (SEEMLESS + Parakeet)
Epoch Latency	60 s	60 s
Throughput	~1000 update/s	~1000 update/s
Lookup proof / verification	20KB / 29ms	2-4 KB / 0.1 ms
Lookup latency (server)	33ms (more computation)	0.5 s < (no computation)
Audit proof / verification	256B / 2ms	80 MB / 0.7s

Concrete Instantiation?

- Aegon = generic over any polynomial commitment schemes. We initiate it with KZH [Kad+25] that requires a powers-of-tau setup \implies MPC ceremony setup

KZH-Fold:

Accountable Voting from Sublinear Accumulation

George Kadianakis¹, Arantxa Zapico², Hossein Hafezi³, and Benedikt Bünz⁴

^{1,2}Ethereum Foundation

^{3,4}New York University

- **Future Work:** Design alternative polynomial commitments without a trusted setup.

Details of IronDict and Aegon?

- Check out the eprint
 - IronDict: <https://ia.cr/2025/1580>
 - Aegon: Would be online in less than a month!
- Talk on Details of IronDict and Aegon:
 - Check out talk we gave at Microsoft Research (thanks to **Melissa Chase!**)