
+
◦ •

Key Transparency *from theory to practice*

Esha Ghosh

Microsoft Research, Redmond

End-to-End Encrypted Communications

- Signal
- WhatsApp
- Messenger
- iMessage
- Skype
- Teams
- Zoom
- Webex
- ...



Verify Safety Number



24032 97690 11537 73387
06161 13594 43384 64880
19955 70000 97677 43237

To verify the security of your end-to-end encryption with [redacted]
[redacted] compare the numbers above with their device.

✓ You have not verified your safety number with [redacted]

Mark as verified


Verify Security Code

If these numbers are the same on everyone's screen, then this meeting is end-to-end encrypted.

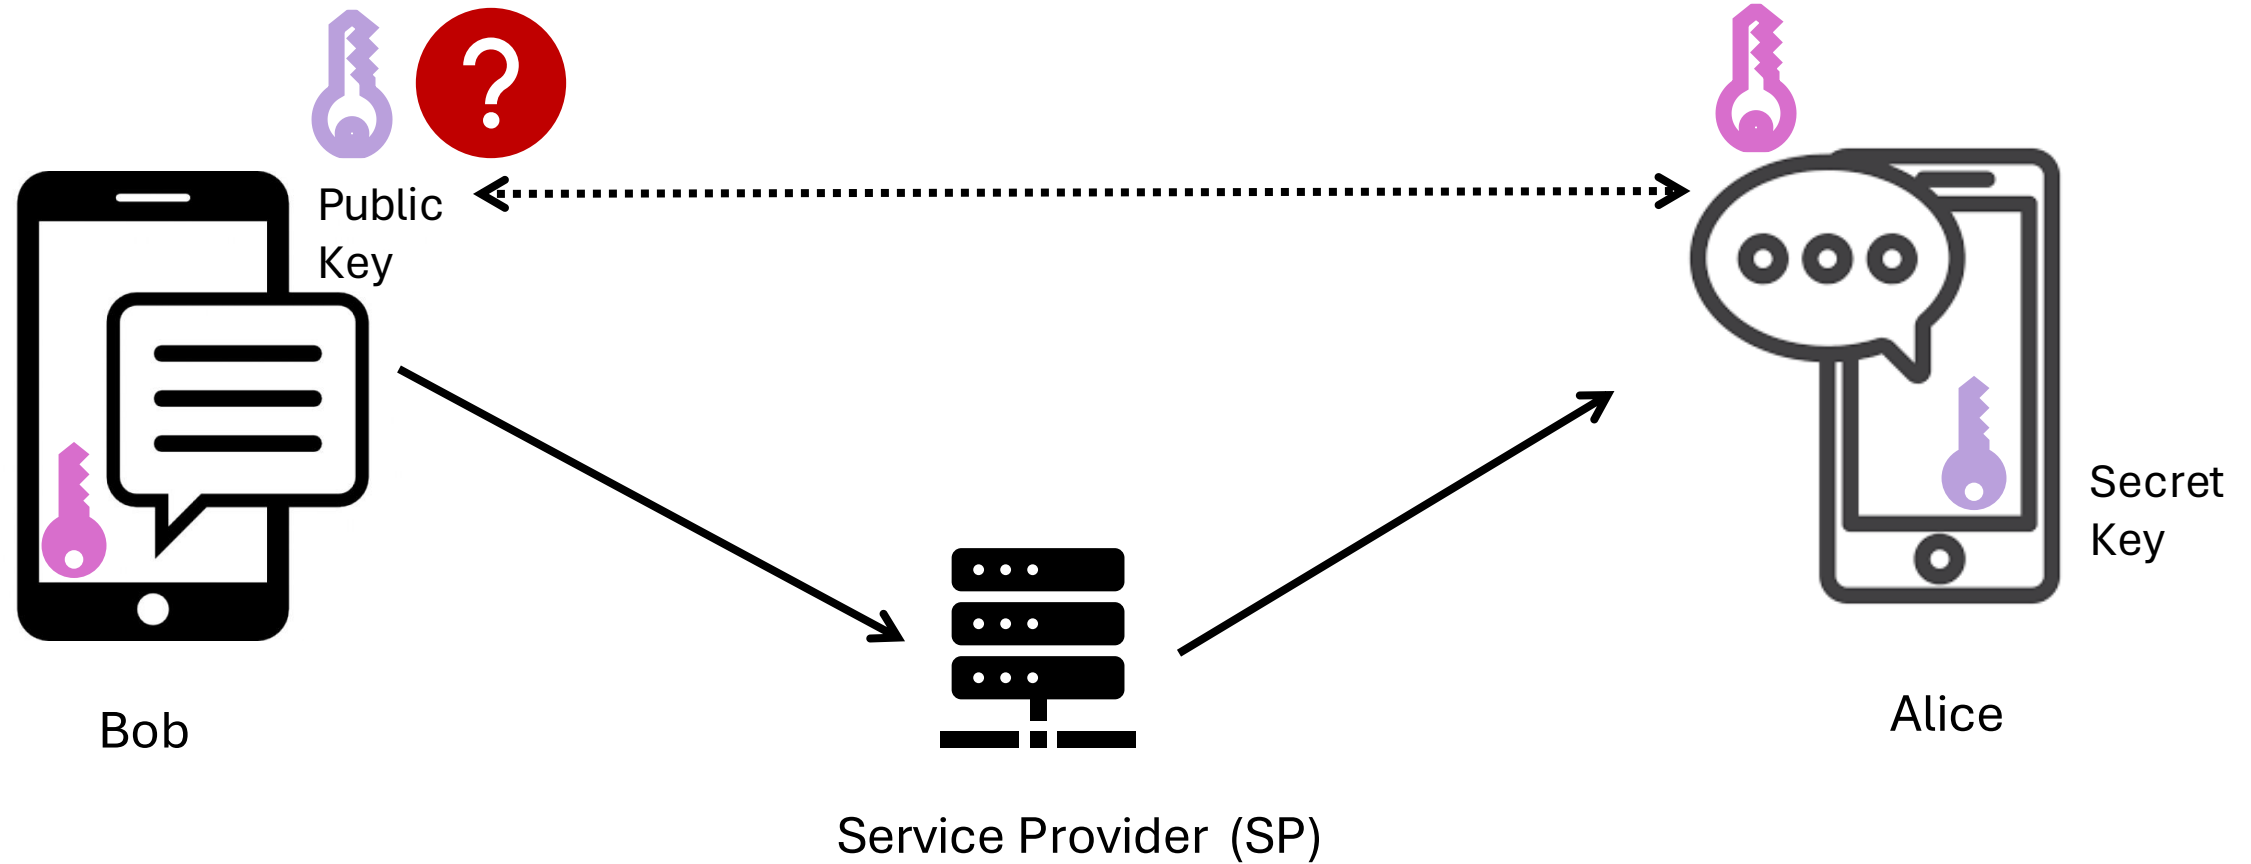
87221 52882 27158 27158

19373 40183 31482 12980

End-to-end encrypted messaging

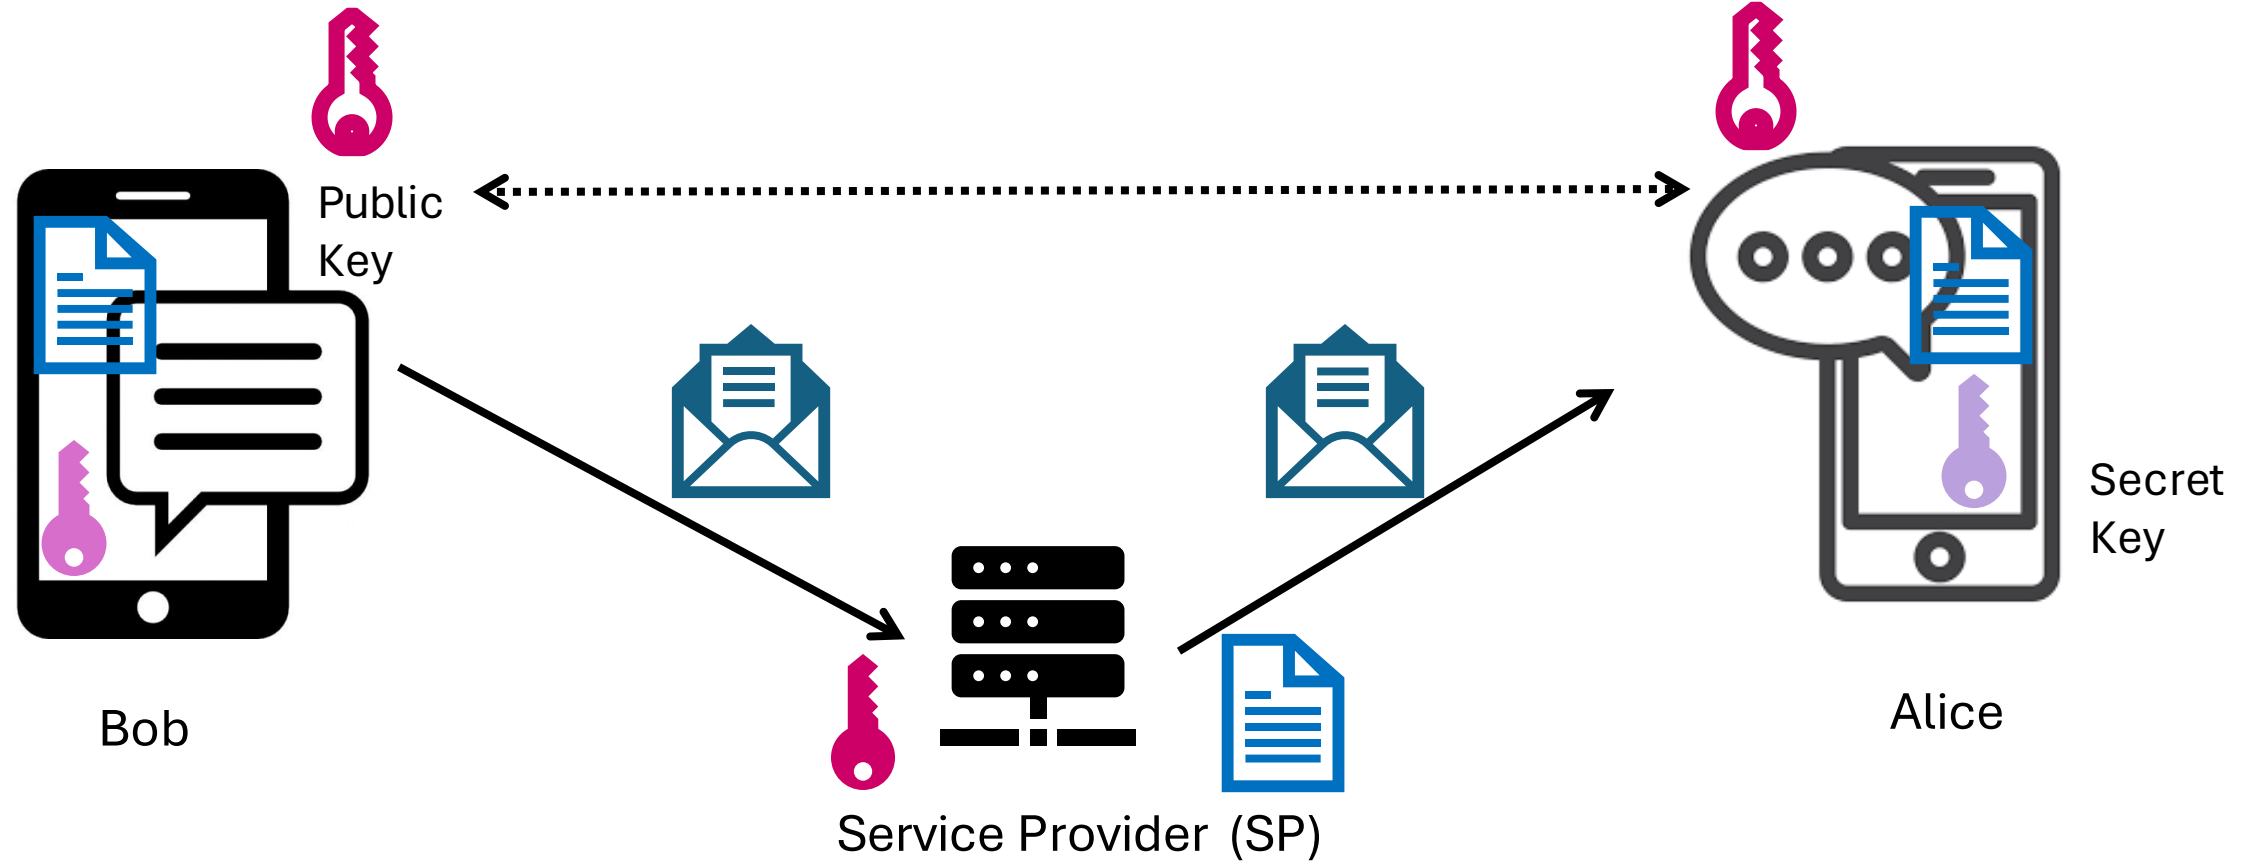


How does Bob get Alice's public key?

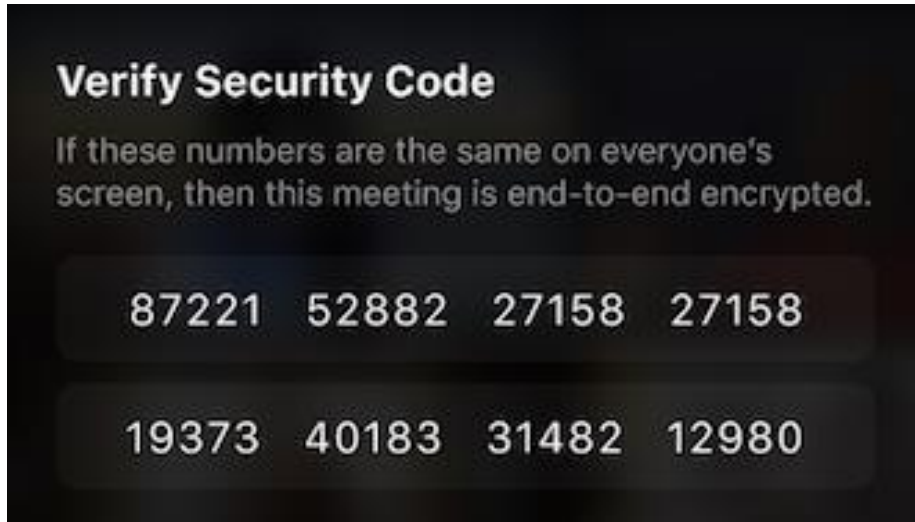


Service Provider manages public key directory and distribution

Undetectable Meddler-in-the-Middle



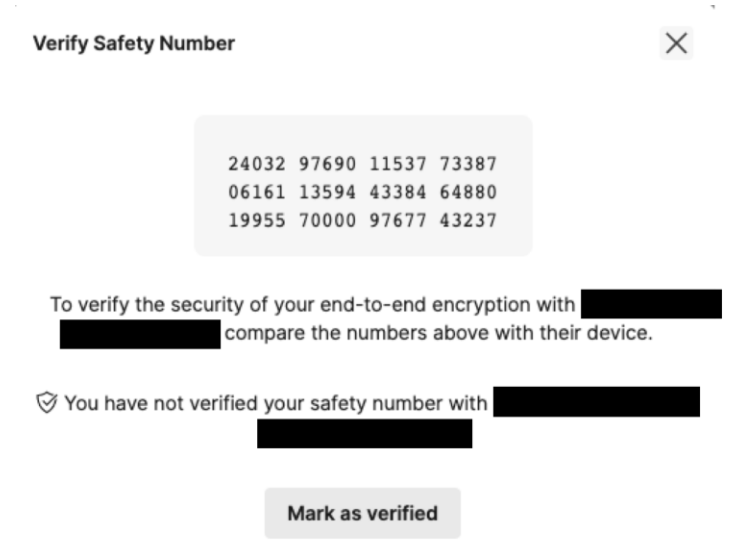
Traditional Approaches



Security Code



QR Code



Safety String

Cryptographic commitments

Commit Phase



Opening Phase




Key Transparency: Automating the checks

- Lifts the burden from users by automating the checks
- Requires the service provider to *commit* to its key directory periodically
- When Bob requests Alice's key, it comes with a proof with respect to the latest commitment
- Alice's client software can monitor her key in the directory:
 - *Regularly requests her own key* (and proof) and checks that it is correct.




Landscape

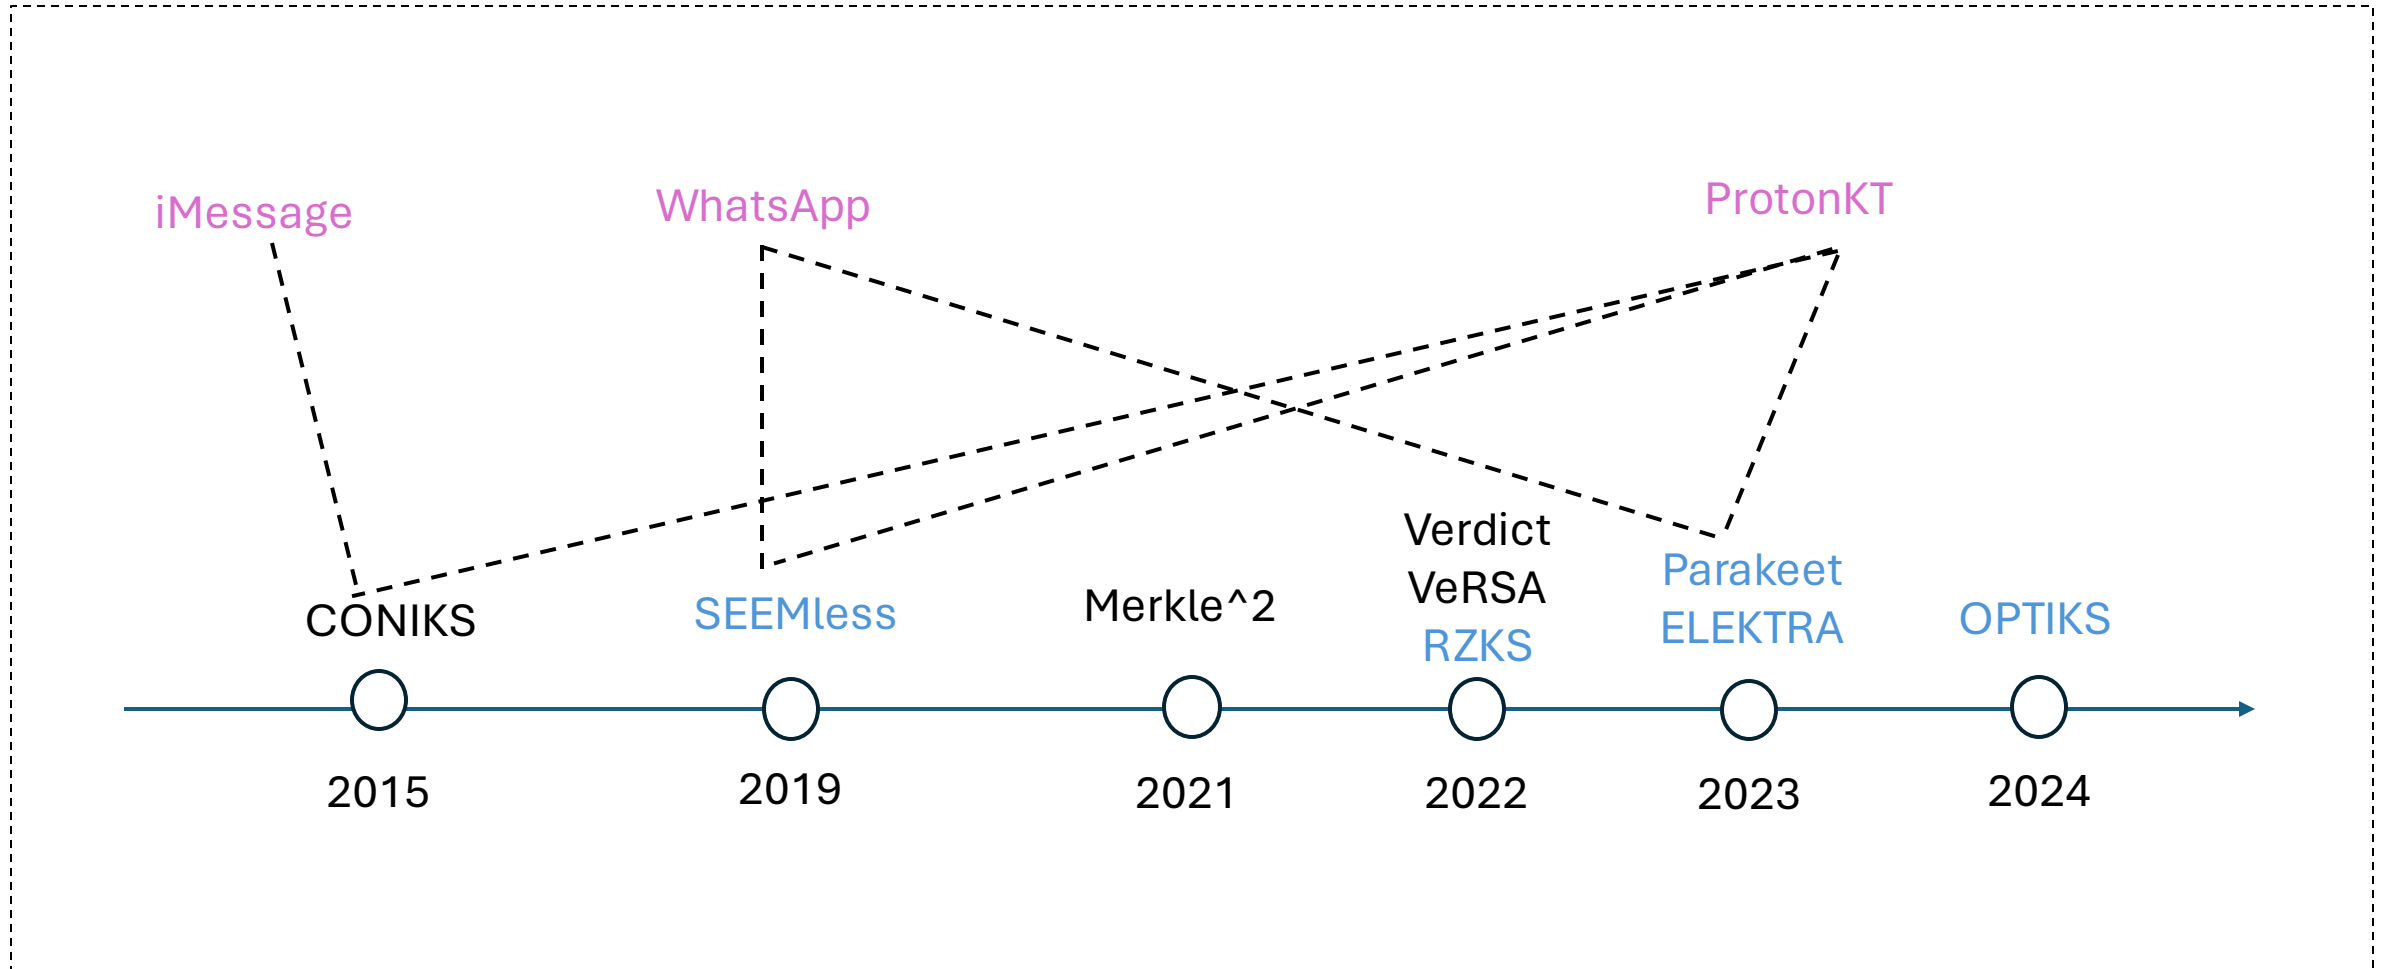
Academic Literature



Industry Deployments and Standardizations



Industry Deployments



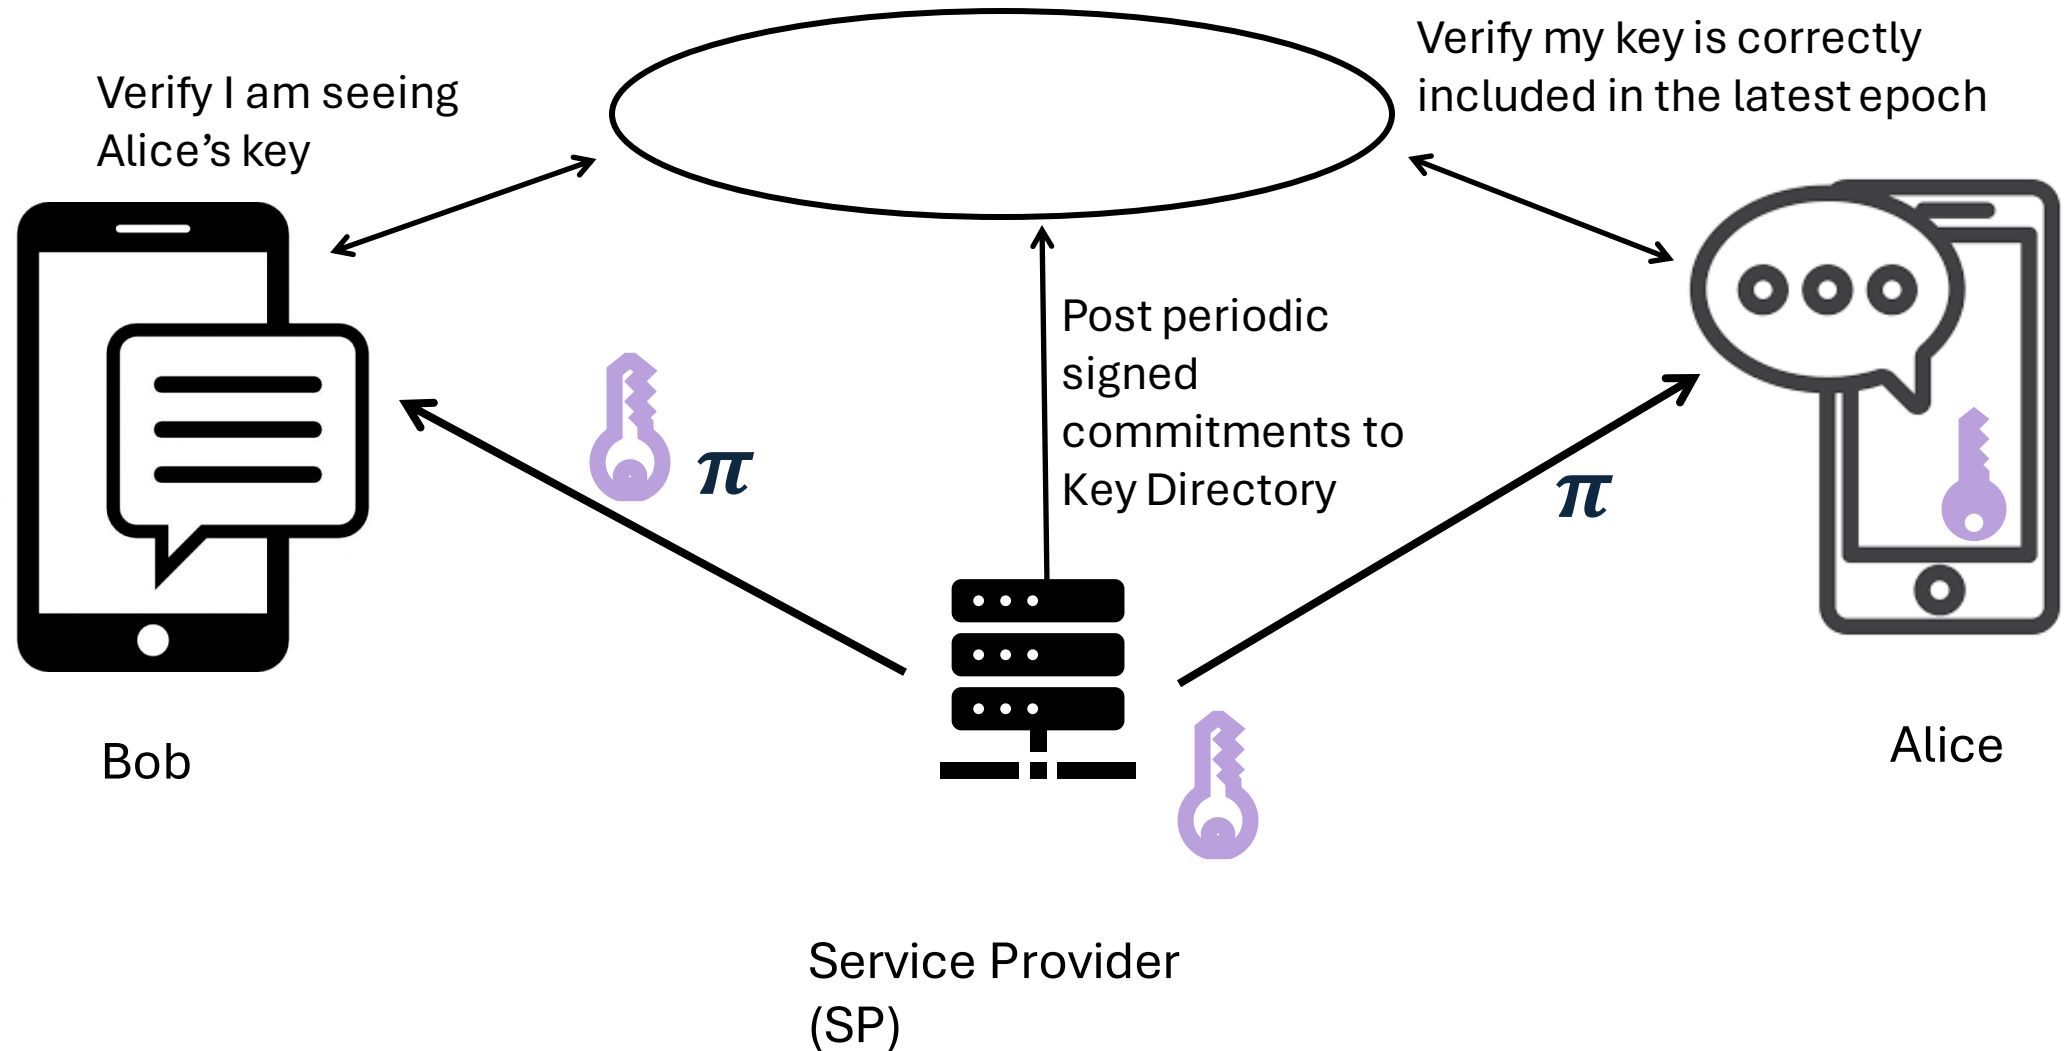
Overview



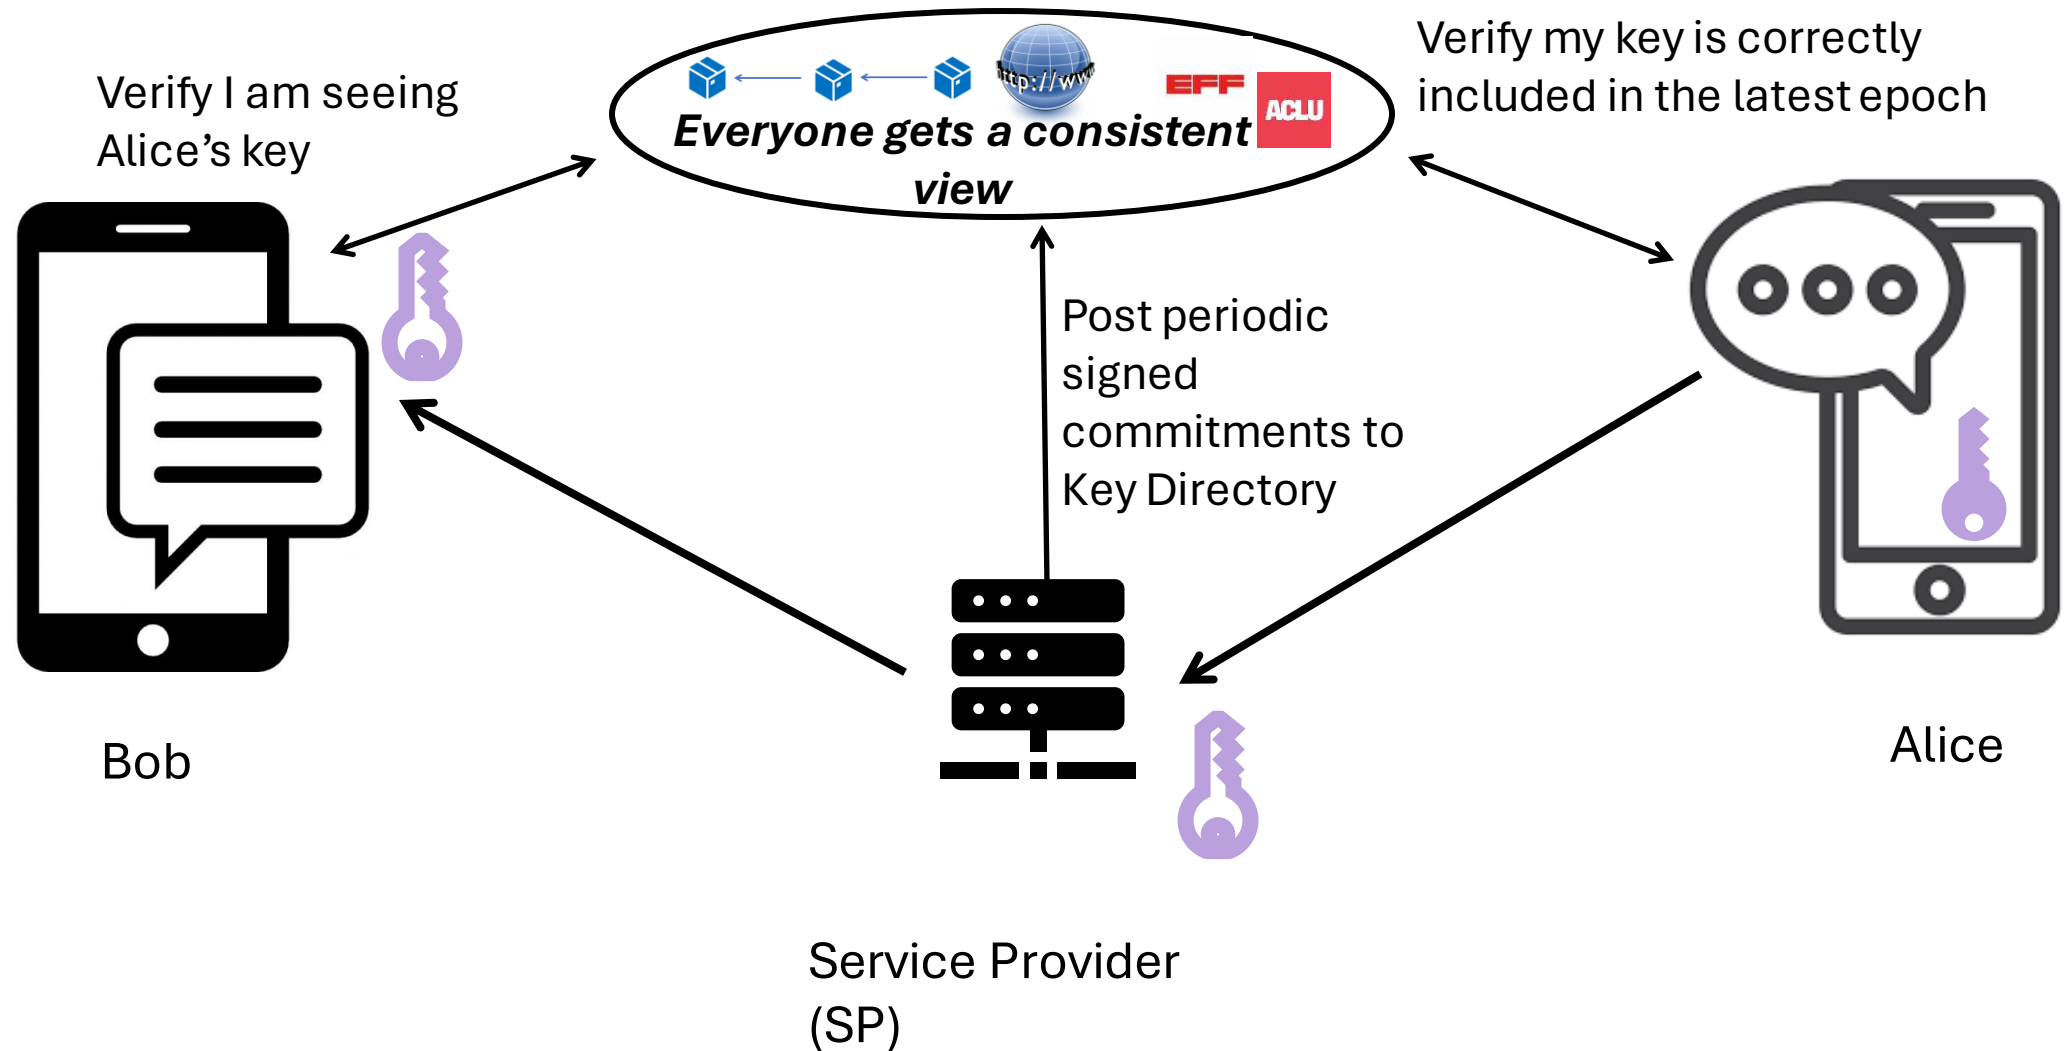
Directory Commitment

- A primitive to generate a **short commitment to a collection of (username, public key) pairs** where the **labels form a set S**
- The Service Provider (SP) generates the short commitment c to the collection
- SP can prove statements about membership of any element $x \in S$ ($x \notin S$) later **non-interactively** with respect to c **without revealing anything else about S**

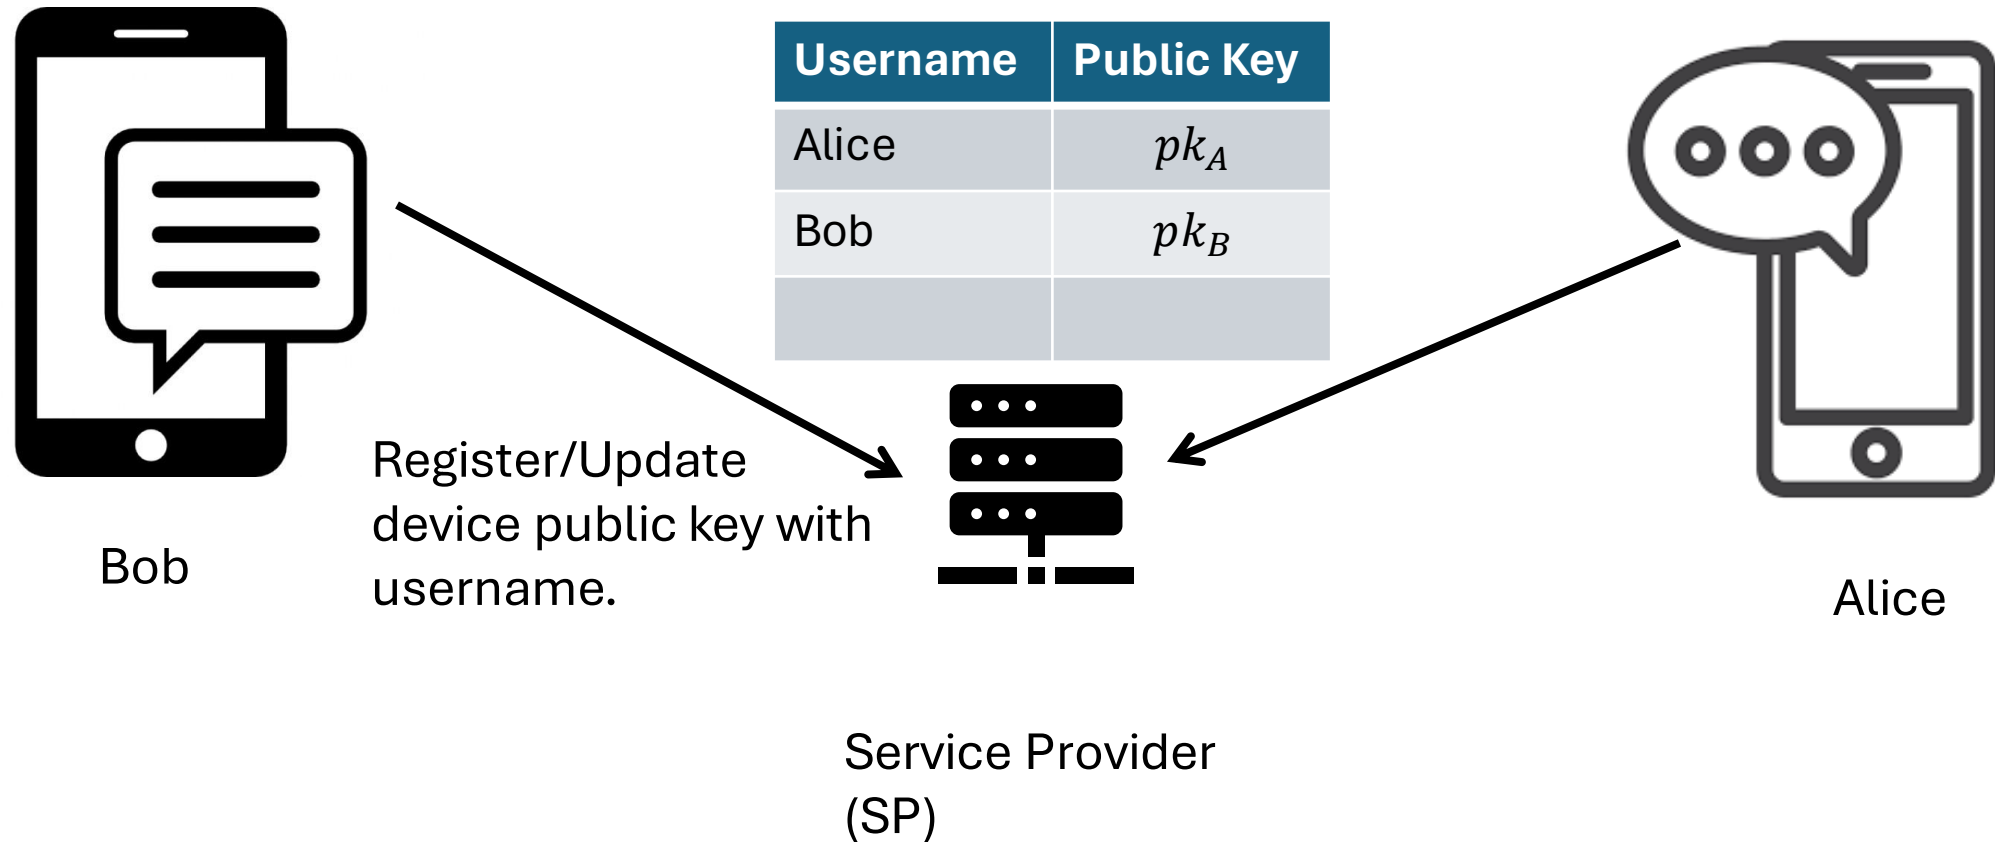
Key Transparency



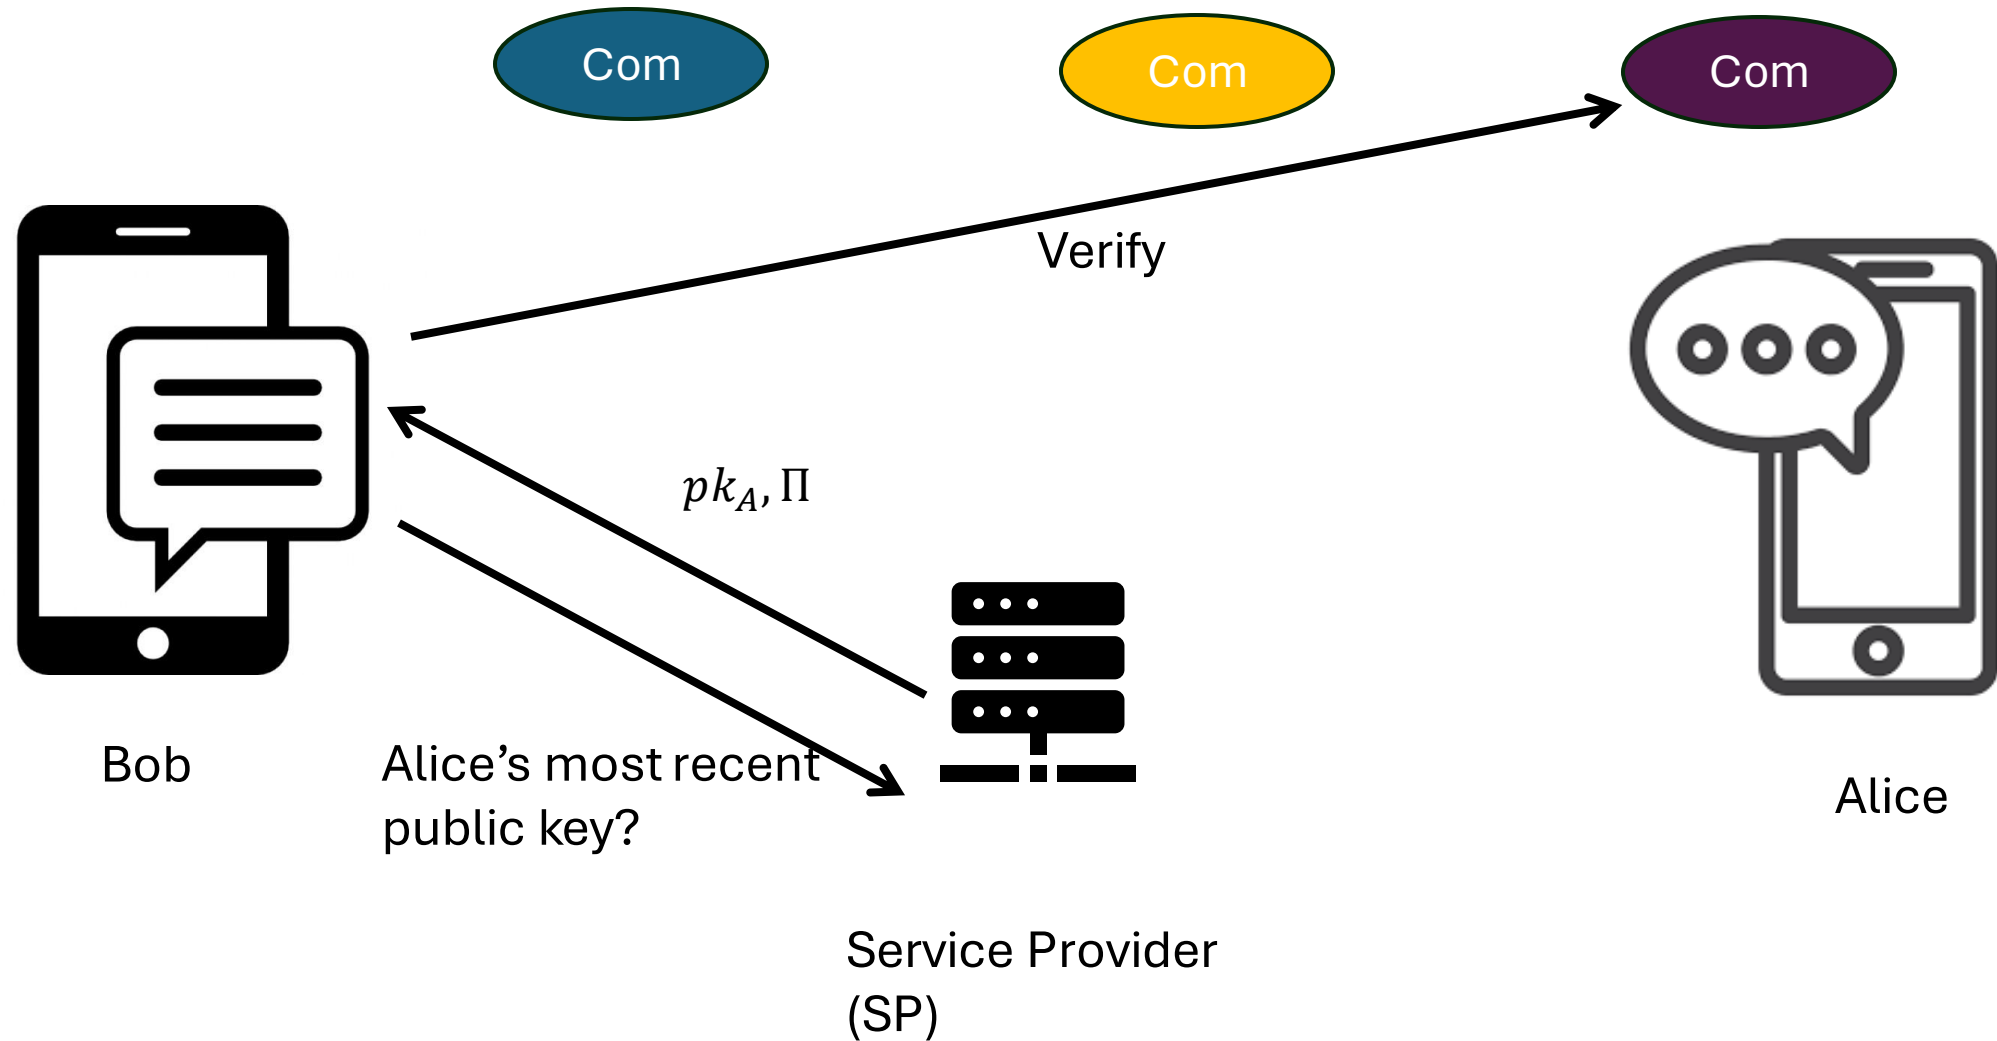
Key Transparency



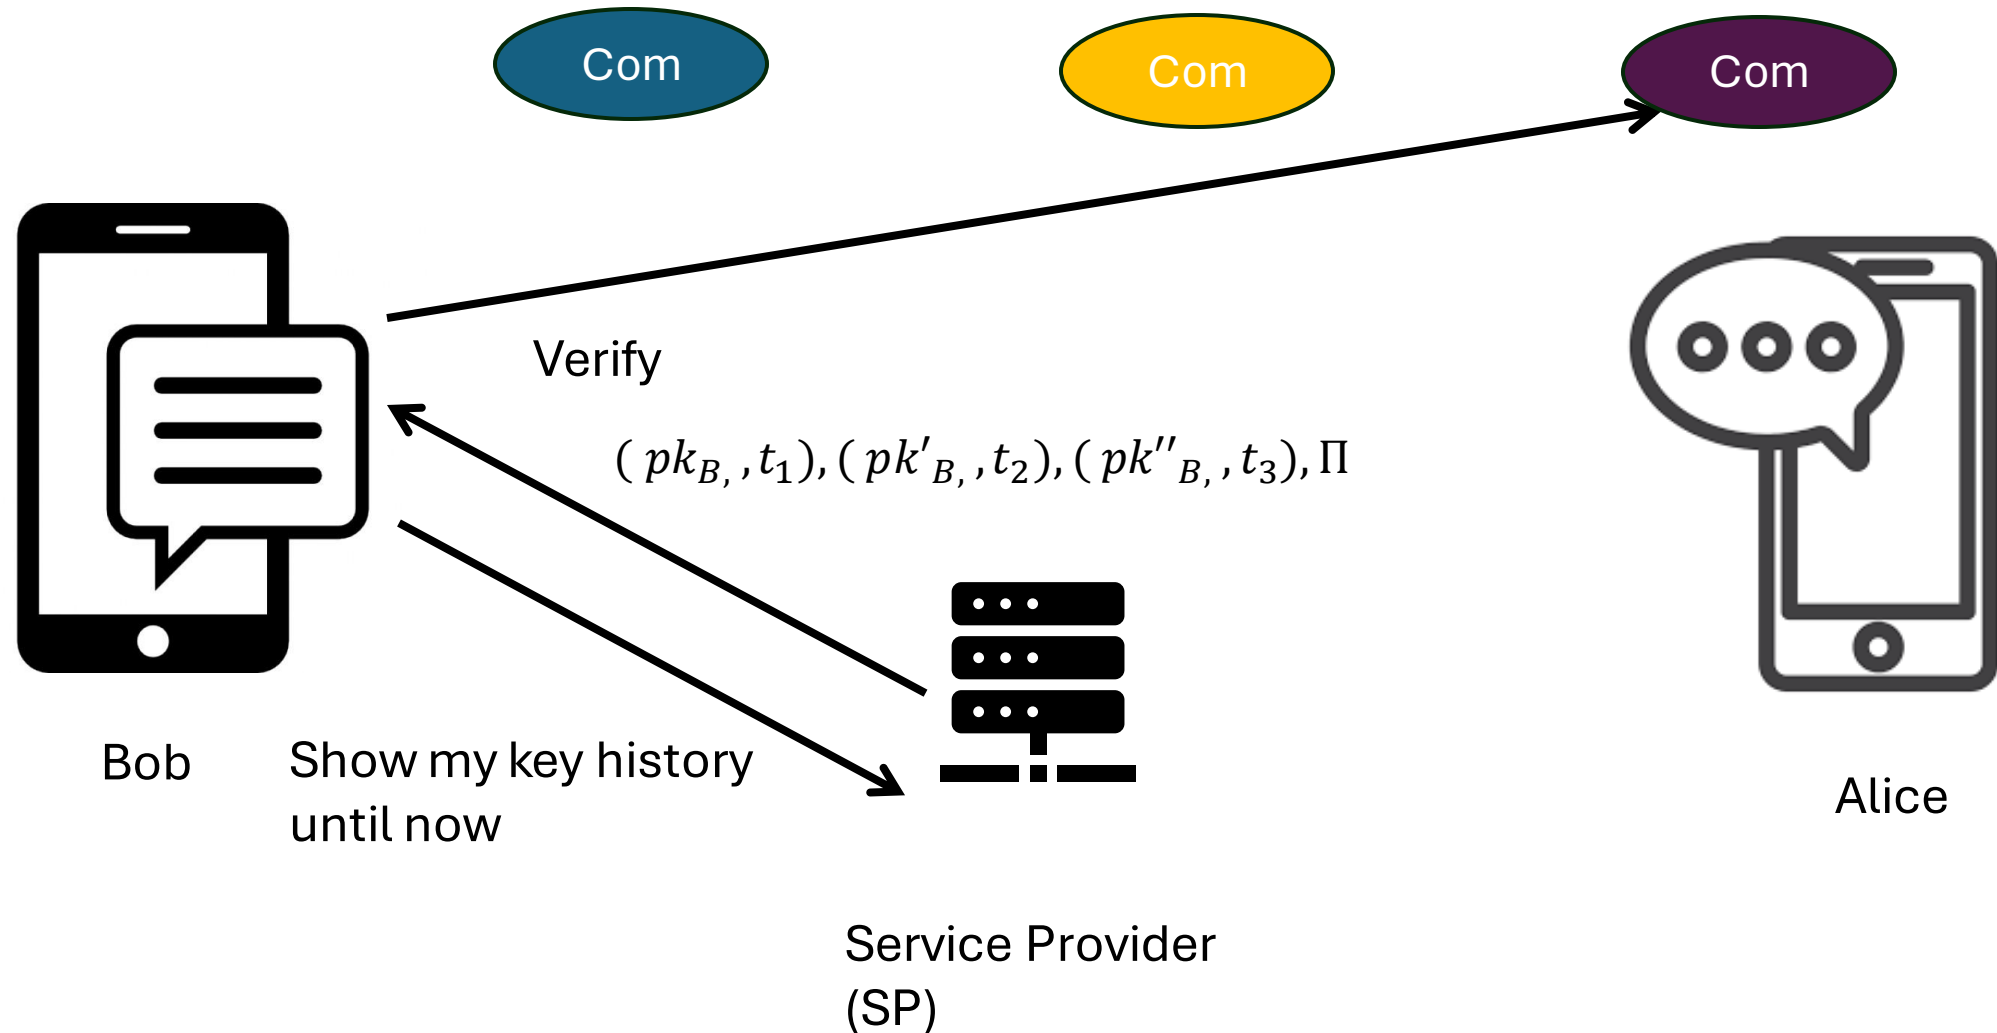
Registration/Update



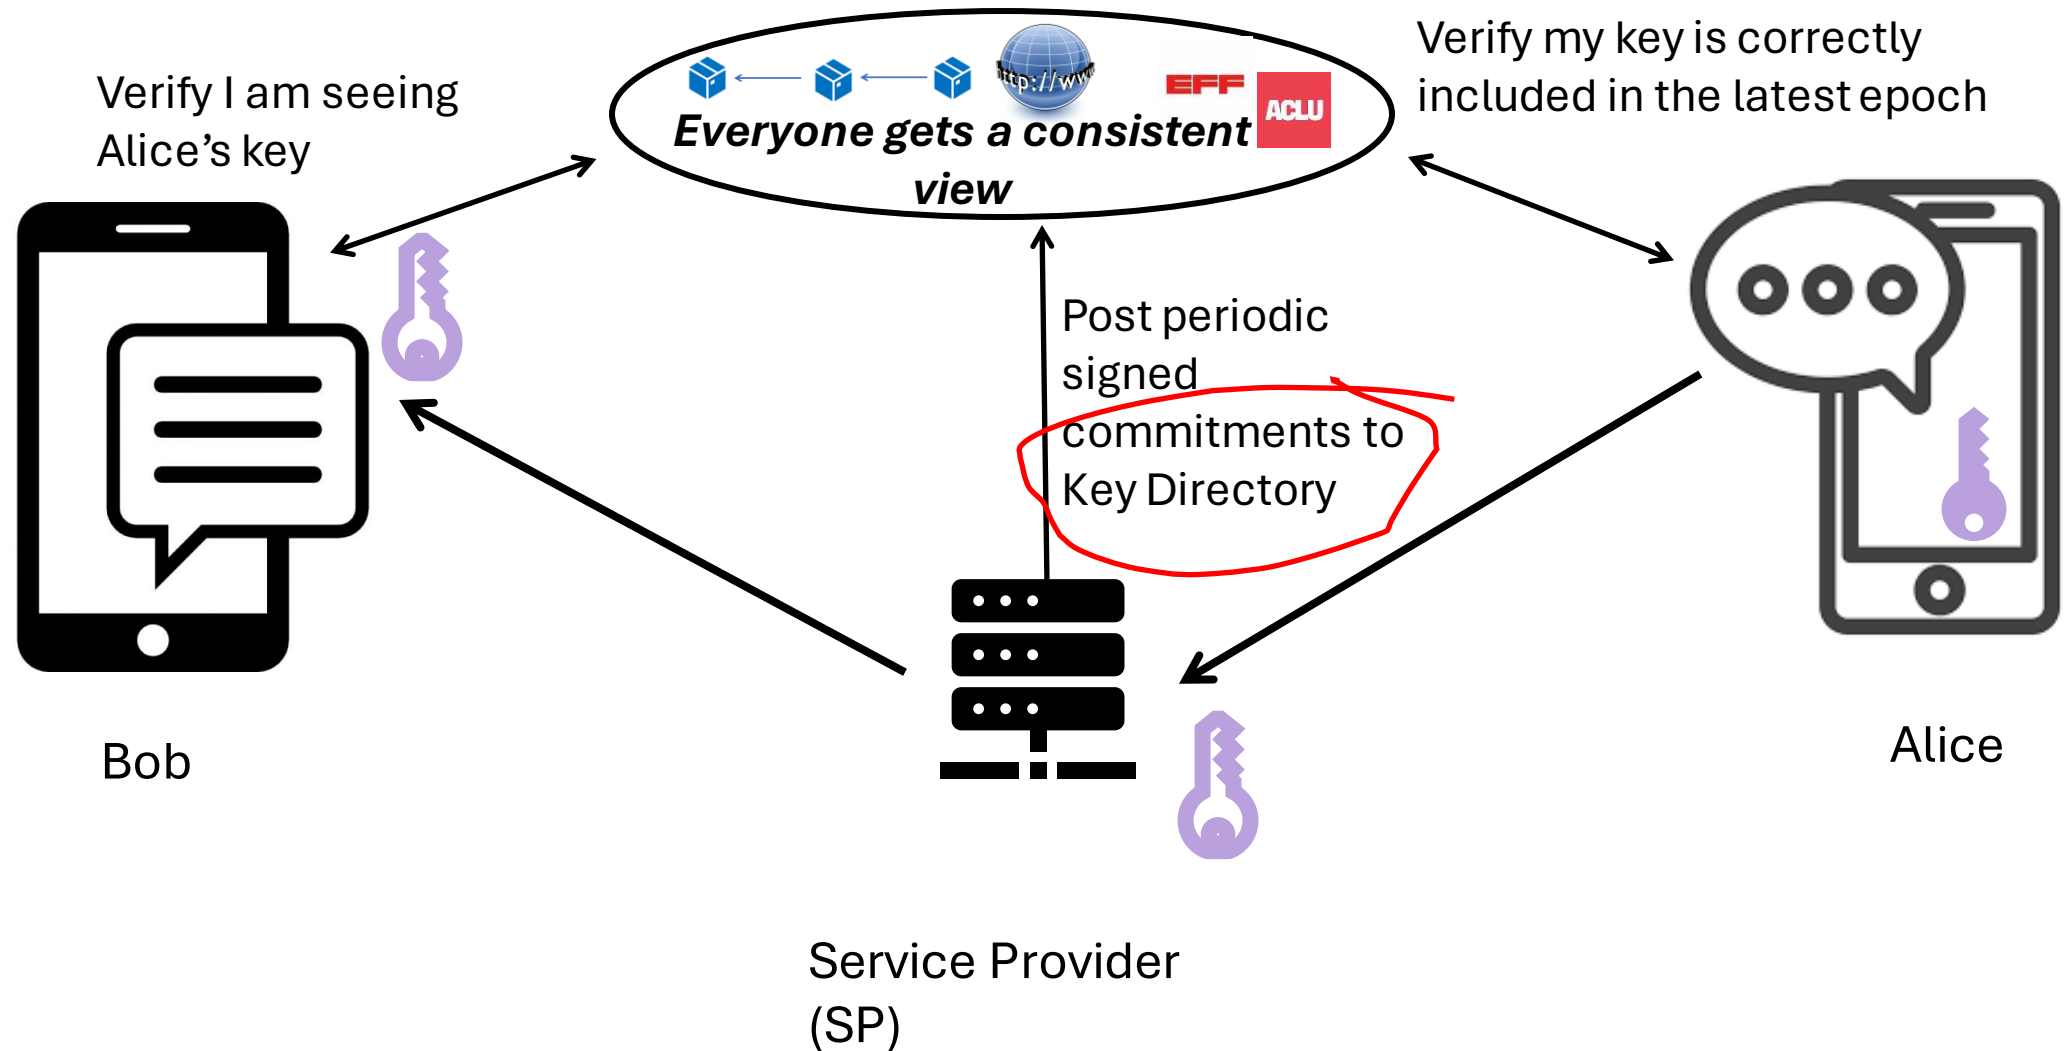
Query for other user's public key



Query for own key history



Key Transparency





Core Data Structure

Key Transparency

- Recall: commitment to a directory.
 - Map usernames to public keys
 - Must be a unique entry per username
- Sparse Merkle Tree construction
 - First used in Keybase
 - Approach underlies many current proposals: CONIKS, SEEMless, Merkle², Parakeet, ELEKTRA, OPTIKS + All the industry deployments

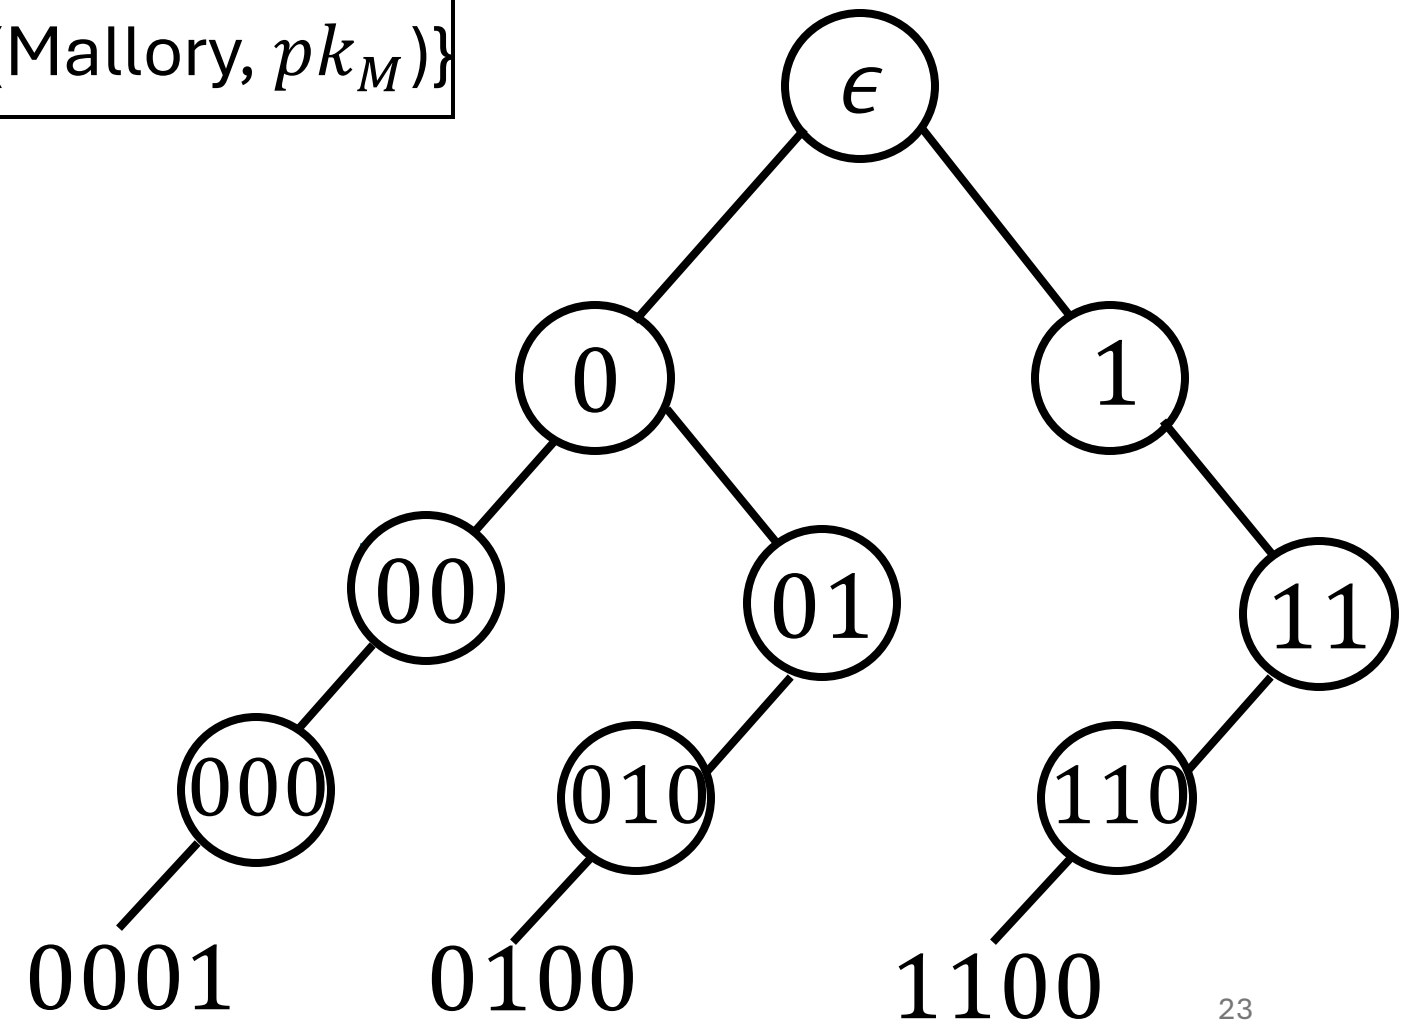
Sparse Merkle Tree construction

$S = \{ (Alice, pk_A), (Bob, pk_B), (Mallory, pk_M) \}$

$H(Alice) = 0001$

$H(Bob) = 0100$

$H(Mallory) = 1100$




Sparse Merkle Tree construction

$S = \{ (Alice, pk_A), (Bob, pk_B), (Mallory, pk_M) \}$

$H(Alice) = 0001$

$H(Bob) = 0100$

$H(Mallory) = 1100$




Sparse Merkle Tree construction

$S = \{ (Alice, pk_A), (Bob, pk_B), (Mallory, pk_M) \}$

$H(Alice) = 0001$

$H(Bob) = 0100$

$H(Mallory) = 1100$




Sparse Merkle Tree construction

$S = \{ (Alice, pk_A), (Bob, pk_B), (Mallory, pk_M) \}$

$H(Alice) = 0001$

$H(Bob) = 0100$

$H(Mallory) = 1100$



Sparse Merkle Tree construction


$S = \{ (Alice, pk_A), (Bob, pk_B), (Mallory, pk_M) \}$

$H(Alice) = 0001$

$H(Bob) = 0100$

$H(Mallory) = 1100$

$com = h_{root}$



Sparse Merkle Tree construction


$S = \{ (Alice, pk_A), (Bob, pk_B), (Mallory, pk_M) \}$

$H(Alice) = 0001$

$H(Bob) = 0100$

$H(Mallory) = 1100$

$com = h_{root}$



Sparse Merkle Tree construction

$S = \{ (Alice, pk_A), (Bob, pk_B), (Mallory, pk_M) \}$


$H(Alice) = 0001$

$H(Bob) = 0100$

$H(Mallory) = 1100$

$H(Carol) = 0011$

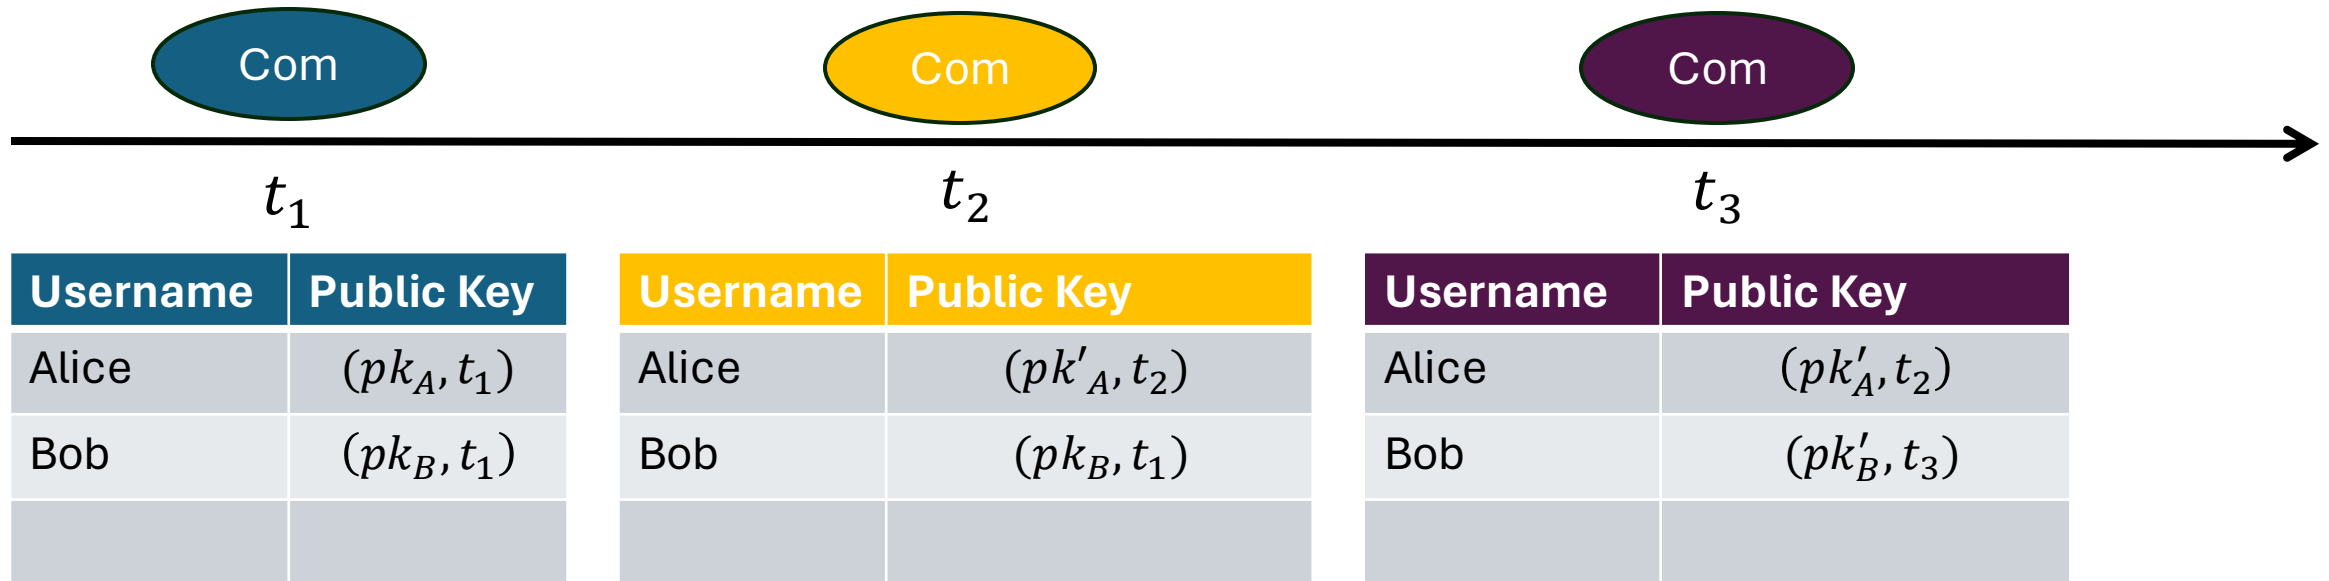
$com = h_{root}$



A hand holding a set of keys against a blurred background of a building. The image is overlaid with a blue-to-orange gradient. A white horizontal line is positioned near the top. In the top right corner, there are three small white icons: a plus sign, an open circle, and a solid circle. The main text is in white, bold font.

Building up to a Key Transparency system

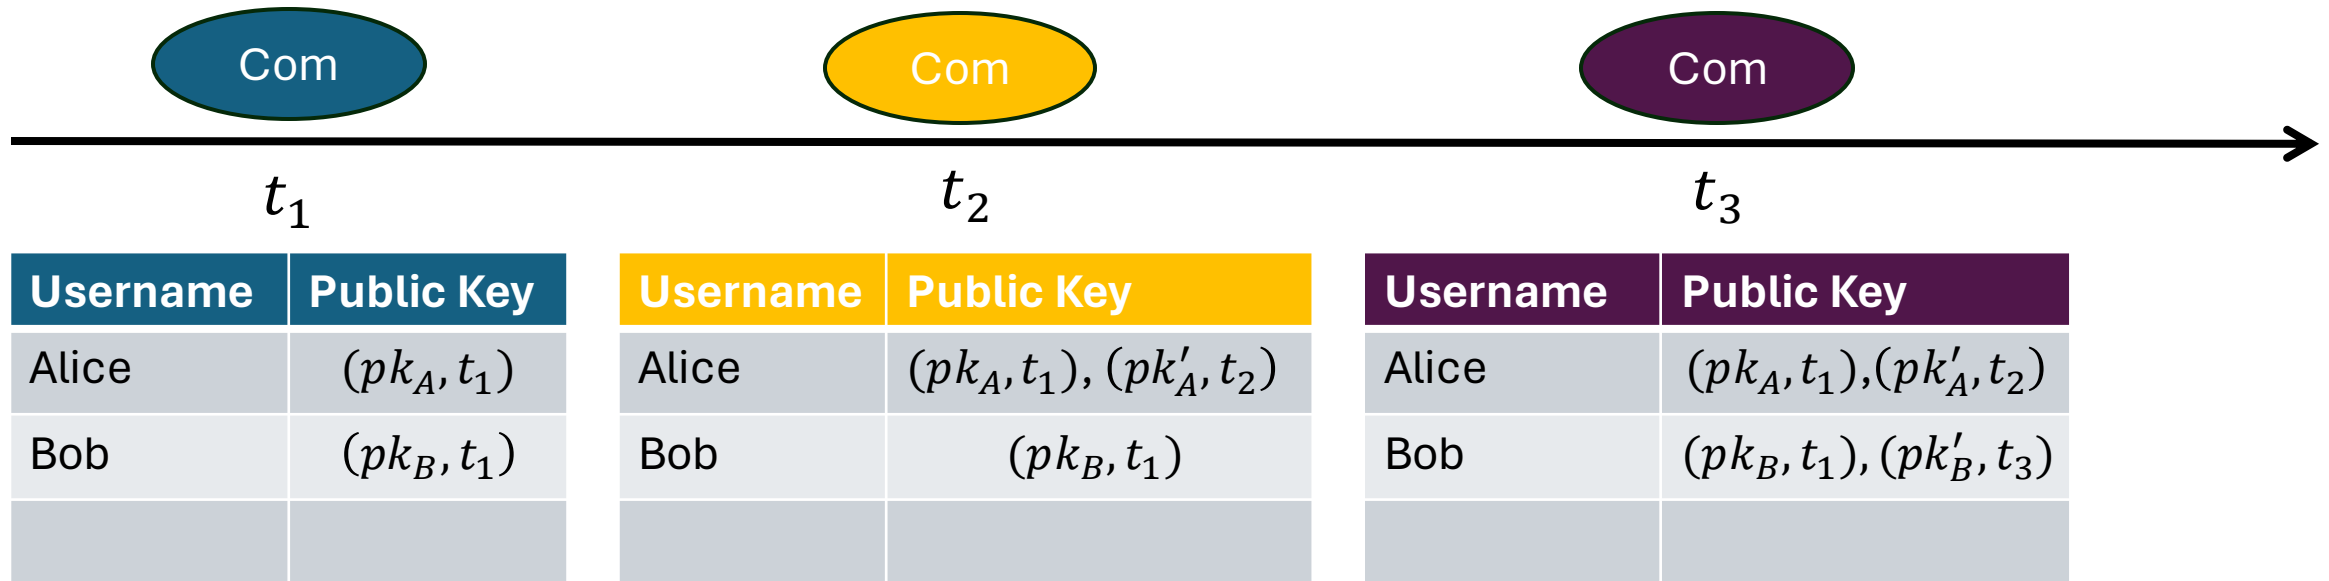
Periodic publish with Sparse Merkle Tree



Limitation: Client History Cost

- Alice's client has to check her key every epoch
- If Alice goes offline, she must check all missed commitments when she comes back online
- Limits how short epochs can be
- Note key changes aren't incorporated until the next epoch, so we want epochs to be *short*

Solution: Append-only Tree




Solution: Append-only Tree

$H(\text{Alice}) = 0001$

$H(\text{Bob}) = 0100$

$H(\text{Mallory}) = 1100$

Requires: an honest auditor to check
no commitment is ever deleted from
the leaves



$$h_1 = H(0001 | com_1, com_2)$$
$$com_1 = com(pk_A), com_2 = com(pk'_A)$$

Limitation: Privacy

- Brute force search reveals usernames
- Reveals information about when keys are updated
 - Even if they're never queried

What if usernames are phone numbers or email addresses?

Info about when user:

- Is Compromised
- Gets new device
- Travels out of country

Sparse Merkle Tree construction – Privacy?

$S = \{ (Alice, pk_A), (Bob, pk_B), (Mallory, pk_M) \}$


Querying Bob reveals Bob has a sibling, so it reveals info about $H(Alice)$!

$H(Alice) = 0001$

$H(Bob) = 0100$

$H(Mallory) = 1100$

$com = h_{root}$



Sparse Merkle Tree construction – Privacy?

$S = \{ (Alice, pk_A), (Bob, pk_B), (Mallory, pk_M) \}$

Querying Bob reveals Bob has a sibling, so it reveals info about $H(Alice)$!

$H(Alice) = 0001$


$H(Bob) = 0100$

$H(Mallory) = 1100$

$H(Carol) = 0011$

$com = h_{root}$

Querying Bob reveals whether Alice has changed her key!



Sparse Merkle Tree construction – Privacy?

$S = \{ (Alice, pk_A), (Bob, pk_B), (Mallory, pk_M) \}$

Querying Bob reveals Bob has a sibling, so it reveals info about $H(Alice)$!

$H(Alice) = 0001$

$H(Bob) = 0010$


$H(Mallory) = 0100$

Auditors will learn exactly which user positions are updating their keys without querying!

$H(Carol) = 0011$

$com = h_{root}$

Querying Bob reveals whether Alice has changed her key!



Limitation: Privacy

- Brute force search reveals usernames
- Auditor learns when each leaf is updated
- Reveals information about when keys are updated
 - Even if they're never queried

What if usernames are phone numbers or email addresses?

Info about when user:

- Is Compromised
- Gets new device
- Travels out of country

Privacy

- Content Privacy
- Metadata Privacy
- Post Compromise Privacy

Content Privacy



Verifiable Random Function (VRF)

- KeyGen $\rightarrow k, pk$
- Evaluate: $F_k(x) \rightarrow y$
- Prove(k, x) $\rightarrow y, \pi$
- Verify (pk, x, y, π) $\rightarrow \textit{accept/reject}$

Content Privacy: Hiding Usernames

CONIKS [MBBFF15] : To hide usernames, add a VRF[MRV99]


$$F_k(\text{Alice}) = 0001$$

$$F_k(\text{Bob}) = 0100$$

$$F_k(\text{Mr. Mary}) = 1100$$

Service provider proves
this is computed correctly

$$\text{com} = h_{root}, \text{VRF } pk$$



Content Privacy: Hiding Usernames

CONIKS [MBBFF15] : To hide usernames, add a VRF[MRV99]


$$F_k(\text{Alice}) = 0001$$

$$F_k(\text{Bob}) = 0100$$

$$F_k(\text{Mallory}) = 1100$$

$$F_k(\text{Carol}) = ?$$

$$\text{com} = h_{root}, \text{VRF } pk$$



Metadata Privacy


- +
-
-

Metadata Privacy: Hiding when user updates their keys

$$F_k(\text{Alice}) = 0001$$

$$F_k(\text{Bob}) = 0100$$

$$F_k(\text{Mallory}) = 1100$$

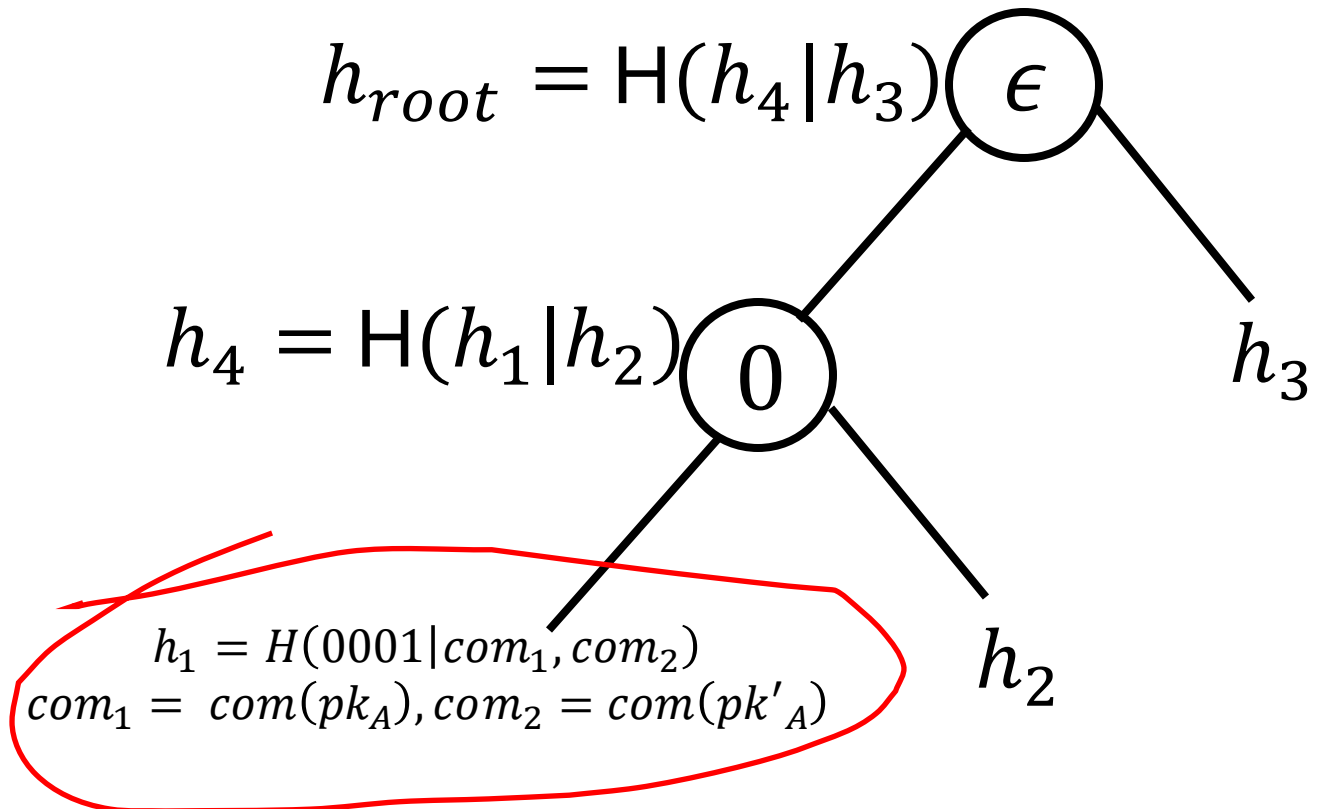


Metadata Privacy: Hiding when user updates their keys

$$F_k(\text{Alice}) = 0001$$

$$F_k(\text{Bob}) = 0100$$

$$F_k(\text{Mallory}) = 1100$$



Metadata Privacy: Hiding when user updates their keys

SEEMless [CDGM19]: label = username||version


$$F_k(\text{Alice}|\mathbf{v.1}) = 0001$$

$$F_k(\text{Bob}|\mathbf{v.3}) = 0100$$

$$F_k(\text{Mallory}|\mathbf{v.6}) = 1100$$

Tree includes all versions

- Leaf includes epoch it was added



Metadata Privacy: Hiding when user updates their keys

SEEMless [CDGM19]: label = username||version

$$F_k(\text{Alice}|\mathbf{v.1}) = 0001$$

$$F_k(\text{Bob}|\mathbf{v.3}) = 0100$$

$$F_k(\text{Mallory}|\mathbf{v.6}) = 1100$$

Tree includes all versions

- Leaf includes epoch it was added

Auditor checks:

- Leaves of old tree are subset of new one
- New leaves have correct epoch t

Privacy:

- Updates add new pseudorandom nodes
- No changing nodes
- No link between previous and new nodes

h_1

$$h_2 = H(0100 | \boxed{pk_B} | t)$$

Post Compromise Privacy



Post Compromise Privacy: Server's key leakage

Rotatable Zero Knowledge Sets [CDGGKMM22]

$$F_{k_1}(\text{Alice}|\mathbf{v.1}) = 0001$$

$$F_{k_1}(\text{Bob}|\mathbf{v.3}) = 0100$$

$$F_{k_1}(\text{Mallory}|\mathbf{v.6}) = 1100$$

$$F_{k_2}(\text{Alice}|\mathbf{v.1}) = 1001$$

$$F_{k_2}(\text{Bob}|\mathbf{v.3}) = 0011$$

$$F_{k_2}(\text{Mallory}|\mathbf{v.6}) = 1010$$

VRF pk_1

VRF pk_2

Verify Rotation: $(pk_1, pk_2, y_1, \dots, y_n, y'_1, \dots, y'_n, \pi_{rotation}) \rightarrow \text{accept/reject}$



A Simple Key Transparency System

OPTIKS-core [LCGKM24]

Commitment as described

Lookup(Bob):

$F_k(\text{Bob}|\mathbf{v.1})$, π_{VRF} , π_{member}


$F_k(\text{Bob}|\mathbf{v.2})$, π_{VRF} , π_{member}

$F_k(\text{Bob}|\mathbf{v.3})$, π_{VRF} , π_{member} ,
 pk_B , $open_3$, t_3

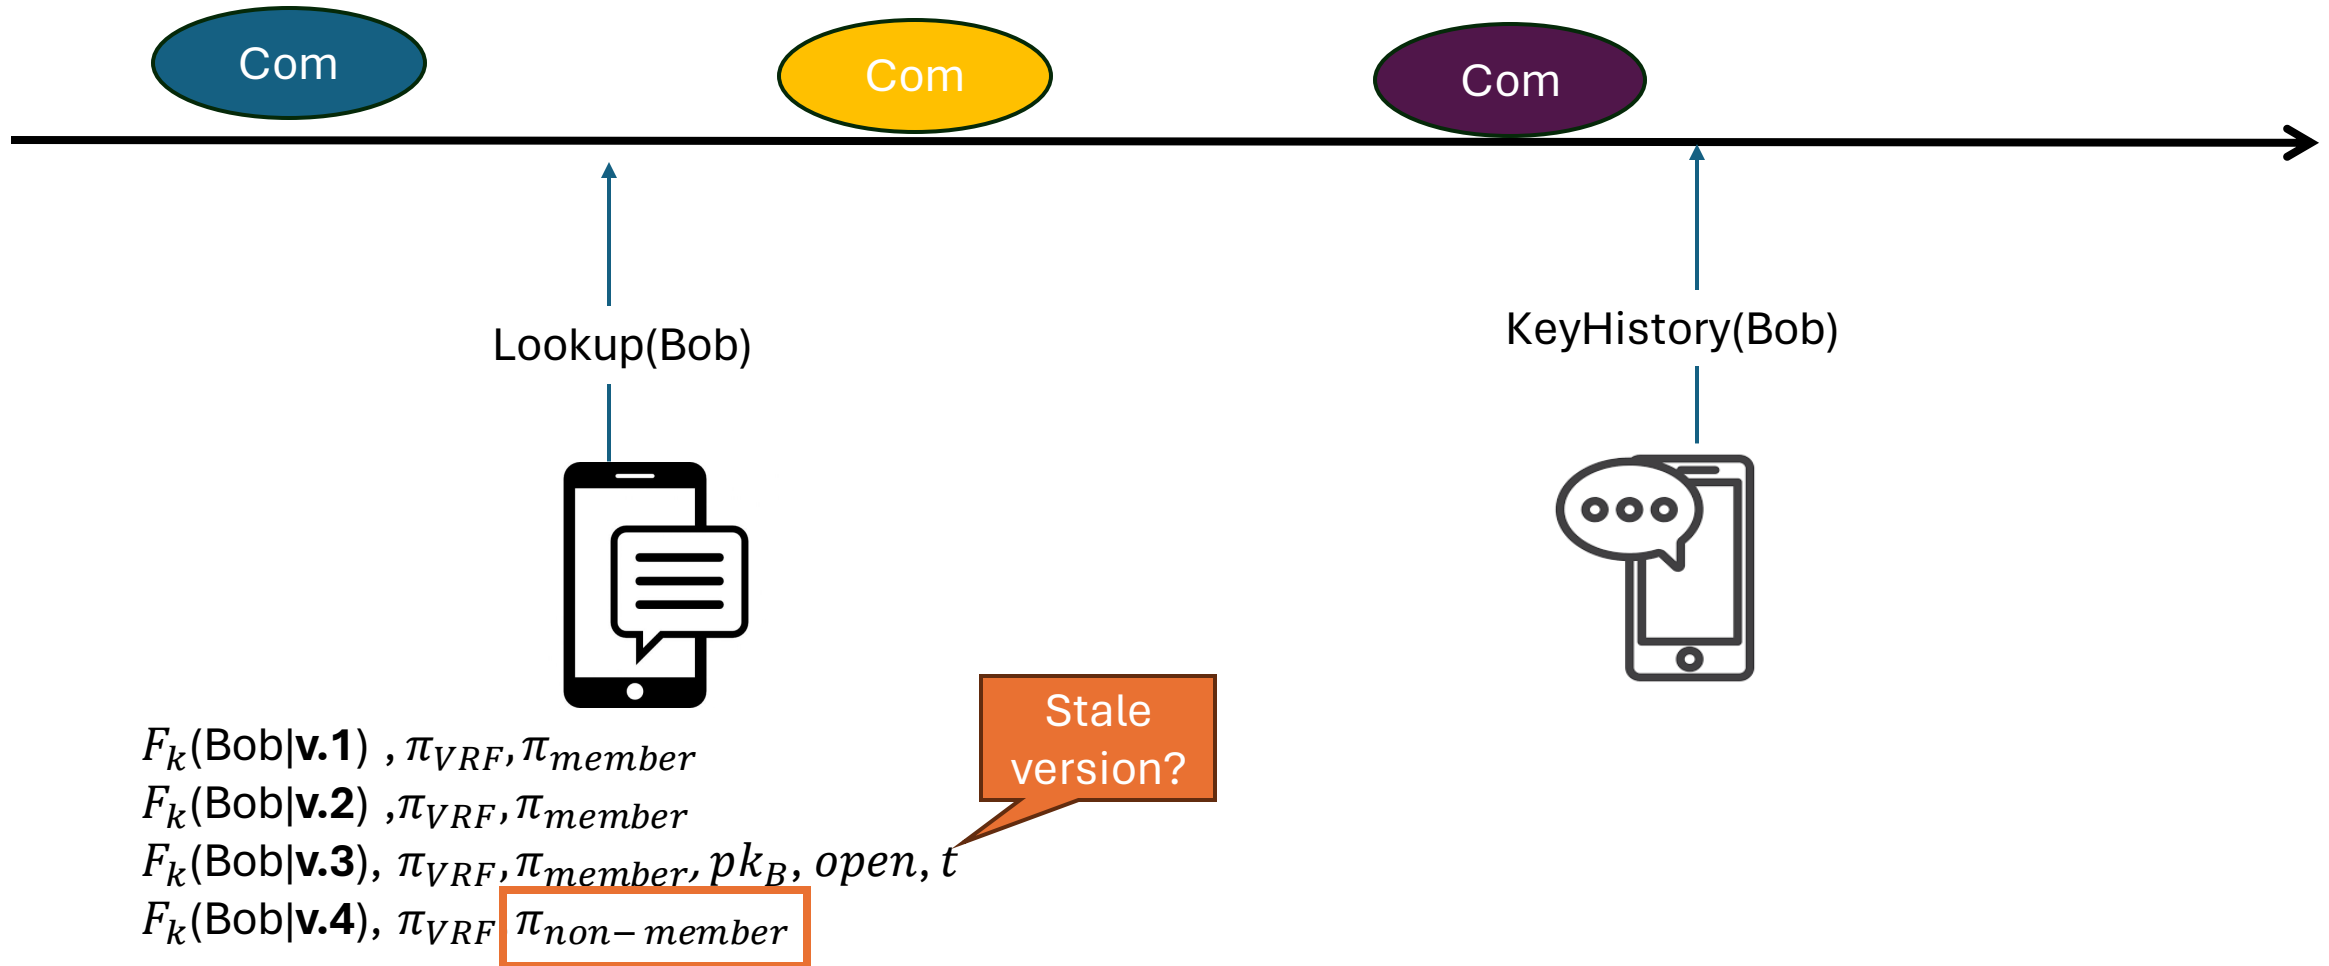
$F_k(\text{Bob}|\mathbf{v.4})$, π_{VRF} , $\pi_{non-member}$

History(Bob):

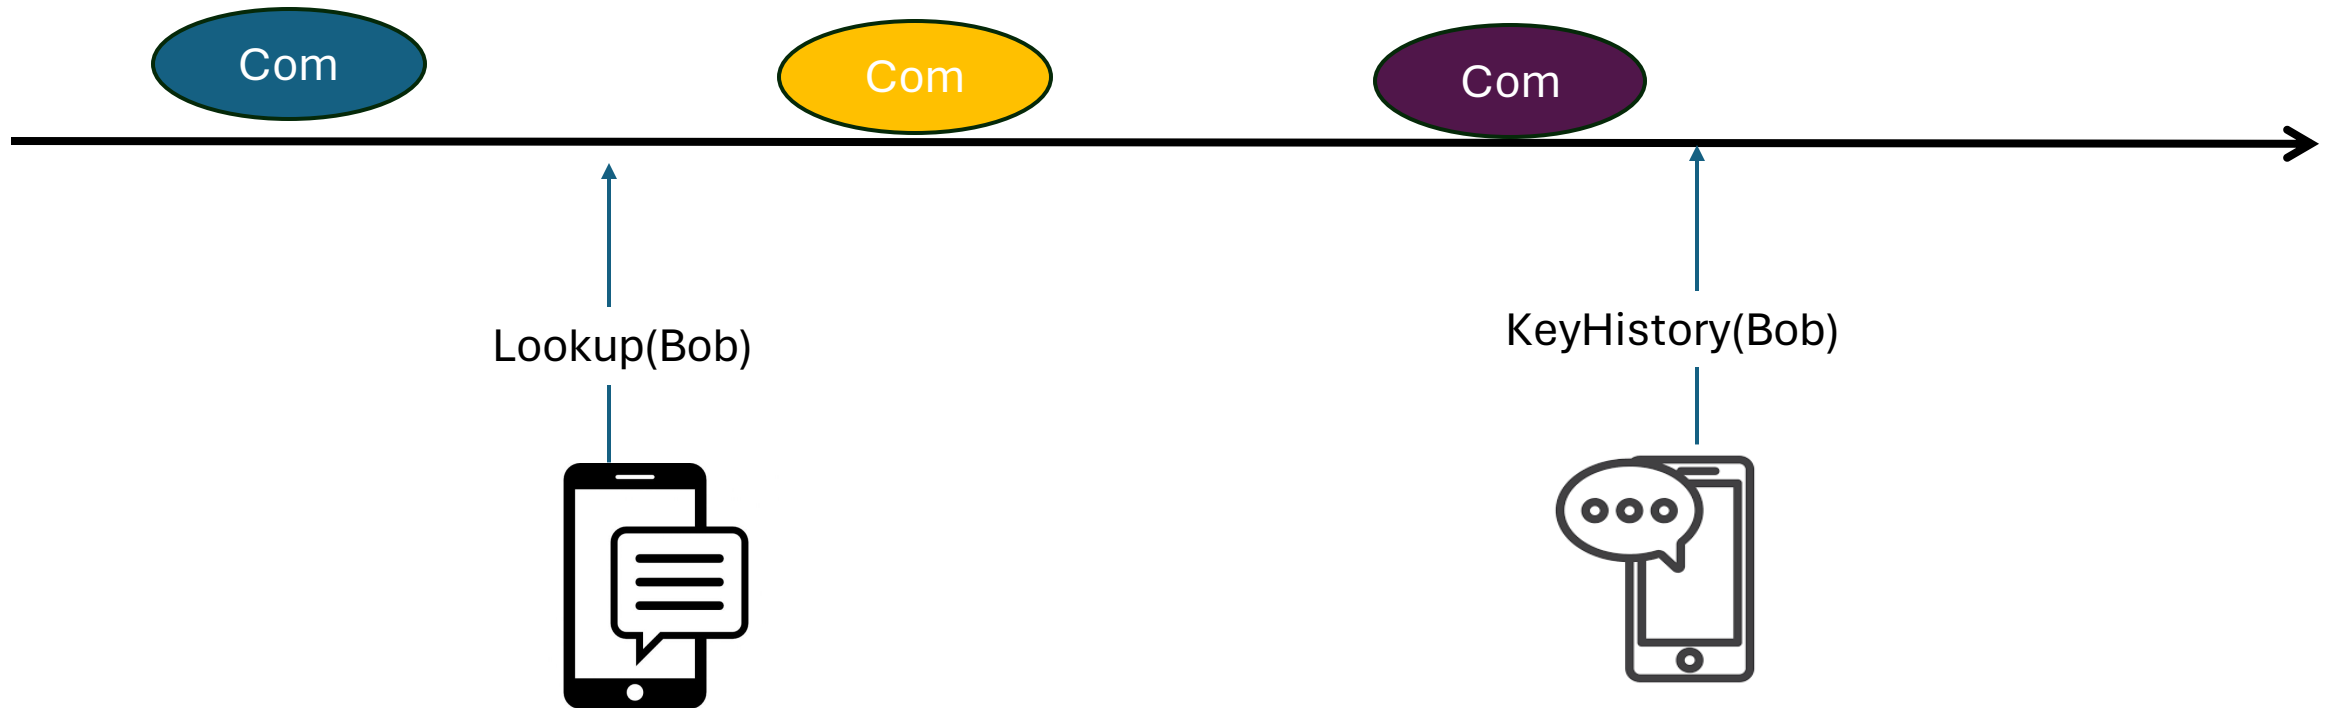
- similar but with $(pk_B, open, t)$ for each version



Soundness



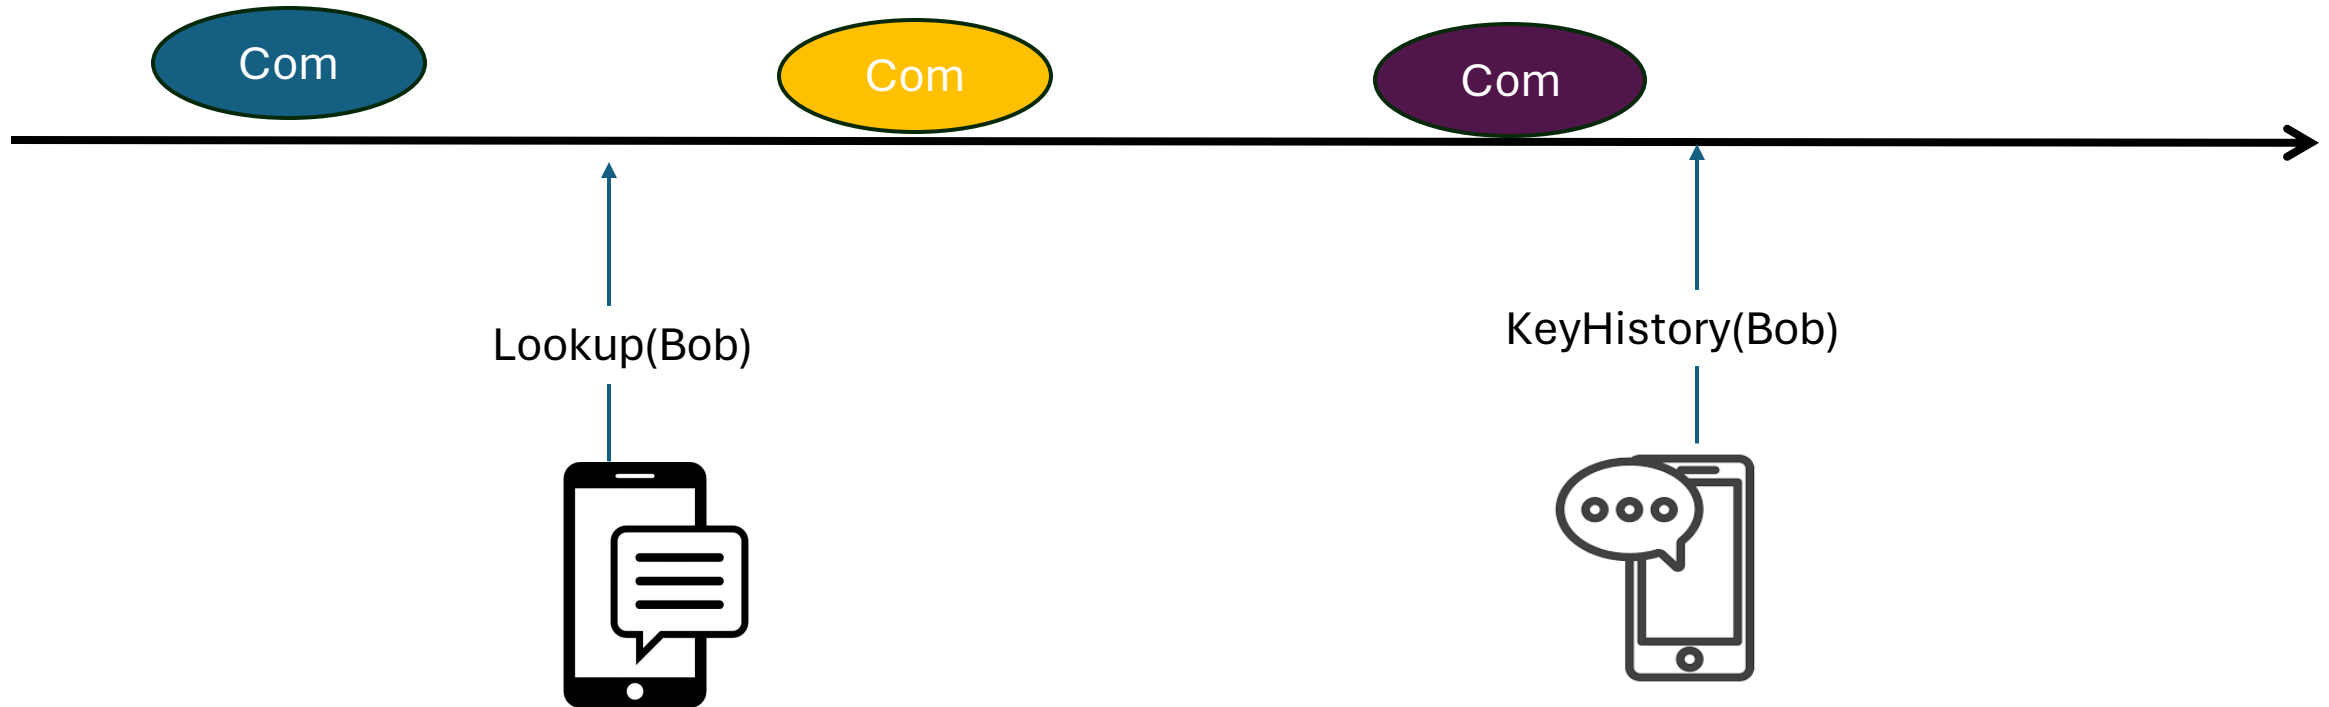
Soundness



$F_k(\text{Bob}|\mathbf{v.1}), \pi_{VRF}, \pi_{member}$
 $F_k(\text{Bob}|\mathbf{v.2}), \pi_{VRF}, \pi_{member}$
 $F_k(\text{Bob}|\mathbf{v.3}), \pi_{VRF}, \pi_{member}, pk_B, open_3, t_3$
 $F_k(\text{Bob}|\mathbf{v.4}), \pi_{VRF}, \pi_{non-member}$

$F_k(\text{Bob}|\mathbf{v.1}), \pi_{VRF}, \pi_{member}, pk_{B1}, open_1, t_1$
 $F_k(\text{Bob}|\mathbf{v.2}), \pi_{VRF}, \pi_{non-member}$

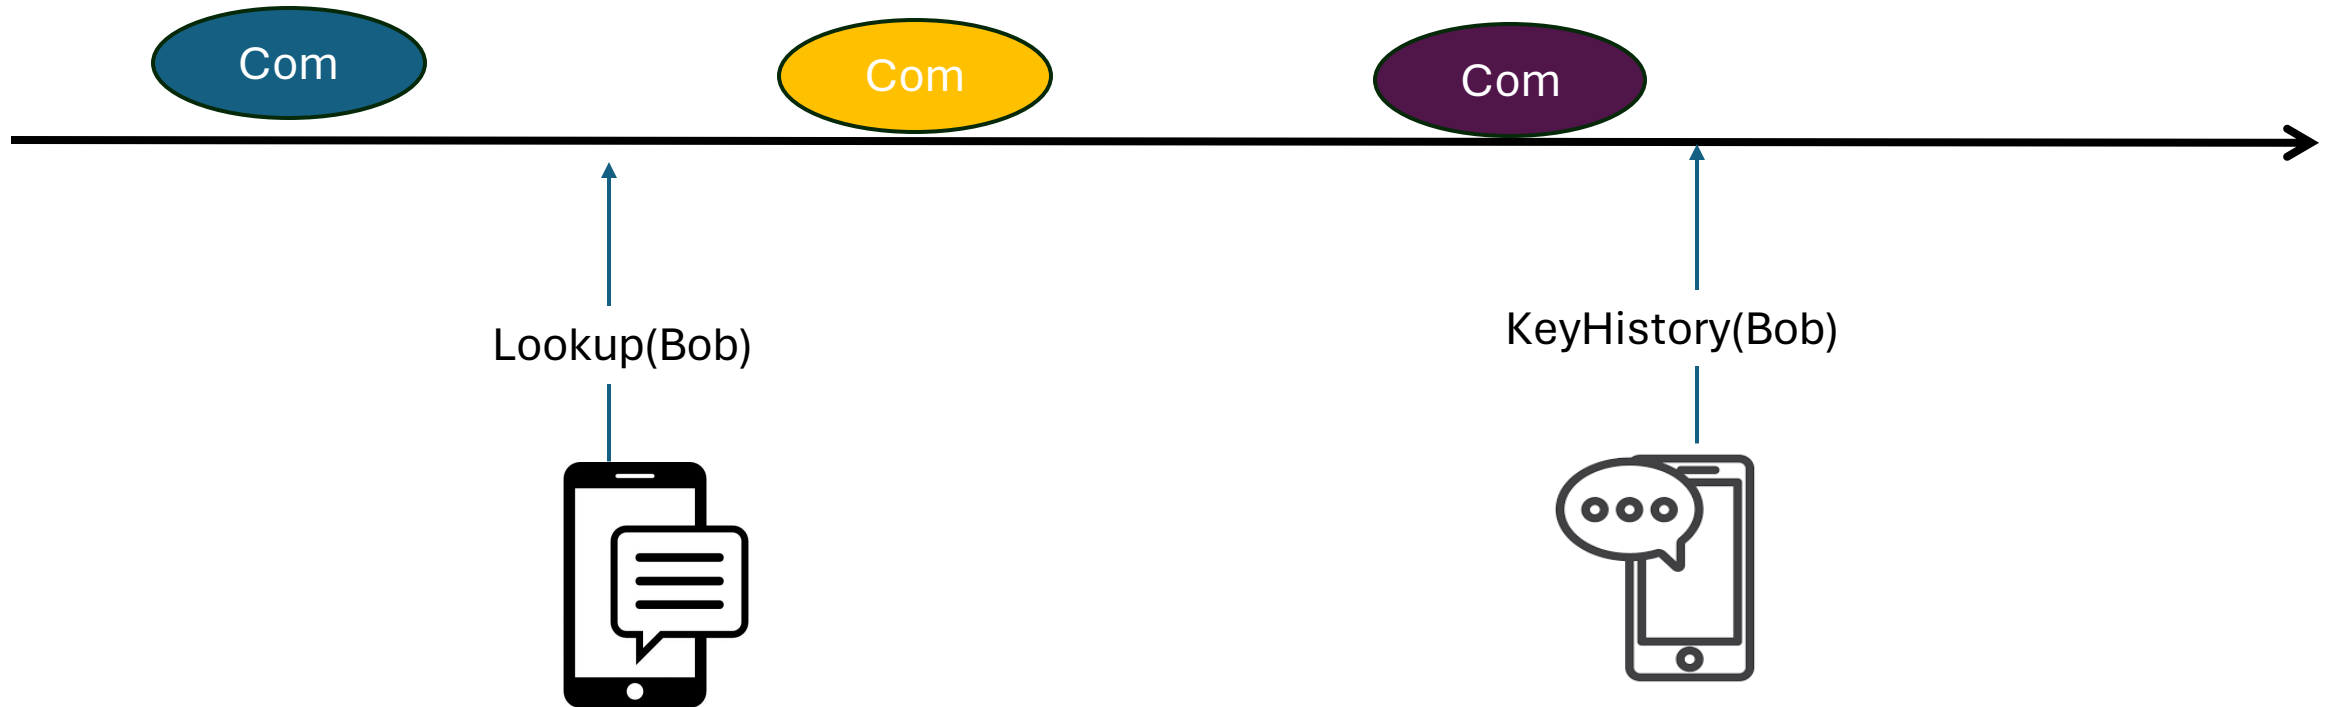
Soundness



$F_k(\text{Bob}|\mathbf{v.1}), \pi_{VRF}, \pi_{member}$
 $F_k(\text{Bob}|\mathbf{v.2}), \pi_{VRF}, \pi_{member}$
 $F_k(\text{Bob}|\mathbf{v.3}), \pi_{VRF}, \pi_{member}, pk_B, open_3, t_3$
 $F_k(\text{Bob}|\mathbf{v.4}), \pi_{VRF}, \pi_{non-member}$

$F_k(\text{Bob}|\mathbf{v.1}), \pi_{VRF}, \pi_{member}, pk_{B1}, open_1, t_1$
 $F_k(\text{Bob}|\mathbf{v.2}), \pi_{VRF}, \pi_{member}, pk_{B2}, open_2, t_2$
 $F_k(\text{Bob}|\mathbf{v.3}), \pi_{VRF}, \pi_{non-member}$

Soundness



$F_k(\text{Bob}|\mathbf{v.1}), \pi_{VRF}, \pi_{member}$
 $F_k(\text{Bob}|\mathbf{v.2}), \pi_{VRF}, \pi_{member}$
 $F_k(\text{Bob}|\mathbf{v.3}), \pi_{VRF}, \pi_{member}, pk_B, open, t$
 $F_k(\text{Bob}|\mathbf{v.4}), \pi_{VRF}, \pi_{non-member}$

$F_k(\text{Bob}|\mathbf{v.1}), \pi_{VRF}, \pi_{member}, pk_{B1}, open_1, t_1$
 $F_k(\text{Bob}|\mathbf{v.2}), \pi_{VRF}, \pi_{member}, pk_{B2}, open_2, t_2$
 $F_k(\text{Bob}|\mathbf{v.3}), \pi_{VRF}, \pi_{member}, pk_{B3}, open_3, t_3$
 $F_k(\text{Bob}|\mathbf{v.4}), \pi_{VRF}, \pi_{non-member}$



Industry Deployments

Core Data Structure and Privacy

	Keybase	WhatsApp	iMessage	ProtonKT
Sparse Merkle Tree	✓	✓	✓	✓
Content Privacy (with VRF)	✗	✓	✓	✓
Metadata Privacy (unlinking key updates)	✗	✓	✗	✗
Post Compromise Privacy	✗	✗	✗	✗



Open Challenge: Auditor-Free Key Transparency

Auditing

Auditors check that commitment preserves previous history

- Necessary to reduce client monitoring

Proofs do not reveal information about individual users/updates

- Auditors do not need to be trusted for privacy

Auditors can be

- Users?
- 3rd parties?

Need Auditors?

Keybase	WhatsApp	iMessage	ProtonKT
✓	✓	[Additional assumption: tamper proof secret storage on server] ?	✓

Research Directions

Reducing auditing work with SNARKS (Verdict[TKPS22], Versa[TFZBT22])

Client monitoring with weak consistency (Versa[TFZBT22]):


- Everytime Alice comes online she checks log commitments
- If Alice and Bob both do these checks, they either
 - get correct keys for one another
 - or detect misbehavior

Open problem: compress audit proofs and verification cost/optimized client auditing



Preliminary Result: Auditor-Free Client Monitoring Mode [GC24]

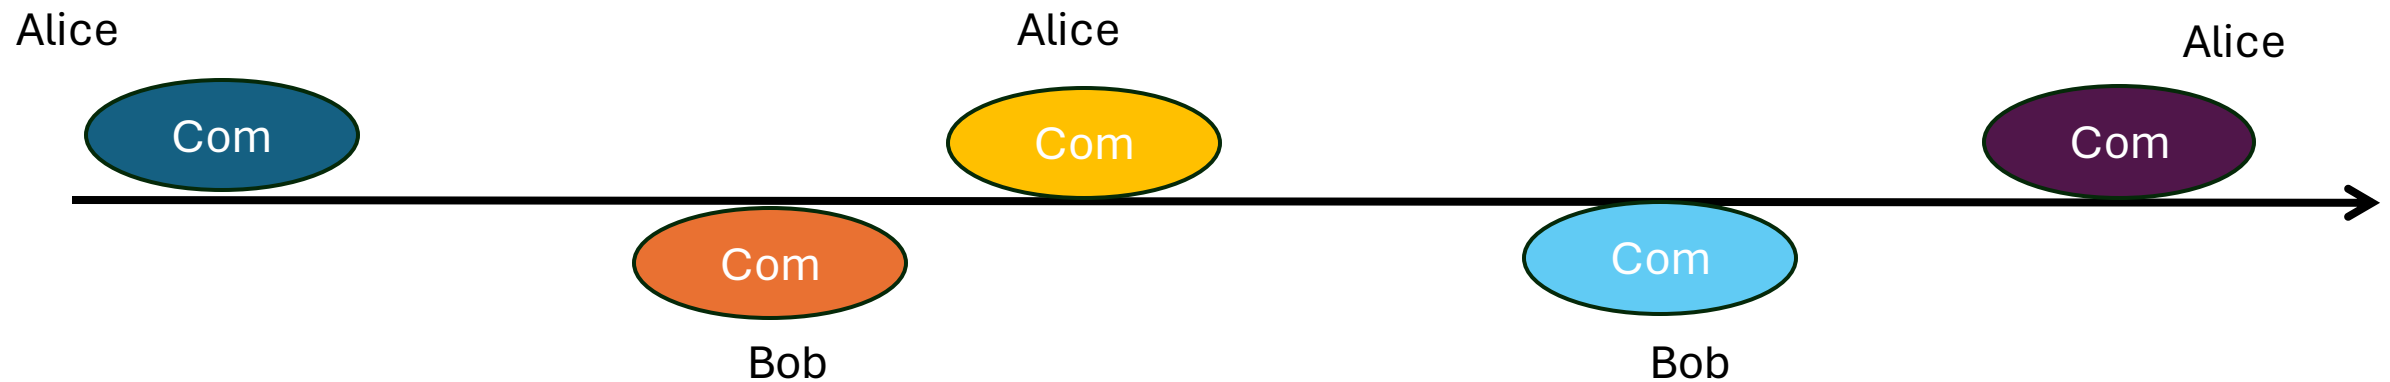
In the absence of a trusted auditor



No append-only checks between commitments

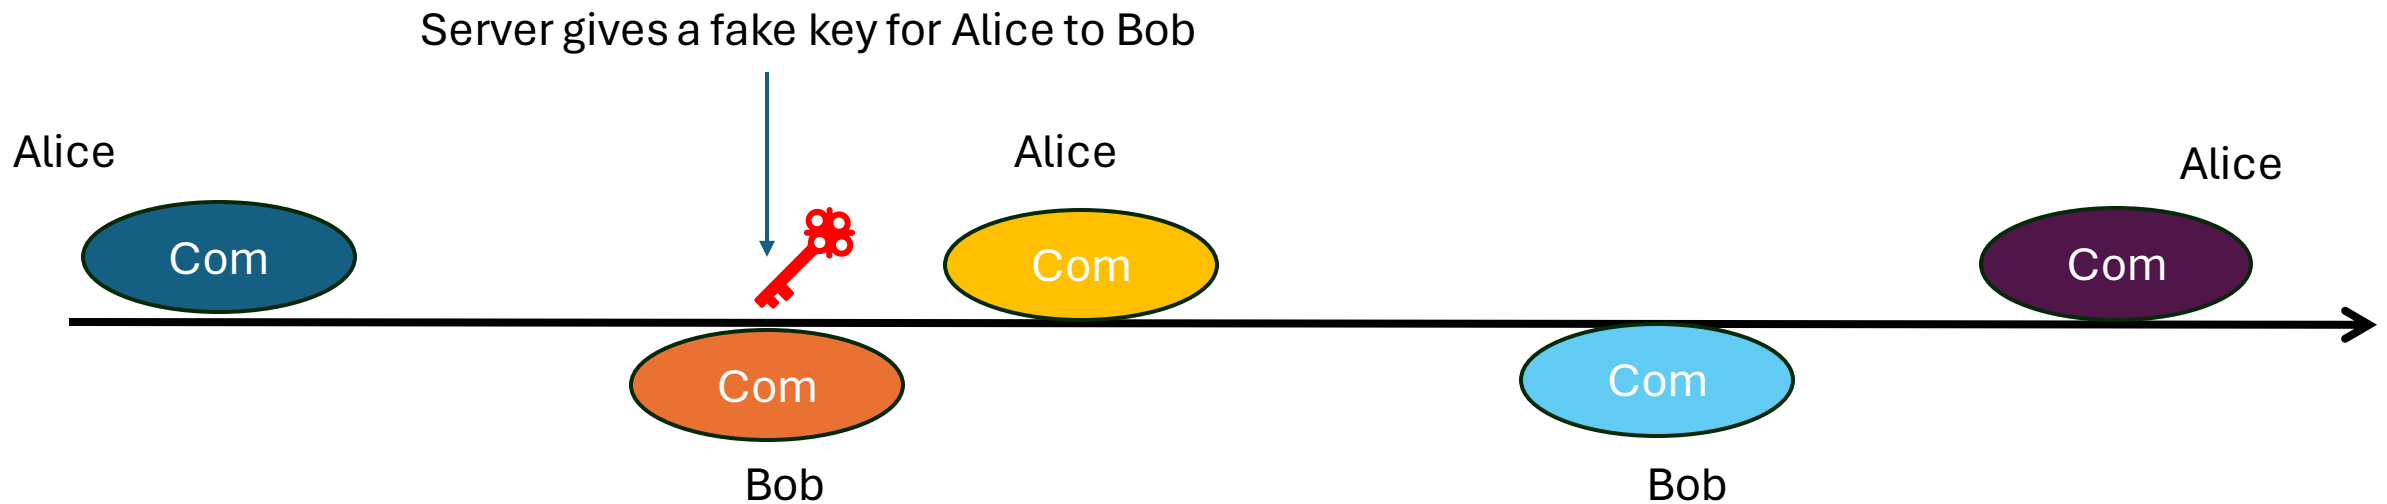
Can Alice and Bob check some additional proofs?

- A new algorithm: Monitor
- Alice and Bob performs monitor sporadically



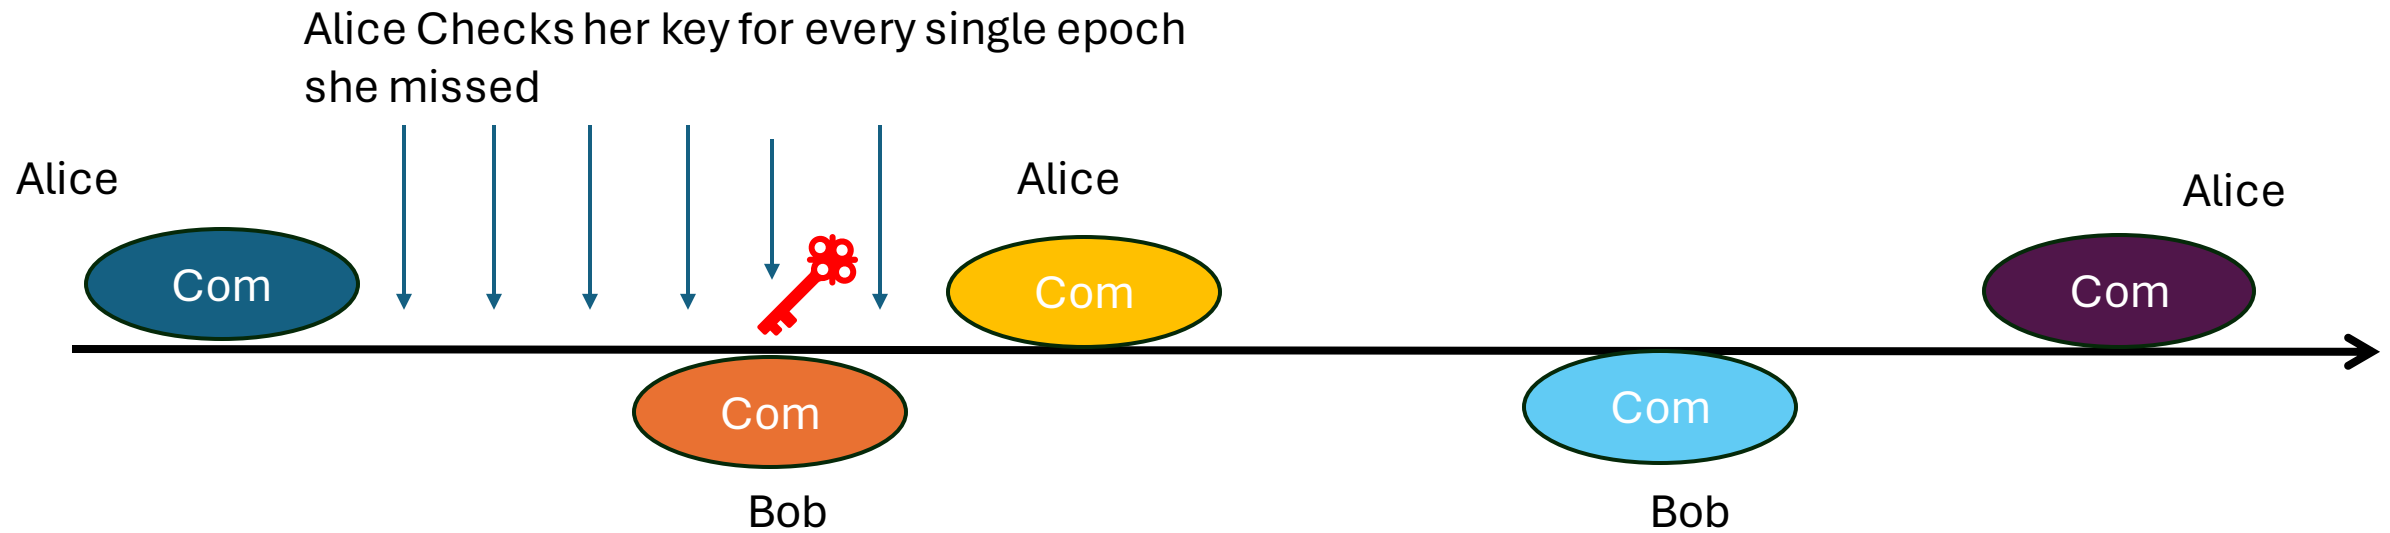
Can Alice and Bob check some additional proofs?

- A new algorithm: Monitor
- Alice and Bob performs monitor sporadically



We want this to be detectable even if no auditor is present

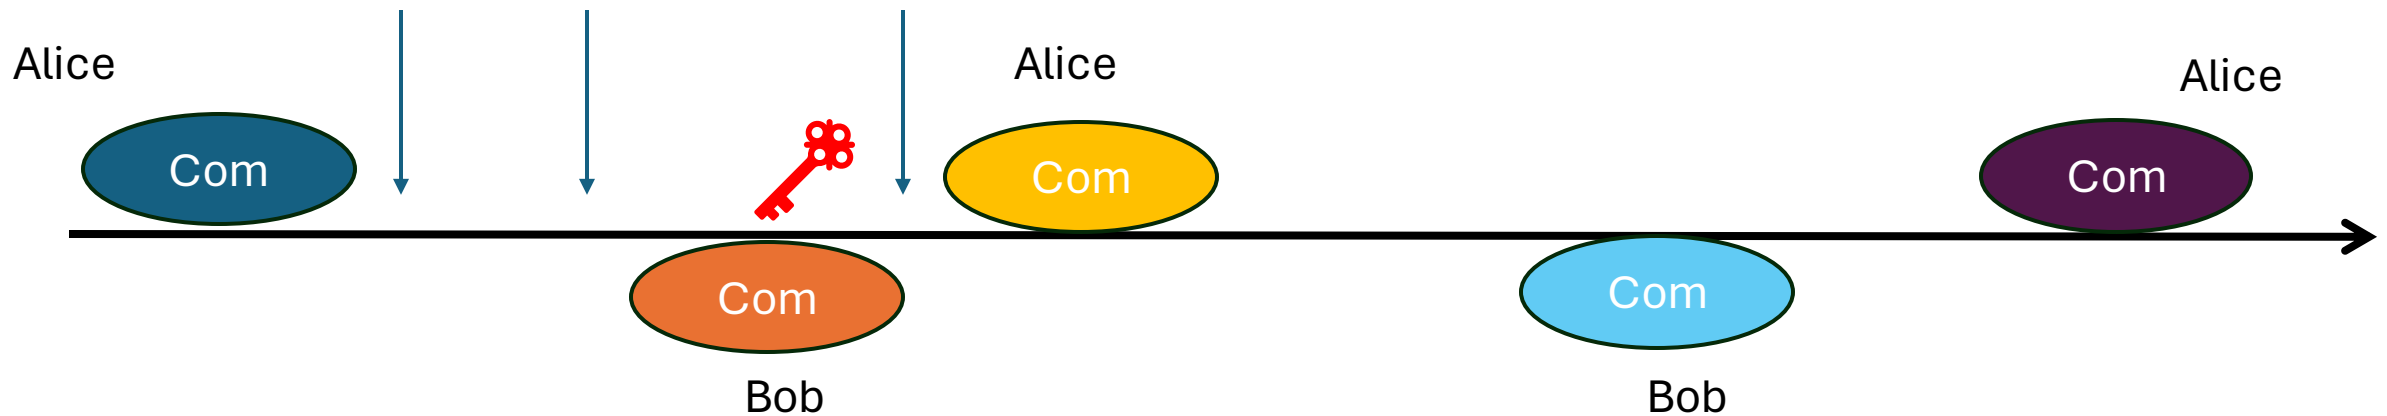
Inefficient Solution: Monitor Checks every epoch



Can we do better?


Not trivially, since Alice can easily miss the epoch in which the fake key was given out!

Can we achieve the detectability even when Alice checks log (number of missed epochs)?




Checkpointing

A deterministic algorithm that picks logarithmic number of points from on an integer interval [TFZBT22]



Checkpointing

Guarantee: This checkpoint does not necessarily coincide with the bad intersection of epochs!



Solution

- At each checkpoint:
 - Run a OPTIKS type check for each of the keys
 - Check that the history of keys returned at each checkpoint are consistent

Full protocol specification in: [eprint 2024/796](#)



Challenges and Open Problems

Instantiating the Bulletin Board (BB)

- BB must be consistently viewed by all users
- Some alternatives:
 - 3rd party
 - Blockchain
 - Certificate Transparency Logs
 - Trusted hardware (e.g. Nimble[ABCJLSS23])
- Can we remove the BB?
 - Out-of-Band (OOB) gossip among users/their representatives
 - **Open problem: Can we use an in-band gossip?**
 - First suggested in [YGHZS22]
 - Apple may be doing this, but no analysis as far as we know
 - We have initial feasibility results, not yet efficient protocol

Stronger security

Service provider can replace user's key

- User will detect this but can't prevent it

Stronger security with multi devices (Keybase, ELEKTRA[LCGJKM23])

- Directory maps username to set of device keys
- Addition/removal must be signed by an existing device
- Service provider can't add/remove device by themselves

- Except with "reset"
 - added for usability
 - Rare, so can come with stronger UI messaging

Open problem: can we do any better?

Related Cryptographic Questions

Rotatable Zero Knowledge Sets/Verifiable Random Function

- Do we need GGM in order to achieve RZKS/RVRFs? Why or why not?
- Improve privacy of RVRF/RZKS

Better optimized SNARKs for compressing audit proofs?

Performant Post Quantum Secure VRF

Other Questions

How can we reduce DB cost?

What is the best way to do UI/explain KT?

How does the user trust client code?

- Code transparency? Verified client code?

What if there is a bug in the server code?

- Verified server code?

How can user tell account compromise from server misbehavior?

- What should the user do if they detect a fake key update?

Thanks!

esha.ghosh@microsoft.com

- + • SEEMless: eprint 2018/607 (*CCS*)
- • Rotatable Zero Knowledge Sets: eprint 2022/1264 (*Asiacrypt*)
- Parakeet: eprint 2023/081 (*NDSS*)
- ELEKTRA: eprint 2024/107 (*CCS*)
- OPTIKS eprint 2023/1515 (*USENIX*)
- Client Auditability in OPTIKS: eprint 2024/796