Gossip: decentralized pseudonymous private messaging protocol with quantum-resistant double ratchet and sealed metadata

Damir Vodenicarevic
dv@massa.net

2025

Contents

1 Overview **1**

2 Global architecture **2**

3 Cryptographic primitives **2**

4 Pseudonymous identities **4**

5 Gossip network assumptions **5**

6 Agraphon: post-quantum secure messaging algorithm **6**

7 Possible extensions **13**

1 Overview

Design goals

Gossip is a new type of mobile-first instant messaging system built around the following pillars:

- **Privacy by default:** all your messages are encrypted before they leave your device, and only yourself and the person you are talking to can see them in the clear. Perfect forward secrecy and post-compromise security ensure that even in the case of key leaks, previous and future messages remain secret. Metadata is also encrypted, ensuring that no outsider knows who is talking to whom. Post-quantum encryption and authentication ensures that even attackers with powerful quantum computers can't read your messages or impersonate Gossip users.

- **No personal linkability:** users authenticate based on a passphrase they choose and are identified by a unique pseudonymous ID. No phone number or personally identifying information is required, new identities are easy to create and manage. The system also provides plausible deniability, ensuring that even if the user you are talking to leaks your conversation, you can still deny ever having sent those messages, since there is no cryptographic proof that they did not just forge them because both participants are aware of all message encryption secrets and messages are not signed.
- **Censorship resistance:** Gossip does not depend on central servers, it uses a decentralized network of nodes, thus ensuring network resilience, availability and strong resistance to censorship. Gossip also provides an alternative Progressive Web App hosted on the Massa Decentralized Web (DeWeb) that remains available in case the app is censored on mobile app stores. Even if Gossip's authors disappear or try to shut down the project, Gossip will still function.
- **Accessibility:** Gossip aims to be as feature-complete and usable as possible, hiding the underlying cryptographic complexity to ensure a smooth user experience. Gossip aims to be politically and socially neutral, welcoming anyone to use it.
- **Trustlessness:** you don't need to trust Gossip's creators to feel secure. Gossip's source code has been publicly available since the project's inception. It aims to be independently audited, and formally verified as much as possible. Only cryptographic primitives with strong public scrutiny are used.

2 Global architecture

Gossip relies on a decentralized network of nodes that propagate encrypted messages to each other and temporarily store them locally in their cache for 2 weeks. Anyone can retrieve data from any node. Anyone with a computer can easily become a node of the network.

The illustration below explains how the network functions, assuming Alice wants to send an encrypted message to Bob:

- The Gossip app on Alice's phone encrypts the message in a way that only herself or Bob can decrypt it, and sends it to one or more nodes in the Gossip network.
- The nodes of the Gossip network propagate the encrypted message and temporarily store it in their local cache.
- The Gossip app on Bob's phone queries one or more nodes of the network to retrieve the message, decrypts it, stores it in a secure encrypted storage space within Bob's phone, which allows Bob to read the message.

3 Cryptographic primitives

Password KDF

When deriving cryptographic keys from a user-specified password, Gossip uses a password-safe key derivation function (Password KDF) that makes it difficult for attackers to brute-force the

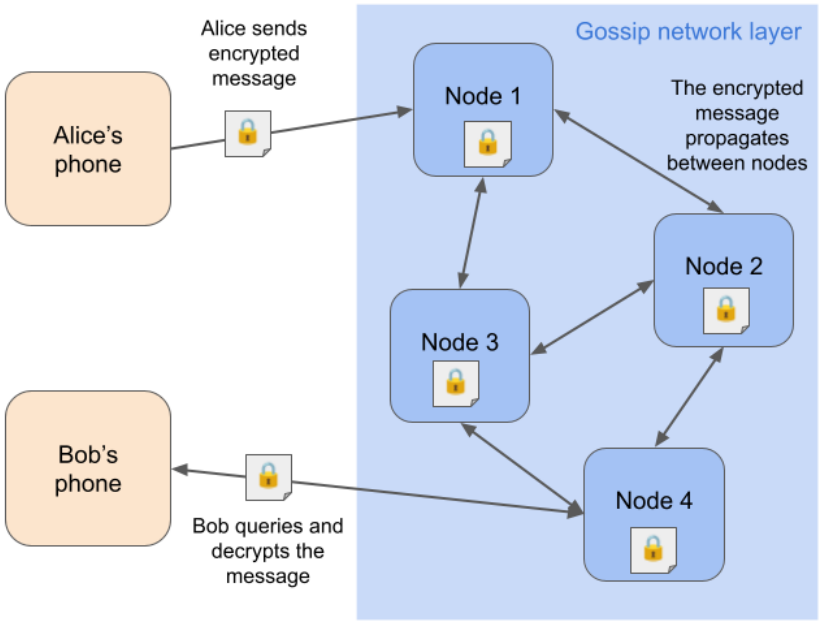


Figure 1: Gossip network layer and message propagation.

passphrase by computing many possible ones.

For this purpose, Gossip uses the Argon2id password-safe, post-quantum KDF with the following parameters:

- Version: V0x13 – The latest version of Argon2id
- Memory cost: 32 MiB – Reasonable for mobile devices.
- Parallelization cost: 1 – Single-threaded for mobile compatibility.
- Time cost: 4 iterations – Higher number of iterations to compensate for single threading.
- Domain-specific salt to avoid collisions with other domains or applications

KDF

When deriving cryptographic material from high-entropy vectors of bytes, a Password KDF is not needed, and a classical key derivation function (KDF) is significantly lighter and faster, which is critical for smooth UX, less memory usage, and improved battery life.

For this purpose, Gossip uses the post-quantum HKDF-SHA256 KDF with input length prefixing against length extension attacks and ambiguous encoding between multiple variable-length inputs causing hash collisions.

A domain-specific salt is used when initializing the KDF, and domain-specific information specific to a given output is mixed in when extracting that output.

RNG

Gossip relies on the device's native cryptographically secure random number generator as provided by the operating system. The OS entropy source *must* be secure enough for the cryptography used in Gossip to make sense (e.g., for seeding KDFs, nonces, and key generation).

AEAD

Authenticated Encryption with Associated Data is used to encrypt and decrypt messages given the same key, while detecting any tampering attempt by a man-in-the-middle. Gossip uses AES-256-SIV which is well studied, widely used, has minimal overhead, offers quantum resistance, as well as resistance to nonce reuse.

KEM

Key encapsulation enables interactively deriving a fresh shared secret based on the peer's public key. Gossip uses the Module Lattice-based Key Encapsulation Mechanism (ML-KEM) which is a NIST-standard post-quantum key encapsulation mechanism. The variant in use is ML-KEM 768.

DSA

Digital signature algorithms (DSA) are used to sign and authenticate data. Gossip uses a dual DSA combining a Massa signature (ed25519 EDDSA) which is tried and proven but not quantum-resistant, and Module Lattice-based Digital Signature Algorithm (ML-DSA) which is a fresh NIST-standard post-quantum digital signature algorithm. The ML-DSA variant in use is ML-DSA 65, which is the most common one. Both the Massa and ML-DSA signatures need to be verified on the same data for a signature to be considered verified.

4 Pseudonymous identities

At launch, the Gossip app invites users to choose between:

- Creating a fresh account with locally generated secrets.
- Importing a full state from another device: this is used to migrate all keys and conversations from one device to another.
- If the user is already logged in, the app simply invites them to unlock the encrypted state on the device using a pin or passkey/biometrics.

The diagram below shows how the following key pairs are derived from a user-provided passphrase:

- Massa DSA key pair
- ML-DSA key pair

- ML-KEM key pair

The `user_id` is then derived by Blake3 hashing the combined serialization of all the public keys.

Because the user ID is computed by hashing all the public keys, no user can impersonate another user.

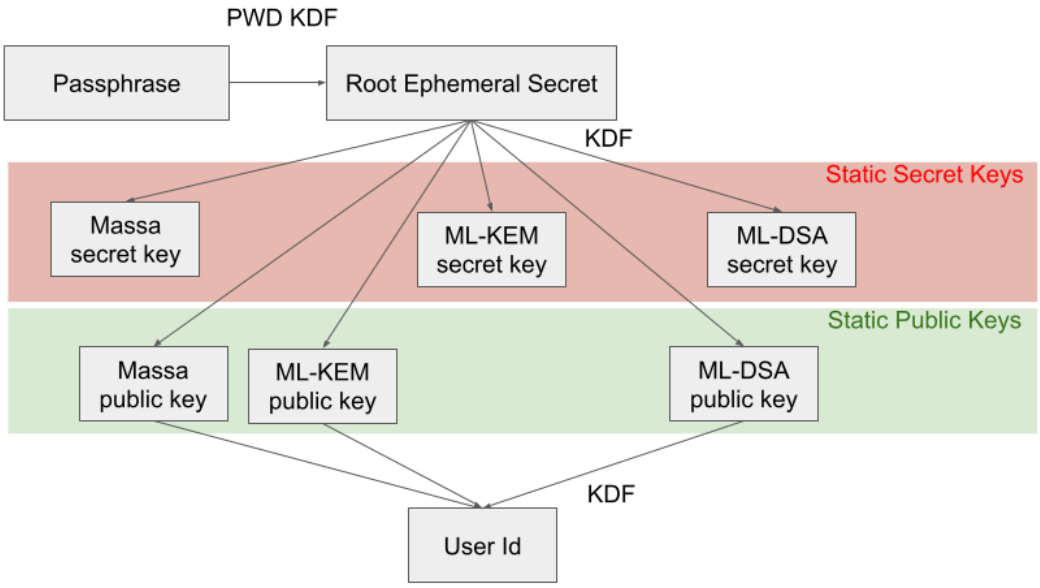


Figure 2: Derivation of pseudonymous identities and keys.

5 Gossip network assumptions

Gossip requires its transport nodes to implement three types of message pools in their cache:

- **Counter-indexed binary blobs:** a pool of encrypted binary blobs indexed only by an ever-increasing counter. The goal here is to be able to query the entries that "we have not queried yet" in order of insertion. Having different nodes index the entries differently depending on the subjective order they have received them in is not an issue: Gossip is robust to it as it rescans from scratch whenever it changes the node it is connected to.
- **DSA-authenticated binary key-value store:** a pool of binary blobs indexed by a Massa address (hash of a Massa public key) concatenated to an arbitrary binary key (max 128 bytes). The value contains the binary blob, the Massa DSA public key and the Massa DSA signature. This allows the reader to verify that the blob is indeed signed by the expected sender, and prevents denial of service by index squatting. Note that this does not break plausible deniability because both sides of the discussion know the ephemeral signing secret. **WARNING:** this feature is not quantum-resistant, but message confidentiality does not depend on it.

- **Hash-indexed binary blobs:** binary blobs indexed by the Blake3 hash of the binary blob. The purpose of this pool is typically to retrieve a set of published public keys based on an `UserID`, because the `UserID` is the hash of those public keys. The reader can easily verify that the returned data hashes to the queried hash in order to ensure data integrity and authenticity.

Binary blobs are limited to 10 megabytes and nodes are required to keep them for 2 weeks. Gossip does not trust those nodes for message validity or integrity, and the nodes are blind to the contents of encrypted messages.

Gossip does not assume anything else about the underlying network, allowing it to work on central servers as well as on P2P mesh networks.

6 Agraphon: post-quantum secure messaging algorithm

Overview

Gossip uses a custom algorithm called Agraphon to achieve secure messaging with the following properties:

- post-quantum security for authentication and encryption
- perfect forward secrecy and post-compromise security
- plausible deniability

Agraphon operates in two steps: announcements to establish a session with a peer, then messaging with the peer.

Announcements

Starting an Agraphon session between Alice and Bob requires both Alice and Bob to have exchanged mutual announcements. We assume that Alice and Bob already know each other's public keys and user IDs.

The process Alice's Gossip app follows to generate an announcement towards Bob's Gossip app is the following:

1. First, Alice's app computes (with `randomness` being 32 bytes of freshly generated public randomness used to seed the KDF):

$$
\begin{aligned}
&\texttt{randomness} = \mathrm{RNG}(32)\\
&(\texttt{kem_ct}, \texttt{kem_ss}) = \mathrm{KEM.encapsulate}(\texttt{bob.kem_pubkey})\\
&(\texttt{cipher_nonce}, \texttt{cipher_key}, \texttt{k_next}, \texttt{auth_pre_key})\\
&\qquad = \mathrm{KDF}\Big(\texttt{randomness}, \texttt{kem_ss}, \texttt{kem_ct}, \texttt{bob.kem_public_key}\Big)\\
&(\texttt{sk_next}, \texttt{pk_next}) = \mathrm{KEM.generate}()\\
&\texttt{auth_key} = \mathrm{KDF}(\texttt{auth_pre_key}, \texttt{pk_next})
\end{aligned}
$$

$$
\begin{aligned}
&\texttt{seeker_seed} = \mathrm{RNG}(32)\\
&\texttt{unix_timestamp_millis} = \mathrm{time.now}()\\
&\texttt{session_init_payload} = \mathrm{serialize}\{\texttt{seeker_seed}, \texttt{unix_timestamp_millis}\}\\
&\texttt{signature} = \mathrm{DSA.sign}(\texttt{alice.secret_keys},\\
&\qquad \mathrm{KDF}\{\texttt{alice.user_id}, \texttt{session_init_payload}, \texttt{auth_key}\})\\
&\texttt{auth_blob} = \{\texttt{alice.public_keys}, \texttt{session_init_payload}, \texttt{signature}\}\\
&\texttt{auth_payload} = \{\texttt{auth_blob}, \texttt{Alice.nickname}\}\\
&\texttt{ciphertext}\\
&\qquad = \mathrm{AEAD.encrypt}\Big(\texttt{cipher_key}, \texttt{cipher_nonce},\\
&\qquad\quad \mathrm{plaintext} = \mathrm{serialize}\{\texttt{pk_next}, \texttt{auth_payload}\},\ \mathrm{ad} = \{\}\Big)\\
&\texttt{announcement_bytes} = \mathrm{concatenate}\,[\texttt{randomness}, \texttt{kem_ct}, \texttt{ciphertext}]
\end{aligned}
$$

2. Then Alice's app sends `announcement_bytes` towards the counter-indexed gossip cache pool, and safely deletes all data except `k_next` and `sk_next` that is kept in safe storage.

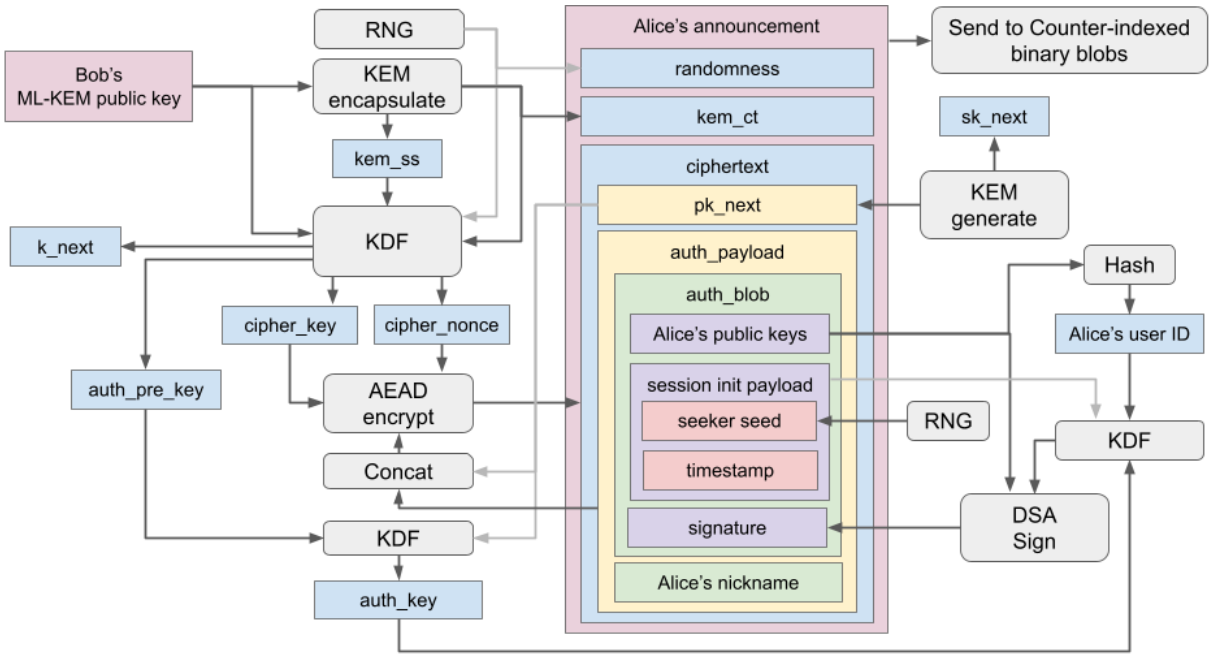


Figure 3: Announcement construction and authenticated payload.

Bob scans the counter-indexed gossip cache pool, starting from the last index he scanned. For each entry he tries the following (ignoring the entry in case of failure):

- decapsulate the `kem_ct` using Bob's KEM secret key
- perform the same initial KDF to deduce the encryption key and nonce
- attempt to decipher the AEAD ciphertext

- extract the origin public keys and signature, checks the signature
- reads the nickname
- checks that the timestamp is more recent than Alice's latest announcement Bob knows about, and not too much in the future either
- reads the seeker seed and `pk_next`
- computes `k_next`

Upon receiving Alice's announcement, Bob's app checks if Bob wants to initiate a session with Alice. And if so, Bob's app also emits a corresponding announcement of the same form towards Alice.

Sessions

Overview

Given a valid bipartite announcement – one announcement from Alice to Bob and another one from Bob to Alice – it becomes possible to initiate a Session between Alice and Bob.

Consider a session between Alice and Bob. The central component of the Session system is the message history. Each message needs to be attached to the two latest messages the sender has seen, one on each side (Alice's and Bob's).

From the standpoint of Alice, the Session contains two histories:

- one history of messages Alice sent, with the following data:
  - the seeker key that acts as a unique identifier of the message
  - the message height on Alice's side (message counter)
  - `sk_next` which is the secret ML-KEM key associated with the `pk_next` in the message
  - `k_next` as derived in the KDF within the message
- one history of messages Bob sent and Alice received, with the following data:
  - height of Alice's parent that the message is linked to
  - the `pk_next` of the message
  - the `k_next` of the message

Bob keeps a similar but mirrored history.

There are rules as to how to choose parents:

- Ancestry must be consistent: for a message M, Alice can't choose a parent on Bob's side that is older (in terms of ancestry) than the one chosen by M's parent on Alice's side. In other terms, M's parent on Bob's side p_{bob} can only be a descendant of a message that is the ancestry of M's parent on Alice's side p_{alice} (including p_{alice} itself). The same applies to Bob's side.
- Either side is not allowed to "fork": two messages on a given side can't have the same parent on that same side.

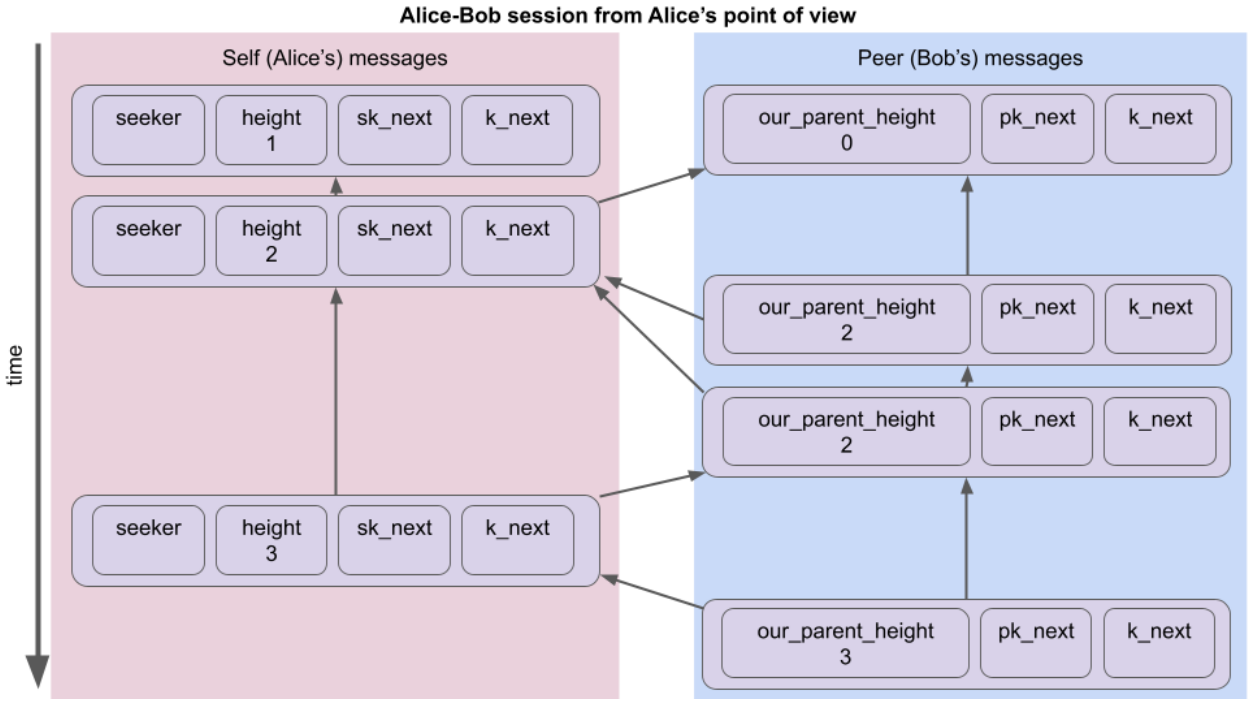


Figure 4: Alice–Bob session view and parent links over time.

- Latency is allowed: Alice might receive Bob's messages late and reference older Bob's messages as parents in the meantime.

The very first messages on either side are the announcements (height 1) sent by that side.

Seekers

Contrary to the announcements which are sent to the Counter-indexed binary blob pool and require systematic scanning, messages are sent to a separate pool: the DSA-authenticated binary key-value store. There, each non-announcement message is indexed by a *seeker* formed by a Massa address concatenated to a constant tag indicating it is a standard message. Those messages are signed by the public key associated with that address in order to avoid front-running and impersonation.

The Massa keypair that yields Alice's first non-announcement message seeker is derived using:

Massa.derive_keypair(KDF([`alice_announcement.seeker_seed`, `bob_announcement.seeker_seed`]))

And for Bob's first message:

Massa.derive_keypair(KDF([`bob_announcement.seeker_seed`, `alice_announcement.seeker_seed`]))

For subsequent messages on each side, the raw Massa keypair together with the message timestamp are added to the message's payload so that the receiver can deduce the seeker key at which to look for the peer's next message.

The fact that both sides always know each other's seeker secret keys ensures plausible deniability.

Message construction and emission

To construct a new message, Alice performs the following steps:

1. Alice first chooses the latest parent message on her own side (`p_alice`) and the latest one she has received from Bob (`p_bob`) as the parent on Bob's side for her new message.
2. Then, Alice computes:

$$
\begin{aligned}
&\texttt{randomness} = \text{RNG}(32)\\
&(\texttt{msg_ct}, \texttt{msg_ss}) = \text{KEM.encapsulate}(\texttt{p_bob.pk_next})\\
&(\texttt{msg_ct_static}, \texttt{msg_ss_static}) = \text{KEM.encapsulate}(\texttt{bob.kem_pubkey})\\
&(\texttt{cipher_key}, \texttt{cipher_nonce}, \texttt{k_next})\\
&\qquad = \text{KDF}\Big(\texttt{randomness}, \texttt{p_alice.k_next}, \texttt{p_bob.k_next},\\
&\qquad\quad \texttt{msg_ct}, \texttt{msg_ss}, \texttt{msg_ct_static}, \texttt{msg_ss_static}\Big)\\
&(\texttt{sk_next}, \texttt{pk_next}) = \text{KEM.generate}()\\
&\texttt{seeker_massa_keypair} = \text{Massa.generate}()\\
&\texttt{payload} = \{\texttt{unix_timestamp_millis}, \texttt{seeker_massa_keypair}, \texttt{content_bytes}\}\\
&\texttt{ciphertext}\\
&\qquad = \text{AEAD.encrypt}\Big(\texttt{cipher_key}, \texttt{cipher_nonce},\\
&\qquad\quad \text{plaintext} = \text{serialize}\{\texttt{pk_next}, \texttt{payload}\},\ \text{ad} = \{\}\Big)\\
&\texttt{message_bytes} = \text{serialize}\{\texttt{randomness}, \texttt{msg_ct}, \texttt{msg_ct_static}, \texttt{ciphertext}\}
\end{aligned}
$$

   where `seeker_massa_keypair` is defined as:

   - `p_alice.seeker_massa_keypair` if `p_alice` is not an announcement;
   - Massa.derive_keypair(KDF(s_a, s_b)) if `p_alice` is an announcement,

   where s_a and s_b are respectively the values of `seeker_seed` carried by Alice's and Bob's announcements.

$$
\begin{aligned}
&\texttt{seeker} = \text{concat}(\text{Massa.derive_address}(\texttt{seeker_massa_keypair}), \texttt{MESSAGE_TAG})\\
&\texttt{seeker_signature} = \text{Massa.sign}\big(\texttt{seeker_massa_keypair},\\
&\qquad \text{hash}(\text{serialize}\{\texttt{seeker}, \texttt{message_bytes}\})\big)\\
&\texttt{binary_blob} = \text{serialize}\{\texttt{seeker_massa_keypair.public_key},\\
&\qquad \texttt{signature}, \texttt{message_bytes}\}
\end{aligned}
$$

3. From there, Alice can post the message on the DSA-authenticated binary key-value store under the key `seeker` and value `binary_blob`. Nodes will authenticate the message before storing and propagating it.

4. Alice then appends the following new entry at the end of her history for her own side:
    - `height` = `p_alice.height` + 1
    - `sk_next`
    - `k_next`
    - `seeker`

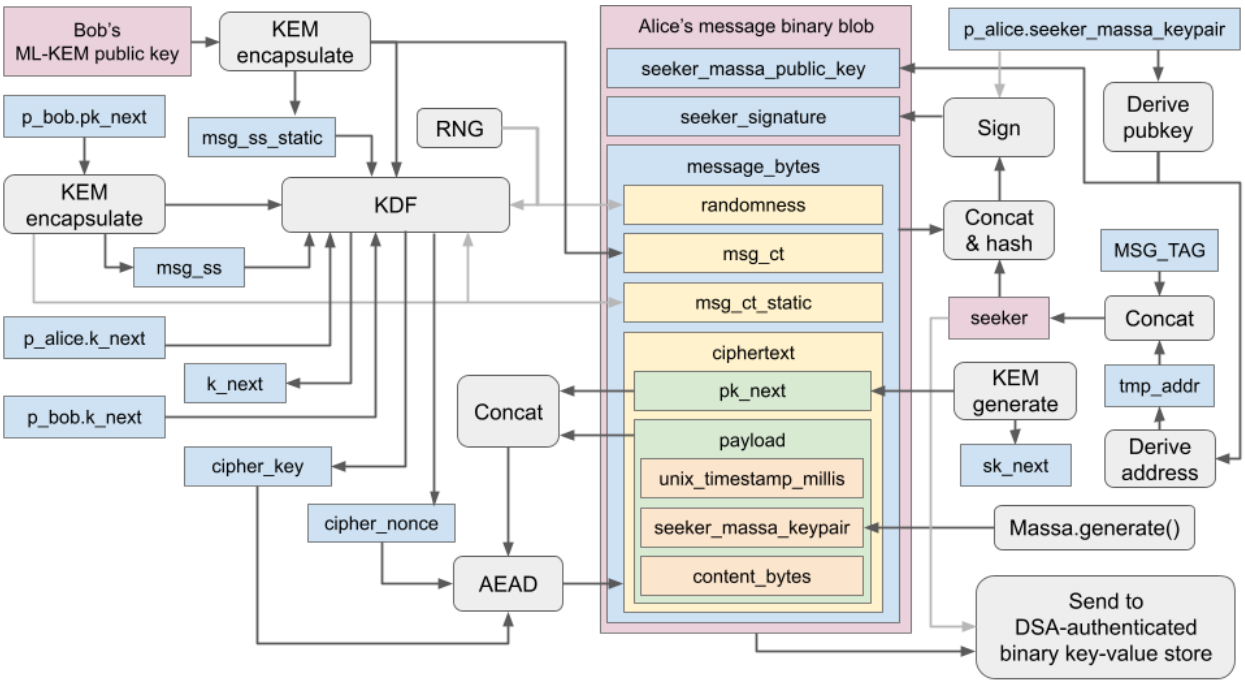


Figure 5: Message construction and emission.

Bob performs similar but mirrored steps when sending a message to Alice.

Message reception and readout

Given the seeker system, Bob always knows at which key to query Alice's next message within the DSA-authenticated binary key-value store, so Bob listens to that key or polls it periodically to read the message.

When Bob receives a DSA-authenticated binary blob, he can autonomously verify the seeker signature. If the seeker signature is wrong, it means that the node Bob is connected to is malicious. Bob switches to another node and retries.

Bob then needs to be able to perform the root KDF to obtain the cipher key. However, the KDF requires Bob knowing `p_bob.k_next`, which implies knowing which one of Bob's messages Alice has chosen as the parent on Bob's side for her message. To solve this, Bob attempts to decipher the message with each one of the allowed parent messages on Bob's side, and stops if one of them succeeds.

On successful decryption, Bob ensures that Alice's message has a consistent timestamp and parent choice on Alice's side. Bob obtains the `seeker_massa_keypair` that will allow him to query the

next message from Alice, without anyone other than Alice being able to know that seeker. Alice and Bob being connected to different nodes ensures that nobody other than Alice and Bob can know they are talking together, nor see any metadata. Note that both Alice and Bob always knowing the seeker secret keys in their mutual conversation helps with plausible deniability despite the presence of digital signatures.

Bob has now successfully extracted the contents of Alice's message, and can therefore append the following entry to his message history on Alice's side:

- `our_parent_height` = `p_bob.height`
- `pk_next`
- `k_next`

Note that in practice, Bob's app must only keep the latest message he has received from Alice in the history on Alice's side and can safely discard any previous messages on Alice's side to minimize storage space, while still saving the message contents and metadata elsewhere in the app for user convenience.

By knowing `p_bob.height` in Alice's message given the parent on Bob side that Alice has chosen, Bob also obtains the information that all of his messages to Alice up until the one of `height = p_bob.height` included have all been received and correctly processed by Alice's app. This is helpful for message delivery status updates. It also has another consequence for optimization: Bob's app can safely discard all of Bob's messages in his history up until the one of `height = p_bob.height` excluded – keeping the message contents and metadata saved elsewhere in the app for user convenience – since they can't be referenced as parents by future messages from Alice given the message parent link rules.

Session management

Sessions ensure perfect forward and backward secrecy: single key leaks do not compromise previous or following messages. Alice keeps at most one open session with each one of her contacts. A session can time out if one side shows no activity for two weeks, if invalid messages are detected, or if there is too high of a lag (eg. Alice's app has sent too many many messages that were never acknowledged by Bob's app). Since stale sessions time out after two weeks, the app regularly needs to send keep-alive messages on active sessions.

If a session dies for any reason, either side can decide to initiate a new session by initiating a new announcement process.

Contact exchange

Sending a contact to a peer can be done simply by sending the contact's static public keys, together with a claimed nickname.

However, post-quantum public keys are very heavy (multiple kilobytes) which makes it inconvenient to share a contact through a direct link or QR-Code for example. To fix this issue, Gossip allows users to advertise their static public keys directly on the Hash-indexed binary blobs index, which allows associating a user ID (the hash of the public keys) to its raw public keys. That way, it

becomes possible to use a 32-byte user ID in order to add a contact, since their public keys can be queried from the gossip cache. However, elements in the cache disappear after 2 weeks without renewal, which requires the app to regularly refresh its entry to avoid it expiring at a moment when this feature is needed.

7 Possible extensions

Latency-resilient parent selection

Motivation

In the protocol described above, when a participant constructs a new message they always reference, as the parent on their own side, the latest message they themselves have sent. While this maximises forward secrecy by ratcheting the chain key forward at every step, it makes the session vulnerable to a single in-transit loss on the sending side: if Alice sends A_1, A_2, A_3 and only A_1 reaches Bob, then Bob cannot derive A_2.`k_next` (which depends on receiving A_2) nor A_3.`k_next`, and any further message from Alice that chains forward from A_3 remains undecipherable until the missing intermediates are recovered.

In practice, the gossip transport layer caches messages for two weeks and allows arbitrary querying, so transient losses are usually masked by repeated polling. However, a more pessimistic deployment scenario – mesh networks, hostile or partitioned transport, or adversarial node sets that drop a fraction of messages – exposes this fragility.

Mechanism

We propose an alternative parent selection rule: when constructing a new message, the sender chooses as its parent on its own side not the latest message they sent, but the latest of their own messages that has been *explicitly acknowledged* as parent by a message they have received from the peer. We refer to this scheme as *latency-resilient parent selection* (LRPS).

Concretely, Alice maintains, in addition to her own message history, the quantity `ack_height`, defined as the height of the latest of Alice's messages that Bob has chosen as parent on Alice's side in any of his messages Alice has received. This quantity is already tracked by the protocol as the `our_parent_height` field of the latest peer message. Under LRPS, Alice's next outgoing message uses the message at height `ack_height` as its parent on Alice's side, rather than the message at the back of her own history.

Bob's behaviour mirrors Alice's, with `ack_height` tracked symmetrically.

Consequences for the parent rule

The original "no fork" rule – two messages on a given side cannot share the same parent on that side – must be relaxed. Under LRPS, several of Alice's messages may share the same parent on Alice's side as long as no later acknowledgement has rotated `ack_height` forward. This forms a *fan* of sibling messages, all rooted at the latest acknowledged point. Once Bob acknowledges any one of

those siblings, `ack_height` advances and the next outgoing message starts a new fan at the new acknowledgement point.

The ancestry-consistency rule (Alice's choice of parent on Bob's side must descend from the corresponding choice in p_{alice}'s message) is unaffected.

Consequences for seeker derivation

The base scheme derives the next seeker on Alice's side deterministically from the previous one in Alice's chain, which means siblings would otherwise compete for the same seeker on the DSA-authenticated key-value store. Two natural fixes:

- *Sibling-indexed seekers.* The seeker for the i-th sibling of a given parent is derived as

  $$\text{Massa.derive_keypair}(\text{KDF}(\texttt{parent.seeker_secret}, \texttt{sibling_index}_i)).$$

  The receiver scans a small bounded range of sibling indices (e.g. the last K_{fork} values) at each parent. The sibling index is also embedded in the message payload so the receiver can confirm and prune.

- *Carry-forward seekers.* Each message carries, in its encrypted payload, the seeker keypair to be used for the next sibling of the same parent (in addition to the next-in-chain seeker keypair already carried). The receiver, upon decrypting any message, learns both the next-in-chain and next-sibling seekers and registers polling for both.

The first option costs the receiver a bounded amount of additional polling per unacknowledged window; the second option costs each message a few additional kilobytes of payload but eliminates the polling overhead.

Security analysis

Forward secrecy. Forward secrecy of past messages is unchanged: each individual message still mixes fresh randomness, a fresh KEM encapsulation against the parent's `pk_next`, and a fresh KEM encapsulation against the peer's static key into its root KDF, and derives its own `k_next`. Compromise of the state held by Alice or Bob *after* the ratchet has advanced past a given message reveals nothing about that message's plaintext.

Post-compromise security. PCS is unchanged in steady state: any successfully delivered message advances both parties' state through fresh KEM material on both the ephemeral and static channels.

Compromise window for unacknowledged tails. The sole new exposure is the lifetime of the parent's `k_next` and `sk_next` in the sender's working state. Under the base protocol, those secrets are consumed by exactly one outgoing child and can in principle be erased shortly after. Under LRPS, they are reused by every sibling message until `ack_height` advances. If an attacker compromises the sender during this window and recovers them, all sibling messages forked from that parent become decryptable in the attacker's hands (rather than at most one). The window's

duration is bounded by the round-trip time before the next acknowledgement arrives; in normal operation this is on the order of seconds. The exposure is also bounded by the cardinality of the sibling fan, which a sender can rate-limit.

Plausible deniability. Plausible deniability is unaffected: both parties continue to hold all message encryption secrets, and there is no signature on the message content – only on its seeker registration, whose secret key both parties also hold.

Authentication and integrity. Per-message AEAD authentication, root-KDF binding to KEM ciphertexts and shared secrets, and seeker DSA signing are all unchanged.

Hybrid deployment strategy

A pragmatic implementation can preserve the maximum-FS guarantee of the base protocol in normal operation and fall back to LRPS only when delivery appears unreliable. For example:

- Use latest-message parenting (the original rule) by default.
- If `self_lag_length` exceeds a configurable threshold, or a send-without-acknowledgement timeout elapses, switch to LRPS for the next message.
- Once a fresh acknowledgement arrives that advances `ack_height`, resume latest-message parenting from the new acknowledgement point.

This bounds the FS-degradation window to actual loss episodes while preserving the optimal ratchet rate during normal exchanges.

Cost summary

- *Sender state*: unchanged; the sender already retains unacknowledged messages in its self history.
- *Receiver state*: unchanged.
- *Network bandwidth*: $O(1)$ extra bytes per message in the carry-forward variant; none in the sibling-indexed variant.
- *Receiver polling*: bounded multiplicative overhead during unacknowledged windows in the sibling-indexed variant; unchanged with carry-forward seekers.
- *Forward secrecy*: degraded only within the unacknowledged window, bounded by the sibling fan size.

LRPS thus offers a controllable trade-off between maximal forward-secrecy ratcheting and resilience to in-transit loss; the hybrid strategy above makes this trade-off automatic and load-adaptive.