

17/WAKU2-RLNRELAY: Privacy-Preserving Peer-to-Peer Economic Spam Protection

Sanaz Taheri Boshrooyeh ✉ 🏠 

Vac Research
Status Research and Development, Singapore

Oskar Thorén ✉ 🏠

Vac Research
Status Research and Development, Singapore

Barry Whitehat ✉

Unaffiliated

Wei Jie Koh ✉ 🏠

Independent

Onur Kilic ✉

Unaffiliated

Kobi Gurkan ✉ 🏠

cLabs

Abstract

This presentation gives an overview of 17/WAKU2-RLNRELAY, a privacy-preserving peer-to-peer economic spam protection mechanism. 17/WAKU2-RLNRELAY is an extension of 11/WAKU2-RELAY i.e., the transport layer of 10/WAKU2 messaging protocol stack. 11/WAKU2-RELAY is, a gossip-based pubsub protocol and is an extension of libp2p GossipSub. In 17/WAKU2-RLNRELAY, we utilize Rate Limiting Nullifiers (RLN) and zkSNARKs to enable this p2p privacy-preserving and economic spam protection mechanism on top of 11/WAKU2-RELAY. 17/WAKU2-RLNRELAY addresses the performance and privacy issues of the state-of-the-art p2p spam prevention techniques including peer scoring (utilized by libp2p), and proof-of-work (used by e.g. Whisper). The spam protection works by limiting the messaging rate of each network participant through rate-limiting nullifiers. To enforce the rate limit, we adopt the suggested framework of Semaphore, however, we modify that framework to properly address the unique requirements of a network of p2p resource-restricted users. The current work dives into the end-to-end integration of Semaphore into 17/WAKU2-RLNRELAY, the modifications required to make it suitable for resource-limited users, and the open problems and future research directions. We also provide a proof-of-concept implementation of 17/WAKU2-RLNRELAY (available in the nim-waku codebase [8]), specifications [22] together with a rough performance evaluation.

2012 ACM Subject Classification Information systems → Spam detection; Networks → Peer-to-peer protocols; Networks → Peer-to-peer networks; Security and privacy → Pseudonymity, anonymity and untraceability; Information systems → Chat; Networks → Routing protocols; Security and privacy → Privacy-preserving protocols

Keywords and phrases P2P, Spam Protection, Messaging, zkSNARKs, Zero-Knowledge, Anonymity, Routing, Pub/Sub, Gossipsub

Digital Object Identifier 10.4230/OASIScs.Tokenomics.2021.1

1 Introduction


10/WAKU2 [18] is a family of peer-to-peer (p2p) protocols for anonymous and privacy-preserving communication. It is designed to be able to run in resource-restricted environments. Being p2p means that 10/WAKU2 relies on no central server. Instead, peers collaboratively



© Jane Open Access and Joan R. Public;
licensed under Creative Commons License CC-BY 4.0

3rd International Conference on Blockchain Economics, Security and Protocols (Tokenomics 2021).

Editors: John Q. Open and Joan R. Access; Article No. 1; pp. 1:1–1:13

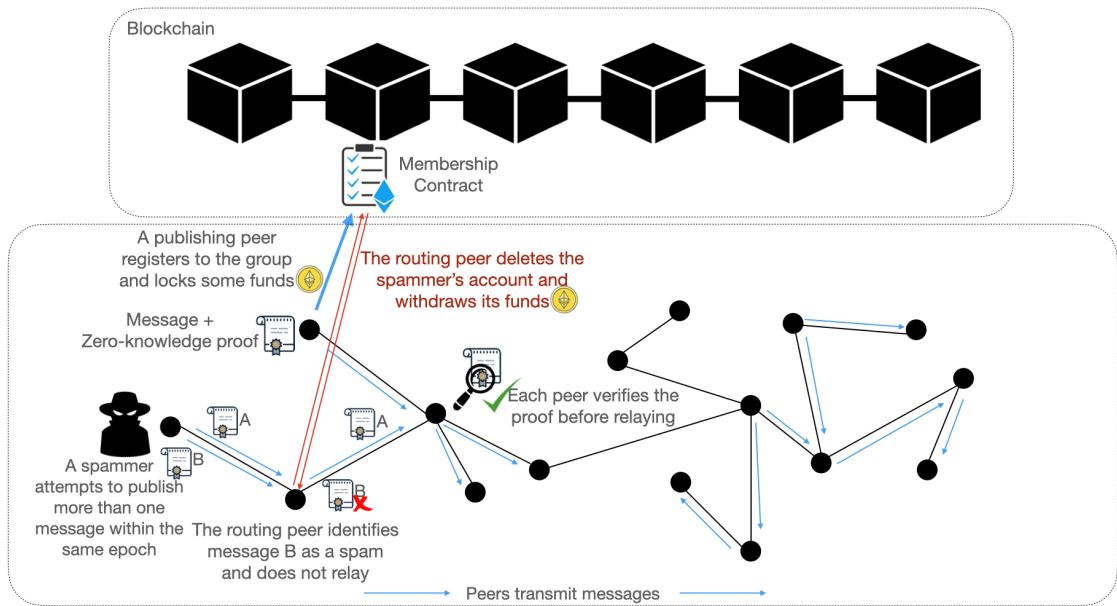


Open Access Series in Informatics

Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

45 deliver messages in the network. 11/WAKU2-RELAY constitutes the transport layer of
 46 10/WAKU2 and aims at being privacy-preserving in which no one knows the owner and
 47 the receiver of a message except the two ends of the communication. Many of the design
 48 choices in this layer are centered around the anonymity requirement. 11/WAKU2-RELAY
 49 is a gossip-based pubsub protocol and a thin layer over the libp2p GossipSub [1] routing
 50 protocol. Its gossip-based structure enables a great level of receiver anonymity [32]. Peers in
 51 11/WAKU2-RELAY congregate around topics they are interested in and can send messages
 52 to topics. Each message gets delivered to all peers subscribed to the topic. Each peer has a
 53 constant number of direct connections/neighbors. To publish a message, the author forwards
 54 its message to a subset of neighbors. The neighbors proceed similarly till the message gets
 55 propagated in the network of the subscribed peers.

56 In addition to 11/WAKU2-RELAY, 10/WAKU2 features other types of protocols for
 57 running in resource-restricted environments. Among which are the request/response protocols
 58 of 12/WAKU2-FILTER [20] which is a light-weight version of 11/WAKU2-RELAY for devices
 59 with limited bandwidth, and 13/WAKU2-STORE [21] by which resourceful peers can persist
 60 and offer historical messages to the querying nodes. Details of these protocols are out of the
 61 scope of this paper and can be found in the Vac RFCs [17].



■ **Figure 1** An overview of privacy-preserving p2p economic spam protection in 17/WAKU2-RLNRELAY protocol.

62 As an open messaging network, 11/WAKU2-RELAY is prone to spam messages. A
 63 spammer usually indicates an entity that uses the messaging system to send an *unsolicited*
 64 message (spam) to *large numbers of recipients*. However, in 11/WAKU2-RELAY with an
 65 open and gossip-based structure, spam messages not only affect the recipients but also all
 66 the other peers involved in the routing process where they have to spend their resources e.g.,
 67 computational power, bandwidth, and storage capacity on processing spam messages. As
 68 such, we define a spammer as an entity that uses the messaging system to publish a large
 69 number of messages in a short amount of time, in other words, has a high messaging rate.

70 We define messages issued in this way as spam. Furthermore, this definition disregards the
71 intention of the spammer as well as the content of the message and the number of intended
72 recipients.

73 The state-of-the-art p2p spam protection techniques for messaging systems i.e., *Proof*
74 *of Work* (POW) [27] deployed by e.g, Whisper [2] and *Peer scoring* [7] method adopted by
75 libp2p are not suitable for resource-constrained environments and also demonstrate privacy
76 issues. The PoW technique imposes a high computational cost for messaging hence devices
77 with limited resources won't be able to participate and benefit from the messaging system.
78 The peer scoring method is prone to censorship and is also subject to inexpensive attacks
79 where the spammer can send bulk messages by deploying millions of bots. The centralized
80 spam protection methods exhibit privacy issues where they usually ask users to disclose and
81 commit to some piece of personally identifiable information e.g., phone number and email
82 address at the registration time. In addition to this, the central provider is aware of messages
83 owned and received by a particular user which is against privacy.

84 The economic-incentive spam protection mechanism of 17/WAKU2-RLNRELAY aims at
85 coping with the aforementioned issues where

- 86 1. It suits **p2p** systems and does not rely on any central entity.
- 87 2. It is **efficient** i.e., with no unreasonable computational, storage, memory, and bandwidth
88 requirement, as such, it fits the network of *heterogeneous* peers with limited resources.
- 89 3. It respects users' **privacy** unlike reputation-based and centralized methods. It also
90 allows identification and removal of spammers while not requiring personally identifiable
91 information from the participants.
- 92 4. It deploys **economic-incentives** to contain spammers' activity. Namely, there is a
93 financial sacrifice for those who want to spam the system. This disincentives spam
94 activity. Additionally, there is financial incentive to monitor the network for spam and
95 remove the offending user.

96 At a high level, 17/WAKU2-RLNRELAY guarantees no one can publish more than
97 a certain number of messages, namely, M messages per epoch, without being financially
98 charged. In addition to financial punishment, the spammer gets removed from the network
99 and will not be able to publish further messages. In 17/WAKU2-RLNRELAY, we set M to
100 1. The epoch can be every second, as defined by UTC date-time $\pm 20s$. The margin of ± 20
101 seconds is to account for the network transmission delay. Nevertheless, the values for M and
102 epoch are not fixed and can be configured according to the application's needs.

103 2 Related Work

104 The studies around spam protection resulted in the development of various techniques and
105 methods. However, each has its own trade-offs and use-case limitations. An overview of the
106 state-of-the-art spam prevention methods is presented below. In this overview, we distinguish
107 between techniques that are targeted for centralized messaging systems and those for p2p
108 ecosystems.

109 2.1 Centralized Messaging Systems

110 In traditional centralized messaging systems, spam usually signifies unsolicited messages sent
111 in bulk or messages with malicious content like malware. Protection mechanisms generally
112 aim at making it expensive and difficult for the user to register or to send bulk messages.
113 Such techniques include:

- 114 ■ Authentication through some piece of personally identifiable information e.g., phone
- 115 number.
- 116 ■ Checksum-based filtering to protect against messages sent in bulk.
- 117 ■ Challenge-response systems.
- 118 ■ Content filtering on the server or via a proxy application.

119 These methods exploit the fact that the messaging system is centralized and a global
120 view of the users' activities is available based on which spamming patterns can be extracted
121 and defeated accordingly. Moreover, users are associated with an identifier e.g., a username
122 which enables the server to profile each user e.g., to detect suspicious behavior like spamming.
123 Such profiling possibility is against user anonymity and privacy and is easier to censor.

124 Among the techniques enumerated above, authentication through phone numbers is a
125 somewhat economic incentive measure as providing multiple valid phone numbers will be
126 expensive for the attacker. Notice that while using an expensive authentication method
127 can reduce the number of accounts owned by a single spammer, cannot address the spam
128 issue entirely. This is because the spammer can still send bulk messages through one single
129 account. For this approach to be effective, a centralized mediator is essential. That is why
130 such a solution would not fit decentralized and p2p environments where no central control
131 exists.

132 2.2 P2P Systems

133 The state-of-the-art p2p spam prevention methods in messaging systems are *Proof of Work*
134 (*POW*) [27] deployed by e.g., Whisper [2] and *Peer scoring* [7] method (namely reputation-
135 based approach) adopted by libp2p. However, each of these solutions has its own shortcomings
136 for real-life use-cases as explained below.

137 2.2.1 Proof of work

138 The idea behind the Proof Of Work (POW) [27] is to make messaging a computationally
139 costly operation hence lowering the messaging rate of all the peers including the spammers.
140 Specifically, the message publisher has to solve a puzzle and the puzzle is to find a nonce
141 such that the hash of the message concatenated with the nonce has at least z leading zeros.
142 z is known as the difficulty of the puzzle. Since the hash function is one-way, peers have to
143 brute-force to find a nonce. Hashing is a computationally heavy operation so is brute force.
144 While solving the puzzle is computationally expensive, it is comparatively cheap to verify
145 the solution.

146 POW is also used as the underlying mining algorithm in Ethereum and Bitcoin blockchain.
147 There, the goal is to contain the mining speed and allow the decentralized network to come
148 to a consensus, or agree on things like account balances and the order of transactions.

149 While the use of POW makes perfect sense in public blockchains such as Ethereum and
150 Bitcoin, it shows practical issues in heterogeneous p2p messaging systems with resource-
151 restricted peers. Some peers won't be able to carry the designated computation and will be
152 effectively excluded. Such exclusion showed to be practically an issue in applications like
153 Status [24], which used to rely on POW for spam protection, to the extent that the difficulty
154 level had to be set close to zero.

155 2.2.2 Peer Scoring

156 The peer scoring method [7] that is utilized by libp2p is to limit the number of messages

157 issued by a peer in connection to another peer. That is each peer monitors all the peers to
158 which it is directly connected and adjusts their messaging quota i.e., to route or not route
159 their messages depending on their past activities. For example, if a peer detects its neighbor
160 is sending more than x messages per month, can drop its quota to $z \cdot x$ where z is less than
161 one. The shortcoming of this solution is that scoring is based on peers' local observations and
162 the concept of the score is defined in relation to one single peer. This leaves room for attacks
163 where a spammer can make connections to k peers in the system and publishes $k \cdot (x - 1)$
164 messages by exploiting all of its k connections. Another attack scenario is through botnets
165 consisting of a large number of e.g., a million bots. The attacker rents a botnet and inserts
166 each of them as a legitimate peer to the network and each can publish $x - 1$ messages per
167 month [10]. Another issue with the peer scoring method is that it is prone to censorship
168 where malicious peers give arbitrary scores to their direct connections to make them look
169 like spammers and prevent their messages from reaching the rest of the network.

170 **3 Preliminaries**

171 **3.1 Semaphore**

172 Semaphore [29] is a zero-knowledge signaling framework on Ethereum. It allows a set of users
173 to broadcast arbitrary signals (where signal is any value like a string, vote, etc.) while proving
174 they are among a group of authorized users without disclosing their identities. Use-cases
175 of Semaphore are anonymous authentication and private voting. It also utilizes external
176 nullifiers to prevent double-signaling. External nullifier can be seen as a voting booth where
177 each user can only cast one vote [23]. Casting a second vote for the same booth will be
178 rejected. Similarly, each signal is bound to an external nullifier, and each group member is
179 allowed to signal only once for that external nullifier. While nullifiers can limit the signal
180 rate per user, there is no way to identify and remove users who violate this rate limit. In an
181 attempt to address this shortcoming, an extended version using Shamir Secret Sharing (SSS)
182 [31] has been proposed. In the extended variant [4], if a user attempts more than one signal
183 for the same external nullifier, it reveals its identity/private key by which it has registered to
184 the group. As such, the identified identity key can be removed from the group and the user
185 won't be able to signal any more. The removed user will be also financially punished. The
186 user initially deposits some fund when joining the group, and the fund will be rewarded to
187 anyone who identifies double signaling of that user.

188 **3.2 Semaphore with Shamir Secret Sharing**

189 In this section, we present an overview of the extended version of Semaphore that utilizes
190 Shamir secret sharing [31] to enable identification of spammers.

191 Each member has a private key sk and a public key $pk = H(sk)$ where H is a Cryptographic
192 hash function. These are also called identity key and the identity commitment,
193 respectively. The list of group members are stored in an *Identity Commitment Tree* which is
194 a Merkle tree whose leaves are members public keys. The root of tree is denoted by τ . The
195 membership of a user in the identity commitment tree is proven by providing a branch of
196 the tree that connects the root to that leaf corresponding to the user's pk . This branch is
197 called *authentication path* and we denote it by $auth$. The tree is stored on a smart contract
198 deployed on the Ethereum blockchain. Each user additionally needs to deposit some funds in
199 the contract at the time of registration.

200 Signaling is bound to a publicly known *external nullifier* denoted by \emptyset . When publishing

201 signal m , a group member utilizes (2,n)-Shamir secret sharing to derive a share (x, y) of its
 202 private identity key sk where $x = H(m)$ and $y = sk + H(sk, \emptyset) * x$. (x, y) will be published
 203 alongside with the signal m .

204 The publishing user also calculates an *internal nullifier* ϕ as $\phi = H(H(sk, \emptyset))$ and
 205 publishes it together with the signal m . Finally, a user needs to prove in zero-knowledge
 206 manner that

- 207 1. its secret key sk belongs to the identity commitment tree namely i.e., provides proof of
- 208 membership.
- 209 2. (x, y) is a valid share of its identity key
- 210 3. the internal nullifier ϕ is correctly calculated

211 The above items are proven through zkSNARKs where the circuit represents the afore-
 212 mentioned constraints. The public inputs to the zero-knowledge proof system are the external
 213 nullifier \emptyset , internal nullifier ϕ , the share of identity secret key (x, y) and the tree root τ
 214 whereas the private inputs provided by the message owner are the identity secret key sk and
 215 the authentication path $auth$. In 17/WAKU2-RLNRELAY, we utilize Groth16[28] for the
 216 proof system. The parameter generation is done through a multi-party setup [30, 16, 6, 14].

217 The user submits its signal m , identity key share (x, y) and internal nullifiers ϕ to the
 218 contract to be stored and accessible by the rest of group members. In fact the contract
 219 maintains the state of all the group signals together with all the metadata that can be used
 220 to identify double signaling. If a user attempts two different signals for the same external
 221 nullifier \emptyset , their internal nullifiers will collide, which signifies a double signaling attempt.
 222 Furthermore, the two shares of the user's identity key can be used to reconstruct its sk and
 223 remove the user from the group. The sk reconstruction stems in the fact that each line
 224 (polynomial of degree 1) can be uniquely reconstructed by using two distinct points on it. In
 225 the case of double signaling, the spammer reveals two distinct shares (x, y) and (x', y') on
 226 line $y = sk + H(sk, \emptyset) * x$, which enables the reconstruction of the line and its evaluation
 227 at $x = 0$ which is sk . The user who reconstructs the sk also gets rewarded by a portion of
 228 the slashed user's stake. This can be done by passing the recovered sk to a function of the
 229 contract.

230 **4 Construction**

231 In this section, we describe the flow of the economic-incentive spam detection mechanism in
 232 17/WAKU2-RLNRELAY from the viewpoint of a single peer. An overview of this flow is
 233 provided in Figure 2.

234 **4.1 Overview of 17/WAKU2-RLNRELAY vs Semaphore**

235 17/WAKU2-RLNRELAY adopts the extended variant of Semaphore which provides the
 236 additional capability of identifying spammers. However, further adjustments are required
 237 to make it fit a p2p routing system. Following is the list of these adjustments that mainly
 238 affect the state of the contract as detailed next.

- 239 1. In 17/WAKU2-RLNRELAY the state of the contract keeps a simple ordered list of users'
 240 identity commitments (instead of Merkle tree) and the Merkle tree is kept off-chain by
 241 individual peers. This is in contrast to the original setting of Semaphore where the
 242 contract holds the entire commitment tree. The reason for this design shift is to mitigate
 243 the significant computational cost/ gas consumption associated with member insertion
 244 and deletion which is logarithmic in the number of registered members. While the cost

245 associated with insertion in Semaphore could be amortized by using batch insertion, this
246 solution is not applicable on the deletion since deletion affects random leaves of the tree
247 which cannot be necessarily batched together. 17/WAKU2-RLNRELAY mitigates this
248 problem as insertion and deletion modify a single item of the list.

249 2. In 17/WAKU2-RLNRELAY, users' messages and their metadata are not stored in the
250 contract which is in contrast to Semaphore. In 17/WAKU2-RLNRELAY messages are
251 stored off-chain and are distributed through the 11/WAKU2-RELAY routing protocol.
252 It has the benefits of 1) having higher message propagation speed and 1) being more
253 economic (messaging is for free) as opposed to the on-chain message store. In the on-
254 chain message storage of Semaphore, published messages will not be visible until blocks
255 containing those message transactions get mined. This results in an unnecessary and
256 undesirable delay which is not acceptable for messaging systems with 1.1 million messages
257 per second [26]. In an attempt to address this issue, 17/WAKU2-RLNRELAY decouples
258 the message propagation and storage from the contract state and provides a p2p and
259 off-chain medium for message transportation i.e., 11/WAKU2-RELAY [19] and storage i.e.,
260 13/WAKU2-STORE [21]. The off-chain storage, adopted by 17/WAKU2-RLNRELAY,
261 has another advantage of being more economic. That is it saves the financial cost (related
262 to the gas consumption) associated with message insertion into the contract state. This
263 cost for one-time messaging scenarios like voting systems can be tolerable, however, it
264 is far from practical in a messaging application where there are millions of messages
265 transmitted per second [26].

266 Due to these adjustments, sending messages in 17/WAKU2-RLNRELAY is for free i.e.,
267 does not need gas consumption. Furthermore, message transmission is not affected by
268 the underlying blockchain and its consensus layer. Such separation allows utilization of
269 various optimization techniques on the message transmission delay which would be otherwise
270 impossible due to reliance on the blockchain.

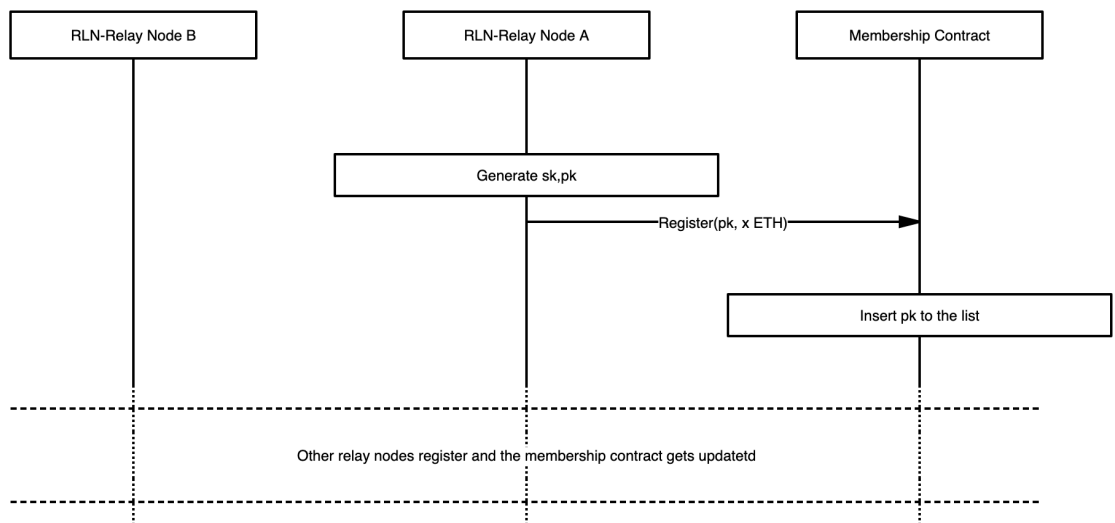
271 4.2 Setup and Registration

272 A peer willing to publish a message is required to register. Registration is moderated through
273 a smart contract deployed on the public Ethereum blockchain. The state of the contract
274 contains the list of registered members' public keys. An overview of registration is illustrated
275 in Figure 2.

276 For the registration, a peer creates its identity private key sk and its commitment
277 $pk = H(sk)$ and sends a transaction to the contract to registers its identity commitment pk
278 in the group. The transaction also transfers the v amount of Ether to the contract. This
279 amount is deposited on the contract to prevent spam activity.

280 4.3 Maintaining the identity commitment tree

281 The construction and maintenance of the identity commitment tree, unlike the original
282 proposal of Semaphore [29], is delegated to the peers. As we previously mentioned, the reason
283 is that the cost associated with member deletion and insertion is high and unreasonable. As
284 such, each peer needs to build the tree locally and listens to the contract's events i.e., peer
285 insertion and deletion and updates its tree accordingly.



■ **Figure 2** Overview of registration process.

286 4.4 External nullifier

287 We use *epoch* as the external nullifier which is defined as the number of T seconds that
 288 elapsed since the Unix epoch. The *epoch* is incremented every T seconds and is calculated
 289 as $epoch = UnixEpoch/T$ where $UnixEpoch$ is the number of seconds since the Unix epoch.
 290 Each peer locally keeps track of the current *epoch*. Peers are allowed to publish one message
 291 per *epoch*.

292 4.5 Publishing

293 Each peer is allowed to send one message m per *epoch* without being slashed/financially
 294 punished. The financial punishment is that the fund deposited by the spammer is rewarded
 295 to the peer reporting spammer's activity. In order to prove that the peer is not a spammer
 296 and has not violated the messaging rate for the current epoch, the peer is required to generate
 297 some metadata and send them alongside the message m . The metadata includes all the
 298 public inputs to the zkSNARKs as explained in 3.2 i.e., a share of the peer's identity secret
 299 key i.e., (x, y) , the internal nullifier ϕ , and the external nullifier $\emptyset = epoch$ and the root of
 300 identity commitment tree τ and zkSNARKs proof π . The peer then sends the message bundle
 301 $(m, (x, y), \phi, epoch, \tau, \pi)$ to its direct connections as instructed by the routing algorithm i.e.,
 302 11/WAKU2-RELAY. An overview of the publishing procedure is provided in Figure 3.

303 Due to the privacy concerns, the publishing peers must always stay updated with the
 304 state of the group (the current members) and make sure they use the latest tree root as well
 305 as the authentication path w.r.t. that root for the proof generation. Using an old tree root
 306 allows inference about the index of the peer's pk in the tree hence its authentication path.

307 4.6 Routing and Slashing

308 Peers follow the regular routing protocol of 11/WAKU2-RELAY and in addition check the
 309 metadata of each relayed message to identify and spot spam messages and slash spammers.
 310 The routing procedure is depicted in Figure 3.

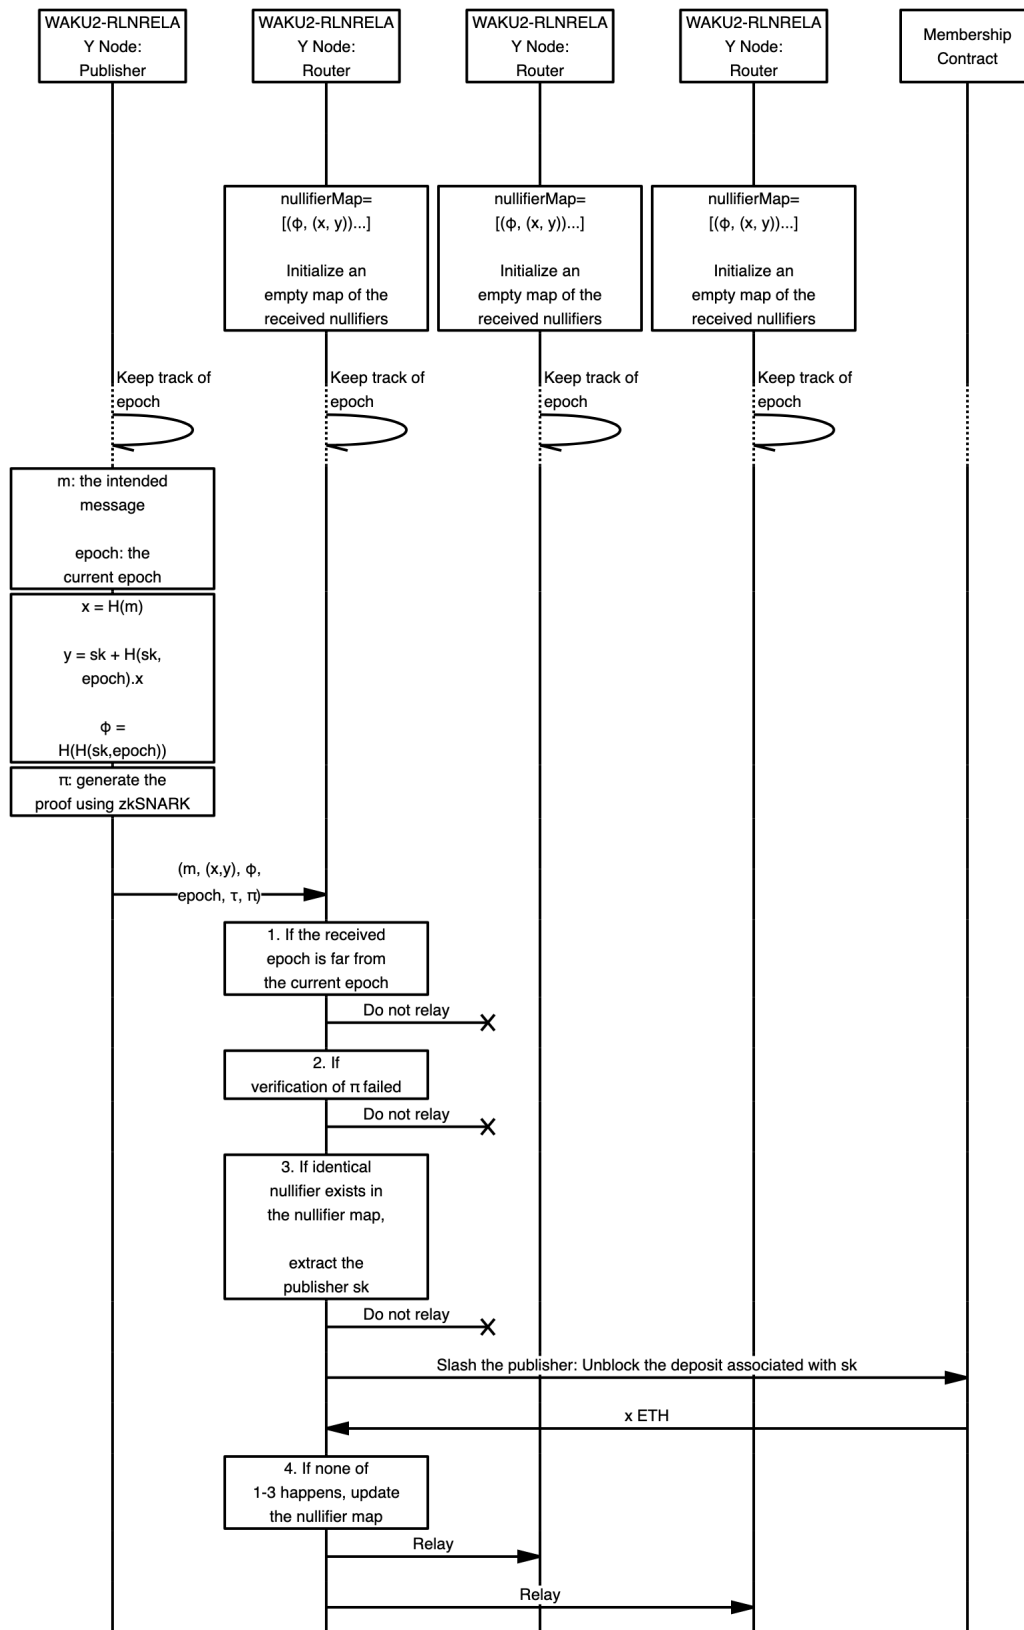


Figure 3 Overview of publishing, routing and slashing.

311 Upon the receipt of a message $(m, (x, y), \phi, epoch, \tau, \pi)$, the routing peer needs to decide
 312 whether to route it or not. The decision relies on the following factors:

- 313 1. If the *epoch* value attached to the message has more than *Thr* gap with the routing peer's
 314 current epoch, the message is considered invalid and must be dropped. This is to prevent
 315 a newly registered peer from spamming the system by messaging for all the past epochs.
- 316 2. The message must contain valid proof π that gets verified by the routing peer.
- 317 3. The messaging rate is no violated.

318 If the preceding checks are passed successfully, then the message is relayed. In case of invalid
 319 proof, the message is dropped. If spamming is detected, the publishing peer gets slashed.

320 In order to identify spam messages, each routing peer keeps a local record of the identity
 321 key share (x, y) and the internal nullifier ϕ of all of its valid incoming message bundles for
 322 the past *Thr* epochs. This list is called nullifier map. The routing peer checks every new
 323 message against this list to spot spam messages. This list does not have to capture the entire
 324 history because any incoming message whose epoch is older than the last *Thr* epochs is
 325 discarded by default (as explained in item 1 of section 4.6). The routing peer utilizes this
 326 list to locally identify spam messages and spammers as follows.

- 327 1. The routing peer first verifies the π and discards the message if not verified.
- 328 2. It checks for the presence of another past message with identical internal nullifier i.e., ϕ .
 329 If nothing is found, then the message gets relayed, otherwise:
 - 330 a. If the identity share i.e., (x', y') of the older message is different from the incoming
 331 message share i.e., $(x, y) \neq (x', y')$, then slashing takes place
 - 332 b. If $(x, y) = (x', y')$, then the message is a duplicate and should be discarded.

333 There is a possibility of race condition to slash an identified spammer. The race happens
 334 when a peer submits the identified identity key of the spammer to the contract in plain format
 335 and at the same time, another peer observes and copies the same information to reclaim the
 336 spammer's fund. This yields a race issue. To avoid this, the commit and reveal technique can
 337 be utilized, i.e., the routing peer sends a commitment to the spammer's identity key (instead
 338 of the plain identity key) to the contract so that no one else can replicate and reclaim it.
 339 Later it opens the commitment and reveals its knowledge of the spammer's identity key.

340 **5 Performance Evaluation**

341 **5.1 Computation overhead**

342 17/WAKU2-RLNRELAY utilizes the RLN library [5] for identity key generation and com-
 343 mitment, Shamir secret sharing, zkSNARKs circuit representation, proof generation, and
 344 verification. The underlying Elliptic Curve is BN254 and the instantiated hash function is
 345 Poseidon with the security level of 128 bits. According to the benchmarking report [15], for a
 346 group size of 2^{32} (Merkle tree depth of 32) the proof generation on an iPhone 8 takes almost
 347 0.5 seconds. This is sufficiently fast for many messaging applications, but may not be low
 348 enough for e.g. real-time communications.

349 **5.2 Storage overhead**

350 A peer has to persist its identity key as well as the prover and verification keys. The former
 351 is of size 32 bytes and the latter is ≈ 3.89 MB [15]. Furthermore, for a group of size N , each
 352 peer has to dedicate $2N - 1 \cdot H$ (H is the hash size (Poseidon hash function)) storage space
 353 to maintain the Merkle tree. Let $N = 2^{32}$ and Hash output be $32 = 2^5$ bytes. Under this

354 setting, the size of the Merkle tree would be $(2 * 2^{32} - 1) * 32 \approx 2^{38} B \approx 2^8$ GB. For smaller
355 size groups e.g., $N = 128$ the tree size drops to almost 8 Kilobytes. In section 6.1, we will
356 discuss potential solutions that can reduce the storage requirement.

357 **6 Conclusion, Future Work and Open Problems**

358 17/WAKU2-RLNRELAY is a routing protocol that features privacy-preserving economic
359 spam protection through rate-limiting nullifiers. The idea is to financially discourage peers
360 from publishing more than one message per epoch. In specific, exceeding the messaging
361 rate results in a financial charge. Those who violate this rule are called spammers and their
362 messages are spam. The identification of spammers does not rely on any central entity. Also,
363 the financial punishment of spammers is cryptographically guaranteed. In this solution,
364 privacy is guaranteed since 1) Peers do not have to disclose any piece of personally identifiable
365 information in any phase i.e., neither in the registration nor in the messaging phase 2) Peers
366 can prove that they have not exceeded the messaging rate in a zero-knowledge manner
367 and without leaving any trace to their identity public keys. Furthermore, it imposes light
368 computational overhead to the routing peers which makes it suitable for resource-limited
369 devices. The proof generation time, as the most recurring and expensive operation in this
370 design, is almost half a second [15] for a group of $2^{32} \approx 4$ billion peers.

371 **6.1 Future Work**

372 17/WAKU2-RLNRELAY is currently a Proof of Concept (POC), and its development is in
373 progress. The following are some of the future work that we would like to peruse further.

374 **Evaluating Merkle tree computation overhead:** We would like to evaluate the
375 running time associated with the Merkle tree operations. Indeed, the need to locally store
376 Merkle tree on each peer was one of the unknowns discovered during this POC and yet the
377 concrete benchmarking result in this regard is not available.

378 **Enhancing performance by off-chain solutions:** Another possible improvement is
379 to replace the membership contract with a distributed group management scheme e.g.,
380 through distributed hash tables. This is to address possible performance issues that the
381 interaction with the public Ethereum blockchain may cause. For example, the registration
382 transactions are subject to delay as they have to be mined before being visible in the state
383 of the membership contract. This means peers have to wait for some time before being able
384 to publish any message. The same issue exists for slashing and other smart contract related
385 functions. The use of state-channels and optimistic Rollups are other ways to overcome
386 transaction delays and achieve better performance.

387 **Lowering the storage overhead per peer:** Currently, peers are supposed to maintain
388 the entire tree locally and it imposes storage overhead which is linear in the size of the group
389 [12]. One way to cope with this is to use the light-node and full-node paradigm in which only
390 a subset of peers who are more resourceful retain the tree whereas the light nodes obtain
391 the necessary information by interacting with the full nodes. Another way to approach this
392 problem is through a more storage efficient method [9] where peers store a partial view of
393 the tree instead of the entire tree. Despite having a partial view, peers are able to construct
394 and update the tree root and their authentication based on the dynamic state of the group.
395 Keeping the partial view lowers the storage complexity to $O(\log(N))$ where N is the size of
396 the group. The use of Verkle tree [25]/ polynomial commitments [3] is also another path to
397 follow. As our future work, we would like to investigate such solutions.

398 **Cost-effective way of member insertion and deletion:** Depending on Ethereum
 399 gas costs, the cost associated with 17/WAKU2-RLNRELAY membership can be more than
 400 30 USD [11]. We aim at finding a more cost-effective approach. For example, using batch
 401 insertion and deletion can lower this cost to almost 15 USD. Another potential solution is
 402 to leverage layer two solutions or to migrate to alternative cost-efficient blockchains with
 403 support for smart contracts.

404 **Evaluating user experience and epoch value:** The usability of a messaging protocol
 405 is heavily influenced by the messaging speed. In the case of 17/WAKU2-RLNRELAY, this is
 406 impacted by the epoch value. While we recommend the epoch value to be 1 second $\pm 20s$
 407 where 20s is the approximated network delay, an empirical analysis is required to measure
 408 network delay as well as to assess the impact of this recommended epoch value on the
 409 messaging speed and user experience.

410 6.2 Open Problems

411 Below is the list of open problems for which no immediate solution is known hence demand
 412 more long-term research.

413 **Exceeding the messaging rate via multiple registrations:** While the economic-
 414 incentive solution has an economic incentive to discourage spamming, we should note that
 415 there is still expensive attack(s) [13] that a spammer can launch to break the messaging rate
 416 limit. That is, the attacker can pay for multiple legit registrations e.g., k , hence being able
 417 to publish k messages per epoch. We believe that the higher the membership fee is, the less
 418 probable would be such an attack, hence a stronger level of spam protection can be achieved.
 419 Following this argument, the high fee associated with the membership, which is discussed in
 420 section 6.1, can indeed be beneficial for spam prevention.

421 **Escaping punishment by early withdrawal:** A spammer can escape from getting slashed
 422 by withdrawing its fund from the contract before its spam activity gets caught. While this
 423 means the attacker burns its initial membership fund (the fee paid to register its key to the
 424 group), it allows saving the other part of the fund that can be otherwise taken by other peers
 425 for slashing.

426 — References —

- 427 1 <https://docs.libp2p.io/concepts/publish-subscribe/>.
- 428 2 <https://eips.ethereum.org/eips/eip-627>.
- 429 3 <https://ethresear.ch/t/open-problem-ideal-vector-commitment/7421/27>.
- 430 4 <https://ethresear.ch/t/semaphore-rln-rate-limiting-nullifier-for-spam-prevention-in-anonymous-p2p-setting/5009>.
- 431 5 <https://github.com/kilic/rln>.
- 432 6 <https://github.com/kobigurk/phase2-bn254/tree/master/powersoftau>.
- 433 7 <https://github.com/libp2p/specs/blob/master/pubsub/gossipsub/gossipsub-v1.1.md#peer-scoring>.
- 434 8 <https://github.com/status-im/nim-waku>.
- 435 9 <https://github.com/vacp2p/research/blob/master/rln-research/merkle-tree-update.md>.
- 436 10 <https://github.com/vacp2p/research/issues/44>.
- 437 11 <https://github.com/vacp2p/research/issues/56>.
- 438 12 <https://github.com/vacp2p/research/issues/57>.
- 439 13 <https://github.com/vacp2p/rfc/issues/251>.
- 440 14 <https://github.com/weijiekoh/perpetualpowersoftau>.
- 441 15 <https://hackmd.io/tmtlmymtr5eynw2lwk9n1w?>

- 444 16 [https://medium.com/coinmonks/announcing-the-perpetual-powers-of-tau-ceremony-to-](https://medium.com/coinmonks/announcing-the-perpetual-powers-of-tau-ceremony-to-benefit-all-zk-snark-projects-c3da86af8377)
445 [benefit-all-zk-snark-projects-c3da86af8377.](https://medium.com/coinmonks/announcing-the-perpetual-powers-of-tau-ceremony-to-benefit-all-zk-snark-projects-c3da86af8377)
- 446 17 [https://rfc.vac.dev/.](https://rfc.vac.dev/)
- 447 18 [https://rfc.vac.dev/spec/10/.](https://rfc.vac.dev/spec/10/)
- 448 19 [https://rfc.vac.dev/spec/11/.](https://rfc.vac.dev/spec/11/)
- 449 20 [https://rfc.vac.dev/spec/12/.](https://rfc.vac.dev/spec/12/)
- 450 21 [https://rfc.vac.dev/spec/13/.](https://rfc.vac.dev/spec/13/)
- 451 22 [https://rfc.vac.dev/spec/17/.](https://rfc.vac.dev/spec/17/)
- 452 23 [https://semaphore.appliedzkp.org/.](https://semaphore.appliedzkp.org/)
- 453 24 [https://status.im/.](https://status.im/)
- 454 25 [https://vitalik.ca/general/2021/06/18/verkle.html.](https://vitalik.ca/general/2021/06/18/verkle.html)
- 455 26 [https://www.oberlo.ca/blog/whatsapp-statistics.](https://www.oberlo.ca/blog/whatsapp-statistics)
- 456 27 Cynthia Dwork and Moni Naor. Pricing via processing or combatting junk mail. In *Annual*
457 *international cryptology conference*, pages 139–147. Springer, 1992.
- 458 28 Jens Groth. On the size of pairing-based non-interactive arguments. In *Annual international*
459 *conference on the theory and applications of cryptographic techniques*, pages 305–326. Springer,
460 2016.
- 461 29 Kobi Gurkan, Koh Wei Jie, and Barry Whitehat. Community proposal: Semaphore: Zero-
462 knowledge signaling on ethereum. 2020. URL: [https://github.com/appliedzkp/semaphore/](https://github.com/appliedzkp/semaphore/blob/master/spec/\Semaphore%20Spec.pdf)
463 [blob/master/spec/\Semaphore%20Spec.pdf.](https://github.com/appliedzkp/semaphore/blob/master/spec/\Semaphore%20Spec.pdf)
- 464 30 Ian Miers. Scalable multi-party computation for zk-snark parameters in the random beacon
465 model. 2019.
- 466 31 Adi Shamir. How to share a secret. *Communications of the ACM*, 22(11):612–613, 1979.
- 467 32 Nik Unger, Sergej Dechand, Joseph Bonneau, Sascha Fahl, Henning Perl, Ian Goldberg, and
468 Matthew Smith. Sok: secure messaging. In *2015 IEEE Symposium on Security and Privacy*,
469 pages 232–249. IEEE, 2015.