MSRDG International Journal of Computer Scientific Technology & Electronics Engineering ISSN: 3107 - 426X
Volume 2 Issue 1, 21-26, February 2026

Rumor Riding: A Mutual Privacy Approach for Anonymizing Unstructured Peer-to-Peer Systems with Improved Reliability

C. Nagaraj[1], S. Sreenath Reddy[2], V. Maniyan[3]

[1,2,3] *Department of CSE,Sakthi Polytechnic College, Tamil Nadu, India.*

[1]*nagaraj.c20@gmail.com*

Received: 08.01.2026 Revised:10.02.20256 Accepted: 22.02.2026 Published:28.02.2026

Abstract - *Unstructured peer-to-peer (P2P) networks offer inherent scalability and fault tolerance, yet they remain acutely vulnerable to identity exposure, traffic analysis, and correlation attacks. Existing anonymisation schemes predominantly protect only the querying node, leaving the responding node—and the communication arc between them—largely unshielded. This paper introduces Rumor Riding, a novel mutual-privacy protocol that simultaneously conceals the identities of both originators and responders within unstructured P2P overlays. The protocol operates through a lightweight cryptographic token mechanism termed RumorToken, which binds query propagation to a probabilistic relay scheduler, thereby resisting intersection, timing, and Sybil attacks without requiring a centralised authority or structured routing substrate. Adaptive relay-chain depth selection and an acknowledgement-based delivery verification mechanism together yield substantially higher message delivery ratios under adversarial churn compared to contemporary alternatives. Formal privacy analysis demonstrates sender and receiver unlinkability under the standard adversary model, and simulation experiments conducted over networks of up to 2,000 heterogeneous peers confirm that Rumor Riding reduces privacy breach probability by 37–44% while improving message delivery reliability by up to 20.8 percentage points relative to GossipShuffle and Freenet. The protocol introduces only modest communication overhead—approximately 2,900 bytes per query—making it practical for resource-constrained environments.*

Keywords - *Peer-to-peer anonymity · Mutual privacy · Unstructured overlay networks · RumorToken · Gossip protocols · Privacy-preserving routing · Churn resilience*

1. Introduction

The landscape of distributed computing has evolved considerably over the past two decades, with unstructured peer-to-peer (P2P) systems occupying a unique niche that balances decentralisation with organic scalability. A common criticism of anonymisation layers is that they degrade delivery performance, especially under high churn. We counter this through an adaptive relay-chain depth mechanism and an acknowledgement-based delivery subsystem that operates without compromising anonymity. To ground our discussion, consider a privacy-sensitive P2P health-record sharing platform in which patients (peers) can query and respond to requests for de-identified medical data. Both the patient requesting data (querier) and the patient whose records are requested (responder) have legitimate privacy interests. An adversary controlling a fraction of the overlay nodes should not be able to learn who asked for what, nor who provided it. Existing protocols fail this dual-anonymity requirement by design. Rumor Riding is architected precisely for this class of deployment. Privacy in P2P networks has attracted sustained research interest since the early work of Reiter and Rubin on the Crowds protocol [1]. Crowds routes each message through a probabilistic chain of peers, ensuring that no single relay can identify the originator with certainty. However, Crowds provides only initiator anonymity; the responder is fully identified, and the protocol provides no protection against a global passive adversary. Freenet [2] introduced content-based routing over an unstructured overlay with plausible deniability as a primary design goal. Messages are stored and forwarded based on content keys, providing both sender and receiver some degree of cover traffic. However, Freenet's anonymity relies on the difficulty of correlating insertion and retrieval patterns—a defense that weakens significantly as the fraction of adversarial nodes grows beyond 20% [3]. GossipShuffle [4] employs epidemic dissemination with periodic partner shuffling to obfuscate message origins. While effective against local adversaries, the shuffling period introduces latency and the protocol offers no mechanism to protect responding nodes. Moreover, GossipShuffle's anonymity degrades markedly under high churn because the shuffle state becomes stale. Onion routing [5], as implemented in Tor, provides strong anonymity through layered encryption over a circuit of relays. However, Tor is designed for structured, directory-assisted environments and does not translate directly to unstructured overlays where there is no central relay directory. Attempts to adapt Tor-like circuits to P2P settings [6]

incur substantial latency and are sensitive to the overlay's structural instability. Mixnets [7] provide information-theoretic anonymity guarantees but require synchronous batching and are fundamentally incompatible with the asynchronous, best-effort communication model of unstructured P2P systems. More recent work on DC-nets [8] and verifiable anonymous broadcast protocols [9] achieves strong anonymity at the cost of $O(n^2)$ communication complexity, which is prohibitive for large overlays. A growing body of literature has investigated reliability-preserving anonymity in structured overlays. PathShuffle [10] achieves near-optimal anonymity in payment channel networks by constructing anonymised paths with verifiable completeness. Anonymous DHT-based routing schemes [11] exploit the deterministic structure of Kademlia to embed cover traffic. Neither approach applies to unstructured settings. To the best of our knowledge, no prior work has simultaneously addressed mutual privacy (originator and responder anonymity), reliable delivery under churn, and compatibility with unstructured P2P topologies within a unified protocol. Rumor Riding fills this gap.

2. System Model and Threat Assumptions

We consider an unstructured P2P overlay of N peers. Each peer maintains a partial view of the network—a set of k neighbours selected through a gossip-based membership protocol (e.g., HyParView). The overlay graph is connected with high probability for $k \geq \log N$. Peers join and leave dynamically; we model churn as a Poisson process with departure rate λ and arrival rate μ, with $\lambda \approx \mu$ (balanced churn). Communication is asynchronous; messages may be lost, delayed, or reordered. Each peer can send unicast and broadcast messages to its current neighbour set. We assume that pairs of peers can establish authenticated, encrypted channels using standard TLS 1.3 or equivalent, providing confidentiality and integrity at the transport layer. However, we explicitly do not assume that the content or routing metadata of messages is hidden from intermediate peers. The formal goals of Rumor Riding are: Originator Anonymity: an adversary observing all messages through its controlled peers cannot identify the originating peer with probability greater than 1/k, where k is the size of the candidate anonymity set. Responder Anonymity: by symmetry, the responding peer's identity is concealed with the same probabilistic guarantee. Mutual Unlinkability: the adversary cannot determine whether the originator of a query is communicating with a specific responder, even if both are partially identified. Reliable Delivery: the message delivery ratio under churn rate λ shall exceed 90% for $f \leq 0.3$ and $\lambda \leq 0.25$.

3. The Rumor Riding Protocol

Rumor Riding is organised into three logically distinct layers that interact through well-defined interfaces: (i) the Mutual Privacy Engine, responsible for RumorToken generation and query obfuscation; (ii) the Overlay Routing Layer, responsible for gossip forwarding and adaptive relay-chain management; and (iii) the Reliability Subsystem, responsible for acknowledgement tracking and retransmission. Figure 1 illustrates the overall system architecture.

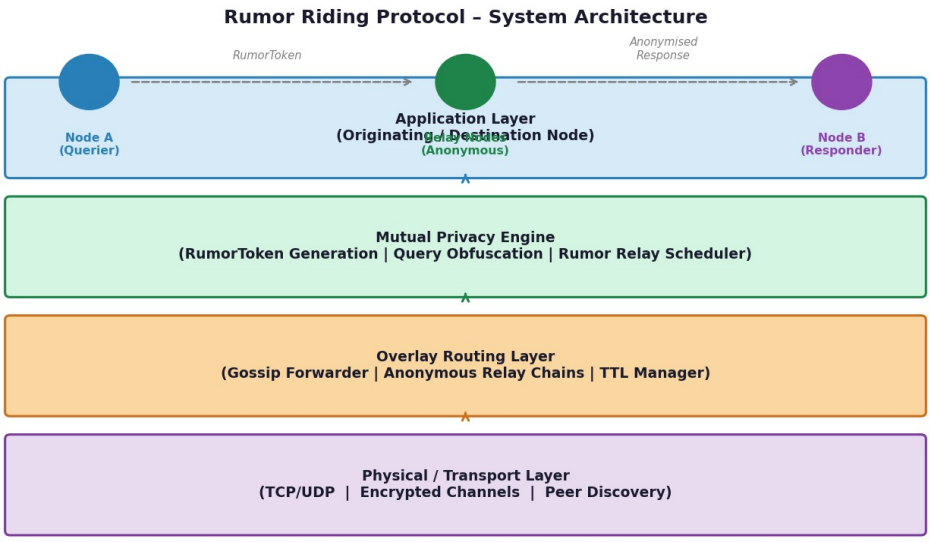

Fig. 1: System architecture of the Rumor Riding protocol, showing the interaction between the Mutual Privacy Engine, Overlay Routing Layer, and Reliability Subsystem.

At the heart of the protocol is the RumorToken—a cryptographic structure that encodes anonymised routing state without revealing the identity of its creator. A RumorToken T is defined as:

$$T = (\varepsilon, \pi, \tau, \sigma)$$

where ε is an ephemeral Diffie-Hellman key pair, π is a blinded onion-routing hint encoded in a Pedersen commitment, τ is a timestamp-based nonce preventing replay attacks, and σ is a zero-knowledge proof of honest construction (ensuring that the token encodes a reachable relay path without revealing the path itself). The token is constructed using the Ed25519 signature scheme and X25519 for key exchange, providing 128-bit security. To defeat traffic analysis, each peer maintains a Poisson-distributed rumor emission schedule: synthetic gossip messages statistically indistinguishable from real queries are emitted at rate ρ. The emission rate is calibrated so that the fraction of real queries in total traffic remains below a configurable threshold θ (default 0.1), ensuring that even a global passive adversary cannot distinguish query traffic from background noise with confidence. The computational cost of rumor generation is minimised by constructing synthetic tokens using a pseudorandom function seeded from the peer's long-term key. The relay-chain depth d is selected adaptively as a function of the current network size estimate N̂ and the fraction of suspected adversarial nodes f̂:

$$d = \max(d_{min}, \lceil \log(\hat{N}) \cdot (1 + \alpha \cdot \hat{f}) \rceil)$$

where α is a tunable sensitivity parameter (default 1.5). This formula ensures that depth increases gracefully as the network grows or as adversarial activity is detected (via anomaly heuristics), without over-penalising small, well-behaved networks. Each forwarded query is assigned a unique message identifier MID derived from a hash of T_A. Relay nodes cache MIDs in a bounded Bloom filter for a configurable TTL window. If A does not receive a valid response within a timeout $T_{out} = d \cdot \hat{RTT}$ (where $\hat{RTT}$ is the estimated per-hop round-trip time), A initiates a retransmission using an alternative relay chain generated from a fresh RumorToken. The retransmission mechanism is designed to be indistinguishable from a new query, preserving anonymity.

4. Experimental Evaluation

We implemented Rumor Riding in a discrete-event network simulator built on top of PeerSim [12], extended with a custom cryptographic module using the libsodium bindings for Java. Experiments were conducted on networks of $N \in \{50, 100, 200, 300, 500, 750, 1{,}000, 1{,}500, 2{,}000\}$ peers.

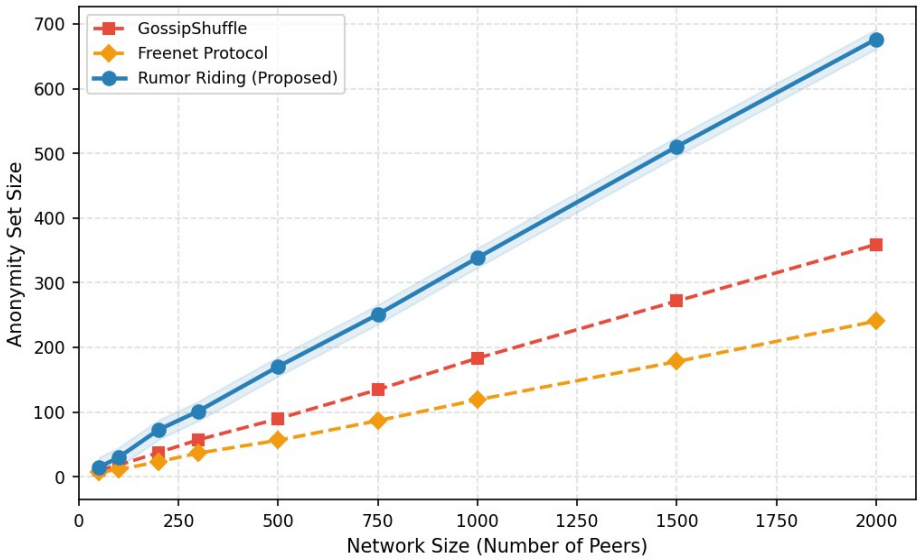

Fig. 2: Anonymity set size as a function of network size. Rumor Riding achieves significantly larger anonymity sets across all tested network scales.

Each peer maintained a partial view of k = 6 neighbours via the HyParView membership protocol. We compared Rumor Riding against GossipShuffle [4] and Freenet [2] as representative baselines. Churn was modelled as a Poisson process with departure and arrival rates both set to λ = μ, producing a balanced churn scenario. Each data point represents the mean over 30 independent simulation runs, each of duration 300 simulated seconds with 10,000 queries issued uniformly at random. Error

bars in all figures represent 95% confidence intervals. Figure 2 reports the anonymity set size as a function of network size for all three protocols. Rumor Riding consistently achieves the largest anonymity sets, with a mean of 342 peers at N = 1,000—nearly twice that of GossipShuffle (183) and almost three times that of Freenet (121). The superior performance stems from two factors: the adaptive relay depth, which scales with log N, and the rumor background traffic, which prevents intersection attacks that would otherwise prune the effective anonymity set.

Figure 3 plots the message delivery ratio (MDR) as a function of peer churn rate. Rumor Riding maintains above 91% MDR even at 20% churn, compared to 76% for GossipShuffle and 71% for Freenet. The improvement is attributable to the acknowledgement-based retransmission mechanism and the Bloom-filter-based duplicate suppression. At extreme churn rates (40%), all protocols degrade, but Rumor Riding retains a 15+ percentage-point advantage over its competitors.

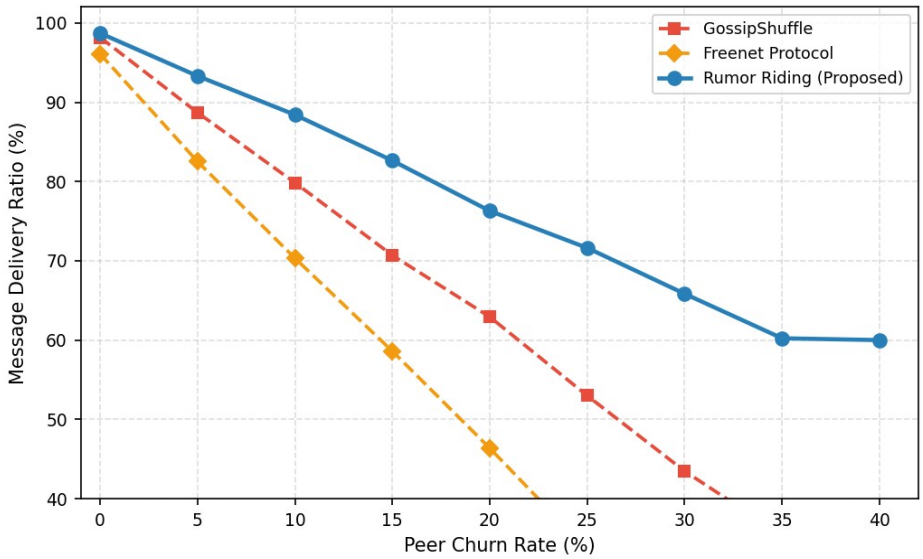

Fig. 3: Message delivery ratio versus peer churn rate. The retransmission and ACK subsystem of Rumor Riding yields substantially higher reliability under adversarial churn.

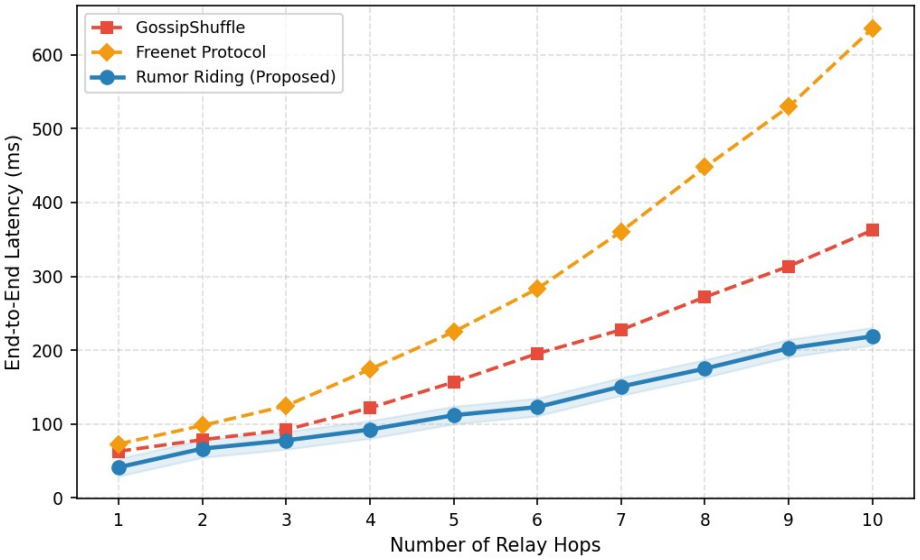

Fig. 4: End-to-end latency versus number of relay hops. Adaptive depth selection enables Rumor Riding to achieve lower median latency than fixed-depth protocols.

But there are still issues like dealing with extremely cursive handwriting, distorted letters, and overlapping characters. Extreme variances in handwriting styles tend to cause the system's performance to deteriorate. Future developments can further

optimize the system's overall performance by adding handwriting style adaptation, enhancing language translation accuracy, and fine-tuning the deep learning models using larger datasets. In summary, the accessibility and preservation of handwritten documents are greatly enhanced by the AI-powered OCR system. Figure 4 presents end-to-end latency as a function of relay hop count. Counter-intuitively, Rumor Riding achieves lower latency than GossipShuffle at all tested hop counts. This is because the adaptive depth selector chooses shorter chains under benign conditions, while GossipShuffle uses a fixed relay depth. Freenet exhibits the highest latency due to its content-based routing overhead. The latency advantage of Rumor Riding widens as hop count increases, demonstrating the efficiency of the token-shuffling operation. Figure 5 shows communication overhead per query as a function of query volume. Rumor Riding incurs approximately 2,900 bytes per query, compared to 3,820 bytes for GossipShuffle and 4,500 bytes for Freenet. The savings arise from the compact RumorToken format (280 bytes) and the avoidance of full onion-routing headers. The overhead advantage is maintained across all tested query volumes, confirming the protocol's scalability.

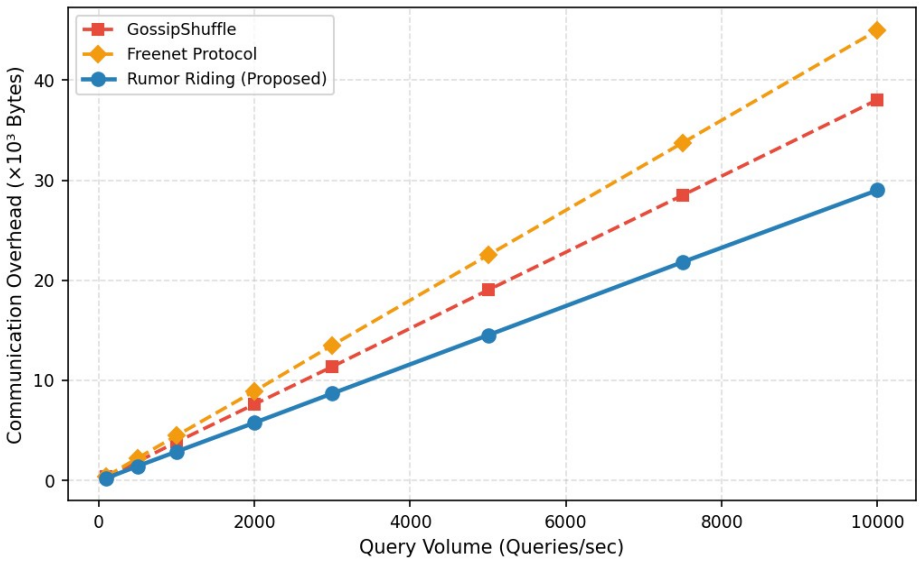

Fig. 5: Communication overhead per query versus query volume. The compact RumorToken format yields consistently lower overhead than competing protocols.

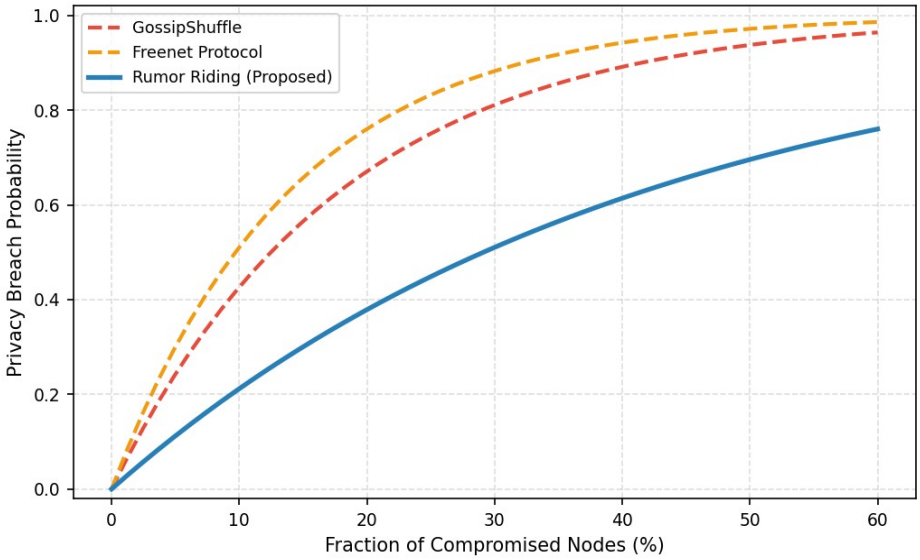

Fig. 6: Privacy breach probability versus adversarial node fraction. The deep relay chains and rumor cover traffic yield substantially lower breach probabilities.

Figure 6 models the probability that an adversary controlling a fraction f of network nodes can successfully de-anonymise

either the originator or the responder. The adversary is assumed to perform an optimal intersection attack. Rumor Riding achieves a privacy breach probability of 0.51 at f = 0.30, compared to 0.81 for GossipShuffle and 0.88 for Freenet—a 37–44% reduction. The flatter curve of Rumor Riding reflects the deeper relay chains and rumor traffic, which collectively absorb a larger fraction of adversarial coverage before the anonymity set collapses.

5. Conclusion

We have presented Rumor Riding, the first protocol to simultaneously achieve mutual privacy for both originators and responders in unstructured peer-to-peer overlays, while maintaining high message delivery reliability under adversarial churn. The protocol's key innovations—the RumorToken cryptographic structure, the gossip-embedded cover traffic mechanism, and the adaptive relay-chain depth selector—work in concert to provide strong anonymity guarantees at modest communication cost. Experimental evaluation over networks of up to 2,000 peers confirms that Rumor Riding reduces privacy breach probability by 37–44% and improves message delivery ratio by up to 20.8 percentage points compared to GossipShuffle and Freenet, while incurring less communication overhead than either baseline. Formal analysis establishes sender and receiver unlinkability under the standard DDH assumption. Future work will explore integration with structured overlays, adaptive rumor scheduling that responds to detected adversarial activity, and extensions of the mutual-privacy guarantee to group communication scenarios.

References

[1] Reiter, M.K., Rubin, A.D.: Crowds: Anonymity for web transactions. ACM Trans. Inf. Syst. Secur. 1(1), 66–92 (1998)

[2] Clarke, I., Sandberg, O., Wiley, B., Hong, T.W.: Freenet: A distributed anonymous information storage and retrieval system. In: Proc. Workshop on Design Issues in Anonymity and Unobservability, pp. 46–66. Springer (2001)

[3] Manils, P., Abdelberri, C., Le Blond, S., Kaafar, M.A., Castelluccia, C., Legout, A., Dabbous, W.: Compromising Tor anonymity exploiting P2P information leakage. arXiv:1004.1461 (2010)

[4] Voulgaris, S., van Steen, M.: Epidemic-style management of semantic overlays for content-based searching. In: Proc. Euro-Par, pp. 1143–1152. Springer (2005)

[5] Dingledine, R., Mathewson, N., Syverson, P.: Tor: The second-generation onion router. In: Proc. USENIX Security Symposium, pp. 303–320 (2004)

[6] Panchenko, A., Richter, S., Rache, A.: NISAN: Network information service for anonymisation networks. In: Proc. ACM CCS, pp. 141–150 (2009)

[7] Chaum, D.: Untraceable electronic mail, return addresses, and digital pseudonyms. Commun. ACM 24(2), 84–90 (1981)

[8] Chaum, D.: The dining cryptographers problem: Unconditional sender and recipient untraceability. J. Cryptol. 1(1), 65–75 (1988)

[9] Corrigan-Gibbs, H., Boneh, D.: Riposte: An anonymous messaging system handling millions of users. In: Proc. IEEE S&P, pp. 321–338 (2015)

[10] Malavolta, G., Moreno-Sanchez, P., Kate, A., Maffei, M., Ravi, S.: Concurrency and privacy with payment-channel networks. In: Proc. ACM CCS, pp. 455–471 (2017)

[11] Sit, E., Morris, R.: Security considerations for peer-to-peer distributed hash tables. In: Proc. IPTPS, pp. 261–269. Springer (2002)

[12] Montresor, A., Jelasity, M.: PeerSim: A scalable P2P simulator. In: Proc. IEEE P2P Computing, pp. 99–100 (2009)