Poster: TopoShot: Uncovering Ethereum's Network Topology Leveraging Replacement Transactions

(presenting a full paper accepted to IMC 2021)

Kai Li
kli111@syr.edu
Syracuse University
Syracuse, NY, USA

Yuzhe Tang✉
ytang100@syr.edu
Syracuse University
Syracuse, NY, USA

Jiaqi Chen
jchen217@syr.edu
Syracuse University
Syracuse, NY, USA

Yibo Wang
ywang349@syr.edu
Syracuse University
Syracuse, NY, USA

Xianghong Liu
xliu317@syr.edu
Syracuse University
Syracuse, NY, USA

ABSTRACT

Ethereum relies on a peer-to-peer overlay network to propagate information. The knowledge of Ethereum network topology holds the key to understanding Ethereum's security, availability, and user anonymity. However, an Ethereum network's topology is stored in individual nodes' internal routing tables, measuring which poses challenges and remains an open research problem in the existing literature.

This work presents TopoShot, a new method uniquely repurposing Ethereum's transaction replacement/eviction policies for topology measurement. TopoShot can be configured to support Geth, Parity and other major Ethereum clients. As validated on local nodes, TopoShot achieves 100% measurement precision and high recall (88% ∼ 97%). To efficiently measure the large Ethereum networks in the wild, we propose a non-trivial schedule to run pair-wise measurements in parallel. To enable ethical measurement on Ethereum mainnet, we propose workload-adaptive configurations of TopoShot to minimize the service interruption to target nodes/network.

We systematically measure a variety of Ethereum networks and obtain new knowledge including the full-network topology in major testnets (Ropsten, Rinkeby and Goerli) and critical sub-network topology in the mainnet. The results on testnets show interesting graph-theoretic properties, such as all testnets exhibit graph modularity significantly lower than random graphs, implying resilience to network partitions. The mainnet results show biased neighbor selection strategies adopted by critical Ethereum services such as mining pools and transaction relays, implying a degree of centralization in real Ethereum networks.

1 INTRODUCTION

A blockchain system relies on an underlying peer-to-peer (P2P) network to propagate information including recent transactions and blocks. Knowing the topology of P2P network is essential to understanding blockchain's availability under network partitions, its security against denial of service attacks (e.g., eclipse attacks [10]), and its protection of user anonymity [6, 11]. The value has motivated a line of measurement studies on the network topology of popular blockchains including Bitcoin [8, 9] and Monero [7]. However, although Ethereum is the second largest blockchain network (after Bitcoin) and the biggest smart-contract platform, measuring Ethereum's network topology remains an open research problem.

Measurement methods: In this work, we propose TopoShot to measure an Ethereum blockchain overlay by repurposing Ethereum's transaction replacement and eviction policies. Briefly, an Ethereum node buffers unconfirmed transactions (prior to mining) in a local data structure named `mempool`, where an unconfirmed transaction can be replaced or evicted by a subsequent transaction at a sufficiently higher Gas price. Transaction replacement and eviction are standard Ethereum features, widely supported by Ethereum clients (including Geth [2], Parity [5] and others [1, 3, 4]), and highly desirable by real-world applications. For instance, a common practice in blockchain-based decentralized applications is that a user having sent a transaction can posthumously speed up its inclusion into the blockchain by sending replacement transactions at higher price per computation unit (or the so-called Gas price). Leveraging these features, TopoShot runs a measurement node M to detect the connection between two remote nodes A and B. In TopoShot, node M propagates a high-priced transaction tx_A on target node A, a low-priced

[*] ✉ Yuzhe Tang is the corresponding author.

transaction tx_B to target node B, and a medium-priced transaction tx_C propagated to all other nodes in the same network. It then observes tx_A's presence on node B and, if so, draws the conclusion that node A is actively connected to node B. To ensure the accurate measurement, when node A is not linked to node B, measurement transaction tx_A should not be propagated and do not reach node B (the so-called "isolation" property [8]). One of the key insights in this work is that Ethereum's transaction replacement policy can be repurposed to enforce isolation property for accurate link measurement. Intuitively, the isolation is ensured by the fact that TopoShot's high-priced tx_A can replace the low-priced tx_B on node B but not the medium-priced tx_C on other nodes, through which tx_A cannot be propagated to reach node B.

To set up the measurement as above, TopoShot further leverages Ethereum's support of transaction eviction and future transactions, that is, to evict an existing unconfirmed transaction on a node by incoming future transactions (the concept of future transaction in Ethereum is similar to orphan transactions in Bitcoin). Specifically, in TopoShot, the measurement node M first propagates tx_C to all nodes, then sends future transactions to evict tx_C (with other existing transactions) on node A and B before sending tx_A and tx_B to node A and B, respectively.

For large-scale measurement on real Ethereum networks, we propose a non-trivial method to parallelize multiple pairwise measurements, reducing the rounds and overall time of measurement.

Contributions: This work makes the following contributions:

• *Novel methods*: We propose a novel method, named TopoShot, to measure Ethereum network links and topology. TopoShot takes a unique approach by exploiting Ethereum's handling of unconfirmed transactions (i.e., transaction replacement and eviction). TopoShot is generic and supports all Ethereum clients (including Geth and Parity). TopoShot is effective and achieves 100% result precision and high recall (88% ∼ 97%).

• *Large-scale measurements*: We address the scalability and ethical challenges raised in measuring large-scale, real Ethereum networks. We propose to schedule pair-wise measurements in parallel for efficiency. We propose workload-adaptive mechanisms to configure TopoShot for minimal service interruption on the target nodes/network.

• *New systematic results*: Without TopoShot, an Ethereum network's topology remains hidden information inside blackbox Ethereum nodes, measuring which stays an open research problem. By systematically conducting measurements against a variety of Ethereum networks, we obtain a series of new knowledge on network topology and its graph-theoretic statistics, ranging from full-network topology in popular testnets (Ropsten, Rinkeby and Goerli) and critical sub-network topology in the mainnet.

REFERENCES

[1] Retrieved May, 2021. Aleth – Ethereum C++ client, tools and libraries. https://github.com/ethereum/aleth.

[2] Retrieved May, 2021. Geth: the Go Client for Ethereum. https://www.ethereum.org/cli#geth.

[3] Retrieved May, 2021. Hyperledger Besu. https://www.hyperledger.org/use/besu.

[4] Retrieved May, 2021. Nethermind Ethereum client. https://nethermind.io/client.

[5] Retrieved May, 2021. Parity Ethereum is now OpenEthereum: Fast and feature-rich multi-network Ethereum client. https://www.parity.io/ethereum/.

[6] Alex Biryukov, Dmitry Khovratovich, and Ivan Pustogarov. 2014. Deanonymisation of Clients in Bitcoin P2P Network. In Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security, Scottsdale, AZ, USA, November 3-7, 2014, Gail-Joon Ahn, Moti Yung, and Ninghui Li (Eds.). ACM, 15–29. https://doi.org/10.1145/2660267.2660379

[7] Tong Cao, Jiangshan Yu, Jérémie Decouchant, Xiapu Luo, and Paulo Veríssimo. 2020. Exploring the Monero Peer-to-Peer Network. In Financial Cryptography and Data Security - 24th International Conference, FC 2020, Kota Kinabalu, Malaysia, February 10-14, 2020 Revised Selected Papers (Lecture Notes in Computer Science, Vol. 12059), Joseph Bonneau and Nadia Heninger (Eds.). Springer, 578–594. https://doi.org/10.1007/978-3-030-51280-4_31

[8] Sergi Delgado-Segura, Surya Bakshi, Cristina Pérez-Solà, James Litton, Andrew Pachulski, Andrew Miller, and Bobby Bhattacharjee. 2019. TxProbe: Discovering Bitcoin's Network Topology Using Orphan Transactions. In Financial Cryptography and Data Security - 23rd International Conference, FC 2019, Frigate Bay, St. Kitts and Nevis, February 18-22, 2019, Revised Selected Papers. 550–566. https://doi.org/10.1007/978-3-030-32101-7_32

[9] Matthias Grundmann, Till Neudecker, and Hannes Hartenstein. 2018. Exploiting Transaction Accumulation and Double Spends for Topology Inference in Bitcoin. In Financial Cryptography and Data Security - FC 2018 International Workshops, BITCOIN, VOTING, and WTSC, Nieuwpoort, Curaçao, March 2, 2018, Revised Selected Papers (Lecture Notes in Computer Science, Vol. 10958), Aviv Zohar, Ittay Eyal, Vanessa Teague, Jeremy Clark, Andrea Bracciali, Federico Pintore, and Massimiliano Sala (Eds.). Springer, 113–126. https://doi.org/10.1007/978-3-662-58820-8_9

[10] Ethan Heilman, Alison Kendler, Aviv Zohar, and Sharon Goldberg. 2015. Eclipse Attacks on Bitcoin's Peer-to-Peer Network. In USENIX Security 2015, Washington, D.C., USA, Jaeyeon Jung and Thorsten Holz (Eds.). USENIX Association, 129–144. https://www.usenix.org/conference/usenixsecurity15

[11] Philip Koshy, Diana Koshy, and Patrick D. McDaniel. 2014. An Analysis of Anonymity in Bitcoin Using P2P Network Traffic. In Financial Cryptography and Data Security - 18th International Conference, FC 2014, Christ Church, Barbados, March 3-7, 2014, Revised Selected Papers (Lecture Notes in Computer Science, Vol. 8437), Nicolas Christin and Reihaneh Safavi-Naini (Eds.). Springer, 469–485. https://doi.org/10.1007/978-3-662-45472-5_30