Miner Dynamics on the Ethereum Blockchain*

Antonio Fernandez Anta[1], Paul Rimba[2], Andres Abeliuk[3], Manuel Cebrian[4], Ioannis Stavrakakis[1,5], An Binh Tran[2], Oluwasegun Ojo[1], Iyad Rahwan[4]

[1]) IMDEA Networks Institute, Madrid, Spain
antonio.fernandez@imdea.org
[2]) Data61, CSIRO, Sydney, Australia
[3]) Information Sciences Institute, University of Southern California, USA
[4]) Scalable Cooperation Group, MIT Media Lab, USA
[5]) National and Kapodistrian University of Athens, Greece

The recent rise of Bitcoin, and blockchain–its underlying technology–has ushered in a wave of other cyptocurrencies and blockchain applications. One such application is the Ethereum, a decentralized software platform that rose to prominence for enabling smart contracts and distributed applications. The Ethereum blockchain is powered by Ether, a platform-specific token used for transactions and paying for services on the network. Like in Bitcoin, miners try to solve cryptographic puzzles in order to get rewards, in this case, Ethers. However, experience has shown that the returns on solo mining have a high variance i.e. returns are unpredictable. Therefore, miners come together to form mining pools with power equivalent to the sum of the power of each member miner. This guarantees more predictable earnings, albeit likely lower, since earnings are shared among miners in the pool according to an agreed scheme.

Miners are allowed to migrate between mining pools as they deem fit which has led to a lot of dynamics between pools. This raises the question about what factors influence miner dynamics, like migrations. Can we understand these factors and their effect on the decentralization of the system? Are miners leaving, joining and changing pools mainly to maximize earnings and returns on their investments (on computing power)? Are there colluding miners or pools? As part of our efforts to answer some of these questions, we obtained data from the Ethereum blockchain system with records of about the 60 most prominent pools in the system (which account for more than 90% of the power), their miners, reward payouts, and miner movements including those that join, leave, or switch pools. These records are broken down into time windows each of four weeks.

Initial analysis seems to suggest a preferential attachment and detachment (when miners decide to join or leave a pool) with fairly strong correlations between pool power and the percentage of miners migrating from and to other pools across all time windows. Furthermore, distributions of relative miner power (inferred from the percentage of payouts they get) in some time windows seem to follow power-law distribution while over other time windows follow a lognormal distribution. To understand miner migrations between pools, we build a network from the payouts records with pools as nodes and miners as edges. Hence, two pools i and j are connected with an edge of weight n if n miners have received a payout from both pools i and j at any point in time. This network revealed a few number of well-connected pools with lots of common miners between them. This may indicate that these pools are the major players in the system, but a lot of common miners may also indicate colluding or related pools, i.e. pools that are seemingly different but are operated by the same entity. Using four different network clustering/community detection methods, namely: fast-greedy, walktrap, edge betweenness, and leading eigenvalues, reveal a community of about 30 pools consistently identified across 3 of the 4 community detection methods. Furthermore, pools that are known to be operated by the same entity are usually clustered in the same community.

As a continuation of this work, we expect to come up with a model that adequately explains the miner dynamics, miner attachments and detachments i.e. how they choose pools to join or leave. Furthermore, we hope to continue to explore the network aspects of this work to gain more insights into the Ethereum ecosystem.

*Partially supported by the Spanish MICIU grant DiscoEdge (TIN2017-88749-R), the Comunidad de Madrid grant EdgeData-CM (P2018/TCS-4499), and the NSF of China grant 61520106005.