An alternative approach to analyze anonymity in cryptocurrencies

Motivation and Overview

The goal of our proposed project is to research how techniques inspired from the area of differential privacy (DP) can be used in order to construct anonymous cryptocurrencies without the need to rely on a trusted setup process.

Our motivation rises from the fact that Monero, one of the most popular private cryptocurrencies that does not require a trusted setup, has been recently empirically analyzed to find that approximately 80% of the transactions provide no or very limited privacy [1]. Monero utilizes ring signatures in order to conceal the source of a transaction. In a nutshell, when a user spends a coin, she first finds other unspent transactions with the same denomination, and uses the public keys corresponding to those transactions to create her ring. The user's identity will be hidden within the set of PKs that are part of the ring. However, the number of Monero mixins is usually rather small and sampled in a way that can be distinguished from the real transaction [1].

Our plan is to investigate whether we can build a protocol for mixing transactions in the spirit of Monero (i.e., utilizing ring confidential transactions), while formally proving that we preserve differential privacy for the users. Specifically, we would like to claim that two neighboring transaction graphs are nearly equally likely to give rise to the same chain. In order to keep the size of each individual ring signature small while providing a large number of potential options for the real transaction, our plan is to have users submit several ring signatures in a sequence of rounds that rope in an ever-increasing number of possible mix-ins.

Technical approach

We propose to construct a cryptosystem that preserves differential privacy with respect to the following definition of neighboring transactions, although we stress that fine-tuning the definition to support our protocol design is within scope.

Initial assumptions: We make two simplifying assumptions at the beginning of this project. Our intent is to remove both of them by the end of the research endeavor. First, we assume that there are n senders and n receivers, with each sender transfering money to exactly one receiver (which also implies that a transaction cannot merge multiple inputs or create multiple outputs). While we hope to weaken this assumption after our initial study, we note that even under this assumption, the result would be useful. In particular, one could use such a mixing protocol for refreshing wallets, with the n senders each send their money to themselves (as classic tumblers do). Our second assumption is that every transfer is of the exact same value, which for simplicity we assume to be 1 coin; in particular, we assume that the senders do not require change transactions.

Definitions: We define a transaction table, π, as a 1-1 mapping from the senders to the receivers, and we say that two transaction tables are *neighboring* if they differ in *exactly one* pair of recipients (i.e., by one transposition). Formally, for senders $(s_1, ..., s_n)$ and receivers $(r_1, ..., r_n)$ we say that π_1 and π_2 are neighboring transaction tables if there exist i and j s.t.:

$$\pi_1(s_i) = \pi_2(s_j),\ \pi_1(s_j) = \pi_2(s_i),$$

and for all k where $k \neq i,\ k \neq j,\ \pi_1(k) = \pi_2(k)$.

Next, we say that a cryptocurrency protocol is (ε, δ)-differentially private if every pair of neighboring transaction tables π_1 and π_2 yield nearly-identical blocks posted to the blockchain. Specifically, let B_1 denote the probability distribution over all blocks that can arise from the transaction graph π_1 within the protocol. Then, we require that

$$\Pr[B_1 \mid \pi_1] \leq e^{\varepsilon} \cdot \Pr[B_1 \mid \pi_2] + \delta$$

for all neighboring π_1 and π_2. We stress that this definition is intended to hide who a sender paid, but not the fact that the sender produced a transaction in the first place. (The sender may choose to pay herself on occasion to obscure when she is paying someone else. We leave the formalization of what is hidden by such actions to future work.)

Warm-up protocol: Let $(s_1, ..., s_n)$ and $(r_1, ..., r_n)$ be the set of senders and receivers, respectively. Let each participant in the protocol create a ring signature over the entire group of n participants. These ring signatures yield a graph of real + mix-in transactions that looks like a complete bipartite graph from senders to receivers, which perfectly hides the true transaction table of which sender paid which receiver. However, this protocol has (at least) two big drawbacks: it scales poorly in n and it requires all senders to know the set of all receivers ahead of time.

High-level idea of our protocol: Instead, we will construct a protocol where the senders must only rope in a (small) constant number of mix-ins with each signature, and where they must only know the identities of other senders (but not of the receivers). The basic idea is to have a transmission protocol that operates in R rounds. In each round, every sender creates a ring signature to send her funds to her own ephemeral sybil-account, roping in other senders randomly as mix-ins. Only in the final round do the senders finally transmit money to the receivers. Our intent is for the composition of the first $R-1$ rounds (from the senders to the sybil-senders) to be a type of tumbling network, such that with high probability all transaction tables will be almost equally likely. Ergo, the R^{th} round payments to the receivers cannot be traced back to the original senders. Note that we do not need R blocks posted in the blockchain for our protocol to be implemented, instead we can include all of our protocol transactions in a single block.

Challenges: We will work towards building a protocol that will (a) provide the right graph structure to be analyzed for differential privacy and (b) will be efficient enough for practical

implementation. We expect that a rather challenging step in our approach will be the security analysis of the protocol, as it involves studying the differential privacy of graphs in a new setting, where each individual contribution influences multiple nodes and edges. Finally, if time permits, we plan on working on a proof of concept implementation of the proposed protocol and providing detailed comparisons with other private cryptocurrencies.

Extensions: While it is common for research into anonymous cryptocurrencies (e.g., TumbleBit) to start with simplifying assumptions akin to ours, as a stretch goal we will attempt to reduce or remove our dependency on the assumptions stated above. First, we believe that RingCT should compose nicely with our envisioned protocol in order to remove the assumption that all users pay the same amount. Second, once we have a rigorous analysis for the base case, we will see how we can adjust the protocol to account for transactions involving multiple senders or multiple recipients (including change transactions). We note that this will require a change to the notion of neighboring transactions as well, albeit in a way that might actually provide stronger anonymity: with multiple recipients per transactions, we can use the 'remove a transaction' notion of differential privacy rather than the 'swap the recipients of two transactions' variant.

Team background and qualifications

Our team consists of a group of cryptographers whose expertise spans many topics that are essential for the current proposal: differential privacy, MPC, and building anonymity solutions for cryptocurrencies.

Foteini Baldimtsi (George Mason University)
Ethan Gertler (George Mason University)
Dov Gordon (George Mason University)
Mayank Varia (Boston University)

Evaluation plan

Our proposed approach will be evaluated in the form of providing formal cryptographic definitions and proofs to showcase the exact level of security and privacy that we achieve. On the practical side, we will provide a proof of concept implementation to showcase the level of efficiency that we achieve.

Security considerations

This project will advance knowledge on the flavors of anonymity that a private cryptocurrency can achieve and will explore the efficiency - privacy trade-offs.

Budget and justification

We request $25K and plan to spend the requested budget towards graduate student stipends, equipment expenditures and research visits among the research team members.

Schedule
(2 months) Work on new DP based definitions for privacy in cryptocurrencies
(3 months) Construction and proofs
(1 month) Discuss trade-offs and comparisons with existing solutions and if time permits provide a proof of concept implementation.

Email address(es) for direct contact
Foteini at gmu dot edu

References
[1] Malte Möser, Kyle Soska, Ethan Heilman, Kevin Lee, Henry Heffan, Shashvat Srivastava, Kyle Hogan, Jason Hennessey, Andrew Miller, Arvind Narayanan, and Nicolas Christin, “An Empirical Analysis of Traceability in the Monero Blockchain”, to appear in PETS 2018