

Christoph Egger, Russell W. F. Lai, Viktoria Ronge, Ivy K. Y. Woo, and Hoover H. F. Yin

On Defeating Graph Analysis of Anonymous Transactions

Abstract: In a ring-signature-based anonymous cryptocurrency, signers of a transaction are hidden among a set of potential signers, called a ring, whose size is much smaller than the number of all users. The ring-membership relations specified by the sets of transactions thus induce bipartite transaction graphs, whose distribution is in turn induced by the ring sampler underlying the cryptocurrency.

Since efficient graph analysis could be performed on transaction graphs to potentially deanonymise signers, it is crucial to understand the resistance of (the transaction graphs induced by) a ring sampler against graph analysis. Of particular interest is the class of partitioning ring samplers. Although previous works showed that they provide almost optimal local anonymity, their resistance against global, e.g. graph-based, attacks were unclear.

In this work, we analyse transaction graphs induced by partitioning ring samplers. Specifically, we show (partly analytically and partly empirically) that, somewhat surprisingly, by setting the ring size to be at least logarithmic in the number of users, a graph-analysing adversary is no better than the one that performs random guessing in deanonymisation up to constant factor of 2.

Keywords: anonymous cryptocurrencies, ring signatures, random directed graph connectivity

DOI 10.56553/popets-2022-0085

Received 2021-11-30; revised 2022-03-15; accepted 2022-03-16.

1 Introduction

In many anonymous systems, a main cryptographic component for providing anonymity is a linkable ring signature (LRS) scheme [8], which is a signature scheme with a restricted anonymity guarantee. The goal of this work is to study the resistance of these systems against graph-based deanonymisation attacks. For concreteness,

we will use privacy-preserving cryptocurrencies as a running example of anonymous systems based on LRS schemes. We emphasise, however, that the techniques introduced in this work are also directly applicable to other applications of LRS, e.g. anonymous voting [15, 23]¹.

1.1 Linkable Ring Signatures

We begin by recalling the basics of LRS schemes. To sign a message μ , e.g. a transaction in a cryptocurrency, the signer first samples a *ring* r , i.e. a set consisting of (the public keys of) the signer itself and decoys, using an external algorithm known as a *ring sampler*, then uses the LRS scheme to produce a signature σ . The tuple (r, μ, σ) is communicated to the verifiers, e.g. by publishing it on the blockchain in the context of cryptocurrencies. In applications, it is common for a human user to own many pairs of public and secret keys. Nevertheless, to simplify terminologies, we will refer to a public key as a “user” and use the notation U to refer to the set of users (public keys) where signers belong to and where decoys are sampled from. Depending on the application, the set U could grow over time.

An LRS scheme is *linkable* in the sense that there exists an efficient public algorithm to determine whether any two given signatures are generated by the same signer, i.e. using the same secret key. Applications of LRS schemes often employ a “single-sign verification rule” which only accepts new signatures which are not linked to any previously accepted ones, so that each user can only perform certain anonymous action once. For example, such anonymous action could be spending a coin in a cryptocurrency, casting a vote in an anonymous voting system, authenticating and redeeming scores in an anonymous credential system, etc. In general, the single-sign rule ensures that, at any time, each signer in the set of users has at most one signature that is accepted by the verifiers.

Christoph Egger, Russell W. F. Lai, Viktoria Ronge:

Friedrich-Alexander-Universität Erlangen-Nürnberg

Ivy K. Y. Woo: Independent

Hoover H. F. Yin: The Chinese University of Hong Kong

¹ Although these schemes as described include all legitimate voters in rings, using smaller rings is more efficient. The analyses provided in this work allow designers to make an informed decision of how smaller ring sizes could be chosen.

The *anonymity* of an LRS scheme guarantees that the tuple (r, μ, σ) leaks no more information (computationally) about the signer creating σ than what is leaked by the ring r sampled by the ring sampler. Typically, for efficiency reasons, the ring size $|r|$ is much smaller than the number of users $|U|$, making it plausible to deanonymise signers just by observing ring membership relations implied by the set of published rings, regardless of how secure the LRS scheme is.² It is therefore important to design ring samplers and choose their parameters in a way that strikes a balance between efficiency and anonymity, which is the central topic of this work.

1.2 Transaction Graphs

To understand deanonymisation attacks of the above kind, we model ring membership relations by *transaction graphs*. Specifically, consider an application of LRS where, at some point in time, the tuples $\{(r_j, \mu_j, \sigma_j)\}_{r_j \in R}$ are accepted by the verifiers, where R is some set of rings and, for all $r_j \in R$, members of the ring r_j were sampled from U . The ring membership relations can be represented by a transaction graph, which is a bipartite graph G with vertex sets U and R , and $u_i \in U$ is connected to $r_j \in R$ if user u_i is a member of ring r_j . Figure 1 is a toy example of a transaction graph consisting of 3 users and 3 rings.

A transaction graph is guaranteed to have a maximum matching involving the vertex set R . Indeed, by the unforgeability of the LRS we can assume that σ_j was issued by some signer $u_j \in r_j$ for each $r_j \in R$, and by the linkability of the LRS and the single-sign verification rule we can assume that all u_j 's are distinct. This means that the set $\{(u_j, r_j)\}_{r_j \in R}$ is a maximum matching.

A transaction graph could have many maximum matchings, each representing a possible assignment of signatures/rings to signers. The union of all maximum matchings of G is known as the Dulmage-Mendelsohn (DM) decomposition [3] or simply the *core* $\text{Core}(G)$ (in the sense of DM), and can be computed in linear time given G [20]. If an edge $(u_i, r_j) \in G$ does not belong to any maximum matching, i.e. $(u_i, r_j) \notin \text{Core}(G)$, then

user u_i cannot have been the signer creating σ_j . Consequently, the signature-signer assignments represented by the edges $G \setminus \text{Core}(G)$ can be ruled out given the knowledge of $\text{Core}(G)$. In extreme cases, where a user u_i is connected to only a single ring r_j in $\text{Core}(G)$, the user u_i is considered completely deanonymised.




Fig. 1. Toy example of transaction graph. Edges correspond to ring memberships, e.g. (u_1, r_3) means user 1 is a member of ring 3. The red edges are the only maximum matching.

Referring to the example in Figure 1, upon knowing that the only maximum matching is $\{(u_j, r_j) : j = 1, 2, 3\}$, in other words $G \neq \text{Core}(G)$ and $G \setminus \text{Core}(G) = \{(u_1, r_3), (u_2, r_3)\}$, all three signers can be deanonymised. Note how user u_3 is deanonymised due to the memberships of the other two rings, although its ring consists of three members. We see that the anonymity of a signer does not only depend on its own ring, but also on the other rings. A global view on the transaction graph is thus required to properly assess the anonymity of signers.

Another richer and more realistic example is given later in Figure 2 (Page 544), which shows a transaction graph G with 8 users and 7 rings, and all rings consist of more than one member. On computing $\text{Core}(G)$, we see that 4 out of the 19 potential signature-signer assignments can be ruled out, and one of the signers, namely user 4, can be completely deanonymised.

1.3 Graph-Based Deanonymisation

Generalising the attack illustrated in Figures 1 and 2, we consider graph-based deanonymisation attacks, where an adversary attempts to identify the signer who sampled $r_{j^*} \in R$, for some j^* chosen by the adversary, given only a transaction graph G representing all rings R . In particular, we consider adversaries which do not attempt to break the LRS scheme and which do not have knowledge about the signing probabilities of the signers. The former is easily justified since the LRS scheme is supposedly cryptographically secure. The latter is sensible

² For example, at the time of writing, Monero mandates a ring size of $|r| = 11$ and has a number of public keys $|U| \geq 16 \times 10^6$. We note that we are considering anonymity at the key level, which is a stronger notion than anonymity at the human user level typically considered for non-anonymous cryptocurrencies such as Bitcoin. Indeed, even if all spenders in Monero transactions are deanonymised, receivers would still be cryptographically anonymous due to the “stealth address” mechanism.

when the signing probabilities are (close to) uniform by heuristics, e.g. when using a partitioning ring sampler to be discussed in Section 1.4. Finally, our security model capturing untargeted attacks is strong, since if an adversary is successful in a targeted attack, then it is also successful in an untargeted one.

A trivial attack strategy is to choose the smallest ring $r_{j^*} \in R$ and output one of the ring members uniformly at random, which has success probability of exactly $1/|r_{j^*}|$. We therefore want to upper-bound the success probability of any graph-analysing adversary such that it is not much greater than that of the trivial strategy. Our strategy is to show that the success probability of an adversary is at most $\Pr[G \neq \text{Core}(G)]$ greater than that of the trivial strategy mentioned above, where G is a transaction graph induced by the ring sampler of interest, i.e. the best non-trivial strategy that an adversary could use is to perform DM decomposition.

Although DM decomposition is a well-known tool in graph theory, the technique seems to be adopted only recently to analyse anonymous cryptocurrencies [21], where it is shown that the analytical deanonymisation attack on Monero based on DM decomposition is at least as effective as existing attacks [7, 12, 25] of the same nature. Indeed, this is as expected since all existing attacks are graph-based and the signature-signer assignments ruled out by these attacks could also be found by DM decomposition. However, a broader understanding of graph-based attack on ring samplers appears to be lacking. In particular, the previous examples of attack lead us to the question: *How should rings be chosen such that the success probability of a graph-based attack can be upper-bounded?*

1.4 Partitioning Samplers

Of particular interest are the partitioning samplers [19], which first publicly partition the set of users into chunks, randomly choose k decoys from the chunk that the signer belongs to, and output the set which contains the signer and the k decoys as the ring. Assuming that for each chunk the signing probabilities of the signers in the chunk are close to each other, a partitioning sampler provides near-optimal *local* anonymity according to an entropy-based measure [19], which we discuss further in both Section 1.7.2 and Appendix A. Furthermore, in the extreme case that all chunks of the partition are of size $k + 1$ – equal to the ring size – then the induced transaction graph G simply consists of disjoint $(k + 1)$ -bicliques and $G = \text{Core}(G)$ trivially. Despite having these features,

little is known about the *global* anonymity, e.g. the resistance against graph analysis, of partitioning samplers for general chunk sizes.

1.5 Our Contributions

In this work, we study the resistance of ring samplers against graph-based deanonymisation attacks. More precisely, let $\mathcal{G}^{\text{Samp}}$ be the distribution of transaction graphs induced by a ring sampler Samp . We derive an upper bound of $\Pr_{G \leftarrow \mathcal{G}^{\text{Samp}}}[G \neq \text{Core}(G)]$ by relating the event $G \neq \text{Core}(G)$ to that of certain digraphs induced by G being not strongly connected. In case Samp is a partitioning sampler, we show that this probability likely upper-bounds the advantage of any adversary performing graph-based deanonymisation attacks.

Specifically, assuming two conjectures on certain distributions of random directed graphs (digraphs) which we support by providing empirical evidence, we show that if the number of decoys k of a partitioning sampler is set to

$$k \geq \ln(2 \cdot |U|) + \sqrt{2 \ln(2 \cdot |U|)},$$

then $\Pr_{G \leftarrow \mathcal{G}^{\text{Samp}}}[G \neq \text{Core}(G)] \leq \frac{1}{k+1}$. In other words, a graph-analysing attack is at most twice as successful as a trivial attack does.

Since graph-based attacks threaten all decoy-based anonymous systems, such as coin-mixing, mix-nets, and voting, not limited to LRS-based cryptocurrencies, our result is broadly applicable: It serves as a guideline for choosing parameters for all such systems to avoid graph-based deanonymisation attacks.

1.6 Technical Overview

For the ease of reading the technical sections, we provide a high-level overview below.

1.6.1 Transaction Graphs and Induced Digraphs

The central objects studied in this work are transaction graphs and their induced digraphs, which are formally defined in Section 2. As described in Section 1.2, a transaction graph is a bipartite graph G with vertex sets (U, R) and edges E . For any transaction graph G , suppose without loss of generality that $M = \{(u_j, r_j)\}_{j=1}^m$ is a maximum matching in G . We can define its induced digraph $\text{id}(G)$ such that (i, j) is an edge in $\text{id}(G)$ whenever

(u_i, r_j) is an edge in G and $i \neq j$. We use $\vec{G} \in \Gamma$ to denote that \vec{G} is strongly connected.

1.6.2 Modelling Graph-Based Deanonimisation

To model the security of ring samplers against graph-based deanonymisation attacks, in Section 3, we first formalise the notion of ring-sampler-induced transaction graphs, then model the security by designing a security experiment.

For any ring sampler Samp and any number of signatures $m \leq |U|$, we define the induced transaction graphs sampler $\mathcal{G}^{\text{Samp}}$ which inputs $(U, 1^m)$ and outputs a tuple (G, M) where $G = (U, R, E)$ with $|R| = m$ is a transaction graph induced by Samp and M is a maximum matching in G .

We say that Samp is ε -secure against graph-based deanonymisation attacks if no adversary, when given a transaction graph G where $(G, M) \leftarrow \mathcal{G}^{\text{Samp}}(U, 1^{|U|})$, could find an edge in M , i.e. a signer-ring assignment, with probability more than ε . The setting of $m = |U|$ in $\mathcal{G}^{\text{Samp}}(U, 1^{|U|})$ is without loss of generality due to Theorem 5.2, to be explained in Section 1.6.4. While the focus of this work is on passive adversaries, we also define a more general notion of security against active adversaries who compromise an admissible subset of users. The generalised notion captures the so-called “black marble attacks” [10, 13, 22] in the literature.

A trivial strategy of the adversary is to pick the smallest ring r^* in G and output a random edge connecting such ring, with success probability $1/|r^*|$. Therefore, a sampler Samp which outputs rings of a fixed size $k + 1$ cannot be ε -secure for $\varepsilon < \frac{1}{k+1}$. Intuitively, the trivial strategy is also the best strategy for the adversary in case $G = \text{Core}(G)$, which we prove to be the case for partitioning samplers in Section 7. Hence, to upper-bound the success probability of any adversary against Samp , it suffices to upper-bound the probability that $G \neq \text{Core}(G)$ for transaction graphs G induced by Samp .

1.6.3 Problem Reduction

Our first step for upper-bounding $\Pr[G \neq \text{Core}(G)]$, carried out in Section 4, is to reduce the problem about $\text{Core}(G)$ of a transaction graph G , a somewhat unwieldy object, to a simpler problem about the induced digraphs of the subgraphs of G . Although the results in Section 4 hold for general transaction graphs, they are motivated by the observation that the transaction graphs induced

by partitioning samplers could be partitioned into subgraphs whose induced digraphs follow some simple-to-describe distributions. The reduction is summarised by Theorem 4.6, which states that $\Pr[G \neq \text{Core}(G)]$ is upper-bounded by a sum of probabilities of some induced digraphs being not strongly connected.

1.6.4 Regular Partitioning Samplers

In Section 5, we move on to identify the transaction graphs induced by a partitioning sampler and their induced digraphs. Intuitively, the more information that is available to an adversary, the better it could perform in deanonymisation attacks, e.g. through graph analysis. Indeed, we show in Theorem 5.2 that for any number of signers $m \leq |U|$, the probability of $G \neq \text{Core}(G)$ where G is transaction graph sampled by a ring sampler is upper-bounded by that when $m = |U|$. This allows us to consider simply the latter case, which corresponds to that all users have signed.

Next, we focus on the partitioning ring samplers proposed in [19], denote by $\text{Samp} = \text{RegSamp}[P, k]$, which are parametrised by a partition P of U and a number of decoys k . The notation RegSamp stands for regular partitioning sampler, whose naming shall become clear shortly below. On input a signer s , $\text{RegSamp}[P, k]$ locates the chunk $C \in P$ which contains the signer s , samples a uniformly random $(k + 1)$ -subset r of C conditioning on $s \in R$, and outputs r as the ring.

A convenient property of a partitioning sampler $\text{RegSamp}[P, k]$ is that its distribution of induced transaction graphs can be naturally partitioned. Going through the reduction established in Section 4, we observe that the induced digraphs of each chunk in the partition follows the uniform distribution over all k -in-degree regular (hence the notation RegSamp) digraphs with $n = |C|$ vertices, denoted by $\vec{\mathcal{G}}_{k,n}^{\text{reg}}$. This, however, presents a challenge to our goal of upper-bounding the probability of $G \neq \text{Core}(G)$, since the distributions $\vec{\mathcal{G}}_{k,n}^{\text{reg}}$ do not appear to be well-studied in random graph theory.

1.6.5 Conjectures and Empirical Evidences

Towards circumventing the above problem, in Section 6, we turn our attention to the distribution $\vec{\mathcal{G}}_{p,n}^{\text{bin}}$ over digraphs with n vertices where each possible edge appears with probability p , with the intuition that the strong

connectivity of $\vec{\mathcal{G}}_{k,n}^{\text{reg}}$ could be estimated by that of $\vec{\mathcal{G}}_{p,n}^{\text{bin}}$ for appropriately chosen (k, p) .³

To relate the two distributions, in Conjecture 6.1, we conjecture that

$$\Pr_{\vec{G} \leftarrow \vec{\mathcal{G}}_{k,n}^{\text{reg}}} [\vec{G} \notin \Gamma] \leq \Pr_{\vec{G} \leftarrow \vec{\mathcal{G}}_{p,n}^{\text{bin}}} [\vec{G} \notin \Gamma]$$

when $p = \frac{k}{n-1}$ and therefore the expected in-degree for $\vec{G} \leftarrow \vec{\mathcal{G}}_{p,n}^{\text{bin}}$ is k . This makes sense intuitively when considering the natures of both digraph models. If the conjecture holds, it allows us to consider the distribution $\vec{\mathcal{G}}_{p,n}^{\text{bin}}$, which is better understood.

Based on the result of Palásti [14], the distribution $\vec{\mathcal{G}}_{p,n}^{\text{bin}}$ was studied by Graham and Pike [6], who proved the limit of $\Pr_{\vec{G} \leftarrow \vec{\mathcal{G}}_{p,n}^{\text{bin}}} [\vec{G} \notin \Gamma]$ under specific choice of p . Using this result, in Conjecture 6.4, we propose our second conjecture, which states that, for $p = \frac{k}{n-1}$,

$$\Pr_{\vec{G} \leftarrow \vec{\mathcal{G}}_{p,n}^{\text{bin}}} [\vec{G} \notin \Gamma] \leq 1 - e^{-2e^{\ln n - pn}},$$

where the expression on the right hand side is heuristically obtained from the result of [6].

Assuming both conjectures and combining all previous results, we conclude a closed-form upper bound for $\Pr[G \neq \text{Core}(G)]$. Although we are unable to prove the conjectures, in Section 6.2, we provide empirical evidences that they seem to hold at least for parameters of interest in the context of cryptocurrencies. In particular, we sampled 8000 random graphs according to either distribution in order to estimate the actual probabilities. We observe that the conjectured inequalities hold for all tested values of k and $n \geq 16$.

1.6.6 Provably Secure Ring Samplers

Putting everything together, in Section 7, we first show that $\text{RegSamp}[P, k]$ is ε -secure for

$$\varepsilon \leq \Pr[G \neq \text{Core}(G)] + \frac{1}{k+1}$$

where G is a random transaction graph induced by $\text{RegSamp}[P, k]$. Together with other established results, we prove that if

$$k \geq \ln(2 \cdot |U|) + \sqrt{2 \ln(2 \cdot |U|)}$$

³ Similar to how a regular partitioning sampler $\text{RegSamp}[P, k]$ relates to the distribution $\vec{\mathcal{G}}_{k,n}^{\text{reg}}$, a “binomial partitioning sampler” $\text{BinSamp}[P, p]$ could be constructed and be related to the distribution $\vec{\mathcal{G}}_{p,n}^{\text{bin}}$. To avoid distraction, we defer a discussion on this to Appendix B.

then $\text{RegSamp}[P, k]$ is ε -secure for $\varepsilon \leq \frac{2}{k+1}$. In other words, for this parameter choice, no graph-analysing adversary is likely to perform better than random guessing up to constant factor of 2.

Finally, we conclude our work by discussing the security of $\text{RegSamp}[P, k]$ against active graph-based attacks.

1.7 Related Work

We conclude the introduction by discussing related works in the areas of graph-based deanonymisation attacks, anonymity metrics, and random graph theory.

1.7.1 Graph-Based Attacks

In recent years, numerous works [7, 12, 25] demonstrated that, by reducing the ring membership relations represented by the transaction graph of Monero, it is possible to completely deanonymise signers of certain transactions. These attacks commonly rely on the fact that, in an early version of Monero, it was not mandatory for a signer to include decoys in a transaction. If such a signer A is chosen as a decoy in a ring sampled by another signer B, the possibility of A being the real signer of the transaction of B can be ruled out easily by the ring membership relation reduction, thereby reducing the anonymity of B. This anonymity reduction effect can be propagated to another signer C if it chooses B as a decoy in its ring, causing a chain reaction.

Recently, Vijayakumaran [21] proposed to use DM decomposition for deanonymising Monero signers, and showed that this is as effective as the prior methods [7, 12, 25]. Indeed, these prior attacks can be seen as finding certain subsets of edges not being in $\text{Core}(G)$ for a transaction graph G , and are therefore subsumed by DM decomposition which computes the entirety of $\text{Core}(G)$.

We remark that the aforementioned works mainly measure the effectiveness of an attack by counting the number of completely deanonymised signers, focusing little on partial deanonymisation. In contrast, the goal of this work is to upper-bound the probability of any partial deanonymisation.

1.7.2 Anonymity Metrics

Yu, Au, and Veríssimo [24] measured the anonymity of a transaction graph using the number of perfect matchings, which is unfortunately $\#P$ -complete to compute.

They also evaluated existing attacks and suggested a partitioning sampler which can be seen as a special case of those proposed in [19] and discussed below.

Beyond graph analysis, general deanonymisation attacks could take the signing probabilities of different signers into consideration. In this setting, Ronge *et al.* [19] proposed to model the anonymity provided by a ring sampler by the min-entropy $H_\infty(s|r)$ of the signer s conditioning on the ring r sampled by the signer s . According to this anonymity measure, the authors also proved that (regular) partitioning samplers are close to optimal assuming that the distribution of signing probabilities in each chunk is close to uniform. The formal definition of this anonymity measure and the corresponding optimality result on (regular) partitioning samplers are recalled in Appendix A.

A major shortcoming of the anonymity measure of Ronge *et al.* [19], however, is that it only captures the *local* anonymity of a single signer given a single ring. In particular, it does not capture *global* attacks such as those based on DM decomposition. Although extensions to the global setting were discussed, it is unclear whether the extended measures are efficiently computable.

We remark, however, that although graph analysis informs us about the anonymity of a ring sampler in the *global* sense, it disregards the distribution of signing probabilities. Consequently, a ring sampler (e.g. the uniform sampler) that behaves well under graph analysis might achieve low anonymity according to the entropy-based measure. We therefore view the two approaches as being complementary with each other.

1.7.3 Random (Di)graph Connectivity

Numerous results have been established for the connectivity problem of random (undirected) graphs. Erdős and Rényi [4] proved the asymptotic probability of a uniform random graph⁴ being connected. Łuczak [9] extended the result to binomial random graphs. Gilbert [5] gave both upper and lower bounds of the probability of a finite binomial random graph being connected. For k -regular random graphs, it is known that such graphs are almost surely connected for $k \geq 2$ [11] and almost surely k -connected for $k \geq 3$ [1].

⁴ A uniform random graph is a graph that is uniformly sampled from the set of all graphs with a fixed vertex set with certain fixed number of edges. A uniform random digraph is defined analogously.

The strong connectivity problem of random digraphs is, however, much worse understood. Among the existing literature, the majority focuses on infinite graphs. Palásti [14] and Graham and Pike [6] proved the asymptotic probability of strong connectedness for a uniform random digraph and a binomial random digraph respectively. Some works studied the asymptotic size of the giant strongly connected component (e.g. [16, 17]). Little seems to be known for finite graphs. The problem of computing (asymptotically) the probability of a k -in(/out)-degree regular random digraph being strongly connected was listed in The Scottish Book [11] in 1981, and in its second edition in 2015 this problem remains open. The reachability problem, which asks the probability that a given node can reach all other nodes in a random digraph, though intuitively simpler than the strong connectivity problem, is proven to be #P-complete [18].

2 Graphs

For $n \in \mathbb{N}$, write $[n] := \{1, 2, \dots, n\}$. A partition P of a set U is a set of disjoint subsets of U , called chunks, satisfying $\bigcup_{C \in P} C = U$. We often use U to denote the set of all users, i.e. potential signers.

We assume the general familiarity of the concepts of bipartite graphs and directed graphs (digraphs). In the following, we recall and establish some concepts which are specific to this work.

2.1 Bipartite Graphs

A bipartite graph $G = (A, B, E)$ consists of the vertex sets (A, B) (whose elements are also called nodes) and a set $E \subseteq A \times B$ of edges. Let $G = (A, B, E)$ and $H = (A', B', E')$ be bipartite graphs. We define the following basic operations and relations:

- Subgraph: H is a subgraph of G , denoted by $H \subseteq G$, if $A' \subseteq A$, $B' \subseteq B$ and $E' \subseteq E$.
- Union: $G \cup H := (A \cup A', B \cup B', E \cup E')$.
- Intersection: $G \cap H := (A \cap A', B \cap B', E \cap E')$.
- Difference: $G \setminus H := (A^-, B^-, E^-) \subseteq G$ where $A^- = A \setminus A'$, $B^- = B \setminus B'$ and $E^- = E \cap (A^- \times B^-)$.
- Edge elements: if $e \in E$, we sometimes abuse the notation and write $e \in G$.

Our analyses are primarily based on the concept of matchings in bipartite graphs, which we recall below.

Definition 2.1 (Matching). A matching $M \subseteq E$ in a bipartite graph $G = (A, B, E)$ is a subset of edges such that for all edges $(a, b), (a', b') \in M$, it holds that $a \neq a'$ and $b \neq b'$. We say M is a maximum matching, if $|M| \geq |M'|$ for any matching M' in G .

Definition 2.2 (Core). The core of a bipartite graph $G = (A, B, E)$, denoted by $\text{Core}(G) = (A, B, E')$, is a subgraph of G where $E' \subseteq E$ is the union of all maximum matchings in G .

The above concept of core is defined in the sense of Dulmage and Mendelsohn [3]. It should not be confused with the core defined with respect to graph homomorphisms. Tassa [20] gave an algorithm for computing $\text{Core}(G)$ in time linear in the number of nodes and edges of G .

A transaction graph is a bipartite graph $G = (U, R, E)$, where U is a set of users and R is a set of rings⁵, such that $|R| \leq |U|$ and there exists at least one maximum matching of size $|R|$. The edges E capture ring memberships, that is, if user u_i belongs to ring r_j , then $(u_i, r_j) \in E$. The existence of a size- $|R|$ maximum matching captures the assumption that each ring is generated by a distinct signer.

Definition 2.3 (Transaction Graph). A transaction graph $G = (U, R, E)$ is a bipartite graph with a maximum matching M of size $|R| \leq |U|$. We say that G is balanced if $|U| = |R|$. Otherwise it is imbalanced.

By renaming of nodes, we can write $U = \{u_i\}_{i=1}^n$, $R = \{r_j\}_{j=1}^m$, and $M = \{(u_j, r_j)\}_{j=1}^m$ for some $m, n \in \mathbb{N}$ with $n \geq m$ without loss of generality.

Definition 2.4 (Upper Graph). Let $G = (U, R, E)$ be a transaction graph, where $U = \{u_i\}_{i=1}^n$ and $R = \{r_j\}_{j=1}^m$, and $M = \{(u_j, r_j)\}_{j=1}^m$ be a maximum matching in G . The M -upper graph $G^M = (U^M, R, E^M)$ is a balanced transaction subgraph of G where $U^M := \{u_j\}_{j=1}^m$ and $E^M = E \cap (U^M \times R)$. We use G^\wedge to denote an M -upper graph G^M for an arbitrary M chosen deterministically given G .

The left panel of Figure 2 is an example transaction graph G with a maximum matching M . The upper graph G^M is the subgraph of G in the dotted rectangle.




Fig. 2. Example of a transaction graph G (U and R being nodes on left and right respectively) and its induced digraph $\text{id}(G)$. The subgraph in the dotted rectangle is G^\wedge . The yellow, blue and red edges correspond to edges considered in Lemma 4.1 Item 1 to Item 3 respectively, the black edges are none of them.

To capture transaction graphs induced by partitioning ring samplers, we define the notion of transaction graph partitioning.

Definition 2.5 (Transaction Graph Partitioning). Let U be a set of signers and P be a partition of U . Let $G = (U, R, E)$ and $G_C = (C, R_C, E_C)$ be transaction graphs for $C \in P$. We say that $\{G_C\}_{C \in P}$ is a partition of G if $\{R_C\}_{C \in P}$ and $\{E_C\}_{C \in P}$ are partitions of R and E respectively.

Let \mathcal{G} be a distribution of transaction graphs with vertex sets (U, R) . We say that $\{G_C\}_{C \in P}$ is a partition of \mathcal{G} , if G_C is a distribution of transaction graphs with vertex sets (C, R_C) for $C \in P$ and $\mathcal{G} = \bigcup_{C \in P} G_C$, i.e. sampling from \mathcal{G} is equivalent to first independently sampling from G_C for all $C \in P$ and then taking the union.

Clearly, if $\{G_C\}_{C \in P}$ is a partition of G , then the G_C 's have disjoint nodes and edges, and $\bigcup_{C \in P} G_C = G$.

2.2 Digraphs

A digraph $\vec{G} = (V, E)$ consists of a vertex set V and a set $E \subseteq V^2$ of edges. All digraphs considered in this work are without self-loop and parallel edge. The definitions of basic operations and relations for digraphs are analogous to those for bipartite graphs.

⁵ More precisely, R is a set of ring identifiers. This is to handle cases where different signers sample the same ring.

Definition 2.6 (Edge Reachability). Let $\vec{G} = (V, E)$ be a digraph. We say that an edge $e \in E$ is reachable from node $v \in V$ through \vec{G} , if there exists a directed path $P = \{(v_i, v_{i+1})\}_{i=0}^{\ell} \subseteq E$ for $v_0 = v$ and some $\ell \in \mathbb{N}$ such that $e \in P$. Generalising, if $\vec{H} = (W, F) \subseteq \vec{G}$, we say that e is reachable from \vec{H} through \vec{G} if e is reachable from w through \vec{G} for some $w \in W$.

The concepts of connectivity and strongly connected components of digraphs will be repeatedly used, their definitions are as follows.

Definition 2.7 (Strong and Weak Connectivity). A digraph $\vec{G} = (V, E)$ is strongly connected, denoted by $\vec{G} \in \Gamma$, if there exists a directed path from i to j for all distinct $i, j \in V$. The digraph \vec{G} is weakly connected, if there exists an (undirected) path from i to j for all distinct $i, j \in V$ when disregarding edge orientations.

Definition 2.8 (Strongly Connected Component). A strongly connected component (SCC) of a digraph \vec{G} is a subgraph of \vec{G} that is strongly connected, and is maximal with this property – that is, no further node or edge from \vec{G} can be added to it without breaking its strongly connected property.

To reduce problems about the cores of transaction graphs to those about digraphs connectivity, we define the notion of induced digraph $\text{id}(G)$ of a transaction graph G .

Definition 2.9 (Induced Digraph). Let $G = (U, R, E)$ be a transaction graph, where $U = \{u_i\}_{i=1}^n$ and $R = \{r_j\}_{j=1}^m$, and $M = \{(u_j, r_j)\}_{j=1}^m$ be a maximum matching in G . The M -induced digraph of G is defined as $\text{id}_M(G) := ([n], F)$ where $F := \{(i, j) \in [n]^2 : (u_i, r_j) \in E \wedge i \neq j\}$. We use $\text{id}(G)$ to denote an M -induced digraph $\text{id}_M(G)$ for an arbitrary M chosen deterministically given G .

In other words, given a maximum matching M , if we rename the users and rings so that $u_j \in r_j$ for all $r_j \in R$, the induced digraph is constructed by including an edge from node i to node j if user u_i is a member of ring r_j whenever $i \neq j$. Figure 2 gives an example of an induced digraph $\text{id}(G)$ of a transaction graph G .

We further introduce two special types of digraphs which the partitioning samplers will be related to.

Definition 2.10 (k -In-Degree Regular Digraphs). Let $k, n \in \mathbb{N}$ with $k < n$. A k -in-degree regular digraph is a digraph where all nodes have a fixed in-degree k . We

write $\vec{G}_{k,n}^{\text{reg}}$ for (the uniform distribution over) the set of all k -in-degree regular digraphs with the vertex set $[n]$.

Definition 2.11 (p -Binomial Digraphs). Let $p \in [0, 1]$ and $n \in \mathbb{N}$. We write $\vec{G}_{p,n}^{\text{bin}}$ for the distribution obtained by (uniformly) sampling a digraph \vec{G} with the vertex set $[n]$ such that each of the possible $n(n-1)$ edges is included in \vec{G} with probability p independent of any other edges.

3 Ring Samplers

We recall the formal definition of ring samplers [19] and define distributions of transaction graphs which are induced by ring samplers.

Definition 3.1 (Ring Samplers [19]). A ring sampler **Samp** is a (stateless) PPT algorithm which inputs a set of users U and a signer $s \in U$ and outputs a ring r satisfying $s \in r \subseteq U$. Syntactically, we write $r \leftarrow \text{Samp}(U, s)$ where **Samp** is understood to take uniform randomness which is omitted.

Remark 3.2. In general, a ring sampler **Samp** could input a set $s = \{s_1, s_2, \dots\} \subseteq U$ of signers and outputs a ring r with $s \subseteq r \subseteq U$.

Consider the following thought experiment: Let there be a set of users U . At each time j , a uniformly random user s_j who has not signed yet decides to issue a ring signature.⁶ To do so, user s_j samples a ring $r_j \leftarrow \text{Samp}(U, s_j)$ and publishes its ring signature together with the ring r_j . The ring membership relations of the published rings r_1, \dots, r_m form a transaction graph, whose distribution is induced by the randomness used for ring sampling.

Definition 3.3 (Induced Transaction Graphs). An induced transaction graph sampler $\mathcal{G}^{\text{Samp}}$ is an oracle-aided PPT algorithm which is given access to a ring sampler **Samp**, inputs a set of users U and a number $m \in [|U|]$ (in unary) of signers, and outputs a transaction graph G and a maximum matching M in G . The procedures of $\mathcal{G}^{\text{Samp}}$ are as described in Figure 3. Whenever we are only concerned with the transaction graph G sampled and not the maximum matching M , we omit M and write simply $G \leftarrow \mathcal{G}^{\text{Samp}}(U, 1^m)$.

⁶ In practice, in case two users publish their signatures simultaneously, a public tie-breaking rule is in place to decide which signature should be verified and accepted first.


```

 $\mathcal{G}^{\text{Samp}}(U, 1^m)$ 


---


for  $j \in [m]$  do
     $s_j \leftarrow U \setminus \{s_i\}_{i=1}^{j-1}$ 
     $r_j \leftarrow \text{Samp}(U, s_j)$ 
 $R := (r_1, \dots, r_m)$ 
 $E := \{(u, r) \in U \times R : u \in r\}$ 
 $M := \{(s_j, r_j)\}_{j \in [m]}$ 
 $G := (U, R, E)$ 
return  $(G, M)$ 
    
```

Fig. 3. Induced transaction graph sampler.

In the definition of induced transaction graphs, the maximum matching M output by $\mathcal{G}^{\text{Samp}}$ represents the “true” signer-ring assignment, i.e. each $(s_j, r_j) \in M$ represents that the ring r_j is sampled by signer s_j .

A graph-based deanonymisation attack against (a system employing) a ring sampler Samp can be modelled by a security experiment involving an adversary \mathcal{A} . We first consider the case with passive adversaries. The adversary \mathcal{A} is given a transaction graph G , where $(G, M) \leftarrow \mathcal{G}^{\text{Samp}}(U, 1^{|U|})$ is sampled by an induced transaction graph sampler, and is asked to find an edge (s^*, r^*) of G such that $(s^*, r^*) \in M$, i.e. to correctly identify that the ring r^* is sampled by the signer s^* . For the active setting, we additionally allow the adversary \mathcal{A} to corrupt a subset $B \subseteq U$ of users.

Definition 3.4. *Let $\varepsilon > 0$ and Samp be a ring sampler. We say that Samp is ε -secure against graph-based deanonymisation attacks if for any adversary \mathcal{A} and any set of users U ,*

$$\Pr [\text{Exp}_{\mathcal{A}, \text{Samp}}(U)] \leq \varepsilon.$$

Generalising, for any predicate π , we say that Samp is (π, ε) -secure against graph-based deanonymisation attacks if for any adversary \mathcal{A} and any set of users U ,

$$\Pr [\text{Exp}_{\mathcal{A}, \text{Samp}, \pi}(U)] \leq \varepsilon.$$

where the experiments $\text{Exp}_{\mathcal{A}, \text{Samp}}$ and $\text{Exp}_{\mathcal{A}, \text{Samp}, \pi}$ are described in Figure 4.

In Definition 3.4, we assume that all users have signed, i.e. $m = |U|$. This captures worst case security since the security experiment for smaller m can be emulated by the worst case adversary, as we will show in Theorem 5.2. While $\text{Exp}_{\mathcal{A}, \text{Samp}}$ captures passive attacks, $\text{Exp}_{\mathcal{A}, \text{Samp}, \pi}$ further captures active attacks by allowing the adversary to corrupt a subset B of users prior to receiving the transaction graph with the restriction that (U, B) satisfies the

```

 $\text{Exp}_{\mathcal{A}, \text{Samp}}(U)$ 


---


 $B \leftarrow \mathcal{A}(U)$ 
 $(G = (U, R, E), M) \leftarrow \mathcal{G}^{\text{Samp}}(U, 1^{|U|})$ 
 $U \leftarrow U \setminus B$ 
 $E \leftarrow E \cap (U \times U)$ 
 $G = (U, R, E)$ 
 $(u^*, r^*) \leftarrow \mathcal{A}(G)$ 
return  $((u^*, r^*) \in M) \wedge (\pi(U, B) = 1)$ 
    
```

Fig. 4. Experiments for the security of Samp against graph-based deanonymisation attacks. The variant incorporating black marble attacks is in dashed boxes.

predicate π . Setting π to only accept $B = \emptyset$, we recover the passive case.

Note that a trivial strategy for graph-based deanonymisation is to pick r^* with the fewest members, pick a random member $s^* \leftarrow r^*$, and output (s^*, r^*) . Clearly, this strategy has success probability $1/|r^*| = 1/(\min_{r \in R} |r|)$. As we will show in Section 7, conditioned on $G = \text{Core}(G)$, this is in fact the best strategy for attacking against partitioning ring samplers.

4 From Cores to Induced Digraphs

In this section we reduce the problem of upper-bounding $\Pr[G \neq \text{Core}(G)]$ to a problem concerning the strong connectivity of digraphs. We first recall a result from Tassa [20] for general bipartite graphs specialised to the case of transaction graphs.

Lemma 4.1 (Tassa [20]). *Let $G = (U, R, E)$ be a transaction graph, where $U = \{u_i\}_{i=1}^n$ and $R = \{r_j\}_{j=1}^m$, and $M = \{(u_j, r_j)\}_{j=1}^m$ be a maximum matching in G . The core $\text{Core}(G) = (U, R, E')$ is a transaction graph where E' is the union of the following sets:*

1. *The maximum matching M ,*
2. *$\{(u_i, r_j) : (i, j) \text{ is in some SCC of } \text{id}_M(G)\}$, and*
3. *$\left\{ (u_i, r_j) : \begin{array}{l} (i, j) \text{ is reachable from} \\ \text{id}_M(G) \setminus \text{id}_M(G^M) \text{ through } \text{id}_M(G) \end{array} \right\}$.*

Proof. This is a direct summary of the results in Tassa [20], specifically Theorem 2.2 and Algorithm 2 for Item 2, and Proposition 2.4, Theorem 2.7 and Algorithm 3 for Item 3. Item 1 is obvious by definition. \square

In the example given in Figure 2, the edges considered in Lemma 4.1 Items 1 to 3 are coloured yellow, blue,

and red respectively. The black edges are those not in $\text{Core}(G)$, corresponding to impossible signer-signature assignments that can be ruled out.

By Lemma 4.1 Item 3, any edge (u_i, r_j) of an imbalanced transaction graph with $i > m$ is maximum-matchable. Further, note that an edge in a digraph must either be within an SCC or connecting two SCCs, and not both. Hence, from Lemma 4.1 Item 2 and Item 3, any edge (u_i, r_j) not being in $\text{Core}(G)$ implies (u_i, r_j) is an edge connecting two SCCs in \vec{G}^M .

Using Lemma 4.1, we derive in the following a number of lemmas on the probability of $G \neq \text{Core}(G)$, which together will lead to Theorem 4.6.

We begin with Lemma 4.2, which states that if G has a partition P , then the cores $\text{Core}(H)$ of the chunks $H \in P$ collectively tell us everything about $\text{Core}(G)$.

Lemma 4.2. *Let G be a transaction graph and P be a partition of G . It holds that $G = \text{Core}(G)$ if and only if $H = \text{Core}(H)$ for all $H \in P$.*

Proof. Recall that $G = \bigcup_{H \in P} H$. Suppose for the moment that $\text{Core}(G) = \bigcup_{H \in P} \text{Core}(H)$, then we can prove the lemma statement as follows.

Suppose $G = \text{Core}(G)$. We have $\bigcup_{H \in P} H = \bigcup_{H \in P} \text{Core}(H)$. Observe that for distinct $H, H' \in P$ we must have $H \cap H' = \emptyset$. Therefore $H = \text{Core}(H)$ for all $H \in P$.

Suppose $H = \text{Core}(H)$ for all $H \in P$, then $G = \bigcup_{H \in P} H = \bigcup_{H \in P} \text{Core}(H) = \text{Core}(G)$.

It remains to show that $\text{Core}(G) = \bigcup_{H \in P} \text{Core}(H)$.

Let the edge $e \in \text{Core}(G)$, i.e. e belongs to a maximum matching M in G . Since $\text{Core}(G) \subseteq G$, and P is a partition of G , we have $e \in H^*$ for some $H^* \in P$. Since $M \cap H^*$ is a maximum matching in H^* , we have $e \in \text{Core}(H^*)$. This shows that $\text{Core}(G) \subseteq \bigcup_{H \in P} \text{Core}(H)$.

Let the edge $e \in \text{Core}(H)$ for some $H \in P$, i.e. e belongs to a maximum matching Y in H . Let M be a maximum matching in G whose existence is guaranteed since G is a transaction graph. Then $M^* := (M \setminus H) \cup Y$ is also a maximum matching in G . Consequently $e \in M^* \subseteq \text{Core}(G)$, which implies $\bigcup_{H \in P} \text{Core}(H) \subseteq \text{Core}(G)$. \square

As an immediate corollary of Lemma 4.2, Corollary 4.3 states a similar relation concerning distributions of transaction graphs. In particular, it states that the probability of $G \neq \text{Core}(G)$ is upper-bounded by the probability of the existence of a chunk G_C of G with $G_C \neq \text{Core}(G_C)$, which can further be upper-bounded by the union bound.

Corollary 4.3. *Let \mathcal{G} be any distribution of transaction graphs with identical vertex sets and let $\{G_C\}_{C \in P}$*

be a partition of \mathcal{G} . Then

$$\Pr_{G \leftarrow \mathcal{G}} [G \neq \text{Core}(G)] \leq \sum_{C \in P} \Pr_{G_C \leftarrow \mathcal{G}_C} [G_C \neq \text{Core}(G_C)].$$

Proof. By Lemma 4.2, we have

$$\Pr_{G \leftarrow \mathcal{G}} [G \neq \text{Core}(G)] = \Pr_{G \leftarrow \mathcal{G}} [\exists C \in P, G_C \neq \text{Core}(G_C)]$$

where on the right hand side $\{G_C\}_{C \in P}$ is a partition of G . We then arrive at the desired conclusion by applying the union bound. \square

Next, Lemma 4.4 upper-bounds the probability of $G \neq \text{Core}(G)$ by that of $G^\Delta \neq \text{Core}(G^\Delta)$, where we recall that G^Δ is an arbitrary fixed upper graph of G . Note that G^Δ is balanced by definition. Therefore, Lemma 4.4 in some sense means that balanced transaction graphs are the worst cases for how likely transaction graphs are equal to their respective cores.

Lemma 4.4. *Let $G = (U, R, E)$ be a transaction graph. If $G^\Delta = \text{Core}(G^\Delta)$, then $G = \text{Core}(G)$. Consequently, let \mathcal{G} be any distribution of transaction graphs, we have*

$$\Pr_{G \leftarrow \mathcal{G}} [G \neq \text{Core}(G)] \leq \Pr_{G \leftarrow \mathcal{G}} [G^\Delta \neq \text{Core}(G^\Delta)].$$

Proof. Let M be a maximum matching in G such that $G^\Delta = G^M = (U^M, R, E^M)$. It suffices to show that each chunk in the partition $\{E \setminus E^M, E^M \setminus M, M\}$ of E is a subset of the edges in $\text{Core}(G)$.

First, we have $e \in \text{Core}(G)$ for all $e \in M$ by the definition of core. Moreover, by Lemma 4.1 Item 3, $e \in \text{Core}(G)$ for edge $e \in E \setminus E^M$.

It remains to consider $E^M \setminus M$. Given that $G^M = \text{Core}(G^M)$, all $e \in E^M$ are in $\text{Core}(G^M)$. Since G^M is balanced, from Lemma 4.1 we have that all $e \in E^M \setminus M$ are in some SCC of $\text{id}_M(G^M)$. By construction, an SCC in $\text{id}_M(G^M)$ is also an SCC in $\text{id}_M(G)$, so by Lemma 4.1 Item 2 all $e \in E^M \setminus M$ are also in $\text{Core}(G)$. \square

Our last lemma for this section, Lemma 4.5, upper-bounds the probability of $G \neq \text{Core}(G)$ by that of $\text{id}(G)$ being not strongly connected, where we recall that $\text{id}(G)$ is an induced digraph of G with arbitrarily chosen maximum matching.

Lemma 4.5. *Let G be a transaction graph. If $\text{id}(G)$ is strongly connected, then $G = \text{Core}(G)$. Furthermore, if G is both balanced and connected, then the converse also holds. Consequently, let \mathcal{G} be any distribution of transaction graphs, we have*

$$\Pr_{G \leftarrow \mathcal{G}} [G \neq \text{Core}(G)] \leq \Pr_{G \leftarrow \mathcal{G}} [\text{id}(G) \notin \Gamma],$$

and the inequality become equality if \mathcal{G} is a distribution of balanced and connected transaction graphs.

Proof. If $\text{id}(G)$ is strongly connected, then by Lemma 4.1 all edges in G are in $\text{Core}(G)$, hence $G = \text{Core}(G)$.

The second statement is proven by contraposition. Suppose $\text{id}(G)$ is not strongly connected, so it has at least two SCCs \vec{C}_1 and \vec{C}_2 . If G is connected, then $\text{id}(G)$ is by construction weakly connected, and there exists an edge (i, j) in $\text{id}(G)$, where i is a node of \vec{C}_1 and j is a node of \vec{C}_2 . By Lemma 4.1 we have that (u_i, r_j) , which is an edge in G , is not in $\text{Core}(G)$, hence $G \neq \text{Core}(G)$. \square

Note that by construction, $\text{id}(G)$ is strongly connected only if G is balanced. Therefore the inequality in Lemma 4.5 becomes trivial if G is imbalanced.

Chaining together the above lemmas, we arrive at the main theorem of this section, which upper-bounds the probability of $G \neq \text{Core}(G)$ by a sum of probabilities related to the strong connectivity of the induced digraphs of the chunks of G .

Theorem 4.6. *Let \mathcal{G} be any distribution of transaction graphs and let $\{G_C\}_{C \in P}$ be a partition of \mathcal{G} . Then*

$$\Pr_{G \leftarrow \mathcal{G}} [G \neq \text{Core}(G)] \leq \sum_{C \in P} \Pr_{G_C \leftarrow \mathcal{G}_C} [\text{id}(G_C^\Delta) \notin \Gamma].$$

Proof. From Corollary 4.3,

$$\Pr_{G \leftarrow \mathcal{G}} [G \neq \text{Core}(G)] \leq \sum_{C \in P} \Pr_{G_C \leftarrow \mathcal{G}_C} [G_C \neq \text{Core}(G_C)].$$

From Lemmas 4.4 and 4.5 we have

$$\begin{aligned} \Pr_{G_C \leftarrow \mathcal{G}_C} [G_C \neq \text{Core}(G_C)] &\leq \Pr_{G_C \leftarrow \mathcal{G}_C} [G_C^\Delta \neq \text{Core}(G_C^\Delta)] \\ &\leq \Pr_{G_C \leftarrow \mathcal{G}_C} [\text{id}(G_C^\Delta) \notin \Gamma] \end{aligned}$$

for any $C \in P$. Combining the above yields the desired result. \square

5 Induced Transaction Graphs

Our goal in this section is to obtain a candidate upper bound for $\Pr [G \neq \text{Core}(G)]$, where G is a random transaction graph induced by a (regular) partitioning sampler [19]. For this, we first prove a theorem on the sufficiency of considering balanced induced transaction graphs. We then recall the definition of (regular) partitioning samplers [19] and apply the established theorems. We realise that $\Pr [G \neq \text{Core}(G)]$ can be upper-bounded in terms of $\Pr [\vec{G} \notin \Gamma]$ where \vec{G} is sampled from $\vec{\mathcal{G}}_{k,n}^{\text{reg}}$ (recall Definition 2.10).

5.1 Balanced Transaction Graphs

Intuitively, it is easier for an adversary to deanonymise signers when more information about them is available, for example, when more rings sampled by the signers are given. Following this line of thought, an adversary should be successful in deanonymising signers with the highest probability when all users have signed.

To formalise this claim, we first prove a technical lemma which states that, if H is constructed by adding ring nodes to a transaction graph G , then $G \neq \text{Core}(G)$ implies $H \neq \text{Core}(H)$.

Lemma 5.1. *Let $G = (U, R, E)$ and $H = (U, R', E')$ be transaction graphs where $R \subset R'$ and $E = E' \cap (U \times R)$, i.e. H can be constructed from G by adding ring nodes $R' \setminus R$ and edges connecting the new ring nodes to some signer nodes U . If $G \neq \text{Core}(G)$, then $H \neq \text{Core}(H)$.*

Proof. Let $U = \{u_i\}_{i=1}^n$, $R = \{r_j\}_{j=1}^m$, and $M = \{(u_j, r_j)\}_{j=1}^m$ be a maximum matching in G . It suffices to prove the case $|R'| = |R| + 1$, and the lemma follows by induction. We therefore assume from here on $R' = R \cup \{r_{m+1}\}$ where $r_{m+1} \notin R$.

Let $M' := M \cup \{(u_{m+1}, r_{m+1})\}$ be a maximum matching in H . Let $\text{id}_M(G) = ([n], F)$ and $\text{id}_{M'}(H) = ([n], F')$. Note that $F \subseteq F'$ (and hence $\text{id}_M(G) \subseteq \text{id}_{M'}(H)$), with the new edges in $F' \setminus F$ being of the form $(i, m+1)$ where $i \in [n] \setminus \{m+1\}$.

Suppose $G \neq \text{Core}(G)$, so there exists an edge $e^* = (u_{i^*}, r_{j^*})$ in G which is not in $\text{Core}(G)$. From Lemma 4.1 we have $i^* \neq j^*$, therefore $e^* \in \text{id}_M(G) \subseteq \text{id}_{M'}(H)$. We prove in the following that e^* is not in any SCC of $\text{id}_{M'}(H)$, and e^* is not reachable from $\text{id}_{M'}(H) \setminus \text{id}_{M'}(H^{M'})$ through $\text{id}_{M'}(H)$. Hence, by Lemma 4.1, e^* is not in $\text{Core}(H)$, and $H \neq \text{Core}(H)$.

We first show that e^* is not in any SCC of $\text{id}_{M'}(H)$. For this, note that from Lemma 4.1, e^* is an edge which connects two SCCs of $\text{id}_M(G^M)$. Let \vec{C} be an SCC of $\text{id}_M(G^M)$ such that i^* is a node of \vec{C} . Observe that by construction, \vec{C} is also an SCC of $\text{id}_M(G)$. Now, since the vertex set of \vec{C} is subset of $[m]$ (the vertex set of $\text{id}_M(G^M)$), there is no edge $(i, j) \in F' \setminus F$ with node j in \vec{C} (since all edges in $F' \setminus F$ are of the form $(i, m+1)$). Clearly this implies, first, that there is no edge in $F' \setminus F$ with both ends in \vec{C} , and second, that there is no edge in $F' \setminus F$ which connects from any node $v \in \text{id}_{M'}(H) \setminus \vec{C}$ to \vec{C} . Therefore, \vec{C} remains an SCC in $\text{id}_{M'}(H)$ by definition, and it follows that e^* is not in any SCC of $\text{id}_{M'}(H)$.

We next show that e^* is not reachable from $\text{id}_{M'}(H) \setminus \text{id}_{M'}(H^{M'})$ through $\text{id}_{M'}(H)$. We begin by drawing atten-

tion to two points. First, by Lemma 4.1, e^* is not reachable from $\text{id}_M(G) \setminus \text{id}_M(G^M)$ through $\text{id}_M(G)$. Second, e^* is not reachable from node $m+1$ through $\text{id}_{M'}(H)$, since e^* is not reachable from $m+1$ through $\text{id}_M(G)$ and all edges in $F' \setminus F$ are of the form $(i, m+1)$. We proceed to prove the statement by contradiction. Suppose e^* is reachable from $\text{id}_{M'}(H) \setminus \text{id}_{M'}(H^{M'})$ through $\text{id}_{M'}(H)$, then there exists a directed path $P = \{(v_{i-1}, v_i)\}_{i=1}^\ell$ in $\text{id}_{M'}(H)$, where v_0 is a node of $\text{id}_{M'}(H) \setminus \text{id}_{M'}(H^{M'})$, $(v^{\ell-1}, v^\ell) = e^*$, and node $v_i \in R'$ for all $i \in [\ell]$.⁷ However, $v_i \neq m+1$ for all $i \in [\ell]$, otherwise contradicting that e^* is not reachable from node $m+1$ through $\text{id}_{M'}(H)$. Therefore $v_i \in M$ for all $i \in [\ell]$. Since all edges in $F' \setminus F$ are of the form $(i, m+1)$, we now have that all edges in P belong to $\text{id}_M(G)$. In other words, e^* is reachable from $\text{id}_{M'}(H) \setminus \text{id}_{M'}(H^{M'})$ through $\text{id}_M(G)$. Finally, since $\text{id}_{M'}(H) \setminus \text{id}_{M'}(H^{M'}) = ([n] \setminus (R \cup \{m+1\}), \emptyset) \subset ([n] \setminus R, \emptyset) = \text{id}_M(G) \setminus \text{id}_M(G^M)$, we arrive at that e^* is reachable from $\text{id}_M(G) \setminus \text{id}_M(G^M)$ through $\text{id}_M(G)$, a contradiction. \square

From Lemma 5.1 we obtain our next theorem, which states that for any number of signers $m \leq |U|$, $\Pr[G \neq \text{Core}(G)]$ is upper-bounded by that when $m = |U|$, i.e. the case that all users have signed.

Theorem 5.2. *For any ring sampler Samp and any $m \leq |U|$, it holds that*

$$\begin{aligned} & \Pr_{G \leftarrow \mathcal{G}^{\text{Samp}}(U, 1^m)} [G \neq \text{Core}(G)] \\ & \leq \Pr_{H \leftarrow \mathcal{G}^{\text{Samp}}(U, 1^{|U|})} [H \neq \text{Core}(H)]. \end{aligned}$$

Proof. As Samp is stateless, the distributions of the outputs of independent runs of Samp are independent. Hence, referring to Figure 3, sampling H from $\mathcal{G}^{\text{Samp}}(U, 1^{|U|})$ is equivalent to first running the for-loop in $\mathcal{G}^{\text{Samp}}(U, 1^{|U|})$ only up to $j = m$ to sample G , then running the remaining of the loop to sample G' , and outputting $H := G \cup G'$. From Lemma 5.1, we know that $H \neq \text{Core}(H)$ whenever $G \neq \text{Core}(G)$. The claim thus follows immediately. \square

5.2 Regular Partitioning Samplers

We consider a special case of the partitioning samplers defined in [19], where there is only one public partition of U and only one signer per ring. The general case with a

distribution of partitions and more than one signer can be handled with generic techniques [19]. Such partitioning samplers, which we refer to as the regular partitioning samplers $\text{RegSamp}[P, k]$, are parametrised by the partition P of U and a number of decoys $k \in \mathbb{N}$ for each ring, such that $k < |C|$ for each chunk $C \in P$. We recall its definition below.

$\text{RegSamp}[P, k](U, s)$: Initiate $r := \{s\}$. Let $C \in P$ be the unique chunk containing s . Sample a uniformly random k -subset $r' \subseteq C \setminus \{s\}$. Output $r := r \cup r'$.

We observe that a $\text{RegSamp}[P, k]$ -induced transaction graph G takes a special form – it can be partitioned into independent subgraphs $\{G_C\}_{C \in P}$, each representing the induced transaction graph of a chunk in P . Moreover, if a subgraph G_C is balanced, then its induced digraph $\text{id}(G_C)$ is a k -in-degree regular digraph. We therefore arrive immediately at the following lemma.

Lemma 5.3. *Let U be a set of users and P be a partition of U . Let $k \in \mathbb{N}$ such that $k < |C|$ for each $C \in P$. Write $\text{Samp} := \text{RegSamp}[P, k]$. For any $m \leq |U|$,*

$$\Pr_{G \leftarrow \mathcal{G}^{\text{Samp}}(U, 1^m)} [G \neq \text{Core}(G)] \leq \sum_{C \in P} \Pr_{\vec{G} \leftarrow \vec{\mathcal{G}}_{k, |C|}^{\text{reg}}} [\vec{G} \notin \Gamma].$$

$$\begin{aligned} & \text{Proof.} \quad \Pr_{G \leftarrow \mathcal{G}^{\text{Samp}}(U, 1^m)} [G \neq \text{Core}(G)] \\ & \leq \Pr_{G \leftarrow \mathcal{G}^{\text{Samp}}(U, 1^{|U|})} [G \neq \text{Core}(G)] \\ & \leq \sum_{C \in P} \Pr_{G \leftarrow \mathcal{G}^{\text{RegSamp}}[\{C\}, k](C, 1^{|C|})} [\text{id}(G) \notin \Gamma] \\ & = \sum_{C \in P} \Pr_{\vec{G} \leftarrow \vec{\mathcal{G}}_{k, |C|}^{\text{reg}}} [\vec{G} \notin \Gamma], \end{aligned}$$

where the first inequality follows from Theorem 5.2, the second inequality from Theorem 4.6, and the equality follows from direct inspection. \square

Lemma 5.3 relates the probability of $G \neq \text{Core}(G)$ with that of $\vec{G} \notin \Gamma$, where G is a transaction graph induced by a regular partitioning sampler and \vec{G} is a k -in-degree regular digraph. Unfortunately, the strong connectivity of random k -in-degree regular digraphs seems to be a non-trivial problem [11, Problem 38]. While (asymptotic) results on the connectivity of random k -regular (undirected) graphs are established [1], their extensions to the strong connectivity of random k -in(/out)-degree regular digraphs remain open. In the next section, we circumvent this difficulty by estimating the strong connectivity of random k -in-degree regular digraphs by that of random p -binomial digraphs (recall Definition 2.11) for appropriate k and p .

⁷ The condition on the intermediate nodes can be achieved by first considering any path P from node $v_0 \in \text{id}_{M'}(H) \setminus \text{id}_{M'}(H^{M'})$ to e^* through $\text{id}_{M'}(H)$, and then taking the tail of P such that no intermediate node in the tail belongs to $\text{id}_{M'}(H) \setminus \text{id}_{M'}(H^{M'})$.

Remark 5.4. To draw connection between partitioning samplers and random p -binomial digraphs, consider the following “binomial partitioning samplers” construction modified from that of regular partitioning samplers: Instead of sampling a random k -subset of $C \setminus \{s\}$, the modified sampler includes each member of $C \setminus \{s\}$ into the ring independently with some fixed probability p . Correspondingly, a counterpart of Lemma 5.3 for binomial partitioning sampler could be stated. For details, we refer to Appendix B.

6 Conjectures and Experiments

Towards finding the final piece of the puzzle of upper-bounding $\Pr [G \neq \text{Core}(G)]$ for G induced by partitioning samplers, we put forth two conjectures concerning the probabilities of random k -in-degree regular digraphs and random p -binomial digraphs being strongly connected. To gain confidence in these conjectures, we empirically estimate the probabilities for parameters which are reasonable in the context of cryptocurrencies.

6.1 Conjectures

Our first conjecture relates the two digraph distributions $\vec{\mathcal{G}}_{k,n}^{\text{reg}}$ and $\vec{\mathcal{G}}_{k,n}^{\text{bin}}$.

Conjecture 6.1. For $k, n \in \mathbb{N}$ with $n \geq 16$ and $p = \frac{k}{n-1} \leq 1$,

$$\Pr_{\vec{G} \leftarrow \vec{\mathcal{G}}_{k,n}^{\text{reg}}} [\vec{G} \notin \Gamma] \leq \Pr_{\vec{G} \leftarrow \vec{\mathcal{G}}_{p,n}^{\text{bin}}} [\vec{G} \notin \Gamma].$$

The condition $n \geq 16$ stems from our simulation results, which we detail in Section 6.2. Intuitively Conjecture 6.1 makes sense, since for all digraphs in the support of $\vec{\mathcal{G}}_{k,n}^{\text{reg}}$, all nodes must be weakly connected to k other nodes, whereas this is not the case for $\vec{\mathcal{G}}_{p,n}^{\text{bin}}$ with any $p < 1$.

In search of a closed-form upper bound for $\Pr [G \neq \text{Core}(G)]$, we draw on the following result from Graham and Pike [6], which are developed based on the work of Palásti [14].

Lemma 6.2 ([6]). Let $c \in \mathbb{R}$ be a constant and $p(n) := \frac{\ln n + c}{n}$. It holds that

$$\lim_{n \rightarrow \infty} \Pr_{\vec{G} \leftarrow \vec{\mathcal{G}}_{p(n),n}^{\text{bin}}} [\vec{G} \notin \Gamma] = 1 - e^{-2e^{-c}}.$$

Remark 6.3. Graham and Pike [6] considered a different model of digraphs where, unlike ours, self-loops

are allowed. Their result however still holds under our model of digraphs, since self-loops have no effect on the strong connectivity of a digraph.

Lemma 6.2 moves us closer towards a closed-form upper bound for $\Pr [G \neq \text{Core}(G)]$, but unfortunately with two issues. First, the results of Palásti [14] and Graham and Pike [6] seem to crucially rely on setting $p(n) := \frac{\ln n + c}{n}$, and infer nothing about the case with general p . Second, their results concern only about infinite digraphs, but say little about finite digraphs.

To close the gaps, we propose our second conjecture, which is obtained heuristically by plugging in $c = pn - \ln n$ and $p = \frac{k}{n-1}$ back to the limit in Lemma 6.2.

Conjecture 6.4. For $k, n \in \mathbb{N}$ with $n \geq 16$ and $p = \frac{k}{n-1} \leq 1$,

$$\Pr_{\vec{G} \leftarrow \vec{\mathcal{G}}_{p,n}^{\text{bin}}} [\vec{G} \notin \Gamma] \leq 1 - e^{-2e^{\ln n - \frac{k}{n-1} n}}.$$

While we are unable to provide analytical proofs, both of the conjectures hold in our numerical simulations in Section 6.2, where (k, n) are chosen to be realistic in the context of cryptocurrencies.

Finally, taking these two conjectures, we can bridge the established results and arrive at the concluding statement below.

Corollary 6.5. Let U be a set of users and P be a partition of U . Let $k \in \mathbb{N}$ such that $k < |C|$ for each $C \in P$. Let $n := \max_{C \in P} |C| \geq 16$. If Conjectures 6.1 and 6.4 hold, then for any $m \leq |U|$,

$$\Pr_{G \leftarrow \mathcal{G}^{\text{RegSamp}}[P,k](U,1^m)} [G \neq \text{Core}(G)] \leq |P| \left(1 - e^{-2e^{\ln n - k}}\right).$$

Proof.

$$\begin{aligned} & \Pr_{G \leftarrow \mathcal{G}^{\text{RegSamp}}[P,k](U,1^m)} [G \neq \text{Core}(G)] \\ & \leq \sum_{C \in P} \Pr_{\vec{G} \leftarrow \vec{\mathcal{G}}_{k,|C|}^{\text{reg}}} [\vec{G} \notin \Gamma] \\ & \leq \sum_{C \in P} \Pr_{\vec{G} \leftarrow \vec{\mathcal{G}}_{p(C),|C|}^{\text{bin}}} [\vec{G} \notin \Gamma] \\ & \leq \sum_{C \in P} \left(1 - e^{-2e^{\ln |C| - \frac{k}{|C|-1} |C|}}\right) \\ & < \sum_{C \in P} \left(1 - e^{-2e^{\ln |C| - k}}\right) \\ & \leq |P| \left(1 - e^{-2e^{\ln n - k}}\right), \end{aligned}$$

where the first inequality follows from Lemma 5.3, the second follows from Conjecture 6.1 by setting $p(C) = \frac{k}{|C|-1}$ for $C \in P$, and the third follows from Conjecture 6.4. \square




Fig. 5. Plots of $p_{k,n}^{\text{reg}}$, $p_{k,n}^{\text{bin}}$, and $\bar{p}_{k,n}^{\text{bin}}$ against k for selected values of n in both linear- and log-scale.

6.2 Experiments

To support our conjectures, we empirically estimated the probabilities

$$p_{k,n}^{\text{reg}} := \Pr_{\vec{G} \leftarrow \vec{\mathcal{G}}_{k,n}^{\text{reg}}} [\vec{G} \notin \Gamma] \quad \text{and}$$

$$p_{k,n}^{\text{bin}} := \Pr_{\vec{G} \leftarrow \vec{\mathcal{G}}_{p,n}^{\text{bin}}} [\vec{G} \notin \Gamma],$$

where $p = \frac{k}{n-1}$, for values of k ranging from 1 to 16 and values of n from 2^2 to 2^{12} in exponential steps. In each case we sampled 8000 graphs, verified whether $\vec{G} \notin \Gamma$, and compared the average with the upper bound

$$\bar{p}_{k,n}^{\text{bin}} := 1 - e^{-2e^{\ln n - \frac{k}{n-1}n}}$$

in Conjecture 6.4.

In Figure 5, we plotted $p_{k,n}^{\text{reg}}$ (dot mark, “regular”), $p_{k,n}^{\text{bin}}$ (plus mark, “binomial”), and $\bar{p}_{k,n}^{\text{bin}}$ (dashed, “bound”) against k for different values of n in both linear- and log-scale. In the log-scale plot, values smaller than 10^{-3} are omitted for their instability due to the limited sampling size. Similarly, in Figure 6 we plotted the same values against n for different values of k .

From Figures 5 and 6, we observe that both conjectured upper bounds appear to hold for all $n \geq 16$. More specifically, the only cases where they fail to hold are $(k, n) = (1, 4)$ and $(1, 8)$. Upon closer inspection, on the one hand, we observe that the first bound

$$p_{k,n}^{\text{reg}} \leq p_{k,n}^{\text{bin}}$$

becomes tighter as the number of nodes n decreases. This makes sense since the variance of the in-degree of the nodes in the graphs sampled from $\vec{\mathcal{G}}_{p,n}^{\text{bin}}$ decreases as n decreases. On the other hand, we notice that the second conjectured upper bound

$$p_{k,n}^{\text{bin}} \leq \bar{p}_{k,n}^{\text{bin}}$$

becomes tighter as n increases. This is also expected as the bound was heuristically derived from the limit of $p_{k,n}^{\text{bin}}$ as n tends to infinity.

7 Interpretation of Our Results

We conclude our work by stating a ring size for partitioning samplers which is sufficient to defeat graph analysis. We also discuss how our results extend to the setting with an active adversary, who attempts to deanonymise honest signers by injecting fake ones in the so-called “black marble attacks” [10, 13, 22].

7.1 On Defeating Graph Analysis

We discuss what our results mean in the context of (passive) graph-based deanonymisation attacks. We begin by showing that, for transaction graphs G induced by k -regular partitioning samplers, conditioned on $G = \text{Core}(G)$, the trivial deanonymisation strategy described in Section 3 is the best strategy.

Lemma 7.1. *Let $k \in \mathbb{N}$, U be a set of users, P be a partition of U where $|C| > k$ for each $C \in P$. Let $\text{Samp} = \text{RegSamp}[P, k]$. For any adversary \mathcal{A} ,*

$$\Pr [\text{Exp}_{\mathcal{A}, \text{Samp}}(U)] \leq \Pr [G \neq \text{Core}(G)] + \frac{1}{k+1}$$

where the probabilities are taken over the randomness of \mathcal{A} and $(G, M) \leftarrow \mathcal{G}^{\text{Samp}}(U, 1^{|U|})$.

Proof. Observe that

$$\Pr [\text{Exp}_{\mathcal{A}, \text{Samp}}(U)]$$

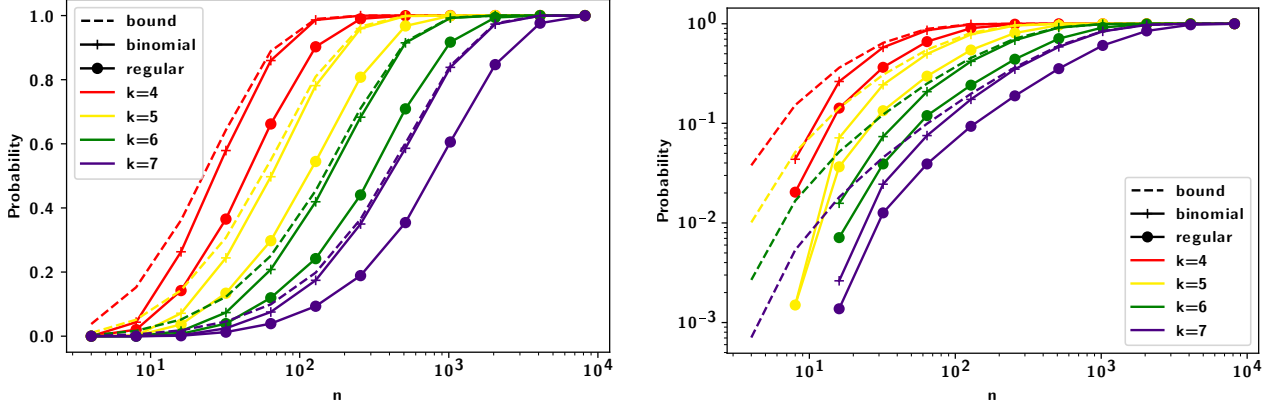


Fig. 6. Plots of $p_{k,n}^{\text{reg}}$, $p_{k,n}^{\text{bin}}$, and $p_{k,n}^{\text{bin}}$ against n for selected values of k in both linear- and log-scale.

$$\leq \Pr[G \neq \text{Core}(G)] + \Pr[\text{Exp}_{\mathcal{A}, \text{Samp}}(U) | G = \text{Core}(G)],$$

which is obtained by applying the law of total probability and upper-bounding two probability terms by 1. Then, it suffices to show that

$$\Pr[\text{Exp}_{\mathcal{A}, \text{Samp}}(U) | G = \text{Core}(G)] \leq \frac{1}{k+1}.$$

Consider the distribution

$$\hat{\mathcal{G}}^{\text{Samp}} := \left\{ (G, M) : \begin{array}{l} (G, M) \leftarrow \mathcal{G}^{\text{Samp}}(U, 1^{|U|}) \\ G = \text{Core}(G) \end{array} \right\}.$$

In other words, $\hat{\mathcal{G}}^{\text{Samp}}$ is the same as $\mathcal{G}^{\text{Samp}}$ conditioning on $G = \text{Core}(G)$. Let $\hat{\text{Exp}}_{\mathcal{A}, \text{Samp}}$ be the same as $\text{Exp}_{\mathcal{A}, \text{Samp}}$, except that the procedure $(G, M) \leftarrow \mathcal{G}^{\text{Samp}}$ is replaced by $(\hat{G}, \hat{M}) \leftarrow \hat{\mathcal{G}}^{\text{Samp}}$. We can rewrite

$$\Pr[\text{Exp}_{\mathcal{A}, \text{Samp}}(U) | G = \text{Core}(G)] = \Pr[\hat{\text{Exp}}_{\mathcal{A}, \text{Samp}}(U)].$$

Since Samp is a partitioning sampler, for any fixed members of any fixed ring, the probability of the ring being sampled by each member is the same. That is, for any fixed ring r in the support of $\bigcup_{u \in U} \text{Samp}(U, u)$, and any fixed $s, s' \in r$, it holds that

$$\Pr[r = \text{Samp}(s)] = \Pr[r = \text{Samp}(s')].$$

Therefore, conditioned on the event $E = ((s, r), (s', r) \in \hat{G})$ for any fixed r, s, s' ,

$$\Pr[(s, r) \in \hat{M} | E] = \Pr[(s', r) \in \hat{M} | E]$$

with probabilities taken over $(\hat{G}, \hat{M}) \leftarrow \hat{\mathcal{G}}^{\text{Samp}}$. Consequently, for any edge (s^*, r^*) output by $\mathcal{A}(\hat{G})$,

$$\begin{aligned} & \Pr[\hat{\text{Exp}}_{\mathcal{A}, \text{Samp}}(U)] \\ &= \Pr[(s^*, r^*) \in \hat{M}] \end{aligned}$$

$$\begin{aligned} &= \Pr[(s^*, r^*) \in \hat{G}] \Pr[(s^*, r^*) \in \hat{M} | (s^*, r^*) \in \hat{G}] \\ &\quad + \Pr[(s^*, r^*) \notin \hat{G}] \Pr[(s^*, r^*) \in \hat{M} | (s^*, r^*) \notin \hat{G}] \\ &= \Pr[(s^*, r^*) \in \hat{G}] \cdot \frac{1}{|r^*|} + 0 \leq \frac{1}{k+1}, \end{aligned}$$

as desired. \square

From Lemma 7.1, if the parameters P and k are set such that $\Pr_{G \leftarrow \mathcal{G}^{\text{RegSamp}}[P, k]}(U, 1^{|U|}) [G \neq \text{Core}(G)] \leq \frac{1}{k+1}$, then $\text{RegSamp}[P, k]$ is ε -secure against graph-based deanonymisation attacks for $\varepsilon = \frac{2}{k+1} = O(1/k)$, which is optimal up to a constant factor of 2. In the next theorem, we give a sufficient condition on k with which this holds.

Theorem 7.2. *Let $k, n \in \mathbb{N}$, U be a set of users, and P be a partition of U where $|C| = n \geq 16$ and $n > k$ for each $C \in P$. If Conjectures 6.1 and 6.4 hold and*

$$k \geq \ln(2|U|) + \sqrt{2 \ln(2|U|)},$$

then $\text{RegSamp}[P, k]$ is $\frac{2}{k+1}$ -secure against graph-based deanonymisation attacks.

Proof. By Lemma 7.1, it suffices to show that $\Pr_{G \leftarrow \mathcal{G}^{\text{RegSamp}}[P, k]}(U, 1^{|U|}) [G \neq \text{Core}(G)] \leq \frac{1}{k+1}$ for the given parameters. Let $k' := k + 1$. If Conjectures 6.1 and 6.4 hold, then by Corollary 6.5 it suffices to set up parameters such that

$$|P| \left(1 - e^{-2e^{\ln n - k}}\right) \leq \frac{1}{k'}$$

or equivalently

$$k \geq \ln \left(\frac{-2n}{\ln \left(1 - \frac{1}{|P|k'}\right)} \right),$$

where $n = \max_{C \in P} |C|$ is the maximum chunk size. Since $\ln \left(1 - \frac{1}{|P|k'}\right) \leq -\frac{1}{|P|k'}$, we have

$$k \geq \ln(2n|P|k')$$

as a sufficient condition, so it suffices to solve k' for

$$k' - \ln k' \geq \ln(2en|P|),$$

the solution of which is

$$k' \geq -W_{-1} \left(\frac{-1}{2en|P|} \right),$$

where $W_{-1}(\cdot)$ is the Lambert W function of branch -1 . From [2] we know that

$$-1 - \sqrt{2x} - x < W_{-1}(-e^{-x-1})$$

for all $x > 0$. Substituting $x = \ln(2n|P|)$, we conclude that it suffices to set

$$k \geq \ln(2n|P|) + \sqrt{2\ln(2n|P|)}.$$

In the particular case stated in the theorem statement, where the set of users is partitioned into chunks of equal size n , i.e. $|U| = n|P|$, it suffices to set

$$k \geq \ln(2|U|) + \sqrt{2\ln(2|U|)}. \quad \square$$

For concreteness, suppose it is believed that the number of all users $|U|$ will never exceed 2^{64} , then Theorem 7.2 suggests that, by setting the number of dummies k to at least 55, the probability that an adversary identifies a signer is at most $\frac{2}{k+1} \leq \frac{1}{28}$. Suppose that users are comfortable with a 1-in- t anonymity for some $t \geq 28$, then it should suffice to set k as such that $\frac{k+1}{2} = t$, yielding a ring size of $2t$.

In the example of Monero, its current recommended ring size of 11 seems far too small under our model. We note, however, that a “correct” level of anonymity is itself a subjective matter. If a Monero user is willing to accept that the anonymity set size will be reduced, say, from 11 to 6 and is comfortable with an anonymity set of size 6, then 11 might still be an acceptable choice. Future empirical study on the actual reduction in anonymity of Monero users could offer useful insights in this direction.

We remark that the above recommendation for the ring size is conservative for several reasons. First, the upper bound of the adversary’s success probability given in Lemma 7.1 is loose in the sense that, while we let $\Pr[\text{Exp}_{\mathcal{A}, \text{Samp}}(U) | G \neq \text{Core}(G)] \leq 1$ in its derivation, having $G \neq \text{Core}(G)$ does not necessarily mean that the adversary immediately has a drastic advantage. Rather, we believe that the anonymity degrades gracefully depending on how close $\text{Core}(G)$ is to G . Second, ring samplers which are secure in our model resist even untargeted attacks against individual signatures. In practice, being able to identify the signer of one random signature

does not seem very useful, especially in the LRS setting where each signing key is only used once. A more meaningful attack, say in the setting of cryptocurrencies, is to identify the signers of a chain of $\ell > 1$ transactions. However, the probability of successfully doing so intuitively decreases exponentially in ℓ .

7.2 On Black Marble Attacks

A type of active deanonymisation attacks is the so-called “black marble attacks” [10, 13, 22], where the adversary actively injects signers, called black marbles, into the set of users, such that including them in rings do not contribute towards the anonymity of honest signers. In the context of cryptocurrencies, injecting black marbles often incur a monetary cost. It is therefore reasonable to assume that the adversary is only able to inject a bounded number of black marbles per some unit of time. Such attacks can be captured by the experiment $\text{Exp}_{\mathcal{A}, \text{Samp}, \pi}$ in Definition 3.4.

For simplicity, suppose that each chunk $C \in P$ is of size $|C| = n$ and contains $\beta \cdot n$ black marbles for some $\beta \in [0, 1]$. Then the “effective” number of users (in the sense of providing anonymity) is given by $(1 - \beta) \cdot |U|$. This is captured by a predicate π which checks that $|B \cap C| \leq \beta \cdot |C|$ for all $C \in P$.

Suppose that $\text{Samp} = \text{RegSamp}[P, k]$. Notice that, after removing the black marbles, the induced digraphs of the chunks of the transaction graphs G are no longer k -in-degree regular, and are tedious to analyse. Fortunately, for the case with binomial partitioning sampler (detailed in Appendix B), we observe that the induced digraphs of the chunks of the transaction graphs G follow the distributions $\left\{ \mathcal{G}_{p, (1-\beta) \cdot n}^{\text{bin}} \right\}_{C \in P}$. That is, injecting black marbles only decreases the size parameter of the p -binomial digraph distribution by a factor of $(1 - \beta)$. We can therefore still apply Conjecture 6.4 and obtain an analogous upper bound for this setting. By replacing $|U|$ with $(1 - \beta) \cdot |U|$ in the proof of Theorem 7.2, we conclude that it suffices to set

$$p \geq \frac{\ln(2 \cdot (1 - \beta) \cdot |U|) + \sqrt{2\ln(2 \cdot (1 - \beta) \cdot |U|)}}{(1 - \beta)n - 1}$$

to defeat graph analysis. Revisiting the setting of $\text{Samp} = \text{RegSamp}[P, k]$, the above heuristically suggests that

$$k \gtrsim \frac{\ln(2 \cdot (1 - \beta) \cdot |U|) + \sqrt{2\ln(2 \cdot (1 - \beta) \cdot |U|)}}{1 - \beta}$$

suffices to defeat graph analysis.

Acknowledgements

This work is supported by Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) as part of the Research and Training Group 2475 “Cybercrime and Forensic Computing” (number 393541319/GRK2475/1-2019) and by the state of Bavaria at the Nuremberg Campus of Technology (NCT).

References

- [1] B. Bollobás and B. Béla. *Random graphs*. 73. Cambridge university press, 2001.
- [2] I. Chatzigeorgiou. “Bounds on the Lambert Function and Their Application to the Outage Analysis of User Cooperation.” In: *IEEE Commun. Lett.* 17.8 (2013), pp. 1505–1508. DOI: 10.1109/LCOMM.2013.070113.130972. URL: <https://doi.org/10.1109/LCOMM.2013.070113.130972>.
- [3] A. L. Dulmage and N. S. Mendelsohn. “Coverings of Bipartite Graphs.” In: *Canadian Journal of Mathematics* 10 (1958), 517–534. DOI: 10.4153/CJM-1958-052-0.
- [4] P. Erdős and A. Rényi. “On random graphs I Publ.” In: *Math. Debrecen* 6 (1959), pp. 290–297.
- [5] E. N. Gilbert. “Random graphs.” In: *The Annals of Mathematical Statistics* 30.4 (1959), pp. 1141–1144.
- [6] A. J. Graham and D. A. Pike. “A note on thresholds and connectivity in random directed graphs.” In: *Atl. Electron. J. Math* 3.1 (2008), pp. 1–5.
- [7] A. Kumar, C. Fischer, S. Tople, and P. Saxena. “A Traceability Analysis of Monero’s Blockchain.” In: *ESORICS 2017, Part II*. Ed. by S. N. Foley, D. Gollmann, and E. Sneekenes. Vol. 10493. LNCS. Springer, Heidelberg, Sept. 2017, pp. 153–173. DOI: 10.1007/978-3-319-66399-9_9.
- [8] J. K. Liu, V. K. Wei, and D. S. Wong. “Linkable Spontaneous Anonymous Group Signature for Ad Hoc Groups (Extended Abstract).” In: *ACISP 04*. Ed. by H. Wang, J. Pieprzyk, and V. Varadharajan. Vol. 3108. LNCS. Springer, Heidelberg, July 2004, pp. 325–335. DOI: 10.1007/978-3-540-27800-9_28.
- [9] T. Łuczak. “On the equivalence of two basic models of random graphs.” In: *Random graphs*. Vol. 87. 1987, pp. 151–157.
- [10] A. Mackenzie, S. Noether, and M. C. Team. *Improving Obfuscation in the CryptoNote Protocol*. Tech. rep. URL: <https://www.getmonero.org/resources/research-lab/pubs/MRL-0004.pdf>.
- [11] R. D. Mauldin. *The Scottish Book*. Vol. 88. 8. Springer, 1981.
- [12] M. Möser, K. Soska, E. Heilman, K. Lee, H. Heffan, S. Srivastava, K. Hogan, J. Hennessey, A. Miller, A. Narayanan, and N. Christin. “An Empirical Analysis of Traceability in the Monero Blockchain.” In: *PoPETs 2018.3* (July 2018), pp. 143–163. DOI: 10.1515/popets-2018-0025.
- [13] S. Noether, S. Noether, and A. Mackenzie. *A Note on Chain Reactions in Traceability in CryptoNote 2.0*. Tech. rep. URL: <https://www.getmonero.org/resources/research-lab/pubs/MRL-0001.pdf>.
- [14] I. Palásti. “On the strong connectedness of directed random graphs.” In: *Studia Sci. Math. Hungar* 1 (1966), pp. 205–214.
- [15] S. Patachi and C. Schürmann. “Eos a Universal Verifiable and Coercion Resistant Voting Protocol.” In: *Electronic Voting - Second International Joint Conference, E-Vote-ID 2017, Bregenz, Austria, October 24-27, 2017, Proceedings*. Ed. by R. Krimmer, M. Volkamer, N. B. Binder, N. Kersting, O. Pereira, and C. Schürmann. Vol. 10615. Lecture Notes in Computer Science. Springer, 2017, pp. 210–227. DOI: 10.1007/978-3-319-68687-5_13. URL: https://doi.org/10.1007/978-3-319-68687-5_13.
- [16] M. D. Penrose. “The strong giant in a random digraph.” In: *Journal of Applied Probability* 53.1 (2016), pp. 57–70.
- [17] B. Pittel and D. Poole. “Asymptotic distribution of the numbers of vertices and arcs of the giant strong component in sparse random digraphs.” In: *Random Structures & Algorithms* 49.1 (2016), pp. 3–64.
- [18] J. S. Provan and M. O. Ball. “The complexity of counting cuts and of computing the probability that a graph is connected.” In: *SIAM Journal on Computing* 12.4 (1983), pp. 777–788.
- [19] V. Ronge, C. Egger, R. W. F. Lai, D. Schröder, and H. H. F. Yin. “Foundations of Ring Sampling.” In: *PoPETs 2021.3* (July 2021), pp. 265–288. DOI: 10.2478/popets-2021-0047.
- [20] T. Tassa. “Finding all maximally-matchable edges in a bipartite graph.” In: *Theoretical Computer Science* 423 (2012), pp. 50–58. ISSN: 0304-3975. DOI: <https://doi.org/10.1016/j.tcs.2011.12.071>. URL: <https://www.sciencedirect.com/science/article/pii/S0304397511010474>.

[21] S. Vijayakumaran. *Analysis of CryptoNote Transaction Graphs using the Dulmage-Mendelsohn Decomposition*. Cryptology ePrint Archive, Report 2021/760. <https://ia.cr/2021/760>. 2021.

[22] D. A. Wijaya, J. Liu, R. Steinfeld, and D. Liu. “Monero Ring Attack: Recreating Zero Mixin Transaction Effect.” In: *TrustCom/BigDataSE 2018*. IEEE, 2018, pp. 1196–1201. DOI: 10.1109/TrustCom/BigDataSE.2018.00165.

[23] B. Yu, J. K. Liu, A. Sakzad, S. Nepal, R. Steinfeld, P. Rimba, and M. H. Au. “Platform-Independent Secure Blockchain-Based Voting System.” In: *ISC 2018*. Ed. by L. Chen, M. Manulis, and S. Schneider. Vol. 11060. LNCS. Springer, Heidelberg, Sept. 2018, pp. 369–386. DOI: 10.1007/978-3-319-99136-8_20.

[24] J. Yu, M. H. A. Au, and P. J. E. Verissimo. “Re-Thinking Untraceability in the CryptoNote-Style Blockchain.” In: *CSF 2019 Computer Security Foundations Symposium*. Ed. by S. Delaune and L. Jia. IEEE Computer Society Press, 2019, pp. 94–107. DOI: 10.1109/CSF.2019.00014.

[25] Z. Yu, M. H. Au, J. Yu, R. Yang, Q. Xu, and W. F. Lau. “New Empirical Traceability Analysis of CryptoNote-Style Blockchains.” In: *FC 2019*. Ed. by I. Goldberg and T. Moore. Vol. 11598. LNCS. Springer, Heidelberg, Feb. 2019, pp. 133–149. DOI: 10.1007/978-3-030-32101-7_9.

A Entropy-Based Anonymity

Ronge *et al.* [19] introduced an anonymity measure for ring samplers based on conditional min-entropy. They also proved that the regular partitioning sampler achieves close to optimal anonymity with respect to this measure under a realistic assumption about the signer distribution. Here we recall the definition of this measure, and prove that the binomial partitioning sampler also achieves close to optimal anonymity with respect to this measure under the same assumption. In this context, a ring sampler Samp is assumed to always sample a ring for some signer s from the set of users, for brevity we omit in the following the input U and write simply $\text{Samp}(s)$.

Definition A.1 (Conditional Min-entropy). *Let \mathcal{X} and \mathcal{Y} be discrete distributions with probability mass functions $p_{\mathcal{X}}$ and $p_{\mathcal{Y}}$ respectively. Let $p_{\mathcal{X}|\mathcal{Y}}$ and $p_{\mathcal{Y}|\mathcal{X}}$ be the corresponding conditional probability mass functions.*

The conditional min-entropy of \mathcal{X} given \mathcal{Y} is defined as

$$\begin{aligned} H_{\infty}(\mathcal{X}|\mathcal{Y}) &:= -\ln \left(\sum_y p_{\mathcal{Y}}(y) \cdot \max_x p_{\mathcal{X}|\mathcal{Y}}(x|y) \right) \\ &= -\ln \left(\sum_y \max_x (p_{\mathcal{Y}|\mathcal{X}}(y|x) \cdot p_{\mathcal{X}}(x)) \right). \end{aligned}$$

Definition A.2 (Signer Distributions [19]). *A signer distribution \mathcal{S} is a distribution over $2^U \setminus \{\emptyset\}$, i.e. each sample of \mathcal{S} is a non-empty subset of U . If all samples of \mathcal{S} are singletons, i.e. $\Pr_{S \leftarrow \mathcal{S}}[|S| = 1] = 1$, we say that \mathcal{S} is a single-signer distribution.*

Definition A.3 (Anonymity [19]). *The anonymity of Samp with respect to a signer distribution \mathcal{S} is defined as*

$$\alpha[\mathcal{S}, \text{Samp}] := H_{\infty}(\mathcal{S}|\text{Samp}(U, \mathcal{S})).$$

Note that the anonymity measure defined in Definition A.3 captures only “local” anonymity since it disregards information about the signer leaked from the rings generated by other users. While the anonymity measure could be generalised to the “global” setting by simply considering the min-entropy of \mathcal{S} conditioned on a sequence of rings, analysing ring samplers with respect to such generalised measure appears to be difficult. Indeed, all analyses done in [19] were with respect to the local measure defined in Definition A.3.

Ronge *et al.* [19] proved that the regular partitioning samplers achieve close to optimal anonymity with respect to the above measure under a mild assumption.

Lemma A.4 ([19, Theorem 6.3]). *Let P be a partition of U . Let \mathcal{S} be a single-signer distribution with probability mass function $p_{\mathcal{S}}$. For each $C \in P$, let μ_C be the mean of $p_{\mathcal{S}}(s)$ over all $s \in C$, i.e. $\mu_C := |C|^{-1} \sum_{s \in C} p_{\mathcal{S}}(s)$. Suppose that for all $C \in P$, all $s \in C$, it holds that $|p_{\mathcal{S}}(s) - \mu_C| \leq \varepsilon_C$ for some $\varepsilon_C \geq 0$. Let $\varepsilon_P := \sum_{C \in P} |C| \varepsilon_C$. Then*

$$\alpha(\mathcal{S}, \text{RegSamp}[P, k]) > \ln k - \ln(\varepsilon_P + 1).$$

B Binomial Partitioning Samplers

Similar to Lemma 5.3 which relates the regular partitioning samplers to the distribution $\tilde{\mathcal{G}}_{k,n}^{\text{reg}}$, we can construct a new type of partitioning samplers – the binomial partitioning samplers – which could be related to the distribution $\tilde{\mathcal{G}}_{p,n}^{\text{bin}}$.

Loosely speaking, a binomial partitioning sampler similarly partitions the set of users into chunks, and

within each chunk the sampler includes each signer as decoy in a ring with some fixed probability independent of all other signers. The independence of signers being chosen as decoys turns out to make the analysis of the corresponding induced transaction graphs much easier than that of the regular partitioning samplers.

As in Section 5.2, we consider the case where there is only one public partition of U and only one signer per ring. A binomial partitioning sampler $\text{BinSamp}[P, p]$, parametrised by the partition P of U and a decoy probability p , is defined as follows.

$\text{BinSamp}[P, p](U, s)$: Initiate $r := \{s\}$. Let $C \in P$ be the unique chunk containing s and, for each $d \in C \setminus \{s\}$, run $r := r \cup \{d\}$ with probability p . Output r .

In the setting where s is a set of signers instead of a single one, a ring could be sampled by repeating the above for each member of s and taking the union.

In case \mathcal{S} is a single-signer distribution, $|C| = n$ for each chunk $C \in P$, and $p = \frac{k}{n-1}$, the binomial partitioning samplers $\text{BinSamp}[P, p]$ are analogous to the regular partitioning samplers $\text{RegSamp}[P, k]$, in the sense that the former has expected ring size $k + 1$ while the latter has fixed size $k + 1$. Furthermore, the numbers of decoys in a ring sampled from $\text{BinSamp}[P, p]$ follow the binomial distribution with mean k and variance $k(1 - p)$.

Similar to the regular partitioning samplers, the distribution of transaction graphs induced by a binomial partitioning sampler is related to some specific distribution. Clearly, the distribution $\mathcal{G}^{\text{BinSamp}[P, p]}$ can be partitioned as $\{\mathcal{G}_C\}_{C \in P}$, each \mathcal{G}_C being independent of each other and representing the distribution of induced transaction graphs of a chunk $C \in P$. Furthermore, each \mathcal{G}_C can be sampled by setting each of the possible edges independently with probability p . We therefore arrive at the following analogy to Lemma 5.3.

Lemma B.1. *Let U be a set of users and P be a partition of U . Let $p \in [0, 1]$. Write $\text{Samp} := \text{BinSamp}[P, p]$. For any $m \leq |U|$,*

$$\Pr_{G \leftarrow \mathcal{G}^{\text{Samp}}(U, 1^m)} [G \neq \text{Core}(G)] \leq \sum_{C \in P} \Pr_{\vec{G} \leftarrow \vec{\mathcal{G}}_{p, |C|}^{\text{bin}}} [\vec{G} \notin \Gamma].$$

Proof. Similar to the proof of Lemma 5.3. \square

Analogous to Lemma 7.1, it is not difficult to show a similar bound for $\text{Samp} = \text{BinSamp}[P, p]$. As the binomial partitioning sampler has variable ring sizes, in the analysis we need to use a tail bound to argue that, with overwhelming (in k) probability, all rings produced by

$\mathcal{G}^{\text{BinSamp}[P, k]}$ have size not far from $k + 1$. Since the argument is tedious but straightforward, we omit it.

For the sake of completeness, we analyse the anonymity of the binomial partitioning samplers according to the entropy-based measure. It turns out that the binomial partitioning samplers have the same near-optimal level of anonymity as the regular partitioning samplers do.

Theorem B.2. *Let P be a partition of U . Let \mathcal{S} be a single-signer distribution with probability mass function $p_{\mathcal{S}}$. Let P and $k \in \mathbb{N}$ be such that $p|C| > k$ for each $C \in P$. For each $C \in P$, let μ_C be the mean of $p_{\mathcal{S}}(s)$ over all $s \in C$, i.e. $\mu_C := |C|^{-1} \sum_{s \in C} p_{\mathcal{S}}(s)$. Suppose that for all $C \in P$, all $s \in C$, it holds that $|p_{\mathcal{S}}(s) - \mu_C| \leq \varepsilon_C$ for some $\varepsilon_C \geq 0$. Let $\varepsilon_P := \sum_{C \in P} |C| \varepsilon_C$. Then*

$$\alpha(\mathcal{S}, \text{BinSamp}[P, p]) > \ln k - \ln(\varepsilon_P + 1).$$

Proof. Let $\text{Samp} = \text{BinSamp}[P, p]$. For any $s \in U$, as the chunk containing s is unique, we know that $\bigcup_{C \in P} (2^C \setminus \{\emptyset\})$ is a superset of the collection of all possible rings. Write $\mathcal{R}_C := 2^C \setminus \{\emptyset\}$ and $\mathcal{R} := \bigcup_{C \in P} \mathcal{R}_C$. Since the ring given by the sampler must contain the signer, we have for all signer s and for all $r \in \mathcal{R}$,

$$\Pr [\text{Samp}(U, s) = r \wedge s \notin r] = 0.$$

If $s \in C \in P$, then each element in $C \setminus \{s\}$ has a probability p to be included in $r \setminus \{s\}$. On the other hand, if $s \notin C \in P$, then we must have $r \notin \mathcal{R}_C$. Therefore, for any $s \in U$, $C \in P$, and $r \in \mathcal{R}_C$, we have

$$\begin{aligned} \Pr [\text{Samp}(U, s) = r \wedge s \in r] &= \begin{cases} p^{|r|-1} (1-p)^{(|C|-1)-(|r|-1)} & s \in C \\ 0 & s \notin C \end{cases} \\ &= \begin{cases} p^{|r|-1} (1-p)^{|C|-|r|} & s \in C \\ 0 & s \notin C. \end{cases} \end{aligned}$$

Now, we analyse the anonymity of the sampler.

$$\begin{aligned} 2^{-\alpha[\mathcal{S}, \text{Samp}]} &= 2^{H_{\infty}(\mathcal{S} | \text{Samp}(U, \mathcal{S}))} \\ &= \sum_{r \in \mathcal{R}} \max_{s \in U} (p_{\text{Samp}(U, \mathcal{S}) | \mathcal{S}}(r | s) \cdot p_{\mathcal{S}}(s)) \\ &\leq \sum_{C \in P} \sum_{r \in \mathcal{R}_C} \max_{s \in C} \left(\Pr [\text{Samp}(U, s) = r \wedge s \in r] \cdot p_{\mathcal{S}}(s) \right) \\ &= \sum_{C \in P} \sum_{r \in \mathcal{R}_C} p^{|r|-1} (1-p)^{|C|-|r|} \max_{s \in C} p_{\mathcal{S}}(s) \\ &= \sum_{C \in P} \frac{1 - (1-p)^{|C|}}{p} \max_{s \in C} p_{\mathcal{S}}(s) \end{aligned}$$

$$\begin{aligned} &\leq \sum_{C \in P} \frac{1 - (1-p)^{|C|}}{p} (\mu_C + \varepsilon_C) \\ &< \sum_{C \in P} \frac{|C|}{k} (\mu_C + \varepsilon_C) \\ &= \frac{\varepsilon_P + 1}{k}. \end{aligned}$$

□