


1 March 2024 Suspected Black Marble Flooding Against Monero:  
2 Privacy, User Experience, and Countermeasures

3 Draft v0.2  
4 Rucknium 


5 March 27, 2024

6 **Abstract**

7 On March 4, 2024, aggregate Monero transaction volume suddenly almost tripled. This note an-  
8 alyzes the effect of the large number of transactions, assuming that the transaction volume is an  
9 attempted black marble flooding attack by an adversary. According to my estimates, mean effective  
10 ring size has decreased from 16 to 5.5 if the black marble flooding hypothesis is correct. At current  
11 transaction volumes, the suspected spam transactions probably cannot be used for large-scale “chain re-  
12 action” analysis to eliminate all ring members except for the real spend. Effects of increasing Monero’s  
13 ring size above 16 are analyzed.

# 1 March 4, 2024: Sudden transaction volume

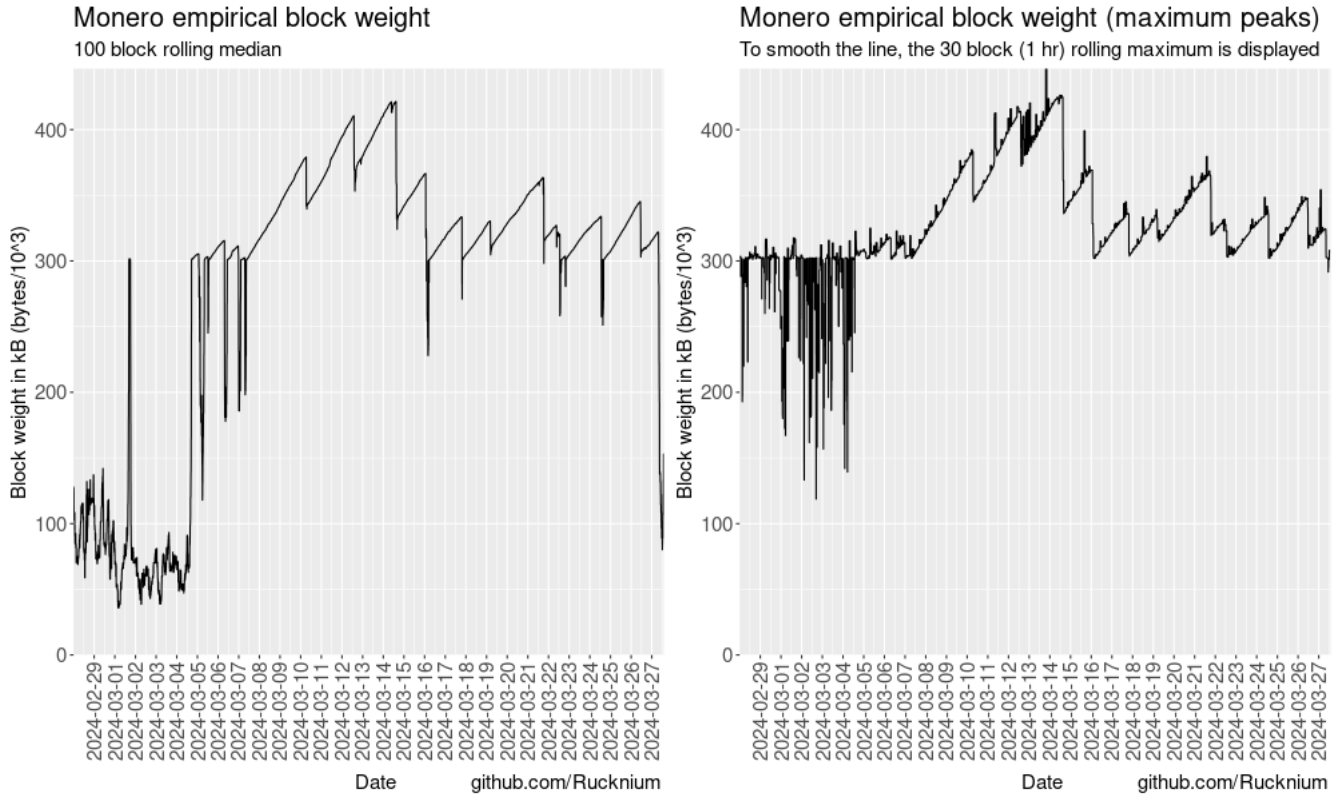
Figure 1: Volume of Monero transactions with spam fingerprint



15 On March 4, 2024 at approximately block height 3097764 (15:21:24 UTC), the number of 1input/2output  
16 minimum fee (20 nanoneros/byte) transactions sent to the Monero network rapidly increased. Figure 1  
17 shows daily volume of this type of transaction increasing from about 15,000 to over 100,000.

18 The large volume of these transactions was enough to entirely fill the 300 kB Monero blocks mined  
19 about every two minutes. Monero’s dynamic block size algorithm activated. The 100 block rolling median  
20 block size slowly increased to adjust for the larger number of transactions that miners could pack in blocks.  
21 Figure 2 shows the adjustment. The high transaction volume raised the 100 block median gradually for  
22 period of time. Then the transaction volume reduced just enough to allow the 100 block median to reset to  
23 a lower level. Then the process would restart. Block sizes have usually remained between 300 kB and 400  
24 kB. Occasionally, high-fee transactions would allow miners to get more total revenue by giving up some  
25 of the 0.6 XMR/block tail emission and including more transactions in a block. The “maximum peaks”  
26 plot shows this phenomenon.

Figure 2: Monero empirical block weight



27 The sudden transaction volume rise may originate from a single entity. The motive may be spamming  
28 transactions to bloat the blockchain size, increase transaction confirmation times for real users, perform  
29 a network stress test, or execute a black marble flooding attack to reduce the privacy of Monero users. I  
30 will focus most of my analysis on the last possibility.

## 31 2 Literature review

32 The very first research bulletin released by the Monero Research Lab described black marble transaction  
33 flooding. [Noether et al., 2014] points out that the ring signature privacy model requires rings to contain  
34 transaction outputs that are could be plausible real spends. If a single entity owns a large share of outputs  
35 (spent or not), it can use its knowledge to rule out ring members in other users' transactions that cannot  
36 be the real spend. Since the entity knows that itself did not spend the output(s) in a particular ring, the  
37 effective ring size that protects other users' privacy can be reduced — even to an effective ring size of 1  
38 when the entity knows the real spend with certainty. Rings with known real spends can be leveraged to  
39 determine the real spend in other rings in a “chain reaction” attack.

40 [Noether et al., 2014] gave the name “black marble” to the outputs owned by an anti-privacy adversary  
41 since they modeled the problem using a marble draw problem with a hypergeometric distribution. When  
42 a specific number of marbles are drawn *without* replacement from an urn containing a specific number of

43 white and black marbles, the hypergeometric distribution describes the probability of drawing a specific  
44 number of black marbles. In my modeling I use the binomial distribution, which is the same as the  
45 hypergeometric except marbles are drawn *with* replacement. The binomial distribution makes more sense  
46 now ten years after [Noether et al., 2014] was written. The total number of RingCT outputs on the  
47 blockchain that can be included in a ring is over 90 million. The hypergeometric distribution converges to  
48 the binomial distribution as the total number of marbles increases to infinity. Moreover, Monero’s current  
49 decoy selection algorithm does not select all outputs with equal probability. More recent outputs are  
50 selected with much higher probability. The hypergeometric distribution cannot be used when individual  
51 marbles have unequal probability of being selected.

52 [Chervinski et al., 2021] simulates a realistic black marble flood attack. They consider two scenarios.  
53 The adversary could create 2input/16output transactions to maximize the number of black marble outputs  
54 per block or the adversary could create 2input/2output transactions to make the attack less obvious. The  
55 paper uses Monero transaction data from 2020 to set the estimated number of real outputs and kB per  
56 block at 41 outputs and 51 kB respectively. The nominal ring size at this time was 11. The researchers  
57 simulated filling the remaining 249 kB of the 300 kB block with black marble transactions. A “chain  
58 reaction” algorithm was used to boost the effectiveness of the attack. In the 2in/2out scenario, the real  
59 spend could be deduced (effective ring size 1) in 11% of rings after one month of spamming black marbles.  
60 Later I will compare the results of this simulation with the current suspected spam incident.

61 [Krawiec-Thayer et al., 2021] analyze a suspected spam incident in July-August 2021. Transactions’  
62 inputs, outputs, fees, and ring member ages were plotted to evaluate evidence that a single entity created  
63 the spam. The analysis concluded, “All signs point towards a single entity. While transaction homogeneity  
64 is a strong clue, a the [sic] input consumption patterns are more conclusive. In the case of organic growth  
65 due to independent entities, we would expect the typically semi-correlated trends across different input  
66 counts, and no correlation between independent users’ wallets. During the anomaly, we instead observed  
67 an extremely atypical spike in 1–2 input txns with no appreciable increase in 4+ input transactions.”

68 TODO: A few papers like [Ronge et al., 2021, Egger et al., 2022] discuss black marble attacks too.

### 69 **3 Black marble theory**

70 The binomial distribution describes the probability of drawing  $x$  number of “successful” items when drawing  
71 a total of  $n$  items when the probability of a successful draw is  $p$ . It can be used to model the number  
72 of transaction outputs selected by the decoy selection algorithm that are not controlled by a suspected  
73 adversary.

74 The probability mass function of the binomial distribution with  $n \in \{0, 1, 2, \dots\}$  number of draws and  
75  $p \in [0, 1]$  probability of success is

$$f(x, n, p) = \binom{n}{x} p^x (1 - p)^{n-x}, \text{ where } \binom{n}{x} = \frac{n!}{x!(n-x)!} \quad (1)$$

76 The expected value (the theoretical mean) of a random variable with a binomial distribution is  $np$ .  
 77 Monero’s standard decoy selection algorithm programmed in `wallet2` does not select outputs with  
 78 equal probability. The probability of selecting each output depends on the age of the output. Specifics  
 79 are in [Rucknium, 2023]. The probability of a single draw selecting an output that is not owned by  
 80 the adversary,  $p_r$ , is equal to the share of the probability mass function occupied by those outputs:  
 81  $p_r = \sum_{i \in R} g(i)$ , where  $R$  is the set of outputs owned by real users and  $g(x)$  is the probability mass  
 82 function of the decoy selection algorithm.

### 83 3.1 Spam assumptions

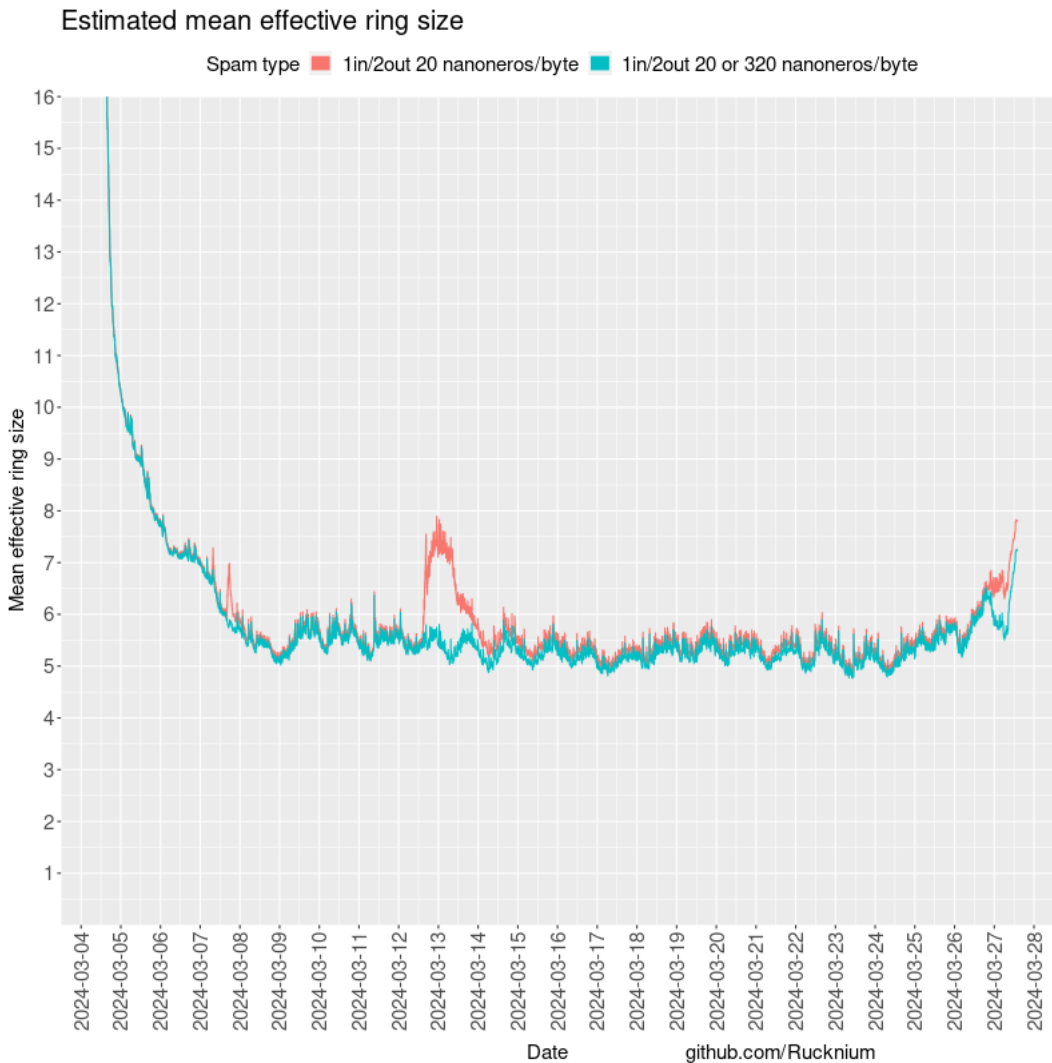
84 There is some set of criteria that identifies suspected spam. The early March 2024 suspected spam trans-  
 85 actions: 1) have one input; 2) have two outputs; 3) pay the minimum 20 nanoneros per byte transaction  
 86 fee. The normal volume of these transactions produced by real users must be estimated. The volume in  
 87 excess of the normal volume is assumed to be spam. I followed this procedure:

- 88 1. Compute the mean number of daily transactions that fit the suspected spam criteria for the four  
 89 weeks that preceded the suspected spam incident. A separate mean was calculated for each day  
 90 of the week (Monday, Tuesday,...) because Monero transaction volumes have weekly cycles. These  
 91 volume means are denoted  $v_{r,m}, v_{r,t}, v_{r,w}, \dots$  for the days of the week.
- 92 2. For each day of the suspected spam interval, sum the number of transactions that fit the suspected  
 93 spam criteria. Subtract the amounts found in step (1) from this sum, matching on the day of the  
 94 week. This provides the estimated number of spam transactions for each day:  $v_{s,1}, v_{s,2}, v_{s,3}, \dots$
- 95 3. For each day of the suspected spam interval, randomly select  $v_{s,t}$  transactions from the set of trans-  
 96 actions that fit the suspected spam criteria, without replacement. This randomly selected set is  
 97 assumed to be the true spam transactions.
- 98 4. During the period of time of the spam incident, compute the expected probability  $p_r$  that one output  
 99 drawn from the `wallet2` decoy distribution will select an output owned by a real user (instead of  
 100 the adversary) when the wallet constructs a ring at the point in time when the blockchain tip is at  
 101 height  $h$ . The closed-form formula of the `wallet2` decoy distribution is in [Rucknium, 2023].
- 102 5. The expected effective ring size of each ring constructed at block height  $h$  is  $1 + 15 \cdot p_r$ . The coefficient  
 103 on  $p_r$  is the number of decoys.

104 Figure 3 shows the results of this methodology. The mean effective ring size settled at about 5.5 by the  
 105 fifth day of the large transaction volume. On March 12 and 13 there was a large increase in the number

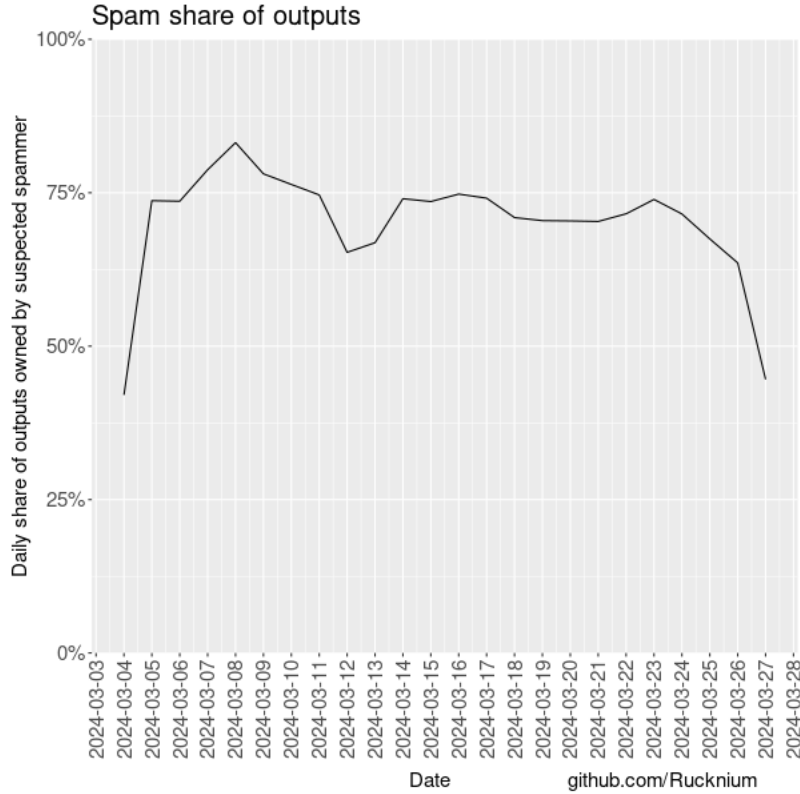
106 of 1in/2out transactions that paid 320 nanoneros/byte (the third fee tier). This could have been  
107 the spammer switching fee level temporarily or a service that uses Monero increasing fees to avoid delays.  
108 I used the same method to estimate the spam volume of these 320 nanoneros/byte suspected spam. The  
109 1in/2out 320 nanoneros/byte transactions displaced some of the 1in/2out 20 nanoneros/byte transactions  
110 because miners preferred to put transactions with higher fees into blocks. Other graphs and analysis will  
111 consider only the 1in/2out 20 nanoneros/byte transactions as spam unless indicated otherwise.

Figure 3: Estimated mean effective ring size



112 Figure 4 shows the daily share of outputs on the blockchain that are owned by the suspected spammer.  
113 The mean share of outputs since the suspected spam started is about 75 percent.

Figure 4: Spam share of outputs



114 **3.2 Long term projection scenarios at different ring sizes**

115 Fix the number of outputs owned by real users at  $r$ . The analysis will let the number  $s$  of outputs owned  
 116 by the adversary vary. The share of outputs owned by real users is

$$p_r = \frac{r}{r + s} \tag{2}$$

117 The 2 expression can be written  $p_r = \frac{1}{r} \cdot \frac{r}{1 + \frac{1}{r}s}$ , which is the formula for hyperbolic decay with the  
 118 additional  $\frac{1}{r}$  coefficient at the beginning of the expression [Aguado et al., 2010].

119 Let  $n$  be the nominal ring size (16 in Monero version 0.18). The number of decoys chosen by the decoy  
 120 selection algorithm is  $n - 1$ . The mean effective ring size for a real user’s ring is one (the real spend) plus  
 121 the ring’s expected number of decoys owned by other real users.

$$E[n_e] = 1 + (n - 1) \cdot \frac{r}{r + s} \tag{3}$$

122 The empirical analysis of Section 3.1 considered the fact that the `wallet2` decoy selection algorithm  
 123 draws a small number of decoys from the pre-spam era. Now we will assume that the spam incident has  
 124 continued for a very long time and all but a negligible number of decoys are selected from the spam era.  
 125 We will hold constant the non-spam transactions and vary the number of spam transactions and the ring

126 size. Figures 5, 6, and 7 show the results of the simulations.

Figure 5: Long-term projected mean effective ring size






Figure 6: Long-term projected mean effective ring size (log-log scale)





Figure 7: Long-term projected share of rings with effective ring size 1



### 127 3.3 Guessing the real spend using a black marble flooder's simple classifier


128 The adversary carrying out a black marble flooding attack could use a simple classifier to try to guess the  
 129 real spend: Let  $n$  be nominal ring size and  $n_s$  be the number of outputs in a given ring that are owned  
 130 by the attacker.  $n_s$  is a random variable because decoy selection is a random process. The adversary  
 131 can eliminate  $n_s$  of the  $n$  ring members as possible real spends. The attacker guesses randomly with  
 132 uniform probability that the  $i$ th ring member of the  $n - n_s$  remaining ring members is the real spend. The  
 133 probability of correctly guessing the real spend is  $\frac{1}{n - n_s}$ . If the adversary owns all ring members except  
 134 for one ring member, which must be the real spend, the probability of correctly guessing the real spend  
 135 is 100%. If the adversary owns all except two ring members, the probability of correctly guessing is 50%.  
 136 And so forth.

137 The mean effective ring size is  $E[n_e]$  from 3. Does this mean that the mean probability of correctly  
 138 guessing the real spend is  $\frac{1}{E[n_e]}$ ? No. The  $h(x) = \frac{1}{x}$  function is strictly convex. By Jensen's inequality,  
 139  $E\left[\frac{1}{n_e}\right] > \frac{1}{E[n_e]}$ . The mean probability of correctly guessing the real spend is

$$E\left[\frac{1}{n_e}\right] = \sum_{i=1}^n \frac{1}{i} \cdot f\left(i - 1, n - 1, \frac{E[n_e] - 1}{n - 1}\right) \quad (4)$$


140  $\frac{1}{i}$  is the probability of correctly guessing the real spend when the effective ring size is  $i$ .  $f$  is the  
 141 probability mass function of the binomial distribution. It calculates the probability of the decoy selection  
 142 algorithm selecting  $i - 1$  decoys that are owned by real users. The total number of decoys to select is  $n - 1$   
 143 (that is the argument in the second position of  $f$ ). The probability of selecting a decoy owned by a real  
 144 user is  $\frac{E[n_e] - 1}{n - 1} = \frac{r}{r + s}$ .

Figure 8: Estimated probability of correctly guessing the real spend



145 The probability of a given ring having all adversary-owned ring members except for the real spend is  
 146  $f\left(0, n - 1, \frac{E[n_e] - 1}{n - 1}\right)$ . Figure 9 plots the estimated share of rings with effective ring size one.

Figure 9: Estimated share of rings with effective ring size of one



## 147 4 Chain reaction graph attacks

148 The effective ring size can be reduced further by applying a process of elimination to related rings. This  
149 technique is called a “chain reaction” or a “graph analysis attack”. Say that the effective ring size in  
150 transaction  $A$  is reduced to two because of a black marble attack. One of the remaining two ring members  
151 is an output in transaction  $B$ . If the output in transaction  $B$  is known to be spent in transaction  $C$   
152 because the effective ring size of transaction  $C$  was one, then that output can be ruled out as a plausible  
153 real spend in transaction  $A$ . Therefore, the adversary can reduce the effective ring size of transaction  $A$   
154 to one.

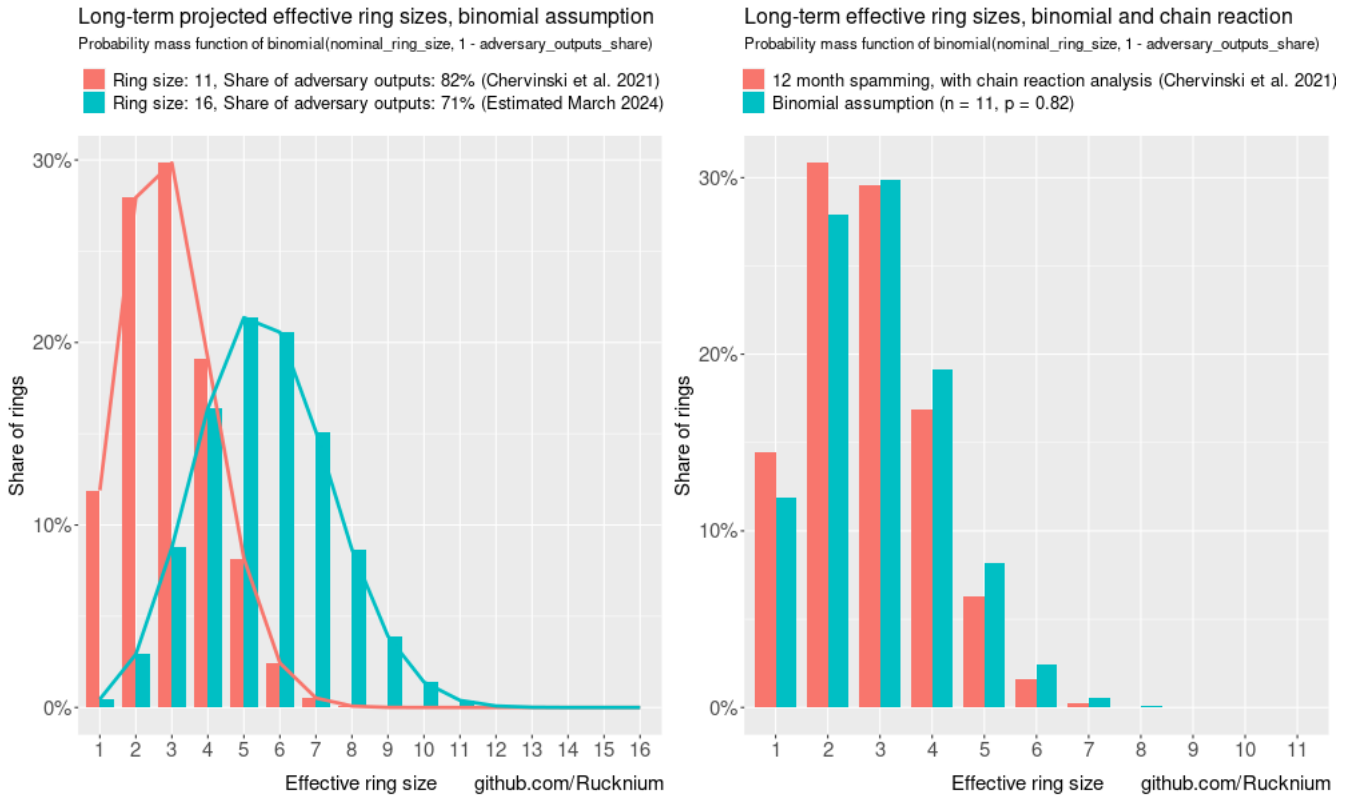
155 Theorem 1 of [Yu et al., 2019] says that a “closed set” attack is as effective as exhaustively checking  
156 all subsets of outputs. The brute force attack is infeasible since its complexity is  $O(2^m)$ , where  $m$  is the  
157 total number of RingCT outputs on the blockchain. [Yu et al., 2019] implements a heuristic algorithm to

158 execute the closed set attack that is almost as effective as the brute force method. [Vijayakumaran, 2023]  
 159 proves that the Dulmage-Mendelsohn (DM) decomposition gives the same results as the brute force closed  
 160 set attack, but the algorithm renders a result in polynomial time. The open source implementation of the  
 161 DM decomposition in [Vijayakumaran, 2023] processes 37 million RingCT rings in about four hours.

162 In practice, how much further can chain reaction attacks reduce the effective ring size when combined  
 163 with a black marble attack? [Egger et al., 2022] suggest some closed-form formulas to compute the vulner-  
 164 ability of different ring sizes to chain reaction attacks. However, [Egger et al., 2022] assume that decoys  
 165 are selected by a partitioning process instead of Monero’s actual mimicking decoy selection algorithm.  
 166 It is not clear how relevant the findings of [Egger et al., 2022] are for Monero’s mainnet. Monte Carlo  
 167 simulations would be a better way to evaluate the risk of chain reactions.

168 [Chervinski et al., 2021] carries out a simulation using the old ring size of 11. In the 2input/2output  
 169 spam scenario, 82% of outputs are black marbles. Assuming only the binomial distribution, i.e. no  
 170 chain reaction analysis, Figure 10 compares the theoretical long-term distribution of effective ring sizes  
 171 in the [Chervinski et al., 2021] scenario and the March 2024 suspected spam on Monero’s mainnet. The  
 172 share of rings with effective ring size 1 in the [Chervinski et al., 2021] scenario is 11.9 percent, but the  
 173 share is only 0.8 percent with the suspected March 2024 spam. The mean effective ring sizes of the  
 174 [Chervinski et al., 2021] scenario without chain reaction and the March 2024 spam estimate are 2.9 and  
 175 5.2, respectively.

Figure 10: Probability mass function of long-term effective ring sizes



176 [Chervinski et al., 2021] executes chain reaction analysis to increase the effectiveness of the attack. The  
177 second plot in Figure 10 compares the long term effective ring size achieved by [Chervinski et al., 2021]  
178 when leveraging chain reaction analysis and the effective ring size when only the binomial distribution is  
179 assumed. [Chervinski et al., 2021] increases the share of ring with effective ring size one from 11.9 to 14.5  
180 percent. Mean effective ring size decreases from 2.94 to 2.76. This is a modest gain of attack effectiveness,  
181 but [Chervinski et al., 2021] appears to be using a suboptimal chain reaction algorithm instead of the  
182 closed set attack.

183 The actual risk from chain reaction analysis in the suspected March 2024 flooding is a gap in our  
184 knowledge. [Vijayakumaran, 2023] provides an open source implementation of the DM decomposition in  
185 Rust and excellent documentation.<sup>1</sup> A Monte Carlo simulation applying the DM decomposition to the  
186 March 2024 black marble estimates should be written.

## 187 **5 Countermeasures**

188 See <https://github.com/monero-project/research-lab/issues/119>

189 TODO

## 190 **6 Estimated cost to suspected spammer**

191 When the `lin/2out 20` nanoneros/byte spam definition is used, the total fees paid by the spam transactions  
192 over the 23 days of spam was 61.5 XMR. The sum total of the transaction sizes of the spam transactions  
193 was 3.08 GB.

194 When the `lin/2out 20 or 320` nanoneros/byte spam definition is used, the total fees paid by the spam  
195 transactions over the 23 days of spam was 81.3 XMR. The sub total of the transaction sizes of the spam  
196 transactions was 3.12 GB.

## 197 **7 Transaction confirmation delay**

198 Monero’s transaction propagation rules are different from BTC’s rules for good reasons, but two of the  
199 rules can make transactions seem like they are “stuck” when the txpool (mempool) is congested. First,  
200 Monero does not have replace-by-fee (RBF). When a Monero node sees that a transaction attempts to  
201 spend an output that is already spent by another transaction in the txpool, the node does not send the  
202 transaction to other nodes because it is an attempt to double spend the output. (Monero nodes do not  
203 know the real spend in the ring, but double spends can be detected by comparing the key images of  
204 ring signatures in different transactions.) Monero users cannot increase the fee of a transaction that they

---

<sup>1</sup><https://github.com/avras/cryptonote-analysis>  
<https://www.respectedsir.com/cna>

205 already sent to a node because the transaction with the higher fee would be considered a double spend.  
206 BTC has RBF that allows a transaction to replace a transaction in the mempool that spends the same  
207 output if the replacement transaction pays a higher fee. One of RBF’s downsides is that merchants cannot  
208 safely accept zero-confirmation transactions because a malicious customer can replace the transaction in  
209 the mempool with a higher-fee transaction that spends the output back to themselves. Without RBF,  
210 Monero users must wait for their low-fee transaction to confirm on the blockchain. They cannot choose to  
211 raise their “bid” for block space even if they were willing to pay more. They have to get it right the first  
212 time. Fee prediction is especially important for Monero users when the txpool is congested because of the  
213 lack of RBF, but very little Monero-specific fee prediction research has been done.

214 Unlike BTC, Monero also does not have child-pays-for-parent (CPFP), which allows users to chain  
215 multiple transactions together while they are still in the mempool. With CPFP, users can spend the  
216 output of the unconfirmed parent transaction and attach a higher fee to the child transaction. Miners  
217 have an incentive to include the parent transaction in the block because the child transaction is only  
218 valid if the parent transaction is also mined in a block. Monero transaction outputs cannot be spent in  
219 the same block that they are confirmed in. Actually, Monero users need to wait at least ten blocks to  
220 spend new transaction outputs because benign or malicious blockchain reorganizations can invalidate ring  
221 signatures.<sup>2</sup>

222 Monero’s transaction propagation rules can create long delays for users who pay the same minimum  
223 fee that the suspected spammer pays. When users pay the same fee as the spam, their transactions are  
224 put in a “queue” with other transactions at the same fee per byte level. Their transactions are confirmed  
225 in first-in/first-out order because the `get_block_template` RPC call to `monerod` arranges transactions  
226 that way.<sup>3</sup> Most miners use `get_block_template` to construct blocks, but P2Pool orders transactions  
227 randomly after they have been sorted by fee per byte.<sup>4</sup>

228 The first plot in Figure 11 shows the mean delay of transaction confirmation in each hour. The  
229 plot shows the mean time that elapsed between when the transaction entered the txpool and when it was  
230 confirmed in a block. Each hour’s value in the line plot is computed from transactions that were confirmed  
231 in blocks in that hour. This data is based on txpool archive data actively collected from a few nodes.<sup>5</sup>  
232 The mean includes transactions with and without the spam fingerprint. Usually mean confirmation time  
233 was less than 30 minutes, but sometimes confirmations of the average transaction were delayed by over  
234 two hours.

---

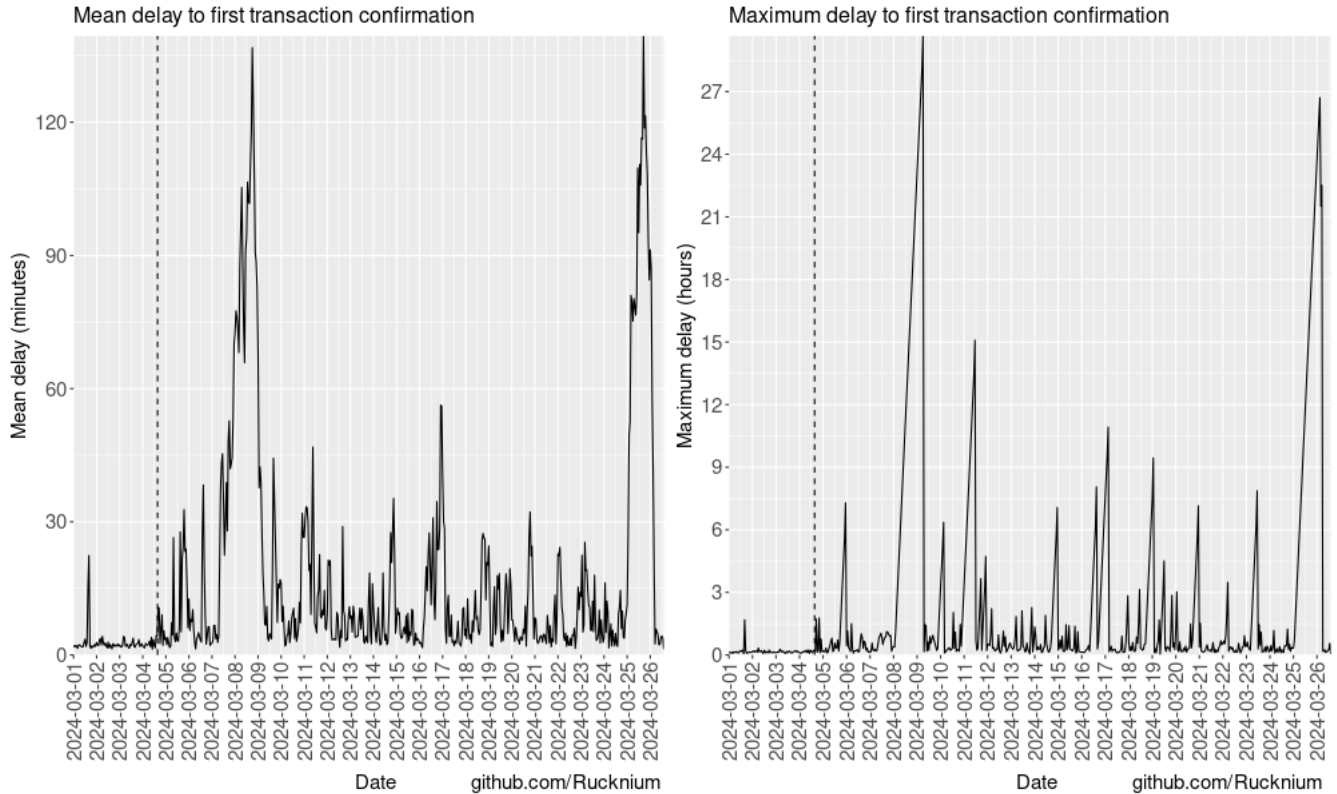
<sup>2</sup>“Eliminating the 10-block-lock” <https://github.com/monero-project/research-lab/issues/95>

<sup>3</sup>[https://github.com/monero-project/monero/blob/9bf06ea75de4a71e3ad634e66a5e09d0ce021b67/src/cryptonote\\_core/tx\\_pool.cpp#L1596](https://github.com/monero-project/monero/blob/9bf06ea75de4a71e3ad634e66a5e09d0ce021b67/src/cryptonote_core/tx_pool.cpp#L1596)

<sup>4</sup>[https://github.com/SChernykh/p2pool/blob/dd17372ec0f64545311af40b976e6274f625ddd8/src/block\\_template.cpp#L194](https://github.com/SChernykh/p2pool/blob/dd17372ec0f64545311af40b976e6274f625ddd8/src/block_template.cpp#L194)

<sup>5</sup><https://github.com/Rucknium/misc-research/tree/main/Monero-Mempool-Archive>

Figure 11: Delay to first transaction confirmation



235 The second plot in Figure 11 shows the *maximum* waiting time for a transaction to be confirmed. The  
236 value of the line at each hour is the longest time that a transaction waited to be confirmed in one of the  
237 block mined in the hour or the amount of time that a transaction was still waiting to be confirmed at the  
238 end of the hour (whichever is greater). There were a handful of transactions that paid fees below the 20  
239 nanoneros/byte tier that the spam was paying. These transactions did not move forward in the queue when  
240 the spam transactions were confirmed. Instead, they had to wait until the txpool completely emptied.  
241 Exactly 100 transactions waited longer than three hours. They paid between 19465 and 19998 piconeros  
242 per byte. Most of the transactions appeared to have set fees slightly lower than 20 nanonerpos per byte  
243 because they had an unusual number of inputs. 92 of them had four or more inputs. The remaining eight  
244 of them had just one input. Those eight may have been constructed by a nonstandard wallet.

## 245 8 Real user fee behavior

246 During the suspected spam, users must pay more than the minimum fee to put their transactions at the  
247 front of the confirmation queue. If users pay more than the minimum fee, usually their transactions would  
248 be confirmed in the next mined block. Monero's standard fee levels are 20, 80, 320, and 4000 nanoneros  
249 per byte. Users are not required to pay one of these fee levels, but all wallets that are based on `wallet2`  
250 do not allow users to choose custom fees outside of the four standard levels because of the privacy risk of



251 unusual transactions.<sup>6</sup>

252 The “auto” fee level of the Monero GUI and CLI wallets is supposed to automatically change the fee  
253 of a transaction from the lowest tier (20 nanoneros/byte) to the second tier (80 nanoneros/byte) when the  
254 txpool is congested. Unfortunately, a bug prevented the automatic adjustment. On March 9, 2024 the  
255 Monero Core Team released the 0.18.3.2 version of Monero and the GUI/CLI wallet that fixed the bug.<sup>7</sup>  
256 Users are not required to upgrade to the latest wallet version, so probably many users still use the version  
257 that is not automatically adjusting fees.

258 The first plot of Figure 12 shows the share of transactions paying each of the four fee tiers. Any  
259 transactions that do not pay in the standard ranges  $\{[18, 22], [72, 82], [315, 325], [3000, 4100]\}$  were not  
260 included in the plot. The 320 nanoneros/byte tier is interesting. About 10 percent of transactions paid  
261 320 nanonero/byte until February 17, 2024. The date could have something to do with Monero being  
262 delisted from Binance on February 20, 2024.<sup>8</sup> Then on March 12-13, 2024 there was a burst of 320  
263 nanonero/byte transactions. The 0.18.3.2 GUI/CLI wallet release could not explain the burst since the  
264 auto fee adjustment would only increase fees from 20 to 80 nanoneros/byte. The burst of 320 nanonero/byte  
265 transactions must have been either from a central service producing fees or from the suspected spammer.

266 The second plot of Figure 12 shows the same data with the suspected spam transactions eliminated  
267 both the 80 and 320 nanoneros/byte transactions with the spam fingerprint were removed. There is a  
268 modest increase in 80 nanonero/byte transactions after the spam started.

---


<sup>6</sup><https://github.com/Rucknium/misc-research/tree/main/Monero-Nonstandard-Fees>

<sup>7</sup>“Monero 0.18.3.2 ‘Fluorine Fermi’ released” <https://www.getmonero.org/2024/03/09/monero-0.18.3.2-released.html>

“wallet2: adjust fee during backlog, fix set priority” <https://github.com/monero-project/monero/pull/9220>

<sup>8</sup><https://decrypt.co/218194/binance-finalizes-monero-delisting>

Figure 12: Share of transactions by fee tier



269 The mempool archive data suggest that merchants using zero-confirmation delivery were still safe  
 270 during the spam incident. Once submitted to the network, transactions did not drop out of the mempool.  
 271 They just took longer to confirm. There were only two transaction IDs in the mempool of one of the  
 272 mempool archive nodes that did not confirm during the spam period. Both occurred on March 8 when the  
 273 mempool was very congested. The the two “disappearing transactions” could happen if someone submits  
 274 a transactions to an overloaded public RPC node, the transactions does not propagate well, and then the  
 275 user reconstructs the transactions with another node. The first transaction will not confirm because it  
 276 is a double spend. Seeing a transaction in the mempool that never confirms happens sometimes during  
 277 normal transaction volumes, too. Single transactions like that appeared on February 14, 17, and 23 and  
 278 March 1 in the mempool archive data.


## 279 9 Evidence for and against the spam hypothesis

280 Is the March 4, 2024 transaction volume a result of many real users starting to use Monero more, or is  
 281 it spam created by a single entity? [Krawiec-Thayer et al., 2021] analyzed the July/August 2021 sudden  
 282 rise in transaction volume. We concluded that it was likely spam. Our evidence was: 1) There was a  
 283 sharp increase of 1in/2out and 2in/1out transactions, but the volume of other transaction types did not  
 284 increase, 2) All the suspected spam paid minimum fees, 3) The distribution of ring members became much

285 younger, suggesting that the spammer was rapidly re-spending outputs as quickly as possible.

286 Available time has not permitted a full run of the [Krawiec-Thayer et al., 2021] analysis on the March  
287 2024 suspected spam data. It is easy to do a quick check of transaction volume by input/output type.  
288 Figure 13 plots the eight most common in/out transaction types on a log scale. Only the volume of  
289 1in/2out transactions increased on March 4, supporting the spam hypothesis.

Figure 13: Transaction volume by number of inputs and outputs (log scale)



290 More can be done to generate evidence for or against the spam hypothesis. [Krawiec-Thayer et al., 2021]  
291 analyzed the age of all ring members. Using the OSPEAD techniques, the distribution of the age of the  
292 real spends can be estimated.<sup>9</sup> The Monero node network can be actively crawled to see if the spam  
293 transactions originate from one node. Dandelion++ can defeat attempts to discover the origin of most  
294 transaction because the signal of the real transaction is covered by the Dandelion++ noise. When the  
295 signal is huge like the spam, some statistical analysis could overcome the Dandelion++ protection. In-

<sup>9</sup><https://github.com/Rucknium/OSPEAD>

296 vestigatory nodes could use `set_log net.p2p.msg:INFO` to view which neighboring nodes the suspected  
297 spam is coming from. Then the investigatory node could crawl the network in the direction of the high-  
298 est incoming volume. The techniques of [Sharma et al., 2022] are useful at extremely high transaction  
299 volumes, like in the spam case, and could be used.

## 300 References

301 [Aguado et al., 2010] Aguado, J., Cid, C., Saiz, E., & Cerrato, Y. (2010). Hyperbolic decay of the dst  
302 index during the recovery phase of intense geomagnetic storms. *Journal of Geophysical Research: Space*  
303 *Physics*, 115(A7). <https://doi.org/https://doi.org/10.1029/2009JA014658>

304 [Chervinski et al., 2021] Chervinski, O. J., Kreutz, D., & Yu, J. (2021). Analysis of transaction flood-  
305 ing attacks against monero. *2021 IEEE International Conference on Blockchain and Cryptocurrency*  
306 *(ICBC)*, 1–8. <https://doi.org/10.1109/ICBC51069.2021.9461084>

307 [Egger et al., 2022] Egger, C., Lai, R. W. F., Ronge, V., Woo, I. K. Y., & Yin, H. H. F. (2022). On  
308 defeating graph analysis of anonymous transactions. *Proceedings on Privacy Enhancing Technologies*,  
309 2022(3). <https://petsymposium.org/2022/files/papers/issue3/popets-2022-0085.pdf>

310 [Krawiec-Thayer et al., 2021] Krawiec-Thayer, M. P., Neptune, Rucknium, Jberman,  
311 & Carrington (2021). *Fingerprinting a flood: forensic statistical analysis of the*  
312 *mid-2021 monero transaction volume anomaly*. [https://mitchellpkt.medium.com/](https://mitchellpkt.medium.com/fingerprinting-a-flood-forensic-statistical-analysis-of-the-mid-2021-monero-transaction-volume-a19cbf41ce60)  
313 [fingerprinting-a-flood-forensic-statistical-analysis-of-the-mid-2021-monero-transaction-volume-a](https://mitchellpkt.medium.com/fingerprinting-a-flood-forensic-statistical-analysis-of-the-mid-2021-monero-transaction-volume-a19cbf41ce60)  
314 Available at [https://mitchellpkt.medium.com/fingerprinting-a-flood-forensic-statistical-analysis-of-the-](https://mitchellpkt.medium.com/fingerprinting-a-flood-forensic-statistical-analysis-of-the-mid-2021-monero-transaction-volume-a19cbf41ce60)  
315 [mid-2021-monero-transaction-volume-a19cbf41ce60](https://mitchellpkt.medium.com/fingerprinting-a-flood-forensic-statistical-analysis-of-the-mid-2021-monero-transaction-volume-a19cbf41ce60)

316 [Noether et al., 2014] Noether, S., Noether, S., & Mackenzie, A. (2014). *A note on chain reactions in trace-*  
317 *ability in cryptonote 2.0*. Research Bulletin. [https://www.getmonero.org/resources/research-lab/](https://www.getmonero.org/resources/research-lab/pubs/MRL-0001.pdf)  
318 [pubs/MRL-0001.pdf](https://www.getmonero.org/resources/research-lab/pubs/MRL-0001.pdf)

319 [Ronge et al., 2021] Ronge, V., Egger, C., Lai, R. W. F., Schröder, D., & Yin, H. H. F. (2021). Foundations  
320 of ring sampling. *Proceedings on Privacy Enhancing Technologies*, 2021(3), 265–288. [https://doi.org/](https://doi.org/doi:10.2478/popets-2021-0047)  
321 [doi:10.2478/popets-2021-0047](https://doi.org/doi:10.2478/popets-2021-0047)

322 [Rucknium, 2023] Rucknium (2023). *Closed-form expression of monero’s wallet2 de-*  
323 *coy selection algorithm*. [https://github.com/Rucknium/misc-research/tree/main/](https://github.com/Rucknium/misc-research/tree/main/Monero-Decoy-Selection-Closed-Form/pdf)  
324 [/Monero-Decoy-Selection-Closed-Form/pdf](https://github.com/Rucknium/misc-research/tree/main/Monero-Decoy-Selection-Closed-Form/pdf). Available at [https://github.com/Rucknium/misc-](https://github.com/Rucknium/misc-research/tree/main/Monero-Decoy-Selection-Closed-Form/pdf)  
325 [research/tree/main/Monero-Decoy-Selection-Closed-Form/pdf](https://github.com/Rucknium/misc-research/tree/main/Monero-Decoy-Selection-Closed-Form/pdf)

- 326 [Sharma et al., 2022] Sharma, P. K., Gosain, D., & Diaz, C. (2022). *On the anonymity of peer-to-peer*  
327 *network anonymity schemes used by cryptocurrencies*. <https://doi.org/10.48550/ARXIV.2201.11860>
- 328 [Vijayakumaran, 2023] Vijayakumaran, S. (2023). Analysis of cryptonote transaction graphs using the  
329 dulmage-mendelsohn decomposition. *5th Conference on Advances in Financial Technologies (AFT*  
330 *2023)*, volume 282 of *Leibniz International Proceedings in Informatics (LIPIcs)*. [https://aftconf.](https://aftconf.github.io/aft23/program.html)  
331 [github.io/aft23/program.html](https://aftconf.github.io/aft23/program.html)
- 332 [Yu et al., 2019] Yu, Z., Au, M. H., Yu, J., Yang, R., Xu, Q., & Lau, W. F. (2019). New empirical  
333 traceability analysis of cryptonote-style blockchains. *Financial Cryptography and Data Security*, 133–  
334 149. [https://link.springer.com/chapter/10.1007/978-3-030-32101-7\\_9](https://link.springer.com/chapter/10.1007/978-3-030-32101-7_9)