

1 Article

2 *k*-root-*n*: An Efficient Algorithm for Avoiding Short 3 Term Double-Spending Alongside Distributed 4 Ledger Technologies Such as Blockchain

5 Zvi Schreiber

6 Individual, London, UK; zvi@zvi.net

7 **Abstract** Blockchains such as the bitcoin blockchain depend on reaching a global consensus on the
8 distributed ledger; therefore, they suffer from well know scalability problems. This paper proposes
9 an algorithm that avoids double-spending in the short term with just $O(\sqrt{n})$ messages; each node
10 receiving money off-chain performs the due diligence of consulting $k\sqrt{n}$ random nodes to check if
11 any of them is aware of double-spending. Two nodes receiving double-spent money will in this way
12 consult at least one common node with very high probability, due to the 'birthday paradox', and
13 any common honest node consulted will detect the fraud. Since the velocity of money in the real
14 world has coins circulating through at most a few wallets per day, the size of the due diligence
15 communication is small in the short term. This '*k*-root-*n*' algorithm is suitable for an environment
16 with synchronous or asynchronous (but with fairly low latency) communication and with Byzantine
17 faults. The presented *k*-root-*n* algorithm should be practical to avoid double-spending with
18 arbitrarily high probability, while feasibly coping with the throughput of all world commerce. It is
19 resistant to Sybil attacks even beyond 50% of nodes. In the long term, the *k*-root-*n* algorithm is less
20 efficient. Therefore, it should preferably be used as a complement and not a replacement to a global
21 distributed ledger technology.

22 **Keywords:** blockchain; bitcoin; cryptocurrency; distributed ledger technology

23

24 1. Introduction

25 In blockchains such as bitcoin, all n nodes reach Nakamoto consensus [1] on each block of
26 transactions, thereby creating a scalability problem [2] [3] [4] that notoriously limits the entire bitcoin
27 network to a few transactions per second while consuming massive power [5]. Bitcoin is considered
28 the first digital currency algorithm to solve the double-spending problem without the need for a
29 trusted authority or central server. Additionally, it can cope with Byzantine faults (e.g. [6] [7])
30 including a Sybil attack [8] of up to 50% dishonest nodes. However, it requires $O(n)$ communication
31 messages per transaction that limit its scale.

32 Practically, bitcoin transactions suffer from a lag time of 15 minutes to several hours before being
33 included in a block on the bitcoin blockchain [9] (this lag time has a complex dependency on how
34 high a fee is offered by the transaction participants to the miner [10]), and then an hour longer to
35 reach the generally desired threshold of 6-block confirmation [11]. Thus, before received bitcoin
36 transfers are confirmed and are safe to re-spend, there is typically a several hours latency.

37 At the time of writing, the typical fee paid to the miner for a single bitcoin transaction is tens of
38 thousands of Satoshi or about US\$0.50-\$5 [9]; this is not cheaper than most domestic bank
39 transactions with fiat currency.

40 Therefore, efforts are being made to redesign the blockchain algorithm itself to ensure greater
41 scalability, such as SCP [12], Algorand [13], bitcoin next generation (bitcoin-NG) [14], which all

42 involve selecting a subset of users (committees or a rotating leader) in various configurations to
43 reduce the number of messages required to reach consensus. In an alternative approach, a subset of
44 nodes transacts with each other off-chain for a time, [15] as in the lightning network [16].

45 In this paper we consider an approach for protecting against double-spending [17] possibly
46 combined with a Sybil attack without the need for a global ledger consensus. In this paper, we did
47 not consider other forms of attacks, such as eclipse attacks [18], routing attacks [19], attacks based on
48 time advantage [20], incentive attacks [21] and quantum computing attacks [22]. Relevant surveys of
49 other attack types are [23] [24]. Some previously research on avoiding double-spending are [25] [26].
50 This study focuses on the protection against the most central vulnerability of cryptocurrencies,
51 namely double-spending covered up by a Sybil attack of malicious nodes.

52 We propose a scalable low-latency algorithm that can run off-chain in parallel to a global ledger
53 consensus mechanism, such as blockchain, protecting against double-spending in the short term.
54 Thus, commerce may continue at a high pace even while the n nodes are working to reach consensus
55 on transactions possibly with a lag of some hours from the transaction time. By applying this
56 algorithm, we accept a situation where consensus is infrequently achieved. Therefore, we can accept
57 longer blockchain blocks that are created every hour, or every few hours, instead of bitcoin's current
58 average of 10 minutes, thereby increasing blockchain's throughput of transactions per second [27],
59 while compensating for longer latency with our complementary off-chain algorithm for preventing
60 short-term double-spending.

61 For example, each morning the nodes may reach a consensus on the valid transaction histories
62 and wallet balances as of the preceding midnight Greenwich Mean Time, and they may
63 asynchronously do so, reaching the consensus by, say, 6 a.m. the next morning. For example, in the
64 particular case of bitcoin, all the transactions by 6 a.m. from the previous day would typically have
65 achieved six-block verification and may be considered final. In this case, the role of our algorithm is
66 to allow fast and safe transactions in the 30 hours say from Sunday midnight to Tuesday 6 am when
67 consensus is finalized for the ledger as of Monday midnight. Thus, in this example, to reach a
68 consensus only every 24 hours with a 6-hour lag, we allow for a situation where the distributed ledger
69 is relaxed relative to the current configuration of bitcoin. Therefore, this enables the transaction rate
70 of the ledger to increase. Furthermore, our proposed solution can allow people to trade, particularly
71 to safely pass on the received coins with next to zero latency.

72 The proposed algorithm, which is called ' $k\sqrt{n}$ ' or ' k -root- n ', avoids double-spending in the short
73 to medium term, while there is no global ledger consensus with an arbitrarily high probability of
74 detecting double-spending requiring just $O(\sqrt{n})$ messages per transaction. This is based on the
75 assumption that specific money balances only circulate through $O(\text{constant})$ wallets in 24 hours. This
76 assumption is realistic since money circulates in the real economy with a velocity measured in one or
77 two transactions per month [28], and bitcoin is already practically constrained by the transaction
78 confirmation lag times to circulate a few times per day, and in practice rarely more than once or twice
79 a day.

80 In this algorithm, every transaction should eventually be on-chain. The initial transaction
81 verification is off-chain, thereby allowing transactions to continue off-chain at high speed and waiting
82 for the blockchain to catch up. The algorithm only involves $O(\sqrt{n})$ nodes and messages per
83 transaction; we typically choose $10\sqrt{n}$.

84 2. Overview of $k\sqrt{n}$ random double-spending detection

85 Suppose there are n nodes, in which each node is also a wallet, connected to a network, and they
86 achieved consensus on the global distributed ledger (or at least on the balance of each node) using
87 blockchain (or another algorithm) sometime recently, a time we shall call the global ledger consensus
88 (in the example above that would occur at 6 a.m. daily).

89 For now, we assume that every wallet is also a node, which is usually online and available, and
90 which also provides basic verification services to the network. The central idea is that any honest
91 node that wants to verify whether the funds it receives have not been double-spent can demand that
92 the sender disclose the pedigree of the transferred funds, namely the sender's transaction history
93 since the last global ledger consensus. In case the sender depends on incoming funds to have
94 sufficient balance to cover the transaction, the receiver can recursively demand disclosure of the
95 source of funds right back to funds that were available as of the last global ledger consensus.

96 Since querying all the nodes to verify each transaction is prohibitively expensive, an honest node
97 will perform its due diligence on the pedigree of each inbound transaction with a random $k\sqrt{n}$ other
98 nodes. If two nodes query $k\sqrt{n}$ random nodes, we will show that the probability of zero common
99 nodes is extremely small for suitable k even if a substantial proportion of nodes are failing or
100 malicious. Therefore, if two honest nodes receive the same double-spent coins, they may consult at
101 least one common node and detect the fraud with a very high probability. This is the main idea of the
102 algorithm, and it is based on the famous birthday paradox [29], where for example just 40 people (\approx
103 $2\sqrt{365}$) have about a 90% chance that at least two of them have the same birthday.

104 Each honest node provides verification services by keeping a history of the transaction pedigrees
105 it was asked to verify; when two honest nodes query random nodes, any common queried honest
106 node can immediately raise the alarm if the two receiving nodes are victims of a double-spending
107 attempt, i.e. if they were given inconsistent transaction histories.

108 Let k be a small number greater than 1. We will generally choose $k = 10$. Assume that we are in
109 an environment with Byzantine faults, say about 10% of nodes may not respond at any time due to
110 node or network failure; assume that about 50% of the nodes are malicious, we would then have an
111 effective $\underline{k} = 4.5$, i.e. $\underline{k}\sqrt{n}$ responsive and honest nodes. When each of two honest nodes receives funds
112 and successfully each query $4.5\sqrt{n}$ honest nodes, this $\underline{k} = 4.5$ is sufficient to ensure an expected value
113 of more than twenty common, honest and responsive nodes. There is a probability of just
114 approximately 10^{-9} of zero common, honest and responsive nodes. Therefore, the chances of getting
115 away with double-spending are negligible, and there is a probability very close to 1 that any double-
116 spending can be detected as soon as both branches of the spend reach honest nodes.

117 The penalty for double-spending is forfeiting the wallet, so if each wallet has a minimum stake
118 m of \$1 and each transaction is limited to well under \$1 billion, say to a maximum $M = \$1,000,000$,
119 then there is a negative expected return from any double-spending attempt since there is a probability
120 of just approximately 10^{-9} of not being caught.

121 A dishonest node may not be checking its inbound transactions or may be maliciously
122 collaborating with other nodes. This is why the receiving honest node should check not only the
123 transaction history of the immediate sender for forked history/double-spending, but also to
124 recursively check any of sender's sender's transactions, to the extent that the immediate sender
125 depends on the sender's sender (recursively) payment to have balance for covering the current
126 transaction. This recursive tree of inbound transactions is called the pedigree of the transaction, that

127 is the recursive list of transactions that it depends on. This recursion is why the k -root- n algorithm is
 128 less efficient for long term use since the recursive pedigree of transactions may become large over a
 129 long time period.

130 As a motivation for the $O(\sqrt{n})$ algorithm, we briefly explore how a $10\sqrt{n}$ algorithm scales by
 131 assuming $n = 10$ billion people (the projected world population for 2050 [30] and much more than
 132 bitcoin's current 32m wallets [31]). Suppose people are each transacting once per hour on the average,
 133 i.e. 24-hours per day (higher than the average rate of commerce). Each transaction involves messages
 134 to $10\sqrt{n} = 10^6$ nodes. We will see that this gives a probability of just $p \approx 10^{-9}$ of getting away with
 135 double-spending even if half of the nodes are fraudulent and 10% of the nodes are unavailable (i.e.
 136 $4.5\sqrt{n}$ honest, responsive validating nodes). Thus, each transaction only burdens 1 out of 10,000
 137 nodes, and with 10 billion transactions per hour globally, or 2.77million transactions per second
 138 globally, each node should be involved in just 278 transactions per second. This transaction
 139 throughput is feasible for a modern computer (especially in 2050).

140 Thus, it seems practical that the algorithm could securely handle not only Visa/Mastercard
 141 volumes, but in fact all the commerce in today's world and the foreseeable future. Visa's volumes
 142 have been widely misquoted in bitcoin articles as 24,000 per second, although that appears to be
 143 mythical [32] with apparently more reliable sources estimating about 78.95 billion Visa transactions
 144 in the first half of 2018 [33] which averages 5,000 per second, although peak time transaction rates
 145 would presumably be higher. In any event, the current algorithm could feasibly handle transaction
 146 volumes orders of magnitude larger than Visa.

147 Whilst the idea of depending on probability to secure commerce may at first seem strange, it is
 148 noteworthy that all commerce already depends on probability. For example, every credit card
 149 transaction is accepted based on a probabilistic evaluation that it isn't fraudulent.

150 We now introduce some definitions and formally present the k -root- n algorithm.

151 3. Preliminaries

152 **Definition 1.** A transaction $T = (x, t, S, R, s_s, s_r)$ is an agreement to transfer a balance from a sender
 153 node/wallet to a receiver node/wallet; the tuple comprises a positive amount $x = x[T]$, a time stamp t
 154 $= t[T]$, identifiers (public keys) of the sender user $S = S[T]$ and the receiver user $R = R[T]$.
 155 Additionally, s_s and s_r are the respective digital signatures of S and R of the data tuple (x, t, S, R) .

156 A transaction T is only valid if the sender S had a balance (see next definition) of at least $m + x$
 157 immediately before the transaction, where m denotes the agreed minimum wallet balance. No sender
 158 or receiver can participate in two transactions with identical time stamps t .

159 A potential transaction \underline{T} is a transaction that has not yet been signed by the receiver

160 $\underline{T} = (x, t, S, R, s_s)$. \square

161 **Definition 2.** The balance $b[u, t_1]$ of a user u at time t_1 , given that s/he had a balance of b_0 at the
 162 time t_0 of the last known global ledger consensus, is defined by

$$163 \quad b_0 + \sum_{t_0 < t[T_i] < t_1, R[T_i]=u} b[T_i] - \sum_{t_0 < t[T_i] < t_1, S[T_i]=u} b[T_i],$$

164 i.e. the last known global ledger consensus balances plus all received amounts, minus all spent
 165 amounts. \square

166 **Definition 3.** The Lineage $LIN[T]$ of a transaction or potential transaction T is the transaction
 167 history for the sender $u[T]$ from the last known global ledger consensus at time t_0 before $t[T]$ and up
 168 to the time of T :

$$169 \quad LIN[T] = \{T_i \mid t_0 < t[T_i] < t[T], S[T_i] = S[T] \vee R[T_i] = S[T]\}. \square$$

170 These are the transactions relevant to establishing that the sender u has sufficient balance to
 171 afford T . However, not all of this history is necessarily required to establish sufficient balance, so for
 172 the sake of efficiency we now introduce a narrower transaction history.


173 **Definition 4.** The Critical Lineage $CLIN[T]$ of a transaction or potential transaction T is a set of
 174 transactions whose elements are a subset of $LIN[T]$, comprising a minimal subset of inbound
 175 transactions critical to provide the balance that allows the sender to afford T . Formally, suppose that
 176 a sender S makes a payment of amount x in a transaction or potential transaction T ; suppose that S 's
 177 last known global ledger consensus balance was b_0 . Suppose that the set of transactions in which S
 178 participated since the last global ledger consensus, $LIN[T]$, includes inbound payments, $T_1...T_n$, in
 179 descending order of amount (and ordered chronologically when amounts are equal) and outbound
 180 payments, $U_1...U_m$. Now, the critical inbound payments are the subset $CLIN[T] = \{T_1...T_j\}$ of inbound
 181 payments with j minimal such that $b + \sum_i^j T_i - \sum_i^m U_i \geq x + m$. \square

182 Thus, given that the sender has opening balance b and has spent the U_i , then, $\{T_1...T_j\}$ is a minimal
 183 subset of inbound transactions that are enough to provide balance coverage for this payment of x .
 184 Even if any of the other inbound payments, $T_{j+1}...T_n$, is derived directly or indirectly from fraud, the
 185 validation by the receiver of these critical inbound payments $CLIN[T]$ of the sender is sufficient due
 186 diligence to ensure that the sender can afford x . Therefore, the minimum due diligence of the receive
 187 $R[T]$ is to check the following: (A) $LIN[T]$ is complete and (B) the lineage of each transaction in
 188 $CLIN[T]$ is complete. We now formalize this recursive set of transactions.

189 **Definition 5.** The Pedigree $PED[T]$ of a transaction or potential transaction T is the recursive
 190 closure of the set $\{T\}$ under the CLIN operator. To compute this:

- 191 • Start with the set of $PED[T] = \{T\}$.
 192 • For each T_i in $PED[T]$, add any element of $CLIN[T_i]$ which is not already in $PED[T]$ to $PED[T]$.
 193 Repeat until there is nothing to add.
 194 • $PED[T] = PED[T]$. \square

195 It is also helpful to think of $PED[T]$ as the nodes of a directed acyclic graph for each transaction,
 196 recursively showing the inbound transactions that the sender depended on to cover the transaction
 197 since the last known global ledger consensus. Figure 1 depicts a \$10 transaction and its PED pedigree.
 198 Here, the opening balances are shown on the nodes, and we assume a minimum balance of \$1. The
 199 \$10 transaction depends on \$2 that the sender already had (over and above the \$1 minimum) plus
 200 two received amounts of \$4 each, of which one, in turn, depended on a received \$3. The dashed lines
 201 represent other received amounts that are not critical to covering the transaction balances so are
 202 excluded from the PED .



203 **Figure 1.** An example of the $PED[T]$ pedigree of a \$10 transaction
 204
 205

206 **Definition 6.** A Disclosure $DIS[\underline{T}]$ for a potential transaction \underline{T} is a communication from sender
207 $S[\underline{T}]$ to receiver $R[\underline{T}]$ of $PED[\underline{T}]$ and $LIN[T_1]$ for every T_1 in $PED[\underline{T}]$. \square

208 **Definition 7.** A Fraudulent transaction T is any transaction wherein the sender $S[T]$ provides
209 the receiver with an incomplete $LIN[T]$ in the disclosure. Additionally, a transaction is considered a
210 fraudulent transaction in retrospect if, at some time between $t[T]$ and the next global ledger
211 consensus, the same sender $S[T]$ is the sender of another transaction T_1 and fails to disclose T when
212 disclosing $LIN[T_1]$. \square

213 Thus, if Malory sends money to Alice and later double spends by sending the same money to
214 Bob without disclosing to Bob the earlier payment made to Alice, then both payments are considered
215 fraudulent. It is insufficient to cancel the second transaction; the one which was directly involved in
216 the fraud like Alice may be a co-conspirator of Malory, while Bob is the only victim. The cancellation
217 of both transactions ensures there is a significant penalty for fraud. In theory, this does mean that Bob
218 would lose out since he was a victim of double-spending in retrospect, but practically this
219 arrangement ensures that double-spending has a negative expected value and is very unlikely to
220 occur at all.

221 **Definition 8.** An Invalid transaction is a transaction that is not fraudulent but wherein the
222 sender in retrospect did not have balance to cover the transaction after removing fraudulent
223 transactions. Equivalently, these are transactions that turn out to have a fraudulent transaction in
224 their pedigree. \square

225 **Definition 9.** Due diligence for a potential transaction \underline{T} is the process of receiver $R[\underline{T}]$
226 communicating the offered disclosure $DIS[\underline{T}]$ with a random selection of $k\sqrt{n}$ (strictly $\lceil k\sqrt{n} \rceil$ i.e.
227 $k\sqrt{n}$ rounded up to the nearest integer) nodes called the validating nodes and confirms that none of
228 them has seen an alternative version of $LIN[T_1]$ for any T_1 in $PED[\underline{T}]$. \square

229 4. k -root- n algorithm

230 Suppose we have n nodes, each of which is also a wallet, connected to a network, and the nodes
231 achieve consensus on the global ledger (or at least on the balance of each node) at some time t_0 in the
232 recent past; we shall call this time the global ledger consensus. Each global ledger consensus may
233 become known at time $t_1 > t_0$, that is with some latency after the time t_0 which it relates to (e.g. in
234 bitcoin, we may wait some hours for transactions to be included in a block and then for 6-block
235 confirmation before trusting that consensus was achieved). When we refer to the last global ledger
236 consensus before time t , we mean the last one known before time $t > t_1$.

237 The nodes transfer balances to each other by mutually digitally signing transactions. Based on
238 the algorithm, each receiving honest node will perform the following steps before accepting and
239 signing a potential transaction \underline{T} .

- 240 • Demand that the sender transmits the disclosure $DIS[\underline{T}]$ that includes the recursive list of
241 dependent transactions $PED[\underline{T}]$ and the transaction history $LIN[T_1]$ for each T_1 in $PED[\underline{T}]$.
- 242 • Validate that each transaction T_1 in $PED[\underline{T}]$ was properly formed and signed by known nodes,
243 and each node had the balance to cover T_1 based on the last known global ledger consensus
244 balance of $S[T_1]$ plus the provided $LIN[T_1]$.
- 245 • Due diligence: Randomly choose $\lceil k\sqrt{n} \rceil$ validating nodes on the network; send each (directly or
246 by communicating through a cascading tree of nodes) $DIS[\underline{T}]$ and ask the validating nodes to
247 validate that they have not seen any alternative LIN history for any transaction in $PED[\underline{T}]$.

- 248 • Any honest node receiving this due diligence request will validate that they have never
 249 previously seen a contradictory *LIN* history for any of the transactions T_1 in $PED[\underline{T}]$. Else, they
 250 will confirm their validation by digitally signing $DIS[\underline{T}]$ and transmitting that back to the receiver
 251 $R[\underline{T}]$. Each validating node must then store every *LIN* transaction history they are asked to
 252 validate until the next achieved global ledger consensus, for future validation.
- 253 • If any validating node has in fact seen an alternative *LIN* transaction history for some T_1 in
 254 $PED[\underline{T}]$, it informs the receiver $R[\underline{T}]$ who rejects potential transaction \underline{T} . Proof of fraud, which
 255 comprises two alternative histories $LIN[T_1]$ and $LIN'[T_1]$, should be broadcast to all nodes on the
 256 network by the validating node. The wallet $S[T_1]$ will be blacklisted and any balance forfeited.
- 257 • The receiver $R[\underline{T}]$ should also broadcast the list of the other $k\sqrt{n}$ validating nodes that failed to
 258 raise an alarm. In case $R[T_1]$ or any other nodes had previously disclosed $LIN'[T_1]$ to one of these
 259 validating nodes V , then this node should also be blacklisted with proof of validation fraud, that
 260 is proof that V signed the current validation which included $LIN[T_1]$ while previously validating
 261 a transaction disclosure that included the forked lineage $LIN'[T_1]$. Thus, the validating nodes are
 262 also held accountable.
- 263 • If on the other hand all the validating nodes validate the transaction, the transaction is accepted
 264 and signed by the receiver.

265 All nodes periodically take time to reach a global ledger consensus on the distributed ledger,
 266 e.g. using Nakamoto consensus. All recipients will request that the transactions they received should
 267 be added to the global ledger. In the case that a node was caught in a fraudulent transaction, it will
 268 be disqualified. All fraudulent transactions are iteratively removed from the ledger.

269 After removing fraudulent transactions, invalid transactions must be iteratively identified until
 270 they are excluded from the global ledger. This process must be iterative since invalidating one
 271 transaction may cause the receiver not to cover for subsequent spends, thereby invalidating further
 272 transactions.

273 In the k -root- n algorithm, there is a need for honest nodes to be online almost all the time. It is
 274 recommended to have a protocol wherein an honest node commits to a service level agreement (SLA)
 275 (e.g. [34] [35]) of say $u = 90\%$ uptime, and a node that does not comply may receive warnings and
 276 eventually financial penalties or disqualification by consensus of all the nodes. A node which tries to
 277 consult $k\sqrt{n}$ nodes and receives less than $uk\sqrt{n}$ responses in a specified target latency time should pick
 278 other nodes and retry until it receives the target $uk\sqrt{n}$ validations.

279 5. An example of detection of double-spending using the k -root- n algorithm

280 Suppose during Monday morning the network reaches a consensus that as of Sunday midnight
 281 the balances on the distributed ledger after all valid transactions were as follows:


- 282 • Chuck (malicious) \$100
 283 • Mallory (malicious) \$100
 284 • Alice (honest) \$100
 285 • Bob (honest) \$100

286 The ledger also showed that there were a total of n valid nodes, in which each has at least a
 287 minimum stake of $m = \$1$.

288 We analyze the scenario where Chuck conspires with Mallory to double-spend, by giving the
 289 same money to Alice and also to Bob. To conceal the double-spend, the payment to Bob is passed by
 290 Chuck via co-conspirator Mallory.

8 of 17

291



292

293

294

Figure 2. An example of transactions between two dishonest and two honest nodes, including attempted double-spending

295

296

297

298

299

300

301

302

- Mallory sends \$99 to Bob (in exchange for some goods, services or other currency). Mallory discloses her transaction history from the last consensus, which is empty, so she has \$99 to spend. Bob first confirms that Mallory had \$100 as of the last known global ledger consensus. Then, Bob being honest does his due diligence and queries $k\sqrt{n}$ network nodes (either directly or through a cascading tree of nodes) to confirm that none of them has seen Mallory signing any other transactions since consensus. They have not, and Bob, therefore, accepts the \$99 and counter-signs the transaction and submits it for eventual inclusion on the main distributed ledger.

303

304

305

306

- Chuck sends \$99 to Mallory. Mallory being malicious and complicit with Chuck tells no one about this transaction. They both sign the transaction and may or may not submit it to the main ledger. Chuck sends this \$99 through Mallory attempting to mask the double-spending he is planning. He might potentially pass this money through further nodes.

307

308

309

310

311

312

313

- Chuck now sends \$99 to Alice in exchange for some value. This is a fraudulent double-spend. He informs Alice fraudulently that he has no other transactions since the last consensus. Alice being honest does her due diligence and queries $k\sqrt{n}$ random validating nodes with the declared disclosure. They all inform Alice that they are unaware of any forked transaction lineages (since transaction #2 was not broadcast) for Chuck, and so Alice accepts the payment. Thus, the double-spend is not yet detected (until both instances of double-spent money reach honest nodes).

314

315

316

317

318

- Mallory sends another \$99 to Bob in exchange for some value. Bob being honest demands disclosure, and Mallory provides Bob with a copy of her transaction history since consensus, namely transaction #1 (-\$99 that Bob already knows about) and transaction #2 (+\$99), thereby evidencing Mallory's balance of \$100 allowing Mallory to spend \$99. At this juncture, Chuck's double-spent money has, via Mallory, reached the honest Bob.

319

320

321

322

323

324

325

326

- Bob notices that Mallory has \$100 but contingent on the money from Chuck (so transaction #2 is in the critical lineage $CLIN[\#4]$ and therefore in $PED[\#4]$). Thus, Bob validates transaction #2 by requiring Mallory to provide Chuck's transaction history $LIN[\#2]$ as part of the pedigree. (Further, in case Chuck, in turn, was depending on the incoming transactions for his balance in transaction #2, which is not the case here, Bob would recursively require the sender's sender's sender's transaction history until he has a transaction history for every transaction since the last consensus which is sufficiently required to justify Mallory's balance to cover the current transaction).

9 of 17

- 327 ▪ Chuck now does his due diligence and queries $k\sqrt{n}$ random network nodes by asking
 328 them to validate the pedigree, which includes both $LIN[\#4]$ and $LIN[\#2]$.
- 329 ▪ Some of these nodes (k^2 on average, but at least 1 with an extremely high probability)
 330 had previously been told about Chuck's alternative transaction history of transaction
 331 #3 where he gave \$99 to Alice. Then, they raise the alarm of double-spending and
 332 broadcast a fraud-proof, namely that transaction #3 was not disclosed in $LIN[\#4]$. The
 333 fraud-proof comprises two divergent transaction histories that were both signed by
 334 Chuck.
- 335 ▪ Bob rejects the fraudulent transaction.
- 336 ○ Chuck has his wallets blacklisted and forfeits his \$1 minimum stake.
- 337 ○ The fraudulent transaction #2 from Chuck to Mallory (which was later
 338 hidden from Alice) is also rejected from the distributed ledger. Therefore,
 339 transaction #4 is invalid since it depends on a fraudulent transaction #2.
- 340 ○ The network preferably should ask Mallory to show that she queried $k\sqrt{n}$
 341 validating nodes. When she fails to do so, Mallory should also be blacklisted
 342 and forfeit her balance. (This is optional extra protection called no due
 343 diligence fraud although this is not required for the algorithm to work and
 344 cannot be enforced in case Mallory had $k\sqrt{n}$ co-conspirators and pretends to
 345 select them at random).
- 346 ○ Alice and Bob compare notes and find all the $\sim k^2$ common validating nodes
 347 they had consulted in #3 and #4 and ensure none of them failed to report the
 348 double-spending. If any common node failed to raise the alarm, then such
 349 node would also be blacklisted for fraud, with the validation fraud-proof
 350 showing that the node received two alternative histories from Chuck and in
 351 both cases approved them (such approvals being signed by the verifying
 352 node comprises verification fraud-proof).

353 The next morning consensus is established again around the following end balances:

- 354 • Chuck (malicious) \$100 (blacklisted with balance forfeited for double-spending)
 355 • Mallory (malicious) \$1 (may be blacklisted for failing to do due diligence on #2)
 356 • Alice (honest) \$100
 357 • Bob (honest) \$199

358 Once this new global ledger consensus is known, future senders only need to provide shorter
 359 transaction histories back to the newer global ledger consensus. In this scenario, Alice has clearly lost
 360 out as a victim of fraud, but the algorithm ensures that the fraudsters have significant losses with
 361 very high probability, meaning that such fraud is very unlikely in the first place.

362 6. Algorithm correctness

363 **Theorem 1.** *For any two honest nodes receiving and successfully validating payments with $k\sqrt{n}$ random
 364 nodes each, there are an average of k^2 common nodes queried by both honest nodes (any one of which can detect
 365 double-spending and raise the alarm).*

366 **Proof of Theorem 1.** The first honest node randomly queries $k\sqrt{n}$ nodes representing a
 367 proportion k/\sqrt{n} of all n nodes. Therefore, when the second honest node queries $k\sqrt{n}$ random nodes, a
 368 proportion of $k\sqrt{n} * (k/\sqrt{n}) = k^2$ on average will overlap. □

10 of 17

369 This result is the key strength of the algorithm since both transactions involve $O(\sqrt{n})$ validating
 370 nodes, and the expected value of the number of overlapping nodes is significant, thereby allowing
 371 any double-spending to be detected.

372 However, since it only requires one common node to detect fraud, what we are not really
 373 interested in the expected number of common validating nodes but rather in the probability of at
 374 least one node in common, versus that of zero common nodes, which we require to be very small.

375 **Lemma 2.** *The probability $p_0(\mathbf{n}, \mathbf{r})$ of zero clashes (zero common nodes) between two random sets of*
 376 *$r = k\sqrt{n}$ nodes satisfies $p_0(\mathbf{n}, \mathbf{r}) < e^{-k^2}$ with $p_0(\mathbf{n}, \mathbf{r}) \approx e^{-k^2}$ for large n and $r \ll n$.*

377 **Proof of Lemma 2.** First, $p_0(\mathbf{n}, \mathbf{r}) = \frac{\binom{n-r}{r}}{\binom{n}{r}}$ since there are $\binom{n}{r}$ ways for the second node to choose
 378 r validating nodes from n nodes in which $\binom{n-r}{r}$ combinations involve zero of the r validating
 379 nodes that the first node chose. Thus,

$$380 \quad p_0(\mathbf{n}, \mathbf{r}) = \frac{\binom{n-r}{r}}{\binom{n}{r}} = \frac{(n-r)!r!(n-r)!}{r!(n-2r)!n!} = \frac{(n-r)\dots(n-2r+1)}{n\dots(n-r+1)} < \left(\frac{n-r}{n}\right)^r = \left(1 - \frac{r}{n}\right)^r \quad \text{with } \approx \text{ for } r \ll n.$$

381 Now, let $r = k\sqrt{n}$. Then, we can approximate

$$382 \quad p_0(\mathbf{n}, \mathbf{r}) = \left(1 - \frac{k\sqrt{n}}{n}\right)^{k\sqrt{n}} = \left(\left(1 - \frac{k}{\sqrt{n}}\right)^{\sqrt{n}}\right)^k < (e^{-k})^k = e^{-k^2}$$

383 again with \approx for the limit of large n \square

384 Therefore, e^{-k^2} is a safe upper bound for p_0 and a good approximation in the realistic case of
 385 large n and small k . By substitution, we can see that $\underline{k} = 4.5$ gives $p_0 \sim 10^{-9}$ for all large n . For
 386 convenience, we typically recommend that $\underline{k} = 4.5$. k must be large enough to allow for the level of
 387 Byzantine faults in the network. A typical practical value would be $k = 10$ to allow for 10% unavailable
 388 nodes and 50% fraudulent nodes, so we have $\underline{k} = 4.5$ of honest available nodes. 10% unavailability
 389 seems generous for most modern networks, while 50% of fraudulent nodes is typically the most
 390 supported by the distributed ledger technology used for global ledger consensus.

391 Appendix A shows values of $p_0(\mathbf{n}, \mathbf{r})$ and confirms that the approximation is excellent for large
 392 n and small k while providing a valid upper bound in all cases.

393 Now, double-spending with co-conspirators is in itself of no value as the co-conspirators will
 394 not provide any value in return for a payment that they know is fraudulent and may be later rejected
 395 from the global ledger. Therefore, the algorithm depends on ensuring a negative expected value
 396 when double-spent money directly or indirectly arrives at honest nodes, catching the fraud in time
 397 before honest nodes naively provide value in exchange for fraudulent payments.

398 **Theorem 3.** *Let M be the maximum transaction amount, m be the minimum wallet balance, n be the*
 399 *number of valid nodes as of the last global ledger consensus, h be the proportion of nodes assumed to be honest*
 400 *and u be the proportion of uptime required from nodes. Assume that the network is designed such that*

$$401 \quad p_0(\mathbf{n}, \underline{k}\sqrt{\mathbf{n}}) < \frac{m}{M+m} \approx \frac{m}{M},$$

402 *where $\underline{k} = khu$. Then, the expected return on any combination of double-spending to honest nodes is*
 403 *negative.*

404 **Proof of Theorem 3.** The maximum amount of a double-spend transaction is the maximum
 405 transaction amount M . By utilizing Lemma 2, the probability of two honest nodes not detecting a
 406 double-spend transaction is $p_0(\mathbf{n}, \underline{k}\sqrt{\mathbf{n}})$, in which case there is a gain of M . In the case that a double-
 407 spend transaction is detected, the double-spender will at least forfeit the minimum wallet balance m .

11 of 17

408 Therefore, the maximum expected gain is given by

$$409 \quad p_0(n, \underline{k}\sqrt{n})M - (1 - p_0(n, \underline{k}\sqrt{n}))m.$$

410 Given $p_0(n, \underline{k}\sqrt{n}) < \frac{m}{M+m} \approx \frac{m}{M}$, this expected value is negative. \square

411 As discussed, practical values are $m = \$1$, $M = \$1,000,000$ and $\underline{k} = 4.5$ which give $p_0 \sim 10^{-9} \ll \frac{m}{M}$.
 412 Assuming $h = 50\%$ honest nodes (the algorithm can handle even fewer than 50% honest nodes but the
 413 blockchain cannot) and $u = 90\%$ uptime, we need $k = 10$ to ensure a negative expected value of any
 414 double-spend.

415 7. Algorithm message space complexity

416 The number of messages per transactions is $k\sqrt{n}$. These may be transmitted directly or cascaded
 417 through a tree of nodes to avoid the receiving node becoming a network bottleneck.

418 Before discussing message size, we present some definitions.

419 **Definition 8.** The critical inbound transaction size $j[T]$ is the number of inbound transactions
 420 (since the last global ledger consensus) which a sender depends on for their balance when spending
 421 money in a transaction t , that is $j[T] = |CLIN[T]|$. \square

422 An upper bound for $j[T]$ is the total number of inbound transactions $|LIN[T]|$ that the node has
 423 participated in since the last global ledger consensus. In most transactions, j will likely be zero. A
 424 person spending money most often already had the money that morning. Although in extreme cases,
 425 v may be large. For example, a grocery store starting the day with zero balance, accepting hundreds
 426 of small transactions, and spending all their accumulated money on a large capital item or payroll
 427 that same evening. The large spend may depend on every one of the small inbound transactions.

428 **Definition 9.** The critical velocity of money $v[T]$ of a transaction T is the maximum depth of the
 429 recursion in $PED[T]$. \square

430 $v[T]$ denotes the number of nodes that the specific balance of coin circulated through in the
 431 period between global ledger consensus until it landed in transaction T , where it is only considered
 432 circulation if the n th transaction depended on its balance for the $(n-1)$ th. v is generally assumed to be
 433 small, most often 1 and rarely more than 2. It is defined more narrowly than the economic concept of
 434 the velocity of money [36] that includes all circulation of currency whether critical to the spender's
 435 balance or not. The velocity of fiat money tends to average a very modest rate of 4–11 per year [37],
 436 so $v > 2$ in a single day between global ledger consensus would be rather rare. That is, in say 24
 437 hours of real commerce, a specific amount of currency will rarely change hands more than once or
 438 twice and at most a few times.

439 The number of transactions in the pedigree of a transaction T is the size of all transactions in the
 440 pedigree, and we can observe that $|PED[T]| = o(\bar{j}^{\bar{v}})$, where \bar{j} and \bar{v} denote the maximum values
 441 of j and v for transactions in the pedigree recursion, respectively.

442 In the majority of the transactions, we expect $v = 1$ and occasionally $v = 2$ but rarely more, while
 443 j is most often 0 but may occasionally hit a few hundred. Thus, the message sizes in realistic commerce
 444 may typically be small; on rare occasions, we may have tens thousands of transactions and require
 445 some megabytes of message size, which is still quite a practical message size for a modern network.

446 However, if the algorithm is continuously used for days and weeks without a global ledger
 447 consensus, the message size v may grow prohibitively large.

448 8. Sybil attack

449 In a Sybil attack, a fraudster develops numerous fraudulent nodes hoping to reduce the chance
 450 of two honest nodes detecting the fraudster's double-spending. We already saw that a 50% attack
 451 does not provide a positive expected value of double-spending with the recommended network
 452 parameters. What about a still larger attack?

453 It is noteworthy that the global ledger consensus algorithm will typically fail with a 51% attack
 454 [38]; regardless we investigate whether such an attack could pay off in the k -root- n algorithm.

455 Suppose again that $m = \$1$, $M = \$10^6$ and $k = 10$. Assume that there are initially n honest nodes,
 456 and the fraudster creates another n fraudulent nodes to control 50%, and suppose further that 10% of
 457 the nodes are unavailable. We already saw that $\underline{k} = 4.5$ and the fraudster has successfully reduced
 458 $p_0 \approx 10^{-44}$ to $p_0 \approx 10^{-9}$. But with $M/m=10^6$ there is no incentive to double-spend with $p_0 \approx 10^{-9}$.

459 In fact, the appendix shows that the fraudster needs to approximately get $\underline{k} < 3.5$ to obtain $p_0 > 10^{-6}$
 460 and achieve a positive expected value for double-spending. Therefore, the criminal would need
 461 approximately $2n$ fraudulent nodes. However, the fraudster then faces another challenge. The loss
 462 from a single unsuccessful double-spending is not limited to forfeiting the double-spending wallet,
 463 but also the loss of all the nodes that failed to detect the double-spend and thereby committed a
 464 verification fraud. According to Theorem 1 the expected number of common nodes is $(10 - 3.5)^2 =$
 465 42.25 nodes for an average loss of at least $42.25 m$. Thus, even in this case, double-spending will have
 466 a negative expected value.

467 Consider more generally that a fraudster creates $(f - 1)n$ fraudulent nodes for a total of $n = fn$
 468 nodes. When a user consults $k\sqrt{n} = k\sqrt{fn}$ nodes, a proportion of $1/f$ of the nodes or $k\sqrt{(n/f)}$ nodes, will
 469 be genuine, this being a proportion k/\sqrt{fn} of all the n honest nodes. Two honest nodes will therefore
 470 have an expectation of consulting $(k/\sqrt{fn})(kn/\sqrt{f}) = k^2/f$ common honest nodes, i.e. $k = k/\sqrt{f}$. They
 471 will also consult on average $k^2 - k^2/f$ dishonest nodes, and if the double-spending is caught these $k^2(1 -$
 472 $1/f)$ nodes will be disqualified.

473 Therefore, the expected payoff from a single double-spending is $p_0(k, n)M \approx e^{-k^2/f}M$, against
 474 the expected cost of $(k^2(1 - 1/f) + 1)m \approx k^2m$ (the fraudulent nodes which fail to report the double-
 475 spending, plus one for the double-spending wallet) for every unsuccessful transaction against a setup
 476 cost of $(f - 1)nm$.

477 Practically, an extremely large number of fraudulent nodes are required for the double-spending
 478 to payoff. Since $k = 10$, we can find numerically that we need approximately $f > 10.5$ for a positive
 479 payoff. Suppose $f = 11$; the fraudster has to create a massive $10n$ fake nodes to control $\sim 91\%$ of the
 480 network. This is done at a cost of $\$10nm$, assuming $\$10m$ for $n = 1$ million. Now, $k = k/\sqrt{f} \approx 3$ and $p_0 =$
 481 $e^{-k^2/f} = e^{-9} \approx 0.00012$. Thus, a double-spending of $M = \$1,000,000$ would have an expected value of
 482 $\$120$, while the expected cost would be losing $k^2(1 - 1/f) + 1 = 91$ nodes at a cost of $91m = \$91$, giving
 483 an expected profit of $\$29$. After such an extreme attack, a single fraudulent transaction would have a
 484 positive expected value.

485 However, even this strategy is doomed to fail in realistic scenarios. If $n = 10^6$, it costs $\$10$ million
 486 to setup 10 million nodes. The user would have to repeat the double-spending three hundred
 487 thousand times to recoup the initial investment. However, they would lose 91 nodes on average
 488 each time they fail, meaning that they would lose the vast majority of the fraudulent nodes before
 489 recovering their investment, so the whole scheme is not feasible.

490 Now, if we increase f further say $f \approx k^2 = 100$, then the fraud can pay off. p_0 gets closer to 1 and
491 the fraudster earns a payback that tends to M as f increases. However, this requires creating $O(100n)$
492 nodes to dominate ~99% of the network. Various strategies that can help to defend against such an
493 extreme attack include the following.

- 494 Reducing M/m : Reducing M can force honest people to have more wallets that increase n .
- 495 • Increasing k .
 - 496 • Biasing the $k\sqrt{n}$ random nodes towards the nodes that have been around for longer or have higher
497 balances.
 - 498 • Monitoring for suspicious behavior such as the creation a huge number of wallets with close to a
499 minimum stake.

500 In summary, with the recommended parameters of k , m and M , we observed that the algorithm
501 is immune to 51% attacks and even 90% attacks, and in fact resilient to all but the most extreme of
502 Sybil attacks.

503 9. Variations on the algorithm for further research

504 9.1. k -root- n without global ledger consensus

505 There would be an option of running k -root- n as the sole algorithm without any common
506 consensus on a ledger. As time goes on, j slowly increases, and the verifications may exponentially
507 become heavier. Each node can should cache everything it knows about other nodes' verified
508 transaction histories. Over time, if money circulates throughout the entire network, every node will
509 end up verifying every transaction at some time or another just once with $k\sqrt{n}$ other nodes, creating
510 in the long term a complexity of $kn\sqrt{n}$ per transaction that seems unattractive. However, there is room
511 for optimizations which could make this approach of standalone k -root- n feasible.

512 9.2. Nodes versus wallets

513 The assumption so far is that every wallet is a node, and the provision of node verification
514 services is part of the cost of being a wallet. This may be feasible as we rapidly move to a world where
515 all devices are online almost all the time, but it could also be a limitation.

516 There could be an alternative variation of the algorithm in which not every wallet is a node. This
517 may be helpful since people may want their wallets to be offline or to be stored on a machine with
518 limited processing power, bandwidth or memory. In this scenario, nodes may be paid a fee to provide
519 verification services with a penalty for failing to report a double-spend they were aware of. Moreover,
520 the nodes could be the same machines as the nodes of the underlying blockchain. Further research is
521 required to formally define such a network.

522 9.3. Forced validation

523 An alternative idea may be considered where even dishonest nodes are forced to consult $O(\sqrt{n})$
524 nodes. In this situation, the dishonest nodes are not given the opportunity to select which nodes they
525 consult since they could pick collaborating dishonest nodes. Therefore, we may introduce a
526 pseudorandom formula to dictate the nodes that are consulted, as well as also ensuring balancing the
527 load between all nodes. In this situation, the idea is that if Alice sends money to Bob and Bob sends
528 the same money to Charlie, then Charlie will again ask $k\sqrt{n}$ nodes to validate that Bob did not double-
529 spend. However, Charlie will not need to ask the network to validate the transaction from Alice to
530 Bob. Charlie can instead ask Bob to see the $k\sqrt{n}$ digital signatures for the appropriate nodes that signed

531 off on the transaction with Alice. Thus, Bob can verify that Bob indeed consulted the right set of $k\sqrt{n}$
532 and received all their approval, creating less traffic and processing demands on the network.

533 We therefore propose that when two people do a transfer, they must notify a formulaically
534 determined pseudorandom selection of $k\sqrt{n}$ other nodes and obtain each of their digitally signed
535 approval. Moreover, the pseudorandom selection is based on a predetermined formula which is
536 known to all and takes as input e.g. sender's ID and the time of the transaction. To reduce the sender's
537 ability to pick and choose a specific time, we preferably take time stamps to be at a resolution of a
538 second or minute (rather than a more fine-grained time slot) when the pseudorandom formula
539 happens to limit their ability to select numerous fraudulent nodes. As before, each of those nodes that
540 are honest may check that the sender has not double-spent.

541 Now for the recursive check of sender's sender and so on, the receiver can check that all the
542 recursive transactions have the necessary sign-off from all the nodes as determined by the
543 pseudorandom formula. Thus, the receiver does not have to burden the network by validating the
544 recursive transactions. Such an algorithm may scale better over longer periods.

545 This algorithm suffers from some clear vulnerabilities. In a real network, there is a high chance
546 that some nodes may not be available, so the sender could feasibly calculate which k^2 nodes would
547 detect his double-spending and simply claim that those particular nodes were not available. This
548 would have to be mitigated by common monitoring of node availability or the honest nodes may self-
549 monitor, so anyone can later validate the claim that a certain node was unavailable at a certain time.

550 The sender may also have multiple wallets and multiple available time slots allowing them some
551 choice of the sending node and time slot, giving them some leeway to plan a double-spend without
552 any clashes of verifying nodes by choosing the particular sending wallet and time slots wherein the
553 pseudorandom formula happens to pick many of their own complicit nodes. If we choose k large
554 enough, we can make this infeasible, for example with $k = 10$ and $p_0 = 3.70 \times 10^{-44}$, the user must
555 consider $O(10^{44})$ combinations of wallets and time slots to obtain one with no clashes to an earlier
556 transaction, which is not feasible.

557 Therefore, it should be feasible to design an algorithm wherein each node (honest or not) is
558 forced to consult a particular set of $k\sqrt{n}$ nodes based on a function of the sender and time slot, and
559 wherein there is no feasible way to create a positive expected value of double-spending.

560 10. Conclusion

561 For a distributed ledger, reaching consensus is expensive and may involve a long lag time. In
562 this paper, we have explored a two-tier system in which the primary algorithm ensures that the global
563 ledger consensus is reached for the distributed ledger, but perhaps only periodically and with high
564 latency. In the meantime, a secondary k -root- n algorithm allows parties to transact rapidly and
565 protect against double-spending with a more efficient $O(\sqrt{n})$ probabilistic algorithm which involves
566 validating each transaction and recursively the transactions it depends on (the transaction's pedigree)
567 with a random selection of $k\sqrt{n}$ nodes. Moreover, we showed that it is feasible for such a network to
568 handle all the world's commerce, while always having a negative expected value of double-spending
569 and being resistant to even aggressive Sybil attacks.

570 Further research is required to investigate the practicality of each wallet being a highly available
571 node or develop the idea of separating wallets and nodes. Further research is required to check the
572 feasibility of the alternative idea of formulaically dictating validating nodes.

573 **Acknowledgments** My thanks to Joshua Fox, Benjamin Fox, Aviv Zohar, for their valuable,
574 constructive feedback.

575 11. Appendix A: Table of k and p_0

576 This table shows p_0 , the chances of zero clashes when two honest nodes each consult $k\sqrt{n}$ random
577 nodes, for various values of k and a couple of values of n .

578 **Table 1.** Values of $p_0(n, k\sqrt{n})$

k	$p_0(n=10^4, k\sqrt{n})$	$p_0(n=10^{10}, k\sqrt{n})$	e^{-k^2}
1	0.36	0.37	0.37
1.5	0.1	0.11	0.11
2	0.017	0.018	0.018
2.5	0.0016	0.0019	0.0019
3	9.30E-05	1.20E-04	1.20E-04
3.5	3.10E-06	4.80E-06	4.80E-06
4	5.80E-08	1.10E-07	1.10E-07
4.5	6.10E-10	1.60E-09	1.60E-09
5	3.70E-12	1.40E-11	1.40E-11
5.5	1.20E-14	7.30E-14	7.30E-14
6	2.30E-17	2.30E-16	2.30E-16
6.5	2.30E-20	4.50E-19	4.50E-19
7	1.30E-23	5.20E-22	5.20E-22
7.5	3.70E-27	3.70E-25	3.70E-25
8	5.60E-31	1.60E-28	1.60E-28
8.5	4.60E-35	4.20E-32	4.20E-32
9	1.90E-39	6.60E-36	6.60E-36
9.5	4.10E-44	6.30E-40	6.40E-40
10	4.40E-49	3.70E-44	3.70E-44

579 12. References

- 580 1. S. Nakamoto, "Bitcoin: A Peer-to-Peer Electronic Cash System," 2009.
581 2. S. S. R. Network, "The Limits to Blockchain? Scaling vs. Decentralization.," Cybersecurity,
582 Privacy & Networks eJournal, 2019.
583 3. "Scalability <https://en.bitcoin.it/wiki/Scalability>," Bitcoin wiki, 2015.
584 4. J. Garzik, "Making decentralized economic policy [http://gtf.org/garzik/bitcoin/BIP100-](http://gtf.org/garzik/bitcoin/BIP100-blocksizechangeproposal.pdf)
585 [blocksizechangeproposal.pdf](http://gtf.org/garzik/bitcoin/BIP100-blocksizechangeproposal.pdf)," 2015.
586 5. A. d. Vries, "Bitcoin's Growing Energy Problem," Joule, vol. 2, no. 15, pp. 801-805, 2018.
587 6. R. Canetti and T. Rabin, "Optimal Asynchronous Byzantine Agreement," Technical
588 Report #92-15, Computer Science Department, Hebrew University, 1992.
589 7. M. Castro and B. Liskov, "Practical Byzantine Fault Tolerance," Proceedings of the Third
590 Symposium on Operating Systems Design and Implementation, New Orleans, 1999.
591 8. J. R. Douceur, "The Sybil Attack," Peer-to-Peer Systems. Lecture Notes in Computer
592 Science., vol. 2429, p. 251-60, 2002.

- 593 9. L. S.A., "Blockchain Unconfirmed Transactions.
594 <https://www.blockchain.com/btc/unconfirmed-transactions>".
- 595 10. D. Azzolini, F. Riguzzi and E. Lamma, "Studying Transaction Fees in the Bitcoin
596 Blockchain with Probabilistic Logic Programming," *Information*, 2019.
- 597 11. "Irreversible Transactions https://en.bitcoin.it/wiki/Irreversible_Transactions," Bitcoin
598 Wiki.
- 599 12. L. e. a. Luu, "SCP: A Computationally-Scalable Byzantine," *IACR Cryptology*, 2015.
- 600 13. Y. Gilad, R. Hemo, S. Micali, G. Vlachos and N. Zeldovich, "Algorand: Scaling Byzantine
601 Agreements," *SOSP '17 Proceedings of the 26th Symposium on Operating Systems Principles*, pp.
602 51-68, 2017.
- 603 14. I. Eyal, A. Efe Gencer, E. Gün Sirerr and R. van Renesse, "Bitcoin-NG: A Scalable
604 Blockchain Protocol," e *Proceedings of the 13th USENIX Symposium on Networked Systems Design
605 and Implementation (NSDI '16)*, 2016.
- 606 15. L. Apeltin, "A CryptoCubic Protocol for Hacker-Proof Off-Chain Bitcoin Transactions,"
607 Cornell University, vol. arXiv:1408.2824, 2014.
- 608 16. J. P. a. T. Dryja, "The Bitcoin Lightning Network: Scalable Off-Chain Instant Payments,"
609 2019.
- 610 17. U. W. Chohan, "The Double Spending Problem and Cryptocurrencies," *Banking &
611 Insurance Journal*, Social Science Research Network (SSRN), 2017.
- 612 18. E. Heilman, A. Kendler, A. Zohar and S. Goldberg, "Eclipse attacks on Bitcoin's peer-to-
613 peer network," *SEC'15 Proceedings of the 24th USENIX Conference on Security Symposium*, pp.
614 129-144, 2015.
- 615 19. M. Apostolaki, A. Zohar and L. Vanbever, "Hijacking Bitcoin: Routing Attacks on
616 Cryptocurrencies," *IEEE Symposium on Security and Privacy*, 2017.
- 617 20. C. Pinzón and C. Rocha, "Double-spend Attack Models with Time Advantage for
618 Bitcoin," *Electronic Notes in Theoretical Computer Science*, vol. 329, 2016.
- 619 21. E. Ittay, G. Peter, J. Aljosha, M. Sarah, S. Nicholas, T. Itay, W. Edgar and Z. Alexei, "Pay-
620 To-Win: Incentive Attacks on Proof-of-Work Cryptocurrencies," *IACR Cryptology ePrint Archive*,
621 2019.
- 622 22. I. Stewart, D. Ilie, A. Zamyatin, S. Werner, M. F. Torshizi and W. J. Knottenbelt,
623 "Committing to quantum resistance: a slow defence for Bitcoin against a fast quantum computing
624 attack," *Royal Society Open Science*, 2018.
- 625 23. M. Saad, J. Spaulding, L. Njilla, C. Kamhoua, S. Shetty, D. Nyang and A. Mohaisen,
626 "Exploring the Attack Surface of Blockchain: A Systematic Overview," arXiv, vol. 1904.03487, 2019.
- 627 24. G. O. Karame, E. Androulaki, M. Roeschlin, A. Gervais and S. Čapkun, "Misbehavior in
628 Bitcoin: A Study of Double-Spending and Accountability," *ACM Transactions on Information and
629 System Security (TISSEC)*, vol. 18, no. 1, 2015.
- 630 25. C. Pérez-Solà, S. Delgado-Segura, G. Navarro-Arribas and J. Herrera-Joancomartí,
631 "Double-spending prevention for Bitcoin zero-confirmation transactions," *International Journal of
632 Information Security*, vol. 18, no. 4, p. 451-463, 2019.
- 633 26. K.-Y. Kang, "Cryptocurrency and Double Spending History: Transactions with Zero
634 Confirmation," *Munich Personal RePEc Archive (MPRA)*, vol. 96875, 2019.
- 635 27. J. Göbel and A. E. Krzesinski, "Increased block size and Bitcoin blockchain dynamics,"
636 27th International Telecommunication Networks and Applications Conference (ITNAC), 2017.

17 of 17

- 637 28. N. R. Ericsson, D. F. Hendry and S. B. Hood, "Milton Friedman and Data Adjustment,"
638 IFDP Notes, Board of Governors of the Federal Reserve System.
- 639 29. W. W. R. Ball, "Probability," in *Mathematical Recreations and Essays*, Macmillan, 1960, *p.*
640 45.
- 641 30. "The World Population Prospects 2019: Highlights," Population Division of the UN
642 Department of Economic and Social Affairs, 2019.
- 643 31. A. Lielacher, "How Many People Use Bitcoin in 2019?" *Bitcoin Market Journal*, 2019.
- 644 32. K. Sedgwick, "No, Visa Doesn't Handle 24,000 TPS and Neither Does Your Pet
645 Blockchain," [https://news.bitcoin.com/no-visa-doesnt-handle-24000-tps-and-neither-does-your-pet-](https://news.bitcoin.com/no-visa-doesnt-handle-24000-tps-and-neither-does-your-pet-blockchain/)
646 [blockchain/](https://news.bitcoin.com/no-visa-doesnt-handle-24000-tps-and-neither-does-your-pet-blockchain/).
- 647 33. "Global General Purpose Cards - Midyear 2018," *The Nilson Report*, vol. 1140, no.
648 https://nilsonreport.com/upload/issues/1140_0321.pdf, *p.* 7, October 2018.
- 649 34. J. Skene, D. D. Lamanna and W. Emmerich, "Precise Service Level Agreements," *ICSE '04*
650 *Proceedings of the 26th International Conference on Software Engineering*, pp. 179-188, 2004.
- 651 35. D. Lamanna, J. Skene and W. Emmerich, "SLang: A Language for Defining," *The Ninth*
652 *IEEE Workshop on Future Trends of Distributed Computing Systems (FTDCS)*, 2003.
- 653 36. I. Fisher, *The Purchasing Power of Money*, 1911.
- 654 37. Federal Reserve Bank of St. Louis, "Velocity of M1 Money Stock," *Fred Economic Data*, no.
655 <https://fred.stlouisfed.org/series/M1V>, 2019.
- 656 38. M. Bastiaan, "Preventing the 51 %-Attack: a Stochastic Analysis of Two Phase Proof of
657 Work in Bitcoin," no. <https://www.hmbastiaan.nl/martijn/docs/2015/preventing-the-majority.pdf>,
658 2015.