

# The Attack of the Clones Against Proof-of-Authority

Parinya Ekparinya

University of Sydney, Australia  
Email: pekp6601@uni.sydney.edu.au

Vincent Gramoli\*

University of Sydney, Australia  
CSIRO-Data61, Sydney, Australia  
Email: vincent.gramoli@sydney.edu.au

Guillaume Jourjon

CSIRO-Data61, Sydney, Australia  
Email: guillaume.jourjon@data61.csiro.au

**Abstract**—The vulnerability of traditional blockchains have been demonstrated at multiple occasions. Various companies are now moving towards Proof-of-Authority (PoA) blockchains with more conventional Byzantine fault tolerance, where a known set of  $n$  permissioned sealers, among which no more than  $t$  are Byzantine, seal blocks that include user transactions. Despite their wide adoption, these protocols were not proved correct.

In this paper, we present the Cloning Attack against the two mostly deployed PoA implementations of Ethereum, namely Aura and Clique. The Cloning Attack consists of one sealer cloning its pair of public-private keys into two distinct Ethereum instances that communicate with distinct groups of sealers. To identify their vulnerabilities, we first specify the corresponding algorithms. We then deploy one testnet for each protocol and demonstrate the success of the attack with only one Byzantine sealer. Finally, we propose counter-measures that prevent an adversary from double spending and introduce the necessary number of sealers needed to decide a block depending on  $n$  and  $t$  for both Aura and Clique to be safe.

## I. INTRODUCTION

Ethereum is one of the most popular blockchain systems thanks to the large ecosystem of distributed applications that it executes. Unfortunately, the default Ethereum protocol based on *proof-of-work* (PoW) can fork as it allows distinct blocks to be appended at the same index of the chain. This forking situation can lead to security vulnerabilities, like double spending, if it is not detected early enough [20], [31], [30]. Alternative protocols, called *proof-of-authority* (PoA) protocols, that aim at avoiding forks have recently been integrated in the most widely deployed versions of Ethereum, parity and geth, and are currently used world-wide. PoA has become rapidly popular and is now distributed through major Software-as-a-Service providers and used in several blockchain networks [7], [39], [3]. Yet, to our knowledge, the level of security offered by PoA protocols has not been properly assessed.

These PoA consensus protocols, called Aura and Clique, are said to use a proof-of-authority because they restrict the creation of a block to a fixed set of  $n$  authority nodes, called *sealers*, among which a maximum of  $t < \frac{n}{2}$  can misbehave or be *Byzantine*. They aim at solving the well-known Byzantine

consensus problem [33], where among  $n$  nodes the  $n-t$  honest ones agree on a unique block. More precisely, honest nodes cannot decide different values (agreement) and eventually they must decide (termination). PoA gives the sealers the authority to *seal* a block, which consists of signing cryptographically the block. This set of sealers can possibly change over time if a subset of the participants allow it, hence being well suited for dynamic consortia of participants as in community blockchains.

For these reasons, PoA recently gained rapid momentum in critical applications [39], [3]. Industry, such as Lavaa, propose a tracking service to prevent fraud counterfeiting on top of Aura [39]. Microsoft describes how to deploy Aura “in production” [16]. Amazon Web Services offers PoA through the Clique protocol built in geth to its customers [25]. They implemented a service that aims at maintaining data privacy and integrity in a multi-tenant scenario. Every day, Internet users exchange digital assets through multiple instances of these two protocols. Huawei uses the Apla blockchain platforms based on PoA to develop smart transportation by coupling IoT with blockchain in supply chains and logistics [5]. Rinkeby is a network of 65 participants offering the Clique service across four continents to its users [34]. The xDai DPOS network uses an Ethereum 1.0 sidechain based on the Aura consensus protocol to transfer assets [3]. Sokol and Kovan are other Ethereum testnets running the Aura protocol [21].

An interesting theoretical work by De Angelis et al. [2] indicated that the consistency of Aura can be limited, for example if the clocks are far apart, and that Clique is eventually consistent. To the best of our knowledge, all these implemented algorithms have simply been described in words often in online documentations or white papers and no pseudocode description has been provided. In particular, we noted that some of the stated assumptions are ambiguous and require experts to maintain them during deployment. As a result it is unclear whether an attacker could violate their data integrity.

In this paper, we show that PoA is not secure even under its disambiguated conditions: when a sufficiently large subset  $V$  of sealers among  $n$  must seal a block despite a minority  $t$  of them being malicious. To this end, we design, implement and experiment an attack, called *the Cloning Attack*, against both Ethereum’s Aura and Clique consensus protocols that allows us to steal digital assets with only  $t = (2 - (n \bmod 2))$  failures. The Cloning Attack consists of one or two sealer attacker(s) cloning a private key to convince half of the honest sealers that a transaction is correctly committed before erasing this transaction to double spend its coins. Thanks to the cloning, to convince half of the honest sealers that transactions

---

\*Part of this work was produced while the author was a visiting scientist at École Normale Supérieure, Paris, France.

are committed, the attacker simply needs to delay messages between two halves of honest sealers. Note that this is achieved without knowing the precise locations of sealers.

Our counter-measure inspired by the theory of Byzantine fault tolerance defines precisely the necessary and sufficient number of sealers  $|V|$  needed to sign each block,  $\frac{n+t}{2} < |V| < n-t$ , for PoA protocols to be safe (i.e., satisfy agreement) and live (i.e., satisfy termination). More specifically, the upper bound  $|V| < n-t$  is necessary for PoA protocols to be live whereas the lower bound  $|V| > \frac{n+t}{2}$  is necessary for PoA protocols to be safe. As expected these bounds comply with the well-known impossibility to solve consensus when  $t \geq n/3$  [33]. In fact, if  $t \geq n/3$ , then the aforementioned bounds cannot be met as  $n-t \leq \frac{n+t}{2}$ : to be safe PoA protocols would violate liveness but to be live they would necessarily violate safety. Our counter-measure differs from previously known results in that it relies on a new parameter, the set  $V$  of needed participating sealers that comprises typically less than  $n$  participants.

**Responsible disclosure:** For ethical reasons, we previously communicated the vulnerability we present here to (i) the security team of Parity Technology, (ii) the security team of the Ethereum Foundation and (iii) the Ethereum development community [14]. Both security teams acknowledged the possibility of the attack and the xDai blockchain of the POS-DAO project is currently implementing one of the counter-measures of Section IX-C at <https://github.com/poanetwork/parity-ethereum/pull/109> [3].

We demonstrate the effectiveness of the Cloning attack by double spending in two testnets, one running parity and the other running geth. The application of the Cloning attack to Aura is slower as it consists of the attacker sealing more blocks in one branch while its application to Clique is faster but more subtle as it consists of the attacker disordering the victim sealers to minimize the weight of a branch. Overall, the attack against Clique is about twice faster than Aura’s but its success rate ranges from 60% to 100%. In order to remedy the identified vulnerabilities, we propose to modify these two PoA protocols to preserve their safety. Even though our counter-measures introduce liveness limitations in these algorithms they make them more suitable for critical applications.

Section II presents the related work. Section III describes the model. Section IV lists the pseudocode of the algorithms in the Aura and Clique algorithms as implemented in, respectively, parity and geth. Section V describes the Cloning Attack against both algorithms. Sections VI and VII explain how to exploit it to double spend in Aura and Clique, respectively. We then present in Section VIII our evaluation of the Cloning Attack on both protocols, while Section IX discusses our results and potential countermeasures. Section X concludes the paper.

## II. BACKGROUND

Most of the known *double spending* attacks against blockchains exploit their inherent permissionless mechanism by including in the blockchain a transaction that transfers coins and then discards this transaction, hence allowing to re-spend the previously spent coins in a subsequent transaction. Below

we list some of these attacks to explain the recent raise of alternative protocols based on PoA.

Perhaps the most conventional way to double spend in permissionless blockchains is for an attacker to exploit more than half of the mining power of the system to create a heavier or longer branch that can overwrite transactions that were expected to be sufficiently confirmed or *committed* [35]. In some blockchains, a quarter of the mining power appears enough in theory to attract participants into a coalition whose cumulative mining power reaches strictly more than half of the total mining power [17]. SMARTPOOL [27] copes with the centralisation of mining power into these blockchains and the risk of mining pools to join a coalition of strictly more than half of the total mining power.

To attack permissionless blockchains without a significant mining power, researchers attacked the network. The Eclipse attack against Bitcoin [20] consists of isolating at the IP layer a victim miner from the rest of the network to exploit its resources. The Blockchain Anomaly [29] exploits message reordering in Ethereum to abort transactions that seemed sufficiently confirmed. The Balance Attack [30] partitions the network into groups of similar mining power to influence the selection of the canonical chain. Recently, actual man-in-the-middle attacks were run to demonstrate the feasibility of stealing assets in Ethereum without a significant mining power [13].

To cope with these attacks, some modern blockchains build upon Byzantine agreement [33], sometimes probabilistically [28], [19], sometimes deterministically [22], [9]. Given the long series of research results in the literature about Byzantine agreement protocols, we know that when the network is *synchronous* and messages are delivered in a known bounded time, then  $t < \frac{n}{3}$  is sufficient to reach consensus. If one also assumes authentication, then even  $t < \frac{n}{2}$  becomes sufficient [24]. What is key for critical applications is that these Byzantine fault tolerant blockchains guarantee that no participants double spend even when messages get unexpectedly delayed. Unfortunately, it is impossible to reach consensus when message delays are unbounded [18].

Proof-of-Authority (PoA) was recently proposed as a Byzantine fault tolerant consensus mechanism that integrates with the Ethereum protocols [4]. The Ethereum geth software offers two different PoA consensus protocol, called Clique and Istanbul BFT [26] whereas the Ethereum parity offers the PoA consensus protocol, called Aura. The concept is similar to traditional Byzantine fault tolerant consensus in that only  $n$  sealers are permissioned to create new blocks but requires authentication and strictly less than  $\frac{n}{2}$  Byzantine participants, similarly to the seminal work on Byzantine consensus [24].

The claim that PoA can “tolerate up to 50% of malicious nodes” [4] raised interest from the industry [7], [39], [3]. However, some work recently questions the consistency of PoA [2], [37], [36]. In particular, it was found that unsynchronized clocks could affect Aura’s consistency whereas Clique was only eventually consistent [2], however, no attacks against Aura or Clique have been proposed. Another work [37] mentioned that an attacker could maintain two chains of equal lengths. To be possible, this requires the attacker to falsify block timestamps to violate the policy that new blocks are

appended to the branch whose latest block has the earliest timestamp among all branches. It turns out, that such a violation can be easily detected by other sealers verifying timestamp or header of the blocks. This is probably why the developers did not change the code to remedy this, as the authors mentioned in their paper.

As we explained previously our result matches the impossibility of solving consensus when  $t \geq \frac{n}{3}$  [33], however, it differs from the classic results on Byzantine consensus where no subset of authority nodes are considered. In our case, if the set  $V$  of needed authority nodes or sealers comprise all  $n$  nodes, then the highest number of failures that Aura could tolerate under a synchrony assumption would be  $t = n - 1$ , however, without this synchrony assumption, even with  $t = 1$ , PoA protocols are not guaranteed to be safe. The relation between  $|V|$  and  $t$  is described in Section IX-B.

Our cloning idea shares similarities with the idea to attack a distributed system by falsifying identities that was already discussed in the past. The Sybil Attack [11] presents the attack against peer-to-peer systems by forging multiple identities. By contrast, the Cloning Attack consists of replicating the machines using the same identity rather than forging identities. More specifically, the Cloning Attack leverages the fact that Ethereum accepts two different machines located at different ends of the network to use the same private key.

### III. MODEL

*Distributed system:* We consider a distributed system of  $n$  permissioned sealers whose identifiers are  $p_1, \dots, p_n \in Ids$ . As the blockchain is open, it accepts the requests issued by *nodes* or processes that are not necessarily sealers, hence the overall number of participants can be larger than  $n$ , but only  $n$  participants can propose blocks and *seal* (or *sign*) them.

*Failures:* We consider the *Byzantine* failure model where Byzantine nodes can act in an arbitrary way. Note that this differs from a rational model in which nodes are typically incentivized to act in a certain way so as to maximize their utility function, like a personal reward or reputation. This rational model is however insufficient to characterize the behavior of participants with motivations external to the system: one can think of a blockchain company investing secretly significant resources into a participant running the blockchain of its competitor in order to attack it, regardless of the reward this participant could generate. Following the original tolerance claim of “up to 50% of malicious nodes” [4], we assume, in our analysis that no more than  $t < \frac{n}{2}$  participants can be *Byzantine* but show that  $t = 1$  Byzantine participant is sufficient to double spend. The remaining  $n - t$  participants are *honest*.

*Authentication:* We assume authentication through a public-key cryptosystem that allows participants to easily identify that a block is correctly signed by a sealer so that incorrectly signed blocks are simply ignored. We assume that keys cannot be forged or stolen by Byzantine participants and that appropriate private keys are correctly distributed to the sealers initially. As in the Dolev-Yao model [10], we assume the attacker, who has the control over the Byzantine participants, can intercept messages.

*Partial synchrony:* As it is well known that consensus cannot be reached in an environment where communication is asynchronous in the presence of faults, it appears natural to assume additional synchrony. It is unclear whether PoA protocols can be safe under *partial synchrony*, where messages get delivered in a bounded amount of time that is not known from the algorithm [12] or how long communication can take for these protocols to work. As an example, a preliminary version of Aura was mentioned to require synchrony in a web document [32], however, this information appears outdated as the implementation is closer to another documentation [4] that does not mention this assumption, as we explain in Section IV-A.

The questions we investigate is whether PoA protocols work under partial synchrony, and if not, whether the risk of unexpected message delays is benign (liveness or termination of the consensus is not guaranteed but safety remains guaranteed in that no double spending occur) or can have dramatic consequences (disagreement can occur), hence letting an attacker double spend. As we will explain, our conclusion is that PoA protocols do not work in that even safety is not guaranteed but some countermeasures can remedy this problem.

## IV. POA CONSENSUS ALGORITHMS

In this section we describe the two main variants of PoA algorithms, called Aura and Clique, implemented in the predominant Ethereum software, called parity and geth, respectively. We first discuss two distinct versions of the Aura algorithm that are both publicly available online. We then present the pseudocode for the Clique algorithm.

### A. The Aura consensus algorithms

There exist two distinct versions of the Aura algorithm as documented online, one that corresponds to the current parity implementation of the Ethereum protocol and another [4] that uses rounds to decide whether a consensus decision is reached.

1) *The parity Aura algorithm:* Algorithm 1 depicts the way Aura guarantees that participating nodes reach consensus on the uniqueness of the block at a given index of the blockchain as implemented in Parity-Ethereum-v2.0.8 (v2.0.8 was the latest version at the time we performed our experiments). Every participating node maintains a state comprising a set of *sealers*, its current view of the blockchain  $c_i$  as a directed acyclic graph  $\langle B_i, P_i \rangle$ , a block  $b$  with fields *parent* that links to the parent block, a *sealer* and a *step* indicating the time at which the block is added to the blockchain, as explained below. They are initialized to  $\perp$  meaning “undefined”.

The function `propose()` is invoked in order to propose a block for a particular index of the blockchain. The consensus is reached once the block is decided, which can happen much later as we will explain in the function `is-decided()` (line 29) below. The algorithm discretises time into steps that correspond to consecutive periods of step-duration time, as specified in a configuration file. Each sealer executes an infinite loop that periodically checks whether the `clock-time()` indicates that this is its turn to propose a block (line 13). When it is its turn (line 14), a sealer sets the *parent* of the block to the last block of its view and signs it (line 16).



---

**Algorithm 1** The parity Aura algorithm at process  $p_i$ 

---

```
1: State:
2:  $sealers \subseteq Ids$ , the set of sealers
3:  $c_i = \langle B_i, P_i \rangle$ , the local blockchain at node  $p_i$  is a directed
4: acyclic graph of blocks  $B_i$  and pointers  $P_i$ 
5:  $b$ , a block record with fields:
6:  $parent$ , the block preceding  $b$  in the chain, initially  $\perp$ 
7:  $sealer$ , the sealer that signed block  $b$ , initially  $\perp$ 
8:  $step$ , the blockchain step when the block gets added, initially  $\perp$ 
9:  $step-duration$ , the duration of each step as configured
10:
11: propose() $i$ : ▷ sealers keep proposing
12: while true do ▷ infinite loop
13:    $step \leftarrow \text{clock-time}() / \text{step-duration}$  ▷ discretize time
14:   if  $i \in sealers \wedge step \bmod |sealers| = i$  then ▷ my turn
15:      $b.parent \leftarrow \text{last-block}(c_i)$  ▷ link a block
16:      $b.sealer \leftarrow p_i$  ▷ seal the block
17:      $b.step \leftarrow step$ 
18:      $c_i \leftarrow \langle B_i \cup \{b\}, P_i \cup \{b.parent\} \rangle$  ▷ update local view
19:     broadcast( $c_i$ ) ▷ send blocks
20:     sleep( $step-duration$ ) ▷ wait before looping
21:
22:  $score(\langle B_j, P_j \rangle)$ : ▷ compute the score of a branch
23: return  $\text{UINT128\_MAX} \times \text{height}(\langle B_j, P_j \rangle) - \text{step-num}(\langle B_j, P_j \rangle)$ 
24:
25:  $deliver(\langle B_j, P_j \rangle)$  $i$ :
26: if  $score(\langle B_j, P_j \rangle) > score(\langle B_i, P_i \rangle)$  then
27:    $\langle B_i, P_i \rangle \leftarrow \langle B_j, P_j \rangle$  ▷ select the right branch in case of forks
28:
29: is-decided( $b$ ) $i$ :
30:  $V \leftarrow \{b_k.sealer \mid b_k \in B_i \wedge b_k.step \geq b.step\}$  ▷ sealers after  $b$ 
31: return  $(|V| \times 2 > |sealers|)$  ▷ more than majority of sealers signed
```

---

Each `broadcast()` invoked by the `propose()` function sends blocks that get delivered to all other participating nodes that are honest (in reality only the last block is broadcast unless some sealers are unaware of more blocks). The `deliver()` function (line 25) is thus invoked at each honest participating node, regardless of whether it is a sealer, upon reception of the broadcast message. Once a blockchain view is delivered to  $p_i$ , the node compares the score of the blockchain view it maintains to the blockchain view it receives, using the `score` (line 22). The highest blockchain has the greatest score, however, if two blockchains share the same height, then the one that is denser in terms of its number of non-empty slots obtains the highest score. In other words, among many blockchains with the same height, a blockchain whose last block has the lowest index wins. This is indicated by the two functions `height` and `step-num` that represent the height of the blockchain and the number of slots for which there exists a block in the blockchain.

Finally, function `is-decided()` (line 29) takes a block  $b$  as an argument and returns whether  $b$  is considered *decided*, meaning that all the transactions it contains are committed [29]. To this end, the function simply lists the set  $V$  of distinct sealers that have signed a block following  $b$  in the blockchain. If this set contains a majority of the sealers, i.e.,  $|V| \times 2 > |sealers|$  (line 31), then the block  $b$  is decided and `true` is returned.

2) *Round-based variant of the Aura algorithm*: The Aura algorithm implemented in parity is not the only algorithm called, Aura. Another variant is presented in the PoA Network white paper available online [4]. Algorithm 2 presents the different decision technique of this variant, the rest of the pseudocode being identical to Algorithm 1.

In order to know whether a block  $b$  is decided at the end of a successful consensus (Algo 2, line 1), a participant simply has to check whether there exist two consecutive rounds  $round1$  and  $round2$  following block  $b$ , in each of which the blocks are sealed by a majority of the sealers.

---

**Algorithm 2** The round-based variant of Aura at process  $p_i$ 

---

```
1: is-decided( $b$ ) $i$ :
2:  $\ell \leftarrow |sealers|$  ▷ number of validators
3:  $round1 \leftarrow (b.step, b.step + \ell)$  ▷ steps in next round
4:  $round2 \leftarrow (b.step + \ell, b.step + 2 \times \ell)$  ▷ steps in the 2nd next round
5:  $maj1 \leftarrow |\{b' : b'.step \in round1\}| > \ell/2$  ▷ majority in round 1
6:  $maj2 \leftarrow |\{b'' : b''.step \in round2\}| > \ell/2$  ▷ majority in round 1
7: return  $(maj1 \wedge maj2)$  ▷ decided if majority in both rounds
```

---

Note that while presented in some documentation, this alternative Aura specification is not the one used by the mainstream implementation of the protocol. The current definition of the Aura algorithm disregards the rounds and simply requires enough blocks to be sealed [32]. Although the version of Aura we experiment in this paper is the mainstream one (Algorithm 1), the attack we present in Section VI also applies to this more restrictive definition presented in Algorithm 2.

### B. The Clique consensus algorithm

Algorithm 3 depicts the pseudocode of the Clique consensus algorithm. It is the one used currently in Go Ethereum geth-1.8.20-stable.

Every participating node shares the same initial block, the genesis block, which also contains the block-period, the period between consecutive block creations. Similarly to the Aura protocol, each node maintains its own view of the growing blockchain  $c_i$  as a directed acyclic graph  $\langle B_i, P_i \rangle$ . A block  $b$  contains a *number* as an index of the block in the blockchain, a *weight* as a weight of the block, a *parent* field that links to its parent block and a *sealer*.

The `propose()` function runs an infinite loop in order to propose blocks to the blockchain when certain conditions are satisfied. The first condition (line 26) requires the process to wait for blocks from other sealers until none of the last sealer-limit blocks contains its signature. In the current implementation the sealer-limit must be  $\lfloor |sealers|/2 \rfloor + 1$ , which is the smallest majority. As a result of this first condition, the sealers need to take turn to sign blocks. The second condition (line 28) is to wait for block-period.<sup>1</sup> When both conditions are met, the process checks if it is its turn to sign the block (line 29). The process may sign a block right away with *weight* equal to 2; otherwise, it may sign a block with *weight* equal to 1 after a random delay between 0 and  $500 \times \lfloor |sealers|/2 \rfloor + 1$  milliseconds (line 32). The consensus is reached once the block is decided later as we will describe in the function `is-decided()` (line 47). The last step in the loop, `broadcast()`, sends messages to other participants.

Upon reception of the broadcast message, the `deliver()` function (line 43) is invoked at each participating node regardless of whether it is a sealer. The total-weight function (line 40) used by the process compares the weight between

---

<sup>1</sup>The default block-period is 15 seconds as developers suggest the same duration to remain analogous to the proof-of-work blockchain Ethereum.

---

**Algorithm 3** The geth Clique algorithm at process  $p_i$ 

---

```
1: State:
2:  $sealers \subseteq Ids$ , the set of sealers
3:  $c_i = \langle B_i, P_i \rangle$ , the local blockchain at node  $p_i$  is a directed
4: acyclic graph of blocks  $B_i$  and pointers  $P_i$ 
5:  $b$ , a block record with fields:
6:  $parent$ , the block preceding  $b$  in the chain, initially  $\perp$ 
7:  $sealer$ , the sealer that signed block  $b$ , initially  $\perp$ 
8:  $number$ , the index of the block in the chain, initially  $\perp$ 
9:  $weight$ , the weight of a block, initially  $\perp$ 
10:  $block\text{-}period$ , minimum duration in second between timestamps of
11: two consecutive blocks, initially 5 seconds
12:  $majority$ , the number of  $\lfloor \frac{|sealers|}{2} \rfloor + 1$ 
13:  $sealer\text{-}limit$ , max. number of consecutive blocks among which a sealer
14: can sign at most one block, initially set to the majority
15:
16:  $sign\text{-}recently(c_i, n)_i$ :
17:  $\lambda \leftarrow sealer\text{-}limit$ 
18:  $ret = false$ 
19: for  $m = n - \lambda, \dots, n$  do  $\triangleright$  loop through last  $\lambda$  blocks
20:   if  $b_m.number \bmod |sealers| = i$  then  $ret = true$ 
21: return  $ret$ 
22:
23:  $propose()_i$ :
24: while true do
25:    $n \leftarrow last\text{-}block(c_i).number$   $\triangleright$  last block index
26:   wait until  $\neg sign\text{-}recently(c_i, n)$   $\triangleright$  wait until I can seal a block
27:    $T \leftarrow get\text{-}last\text{-}timestamp(c_i)$ 
28:   wait until  $clock \geq T + block\text{-}period$   $\triangleright$  wait  $\geq$  block-period
29:   if  $(n + 1) \bmod |sealers| = i$  then  $\triangleright$  in-tum sealing
30:      $b.weight = 2$   $\triangleright$  block weight 2
31:   else  $\triangleright$  out-of-tum sealing
32:      $sleep(rand([0, 500 \times majority]) ms)$   $\triangleright$  random delay in millisecs
33:      $b.weight = 1$   $\triangleright$  block weight 1
34:      $b.number = n + 1$   $\triangleright$  increment block index
35:      $b.parent \leftarrow last\text{-}block(c_i)$   $\triangleright$  link a block
36:      $b.sealer \leftarrow sign()$   $\triangleright$  seal the block
37:      $c_i \leftarrow \langle B_i \cup \{b\}, P_i \cup \{b.parent\} \rangle$   $\triangleright$  update local view
38:      $broadcast(c_i)$   $\triangleright$  send blocks
39:
40:  $total\text{-}weight(\langle B_j, P_j \rangle)_i$ :  $\triangleright$  total weight
41: return  $\sum_{b \in B_j} b.weight$ 
42:
43:  $deliver(\langle B_j, P_j \rangle)_i$ :
44: if  $total\text{-}weight(\langle B_j, P_j \rangle) > total\text{-}weight(\langle B_i, P_i \rangle)$  then  $\triangleright$  heaviest
45:    $\langle B_i, P_i \rangle \leftarrow \langle B_j, P_j \rangle$ 
46:
47:  $is\text{-}decided(b)_i$ :
48:  $V \leftarrow \{b_k.sealer : b_k \in B; k \geq i\}$   $\triangleright$  sealers in blocks since  $b_i$ 
49: return  $|V| > majority$   $\triangleright$  more than majority of sealers signed
```

---

two blockchain views, a current blockchain that it maintains locally and the one freshly received. The process updates its local view if the received blockchain is heavier; otherwise it keeps the same local blockchain view.

To consider whether a block  $b$  is *decided* (line 47), a process has to check the set of sealers who sign blocks after  $b$ . Only when a majority of sealers have appended subsequent blocks to the chain, can a block be considered decided.

## V. THE CLONING ATTACK

In this section, we present the Cloning Attack to double spend in PoA blockchains. In particular, we present the commonalities between the attacks against Aura and Clique, namely the cloning process that allows an attacker to play different roles in the blockchain, the majority that allows two

groups of sealers to make progress without the other, and the way transactions should conflict to double spend. The difference in how these attacks are applied to Aura and Clique are deferred to Sections VI and VII, respectively.

By assumption, only a minority of the sealers can be malicious, this is the reason why PoA algorithms require a majority of sealed blocks to consider whether a block is decided and its transactions appear to be committed. Intuitively, this should prevent the malicious sealers from forming a coalition that can double spend. In reality, as we explain below,  $(2 - (n \bmod 2))$  attacker(s) cloning their own instance into two clones are sufficient to double spend.

### A. Cloning instances by duplicating keys

The first step necessary in the Cloning Attack is for some attacker to duplicate its Ethereum instance into two clones. *Cloning* consists for a single user of running two instances of the Ethereum protocol with the same *address* or public-private key pair. Note that these two instances could run either on the same machine, using the same IP address, or on distinct machines with distinct IP addresses. We call these two instances *clones* because one has the same information as the other before messages start being delayed. In addition, during the whole duration of the attack, both clones use the same public-private key pair. Interestingly, we noted that Ethereum allows these two cloned instances to both create blocks, however, as they use the same private key to seal blocks, they are considered to act as a unique sealer.

At some point, the attacker exploits message delays (either accidental or as a result of a network attack) between two groups of a minority of  $\lceil n/2 \rceil - 1$  sealers, hence creating a transient partition. At this moment, the two clones may not share exactly the same database content as they may not be aware of the exact same blocks that are present in the blockchain. To maintain the cloning at the start of the partition, the attacker copies the content of the blockchain database of one of the clones to the database of the other clone and connects each of these clones to a different partition. During the time of this partition, the Ethereum protocol readjusts the peering so that sealers within the same group keep communicating.

Note that there are various conditions that can cause a partition in the Ethereum network and only one of these conditions is sufficient for the attack to be successful. First, one can leverage a delay in the network due to BGP table misconfiguration for example. Second, one can maliciously attack the network through DoS, ARP-spoofing, or BGP-hijacking. In particular, BGP-hijacking works by having an attacker advertising to one group wrong routes that reach the other group in order to intercept all traffic between the two groups. Once the traffic is rerouted, the attacker can simply delay the propagation of messages. We refer the interested reader to existing ways of implementing man-in-the-middle attacks in Ethereum [13].

### B. Majority groups to guarantee progress

Clones are exploited in the attack to give the illusion to honest sealers that each group contains a majority of sealers. In order to progress towards a double spending situation, each

group must commit transactions and thus decide blocks, this is why we need  $(2 - (n \bmod 2))$  attackers that clone instances:

- **Case  $n$  is odd.** The honest sealers can be split into two groups of  $(n-1)/2$  sealers, each representing a minority. In order to guarantee progress of the protocol on both sides of the partition, a single attacker can simply add one clone in each minority, hence reaching a majority of  $\lfloor n/2 \rfloor + 1$  sealers on each side. This is the reason why  $(2 - (n \bmod 2)) = 1$  attacker is sufficient when  $n$  is odd.
- **Case  $n$  is even.** A single attacker could split the  $n-1$  honest sealers into two groups of different sizes, one that contains  $n/2$  sealers and another that contains  $n/2 - 1$  sealers. It would however be insufficient to include a clone in the second group to guarantee its progress. This is why  $(2 - (n \bmod 2)) = 2$  attackers are needed.

To conclude, the  $(2 - (n \bmod 2))$  attacker(s) thus partition(s) a network of  $n$  sealers into roughly two halves to which they add clones so that each group contains a majority of at least  $\lfloor n/2 \rfloor + 1$  sealers. This guarantees the progress of the protocol on each group so as to obtain the commit of a transaction  $TX_1$  on one group and the canonical chain containing  $TX_2$  in the other group. For example, there must be at least 5 sealers in each subgroup for a network of  $n = 9$  sealers. Such a condition is required to ensure termination of the consensus algorithm, so that blocks will be decided, or appear to be final, from the viewpoint of both subgroups.

Note that we consider here the necessary time for a partition. In a realistic scenario, the attacker may want the effect of its transaction to take occur before stopping the partition. For example, an attacker buying a good in transaction  $TX_1$  may want to receive the good before the transaction gets discarded from the blockchain.

### C. Conflicting transactions

The most common way of double spending is to make sure a transaction  $TX_1$  ends up being included in one branch of a fork, then convincing the recipient that  $TX_1$  is committed, before resolving the fork by discarding the branch of this transaction  $TX_1$ . Later on, the sender of the transaction  $TX_1$  can simply reuse the coins he initially spent in  $TX_1$  in another transaction  $TX_2$ . Interestingly in Ethereum, if the conflicting transaction  $TX_2$  is not issued early enough, then  $TX_1$  could be re-included in a mempool and committed later on.

The goal is for the clones to leverage the message delays between network partitions to rapidly issue two conflicting transactions. As soon as the blockchain network is divided into two subgroups, the attacker issues a minimum of two conflicting transactions, at least one transaction to each subgroup. A typical example to illustrate the double spending attack is two conflicting transactions:

- $TX_1$  where Alice gives all her coins to Bob in the first transaction sent to one group and
- $TX_2$  where Alice gives all her coins to Carol on the other transaction sent to the other group.

It is clear that committing both transactions would violate the integrity of Alice’s account and would result in a double spending. Once the first transaction appears committed, delivering

the delayed messages or ending the partition will have the effect of discarding one of the two transactions.

In the next two sections, we explain how the majority of sealers in Aura and the order of the sealings in Clique allow to select the transaction to be discarded by the system.

## VI. THE CLONING ATTACK AGAINST AURA

We now present a simple way to apply the Cloning Attack to double spend in Aura. To discard the branch, say the *victim branch*, that contains  $TX_1$  and double spend, the attacker must influence Aura to select the branch containing  $TX_2$ , say the *attacker branch*, as the canonical chain.

As explained earlier in Algorithm 1, the current implementation of Aura simply chooses the longest chain as the canonical chain whenever a fork is detected. So, to influence the selection of the attacker branch as the canonical chain, the attacker simply has to contribute to the attacker branch by sealing more blocks in the group maintaining this branch than the other group.

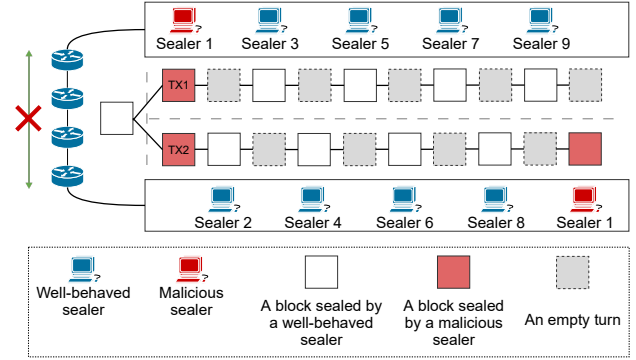


Figure 1. Applying the Cloning Attack to double spend in Aura requires the attacker, “Sealer 1”, to delay messages during  $(n + 1) \times s$  seconds for transaction  $TX_1$  to be committed on the upper branch and for the attacker to seal more blocks on the lower branch than on the upper branch

To seal more blocks in one branch than another, the attacker maintains the partition during  $(n + 1) \times s$  seconds, where  $n$  is the number of sealers and  $s$  is the step duration in seconds that separates consecutive blocks. The reason is twofold.

- First, as mentioned earlier in Algorithm 1, Aura requires  $ns$  delay after a block is created to ensure that it is decided. Deciding a block on the victim side is necessary to make sure that  $TX_1$  gets committed. Given that both the size of the group on each side is  $\lfloor n/2 \rfloor + 1$  and that each sealer seals one after another, the attacker clone must also seal at least one block.
- Second, the attacker must ensure that the attacker branch is longer than the victim branch so that the attacker branch gets selected by Aura as the canonical branch. This can only be done if the attacker seals two blocks on the attacker branch, i.e., one extra block compared to the number of blocks it sealed on the victim side. As a result, the attacker needs to maintain the network partition for  $(n + 1) \times s$  seconds to get at least two turns in which it can seal a block.



*Example with 9 sealers.* For the sake of simplicity, Figure 1 depicts the Cloning Attack against Aura with a network partition where there are  $n = 9$  sealers and where  $(2 - n \bmod 2 = 1)$  sealer is malicious, namely “Sealer 1”. This attacker is thus present in both groups through its two cloned instances and gives the illusion that each group contains a majority of  $\lfloor n/2 \rfloor + 1 = 5$  sealers while one of the sealers in each group is actually a clone. As we can see, this attack translates into having Sealer 1 creating the last block (depicted with the red right-most block in the figure) only on the lower partition before merging the two partitions. By doing so, Sealer 1 makes sure that this branch will be the canonical branch whereas the upper branch will disappear. The attacker is thus guaranteed to double spend successfully.

## VII. THE CLONING ATTACK AGAINST CLIQUE

In this section, we apply the cloning attack against Clique. In Clique, the Cloning Attack does not require to take as long as in Aura. Unlike in Aura, a sealer of Clique can seal a block even when it is not its turn. Depending on their turn, some sealers may have to wait while others do not. These differences impact the way the attacker can influence the selection of one branch of a fork as the canonical chain and allow an attacker to double spend faster than in Aura.

### A. In-turn and out-of-turn sealers

The cloning attack against Clique differs from the one against Aura in the moment at which it starts delaying messages. Because the order of sealing is important in Clique, the attacker should ideally decide to start delaying the messages based on the sealer’s turn to seal a block.

When a sealer seals a block while it is his turn, we call this sealer an *in-turn sealer* and the block an *in-turn block* (cf. Alg. 3, line 29). There is at most one in-turn sealer to seal the current block in each partition of Clique. When a sealer seals a block while it is not his turn, we call this sealer an *out-of-turn sealer* and this block an *out-of-turn block* (cf. Alg. 3, line 32). As a sealer must wait for *sealer-limit* blocks between two blocks it seals, there are at most  $(n - \text{sealer-limit})$  potential out-of-turn sealers to seal a block. The in-turn block contributes a weight of 2 to the weight of its branch whereas the out-of-turn blocks contribute to 1 to the weight of its branch, hence sealing in-turn or out-of-turn impacts the decision regarding the branch selection process.

In addition, an in-turn sealer can append a block to the chain without waiting for any delay as shown in line 29 of Alg. 3. By contrast, an out-of-turn sealer has to wait for a random period as indicated at line 32 of Alg. 3. This mechanism gives the in-turn sealer some time to be the first to seal a block in his turn, but allows out-of-turn sealers to seal a block if the in-turn sealer is lagging.

As the canonical chain is chosen among the branches of a fork by comparing the sum of their block weights, the attacker must have a maximum number of in-turn sealers at the time of the partition to maximize the overall weight. Hence, to influence the selection of the branch as the canonical chain, the attacker must choose the proper turn to start delaying messages. If not done properly, the attacker risks to maximizing the weight of the branch where its transaction was included, limiting the chances of a successful double spending.

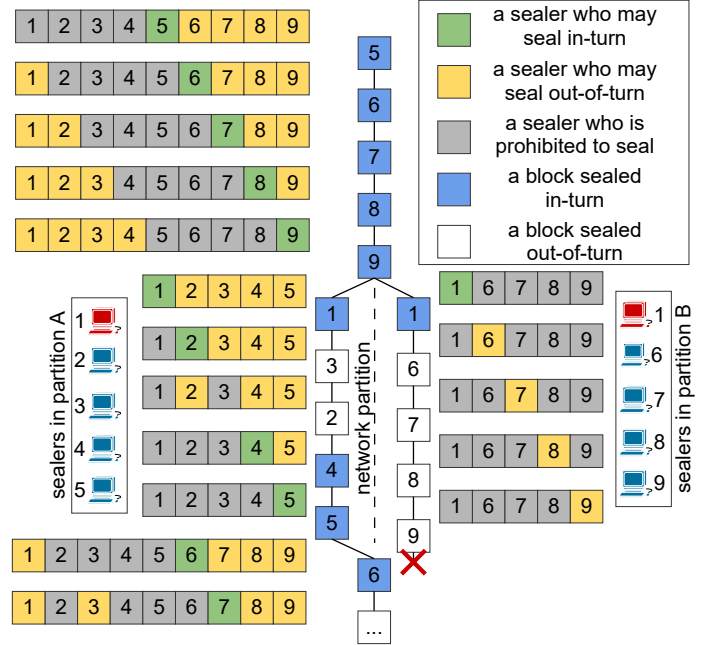


Figure 2. An execution of the in-turn cloning attack against Clique where  $n = 9$  sealers mine the blue blocks in-turn before the messages get delayed, after which each group seals five blocks, 3 in-turn blocks on the left group and 1 in-turn block on the right group

### B. Disorderer sealers to select a branch

Figure 2 depicts the execution of the attack with  $n = 9$  sealers and one attacker (Sealer 1) as time increases from top to bottom. Initially, the blockchain starts with block 5, indicating that the first block is sealed by Sealer 5. As times goes on, Sealers 6, 7, 8 and 9, seal one after the other the subsequent blocks of the blockchain. As there is no partition yet, the in-turn sealers are the first to sign these blocks during their respective turn, hence all blocks are in-turn blocks represented in blue in the figure. Next to each created block is a list of sealers that are either unable to seal (grey), in-turn sealers (green) or out-of-turn sealers (yellow).

Consider that Sealer 1, the attacker, performs the cloning and delays the network messages. Right after Sealer 9 sealed his block, Sealer 1 starts intercepting the messages between the group of sealers 2, 3, 4 and 5 on the left side and the group of sealers 6, 7, 8 and 9 on the right side. Note that Sealer 1 is represented on both sides because of the presence of one of its clones on each side. The resulting partition is indicated in Figure 2 with a fork of the blockchain into two branches. Right after the partition starts, Sealer 1 issues two conflicting transactions  $TX_1$  and  $TX_2$  on each side of the partition that will double spend. The two clones of Sealer 1 allow him to seal one block in each group. Note that these blocks are labelled 1 and represented in blue because Sealer 1 is the in-turn sealer at this point in time. After sealing, Sealer 1 is no longer able to seal any block due to the *sealer-limit*, hence Sealer 1 is depicted in grey in both groups.

On the right side of the partition, we can see that Sealer 6 seals the following block, even though it is not the in-turn sealer at this moment. This is because the in-turn sealer, Sealer 2, cannot communicate with this group as the network

is partitioned. For some reason it might also be the case on the left side of the partition that Sealer 2 is not fast enough to seal the next block and that another sealer, say Sealer 3, manages to seal it before. Note that this can happen as the delay Sealer 3 has to wait before sealing is a random number that can be null (cf. Alg. 3, line 32). However, this last seal from Sealer 3 prevents it from sealing the next block in-turn as it has to wait for the *sealer-limit*, hence the next block is again out-of-turn. The process continues where sealers on the left side seal in-turn whereas sealers on the right side seal out-of-turn.

Finally, the attacker does no longer need to delay the messages and can stop the partition as both transactions  $TX_1$  and  $TX_2$  are now successfully committed. In fact, the transactions are both now in the first block of a series of  $\lfloor n/2 \rfloor + 1 = 5$  consecutive blocks, which is sufficient for all Clique users to consider these transactions as committed because their block is decided as indicated at line 49 of Algorithm 3. We can conclude that the weight gained by the branch on the left side during the partition is  $3 \times 2 + 2 \times 1 = 8$  because it contains 3 in-turn blocks and 2 out-of-turn blocks. By contrast, the weight gained by the branch on the right side during the partition is  $1 \times 2 + 4 \times 1 = 6$  because it contains 1 in-turn block and 4 out-of-turn blocks. It follows from the difference in weight of the two branches that the heaviest branch on the left side is chosen as the canonical branch whereas the lightest branch on the right side is simply discarded by the protocol (cf. Alg. 3, line 44).

### C. Attack regardless of the order of sealers

Note that even if the attacker does not know the topology, there is a way to attack Clique. The attack is slightly different from the previous one as it relies on the possibility for the attacker to become the only sealer able to seal a block on both sides of the partition. The attacker can simply seal a single block on the victim branch, and keep sealing blocks on the attacker branch. In the worst case scenario for the attacker, all the  $\lfloor n/2 \rfloor + 1$  upcoming in-turn sealers end up on the victim side, which will maximize the weight of the branch on the victim side gained during the partition. Recall that the *sealer-limit* is always  $\lfloor n/2 \rfloor + 1$  in Clique (Alg. 3, line 14). Now, if the attacker stops sealing a second block on the victim side, then the maximum weight gained on this side during the partition will be  $(sealer-limit \times 2)$ . The attacker simply needs to keep sealing on the other branch until the gained weight on this branch reaches  $(sealer-limit \times 2 + 1)$ . In this case, the attacker successfully double spends regardless of the sealer turn in each group.

## VIII. EXPERIMENTS

In this section, we present the double spending results of the Cloning Attack in both Aura and Clique. For ethical reasons, we do not experiment our attack against PoA networks that are currently in production. Instead, we first present our experimental setup then detail the risk for an attacker to perform double spending in both Aura and Clique within our network.

### A. Testnet setup

To practically observe the chance of successful double spending using the approaches described in the previous sec-

tions we have created our own PoA blockchain networks, experimented the attacks and measured their success rate empirically.

Our testnet consists of 10 Ubuntu 18.04 Virtual Machines (VMs) on our OpenStack private cloud; each VM is provided with 1 virtual CPU core and 2GB of memory. These VMs are placed into two subnets, 5 VMs each; they are connected through 5 linux virtual routers and a physical Ethernet switch with dedicated VLAN. An instance of either Parity-Ethereum-v2.0.8-stable with Aura or geth-1.8.21-stable with Clique runs on each VM.

All of these instances are peering with each other to form the blockchain network. While we have 10 Ethereum instances in total, our PoA blockchain employed only 9 unique private keys for sealers; the last instance instead uses the same key as the first one as explained in Section V where one instance is seen as a clone of the other. As of writing, neither Aura nor Clique incorporate a mechanism to prevent private key reusing. One can simply copy a key and configuration files from one instance to another in order share the private key, other instances will simply accept the connection from the clone instance in the same way as the original.

In our experiments, the attacker is a Byzantine (or malicious) sealer with the intention to achieve double spending. This attacker is provided the capability to transiently partition the network into two sealer groups: the attacker and the victim group. We refer to the attacker group as the group of sealers whose blocks sealed during the network partition are intended to be adopted as a part of canonical chain, while we refer to the victim group as the group of sealers whose sealed blocks are intended to be discarded after the fork is resolved.

To grant the capability to partition the network, we allow our attacker to cut the network connectivity between two subnets using a firewall feature on the linux routers. Note that the same result is achieved using a man-in-the-middle attack though ARP-spoofing in a local area network or with BGP-hijacking in other networks [13]. The attacker also has control over 2 Ethereum instances (2 VMs) that share the same private key used to seal the blocks.

The attacker aims to partition the network right before their turn to seal the block, where each sealer group must contain one VM that is under the control of the attacker. To begin the attack, our attacker actively checks the owner of the current turn every 10 ms in order to partition the network close to the right timing. Right after the network partition, the attacker issues one transaction to each sealer group; these two conflicting transactions  $TX_1$  and  $TX_2$ , for example Alice is giving all of her coins to Bob in  $TX_1$  and gives the same coins to Carol in  $TX_2$ .

After issuing the transactions, the network partition is maintained during a period that depends on which PoA algorithm as explained below. When the fork is resolved at the end of the network partition, we look at the resulting branch of the fork which has been adopted as a canonical chain as well as the status of transactions.

A double spending is considered successful only if:

- 1) A transaction issued to the victim group is committed before the end of the network partition;



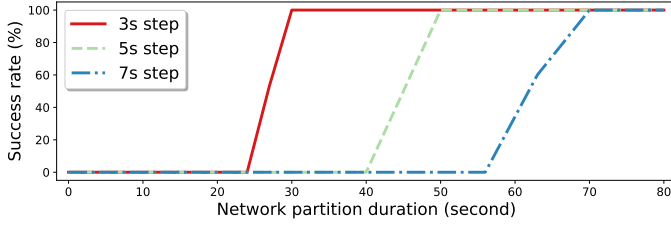


Figure 3. The success rate of double spending with the Cloning Attack in Aura

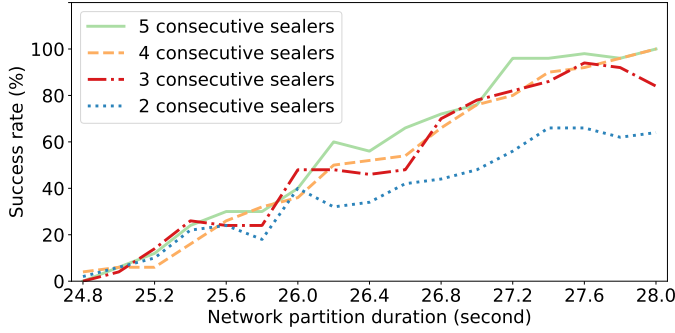


Figure 4. The success rate of the Cloning Attack double spending in Clique as the duration of the network partition increases and for different numbers of consecutive possible in-turn sealers

- 2) the blocks sealed by the attacker group during the fork have been adopted as a part of the canonical chain after the end of the partition; and
- 3) the resulting canonical chain does not contain a transaction issued to the victim group.

### B. Running the Cloning Attack against Aura

We experiment the Cloning Attack in Aura by varying the step duration and network partition duration. We chose Step durations 3, 5, and 7 seconds in order to observe their impact of the minimum partition duration that makes the attack successful.

We maintain the network partition to match the step duration in use, such that for example a 24, 27, and 30 second partition duration corresponds respectively to the 8th, 9th and 10th step for a 3 second step duration, respectively. We divide the sealers into two groups, such that apart from the two attacker instances, the placement of the reminder sealers is randomly but equally balanced between the two partitions. We do ensure, however, that both groups have an equal number of instances, which is 5, and each group contains one of two instances under the control of the attacker. The values plotted for each combination of step duration and network partition duration are the averages over 30 runs.

Figure 3 presents the double spending success rate of 3, 5 and 7 second step durations. In all three cases, the obtained results show a similar trend. As expected, achieving a successful double spending is impossible in all step durations at the 8th step or any earlier step, namely 24, 40 and 56 seconds for 3, 5, and 7 second step durations, respectively. Indeed, these attempts fail because any attack attempt among

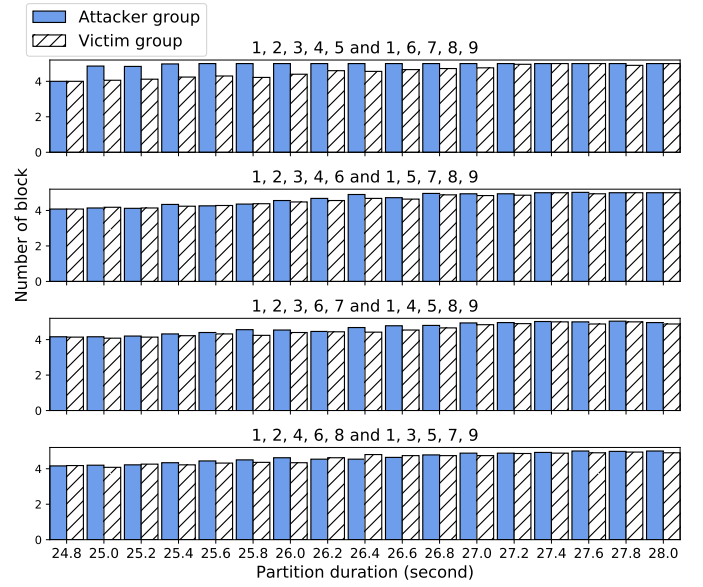


Figure 5. The average number of blocks created in Clique depending on the partition duration and the sealer distribution across partitions

these runs could neither commit the transaction in the victim group nor force the block sealed by the attacker group to be adopted as a canonical chain when the network partition ends.

However, we can observe that as expected, the chances of successful double spending at the 9th step falls within the range between 50-60%. Even though both groups are provided enough time to seal 5 blocks in order to commit the transactions, the attacker still cannot force a particular branch of the fork to be adopted. The variation at this point is due to the randomness of Ethereum instance placement during our experiment.

For all three step durations, at the 10th step and any step thereafter, the attack is always successful (100% chances). This is due to the attack technique in use that allows the attacker to force a branch of the fork to be adopted. Overall, we can see that a longer step duration requires a longer period of network partition in order to achieve a successful double spending, which confirms our expectations.

### C. Running the Cloning Attack against Clique

We experiment the Cloning Attack on Clique while varying the partition duration and the way sealers are distributed between two partitions.

The variations in the sealer divisions are included in the experiments in order to capture the changes in weight of each branch as a result of the sealing sequences. In particular, we experiment with the 4 sealer divisions presented below with different number of consecutive sealers:

- 5 consecutive sealers: 1, 2, 3, 4, 5 in the attacker group and 1, 6, 7, 8, 9 in the victim group;
- 4 consecutive sealers: 1, 2, 3, 4, 6 in the attacker group and 1, 5, 7, 8, 9 in the victim group;
- 3 consecutive sealers: 1, 2, 3, 6, 7 in the attacker group and 1, 4, 5, 8, 9 in the victim group;

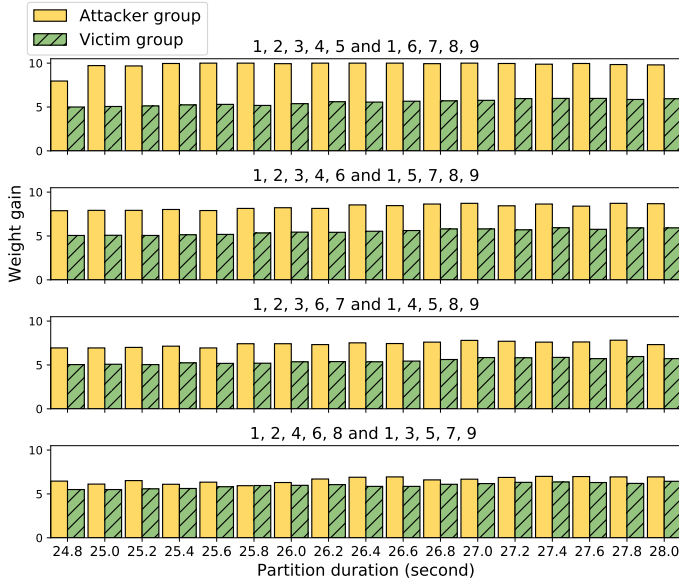


Figure 6. The average weight gained in Clique depending on the partition duration and the sealer distribution across partitions

- 2 consecutive sealers: 1, 2, 4, 6, 8 in the attacker group and 1, 3, 5, 7, 9 in the victim group.

The partition duration is based on the block duration in use, which is fixed to 5 seconds in all our Clique experiments. Since our testnet setup consists of 9 sealers in total, to commit a transaction during a partitioning, at least 5 blocks must be sealed in such a period. In the best case where 5 sealers could seal 5 in-turn blocks, the minimum duration required for the attack to succeed is equal to  $5 \times 5 = 25$  seconds. In other cases where at least 1 out of 5 blocks is sealed out-of-turn, however, the required duration exceeds 25 seconds.

Based on our knowledge of the time necessary for the algorithm to seal 5 blocks, we vary the duration from 24.8 to 28.0 seconds in an incremental step of 200 milliseconds. The range of duration allows to take into account the random delay of out-of-turn sealers as shown in Algorithm 3 and yet to capture the behavior of the system from the point where only 4 blocks can be sealed to the point where 5 blocks can be sealed.

For each run we keep a record of whether the double spending was successful, which block was sealed by which sealer, the weight gained during the partition for each fork, and the number of blocks created during the network partition. The values averaged over 50 runs are depicted in the charts for each partition duration.

Figure 4 reveals the double spending success rate for the four aforementioned sealer divisions while Figure 5 and Figure 6 show the number of sealed blocks and the weight gained during the network partition, respectively. We observe that the success rate in Figure 4, follows a similar trend for all of 4 grouping variations; the longer the partition duration, the higher the chance of successful double spending.

The shortest partition duration value in the chart, 24.8 second, gives the lowest success rate regardless of the sealer

division. This low success rate for short duration can be explained by the number of blocks sealed during the network partition. Indeed, due to the limited partition duration, the victim group is rarely able to seal five blocks during the network partition as shown in Figure 5, thus a transaction issued by the victim group could not be committed and the attack fails. When the partition takes longer, we can see that the victim group is able to seal five blocks.

When the partition duration is less or equal to 26 seconds, there is no noticeable difference between the four different sealer divisions. In the case of two consecutive possible in-turn sealers and when the partition duration is greater than 26 seconds, however, the success rate is lower than the other three divisions. This phenomenon can be explained by the weight gained during the network partition as shown in Figure 6. In fact, in case of 2 consecutive possible in-turn sealers, the difference in weight gained between attacker and victim branches becomes relatively low; this gap narrows with the increase in the partition duration.

## IX. ANALYSIS AND COUNTERMEASURES

We begin this section by comparing the vulnerabilities of Aura and Clique to the Cloning Attack resulting from our experiments in Section VIII. Next, we analyse further the Aura algorithm and discuss its implication to the blockchain safety and liveness. Finally, we present potential countermeasures against the Cloning Attack.

### A. Comparison between Aura and Clique

In this section, we explain why the Cloning Attack against Aura can always be successful whereas the Cloning Attack against Clique is much faster but not always successful.

As detailed in Section IV, one of the main differences between Aura and Clique resides in the predictability of the sequence of sealers. In fact, in Aura the sequence is strictly enforced whereas in Clique this sequence may change depending on the difference between a random number and the network communication delay. This slight algorithmic difference has however significant consequences on consensus algorithms resilience to double spending attacks using our proposed Cloning Attacks.

On the one hand and as we have demonstrated in Section VIII, due to its strict enforcement of sealing order, Aura is vulnerable to the Cloning Attack in case of network partition. Performing the Cloning Attack against Aura, the attacker does not need to know anything about the identity of the sealers nor does it need to know their order. Thus, a malicious sealer only needs to partition the overlay network using classical network attacks such as BGP hijacking to succeed in double spending with a 100% chance of success.

On the other hand, double spending without topology information on Clique is possible, but the attack against Clique is about twice as fast as against Aura when the topology is known. Indeed, as we have presented in Section VIII, the knowledge of potential next in-turn sealer greatly influences the chance of double spending. When the attacker is capable of isolating the next  $\lfloor n/2 \rfloor + 1$  sealers, it is able to perform the double spending attack with 100% success rate. By contrast,

the knowledge of only the next two in-turn sealers only guarantees a success rate of 60% maximum.

Interestingly, when considering the attacks against both Aura and Clique without the knowledge of the topology, it appears that attacking Clique can be even slower than attacking Aura. The reason is that in the worst case scenario where all in-turn sealers are on the victim side, the attacker will have to obtain a branch that is twice as large as the victim branch before it can double spend. Growing this branch would take more time than executing the Cloning Attack on Aura. But overall, even without knowledge of the topology both Aura and Clique consensus algorithms are vulnerable to a malicious sealer aiming at double spending.

### B. Modifications to make Aura and Clique safe and live

Our attack violates the safety of Aura and Clique, in that it leads the system to an undesirable state, where coins can be stolen.<sup>2</sup> As we explain here, this problem can be mitigated by simply selecting the right number of sealers  $|V|$  necessary to decide a block. Determining  $|V|$  to ensure safety may however be insufficient, as it does not ensure that the system makes progress, which is a liveness problem. So we also explain how to choose  $|V|$  in the case of Aura to ensure that it is both safe and live. Changing  $|V|$  can be easily achieved by modifying the boolean condition under which a block is decided at line 31 of Algorithm 1. Note that while we focus on Aura, the same arguments also apply to Clique.

Figure 7 depicts the relation between the desired fault tolerance of the system and the number of sealers  $|V|$  necessary to decide a block to ensure Aura's safety. This analysis is helpful as it allows us, depending on the targeted fault tolerance, to decide the minimal number of sealers that are necessary to decide a block in a synchronous or partially synchronous network so as to ensure that Aura is safe and live. In particular, we consider two distinct cases, whether the network communication is synchronous or partially synchronous. The asynchronous case is ignored here as consensus would be impossible in all interesting cases where  $t > 0$  [18].

**Synchronous network.** In the synchronous case, Aura's safety is guaranteed when  $|V| > t$ , which is represented by the area above the orange dotted line  $|V| = t + 1$  in Figure 7. In fact, to ensure that there exists a block sealed by an honest sealer, one has to wait for the number of blocks sealed by distinct sealers to be greater than the total number of malicious sealers, which requires the number of sealers  $|V|$  to exceed the number  $t$  of malicious sealers.

- **Case  $t = 0$ :** When all the sealers are honest,  $|V| = 1$  is sufficient to consider finality of a block, i.e., the transactions within a block can be considered committed instantly because one may safely assume that a block will be delivered to all the sealers in the known upper bounded time.
- **Case  $0 < t < n$ :** When there is at least one malicious sealer, one must ensure that at least one honest sealer

<sup>2</sup>Safety (resp. liveness) is often referred to as the property of a system to guarantee that "a bad thing will never happen" (resp. "a good thing eventually happens") [23], [1].

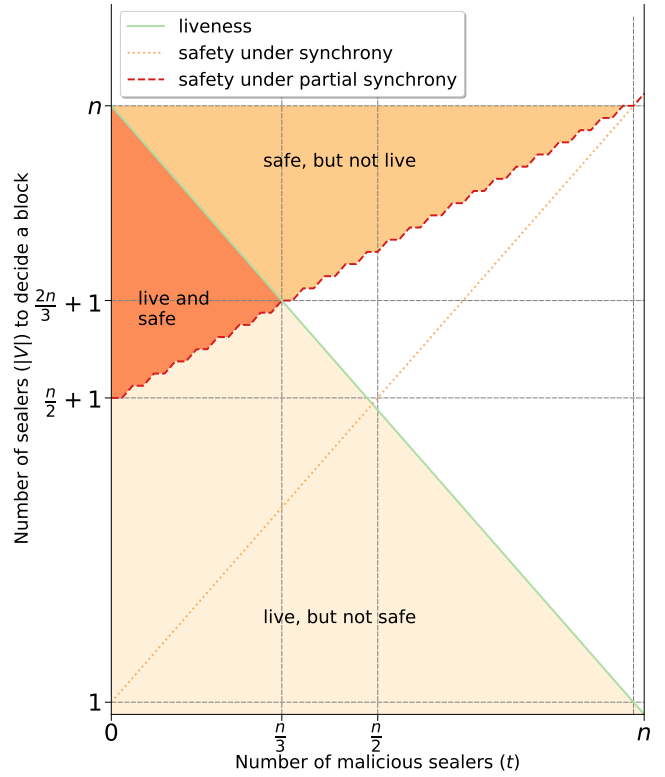


Figure 7. Required number of sealers  $|V|$  to decide a block depending on the number of malicious sealers  $t$ : Aura is safe only when  $|V| > \frac{n+t}{2}$ , Aura is live only when  $|V| < n - t$ , so Aura is both safe and live only when  $\frac{n+t}{2} < |V| < n - t$ .

has recently appended a block to the blockchain before considering finality of its previous blocks because a malicious sealer may violate the protocol by introducing a malformed block or appending a block out-of-turn.

- **Case  $t = n$ :** In this case, we can see that  $|V|$  should be strictly greater than  $n$ , which is impossible by definition. This illustrates that it is impossible for Aura to work when  $t = n$ .

To conclude, the highest fault tolerance  $t$  that Aura can tolerate in a synchronous system is  $t = n - 1$  by requiring blocks from all the  $|V| = n$  sealers to decide a block.

**Partially synchronous network.** Requiring  $\lfloor \frac{n}{2} \rfloor + 1$  sealers to decide one block is insufficient to tolerate unpredictable message delays between two partitions of honest sealers when  $t > 0$ . As an example, recall that Aura aims at tolerating up to a minority  $\lfloor \frac{n}{2} \rfloor - 1$  of malicious sealers, however, if a majority  $\lfloor \frac{n}{2} \rfloor + 1$  of sealers are sufficient to decide a block, then a group of only 2 honest sealers helped with the  $t = \lfloor \frac{n}{2} \rfloor + 1$  malicious sealers would be sufficient to decide one block. This does not prevent the other honest sealers in another partition from deciding a conflicting block with the help of the clones of the  $t$  malicious sealers. This is the reason why the dashed red line on Figure 7 indicates that  $|V| \geq \lfloor \frac{n+t}{2} \rfloor + 1$  is necessary to guarantee that a majority of honest sealers sign a block so that no other conflicting blocks can be decided. More precisely, here are the 3 interesting cases to consider:



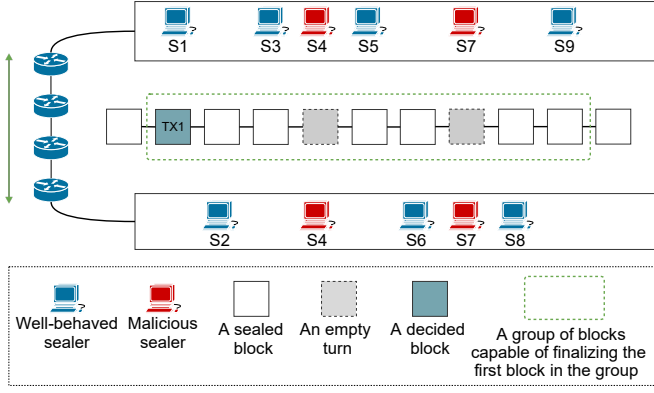


Figure 8. An Aura execution using  $|V| > \frac{2n}{3}$  and  $t < \frac{n}{3}$  ( $n = 9$ ,  $|V| = 7$  and  $t = 2$ ) where transaction  $TX1$  is committed

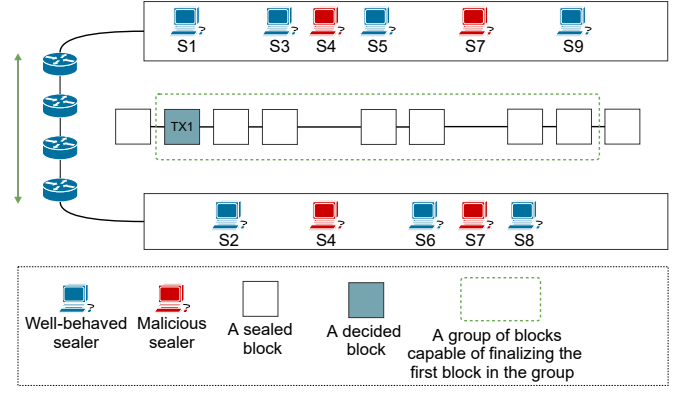


Figure 10. A Clique execution using  $|V| > \frac{2n}{3}$  and  $t < \frac{n}{3}$  ( $n = 9$ ,  $|V| = 7$  and  $t = 2$ ) where transaction  $TX1$  is committed

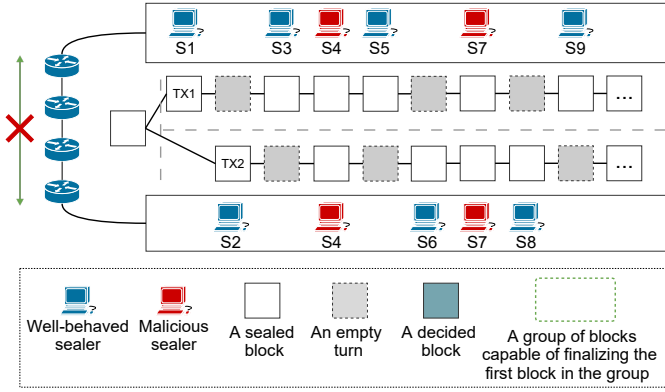


Figure 9. An Aura execution using  $|V| > \frac{2n}{3}$  and  $t < \frac{n}{3}$  ( $n = 9$ ,  $|V| = 7$  and  $t = 2$ ) where neither transaction  $TX1$  nor  $TX2$  are committed

- **Case  $t = 0$ .** When all the sealers are honest, the algorithm only needs at least a majority of sealers,  $|V| \geq \lfloor \frac{n}{2} \rfloor + 1$ , before considering that a block is decided.
- **Case  $0 < t < \frac{n}{3}$ .** To ensure that a majority of the honest sealers seal a block for one block to be decided, we need  $|V| \geq \lfloor \frac{n+t}{2} \rfloor + 1$  sealers to seal blocks. As there are  $n - t$  honest sealers, a majority of them contains  $\lfloor \frac{n-t}{2} \rfloor + 1$  sealers. As there are  $t$  malicious sealers, we need strictly more than half  $\frac{n-t}{2}$  of the honest sealers and the  $t$  malicious ones to seal blocks, which leads to  $|V| \geq \lfloor \frac{n+t}{2} \rfloor + 1$ . The two upper triangles of Figure 7 depict these conditions, under which Aura is safe.
- **Case  $t \geq \frac{n}{3}$ .** Interestingly, when  $t \geq \frac{n}{3}$ , it is impossible to guarantee that at least  $|V| \geq \lfloor \frac{n+t}{2} \rfloor + 1$  will seal blocks. In fact, this would imply that strictly more than  $\frac{2n}{3}$  sealers seal a block. However, as  $t \geq \frac{n}{3}$  this would also mean that at least one malicious sealer seals a block, which cannot be guaranteed as malicious sealers may choose, by definition, to not follow the protocol. The upper-right triangle of Figure 7 depicts the conditions under which Aura is safe because  $|V| \geq \lfloor \frac{n+t}{2} \rfloor + 1$  but not live because  $t \geq \frac{n}{3}$ .

Even with only one malicious sealer ( $t = 1$ ), the current

implementation of Aura cannot guarantee safety. The reason is that Aura claims that a majority of sealers is enough as long as  $t < \frac{n}{2}$ . However, one can find a counter-example where  $|V| = \lfloor \frac{n}{2} \rfloor + 1$  that falls in the unsafe area shown in Figure 7.

### C. Simple safe countermeasure

As explained in the introduction and for ethical reasons, we disclosed early a simple, yet naive, countermeasure to make both Aura and Clique safe. This countermeasure has been implemented in the xDai blockchain as acknowledged in their white paper [3]. This first countermeasure was not as precise as of today and required  $|V| > \frac{2n}{3}$  and  $t < \frac{n}{3}$  [15] to ensure safety but ignoring liveness guarantees. Interestingly, the same conditions of this countermeasure are required by IBFT [7] and recent work suggested that IBFT is not live either [36].

1) *Making Aura safe:* As an example, Figures 8 and 9 demonstrate how a modified version of Aura where  $n = 9$ ,  $|V| = 7$ ,  $t = 2$  where S4 and S7 and malicious. These two sealers are allowed to be silent or even seal the blocks in all groups whenever a network partition occurs. As indicated with a green dash frame in Figure 8, this version of Aura can commit  $TX1$  even though S4 and S7 do not contribute any block to the chain in their turns. The S4 turn can be left empty and S5 may simply continue sealing a block after that; the same goes for S8 after the end of S7 turn. In this case, the other 7 honest sealers alone are sufficient to decide a block. As indicated in the execution of Figure 9, however, it is impossible to have at least 7 sealers on both sides at the same time. Therefore, it is impossible to commit the transactions issued to both partitions concurrently, even though Aura allows their sealers to continue sealing more blocks during the network partition.

2) *Making Clique safe:* The same condition also applies to Clique. To illustrate how a modified Clique with two-thirds consensus works, let us use the same setup with a blockchain network of  $n = 9$  sealers: S1 to S9 with two malicious sealers S4 and S7. As indicated with green dash frame in Figure 10, the blockchain network can still commit  $TX1$ . Instead of S4 and S7, Clique allows other sealers to continue sealing the blocks in an out-of-turn fashion without any gap and eventually reaches the 7 blocks requirement. In another case where the

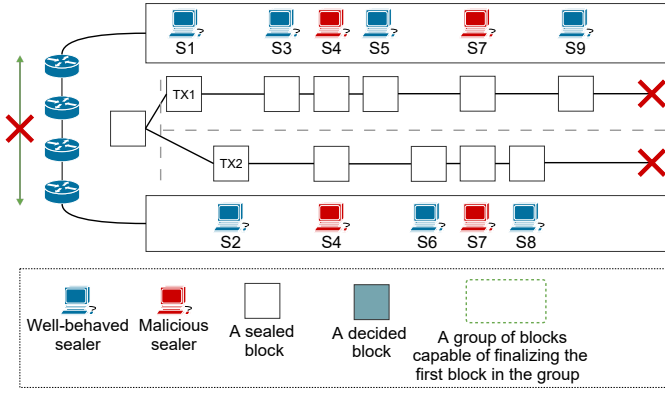


Figure 11. A Clique execution using  $|V| > \frac{2n}{3}$  and  $t < \frac{n}{3}$  ( $n = 9$ ,  $|V| = 7$  and  $t = 2$ ) where neither transaction  $TX1$  nor  $TX2$  are committed

network partition occurs, if the number of sealers in a partition is not greater than two-thirds of all sealers, such a partition will be stuck even before sealing enough blocks to commit a transaction. Figure 11 shows a concrete example where both partitions get stuck due to the lack of sealers; one partitions contains only 6 sealers whereas there are only 5 sealers in the other.

#### D. Replacing the consensus by a formally verified alternative

On Figure 7, it is easy to identify that the simple countermeasure presented above, which requires  $|V| > \frac{2n}{3}$  and  $t < \frac{n}{3}$ , does not ensure liveness. This is because the protocol can reach an undesirable situation depicted by the upper-right triangle of the figure. To remedy this issue, it is sufficient to ensure that  $\frac{n+t}{2} < |V| < n - t$ . Note that this implicitly guarantees that  $t < \frac{n}{3}$  anyway. As a result, under these conditions Aura remains safe and live; likewise, the same conditions also apply to Clique.

Finally, a radically different countermeasure that also offers safety and liveness is to use a deterministic consensus algorithm that is partially synchronous in that it tolerates arbitrary delays. PBFT [6] is one example but is not designed to scale outside a small network as it relies on a leader. DBFT [8] is a leaderless deterministic partially synchronous consensus algorithm that was especially designed to scale to blockchain systems. One of the variant of DBFT has recently been formally proved correct with model checking [38]. In addition, DBFT is time optimal and resilience optimal. It has been shown that DBFT is resilient to double spending attacks, as it is not possible for a blockchain building upon it, like the Red Belly Blockchain [9], to fork.

## X. CONCLUSION

To cope with the drawbacks of proof-of-work (PoW), Byzantine fault tolerance has been introduced in mainstream blockchains in the form of proof-of-authority (PoA) where sufficient sealers  $|V|$  among  $n$  must seal a block despite a minority  $t$  of them being malicious. In this paper, we have detailed three of the most popular PoA algorithms in pseudocode to dissect their executions under various conditions. Although they feature more elaborate consensus algorithms

than the classical longest branch strategy typically used in PoW blockchains, our study shows that PoA fails at offering a safe alternative.

The Cloning Attack allows malicious participants to double spend in Ethereum instances using Clique and Aura consensus protocols. This attack leverages the possibility of cloning public-private key pairs in Ethereum in two distinct instances before executing a network attack between two groups to lure the sealers on the number of votes that occurred in each group.

Our findings inspired by the theory of Byzantine fault tolerance capture precisely the relation between the fault tolerance  $t$  and the amount of sealers  $|V|$  that must sign a block for PoA to be safe and live. Other countermeasures include simply replacing the Aura and Clique algorithms within Ethereum parity and geth by a recent Byzantine deterministic consensus especially designed for blockchains and that has been proved correct.

## ACKNOWLEDGEMENTS

We wish to thank the anonymous reviewers and Srdjan Capkun, our shepherd, for their comments, Kirill Pimenov, Head of Security at Parity Technologies, for going through the technicalities of the Cloning Attack and confirming the theoretical vulnerability of the parity software and to Martin Holst Swende, Security Lead at the Ethereum Foundation for confirming the possibility of the attack on geth. This research is in part supported under Australian Research Council Discovery Projects funding scheme (project number 180104030) entitled “Taipan: A Blockchain with Democratic Consensus and Validated Contracts” and Australian Research Council Future Fellowship funding scheme (project number 180100496) entitled “The Red Belly Blockchain: A Scalable Blockchain for Internet of Things”.

## REFERENCES

- [1] B. Alpern and F. B. Schneider, “Recognizing safety and liveness,” *Distributed computing*, vol. 2, no. 3, pp. 117–126, 1987.
- [2] S. D. Angelis, L. Aniello, R. Baldoni, F. Lombardi, A. Margheri, and V. Sassone, “PBFT vs Proof-of-Authority: Applying the CAP theorem to permissioned blockchain,” in *Proc. 2nd Italian Conf. Cyber Secur.*, 2018, available: <https://eprints.soton.ac.uk/id/eprint/415083>.
- [3] I. Barinov, V. Arasev, A. Fackler, V. Komendantskiy, A. Gross, A. Kolotov, and D. Isakova, “POSDAO: Proof of stake decentralized autonomous organization,” SSRN, Tech. Rep., 2019, available at SSRN: <https://ssrn.com/abstract=3368483> or <http://dx.doi.org/10.2139/ssrn.3368483> as of 3rd Sept. 2019.
- [4] I. Barinov, V. Baranov, and P. Khahuln, “POA network white paper,” Sept. 2018, <https://github.com/poanetwork/wiki/wiki/POA-Network-Whitepaper>.
- [5] “Blockchain Supply Chain Proof of Concept,” Jan. 2019, <https://blog.apla.io/blockchain-supply-chain-proof-of-concept-5a6d361c943c>. [Online]. Available: <https://blog.apla.io/blockchain-supply-chain-proof-of-concept-5a6d361c943c>
- [6] M. Castro and B. Liskov, “Practical Byzantine fault tolerance,” in *Proc. 3rd Symp. Operating Syst. Design and Implementation (OSDI)*, ser. OSDI ’99. Berkeley, CA, USA: USENIX Association, 1999, pp. 173–186. [Online]. Available: <http://dl.acm.org/citation.cfm?id=296806.296824>
- [7] J. Chase, “Quorum whitepaper,” <https://github.com/jpmorganchase/quorum>.
- [8] T. Crain, V. Gramoli, M. Larrea, and M. Raynal, “DBFT: efficient leaderless Byzantine consensus and its application to blockchains,” in *Proc. 17th IEEE Int. Symp. Netw. Comp. and Appl (NCA)*, 2018, pp. 1–8.

- [9] T. Crain, C. Natoli, and V. Gramoli, "Evaluating the Red Belly Blockchain," arXiv, Tech. Rep. 1812.11747, 2018.
- [10] D. Dolev and A. Yao, "On the security of public key protocols," *IEEE Trans. Inf. Theory*, pp. 198–208, 1983.
- [11] J. R. Douceur, "The Sybil Attack," in *Int. Workshop Peer-To-Peer Syst.* Springer, 2002, pp. 251–260.
- [12] C. Dwork, N. Lynch, and L. Stockmeyer, "Consensus in the presence of partial synchrony," *J. ACM*, vol. 35, no. 2, pp. pp.288–323, 1988.
- [13] P. Ekparinya, V. Gramoli, and G. Jourjon, "Impact of man-in-the-middle attacks on Ethereum," in *Proc. 37th IEEE Int. Symp. Reliable Distrib. Syst. (SRDS)*, Oct. 2018.
- [14] —, "The Attack of the Clones against Proof-of-Authority," in *Community Ethereum Development Conference (EDCON'19)*, Apr. 2019.
- [15] —, "The Attack of the Clones against Proof-of-Authority," arXiv, Tech. Rep. 1902.10244v1, Feb. 2019.
- [16] "Ethereum Proof-of-Authority Consortium - Azure," <https://docs.microsoft.com/en-us/azure/blockchain/templates/ethereum-poa-deployment>. [Online]. Available: <https://docs.microsoft.com/en-us/azure/blockchain/templates/ethereum-poa-deployment>
- [17] I. Eyal and E. G. Sirer, "Majority is not enough: Bitcoin mining is vulnerable," in *Proc. FC 2014*, 2014, pp. 436–454.
- [18] M. J. Fischer, N. A. Lynch, and M. S. Paterson, "Impossibility of distributed consensus with one faulty process," *J. ACM*, vol. 32, no. 2, pp. 374–382, Apr. 1985.
- [19] Y. Gilad, R. Hemo, S. Micali, G. Vlachos, and N. Zeldovich, "Algorand: Scaling Byzantine agreements for cryptocurrencies," in *Proc. 26th Symp. Operating Syst. Principles*, 2017, pp. 51–68.
- [20] E. Heilman, A. Kendler, A. Zohar, and S. Goldberg, "Eclipse attacks on bitcoin's peer-to-peer network," in *24th USENIX Secur. Symp.*, 2015, pp. 129–144.
- [21] "Kovan Testnet," <https://kovan-testnet.github.io/website/>. [Online]. Available: <https://kovan-testnet.github.io/website/>
- [22] J. Kwon, "Tendermint, consensus without mining," 2015, unpublished manuscript.
- [23] L. Lamport, "Proving the correctness of multiprocess programs," *IEEE Trans. Softw. Eng.*, vol. 3, no. 2, pp. 125–143, 1977.
- [24] L. Lamport, R. Shostak, and M. Pease, "The Byzantine generals problem," *ACM Trans. Program. Lang. Syst.*, vol. 4, no. 3, pp. 382–401, Jul. 1982.
- [25] "Launch Enterprise-Ready Blockchain Networks on AWS in Minutes with Kaleido—a ConsenSys Solution," May 2018, <https://aws.amazon.com/blogs/apn/launch-enterprise-ready-blockchain-networks-on-aws-in-minutes-with-kaleido-a-consensys-solution/>. [Online]. Available: <https://aws.amazon.com/blogs/apn/launch-enterprise-ready-blockchain-networks-on-aws-in-minutes-with-kaleido-a-consensys-solution/>
- [26] Y.-T. Lin, "Istanbul Byzantine Fault Tolerance · Issue #650 · ethereum/EIPs," <https://github.com/ethereum/EIPs/issues/650>.
- [27] L. Luu, Y. Velner, J. Teutsch, and P. Saxena, "Smartpool: Practical decentralized pooled mining," in *26th USENIX Secur. Symp.*, 2017, pp. 1409–1426.
- [28] A. Miller, Y. Xia, K. Croman, E. Shi, and D. Song, "The honey badger of BFT protocols," in *Proc. 2016 ACM SIGSAC Conf. Comput. and Commun. Secur.* New York, NY, USA: ACM, 2016, pp. 31–42. [Online]. Available: <http://doi.acm.org/10.1145/2976749.2978399>
- [29] C. Natoli and V. Gramoli, "The Blockchain Anomaly," in *2016 IEEE 15th Int. Symp. Netw. Comput. and Applications (NCA)*, Oct. 2016, pp. 310–317.
- [30] —, "The balance attack or why forkable blockchains are ill-suited for consortium," in *47th IEEE/IFIP Int. Conf. Dependable Syst. and Netw. (DSN)*, Jun 2017.
- [31] K. Nayak, S. Kumar, A. Miller, and E. Shi, "Stubborn mining: Generalizing selfish mining and combining with an eclipse attack," in *Proc. IEEE EuroS&P 2016*, 2016, pp. 305–320.
- [32] "Parity Documentation - Aura - Authority Round," <http://wiki.parity.io/Aura.html>. [Online]. Available: <http://wiki.parity.io/Aura.html>
- [33] M. C. Pease, R. E. Shostak, and L. Lamport, "Reaching agreement in the presence of faults," *J. ACM*, vol. 27, no. 2, pp. 228–234, 1980.
- [34] "Rinkeby: Ethereum Testnet," <https://www.rinkeby.io/#stats>. [Online]. Available: <https://www.rinkeby.io/#stats>
- [35] M. Rosenfeld, "Analysis of hashrate-based double spending," 2014.
- [36] R. Saltini and D. Hyland-Wood, "Correctness analysis of ibft," 2019.
- [37] E. Shi, "Analysis of deterministic longest-chain protocols." *IACR Cryptology ePrint Archive*, vol. 2018, p. 1079, 2018.
- [38] P. Tholoniat and V. Gramoli, "Formal verification of blockchain Byzantine consensus," arXiv, Tech. Rep. 1909.07453, 2019. [Online]. Available: <https://arxiv.org/abs/1909.07453>
- [39] I. Weber, Q. Lu, A. B. Tran, A. Deshmukh, M. Gorski, and M. Strazds, "A platform architecture for multi-tenant blockchain-based systems," in *IEEE Int. Conf. on Softw. Architecture*, 2019, pp. 101–110.