

ÆGIS: Shielding Vulnerable Smart Contracts Against Attacks

Christof Ferreira Torres
SnT, University of Luxembourg
Luxembourg, Luxembourg
christof.torres@uni.lu

Mathis Baden
SnT, University of Luxembourg
Luxembourg, Luxembourg
mathis.steichen@uni.lu

Robert Norvill
SnT, University of Luxembourg
Luxembourg, Luxembourg
robert.norvill@uni.lu

Beltran Borja Fiz Pontiveros
SnT, University of Luxembourg
Luxembourg, Luxembourg
beltran.fiz@uni.lu

Hugo Jonker
¹Open University of the Netherlands
Heerlen, Netherlands
²Radboud University
Nijmegen, Netherlands
hugo.jonker@ou.nl

Sjouke Mauw
SnT, University of Luxembourg
Luxembourg, Luxembourg
sjouke.mauw@uni.lu

ABSTRACT

In recent years, smart contracts have suffered major exploits, costing millions of dollars. Unlike traditional programs, smart contracts are deployed on a blockchain. As such, they cannot be modified once deployed. Though various tools have been proposed to detect vulnerable smart contracts, the majority fails to protect vulnerable contracts that have already been deployed on the blockchain. Only very few solutions have been proposed so far to tackle the issue of post-deployment. However, these solutions suffer from low precision and are not generic enough to prevent any type of attack.

In this work, we introduce ÆGIS, a dynamic analysis tool that protects smart contracts from being exploited during runtime. Its capability of detecting new vulnerabilities can easily be extended through so-called attack patterns. These patterns are written in a domain-specific language that is tailored to the execution model of Ethereum smart contracts. The language enables the description of malicious control and data flows. In addition, we propose a novel mechanism to streamline and speed up the process of managing attack patterns. Patterns are voted upon and stored via a smart contract, thus leveraging the benefits of tamper-resistance and transparency provided by the blockchain. We compare ÆGIS to current state-of-the-art tools and demonstrate that our solution achieves higher precision in detecting attacks. Finally, we perform a large-scale analysis on the first 4.5 million blocks of the Ethereum blockchain, thereby confirming the occurrences of well reported and yet unreported attacks in the wild.

KEYWORDS

Ethereum; Smart contracts; Exploit prevention; Security updates

1 INTRODUCTION

Blockchain has evolved greatly since its first introduction in 2009 [25]. A blockchain is essentially a verifiable, append-only list of records in which all transactions are recorded in batches of so-called blocks. Each block is linked to a previous block via a cryptographic hash. This linked list of blocks is maintained by a decentralised peer-to-peer network. The peers in this network follow a consensus protocol that dictates which peer is allowed to append the next block. By introducing the concept of smart contracts, Ethereum [41] revolutionized the way digital assets are traded. As smart contracts

govern more and more valuable assets, the contracts themselves have come under attack from hackers.

Smart contracts are programs that are stored and executed across blockchain peers. They are deployed and invoked via transactions. Deployed smart contracts are immutable, thus any bugs present during deployment [2], or as a result of changes to the blockchain protocol [7], can make a smart contract vulnerable. Moreover, since contract owners are anonymous, responsible disclosure is usually infeasible or very hard in practice. Though smart contracts can be implemented with upgradeability and destroyability in mind, this is not compulsory. As a matter of fact, Ethereum already faced several devastating attacks on vulnerable smart contracts.

In 2016, an attacker exploited a reentrancy bug in a crowdfunding smart contract known as the DAO. The attacker exploited the capability of recursively calling a payout function contained in the contract. The attacker managed to drain over \$150 million [32] worth of cryptocurrency from the smart contract. The DAO hack was a poignant demonstration of the impact that insecure smart contracts can have. The Ethereum market cap value dropped from over \$1.6 billion before the attack, to values below \$1 billion after the attack, in less than a day. Another example happened with the planned Constantinople hard fork in January 2019. Ethereum was scheduled to receive an update intended to introduce a cheaper gas cost for certain smart contract operations. On the eve of the hard fork, a new reentrancy issue caused by this update was detected. It turned out that the reduction of gas costs also enabled reentrancy attacks on smart contracts that were previously secure. This resulted in the update being delayed [7]. A third example is the Parity wallet hack. In 2017, the Parity wallet smart contract was attacked twice due to a bug in the access control logic. The bug allowed anyone to claim ownership of the smart contract and to take control of all the funds. The first attack resulted in over \$30 million being stolen [44], whereas the second attack resulted in roughly \$155 million being locked forever [29].

The manner in which these issues are currently handled is not ideal. At the moment, whenever a major vulnerability is detected by the Ethereum community, it can take several days or weeks for the community to issue a critical update and even longer for all nodes to adopt this update. Such a delay extends the window for exploitation and can have dire effects on the trading value of the underlying cryptocurrency. Moreover, the lack of a standardised procedure to

deal with vulnerable smart contracts, has led to a “Wild West”-like situation where several self-appointed white hats started attacking smart contracts in order to protect the funds that are at risk from other malicious attackers [4]. However, in some cases the effects of attacks can cause a split in the community so contentious that it leads to a hard fork, such as in the case of the DAO hack which led to the birth of the Ethereum classic blockchain [32].

Academia has proposed a plethora of different tools that allow users to scan smart contracts for vulnerabilities prior to deploying them on the blockchain or interacting with them (see e.g. [21, 23, 36, 38]). However, by design these tools cannot protect vulnerable contracts that have already been deployed. Grossman et al. [12] are the first to present `ECFCHECKER`, a tool that allows to dynamically check executed transactions for reentrancy. However, `ECFCHECKER` does not prevent reentrancy attacks. In order to protect already deployed contracts, Rodler et al. [31] propose `SEREUM`, a modified Ethereum client that detects and reverts¹ transactions that trigger reentrancy attacks. `SEREUM` leverages the principle that every exploit is performed via a transaction. Unfortunately, `SEREUM` has three major drawbacks. First, it requires the client to be modified whenever a new type of vulnerability is found. Second, not only the tool itself, but also any updates to it must be manually adopted by the majority of nodes for its security provisions to become effective. Third, their detection technique can only detect reentrancy attacks, despite there being many other types of attacks [2].

In summary, we make the following contributions:

- We introduce a novel domain-specific language, which enables the description of so-called *attack patterns*. These patterns reflect malicious control and data flows that occur during execution of malicious transactions.
- We present `ÆGIS`, a tool that reverts malicious transactions in real-time using attack patterns, thereby preventing attacks on deployed smart contracts.
- We propose a novel way to quickly propagate security updates without relying on client-side update mechanisms, by making use of a smart contract to store and vote upon new attack patterns. Storing patterns in a smart contract ensures integrity, decentralizes security updates and provides full transparency on the proposed patterns.
- We illustrate the effectiveness by providing patterns to prevent the two most prominent hacks in Ethereum, the DAO and Parity wallet hacks.
- Finally, we provide a detailed comparison to current state-of-the-art runtime detection tools and perform a large-scale analysis on 4.5 million blocks. The results demonstrate that `ÆGIS` achieves better precision than current state-of-the-art tools.

2 BACKGROUND

In this section, we provide the necessary background for understanding the setting of our work. We describe the Ethereum blockchain and its capability of executing smart contracts. We focus on Ethereum since it is currently the most prominent blockchain platform when it comes to smart contract deployment. Finally, we also

provide background information on the two most prominent smart contract vulnerabilities, namely, reentrancy and access control.

2.1 Ethereum and Smart Contracts

Ethereum. The Ethereum blockchain is a decentralized public ledger that is maintained by a network of nodes that distrust one another. Every node runs one of several existing Ethereum clients. The clients can operate with different configurations. For instance, nodes who are configured to mine blocks are called miners. Miners execute transactions, include them in blocks and append them to the blockchain. They compete to create a block by solving a cryptographic puzzle. Once they succeed, the block is proposed to the network. Other miners verify the new block and either accept or reject it. A miner whose block is included in the blockchain is rewarded with a block reward and the execution fees from the included transactions.

Transactions. Transactions are used to modify state in Ethereum. As such, they allow users to transfer ether (Ethereum’s cryptocurrency), and to create smart contracts or trigger their execution. Transactions are created using an account. There are two types of accounts in Ethereum, user accounts and contract accounts. Transactions are given a certain amount of gas to execute, called the gas limit. Gas is a unit which is used to measure the use of computing resources. Gas can be converted to ether through the so-called gas price of a transaction. Gas limit and gas price can be chosen by the creator of the transaction. Together they determine the fee that the user is willing to pay for the inclusion of their transaction into the blockchain. Moreover, transactions also contain a destination address. It identifies the recipient of the transaction, and it can be either a user account or a smart contract. Transactions can also carry value that is transferred to the recipient. Once created, transactions are broadcast to the network. Miners then execute the transactions and include them into blocks. Smart contracts (i.e. contract accounts) are created by leaving the destination address of a transaction empty. The bytecode that is provided within the transaction is then copied into the blockchain and it is given a unique address that identifies the smart contract.

Smart Contracts. Smart contracts are fully-fledged programs that are stored and executed across the blockchain. They are developed using a dedicated high-level programming language that compiles into low-level bytecode. This bytecode gets interpreted by the Ethereum Virtual Machine. Smart contracts contain functions that can be triggered via transactions. The name of the function as well as the data to be executed is included in the data field of the transaction. A default function or so-called fallback function is executed whenever the provided function name is not recognized by the smart contract. Moreover, smart contracts can initiate calls to other smart contracts. Thus, a single transaction may interact with several smart contracts that call one another. By default smart contracts cannot be destroyed or updated. It is the task of the developer to implement these capabilities before deploying the smart contract. Unfortunately, many smart contracts are released without destroyability or upgradeability in mind. As a result, many contracts remain vulnerable or active on the blockchain even past their

¹Consuming gas, without letting the transaction affect the state of the blockchain.

```

1  contract A { // Victim contract
2  ...
3  function withdraw() public {
4      if (credit[msg.sender]) {
5          msg.sender.call.value(credit[msg.sender])();
6          credit[msg.sender] = 0;
7      }
8  }
9
10 contract B { // Exploiting contract
11 ...
12 function () public payable {
13     A.withdraw();
14 }
15 }

```

Figure 1: Example of a reentrancy vulnerability.

utility. As mentioned earlier, once deployed, smart contracts are immutable, they cannot be modified and bugs cannot be fixed. Thus, it is not possible to update a smart contract in the later run.

EVM. The Ethereum Virtual Machine (EVM) is a purely stack-based, register-less virtual machine that supports a Turing-complete instruction set of opcodes. These opcodes allow smart contracts to perform memory operations and interact with the blockchain, such as retrieving specific information (e.g., the current block number). Ethereum makes use of gas to make sure that contracts terminate and to prevent denial-of-service attacks. It assigns a gas cost to the execution of an opcode. The execution of a smart contract results in the modification of its state. The latter is stored on the blockchain and consists of a balance and a storage. The balance represents the amount of ether currently owned by the smart contract. The storage is organized as a key-value store and allows the smart contract to store values and keep state across executions. During execution, the EVM holds a machine state $\mu = (g, pc, m, i, s)$, where g is the gas available, pc is the program counter, m represents the memory contents, i is the active number of words in memory and s is the content of the stack. In summary, the EVM is a transaction-based state machine that updates a smart contract based on transaction input data and the smart contract’s bytecode.

2.2 Smart Contract Vulnerabilities

Although, a number of smart contract vulnerabilities exist [2], in this work, we primarily focus on two types of vulnerabilities that have been defined by the NCC Group as the top two vulnerabilities in their Decentralized Application Security Project [13]: *reentrancy* and *access control*.

Reentrancy Vulnerabilities. Reentrancy occurs whenever a contract calls another contract, which then calls back into the original contract, thereby creating a reentrant call. This is not an issue as long as all the state updates that depend on the call from the original contract are performed before the call. In other words, reentrancy only becomes problematic when a contract updates its state after calling another contract. A malicious contract can take advantage of this by recursively calling a contract until all the funds are drained. Figure 1 provides an example of a malicious reentrancy. Contract

B contains a fallback function (line 12-14), a default function that is automatically executed when no other function is called. In this example, the fallback function of contract B calls the withdraw function of contract A . Assuming that contract B already deposited some ether in contract A , contract A now calls contract B to transfer back its deposited ether. However, the transfer results in calling the fallback function of contract B once again, which results in reentering contract A and once more transferring the value of the deposited ether to contract B . This repeats until the balance of contract A becomes zero or the execution runs out of gas.

Reentrancy vulnerabilities were extensively studied by Rodler et al. [31], and can be divided into four distinct categories: *same-function* reentrancy, *cross-function* reentrancy, *delegated* reentrancy and *create-based* reentrancy. Same-function reentrancy occurs whenever an attacker reenters the original contract via the same function (see Figure 1). Cross-function reentrancy builds on the same-function reentrancy. However, here the attacker takes advantage of another function that shares a state with the original function. Delegated reentrancy and create-based reentrancy are similar to same-function reentrancy, but use different opcodes to initiate the call. Specifically, delegated reentrancy can occur using either the DELEGATECALL or CALLCODE opcodes, while create-based reentrancy only occurs when using the CREATE opcode. While the DELEGATECALL and CALLCODE opcodes behave roughly similar to the CALL opcode, the CREATE opcode causes a new contract to be created and the contract constructor to be executed. This newly created contract can then call and reenter the original contract.

Access Control Vulnerabilities. Access control vulnerabilities result from incorrectly enforced user access control policies in smart contracts. Such vulnerabilities allow attackers to gain access to privileged contract functions that would normally not be available to them. The most famous examples of this type of vulnerability are the two Parity MultiSig-Wallet hacks [29, 44]. The issue originates from the fact that the developers of the Parity wallet decided to split some of the contract logic into a separate smart contract named `WalletLibrary`. This had the advantage of reusing parts of the code for multiple wallets allowing users to save on gas costs during deployment. A simplified version of the code can be seen in Figure 2. As can be seen in line 17-20, the initialisation of the wallet is performed via the `initWallet` function located in contract L , which is called by the constructor of contract W . In addition, any unmatched function calls to contract W are caught by the fallback function in line 6-8, which redirects the call to contract L by means of the DELEGATECALL operation. Unfortunately, in the first version of the Parity MultiSig-Wallet, the developers forgot to write a safety check for the `initWallet` function, ensuring that the function can only be called once. As a result an attacker was able to gain ownership of contract W by calling the `initWallet` function via the fallback function. Once in control the attacker withdrew all the funds by invoking the `execute` function (line 32-34).

After the first Parity hack, a new Parity MultiSig-Wallet Library contract was deployed addressing the issue above. In the newly deployed version, the `initWallet` function was not part of the constructor anymore, but had to be called separately after deployment. However, the developers did not call the `initWallet` function after deployment. Hence, contract L remained uninitialised, meaning that

```

1  contract W { // Wallet contract
2  ...
3  function W(address _owner) { // Constructor
4      L.delegatecall("initWallet(address)", _owner);
5  }
6  function () payable {
7      L.delegatecall(msg.data);
8  }
9  }
10
11 contract L { // Library contract
12 ...
13 modifier onlyOwner {
14     if (m_ownerIndex[msg.sender] > 0) _;
15 }
16 ...
17 function initWallet(address[] _owners, uint
18     _required, uint _daylimit) {
19     initDaylimit(_daylimit);
20     initMultiowned(_owners, _required);
21 }
22 function initMultiowned(address[] _owners, uint
23     _required) {
24     ...
25     for (uint i = 0; i < _owners.length; ++i) {
26         ...
27         m_ownerIndex[_owners[i]] = 2+i;
28     }
29     ...
30 }
31 function execute(address _to, uint _value, bytes
32     _data) onlyOwner {
33     _to.call.value(_value)(_data);
34 }
35 function kill(address _to) onlyOwner {
36     suicide(_to);
37 }
38 }

```

Figure 2: Example of an access control vulnerability.

the library contract itself had no owners. As a result, 3 months after deployment a user known as *devops199* was experimenting with the previous Parity hack vulnerability and called the `initWallet` function directly inside contract `L`, marking its address as the owner. The user then called the `kill` function (line 32-34), which removed the executable code of contract `L` from the blockchain² and sent the remaining funds to the new owner. The contract itself contained no funds, however it was used by multiple Parity wallets which had the address of contract `L` defined as a constant in their executable code. As a result any wallet trying to use contract `L` as a library would now receive zero as return value, effectively rendering the wallet unusable and therefore freezing the funds contained in the wallets. This led the user to publicly disclose the steps that led to this tragedy, with the words: “I accidentally killed it.” [9].

3 RELATED WORK

In this section, we discuss some of the works that are most closely related to ours.

²The contract code is technically not removed from the blockchain, however, the contract’s code can no longer be executed on the blockchain, because the contract has been marked as killed.

Security Analysis of Smart Contracts. As with any program, smart contracts may contain bugs and can be vulnerable to exploitation. As discussed in [2], different types of vulnerabilities exist, often leading to financial losses. The issue is made worse by the fact that smart contracts are immutable. Once deployed, they cannot be altered and vulnerabilities cannot be fixed. In addition to that, automated tools for launching attacks exist [21].

Several defense mechanisms have been proposed to detect security vulnerabilities in smart contracts. This includes tools such as ERAYS [46], designed to provide smart contract auditors with a reverse engineered pseudo code of a contract from its bytecode. The interpretation of the pseudo code however remains a slow and gruelling task. More automated tools have also been proposed benefiting from regular expressions [43] and machine learning techniques [34] to detect vulnerabilities.

A wealth of security research has focused on the creation of static analysis tools to automatically detect vulnerabilities in smart contracts. Formal verification has been used together with a formal definition of the EVM [1, 16], or by first converting smart contracts into the formal language F* [5, 11]. Other works focused on analysing the higher level solidity code [10, 35], which limits the scope to those contracts with available source code. Another approach is to apply static analysis on the smart contract bytecode [38]. A technique commonly used for this purpose is symbolic execution, designed to thoroughly explore the state space of a smart contract utilising constraint solving. It has been used to detect contracts with vulnerabilities [23, 28], to find misbehaving contracts [20, 26, 37], or detect integer bugs [19, 36]. Fuzzing techniques have also been applied [15, 18]. In [42] the authors propose HARVEY, a greybox fuzzer that selects appropriate inputs and transaction sequences to increase code coverage. Fuzzing techniques however involve a trade-off between the number of discovered paths and the efficiency in input generation.

While all the listed tools help identify vulnerabilities, they cannot protect already deployed smart contracts from being exploited. Therefore, to deal with the issue of vulnerabilities in deployed smart contracts, [12, 31] propose a modification to the Ethereum client, that would allow detection and prevent exploitation of reentrancy vulnerabilities at runtime. However, these approaches only deal with reentrancy and require all the clients in the network to be modified. This is an issue for the following reasons. On one hand, every update of the vulnerability detection software requires an update of the different Ethereum client implementations. This is true for both bug fixes and functionality upgrades, for example the detection of new vulnerabilities. On the other hand, every modification of the clients needs to be adopted by all the nodes participating in the Ethereum blockchain. This requires time and breaks compatibility between updated and non-updated clients. In this work, we propose a generic solution that protects contracts and users from existing and future vulnerabilities, without requiring client modifications and forks every time a new vulnerable smart contract is found.

Wang et al. [40] propose an approach to detect vulnerabilities at runtime based on two invariants that follow the intuition that most vulnerabilities are due to a mismatch between the transferred amount and the amount reflected by the contract’s internal book-keeping logic. However, this approach has three main drawbacks.

First, it requires the automated and correct identification of bookkeeping variables, which besides being a non-trivial task also does not hold for every contract, since there can be contracts that do not use internal bookkeeping logic but are nevertheless vulnerable. Second, their approach does not model environmental information such as timestamps or block numbers, which does not allow them to detect vulnerabilities such as timestamp dependence or transaction order dependency, whereas our approach models environmental information and allows for the detection of these vulnerabilities. Finally, Wang et al.’s approach can only detect violations of safety properties and not violations of liveness properties such as the Parity Wallet Hack 2. In this work, we demonstrate that our approach is capable of detecting both Parity wallet hacks and therefore violations to safety as well as liveness properties.

Blockchain-Based Voting. Since blockchains provide the means for transparency and decentralization, multiple blockchain-based solutions have been proposed for performing electronic voting [3, 17, 27]. Interestingly, with the recent developments in quantum computers, recent work also has started to focus on the development of quantum-resistant blockchain-based voting schemes [33]. These solutions can all be categorised into two categories: cryptocurrency-based and smart-contract-based.

Cryptocurrency-based solutions focus on using payments as a proxy for votes in an election. When a voter wishes to cast a vote, he or she makes a payment to the address of the candidate. Lee et al. [22] proposed such a system in the Bitcoin network. However, their system requires a trusted third party to perform the ballot counting. Zao et al. [45] were the first to propose a voting scheme using the public Bitcoin network while preserving the privacy of the votes. Another well-known cryptocurrency-based solution is CarbonVote [8]. It was introduced in the aftermath of the DAO hack to allow the Ethereum Foundation to determine if the Ethereum community wanted a hard fork or not. The tallying was performed by counting the amount of ether that each address received. Needless to say, such a system gives a tremendous amount of voting power to users with a large amount of funds.

Smart-contract-based voting relies on a decentralized application to assist the voting process – there is no central entity. McCorry et al. [24] propose a practical implementation of the Open Vote Network [14] in the form of a smart contract deployed on the Ethereum blockchain for boardroom voting. Their implementation is self-tallying and provides, in addition to vote privacy, also transparency. Voting proceeds in several rounds, where the voters first broadcast their voting key, followed by a proof that their vote is binary (a “yes” or “no” vote). A final tally round allows anyone to calculate the total sum of votes, without revealing individual ballots. The voting mechanism described in this paper is inspired by McCorry et al.’s proposed solution and implementation. The limitations of their proposed solution, namely having a binary voting system and limiting the number of voters to less than 50 participants, are acceptable for our purposes.

4 METHODOLOGY

In this section, we present the details of our solution towards a generic and decentralized way to prevent any type of attacks

on already deployed smart contracts. Our idea is to bundle every Ethereum client with a runtime analysis tool, that interacts with the EVM and is capable of interpreting so-called *attack patterns*, and reverting transactions that match these patterns. Attack patterns are described using our domain-specific language (DSL), which is tailored to the execution model of the EVM and which allows to easily describe malicious control and data flows. The fact that we shift the capability of detecting attacks from the client-side implementation to the DSL, gives us the advantage of being able to quickly propose mitigations against new vulnerabilities, without having to modify the Ethereum client. Existing approaches, such as SEREUM for example, require the client-side implementation to be modified whenever a new vulnerability is found.

4.1 Generic Attack Detection

Attacks are detected in our system through the use of patterns, which are described using our DSL. The DSL allows for the definition of malicious events that occur during the execution of EVM instructions. The syntax of our DSL is defined by the following BNF grammar:

```

<instr> ::= CALL | CALLDATALOAD | SSTORE | JUMPI | ...
<exec> ::= depth | pc | address | stack(int) | stack.result |
         | memory(int, int) | transaction.<trans> |
         | block.<block>
<trans> ::= hash | value | from | to | ...
<block> ::= number | gasUsed | gasLimit | ...
<comp> ::= < | > | ≤ | ≥ | = | ≠ | + | - | · | /
<expr> ::= (src.<exec> <comp> <expr>) [∧ <expr>]
         | (<expr> <comp> dst.<exec>) [∧ <expr>]
         | (src.<exec> <comp> src.<exec>) [∧ <expr>]
         | (src.<exec> <comp> dst.<exec>) [∧ <expr>]
         | (dst.<exec> <comp> dst.<exec>) [∧ <expr>]
         | (src.<exec> <comp> int) | (dst.<exec> <comp> int)
<rel> ::= ⇒ | ∼ | →
<pattern> ::= (opcode = <instr>) <rel> (opcode = <instr>) [where
         <expr>]
         | <pattern> <rel> (opcode = <instr>) [where <expr>]
         | (opcode = <instr>) <rel> <pattern> [where <expr>]

```

Figure 3: DSL for describing attack patterns.

A pattern is a sequence of relations between EVM instructions that may occur at runtime. We distinguish between three types of relations, a “control flow” relation (\Rightarrow), a “data flow” relation (\sim), and a “follows” relation (\rightarrow). A control flow relation means that an instruction is control dependent on another instruction. A data-flow relation means that an instruction is data dependent on another instruction. A follows relation means that an instruction is executed after another instruction, without necessarily being control or data dependent on the other instruction. A relation is always between two EVM opcodes: a source opcode (**src**) and a destination opcode

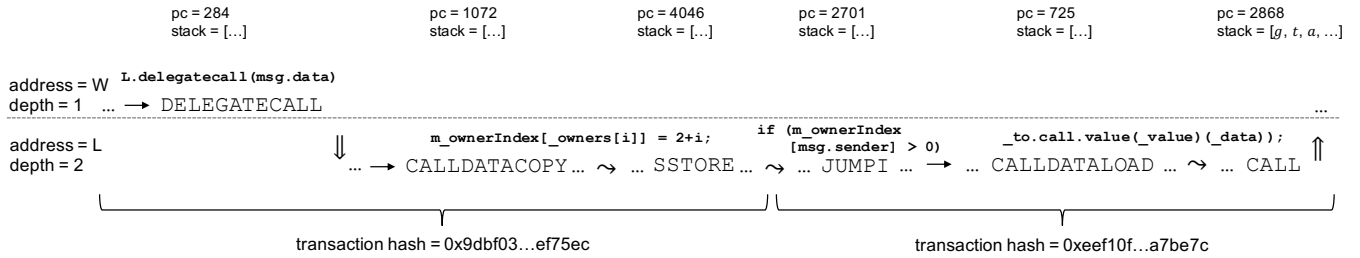


Figure 5: Execution example of an attack on an access control vulnerability. A data flow relation is depicted with \rightsquigarrow . The variables g , t and a are as discussed in Figure 4.

as part of the attacker calling the fallback function of contract W . In the first transaction, we see that at a certain point contract L copies data from the transaction using the `CALLDATACOPY` instruction and stores it into storage via the `SSTORE` instruction. An interesting observation here is that state is shared across transactions through storage. In the second transaction, the data that has previously been stored is now loaded onto the stack and used by a comparison. A comparison is ultimately reflected via the `JUMPI` instruction. Finally, we see that the comparison follows a `CALLDATALOAD` instruction whose data is used by a call `CALL` instruction. Given these observations, we are now able to create the following attack pattern in order to detect and thereby prevent the first Parity wallet hack:

```
(opcode = DELEGATECALL) => (opcode = CALLDATACOPY) ~>
(opcode = SSTORE) ~> (opcode = JUMPI) where
  (src.transaction.hash != dst.transaction.hash) ->
((opcode = CALLDATALOAD) ~> (opcode = CALL)) where
  (dst.stack(2) > 0)
```

The above attack pattern evaluates to true if the following two conditions are met:

- (1) there is a transaction with a control flow relation between a `DELEGATECALL` instruction and a `CALLDATACOPY` instruction, where the data of the `CALLDATACOPY` instruction flows into an `SSTORE` instruction;
- (2) there is another transaction (i.e. `src.transaction.hash != dst.transaction.hash`) where the data that has been previously stored via the `SSTORE` instruction flows into a `JUMPI` instruction and is followed by a `CALLDATALOAD` instruction whose data flows into a `CALL` instruction that sends out funds (i.e. `dst.stack(2) > 0`).

It is worth noting that the Parity wallet attack is a multi-transactional attack and that it is therefore significantly different from a reentrancy attack, that is solely based on a single transaction. For more examples of attack patterns, please refer to Table 5 in Appendix A.

4.2 Decentralized Security Updates

While our approach of using a DSL allows us to have a generic solution for detecting attacks, it still leaves two open questions:

- (1) How do we distribute and enforce the same patterns across all the clients?

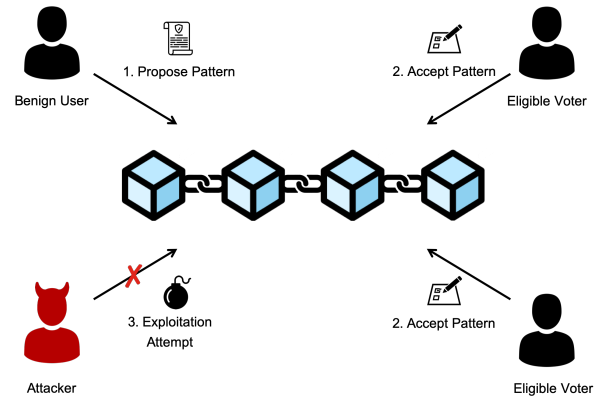


Figure 6: An illustrative example of AEGIS's workflow: step 1) A benign user detects a vulnerability and proposes a pattern (written using our DSL) to the smart contract. Step 2) Eligible voters vote to either accept or reject the pattern. If the majority votes to accept the pattern, then all the clients are updated and the pattern is activated. Step 3) An attacker tries but fails to exploit a vulnerable smart contract due to the voted pattern matching the malicious transaction.

- (2) How do we decentralize the governance of patterns in order to prevent a single entity from deciding which patterns are added or removed?

The answer to these questions is to use a smart contract that is deployed on the blockchain itself. This solves the problem of distributing and enforcing that the same patterns are always used across all clients. Specifically, patterns are stored inside the smart contract and the blockchain protocol itself guarantees that every client knows about the exact same state and therefore has access to exactly the same patterns. The second problem of decentralizing the governance of patterns, is solved by allowing the proposal and voting of patterns via the smart contract as depicted in Figure 6. The contract maintains a list of eligible voters that vote for either accepting or rejecting a new pattern. If the majority has voted with "yes", i.e. to accept the pattern, then it is added to the list of active patterns. In that case, every client is automatically notified through the mechanism of smart contract events, and retrieves the updated list of patterns from the blockchain. In other words, if a pattern is accepted by the voting mechanism, it is updated across all the




Figure 7: Architecture of ÆGIS. The dark gray boxes represent ÆGIS’s main components.

clients through the existing consensus mechanism of the Ethereum blockchain. However, solving the second problem using a voting mechanism opens up a new problem concerning the requirements needed for governing the votes. In voting literature, verifiability and privacy are typically seen as key requirements. *Verifiability* concerns linking the output to the input in a verifiable way. *Privacy* concerns whether a vote can be linked back to a voter. In addition, we argue that the situation here is more akin to boardroom voting than to general elections, because it should be possible to hold voters *accountable*. This means that privacy must be maintained only until the election is over. Finally, the voting system must not be favorable to any voters – e.g., it should not confer an advantage to voters that cast their vote late. This final property is called *fairness*. It is worth noting that fairness requires privacy during the voting phase. This leads to the following three requirements:

- (1) **Verifiability:** The outcome of the vote must be verifiably related to the votes as cast by the voters;
- (2) **Accountability:** Voters can be held accountable for how they voted;
- (3) **Fairness:** No intermediate information must be leaked.

5 IMPLEMENTATION

In this section, we provide the implementation details of our solution called ÆGIS. The code is publicly available³. Figure 7, provides an overview of the architecture of ÆGIS and highlights its main components. ÆGIS is implemented on top of Trinity⁴, an Ethereum client implemented in Python.

5.1 Ethereum Client

EVM. We modified the EVM of Trinity such that it keeps track of all the executed instructions and their states at runtime, in the form of an ordered list. We refer to this list as the execution trace. Each record in this list contains the executed opcode, the value of the program counter, the depth of the call stack, the address of the contract that is being executed, and finally, all the values that were stored on the stack and in memory. This list is passed to the interpreter component of ÆGIS.

Interpreter. The interpreter loops through the list of executed instructions and passes the relevant instructions to the control flow and data flow extractor components. It is also responsible for signalling the EVM a revert in case the execution trace matches an attack pattern.

Control Flow Extractor. The control flow extractor is responsible for inferring control flow information. We do so by dynamically building a call tree from the instructions received by the interpreter. A control flow relation is reported if there exists a path along the call tree, from the source instruction to the destination instruction defined in a given pattern. Thus, control flow relations represent call dependencies between two instructions.

Data Flow Extractor. The data flow extractor is responsible for collecting data flow information. We track the flow of data between instructions by using dynamic taint analysis. Taint is introduced whenever we come across a source instruction and checked whenever we come across a destination instruction. Source and destination instructions are defined by a given pattern. Taint propagation follows the semantics of the EVM [41] across stack, memory and storage. We perform byte-level precision tainting. Taint that is stored across stack and memory is volatile, meaning that it is cleared across transactions. Taint that is stored across storage is persistent, meaning that it remains in storage across transactions. This allows us to perform inter-transactional taint analysis. A data flow relation is given if taint flows from a source instruction into a destination instruction.

Pattern Parser. The pattern parser is responsible for extracting and parsing the patterns from the voting smart contract. We implemented our pattern language using textX⁵, a Python framework providing a meta-language for building DSLs.

5.2 ÆGIS Smart Contract

The ÆGIS smart contract ensures proper curation of the list of active patterns. We implemented our smart contract in Solidity. As previously mentioned, patterns are accepted or removed via a voting mechanism. The contract holds all proposed additions and removals of patterns and allows a vote on them within a set time window. The duration can be configured and updated by the contract owner. Proposals are open to the public and anyone can propose an addition to or removal from the list of patterns.

Fairness. Votes should remain secret until all eligible voters have had sufficient opportunity to vote. Therefore, two time windows are employed. The first window is for sending a commitment that includes a deposit. The second window is for revealing a vote including the return of the committed deposits. The two windows are illustrated in Figure 8. In the figure, t_p represents the point in time when a pattern is proposed and marks the start of the commit window. t_c marks the end of the commit window and the start of the reveal window. Lastly, t_r marks the end of the reveal window and the time when the pattern list is updated in case of a positive vote outcome. A commitment is a hash of the vote ID, the voter’s

³<https://github.com/christofortres/Aegis>

⁴<https://trinity.ethereum.org/>

⁵<https://github.com/textX/textX>

vote and a nonce. The vote ID is a hash of the proposed pattern and identifies the pattern that is being voted on. The voter’s vote is encoded as a boolean. The nonce ensures that commitments cannot be replayed. The smart contract records these commitments, which must be sent with the predefined deposit and within the predefined time window. During the commitment phase no one knows how anyone else has voted on a given pattern, and so cannot be swayed by the decisions of others. However, the process should ultimately be transparent to both voters and non-voters to foster trust in the system. As such, during the second window, the reveal window, all voters reveal how they have voted. They must reveal their vote in order to get their deposit back. No commits may be made once the reveal period has started.




Figure 8: Timeline of the two voting stages.

Tallying. The voting ends either when more than 50% (50%+1 vote) of the total number of votes reaches either accept or reject, or when the time window for revealing expires with less than 50% having been reached. In case the voting has ended but the reveal window has not yet passed, any remaining voters are still eligible to reveal their vote, such that their deposit can be returned. The reveal period is bounded so that patterns are accepted or rejected in a practical amount of time. In the event of a successful vote, the pattern to which the vote pertains is added to or removed from the record held by the contract, according to the proposal. If a vote is unsuccessful, i.e. no majority voted for the proposal, the record of patterns is not changed.

Actors. There are three types of actors: the proposers that submit proposals to add or remove patterns, the voters that vote on proposals, and the admins that govern the list of eligible voters as well as the parameters of the smart contract (e.g. deposit, commit and reveal windows, etc.). The \mathcal{AEGIS} smart contract allows every user on the blockchain to become a proposer by submitting a proposal. Voters then vote on the proposals by first committing their vote and at a later stage revealing it. Not every user is an eligible voter. Voters are only those users whose account address is stored in the list of eligible voters maintained by the smart contract. Admins may update the list of eligible voters. They oversee the proper curation of the smart contract and act as a governing body. Admins are agreed upon off-chain and are represented by a multi-signature wallet. A multi-signature wallet is an account address which only performs actions if a group of users give their consent in form of a signature.

Data Structures. The smart contract consists of several functions and data structures that allow for the voting process to take place. We make use of a number of *modifiers*, which act as checks carried out before specific functions are executed. We use these to check that: 1) a voter is eligible, 2) a vote is in progress, 3) a reveal is

in progress and 4) the associated vote has ended. We use a struct to hold the details of each vote, these include the `patternID`, the proposed pattern and the `startBlock`. These values enable us to record the details needed to check when a vote ends, check that the same pattern has not already been proposed, and count the number of votes. The struct is used in conjunction with a mapping, which maps a 32 bytes value to the details of each vote. The 32 bytes value represents the `voteID` of each vote, created by hashing unique vote information. A constructor is used to define, at contract launch, the value of the necessary deposit and the time windows during which voters can commit or reveal. The former is given in ether, while the latter are given in number of blocks. The deposit is used to ensure that those who committed a vote also reveal their vote. These values can be changed later using the contract’s admin functions.

Functionality. The public functions for the voting process are: `addProposal`, `removeProposal`, `commitToVote` and `revealVote`. Both proposal functions first check if a vote with the same ID already exists, and if not create a new instance of voting details via the mapping. Next, the `commitToVote` function can be used inside the defined number of blocks to submit a unique hash of an eligible voter’s vote. This function makes use of the `canVote` modifier to protect access. The voter’s commitment and vote hash are stored only if the correct deposit amount was sent to the function. Once the vote stage has ended the reveal stage begins. During this window the `revealVote` function, protected by the `canVote` modifier, processes vote revelations and returns deposits. The function checks that the stored hash matches the hash calculated from the parameters passed to it, and if so, returns the voter’s deposit and records the vote. Lastly, it calls an internal function which tallies the votes and adds or removes the pattern if either the for or against vote has reached over 50%. In this way the vote is self tallying. The patterns are ultimately stored in an array that can be iterated over to ensure each node has the full set. Finally, the contract also has two admin functions: `transferOwnership`, `changeVotingWindows`. Both of these are protected by the `isOwner` modifier. The former allows the current owning address to transfer control of the contract to a new address. The latter allows the commit and reveal windows to be changed as well as the amount required as a voting deposit.

6 EVALUATION

In this section, we evaluate the effectiveness and correctness of \mathcal{AEGIS} , by conducting two experiments. In the first experiment we compare the effectiveness of \mathcal{AEGIS} to two state-of-the-art reentrancy detection tools: `ECFCHECKER` [12] and `SERUM` [31]. In the second experiment we perform a large-scale analysis and measure the correctness as well as the performance of \mathcal{AEGIS} across the first 4.5 million blocks of the Ethereum blockchain.

6.1 Comparison to Reentrancy Detection Tools

By analyzing transactions sent to contracts, Rodler et al’s tool `SERUM` flagged 16 contracts as victims of reentrancy attacks. However, after manual investigation the authors found that only 2 out of the 16 contracts have actually become victims to reentrancy attacks. We decided to analyze these 16 contracts and see if we face the same

	CCRB	DAO	0x7484a1	proxyCC	DAC	DSEthToken	0x695d73	EZC	0x98D8A6	WEI	0xbD7CeC	0xF4ee93	Alarm	0x771500	KissBTC	LotteryGameLogic
SEReUM	FP	TP	FP	FP	FP	TP	FP	FP	FP	FP	FP	FP	FP	FP	FP	FP
ÆGIS	TN	TP	TN	TN	TN	TP	TN	TN	TN	TN	TN	TN	TN	TN	TN	TN

Table 1: Comparison between SEReUM and ÆGIS on the effectiveness of detecting reentrancy attacks.

Smart Contract	Reentrancy Type	ECFCHECKER	SEReUM	ÆGIS
<i>VulnBankNoLock</i>	Same-Function	TP	TP	TP
	Cross-Function	FN	TP	TP
<i>VulnBankBuggyLock</i>	Same-Function	TN	FP	TN
	Cross-Function	FN	TP	TP
<i>VulnBankSecureLock</i>	Same-Function	TN	FP	TN
	Cross-Function	TN	FP	TN

Table 2: Comparison between ECFCHECKER, SEReUM and ÆGIS on the effectiveness of detecting same-function and cross-function reentrancy attacks with manually introduced locks.

challenges in classifying these contracts correctly. We contacted the authors of SEReUM and obtained the list of contract addresses. Afterwards, we ran ÆGIS on all transactions related to the contract addresses, up to block number 4,500,000⁶. Table 1 summarizes our results and provides a comparison to the results obtained by SEReUM. From Table 1, we can observe that ÆGIS successfully detects transactions related to the DAO contract and the DSEthToken contract, as reentrancy attacks. Moreover, ÆGIS correctly flags the remaining 14 contracts as not vulnerable. Hence, in contrast to SEReUM, ÆGIS produces no false positives on these 16 contracts. After analyzing the false positives produced by SEReUM, we conclude that ÆGIS does not produce the same false positives because first, ÆGIS does not use taint analysis in its pattern and therefore does not face issues of over-tainting, and secondly, it does not make use of dynamic write locks to detect reentrancy.

6.1.1 Reentrancy with Locks. Besides evaluating SEReUM on the set of 16 real-world smart contracts, Rodler et al. also compared SEReUM to ECFCHECKER, using self-crafted smart contracts as a benchmark [30]. The goal of this benchmark is to provide means to investigate the quality of reentrancy detection tools. The benchmark consists of three functionally equivalent contracts, except that the first contract does not employ any locking mechanism

to guard the reentry of functions (*VulnBankNoLock*), the second contract employs partial implementation of a locking mechanism (*VulnBankBuggyLock*), and the third contract employs a full implementation of a locking mechanism (*VulnBankSecureLock*). As a result, the first contract is vulnerable to same-function reentrancy as well as cross-function reentrancy. The second contract is vulnerable to cross-function reentrancy, but not to same-function reentrancy. Finally, the third contract is safe regarding both types of reentrancy. We deployed these three contracts on the Ethereum test network called Ropsten and ran the three contracts against ÆGIS. Table 2 contains our results and compares ÆGIS to ECFCHECKER and SEReUM. We can see that ECFCHECKER has difficulties in detecting cross-function reentrancy, whereas SEReUM has difficulties in distinguishing between reentrancy and manually introduced locks. This is probably due to the locking mechanism exhibiting exactly the same pattern as a reentrancy attack and SEReUM being unable to differentiate between these two. We found that ÆGIS correctly classifies every contract as either vulnerable or not vulnerable in all the test cases.

6.1.2 Unconditional Reentrancy. Calls that send ether are usually protected by a check in the form of an if, require, or assert. Reentrancy attacks typically try to bypass these checks. However, it is possible to write a contract, which does not perform any check before sending ether. Rodler et al. present an example of such a vulnerability and name it *unconditional reentrancy* (see Appendix B). Moreover, they also find an example of such a contract deployed on the Ethereum blockchain⁷. When SEReUM was published, it was not able to detect this type of reentrancy since the authors assumed that every call that may lead to a reentrancy is guarded by a condition. However, the authors claim to have fixed this issue by extending SEReUM to tracking data flows from storage to the parameters of calls. We cannot verify this since the source code of SEReUM is not publicly available. We run ÆGIS on both examples, the manually crafted example by Rodler et al. and the contract deployed on the Ethereum blockchain. ÆGIS correctly identifies the unconditional reentrancy contained in both examples without modifying the existing patterns. This is as expected, since in contrast to SEReUM’s initial way to detect reentrancy, ÆGIS’s reentrancy patterns do not rely on the detection of conditions (i.e. JUMPI) to detect reentrancy.

⁶This is the maximum block number analyzed by the authors of SEReUM.

⁷<https://etherscan.io/address/0xb7c5c5aa4d42967efc906e1b66cb8df9ceb04f7>

Vulnerability	Contracts	Transactions
Same-Function Reentrancy	7	822
Cross-Function Reentrancy	5	695
Delegated Reentrancy	0	0
Create-Based Reentrancy	0	0
Parity Wallet Hack 1	3	80
Parity Wallet Hack 2	236	236
Total Unique	248	1118

Table 3: Number of vulnerable contracts detected by ÆGIS.

6.2 Large-Scale Blockchain Analysis

In this experiment we analyse the first 4.5 million blocks of the Ethereum blockchain and compare our findings to those of Rodler et al. We started by scanning the Ethereum blockchain for smart contracts that have been deployed until block 4,500,000. We found 675,444 successfully deployed contracts. The deployment timestamps of the found contracts range from August 7, 2015 to November 6, 2017. Next, we replayed the execution history of these 675,444 contracts. As part of the scanning we found that only 12 contracts in our dataset have more than 10,000 transactions. Therefore, to reduce the execution time, we decided to limit our analysis to the first 10,000 transactions of each contract. In addition, similar to Rodler et al., we tried our best to skip those transactions which were involved in denial-of-service attacks as they would result in high execution times⁸.

We ran ÆGIS on our set of 675,444 contracts using a 6-core Intel Core i7-8700 CPU @ 3.20GHz and 64 GB RAM. Our tool took on average 108 milliseconds to analyse a transaction, with a median of 24 milliseconds per transaction. All in all, we re-executed 4,960,424 transactions with an average of 8 transactions per contract. Table 3 summarizes our results. ÆGIS found a total of 1,118 malicious transactions and 248 unique contacts that have been exploited through either a reentrancy or an access control vulnerability. More specifically, ÆGIS found that 7 contracts have become victim to same-function reentrancy, 5 contracts to cross-function reentrancy, 3 contracts to the first Parity wallet hack and 236 contracts to the second Parity wallet hack. Similar to the results of Rodler et al., we did not find any contracts to have become victim to delegated reentrancy or create-based reentrancy. We validated all our results by manually analyzing the source code (whenever it was publicly available) and/or the execution traces of the flagged contracts. Our validation did not reveal any false positives.

Table 4 lists all the contract addresses that ÆGIS detected to have become victim of a same-function reentrancy attack. The block range defines the block heights where ÆGIS detected the malicious transactions. The first and second contract addresses contained in Table 4 are the same as reported by Sereum, and belong to the DSEthToken and DAO contract, respectively. The rows highlighted in gray mark 5 contracts that have been flagged by ÆGIS but not by Sereum. After investigating the transactions of these 5 contracts, we find that the contract addresses 0x26b8af052895080148dabb1007

⁸<https://tinyurl.com/rv1vues>

b3045f023916e and 0xbf78025535c98f4c605f9eaf672999abf19dc1 became victim to same-function reentrancy, but seem to be contracts that have been deployed with the purpose of studying the DAO hack. However, the three other contract addresses seem to be true victims of reentrancy attacks.

7 DISCUSSION

In this section we discuss alternatives to determine eligible voters, highlight some of the current limitations as well as future research directions for this work.

7.1 Determining Eligible Voters

The introduction of new patterns in ÆGIS depends on achieving consensus in a predetermined group of voters. Although it may intuitively make sense to let miners vote, they are not necessarily a good fit. Their interests may differ from those of smart contract users. For example, depending on a pattern’s complexity, it might introduce an overhead in terms of execution time. Miners are then incentivized to prefer simpler patterns that are evaluated quicker, while smart contract users would prefer more secure patterns.

Alternatively, a group of trusted security experts could act as eligible voters⁹. Security experts are (by definition) able to properly evaluate patterns and have the interest in doing so. The voting contract is then controlled by a group of trusted experts who are decided upon off-chain by a group of admins. For transparency, the identity of admins and experts would be exposed to the public by mapping every identity to an Ethereum account. Changes to the list of voters, the deposit, or the commit and reveal windows are then visible to anyone via the blockchain. Through this setup, security experts would be able to organise themselves with the voter list being comprised of a curated group of knowledgeable people. Such groups already exist in reality, for example, the members of the Smart Contract Weakness Classification registry (SWC)¹⁰, and would be a good fit for our system. Moreover, misbehaving or unresponsive experts could be easily removed by the group of admins. Although this approach allows for scalability and control, it has the disadvantage of introducing managing third-parties. That runs counter to the decentralised concept of Ethereum.

⁹Somewhat similar to how CVEs are handled.

¹⁰<https://smartcontractsecurity.github.io/SWC-registry/>

Contract Address	Block Range
0xd654bdd32fc99471455e86c2e7f7d7b6437e9179	1680024 - 1680238
0xbb9bc244d798123fde783fcc1c72d3bb8c189413	1718497 - 2106624
0xf01fe1a15673a5209c94121c45e2121fe2903416	1743596 - 1743673
0x304a554a310c7e546dfe434669c62820b7d83490	1881284 - 1881284
0x59752433dbe28f5aa59b479958689d353b3dee08	3160801 - 3160801
0xbf78025535c98f4c605f9eaf672999abf19dc1	3694969 - 3695510
0x26b8af052895080148dabb1007b3045f023916e	4108700 - 4108700

Table 4: Same-function reentrancy vulnerable contracts detected by ÆGIS. Contracts highlighted in gray have only been detected by ÆGIS and not by Sereum.

Alternatively, there is also an option to select voters, while preserving the decentralised concept of Ethereum. This is to remove the role of admins altogether, and instead follow a self-organizing strategy, similar to Proof-of-Stake. In this case, everyone is allowed to become a voter through the purchase of (not prohibitively priced) voting power. This could be achieved by depositing a fixed amount of ether into the voting smart contract as a form of collateral.

7.2 Adoption and Participation Incentives

The deployment of ÆGIS would require a modification of the Ethereum consensus protocol, which would require existing Ethereum clients to be updated. This could be easily achieved through a major release by including this one-time modification as part of a scheduled hardfork. Another issue concerns the incentives to propose and vote on patterns. While prestige or a feeling of contributing to the security of Ethereum may be sufficient for some, more incentives may be needed to ensure that the protective capabilities of ÆGIS are used to the full extent. A monetary incentive could address this. That is, ÆGIS could be extended with automatically paid rewards. In other words, ÆGIS could be extended to enable bug bounties [6]. ÆGIS's smart contract could be modified such that, owners of smart contracts can register their contract address by sending a transaction to ÆGIS's voting smart contract and deposit a bounty in the form of ether. Then, proposers of patterns would be rewarded automatically with the bounty by ÆGIS's voting smart contract, if their proposed pattern is accepted by the group of voters. Moreover, owners could simply replenish the bounty for their contract by making new deposits to ÆGIS's smart contract.

7.3 Limitations and Future Work

A current limitation of our tool is that proposed attack patterns are submitted in plain text to the smart contract. Potential attackers can view the patterns and use them to find vulnerable smart contracts. To mitigate this, we propose to make use of encryption such that only the voters would be able to view the patterns. However, this would break the current capability of the smart contract being self-tallying. Designing an encrypted and practical self-tallying solution is left for future work. Finally, we intend to make use of parallel execution inside the extractors and the checking of patterns in order to improve the time required to analyse transactions.

8 CONCLUSION

Although academia proposed a number of tools to detect vulnerabilities in smart contracts, they all fail to protect already deployed vulnerable smart contracts. One of the proposed solutions is to modify the Ethereum clients in order to detect and revert transactions that try to exploit vulnerable smart contracts. However, these solutions require all the Ethereum clients to be modified every time a new type of vulnerability is discovered. In this work, we introduced ÆGIS, a system that detects and reverts attacks via attack patterns. These patterns describe malicious control and data flows through the use of a novel domain-specific language. In addition, we presented a novel mechanism for security updates that allows these attack patterns to be updated quickly and transparently via the blockchain, by using a smart contract as means of storing them. Finally, we compared ÆGIS to two current state-of-the-art online

reentrancy detection tools. Our results show that ÆGIS not only detects more attacks, but also has no false positives as compared to current state-of-the-art.

ACKNOWLEDGMENTS

We would like to thank the SEREUM authors, especially Michael Rodler, for sharing their data with us. We would also like to thank the reviewers for their valuable comments as well as Daniel Xiapu Luo for his valuable help. The experiments presented in this paper were carried out using the HPC facilities of the University of Luxembourg [39] – see <https://hpc.uni.lu>. This work is partly supported by the Luxembourg National Research Fund (FNR) under grant 13192291.

REFERENCES

- [1] Sidney Amani, Myriam Bégel, Maksym Bortin, and Mark Staples. 2018. Towards verifying ethereum smart contract bytecode in Isabelle/HOL. In *Proc. 7th ACM SIGPLAN International Conference on Certified Programs and Proofs (CPP'18)*. ACM, 66–77. <https://doi.org/10.1145/3167084>
- [2] Nicola Atzei, Massimo Bartoletti, and Tiziana Cimoli. 2017. A Survey of Attacks on Ethereum Smart Contracts (SoK). In *Proc. 6th International Conference on Principles of Security and Trust - Volume 10204*. Lecture Notes in Computer Science, Vol. 10204. Springer-Verlag, 164–186.
- [3] Ahmed Ben Ayed. 2017. A conceptual secure blockchain-based electronic voting system. *International Journal of Network Security & Its Applications* 9, 3 (2017), 01–09.
- [4] Jordi Baylina. 2019. Verification of the balances rescued from the multisig compromise. <https://github.com/Giveth/WHGBalanceVerification>.
- [5] Karthikeyan Bhargavan, Antoine Delignat-Lavaud, Cédric Fournet, Anitha Gollamudi, Georges Gonthier, Nadim Kobeissi, Natalia Kulatova, Aseem Rastogi, Thomas Sibut-Pinote, Nikhil Swamy, and Santiago Zanella-Béguelin. 2016. Formal Verification of Smart Contracts: Short Paper. In *Proceedings of the 2016 ACM Workshop on Programming Languages and Analysis for Security (PLAS '16)*. ACM, 91–96. <https://doi.org/10.1145/2993600.2993611>
- [6] Lorenz Breidenbach, Phil Daian, Florian Tramèr, and Ari Juels. 2018. Enter the Hydra: Towards Principled Bug Bounties and Exploit-Resistant Smart Contracts. In *Proc. 27th USENIX Security Symposium (USENIX Security'18)*. USENIX Association, 1335–1352. <https://www.usenix.org/conference/usenixsecurity18/presentation/breidenbach>
- [7] ChainSecurity. 2019. Constantinople enables new Reentrancy Attack. <https://medium.com/chainsecurity/constantinople-enables-new-reentrancy-attack-ace4088297d9>.
- [8] Ashu Daniel Lv. 2016. CarbonVote. <https://http://carbonvote.com/>.
- [9] devops199. 2017. anyone can kill your contract #6995. <https://github.com/paritytech/parity-ethereum/issues/6995>.
- [10] Josselin Feist, Gustavo Grieco, and Alex Groce. 2019. Slither: a static analysis framework for smart contracts. In *2019 IEEE/ACM 2nd International Workshop on Emerging Trends in Software Engineering for Blockchain (WETSEB)*. IEEE, 8–15.
- [11] Ilya Grishchenko, Matteo Maffei, and Clara Schneidewind. 2018. A Semantic Framework for the Security Analysis of Ethereum Smart Contracts. In *Proc. 7th International Conference on Principles of Security and Trust (POST'18) (Lecture Notes in Computer Science)*, Lujo Bauer and Ralf Küsters (Eds.), Vol. 10804. Springer, 243–269. https://doi.org/10.1007/978-3-319-89722-6_10
- [12] Shelly Grossman, Ittai Abraham, Guy Golan-Gueta, Yan Michalevsky, Noam Rinetzky, Mooly Sagiv, and Yoni Zohar. 2017. Online detection of effectively callback free objects with applications to smart contracts. *Proceedings of the ACM on Programming Languages* 2, POPL (2017), 48.
- [13] NCC Group. 2018. Decentralized Application Security Project (DASP) Top 10. <https://dasp.co/index.html>.
- [14] Feng Hao, Peter YA Ryan, and Piotr Zielinski. 2010. Anonymous voting by two-round public discussion. *IET Information Security* 4, 2 (2010), 62–67.
- [15] Jingxuan He, Mislav Balunovic, Nodar Ambroladze, Petar Tsankov, and Martin T. Vechev. 2019. Learning to Fuzz from Symbolic Execution with Application to Smart Contracts. In *Proc. 26th ACM SIGSAC Conference on Computer and Communications Security (CCS'19)*, Lorenzo Cavallaro, Johannes Kinder, XiaoFeng Wang, and Jonathan Katz (Eds.). ACM, 531–548. <https://doi.org/10.1145/3319535.3363230>
- [16] Everett Hildenbrandt, Manasvi Saxena, Nishant Rodrigues, Xiaoran Zhu, Philip Daian, Dwight Guth, Brandon Moore, Daejun Park, Yi Zhang, Andrei Stefanescu, et al. 2018. Kevm: A complete formal semantics of the ethereum virtual machine. In *2018 IEEE 31st Computer Security Foundations Symposium (CSF)*. IEEE, 204–217.

- [17] Friðrik Hjálmarsson, Gunnlaugur K Hreioarsson, Mohammad Hamdaqa, and Gisli Hjálmtýsson. 2018. Blockchain-based e-voting system. In *2018 IEEE 11th International Conference on Cloud Computing (CLOUD)*. IEEE, 983–986.
- [18] Bo Jiang, Ye Liu, and WK Chan. 2018. Contractfuzzer: Fuzzing smart contracts for vulnerability detection. In *Proceedings of the 33rd ACM/IEEE International Conference on Automated Software Engineering*. ACM, 259–269.
- [19] Sukrit Kalra, Seep Goel, Mohan Dhawan, and Subodh Sharma. 2018. ZEUS: Analyzing Safety of Smart Contracts. In *Proc. 25th Network and Distributed System Security Symposium (NDSS'18)*. The Internet Society, 1–12.
- [20] Aashish Kolluri, Ivica Nikolic, Ilya Sergey, Aquinas Hobor, and Prateek Saxena. 2019. Exploiting the laws of order in smart contracts. In *Proceedings of the 28th ACM SIGSOFT International Symposium on Software Testing and Analysis*. ACM, 363–373.
- [21] Johannes Krupp and Christian Rossow. 2018. teEther: Gnawing at Ethereum to Automatically Exploit Smart Contracts. In *27th USENIX Security Symposium (USENIX Security'18)*, William Enck and Adrienne Porter Felt (Eds.). USENIX Association, 1317–1333. <https://www.usenix.org/conference/usenixsecurity18/presentation/krupp>
- [22] Kibin Lee, Joshua I James, Tekachew G Ejeta, and Hyoungh J Kim. 2016. Electronic voting service using block-chain. *Journal of Digital Forensics, Security and Law* 11, 2 (2016), 8.
- [23] Loi Luu, Duc-Hiep Chu, Hrishi Olickel, Prateek Saxena, and Aquinas Hobor. 2016. Making Smart Contracts Smarter. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security (CCS '16)*. ACM, New York, NY, USA, 254–269. <https://doi.org/10.1145/2976749.2978309>
- [24] Patrick McCorry, Siamak F. Shahandashti, and Feng Hao. 2017. A Smart Contract for Boardroom Voting with Maximum Voter Privacy. In *Proc. 21st International Conference on Financial Cryptography and Data Security (FC'17) (Lecture Notes in Computer Science)*, Aggelos Kiayias (Ed.), Vol. 10322. Springer, 357–375. https://doi.org/10.1007/978-3-319-70972-7_20
- [25] Satoshi Nakamoto. 2009. Bitcoin: A Peer-to-Peer Electronic Cash System. *Cryptography Mailing list* at <https://metzdowd.com> (03 2009).
- [26] Ivica Nikolić, Aashish Kolluri, Ilya Sergey, Prateek Saxena, and Aquinas Hobor. 2018. Finding the greedy, prodigal, and suicidal contracts at scale. In *Proc. 34th Annual Computer Security Applications Conference (ACSAC'18)*. ACM, 653–663.
- [27] Ryan Osgood. 2016. The future of democracy: Blockchain voting. *COMP116: Information Security* (2016), 1–21.
- [28] Anton Permenev, Dimitar Dimitrov, Petar Tsankov, Dana Drachler-Cohen, and Martin Vechev. 2020. Verx: Safety verification of smart contracts. In *Proc. 41st IEEE Symposium on Security and Privacy (IEEE SP'20)*. IEEE, 18–20.
- [29] Sergey Petrov. 2017. Another Parity Wallet hack explained. <https://medium.com/@ProGer/another-parity-wallet-hack-explained-847ca46a2e1c>.
- [30] Michael Rodler, Wenting Li, Ghassan O. Karame, and Lucas Davi. 2019. Re-Entrancy Attack Patterns. <https://github.com/uni-due-syssec/eth-reentrancy-attack-patterns>.
- [31] Michael Rodler, Wenting Li, Ghassan O. Karame, and Lucas Davi. 2019. Sereum: Protecting Existing Smart Contracts Against Re-Entrancy Attacks. In *Proc. 26th Network and Distributed System Security Symposium (NDSS'19)*. The Internet Society.
- [32] David Siegel. 2016. Understanding The DAO Attack. <https://www.coindesk.com/understanding-dao-hack-journalists/>.
- [33] Xin Sun, Quanlong Wang, Piotr Kulicki, and Mirek Sopek. 2019. A simple voting protocol on quantum blockchain. *International Journal of Theoretical Physics* 58, 1 (2019), 275–281.
- [34] A Tann, Xing Jie Han, Sourav Sen Gupta, and Yew-Soon Ong. 2018. Towards safer smart contracts: A sequence learning approach to detecting vulnerabilities. *arXiv preprint arXiv:1811.06632* (2018).
- [35] Sergei Tikhomirov, Ekaterina Voskresenskaya, Ivan Ivanitskiy, Ramil Takhaviev, Evgeny Marchenko, and Yaroslav Alexandrov. 2018. SmartCheck: Static Analysis of Ethereum Smart Contracts. In *Proc. 1st IEEE/ACM International Workshop on Emerging Trends in Software Engineering for Blockchain (WETSEB@ICSE'18)*. ACM, 9–16. <http://ieeexplore.ieee.org/document/8445052>
- [36] Christof Ferreira Torres, Julian Schütte, and Radu State. 2018. Osiris: Hunting for Integer Bugs in Ethereum Smart Contracts. In *Proceedings of the 34th Annual Computer Security Applications Conference (ACSAC '18)*. ACM, New York, NY, USA, 664–676. <https://doi.org/10.1145/3274694.3274737>
- [37] Christof Ferreira Torres, Mathis Steichen, and Radu State. 2019. The Art of The Scam: Demystifying Honeypots in Ethereum Smart Contracts. In *28th USENIX Security Symposium (USENIX Security 19)*. USENIX Association, Santa Clara, CA, 1591–1607.
- [38] Petar Tsankov, Andrei Dan, Dana Drachler-Cohen, Arthur Gervais, Florian Buenzli, and Martin Vechev. 2018. Securify: Practical security analysis of smart contracts. In *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security*. ACM, 67–82.
- [39] S. Varrette, P. Bouvry, H. Cartiaux, and F. Georgatos. 2014. Management of an Academic HPC Cluster: The UL Experience. In *Proc. of the 2014 Intl. Conf. on High Performance Computing & Simulation (HPCS 2014)*. IEEE, Bologna, Italy, 959–967.
- [40] Haijun Wang, Yi Li, Shang-Wei Lin, Lei Ma, and Yang Liu. 2019. Vultron: catching vulnerable smart contracts once and for all. In *Proc. 41st International Conference on Software Engineering: New Ideas and Emerging Results (ICSE NIER'19)*, Anita Sarma and Leonardo Murta (Eds.). IEEE / ACM, 1–4. <https://doi.org/10.1109/ICSE-NIER.2019.00009>
- [41] Gavin Wood. 2014. Ethereum: A secure decentralised generalised transaction ledger. *Ethereum Project Yellow Paper* 151 (2014), 1–32.
- [42] Valentin Wüstholtz and Maria Christakis. 2019. Harvey: A greybox fuzzer for smart contracts. *arXiv preprint arXiv:1905.06944* (2019).
- [43] Pengcheng Zhang, Feng Xiao, and Xiapu Luo. 2019. SolidityCheck: Quickly Detecting Smart Contract Problems Through Regular Expressions. *arXiv preprint arXiv:1911.09425* (2019).
- [44] Wolfie Zhao. 2017. \$30 Million: Ether Reported Stolen Due to Parity Wallet Breach. <https://www.coindesk.com/30-million-ether-reported-stolen-parity-wallet-breach>.
- [45] Zhichao Zhao and T.-H. Hubert Chan. 2015. How to Vote Privately Using Bitcoin. In *Proc. 17th International Conference on Information and Communications Security (ICICS'15) (Lecture Notes in Computer Science)*, Sihang Qing, Eiji Okamoto, Kwangjo Kim, and Dongmei Liu (Eds.), Vol. 9543. Springer, 82–96. https://doi.org/10.1007/978-3-319-29814-6_8
- [46] Yi Zhou, Deepak Kumar, Surya Bakshi, Joshua Mason, Andrew Miller, and Michael Bailey. 2018. Erays: reverse engineering ethereum's opaque smart contracts. In *27th {USENIX} Security Symposium ({USENIX} Security 18)*. USENIX Association, 1371–1385.

A COMPLETE LIST OF ÆGIS'S ATTACK PATTERNS

Table 5 provides a complete list of vulnerabilities as well as their respective attack patterns that ÆGIS is currently capable to detect.

B UNCONDITIONAL REENTRANCY EXAMPLE

Figure 9 shows an example of a smart contract with an unconditional reentrancy. In this example an attacker first deposits a small amount of ether and then uses a reentrancy attack in order to drain all the ether that every single user has deposited.

```

1 contract VulnBank {
2     mapping (address => uint) public userBalances;
3
4     function deposit() public payable {
5         userBalances[msg.sender] += msg.value;
6     }
7
8     function withdrawAll() public {
9         uint amountToWithdraw = userBalances[msg.sender];
10        msg.sender.call.value(amountToWithdraw)("");
11        userBalances[msg.sender] = 0;
12    }
13 }

```

Figure 9: Example of a contract that is vulnerable to unconditional reentrancy [30].

Vulnerability	Attack Pattern
Same-Function Reentrancy	$(\text{opcode} = \text{CALL}) \Rightarrow (\text{opcode} = \text{CALL}) \text{ where } (\text{src.stack}(1) = \text{dst.stack}(1)) \wedge$ $(\text{src.address} = \text{dst.address}) \wedge (\text{src.pc} = \text{dst.pc}) \rightarrow$ $(\text{opcode} = \text{SSTORE}) \rightarrow (\text{opcode} = \text{SSTORE}) \text{ where } (\text{src.stack}(0) = \text{dst.stack}(0)) \wedge$ $(\text{src.address} = \text{dst.address}) \wedge (\text{src.depth} > \text{dst.depth})$
Cross-Function Reentrancy	$(\text{opcode} = \text{CALL}) \Rightarrow (\text{opcode} = \text{CALL}) \text{ where}$ $(\text{src.stack}(1) = \text{dst.stack}(1)) \wedge (\text{src.address} = \text{dst.address}) \wedge$ $(\text{src.memory}(\text{src.stack}(3), \text{src.stack}(4)) \neq \text{dst.memory}(\text{dst.stack}(3), \text{dst.stack}(4))) \rightarrow$ $(\text{opcode} = \text{SSTORE}) \rightarrow (\text{opcode} = \text{SSTORE}) \text{ where } (\text{src.stack}(0) = \text{dst.stack}(0)) \wedge$ $(\text{src.address} = \text{dst.address}) \wedge (\text{src.depth} > \text{dst.depth})$
Delegated Reentrancy	$(\text{opcode} = \text{DELEGATECALL}) \Rightarrow (\text{opcode} = \text{DELEGATECALL}) \text{ where } (\text{src.stack}(1) = \text{dst.stack}(1)) \wedge$ $(\text{src.address} = \text{dst.address}) \wedge (\text{src.pc} = \text{dst.pc}) \rightarrow$ $(\text{opcode} = \text{SSTORE}) \rightarrow (\text{opcode} = \text{SSTORE}) \text{ where } (\text{src.stack}(0) = \text{dst.stack}(0)) \wedge$ $(\text{src.address} = \text{dst.address}) \wedge (\text{src.depth} > \text{dst.depth})$ <hr/> $(\text{opcode} = \text{CALLCODE}) \Rightarrow (\text{opcode} = \text{CALLCODE}) \text{ where } (\text{src.stack}(1) = \text{dst.stack}(1)) \wedge$ $(\text{src.address} = \text{dst.address}) \wedge (\text{src.pc} = \text{dst.pc}) \rightarrow$ $(\text{opcode} = \text{SSTORE}) \rightarrow (\text{opcode} = \text{SSTORE}) \text{ where } (\text{src.stack}(0) = \text{dst.stack}(0)) \wedge$ $(\text{src.address} = \text{dst.address}) \wedge (\text{src.depth} > \text{dst.depth})$
Create-Based Reentrancy	$(\text{opcode} = \text{CREATE}) \Rightarrow (\text{opcode} = \text{CREATE}) \text{ where } (\text{src.stack}(1) = \text{dst.stack}(1)) \wedge$ $(\text{src.address} = \text{dst.address}) \wedge (\text{src.pc} = \text{dst.pc}) \rightarrow$ $(\text{opcode} = \text{SSTORE}) \rightarrow (\text{opcode} = \text{SSTORE}) \text{ where } (\text{src.stack}(0) = \text{dst.stack}(0)) \wedge$ $(\text{src.address} = \text{dst.address}) \wedge (\text{src.depth} > \text{dst.depth})$
Parity Wallet Hack 1	$(\text{opcode} = \text{DELEGATECALL}) \Rightarrow (\text{opcode} = \text{CALLDATACOPY}) \rightsquigarrow (\text{opcode} = \text{SSTORE}) \rightsquigarrow (\text{opcode} = \text{JUMPI}) \text{ where}$ $(\text{src.transaction.hash} \neq \text{dst.transaction.hash}) \rightarrow$ $((\text{opcode} = \text{CALLDATALOAD}) \rightsquigarrow (\text{opcode} = \text{CALL})) \text{ where}$ $(\text{dst.stack}(2) > 0)$
Parity Wallet Hack 2	$(\text{opcode} = \text{CALLDATACOPY}) \rightsquigarrow (\text{opcode} = \text{SSTORE}) \rightsquigarrow (\text{opcode} = \text{JUMPI}) \text{ where}$ $(\text{src.transaction.hash} \neq \text{dst.transaction.hash}) \rightarrow$ $((\text{opcode} = \text{CALLDATALOAD}) \rightsquigarrow (\text{opcode} = \text{SELFDESTRUCT}))$
Integer Overflow (Addition)	$(\text{opcode} = \text{CALLDATALOAD}) \rightsquigarrow (\text{opcode} = \text{ADD}) \text{ where}$ $((\text{dst.stack}(0) + \text{dst.stack}(1)) \neq \text{dst.stack.result}) \rightsquigarrow (\text{opcode} = \text{CALL})$
Integer Overflow (Multiplication)	$(\text{opcode} = \text{CALLDATALOAD}) \rightsquigarrow (\text{opcode} = \text{MUL}) \text{ where}$ $((\text{dst.stack}(0) * \text{dst.stack}(1)) \neq \text{dst.stack.result}) \rightsquigarrow (\text{opcode} = \text{CALL})$
Integer Underflow	$(\text{opcode} = \text{CALLDATALOAD}) \rightsquigarrow (\text{opcode} = \text{SUB}) \text{ where}$ $((\text{dst.stack}(0) - \text{dst.stack}(1)) \neq \text{dst.stack.result}) \rightsquigarrow (\text{opcode} = \text{CALL})$
Timestamp Dependence	$(\text{opcode} = \text{TIMESTAMP}) \rightsquigarrow (\text{opcode} = \text{JUMPI}) \rightarrow (\text{opcode} = \text{CALL}) \text{ where } (\text{dst.stack}(2) > 0)$
Transaction Order Dependency	$(\text{opcode} = \text{SSTORE}) \rightsquigarrow (\text{opcode} = \text{SLOAD}) \text{ where}$ $(\text{src.block.number} = \text{dst.block.number}) \wedge (\text{src.transaction.from} \neq \text{dst.transaction.from})$

Table 5: List of vulnerabilities and their respective attack patterns.