



The Good, The Bad and The Ugly: Pitfalls and Best Practices in Automated Sound Static Analysis of Ethereum Smart Contracts

Clara Schneidewind^(✉), Markus Scherer^(✉), and Matteo Maffei^(✉)

TU Wien, Vienna, Austria

{clara.schneidewind,markus.scherer,matteo.maffei}@tuwien.ac.at
<https://secpriv.wien/>

Abstract. Ethereum smart contracts are distributed programs running on top of the Ethereum blockchain. Since program flaws can cause significant monetary losses and can hardly be fixed due to the immutable nature of the blockchain, there is a strong need of automated analysis tools which provide formal security guarantees. Designing such analyzers, however, proved to be challenging and error-prone. We review the existing approaches to automated, sound, static analysis of Ethereum smart contracts and highlight prevalent issues in the state of the art. Finally, we overview *eThor*, a recent static analysis tool that we developed following a principled design and implementation approach based on rigorous semantic foundations to overcome the problems of past works.

Keywords: Static analysis · Smart contracts · Formal methods

1 Introduction

Blockchain technologies are revolutionizing the distributed system landscape, providing an innovative solution to the consensus problem leveraging probabilistic guarantees and incentives. In particular, they allow for the secure execution of payments, and more in general computations, among mutually distrustful parties. While some cryptocurrencies, like Bitcoin [27] provide only a limited scripting language tailored to payments, others, like Ethereum [32], support a quasi Turing complete¹ smart contract language, allowing for advanced applications such as trading platforms [25, 28], elections [26], permission management [7, 11], data management systems [5, 29], or auctions [13, 17]. With the growing complexity of smart contracts, however, also the attack surface grows. This is particularly problematic as smart contracts control real money flows and hence constitute an attractive target for attackers. In addition, due to the immutable nature of blockchains, smart contracts cannot be modified once they are uploaded to the blockchain, which makes the effects of security vulnerabilities permanent. This is not only a theoretical threat, but a practical problem, as demonstrated by

¹ Supporting a Turing complete instruction set, Ethereum enforces termination by bounding the number of computation steps based on an prespecified resource limit.

infamous hacks, such as the DAO hack [1] or the Parity hacks [2, 3] which caused losses of several millions of dollars. This state of affairs calls for reliable static analysis tools which are accessible to the users of the Ethereum system, that is, developers, who need to be able to verify their contracts before uploading them to the blockchain, and users interacting with existing smart contracts, who need tool assistance to assess whether or not those contracts (which are published in human unreadable bytecode format on the blockchain) are fraudulent.

State of the Art. The bug-finding tool Oyente [24] (published in 2016) pioneered the (automatic) static analysis of Ethereum smart contracts. This work highlighted, for the first time, generic types of bugs that typically affect smart contracts, and proposed a tool based on symbolic execution for the detection of contracts vulnerable to these bugs.

A particular compelling feature of Oyente is that it is a push-button tool that does not expect any interaction or deeper knowledge of the contract semantics from the user. On the downside, however, Oyente does not provide any guarantees on the reported results, being neither sound (absence of false negatives) nor complete (absence of false positives) and, thereby, yielding only a heuristic indication of contract security. Given the importance of rigorous security guarantees in this context, several approaches have been later proposed for the verification of Ethereum smart contracts. In particular, alongside tools that aim at machine-checked smart contract auditing [6, 9, 18, 19], a line of work focused on automation in the verification process to ensure usability and broad adaptation. Despite four years of intense research, however, until now only four works on such sound and fully automatic static analysis of Ethereum smart contracts have been published. Furthermore, all of these works exhibit shortcomings which ultimately undermine the security guarantees that they aim to provide.

Our Contributions. Motivated by these struggles, we overview the difficulties that arise in the design of sound static analysis for Ethereum smart contracts, and the pitfalls that so far hindered the development of such analyzers. To this end, we first give a short introduction to the Ethereum platform and its native smart contract language (Sect. 2). Afterwards, we illustrate the challenges in automated smart contract verification by overviews existing works in this field with emphasis on their weak spots (Sect. 3 and 4). We conclude the paper summarizing our experiences in the design of *eThor* [30], a static analyzer for EVM bytecode we recently introduced, along with a breakdown of the components we deem crucial to achieve efficiency and soundness (Sect. 5).

2 Ethereum and the EVM Bytecode Language

Ethereum is (after Bitcoin) the second most widely used cryptocurrency with a market capitalization of over 14 billion U.S. dollars². Compared to Bitcoin,

² As of the fourth quarter of 2019, see <https://www.statista.com/statistics/807195/ethereum-market-capitalization-quarterly>.

Ethereum stands out due its expressive scripting language, that enables the execution of arbitrary distributed programs, so called *smart contracts*. Ethereum smart contracts are stored on the blockchain in bytecode that is jointly executed by the network participants (also called *nodes*) according to the Ethereum Virtual Machine (EVM) – Ethereum’s execution environment that is implemented in different clients³. Smart contract execution is part of Ethereum’s consensus protocol: The state of the system is determined by a jointly maintained, tamper-resistant public ledger, the *blockchain*, that holds a sequence of transactions. Transactions do not only indicate money transfers, but can also trigger contract executions. To derive the state of the system, every node locally executes the transactions in the blockchain as specified by the EVM. For advancing the system, nodes broadcast transactions which are assembled into blocks and appended to the blockchain. While the correctness of the system is ensured by the nodes validating all blocks, fairness is established by a proof of work mechanism: Only nodes (called *miners*) capable of solving a computationally hard puzzle are eligible to propose a block. This adds randomness the selection of proposers and prevents a minority from steering the evolution of the system.

Ethereum Ecosystem. The state of the Ethereum system (called *global state*) consists of the state of all virtual accounts and their balances in the virtual currency *Ether*. Smart contracts are special account entities (*contract accounts*) that in addition to a balance hold persistent storage and the contract code. While non-contract (so called *external*) accounts can actively transfer fractions of their balance to other accounts, contract accounts are purely governed by their code: A contract account can be activated by a transaction from another account and then executes its code, possibly transferring money or stimulating other contracts. Similarly, a contract can be created on behalf of an external account or by another contract. We will in the following refer to such inter-contract interactions as *internal transactions*, as opposed to *external transactions* that originate from external accountexternal accounts and are explicitly recorded on the blockchain.

EVM Bytecode. The EVM bytecode language supports designated instructions to account for the blockchain ecosystem. These encompass primitives for different flavors of money transfers, contract invocations, and contract creations. Most prominently, the CALL instruction allows for transferring money to another account while at the same time triggering code execution in case the recipient is a contract account. Other domain-specific bytecodes include instructions for accessing the blockchain environment such as the instructions SSTORE and SLOAD for reading and writing the cells of the persistent contract storage. Further the instruction set contains opcodes for accessing information on the ongoing (internal) transaction (its caller, input, or value transferred along) and for

³ Currently, a Go, a C++ and a Python implementation are distributed by the Ethereum Foundation: <https://github.com/ethereum/wiki/wiki/Clients,-tools,-dapp-browsers,-wallets-and-other-projects>.

computation: The EVM is a stack-based machine supporting standard instructions for arithmetic and stack manipulation. The control flow of a contract is also subject to the stack-based architecture: conditional and unconditional jump instructions allow for resuming execution at another program position that is determined by the value on the stack. While these features would make EVM bytecode Turing-complete, to enforce termination the execution of EVM smart contracts is bounded by an upfront-specified resource called *gas*. Every instruction consumes gas and the execution halts with an exception when running out of gas. The gas budget is set by the initiator of the transaction who will pay a compensation to the miner of the enclosing block for the effectively consumed amount of gas when executing the transaction. Due to the low-level nature of EVM bytecode, smart contracts are generally written in high-level languages (most prominently the *Solidity* [4] language) and compiled to EVM bytecode.

3 Related Work on Automated Sound Static Analysis of Ethereum Smart Contracts

We overview the state of the art in the automated static analysis of Ethereum smart contracts. So far there have been works on four static analyzers published that come with (explicit or implicit) soundness claims: the dependency analysis tool Securify [31] for EVM bytecode, the static analyzer ZEUS [22] for *Solidity*, the syntax-guided *Solidity* analyzer NeuCheck [23], and the bytecode-based reachability analysis tool EtherTrust [14]. By implicit soundness claim, we mean that the tool claims that a positive analysis result guarantees the contract’s security (i.e., absence of false negatives with respect to a specific security property). While Securify, ZEUS, and EtherTrust implement semantic-based analysis approaches, NeuCheck is purely syntax-driven.

Securify supports data and control flow analysis on the EVM bytecode level. To this end, it reconstructs the control-flow graph (CFG) from the contract bytecode and transforms it into SSA-form. Based on this structured format, it models immediate data and control flow dependencies using logical predicates and establishes datalog-style logical rules for deriving transitive dependencies which are automatically computed using the enhanced datalog engine *Soufflé* [21]. For checking different security properties, Securify specifies patterns based on the derived predicates which shall be sufficient for either proving a property (compliance patterns) or for showing a property to be broken (violation patterns).

ZEUS analyzes *Solidity* contracts by first translating them into the intermediate language LLVM bitcode and then using off-the-shelf model checkers to verify different security properties. In the course of the translation, ZEUS uses another intermediate layer for the *Solidity* language which introduces abstractions and that allows for the insertion of assumptions and assertions into the program code which express security requirements on the contract. The security properties supported by ZEUS are translated to such assertion checks, possibly in conjunction with additional property-specific contract transformations.

NeuCheck analyzes *Solidity* contracts by pattern matching on the contract syntax graph. To this end, it translates *Solidity* source code into an XML parse tree. Security properties are expressed as patterns on this parse tree and are matched by custom algorithms traversing the tree.

EtherTrust implements a reachability analysis on EVM bytecode by abstracting the bytecode execution semantics into (particular) logical implications, so called Horn clauses, over logical predicates representing the execution state of a contract. Security properties are expressed as reachability queries on logical predicates and solved by the SMT solver *z3* [12]. EtherTrust was a first prototype that later evolved into the *eThor* analyzer, which we will discuss in Sect. 5.1.

All presented tools focus on generic (contract-independent) security properties for smart contracts. However, the underlying architectures allow for extending the frameworks with further properties. Foremost, the tool *eThor* supports general reachability properties and hence also functional properties characterized in terms of pre- and postconditions. For the soundness considerations in this paper we put the focus on the abstractions of generic security properties.

4 Challenges in Sound Smart Contract Verification

EVM bytecode exposes several domain-specific subtleties that turn out to be challenging for static analysis. Furthermore, even the definition itself of security properties for smart contracts is highly non-trivial and subject to ongoing research. Furthermore, characterizing relevant generic security properties for smart contracts is highly non-trivial and subject to ongoing research. We will examine both of these problems in the following.

4.1 Analysis Design

We summarize below the main challenges that arise when designing a performant and still sound analysis for Ethereum smart contracts:

- *Dynamic jump destinations*: Jump destinations are statically unknown and computed during execution. They might be influenced by the blockchain environment as well as the contract state. As a consequence, the control flow graph of a contract is not necessarily determinable at analysis time.
- *Mismatch between memory and stack layout*: The EVM has a (stack) word size of 256 bits while the memory (heap) is fragmented into bytes and addressed accordingly. Loading words from memory to the stack, and conversely writing stack values to memory, requires (potentially costly) conversions between these two value formats.
- *Exception propagation and global state revocation*: If an internal transaction (as, e.g., initiated by a CALL) fails, all effects of this transaction including those on the global state (e.g., writes to global storage) are reverted. However, such a failure is not propagated to the callee, who can continue execution in the original global state. Modeling calls must thus save the state before calling in order to account for global state revocation.

- *Native support for low-level cryptography:* The EVM supports a designated SHA3 instruction to compute the hash of some memory fraction. As a consequence, hashing finds broad adaption in Ethereum smart contracts, and the Solidity compiler bases its storage layout on a hash-based allocation scheme.
- *Dynamic calls:* The recipient of an (inter-contract) call is specified on the stack and hence subject to prior computation. Consequently, the recipient is not necessarily derivable at analysis time, resulting in uncertainty about the behavior of the callee and the resulting effects on the environment.
- *Dynamic code creation:* Ethereum supports the generation of new smart contracts during transaction execution: A smart contract can deploy another one at runtime. To do so, the creating smart contract reads the deployment code for the new contract from the heap. The newly created contract may hence be subject to prior computation and even to the state of the blockchain.

In order to effectively tackle these challenges, several contributions of independent interest are required, such as domain-specific abstractions (e.g., suitable over-approximations of unknown environment behavior); the preprocessing of the contract to reconstruct its control flow or call graph; (easily checkable) assumptions that restrict the analysis scope (e.g., restriction to some language fragment); and optimizations or simplifications in intermediate processing steps (e.g, contract transformations to intermediate representations). Altogether, these analysis steps enlarge the semantic gap between the original contract semantics and the analysis, making it harder to reliably ensure the soundness of the latter. In the following, we will review in more detail the tension between soundness and performance of the analysis, and how past works stumbled in this minefield.

Soundness. Ensuring the soundness of the analysis requires a rigorous specification of the semantics of EVM bytecode. The original semantics was written in the Yellow Paper [32]. This paper however, from the beginning exhibited flaws [15, 18, 19] and underspecified several aspects of bytecode execution. The ultimate truth of smart contract semantics could therefore only be extracted from the client implementations provided by the Ethereum foundation. In the course of time, several formal specifications of the EVM semantics have been proposed by the scientific community [15, 18, 19], leading the Yellow paper to be replaced by an executable semantics in the K framework [18]⁴.

For the high level language Solidity, despite first efforts within the scientific community [8, 10, 20, 34, 35], there exists at present no full and generally accepted formal semantics. Consequently the semantics of Solidity is only defined by its compilation to EVM bytecode. Since the compiler is subject to constant changes, Solidity constitutes a moving target.

The complexity and uncertainty about the concrete semantics made most works build on ad-hoc simplified versions of the semantics which do not cover all language features and disregard essential aspects of the EVM’s execution model.

⁴ Also called the Jello paper: <https://jellopaper.org>.

```

1 contract Test {
2   bool test = false;
3   function flipper () { if (msg.sender != 0){flip();} }
4   function flip () internal {test = !test;} }

```

Fig. 1. Simple contract highlighting an unsoundness in Securify’s dependency analysis.

ZEUS [22], for instance, defines an intermediate goto language for capturing the core of Solidity. The semantics of this language, however, is inspired by the (ad-hoc) semantic modeling used in Oyente [24], inheriting an essential flaw concerning global state revocation: In case that an internal transaction returns with an exception, the effects on the global state are not reverted as they would be in real EVM (and Solidity) executions. Since the translation to the intermediate language is part of the analysis pipeline of [22], such a semantic flaw compromises the soundness of the whole analysis.

Also Securify [31] introduces an ad-hoc formalism for EVM bytecode semantics. This is not, however, related to the dependency predicates used for the analysis, but just serves for expressing security properties. It is hence unclear to which extent the dependency predicates faithfully reflect the control flow and value dependencies induced by the EVM semantics. Assessing the correctness of this approach is difficult, since no full logical specification of the dependency analysis is provided⁵. Indeed we found empirical indication for the unsoundness of the dependency analysis in the presence of complicated control flow. Consider the example contract depicted in Fig. 1. For better readability, we present the contract in the high-level language *Solidity*, a language inspired by *JavaScript*, that is centered around contracts which are used analogously to the concept of classes in object-oriented programming languages. The depicted contract `Test` has a global boolean field `test`. Global fields are reflected in the persistent storage of the contract and constitute the contract state. The public function `flipper()` allows every account but the one with address 0 to flip the value of the `test` field: For checking the restriction on the calling account, the `flipper()` function accesses the address of the caller using *Solidity*’s `msg.sender` construct. For writing the `test` field, the internal function `flip()` is called. Internal functions are not exposed to the public, but are only accessible by the contract itself and calls to such functions are compiled to local jumps. The use of internal functions consequently substantially complicates the control flow of a contract.

We identified a soundness issue affecting the conditional reachability of contract locations. We identified a correctness issue that affects both the soundness and completeness of Securify. This incorrectness becomes evident in the violation pattern that checks for unrestricted writes. An unrestricted write is a write access of the global storage that can be performed by any caller. The violation pattern states that such an unrestricted write is guaranteed to happen if there is a `SSTORE` instruction whose reachability does not depend on the caller of

⁵ Only an excerpt is presented in [31], and the public implementation at <https://github.com/eth-sri/securify> intermingles specification and implementation.

the contract. This pattern should not be matched by the `Test` contract since the only write access to the `Test`'s sole variable `test` in function `flip()` is only reachable via the function `flipper` where it is conditioned on the caller (`msg.sender`). Hence not every contract can write the `test` variable, but the write access depends on the caller. Still Securify reports this contract to match the violation pattern, consequently proving the wrong statement that there is no dependency between writing the `test` field and the account calling the contract. Note that even though showing up in a violation pattern (hence technically producing false positives), the underlying issue also affects the soundness of the core dependency analysis⁶. Securify specifies a *may dependency* relation to capture (potential) abstract dependencies between program locations. For correctly (i.e. soundly) abstracting the dependencies in real programs, the absence of a may dependency should imply a corresponding independence in the real program. Since the may dependency relation is used in both compliance and violation patterns, without such a guarantee Securify can be neither sound nor complete. The example refutes this guarantee and thereby illustrates the importance of providing clear formal correctness (i.e. soundness) statements for the analysis.

These two examples show how the missing semantic foundations of the presented analysis approaches can lead to soundness issues in the core analysis design itself. These problems are further aggravated once additional stages are added to the analysis pipeline for increasing performance, since such additional stages are often not part of the correctness considerations.

Performance. For performance reasons, it is often unavoidable to leverage well-established and optimized analysis frameworks or solvers. This leaves analysis designers with the challenge to transform their analysis problem into a format that is expressible and efficiently solvable within the targeted framework while preserving the semantics of the original problem.

ZEUS [22] makes use of off-the-shelf model checkers for LLVM bitcode and hence requires a transformation of smart contracts to LLVM bitcode. The authors describe this step to be a ‘faithful expression-to-expression translation’ that is semantics preserving, but omit a proof for this property. The paper itself later contradicts this statement: The authors report on LLVM’s optimizer impacting the desired semantics. This indicates that the semantics of the LLVM bitcode translation does not coincide with the one of the intermediate language, since it would otherwise not be influenced by (semantics-preserving) optimizations.

The Securify tool [31] makes use of several preprocessing steps in order to make EVM bytecode amenable to dependency analysis: First it reconstructs the control flow graph of a contract and based on that transforms the contract to SSA form. The correctness of these steps is never discussed. Indeed we found Securify’s algorithm for control flow reconstruction to be unsound: The algorithm fails when encountering jump destinations that depend on blockchain information. In such a case the control flow should be considered to be non-reconstructable since a jump to such a destination may result in a valid jump at runtime or

⁶ We illustrate the issue with a violation pattern for easier presentation and since the affected compliance pattern turned out not to be implemented in Securify.

```

1 contract DAO{
2   mapping (address => uint) bal;
3
4   function invest () public payable {
5     bal[msg.sender]+= msg.value;};
6
7   function withdraw () public {
8     address a = msg.sender;
9     if (bal[a] > 0){
10      a.call.value(bal[a]);
11      bal[a] = 0;};
12 }

```

```

1 contract Mallory{
2   address DAO_ADDRESS = 0x...;
3   DAO dao = DAO(DAO_ADDRESS);
4
5   function investSmallAmount() public {
6     ① dao.invest().value(1);};
7
8   function() payable{
9     ② dao.withdraw();};
10 }

```

Fig. 2. Simplified DAO contract.

simply fail due to a non-existing jump destination. Securify’s algorithm however does not report an error on such a contract, but returns a (modified) contract that does not contain jumps. Such an unsound preprocessing step again impacts the soundness of the whole analysis tool since it breaks the assumption that the contract semantics is preserved by preprocessing.

4.2 Security Properties

The Ethereum blockchain environment opens up several new attack vectors which are not present in standard (distributed) program execution environments. This is in particular due to the contracts’ interaction with the blockchain which is in general controlled by unknown parties and hence needs to be considered hostile. It is a partly still open research question what characterizes a contract that is robust in such an environment. A well-studied property in this domain is robustness against reentrancy attacks. We will focus on this property in the following to illustrate the challenges and pitfalls in proving a contract to be safe.

Reentrancy Attacks. Reentrancy attacks became famous due to the DAO hack [1] in 2016 which caused a loss of over 60 Million dollars, and ultimately led to a hard fork (a change in the consensus to explicitly ignore this particular incident) of the Ethereum blockchain. The DAO was a contract implementing a crowd-funding mechanism which allowed users to invest and conditionally withdraw their invested money from the contract. An attacker managed to exploit a bug in the contract’s withdraw functionality for draining the whole contract, stealing the money invested by other participants. We illustrate the main workings of this attack with a simplified example in Fig. 2.

The depicted DAO contract has a global field `bal` which is a mapping from account addresses to the payments that they made so far. The two (publicly accessible) functions of the contract allow arbitrary entities to invest and withdraw money from the contract. If an account with address `a` calls the `invest` function, the money transferred with this invocation is registered in the `bal` mapping. Similar to `msg.sender`, *Solidity* provides the variable `msg.value` to access the value transferred with the currently executed (internal) transaction.

The `withdraw` function when being called by a , will check the amount of money invested by a so far and in case of non-zero investments, transfer the whole amount of Ether (as recorded in `bal[a]`) back to a . This is done using Solidity's `call` construct for function invocations: it initiates a transaction to the specified address (here `a`) and allows for the specification of the value to be sent along (using `.value()`). The attack on the DAO contract can be conducted by an attacker that deploys a malicious contract `Mallory` to first make a small investment to the DAO contract (①) that they later withdraw (②). When the `withdraw` function of the DAO contract calls back to the sender (`Mallory`, ③), not only the corresponding amount of Ether is transferred, but also code of `Mallory` is executed. This is as in the case that not a specific (*Solidity*) function gets invoked with a contract call, the contract's *fallback function* (a function without name and arguments) is executed. `Mallory` implements this function to call the DAO's `withdraw` function (④). Since at this point the balance of `Mallory` in the `bal` mapping has not been updated yet, another value transfer to `Mallory` will be initiated (⑤). By proceeding in this way, `Mallory` can drain all funds of the DAO contract.

The depicted attack is an example of how standard intuitions from (sequential) programming do not apply to smart contracts: In Ethereum one needs to consider that an internal transaction hands over the control to a (partly) unknown environment that can potentially schedule arbitrary contract invocations.

Formalizing Security Properties. While bug-finding tools typically make use of heuristics to detect vulnerable contracts, there have been two systematic studies that aim at giving a semantic characterization of what it means for a contract to be resistant against reentrancy attacks: The resulting security definitions are call integrity [15] and effective callback freedom [16].

Call integrity follows non-interference-style integrity definitions from the security community. It states that two runs of a contract in which the codes of the environment accounts may differ, should result in the same sequences of observable events (in this case outgoing transactions). In simpler words, another contract should not be able to influence how a secure contract spends its money. Intuitively, this property is violated by the DAO contract since an attacker contract can make the contract send out more money than in an honest invocation.

In contrast, effective callback freedom is inspired by the concept of linearizability from concurrency theory: It should be possible to mimic every (recursive) execution of a contract by a sequence of non-recursive executions. The DAO contract violates this property since the attack is only possible when making use of recursion (or callbacks respectively). After each callback-free execution, the `investments` mapping will be updated, so that a subsequent execution will prevent further withdrawals by the same party.

While [15] shows how to over-approximate the hyperproperty call integrity by three simpler properties (the reachability property single-entrancy and two dependence properties), [16] does not indicate a way of statically verifying effective callback freedom, but proves this property to be undecidable. This leaves sound, and (efficiently) verifiable approximations an open research question.

```

1 library Lib {
2   struct Data { mapping (address => uint) map;}
3   function write(Data storage self, address a, uint v) { self.map[a] = v;}
4   function get(Data storage self, address a) returns (uint) {
5     return (self.map[a]);} }
6
7 contract DAO {
8   Lib.Data bal;
9   function invest() public payable {
10    Lib.write(bal, msg.sender, Lib.get(bal, msg.sender) + msg.value);}
11  function withdraw () public {
12    address a = msg.sender;
13    if (Lib.get(bal, a) > 0){
14      a.call.value(Lib.get(bal, a));
15      Lib.write(bal, a, 0);}} }

```

Fig. 3. Simplified DAO contract using a library

Checking Security Properties. The state-of-the-art sound analyzers discussed so far do not build on prior efforts of semantically characterizing robustness against reentrancy attacks, but come up instead with own ad-hoc definitions.

Securify. Securify expresses security properties of smart contracts in terms of compliance and violation patterns over data flow and control flow dependency predicates. In [31] it is stated that Securify supports the analysis of a property called ‘no writes after call’ (NW) which is different from (robustness against) reentrancy, but still aims at detecting bugs similar to the one in the DAO. The NW property is defined using an ad-hoc semantic formalism, and it states that for any contract execution trace, the contract storage shall not be subject to modifications after performing a CALL instruction. Intuitively, this property should exclude reentrancy attacks by preventing that the guards of problematic money transfers are updated only after performing the money transferring call. However, this criterion is not sufficient e.g., since reentrancies can also be triggered by instructions other than CALL. For proving the NW property, the compliance pattern demands that a CALL instruction may not be followed by any SSTORE instruction. We found this pattern not to be sufficient for ensuring compliance with the NW property (nor robustness against reentrancy). We will illustrate this using a variation of the DAO contract in Fig. 3. This contract implements the exact same functionality as the one in Fig. 2. The only difference is that the access to the balance mapping is handled via the library contract Lib. Ethereum actively supports the use of library contracts in that it provides a specific call instruction, called DELEGATECALL, that executes another contract’s code in the environment of the caller. When calling Lib.write in the withdraw function, such a delegated call to the (external) library contract is executed. Executing write in the context of contract DAO will then modify DAO’s storage (instead of the one of the Lib contract). In order to let the write and the get functionality access the right storage position (where DAO stores the bal mapping), these functions take as first argument the reference to the corresponding storage location. Same as the version in Fig. 2, this contract is vulnerable to a reentrancy bug. Also, it violates the NW property: The storage of the contract

```

1 contract DAO{
2   mapping (address => uint) bal;
3   uint lock = 0;
4   function withdraw () public {
5     if(lock ==1){throw;}
6     lock=1;
7     address a = msg.sender;
8     a.call.value(bal[a]);
9     bal[a] = 0;
10    lock=0;}
11  function switchLock () {
12    lock = 1-lock;} }

```

```

1 contract DAO{
2   mapping (address => uint) bal;
3   uint lock = 0;
4   function withdraw () public {
5     if(lock ==1){throw;}
6     lock=1;
7     address a = msg.sender;
8     a.call.value(bal[a]);
9     bal[a] = 0;
10    lock=0;}

```

Fig. 4. Simple versions of the DAO contract with reentrancy protection.

can be changed after executing the call (when writing the `bal`) mapping. Still, this contract matches the compliance pattern (which should according to [31] guarantee the contract to satisfy the NW property), since it does not contain any explicit `SSTORE` instruction. This example illustrates how without a proven connection between a property and its approximation, the soundness of an analyzer can be undermined. This issue does not only constitute a singular case, but is a structural problem: There are counter examples for the soundness of 13 out of the 17 patterns presented in [31], as we detail out in [30].

ZEUS. In [22], the property to rule out reentrancy attacks is only specified in prose as a function being vulnerable ‘if it can be interrupted while in the midst of its execution, and safely re-invoked even before its previous invocations complete execution.’ This definition works on the level of functions, a concept which is only present on the *Solidity* level, and leaves open the question what it means for a *contract* to be robust against reentrancy attacks. The authors distinguish between ‘same-function-reentrancy’ and ‘cross-function-reentrancy’ attacks, but do not consider cross-function reentrancy (where a function reenters another function of the same contract) in the analyzer. We found that without excluding cross-function reentrancy also single-function reentrancy cannot be prevented.

Consider the versions of the DAO contract depicted in Fig. 4 that aim to prevent reentrancy using a locking mechanism. The global `lock` field tracks whether the `withdraw` function was already entered (indicated by value 1). In that case, the execution of `withdraw` throws an exception. Otherwise the `lock` is set and only released when concluding the execution of `withdraw`. While the two depicted contracts implement the exact same `withdraw` function, the first contract’s function is vulnerable to a reentrancy attack, while the second one is safe. This is as the first contract implements a public `switchLock()` function that can be used by anyone to change the `lock` value. An attacker could hence mount the standard attack with the only difference that they would need to invoke the `switchLock()` function once before reentering to disable the reentrancy protection in line 5. Without exposing such functionality, the second contract is safe, since every reentering execution will be stopped in line 5. This example

shows that ZEUS’ approach of analyzing functions in isolation to exclude ‘same-function-reentrancy’ is not sound.

Another issue in the reentrancy checking of ZEUS is caused by the reentrancy property exceeding the scope of the analysis framework. For proving a function resistant against reentrancy attacks, ZEUS checks whether it is ever possible to reach a call when a function is recursively invoked by itself. However, the presented translation to LLVM bitcode only models non-recursive executions of a function. Consequently, the reentrancy property cannot be expressed as a policy (which could be translated to assertions in the program code), but requires to rewrite the contract under analysis to contain duplicate functions that mimic reentering function invocations. This contract transformation is not part of any soundness considerations. As a result, not only the previously discussed unsoundness due to the lacking treatment of cross-function reentrancies is missed, but it is also disregarded that *Solidity*’s `call` construct is not the only way to reinvoke a function. Indeed there are several other mechanisms (e.g., direct function calls) that allow for the same functionality. Still, ZEUS classifies contracts that do not contain an explicit invocation of the `call` construct to be safe by default.

NeuCheck. The NeuCheck tool formulates a syntactic pattern for detecting robustness against reentrancy attacks. The pattern checks for all occurrences of the `call` function whether they are followed by the assignment of a state variable. As discussed for Securify, the absence of explicit writes to the storage does not imply that the storage stays unchanged. Hence the example in Fig. 3 would also serve as a counter example for the soundness claim of NeuCheck. Also, as discussed for ZEUS, `call` is not the only way of invoking another contract, what reveals another source of unsoundness in this definition. Furthermore, neither the security properties that the tool aims for are specified nor any justifications for the soundness of this syntactic analysis approach are provided.

5 How to Implement a Practical, Sound Static Analysis?

After exposing the problems that can arise when designing a practical, sound static analysis, we discuss how we tackled them in developing *eThor* and the underlying static analysis specification framework *HoRSt* [30]. We then present the elements we identified as essential for designing an automated sound analysis: A semantic foundation, sound abstractions, and a principled implementation.

5.1 Overview of *eThor*

eThor was preceded by an earlier prototype, called *EtherTrust* [14]. *EtherTrust* implemented the rules of a formal abstract semantics in *JavaTM* and exported them to *z3*. While this design showed promising preliminary results, it turned out to be too inflexible for our purposes: Changes in the abstract semantics had to be tediously translated to *JavaTM* code; the non-declarative manner of specifying rules made them hard to write and review; and the lack of a proper

intermediate representation made it difficult to implement custom optimizations before passing the verification task to $z3$.

These limitations are addressed by *HoRSt* [30], a dedicated high-level language for designing Horn clause based static analyses. *HoRSt* allows for the specification of Horn Clause based semantic rules in an declarative language and can apply different optimizations before translating them to $z3$ formulae. Thus, the semantics specification and the tool implementation are logically separated and systematic experiments with different versions of the semantics are possible. Additionally, optimizations can be implemented independently from specific semantics, improving the overall performance in a robust fashion.

eThor [30] combines *HoRSt* with an abstract EVM semantics, a parser to read EVM bytecodes, and a set of EVM-specific preprocessing steps, including the reconstruction of the control flow and the propagation of constants. It supports general reachability analysis and in particular allows for (soundly) verifying that a contract is single-entrant (following the definition in [15]).

What distinguishes *eThor* from prior work discussed in Sect. 3 is its well defined analysis specification that is supported by rigorous formal soundness proofs, as well as its principle implementation design. Prior works do not come with thorough formalization and proofs what ultimately leads to soundness issues in the analyzers, as we confirmed empirically. In contrast, *eThor* lives up to its theoretical soundness guarantees in an extensive evaluation while still being practical in terms of runtime and precision. In the following, we will discuss in detail the semantic foundations, modular design and implementation of *eThor* as well as its empirical performance evaluation.

5.2 Semantic Foundations

A formal soundness guarantee requires a formal semantics of the system under analysis. Such a semantics might be specified on paper [33] or in an executable fashion [15, 18], but in any case has to be precise enough to unambiguously capture all relevant aspects of the system. While semantics defined in prose tend to be more readable, executable semantics lend themselves to automated testing or tooling (e.g., the generation of interpreters or symbolic debuggers [18]). *eThor* builds on the semantics presented in [15] which consists of a logical specification as well as an executable F^* semantics that was rigorously tested for its compliance with the Ethereum client software.

Using a formal semantics, security properties can be precisely characterized. *eThor* bases its analysis for the absence of reentrancy attacks on the notion of single-entrancy [15]. Single-entrancy captures that the reentering execution of a contract should not initiate any further internal transactions. This property rules out reentrancy attacks and also contributes to the proof strategy for the more general call integrity property as detailed out in [15].

However, an executable semantics combined with precisely defined security properties alone does not yield a useful analysis tool. While these components allow experts to semi-automatically verify contracts (using frameworks such as [6, 15, 18, 19]), automation generally requires abstractions to be feasible.

5.3 Sound Abstractions

A first step to reduce the complexity of the analysis problem, and hence to make it amenable to automation, is to over-approximate the target property. For *eThor*, we over-approximate the single-entrancy property by the simpler *call unreachability* [14] property. Call unreachability breaks down single-entrancy to a simple criterion on the execution states of a single contract, as opposed to reasoning about the structure and evolution of whole call stacks. Such over-approximations have to be proven sound – every program fulfilling the over-approximated property also has to fulfill the original property. A corresponding proof for single-entrancy is conducted in [14].

To further simplify the analysis task, the relevant parts of a contract’s execution behavior need to be abstracted in a sound manner. In *eThor* for this purpose we devised an abstract semantics based on Horn clauses that we proved in [30] to soundly over-approximate the small step semantics in [15]. The abstract semantics simplifies and summarizes complex execution scenarios that may emerge due to the uncertain blockchain environment, as we exemplify in the following.

In the largely unknown blockchain environment it is infeasible to track constraints on all unknown values. Instead, following a standard technique in abstract interpretation, we enriched our domain of concrete computation values with a new value \top , signifying *all* possible values. This designated symbol over-approximate under-specified values while dropping constraints on them. Some computations, such as the SHA-3-computations or unaligned memory accesses in EVM, are due to their complexity over-approximated by \top in *eThor*.

Further, we abstract the initial invocation of a contract and its reentering executions as they might be scheduled by (unknown) contracts which are called during execution. In *eThor* we ignore the call stack up until the first execution of the analyzed contract, and assume a contract to be called in an arbitrary environment instead. Also, we only distinguish between the first execution of a contract under analysis in a call chain and a reentering execution of such a contract. In this way we collapse all reentering executions while modeling relevant storage invariants with a sophisticated domain-specific abstraction.

For an extended discussion of the abstractions used in *eThor*, including those for inter-contract calls, gas treatment, and memory layout, we refer to [30].

In summary, *eThor* provides a reliable soundness guarantee for the single-entrancy property, proving that a contract labeled secure by *eThor* satisfies single-entrancy. This guarantee stems from the soundness of the abstract semantics with respect to the rigorously tested small step semantics and from the proof that call unreachability (formulated in terms of the abstract semantics) soundly approximates single-entrancy. The soundness of the abstract semantics further enables the sound verification of arbitrary reachability properties that are expressed in terms of the abstract semantics. In particular this holds for functional contract properties phrased as pre- and postconditions: *eThor* can prove that a contract starting in a state satisfying the precondition is never able to reach a state that does not satisfy the postcondition.

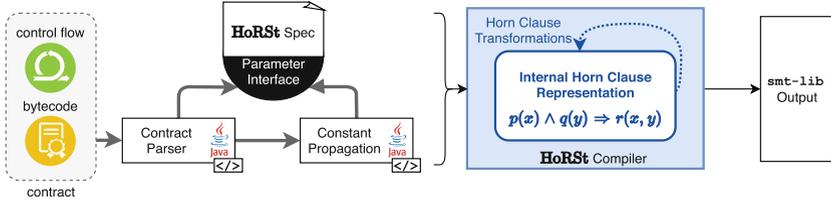


Fig. 5. Architecture of *eThor*

5.4 Implementation Strategies

To arrive at a fast and stable analysis implementation, the analysis task is usually reduced to a known problem supported by performant and well-maintained solvers. This does not only save implementation time and help performance, but it also adds an abstraction layer that facilitates reasoning. For *HoRSt* we decided to use *z3*, respectively Constrained Horn Clauses over *smt-lib*'s linear integer arithmetic fragment, as translation target. We chose *z3* since it is a state-of-the-art solver for program analysis and the fragment suffices to formulate reachability properties.

Architecture. In *eThor*, the generation of *smt-lib* code that models the abstract semantics of a contract is structured into separate and well-defined phases. As can be seen in Fig. 5, the input of *eThor* consists of a contract with reconstructed control flow. The bytecode of the contract is then parsed and constants are propagated within basic blocks. With this information, the abstract semantics (provided as a *HoRSt* specification) is instantiated to a set of Horn clauses which are, after several transformation steps, translated to *smt-lib* formulae.

Optimizations. The performant analysis of real-world programs might require the usage of different optimizations, such as leveraging domain-specific knowledge in a pre-processing step. Such preprocessing may include propagation of constants [30, 31], reconstruction of the control flow [30, 31], computation of memory offsets [31], and pruning of irrelevant parts of the input [30, 31].

As mentioned in Sect. 4.1, unsoundness introduced in any optimization or pre-processing step (e.g., by using an unsound decompiler) immediately affects the soundness of the whole analysis. It is hence crucial to formally reason over each step. In *eThor* the control flow graph reconstruction of a smart contract is realized by symbolically computing the destinations of all jump instructions based on a simplified version of the sound abstract semantics used for the later reachability analysis. Therefore, all soundness considerations from the full abstract semantics carry over to the preanalysis. Since this version of the semantics falls into the datalog solvable fragment as implemented by the *Soufflé* solver, we encoded this simple abstract semantics as a *Soufflé* program. To automate the generation of such preprocessing steps in the future we plan to extend *HoRSt* with *Soufflé* as additional compilation target.

Evaluation. To ensure the correctness and performance of an analysis tool, it is inevitable to extensively and systematically test the tool implementation. To this end synthetic, well-understood inputs help to identify problems regarding precision, performance, and correctness early. These, however, may not be representative of the challenges that are found in real-world contracts. Data gathered from a real-world setting, on the other hand, might be difficult to classify manually (i.e. check for presence or absence of properties), making it difficult to check for correctness of the implementation, and may overtax earlier, non-optimized iterations of the analysis tool. be of uncertain ground truth or too complex to give guidance in early stages of the development. In our experience, an automated test suite with corpus of synthetic and real-world inputs is a significant help while experimenting with different formulations and optimizations, as implementation bugs can be found already at an early stage.

For *eThor* we leveraged the official EVM test suite and our own property-based test suite for assessing the correctness of the abstract semantics and abstract properties. Out of 604 relevant EVM test cases, we terminated on 99%. All tests confirmed the tool’s soundness and the possibility of specifying the test suite within *eThor* confirmed the versatility of our approach beyond reentrancy.

The correctness and precision of *eThor* for the single-entrancy property were assessed on a benchmark of 712 real-world contracts. Within a 10 min timeout, *eThor* delivered results for 95% of the contracts, with all of them confirming soundness, and yielding a specificity of 80%, resulting in an F-measure of 89%. These results do not only demonstrate *eThor*’s practicability on real-world contracts, but also clearly improve over the state-of-the-art analyzer ZEUS. When run on the same benchmark, ZEUS yields a specificity of only 11.4% (challenging its soundness claim) and a specificity of 99.8%, giving an F-measure of 20.4%⁷.

6 Future Challenges

To bring forward the robust design and implementation of sound static analyzers, we plan on extending *HoRSt* in multiple ways: We want to integrate *HoRSt* with proof assistants in order to streamline and partially automate soundness proofs. Further, we want to add support for additional compilation targets, and enrich the specification language and compilation to go beyond reachability analysis, and to support restricted classes of hyperproperties.

For the particular case of *eThor*, we want to improve the precision of the analysis, e.g., to include a symbolic treatment of hash values, and to enable the joint verification of multiple interacting contracts. Further, we strive to create a public benchmark of smart contracts exhibiting different security vulnerabilities, as well as mitigations. This would enable the community to systematically compare the performance, correctness, and precision of different tools.

⁷ *eThor* was evaluated against ZEUS since this is the only tool to implement a property similar to single-entrancy.

Beyond that, we plan to transfer the presented techniques to other smart contract platforms, such as Libra, EOS, or Hyperledger Fabric, which exhibit domain-specific security properties and different semantics.

Acknowledgements. This work has been partially supported by the European Research Council (ERC) under the European Union’s Horizon 2020 research (grant agreement 771527-BROWSEC); by the Austrian Science Fund (FWF) through the projects PROFET (grant agreement P31621) and the project W1255-N23; by the Austrian Research Promotion Agency (FFG) through the Bridge-1 project PR4DLT (grant agreement 13808694) and the COMET K1 SBA; and by the Internet Foundation Austria (IPA) through the netidee project EtherTrust (Call 12, project 2158).

References

1. The DAO smart contract (2016). <http://etherscan.io/address/0xbb9bc244d798123fde783fcc1c72d3bb8c189413#code>
2. The parity wallet breach (2017). <https://www.coindesk.com/30-million-ether-reported-stolen-parity-wallet-breach/>
3. The parity wallet vulnerability (2017). <https://paritytech.io/blog/security-alert.html>
4. Solidity (2019). <https://solidity.readthedocs.io/>
5. Adhikari, C.: Secure framework for healthcare data management using ethereum-based blockchain technology (2017)
6. Amani, S., Bégel, M., Bortin, M., Staples, M.: Towards verifying ethereum smart contract bytecode in isabelle/hol. In: Proceedings of the 7th ACM SIGPLAN International Conference on Certified Programs and Proofs, pp. 66–77 (2018)
7. Azaria, A., Ekblaw, A., Vieira, T., Lippman, A.: Medrec: using blockchain for medical data access and permission management. In: International Conference on Open and Big Data (OBD), pp. 25–30. IEEE (2016)
8. Bartoletti, M., Galletta, L., Murgia, M.: A minimal core calculus for solidity contracts. In: Pérez-Solà, C., Navarro-Arribas, G., Biryukov, A., Garcia-Alfaro, J. (eds.) DPM/CBT -2019. LNCS, vol. 11737, pp. 233–243. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-31500-9_15
9. Bhargavan, K., et al.: Formal verification of smart contracts: short paper. In: Proceedings of the 2016 ACM Workshop on Programming Languages and Analysis for Security, pp. 91–96 (2016)
10. Crafa, S., Di Pirro, M., Zucca, E.: Is solidity solid enough? In: Bracciali, A., Clark, J., Pintore, F., Rønne, P.B., Sala, M. (eds.) FC 2019. LNCS, vol. 11599, pp. 138–153. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-43725-1_11
11. Cruz, J.P., Kaji, Y., Yanai, N.: RBAC-SC: role-based access control using smart contract. IEEE Access **6**, 12240–12251 (2018)
12. de Moura, L., Bjørner, N.: Z3: an efficient SMT solver. In: Ramakrishnan, C.R., Rehof, J. (eds.) TACAS 2008. LNCS, vol. 4963, pp. 337–340. Springer, Heidelberg (2008). https://doi.org/10.1007/978-3-540-78800-3_24
13. Galal, H.S., Youssef, A.M.: Verifiable sealed-bid auction on the ethereum blockchain. In: Zohar, A., et al. (eds.) FC 2018. LNCS, vol. 10958, pp. 265–278. Springer, Heidelberg (2019). https://doi.org/10.1007/978-3-662-58820-8_18

14. Grishchenko, I., Maffei, M., Schneidewind, C.: Foundations and tools for the static analysis of ethereum smart contracts. In: Chockler, H., Weissenbacher, G. (eds.) CAV 2018. LNCS, vol. 10981, pp. 51–78. Springer, Cham (2018). https://doi.org/10.1007/978-3-319-96145-3_4
15. Grishchenko, I., Maffei, M., Schneidewind, C.: A semantic framework for the security analysis of ethereum smart contracts. In: Bauer, L., Küsters, R. (eds.) POST 2018. LNCS, vol. 10804, pp. 243–269. Springer, Cham (2018). https://doi.org/10.1007/978-3-319-89722-6_10
16. Grossman, S., et al.: Online detection of effectively callback free objects with applications to smart contracts. Proc. ACM Program. Lang. **2**(POPL), 1–28 (2017)
17. Hahn, A., Singh, R., Liu, C.C., Chen, S.: Smart contract-based campus demonstration of decentralized transactive energy auctions. In: 2017 IEEE Power & Energy Society Innovative Smart Grid Technologies Conference (ISGT), pp. 1–5. IEEE (2017)
18. Hildenbrandt, E., et al.: KEVM: a complete formal semantics of the ethereum virtual machine, pp. 204–217. IEEE (2018). <https://doi.org/10.1109/CSF.2018.00022>. <https://ieeexplore.ieee.org/document/8429306/>
19. Hirai, Y.: Defining the ethereum virtual machine for interactive theorem provers. In: Brenner, M., et al. (eds.) FC 2017. LNCS, vol. 10323, pp. 520–535. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-70278-0_33
20. Jiao, J., Kan, S., Lin, S.W., Sanan, D., Liu, Y., Sun, J.: Executable operational semantics of solidity. arXiv preprint [arXiv:1804.01295](https://arxiv.org/abs/1804.01295) (2018)
21. Jordan, H., Scholz, B., Subotić, P.: SOUFFLÉ: on synthesis of program analyzers. In: Chaudhuri, S., Farzan, A. (eds.) CAV 2016. LNCS, vol. 9780, pp. 422–430. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-41540-6_23
22. Kalra, S., Goel, S., Dhawan, M., Sharma, S.: ZEUS: analyzing safety of smart contracts. Internet Society (2018). <https://doi.org/10.14722/ndss.2018.23082>. https://www.ndss-symposium.org/wp-content/uploads/2018/02/ndss2018.09-1_Kalra_paper.pdf
23. Lu, N., Wang, B., Zhang, Y., Shi, W., Esposito, C.: Neuchek: a more practical ethereum smart contract security analysis tool. Pract. Exp. Softw.(2019)
24. Luu, L., Chu, D.H., Olickel, H., Saxena, P., Hobor, A.: Making smart contracts smarter. In: Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security, pp. 254–269 (2016)
25. Mathieu, F., Mathee, R.: Blocktix: decentralized event hosting and ticket distribution network (2017). <https://blocktix.io/public/doc/blocktix-wp-draft.pdf>
26. McCorry, P., Shahandashti, S.F., Hao, F.: A smart contract for boardroom voting with maximum voter privacy. In: Kiayias, A. (ed.) FC 2017. LNCS, vol. 10322, pp. 357–375. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-70972-7_20
27. Nakamoto, S.: Bitcoin: a peer-to-peer electronic cash system (2008). <http://bitcoin.org/bitcoin.pdf>
28. Notheisen, B., Gödde, M., Weinhardt, C.: Trading stocks on blocks - engineering decentralized markets. In: Maedche, A., vom Brocke, J., Hevner, A. (eds.) DESRIST 2017. LNCS, vol. 10243, pp. 474–478. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-59144-5_34
29. Panescu, A.T., Manta, V.: Smart contracts for research data rights management over the ethereum blockchain network. Sci. Technol. Libr. **37**(3), 235–245 (2018)
30. Schneidewind, C., Grishchenko, I., Scherer, M., Maffei, M.: ethor: practical and provably sound static analysis of ethereum smart contracts. arXiv preprint [arXiv:2005.06227](https://arxiv.org/abs/2005.06227) (2020)

31. Tsankov, P., Dan, A., Drachsler-Cohen, D., Gervais, A., Bünzli, F., Vechev, M.: Securify: practical security analysis of smart contracts, pp. 67–82. ACM (2018). <https://doi.org/10.1145/3243734.3243780>
32. Wood, G.: Ethereum: a secure decentralised generalised transaction ledger. Ethereum Proj. Yellow Paper **151**, 1–32 (2014)
33. Wood, G.: Ethereum: a secure decentralised generalised transaction ledger (2014)
34. Yang, Z., Lei, H.: Lolisa: formal syntax and semantics for a subset of the solidity programming language. arXiv preprint [arXiv:1803.09885](https://arxiv.org/abs/1803.09885) (2018)
35. Zakrzewski, J.: Towards verification of ethereum smart contracts: a formalization of core of solidity. In: Piskac, R., Rümmer, P. (eds.) VSTTE 2018. LNCS, vol. 11294, pp. 229–247. Springer, Cham (2018). https://doi.org/10.1007/978-3-030-03592-1_13