



EVMPatch: Timely and Automated Patching of Ethereum Smart Contracts

Michael Rodler, *University of Duisburg-Essen*; Wenting Li and Ghassan O. Karame,
NEC Laboratories Europe; Lucas Davi, *University of Duisburg-Essen*

<https://www.usenix.org/conference/usenixsecurity21/presentation/rodler>

**This paper is included in the Proceedings of the
30th USENIX Security Symposium.**

August 11-13, 2021

978-1-939133-24-3

**Open access to the Proceedings of the
30th USENIX Security Symposium
is sponsored by USENIX.**

EVMPatch: Timely and Automated Patching of Ethereum Smart Contracts

Michael Rodler
University of Duisburg-Essen

Wenting Li
NEC Laboratories Europe

Ghassan O. Karame
NEC Laboratories Europe

Lucas Davi
University of Duisburg-Essen

Abstract

Recent attacks exploiting errors in smart contract code had devastating consequences thereby questioning the benefits of this technology. It is currently highly challenging to fix errors and deploy a patched contract in time. Instant patching is especially important since smart contracts are always online due to the distributed nature of blockchain systems. They also manage considerable amounts of assets, which are at risk and often beyond recovery after an attack. Existing solutions to upgrade smart contracts depend on manual and error-prone processes. This paper presents a framework, called EVMPATCH, to instantly and automatically patch faulty smart contracts. EVMPATCH features a bytecode rewriting engine for the popular Ethereum blockchain, and transparently/automatically rewrites common off-the-shelf contracts to upgradable contracts. The proof-of-concept implementation of EVMPATCH automatically hardens smart contracts that are vulnerable to integer over/underflows and access control errors, but can be easily extended to cover more bug classes. Our evaluation on 14,000 real-world contracts demonstrates that our approach successfully blocks attack transactions launched on contracts, while keeping the intended functionality of the contract intact. We perform a study with experienced software developers, showing that EVMPATCH is practical, and reduces the time for converting a given Solidity smart contract to an upgradable contract by 97.6%, while ensuring functional equivalence to the original contract.

1 Introduction

Smart contracts are used in modern blockchain systems to allow nearly arbitrary (Turing-complete) business logic to be implemented. They enable autonomous management of cryptocurrency or tokens and have the potential to revolutionize many business applications by removing the need for a trusted (potentially malicious) third party, e.g., in applications for payments, insurances, crowd funding, or supply chains. Due to their ease of use and the high monetary value (cryptocurrency)

some of these contracts hold, smart contracts have become an appealing target for attacks. Programming errors in smart contract code can have devastating consequences as an attacker can exploit these bugs to steal cryptocurrency or tokens.

Recently, the blockchain community has witnessed several incidents due smart contract errors [7, 39]. One especially infamous incident is the “TheDAO” reentrancy attack, which resulted in a loss of over 50 million US Dollars worth of Ether [31]. This led to a highly debated hard-fork of the Ethereum blockchain. Several proposals demonstrated how to defend against reentrancy vulnerabilities either by means of offline analysis at development time or by performing run-time validation [16, 23, 32, 42]. Another infamous incident is the parity wallet attack [39]. In this case, an attacker moved a smart contract into a state, where the currency held by the contract could not be accessed anymore. This resulted in a total of about 500,000 Ether to be stuck in smart contracts due to an access control error [38]. Automatic detection of such access control vulnerabilities has been previously studied in the context of automated exploit generation [20, 28]. Further, integer overflow bugs constitute a major vulnerability class in smart contracts. Such bugs occur when the result of an arithmetic operation has a longer width than the integer type can hold [34]. According to a study by Torres et al. [13] more than 42,000 contracts suffer from an integer bug. They especially affect so-called ERC-20 Token contracts, which are leveraged in Ethereum to create subcurrencies. Interestingly, several of the disclosed vulnerabilities were actually exploited leading to substantial token and ether losses.

These attacks have fueled interest in the community to enhance the security of smart contracts. In this respect, a number of solutions ranging from devising better development environments to using safer programming languages, formal verification, symbolic execution, and dynamic runtime analysis have been proposed in the last few years [19, 23, 32]. We point out that all these solutions only aim to prove the correctness or absence of a certain type of vulnerability [19, 23, 42] and as such cannot be used to protect already deployed (legacy) contracts. Although some contracts integrate upgrade

mechanisms (see § 2), once a particular contract has been flagged as vulnerable, it is unclear how to automatically patch it and test the effectiveness of the patched contract. Even though manually patching contracts on the source-code level seems plausible, the patch may unexpectedly break compatibility and make the upgraded contracts unusable. For example, given the special storage layout design of Ethereum, the delegatecall-proxy pattern requires developers to ensure that the patched version of the contract is *compatible* with the previously deployed version. Even small changes like changing the ordering of variables in the source code can break this compatibility. This additionally poses the challenge that developers must adhere to strict coding standards [46] and have to use the same exact compiler version. As a result, patching smart contract errors is currently a *time-consuming, cumbersome, and error-prone process*. For instance, while patching the Parity multisig wallet contract, a vulnerability was introduced. An attacker was able to become the owner of the newly deployed library contract. This allowed the attacker to destroy the contract and break all contracts that depend on the multisig wallet library contract. As a result, a considerable amount of Ether is now locked in those broken contracts [38]. On top of that, patching smart contract bugs is highly *time-critical*. In contrast to errors discovered in PC or mobile software, smart contract errors are unique from an attacker’s point of view as (1) *smart contracts are always online on the blockchain*, (2) *they usually hold a significant amount of assets*, and (3) *an attacker does not need to consider other environmental variables* (e.g., software and library version, network traffic analysis, spam or phishing mails to trigger the exploit through a user action).

Contributions. In this paper, we address the problem of automated and timely patching of smart contracts to aid developers to instantly take action on reported smart contract errors. We introduce a novel patching framework (in § 3) that features a *bytecode-rewriter* for Ethereum smart contracts, is independent of the source programming language and works on unmodified contract code. Our framework, dubbed EVMPATCH, utilizes the bytecode-rewriting engine to ensure that patches are minimally intrusive and that the newly patched contract is compatible with the original contract. In particular, our framework automatically *replays transactions* on the patched contract to

1. test the functional correctness of the patched contract with respect to previous transactions pertaining to the contract,
2. identify potential attacks, i.e., developers can determine whether their vulnerable contract has been attacked in the past.

EVMPATCH uses a best effort approach to ensure the introduced patch does not break functionality by testing with previously issued transactions to the contract and optionally also developer provided unit tests. While such a differential testing approach cannot provide a formal proof on the correctness of the patched contract, it works without requiring a formal specification. Our experiments (see § 5.2.1) show that this approach is sufficient in practice to identify broken patches.

By applying patches on the bytecode level, EVMPATCH is independent of the used programming language/compiler and compiler version. That is, EVMPATCH supports any off-the-shelf Ethereum smart contract code. We employ bytecode writing to ensure minimally intrusive patches, that are compatible by design with the contract’s storage layout, We argue that source-level patching is not easily usable in an automated patching process that we propose. However, as for any approach working on either the binary or bytecode-level, we had to tackle several technical challenges (§ 4). Furthermore, EVMPATCH automatically converts the original contract to use the delegatecall-proxy pattern. As such, EVMPATCH is able to automatically deploy newly patched contracts in a fully automated way without requiring any developer intervention.

While in principle EVMPATCH can support patching of different classes of vulnerabilities (see § 4.5), our proof-of-concept implementation targets the two major classes of access control and integer overflow (§ 5) bugs. The latter have been repeatedly exploited in high-value ERC-20 contracts [30], whereas the former has been abused in the Parity wallet attack [39].

To evaluate EVMPATCH in terms of performance, effectiveness, and functional correctness, we apply EVMPATCH to 14,000 real-world vulnerable contracts. To this end, we used the patch testing component of the EVMPATCH framework to re-play all existing transactions to the original contract on the patched contract. This allows us to provide in-depth investigation of several actively exploited smart contracts, e.g., token burning and history of attack transactions (before and after public disclosure). For a number of contracts we investigated in our evaluation, we found that EVMPATCH would have blocked several attacks that happened after public disclosure of the vulnerability. This shows that even though those contracts were officially deprecated, they were still used by legitimate users and exploited by malicious actors. As such, there is an immediate need for tooling, as provided by EVMPATCH, which allows the developers of smart contracts to efficiently patch their contracts. Our evaluation also covers important practical aspects such as gas and performance overhead (i.e., the costs for executing transactions in Ethereum). The gas overhead for all our patched contracts was below 0.01 US\$ per transaction and the performance overhead negligible.

To assess the usefulness of EVMPATCH, we conducted a sophisticated developer study¹ that focuses on comparing the usability of patching and deploying an upgradable contract with and without EVMPATCH (§ 5.3). Our study reveals that developers required 62.5 min (median) to manually (without EVMPATCH) convert a simple smart contract, which implements common Wallet functionality in about 80 lines of code, into an upgradable smart contract. In spite of this considerable time, none of them performed a *correct* conversion, leading to broken and potentially vulnerable contracts. As such, this time

¹ See github.com/uni-due-syssec/evmpatch-developer-study for details

measurements must be seen as a lower bound, as correctly converting a more complex contract will take even more time. In contrast, the same task was performed by the developers using `EVMPATCH` in 1.5 min (median)—a reduction by 97.6%—while producing a correct upgradable contract.

2 Background

In this section, we provide background information on the Ethereum Virtual Machine (EVM), binary rewriting, and some common contract upgrade strategies.

EVM & Smart Contracts: At the core of the Ethereum blockchain system lies a custom virtual machine, dubbed Ethereum Virtual Machine (EVM), which executes Ethereum smart contracts. EVM consists of a simple stack-based virtual machine with a custom instruction format. Every instruction is represented as a one-byte opcode. Arguments are passed on the data stack. The only exception are the push instructions, which are used to push constants onto the stack. These constants are encoded directly into the instruction bytes. Furthermore, the EVM follows the Harvard architecture model and separates code and data into different address spaces. In fact, the EVM features different address spaces for different purposes: the code address space, which contains a smart contract's code and is considered immutable, the storage address space for storing global state, and the memory address space for temporary data.

In the Ethereum network, a smart contract must be executed by every miner and every full node in the network to compute and verify the state before and after a block. Ethereum features a mechanism to limit the execution time per smart contract and reward miners for executing smart contracts: the so-called *gas*. Every EVM instruction requires a certain gas budget to execute. The transaction sender selects the price per gas unit in Ether and when a transaction is included into a block the corresponding Ether is transferred to the miner as a reward. Minimizing the gas required for executing a contract is important as it indirectly minimizes the cost of operating a smart contract in Ethereum.

Smart contracts are developed in an object-oriented fashion, i.e., every smart contract has a defined interface of functions: the contract's ABI (Application Binary Interface). Whenever a smart contract calls another smart contract, it utilizes one of the call instructions, such as `CALL` or `STATICCALL`. The called contract will then process the provided input and update its own state accordingly. A contract cannot directly access the state (i.e., the storage area) of other contracts and must always use function calls according to the ABI to retrieve any data from another contract.

In contrast to the regular `CALL` instruction, the `DELEGATECALL` instruction will execute the called contract's code in the context of the caller contract. This instruction was introduced to implement *library contracts*, i.e., common functionality can be deployed once to the blockchain and multiple contracts can rely on one library contract. This means that the callee, i.e., the library contract, has full access

to the state (the storage) and the Ether funds of the caller. As such, a contract that utilizes a `DELEGATECALL` instruction must fully trust the callee.

Binary Rewriting: Binary rewriting is a well-known technique to instrument programs after compilation. Binary rewriting has also been applied to retrofit security hardening techniques such as *control-flow integrity*, to compiled binaries [8], but also to dynamically apply security patches to running programs [29]. For binary rewriting on traditional architectures two flavors of approaches have been developed: static and dynamic rewriting.

Dynamic approaches [22] rewrite code on-the-fly, i.e., while the code is executing. This avoids imprecise static analysis on large binaries. However, dynamic binary rewriting requires an intermediate layer, which analyzes and rewrites code at runtime. Since the EVM does not support dynamic code generation or modification, it is not possible to apply this approach efficiently in Ethereum. In contrast, static binary rewriting [5, 21] is applicable to Ethereum as it works completely offline. It relies on static analysis to recover enough program information to accurately rewrite the code.

Contract Upgrade Strategies: Ethereum treats the code of smart contracts as immutable once they are deployed on the blockchain². To remedy this, the community came up with strategies for deploying upgraded smart contracts [11, 41, 45]. The most naive approach is to deploy the patched contract at a new address and migrate the state of the original contract to it. However, state migration is specific to the contract and must be manually implemented by the developers of the contract. It requires the contract developers to have access to all the internal state of the old contract, and a procedure in the new contract to accept state transfers. To avoid state migration, developers can also use a separate contract as a data storage contract, which is sometimes referred to as the eternal storage pattern [10, 45]. However, this adds additional gas overhead since every time the logic contract needs to access data it must perform a costly external call into the data storage contract.

A more common strategy is to write contracts with the proxy-pattern, with the most favorable version being the *delegatecall-proxy* pattern. Here, one smart contract is split into two different contracts, one for the code and one for data storage: i) an immutable *proxy contract*, which holds all funds and all internal state, but does not implement any business logic; ii) a *logic contract*, which is completely stateless and implements all of the actual business logic, i.e., this contract contains the actual code that governs the actions of the contract. The proxy contract is the entry point of all user transactions. It has immutable code and its address remains constant over the lifetime of the contract. The logic contract implements the rules, which govern the behavior of the smart contract. The proxy contract forwards all function calls to the registered logic contract using the `DELEGATECALL` instruction. This instruction is used to give

²Except for the selfdestruction mechanism to kill a smart contract.

the logic contract access to all internal state and funds stored in the proxy contract. To upgrade the contract, a new logic contract is deployed and its address is updated in the proxy contract. The proxy contract then forwards all future transactions to the patched logic contract. As a result, deploying upgraded contracts does not require any data migration, as all data is stored in the immutable proxy contract. Moreover, the upgrading process is also *transparent* to users, as the contract address remains the same. Although existing blockchain platforms do not provide mechanisms to upgrade smart contracts, the usage of this proxy pattern allows EVMPATCH to quickly upgrade a contract with negligible costs (in terms of gas consumption).

3 Design of EVMPatch

In this section, we introduce the design of our automated patching framework to timely patch and harden smart contracts. Our framework operates on unmodified smart contracts and is independent of the source code programming language, as it does not require source code. At its core, our framework utilizes a bytecode rewriter to apply minimally intrusive patches to EVM smart contracts. Combined with a proxy-based upgradable smart contract, this bytecode rewriting approach allows the developer to automatically introduce patches and deploy them on the blockchain. One major advantage of this approach is that when new attack types are discovered or bug finding tools improve, the contract can be automatically re-checked, patched, and re-deployed in a short amount of time and with minimal developer intervention. EVMPATCH is typically executed on a developer’s machine and is continuously running new and updated vulnerability detection tools. This can also include dynamic analysis tools, which analyze transactions that are not yet included in a block, but already available to the Ethereum network. Whenever one of the analysis tools discovers a new vulnerability, EVMPATCH automatically patches the contract, tests the patched contract and deploys it.

3.1 Design Choices

The proxy-pattern makes it possible to easily deploy a patched smart contract in Ethereum. However, it neither generates a patched version nor features functional tests on the patched contract. EVMPATCH fills this gap by providing a comprehensive framework and toolchain to automatically and timely patch and test the effectiveness of the generated patch.

As shown in Table 1, there are two possible strategies for automatically generating a patch in Ethereum: static rewriting of source or EVM bytecode. At first glance, *source-code patching* seems to be the option of choice as developers have access to source code, they are able to inspect the source code changes, and can even do adjustments if the automated approach introduces undesired changes. However, in Ethereum, there is one major challenge when applying source code rewriting: one needs to carefully preserve the storage layout. Otherwise, the

patched contract will corrupt its memory and fail or (worse) introduce dangerous bugs. Namely, some changes in the source code can break the contract compatibility, even though the changes do not break the logic of the contract.

To put things into context, statically-sized variables are laid out contiguously in storage starting from address 0; and contiguous variables with size less than 32 B can be packed into a single 32 B storage slot [9]. As a result, any changes to re-order, add, or remove variables in the source-code may look harmless, but on the memory level, such changes will lead to mapping of variables to wrong and unexpected storage addresses. In other words, changes in variable declaration corrupt the internal state of the contract, as the legacy contract and the patched contract have different storage layouts.

In contrast, *bytecode rewriting* does not suffer from this deficiency as many bug classes only require changes on the level of EVM instructions (see §5) avoiding any error-prone storage-layout changes. Another reason to opt for bytecode rewriting are existing smart contract vulnerability detection tools. As of now, the majority of them operate on the EVM level [13, 20, 23, 24, 32] and report their findings on the EVM level. A bytecode rewriting approach can exploit the reports of these analysis tools to directly generate an EVM bytecode-based patch. Finally, if source-code rewriting is utilized, the developer has limited possibilities to perform thorough testing on the effectiveness of the patched contract. In particular, checking the patched contract against old transactions (including transactions that encapsulate attacks) are more feasible on bytecode level. That is, transaction testing naturally would still require analysis on the bytecode level to reverse-engineer the attack transactions and how they fail against the patched contract. Bytecode-rewriting allows developers to directly match the rewritten bytecode instructions to the attack transactions making forensic analysis feasible. Given all these reasons, we decided to opt for bytecode rewriting.

3.2 Framework Design

Our framework depicted in Figure 1 consists of the following major components: (1) the *vulnerability detection* engine consisting of automatic analysis tools and public vulnerability disclosures, (2) *bytecode rewriter* to apply the patch to the contract, (3) the *patch testing* mechanism to validate the patch

Table 1: Comparison of rewriting strategies in Ethereum

Source Rewriting	Bytecode Rewriting
Corrupts storage-layout	Preserves storage layout
Checking modifications by human analyst feasible	Human analysis of bytecode changes challenging
Limited tool support for vulnerability analysis	Easy integration of vulnerability analysis tools
Patch testing based on prior transactions challenging	Easy patch testing with prior transactions

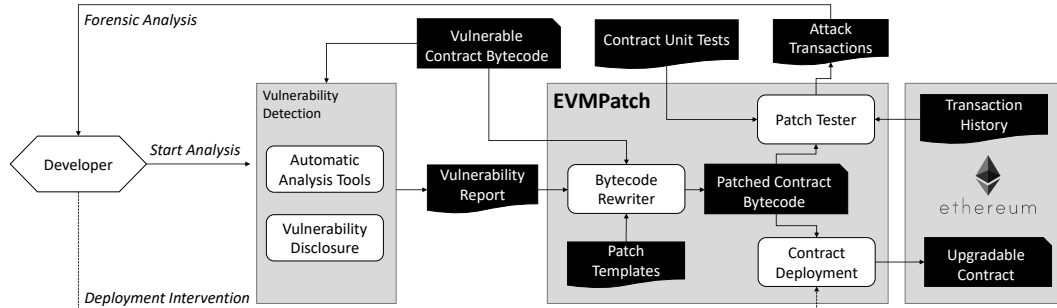


Figure 1: Architecture of EVMPATCH

on previous transactions, and (4) the *contract deployment* component to upload the patched version of the contract. At first, the vulnerability detection engine identifies the location and type of the vulnerability. This information is then passed to the bytecode rewriter, which patches the contract according to previously defined patch templates. The patched contract is thereafter forwarded to the patch tester, which replays all past transactions to the contract. That said, we do not only patch the contract, but we allow the developer to retrieve a list of transactions that exhibit a different behavior and outcome between the original and patched contract. These transactions serve as an indicator for potential attacks on the original contract. If the list is empty, our framework automatically deploys the patched contract instantly on the Ethereum blockchain. Next, we will provide a more detailed description of the four major components of our design.

Vulnerability Detection. Before being able to apply patches, our framework needs to identify and detect vulnerabilities. To do this, our framework leverages existing vulnerability detection tools such as [13, 16, 20, 23, 28, 32, 42]. For vulnerabilities that are not detected by any existing tool, we require that a developer or a security consultant creates a vulnerability report. In our system, the vulnerability detection component is responsible to identify the exact address of the instruction, where the vulnerability is located, and the type of vulnerability. This information is then passed to the bytecode rewriter, which patches the contract accordingly.

Bytecode Rewriter. In general, static binary rewriting techniques are well suited for applying patches in Ethereum since smart contracts have comparably small code size: typically in the range of about 10 KiB. Furthermore, EVM smart contracts are always statically linked to all library code. It is not possible for a contract to dynamically introduce new code into the code address space. This makes the reliance on binary rewriting techniques simpler compared to traditional architectures, where dynamically linked libraries are loaded at runtime. However, some smart contracts still utilize a concept similar to dynamically linked libraries: dedicated EVM call instructions allow a contract to switch to a different code address space. We tackle this peculiarity by applying our bytecode rewriter to both the contract itself and the library contract.

The stack-based architecture of the EVM requires special attention when implementing a patch: all address-based references to any code or data in the code address space of the smart contract must be either preserved or updated when new code is inserted into the code address space. Such references cannot be easily recovered from the bytecode. To tackle this challenge, EVMPATCH utilizes a trampoline-based approach for adding new EVM instructions into empty code areas. The implementation details will be described in § 4.

To implement a patch, the bytecode rewriter processes the bytecode of the vulnerable contract as well as the vulnerability report. The rewriting is based on a so-called patch template which is selected according to the vulnerability type and adjusted to work with the given contract.

Patch Templates. In EVMPATCH, we utilize a template-based patching approach: for every supported class of vulnerabilities, a patch-template is integrated into EVMPATCH. This patch template is automatically adapted to the contract that is being patched. We create generic patch templates such that they can be easily applied to all contracts. EVMPATCH *automatically* adapts the patch template to the contract at hand by replacing contract-specific constants (i.e., code addresses, function identifier, storage addresses). Patch templates for common vulnerabilities, such as integer overflows, are shipped as part of EVMPATCH, and a typical user of EVMPATCH will never interact with the patch templates. However, optionally, a smart contract developer can also inspect or adapt existing patch templates or even create additional patch templates for vulnerabilities that are not yet supported by EVMPATCH.

Patch Tester. As smart contracts directly handle assets (such as Ether), it is critical that any patching process does not impede the actual functionality of a contract. As such, any patch must be tested thoroughly. To address this issue, we introduce a patch testing mechanism which is based (1) on the transaction history recorded on the blockchain and (2) optional developer supplied unit tests. At this point, we exploit the fact that any blockchain system records all previous executions of a smart contract, i.e., transactions in Ethereum. In our case, the patch tester re-executes all existing transactions and optionally any available unit test and verifies that all transactions of the old legacy and the newly patched contract behave consistently.

The patch tester detects any behavioral discrepancy between the old legacy and the newly patched contract and reports a list of transactions with differing behavior to the developer. That said, as a by-product, our patch testing mechanism can be used as a forensic attack detection tool. Namely, while executing the patching process, the developer will also be notified of any prior attacks that abuse any of the patched vulnerabilities and can then act accordingly. In case both versions of the contract behave the same way, the patched contract can be automatically deployed. Otherwise, the developer must investigate the list of suspicious transactions and thereafter invoke the contract deployment component to upload the patched contract. The list of suspicious transactions may not only serve as an indicator of potential attacks, but may reveal that the patched contract is not functionally correct, i.e., the patched contract shows a different behavior on benign transaction. In § 5, we provide a thorough investigation on real-world, vulnerable contracts to demonstrate that EVMPATCH successfully applies patches without breaking the original functionality of the contract.

Contract Deployment. As discussed in § 2, the delegatecall-proxy based upgrade scheme is the option of choice to enable instant contract patching. Thus, EVMPATCH integrates this deployment approach utilizing a proxy contract as the primary entry point for all transactions with a constant address. Before the first deployment, EVMPATCH transforms the original unmodified contract code to utilize the delegatecall-proxy pattern. This is done by deploying a proxy contract, which is immutable and assumed to be implemented correctly³. The original bytecode is then converted to a logic contract using the bytecode rewriter with only minor changes to the original code. The logic contract is then deployed alongside the proxy contract.

Patch Deployment. Finally, when the contract is patched and after the patch is tested by the patch tester component, EVMPATCH can deploy the newly patched contract. Our upgrade scheme deploys the newly patched contract code to a new address and issues a dedicated transaction to the previously deployed proxy contract, which switches the address of the logic contract from the old vulnerable version to the newly patched version. Any further transactions are now handled by the patched logic contract.

Human Intervention. EVMPATCH is designed to be fully automated. However, there are a few scenarios, where developer intervention is needed if (1) the vulnerability report relates to a bug class that is not yet supported by EVMPATCH, or (2) the patch tester reports at least one transaction that fails due to the newly introduced patch and the failing transaction is not a known attack transaction, (3) the patch tester reports that at least one known attack transaction is not prevented by the newly introduced patch.

If a bug class is not supported, EVMPATCH informs the developer about the unsupported vulnerability class. Since EVMPATCH is extensible, it easily allows developers to pro-

³EVMPATCH comes with a well audited default proxy contract that is only 80 lines of Solidity code.

vide custom patch templates thereby allowing quick adaption to new attacks against smart contracts. More specifically, EVMPATCH supports multiple formats for custom patch templates: EVM instructions, a simple domain-specific language that resembles Solidity expressions and allows developers to enforce pre-conditions on functions (similar to Solidity modifiers). We performed a developer study in Section 5.3 to demonstrate that writing a patch template is feasible and more successful than manually patching a contract.

If the patch tester finds a new failing transaction, the developer has to analyze whether a new attack transaction has been discovered or a legitimate transaction has failed. For a newly discovered attack transaction, EVMPATCH adds this transaction to the list of attacks and proceeds. Otherwise, the developer investigates why the legitimate transaction failed. As our evaluation in § 5.2.2 shows, such cases typically occur due to inaccurate vulnerability reports, i.e., wrongly reported vulnerabilities rather than faulty patching. Thus, the developer can simply blacklist the wrongly reported vulnerable code locations to avoid patching at these locations.

These manual interventions typically only need quick code reviews or debugger sessions. We believe even moderately experienced Solidity developers can perform these tasks as no detailed knowledge about the underlying bytecode rewriting system is needed (see also § 5.3 on our developer study). As such, EVMPATCH positions itself as a tool to enable more developers to securely program and operate Ethereum smart contracts.

4 EVMPatch Implementation

In this section, we describe the implementation of EVMPATCH: in § 4.1, we discuss engineering challenges for bytecode rewriting in Ethereum. Thereafter, we describe the implementation of the bytecode rewriter (§ 4.2), the patch testing feature (§ 4.3), and the contract deployment mechanism (§ 4.4). We conclude this section with a discussion on possible applications regarding smart contract errors in § 4.5.

4.1 Challenges of Bytecode Rewriting

There are several unique challenges that must be solved when rewriting EVM bytecode: we need to handle static analysis of the original EVM bytecode, and tackle several particularities of Solidity contracts and the EVM.

Similar to traditional computer architectures, EVM bytecode uses addresses to reference code and data constants in the code address space. Hence, when modifying the bytecode, the rewriter must ensure that address-based references are correctly adjusted. To do so, a rewriter typically employ two static analysis techniques: control-flow graph (CFG) recovery and subsequent data-flow analysis. The latter is necessary to determine which instructions are the sources of any address constants utilized in the code. For the EVM bytecode, two

classes of instructions are relevant in this context: code jumps and constant data references.

Code Jumps. The EVM features two branch instructions: `JUMP` and `JUMPI`. Both take the destination address from the stack. Note that function calls inside the same contract also leverage `JUMP` and `JUMPI`. That said, there is *no explicit difference* between local jumps inside a function and calls to other functions. The EVM also features dedicated call instructions, but these are only used to transfer control to a completely separate contract. Hence, they do not require modification when rewriting the bytecode.

Constant Data References. The so-called `CODECOPY` instruction is leveraged to copy data from the code address space into the memory address space. A common example use-case are large data constants such as strings. Similar to the jump instructions, the address from which memory is loaded is passed to the `CODECOPY` instruction via the stack.

Handling both types of instructions is challenging due to the stack-based architecture of the EVM. For instance, the target addresses of jump instructions are always provided on the stack. That is, every branch is indirect, i.e., the target address cannot be simply looked up by inspecting the jump instruction. Instead, to resolve these indirect jumps, one needs to deploy data-flow analysis techniques to determine where and which target address is pushed on the stack. For the majority of these jumps, one can analyze the surrounding basic block⁴ to trace back where the jump target is pushed on the stack. For example, when observing the instructions `PUSH2 0xdb1; JUMP`, we can recover the jump target by retrieving the address (0xdb1) from the push instruction.

However, many contracts contain more complicated code patterns, primarily because the Solidity compiler also supports calling functions internally without utilizing a call instruction. Recall that, in the EVM, a call instructions perform similarly to remote-procedure calls. To optimize code size and facilitate code re-use, the Solidity compiler introduced a concept where functions are marked as *internal*. These functions cannot be called by other contracts (private to the contract) and follow a different calling convention. Since there are no dedicated return and call instruction for internal functions, Solidity utilizes the jump instruction to emulate both. As such, a function return and a normal jump cannot be easily distinguished. This makes it challenging to (1) identify internal functions and (2) build an accurate control-flow graph of the contract.

When rewriting an EVM smart contract, both the jump instructions and the codecopy instruction need to be considered in the bytecode rewriter. The obvious strategy to rewrite smart contracts is to fix-up all constant addresses in the code to reflect the new addresses after inserting new instructions or removing old instructions. However, this strategy is challenging because it requires accurate control-flow graph recovery and data-flow

⁴A basic block is sequence of EVM instructions that terminate in a branch. The branch connects one basic block to subsequent basic blocks in the CFG of the EVM code.

analysis, which needs to deal with particularities of EVM code, such as internal function calls. In the research area of binary rewriting of traditional architectures, a more pragmatic approach has been developed: the so-called trampoline concept [8, 21]. We utilize this approach in our rewriter and avoid adjusting addresses. Whenever our rewriter must perform changes to a basic block, e.g., inserting instructions, our rewriter replaces the basic block with a trampoline that immediately jumps to the patched copy. Hence, any jump target in the original code stays the same and all data constants are kept at their original addresses. We describe this process in more detail in the subsequent section.

4.2 Bytecode Rewriter Implementation

We implemented a trampoline-based rewriter in Python and utilize the `pyevmasm`⁵ library for disassembling and assembling raw EVM opcodes. Our trampoline-based bytecode rewriter works on the basic block level. When an instruction needs to be instrumented, the whole basic block is copied to the end of the contract. The patch is then applied to this new copy. The original basic block is replaced with a trampoline, i.e., a short instruction sequence that immediately jumps to the copied basic block. Whenever the contract jumps to the basic block at its original address, the trampoline is invoked redirecting execution to the patched basic block by means of a jump instruction. To resume execution, the final instruction of the instrumented basic block issues a jump back into the original contract code. While the trampoline-based approach avoids fixing up any references, it introduces additional jump instructions. However, as we will show, the gas cost associated with these additional jumps is negligible in practice (see § 5).

To ensure correct execution, we must still compute at least a partial control-flow graph, starting from the patched basic block. This is necessary to recover the boundaries of the basic blocks that are patched and the following basic blocks that are connected by a so called fall-through edge. Not all basic blocks terminate with an explicit control-flow instruction: Whenever a basic block ends with a conditional jump instruction (`JUMPI`) or simply does not end with a control-flow instruction, there is an implicit edge (i.e., fall-through) in the control-flow graph to the instruction at the following address.

Handling Fall-Through Edge. To handle the fall-through edge, two cases must be considered. When the basic block targeted by the fall-through edge starts with a `JUMPDEST` instruction, the basic block is marked as a legitimate target for regular jumps in the EVM. In this case, we can append an explicit jump to the rewritten basic block at the end of the contract and ensure that execution continues at the beginning of the following basic block in the original contract code. In case that the following basic block does not begin with a `JUMPDEST` instruction, the EVM forbids explicit jumps to this address. In the control-flow graph, this means that this basic block can

⁵github.com/crytic/pyevmasm

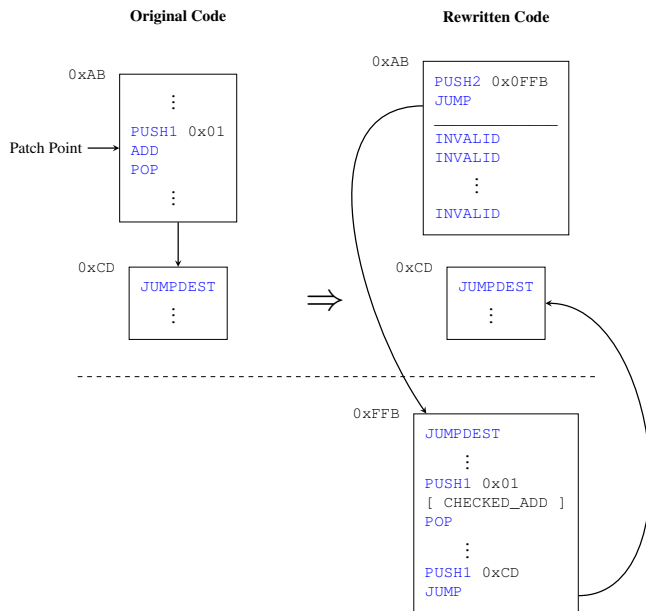


Figure 2: Control-flow graph of original and rewritten code.

only be reached with a fall-through edge. To handle this case, our rewriter copies the basic block to the end of the contract right behind the rewritten basic block constructing another fall-through edge in the control-flow graph of the rewritten code.

Figure 2 shows an example for how our rewriter changes the control-flow graph of the original contract. The `ADD` instruction is replaced with a checked add routine that additionally performs integer overflow checks. We call the address of the `ADD` instruction the *patch point*. The basic block, which contains the patch point, is replaced with a trampoline. In this case, it immediately jumps to the basic block at `0xFFB`. This basic block, which is placed at the end of the original contract, is a copy of the original basic block at `0xAB`, but with the patch applied. Since the basic block is now at the end of the contract, the bytecode rewriter can insert, change, and remove instructions in the basic block without changing any address in code that is located at higher-numbered addresses. We fill the rest of the original basic block with the `INVALID` instruction to ensure the basic block has the exact same size as the original basic block. The basic block at `0xCD` is connected to the prior basic block by means of a fall-through edge. However, this basic block starts with a `JUMPDEST` instruction and as such is a legitimate jump target. Hence, the rewriter then appends a jump to the patched basic block at `0xFFB` which ensures execution continues in the original contract’s code at address `0xCD`.

Adapting to EVM. The EVM has some particularities that must be considered when implementing a bytecode rewriter. Namely, the EVM enforces some separation of code and data in the code address space. EVM implementations prevent jumps into the data constants that are embedded into `PUSH` instructions. The constant operands of the push instructions follow directly after the byte of the push instruction opcode.

Such a constant operand can accidentally include the byte for the `JUMPDEST` instruction. Then, the constant would be a legitimate jump target and a new unintended instruction sequence would occur. To avoid such unintended instruction sequences, EVM implementations perform a linear sweep over the code section to find all push instructions. The constants that are part of those push instructions are then marked as data and therefore as invalid jump targets, even if they contain a byte equivalent to the `JUMPDEST` instruction. However, due to performance reasons, EVM implementations ignore control-flow information when marking data. As such, the push instructions opcode byte itself can be part of some data constant, such as a string or other binary data. For this reason, smart contract compilers accumulate all data constants at addresses strictly larger than any reachable code, avoiding any conflicts between the generated code and data encoded into the code address space. However, our trampoline-based rewriter does append code behind the data constants of the smart contracts. To avoid that code appended by the rewriter is accidentally marked as an invalid jump destination due to a preceding push opcode byte, we carefully insert padding between the data of the original contract and the newly appended code.

Applicability of Trampoline Approach. The trampoline-based approach to rewriting requires only minimal code analysis and works for most use cases. However, this approach faces two problems. First, instructions can only be patched in basic blocks that are large enough (in terms of size in bytes) to also contain the trampoline code. However, a typical trampoline requires 4 to 5 bytes and typically basic blocks that perform some meaningful computation are large enough to contain the trampoline code. Second, due to the copying of basic blocks the code size increases depending on the basic block that is patched thereby increasing deployment cost. However, our experiments show that the overhead during deployment is negligible (on average US\$0.02 per deployment, see § 5).

No reliance on accurate control-flow graph. Recovering an accurate control-flow graph given only EVM bytecode is a challenging and open problem. However, our trampoline based approach does not require an accurate and complete control-flow graph. Instead, we only need to recover basic block boundaries given the program counter of the instruction, where the patch needs to be applied. In doing so, recovering the basic block boundaries is tractable, since the EVM has an explicit marker for basic block entries (i.e., the `JUMPDEST` pseudo-instruction). Furthermore, our rewriter only needs to recover the end of the basic block and any following basic blocks that are connected via fallthrough edges in the control-flow graph.

4.3 Patch Testing

While the insertion of trampolines into the original code does not change the functionality of the contract, the patch template itself can perform arbitrary computations and could potentially violate the semantics of the patched contract. To test the

patched contract, EVMPATCH utilizes a differential testing approach. That is, we re-execute all transactions of the contract to determine if the behavior of the original, vulnerable code and the newly, patched code differ. EVMPATCH utilizes past transactions to the contract retrieved directly from the blockchain. If the contract comes with unit tests, EVMPATCH also utilizes the unit tests to test the newly patched contract. This differential testing approach cannot guarantee formal correctness of the contract. Contracts with a low number of available transactions are prone to low test coverage. However, our experiments (see § 5.2.1) show that the differential testing approach works well enough in practice to show that the patches do not break functionality. Given the availability of a formal specification of the contract's functionality, EVMPATCH could also leverage a model checker to validate a patched contract more rigorously.

During differential testing, we first retrieve a list of transactions to the vulnerable contract from the blockchain. Second, we re-execute all those transactions and retrieve the execution trace for each transaction. Then, we then re-execute the same transactions, but replace the code of the vulnerable contract with the patched contract code, to obtain the second execution trace. We use a modified Ethereum client, based on the popular *go-ethereum* client⁶, since the original client does not support this functionality. Finally, we compare both execution traces and the patch tester produces a list of transactions, where the behavior differs. If there are no such transactions, then we assume that the patch does not inhibit the functionality of the contract and proceed with deploying the patched contract.

The execution traces of the original and patched contracts are never equal since patching changes control flow and inserts instructions. Hence, we examine only potentially state-changing instructions, i.e., instructions that either write to the storage area (i.e., a `SSTORE`) or transfer execution flow to another contract (e.g., a `CALL` instruction). We then compare the order, parameters, and result of all state-changing instructions and find the first instruction where the two execution traces differ. Currently, we assume that the introduced patches do not result in any new state-changing instructions. This assumption holds for patches that introduce input-validation code and revert when invalid input is passed. However, the trace difference computation can be adapted to become aware of potential state changes that a patch introduces. Reported transactions that fail in the code, which is part of the patch, are marked as potential attack transactions. If the reported transaction failed due to out-of-gas in the patched code, we re-run the same transaction with an increased gas budget. We issue a warning since users will have to account for additional gas cost introduced by the patch. Finally, the developer must examine the reported transactions to decide whether the given list of transactions are legitimate or malicious. As a side-effect, this makes our patch tester an attack detection tool for the vulnerable contract allowing developers to quickly find prior attack transactions.

⁶We utilized version 1.8.27-stable-3e76a291

4.4 Deployment of Patched Contracts

As described in § 3, EVMPATCH utilizes the `delegatecall`-proxy based upgrade pattern to deploy the patched contract. To achieve this, EVMPATCH splits the smart contract to two contracts: a proxy contract and a logic contract. The proxy contract is the primary entry point and stores all data. By default, EVMPATCH utilizes a proxy contract that is shipped with EVMPATCH. However, EVMPATCH can also re-use existing upgradable contracts, such as contracts developed with the ZeppelinOS framework [46]. Users interact with the proxy contract, which is located at a fixed address. To facilitate the upgrade process, the proxy contract also implements functionality to update the address of the logic contract. To prevent malicious upgrades, the proxy contract also stores the address of an owner, who is allowed to issue upgrades. The upgrade then simply consists of sending one transaction to the proxy contract, which will (1) check whether the caller is the owner and (2) update the address of the logic contract.

The proxy contract retrieves the address of the new logic contract from storage and simply forwards all calls to that contract. Internally, the proxy contract utilizes the `DELEGATECALL` instruction to call into the logic contract. This allows the logic contract to gain full access to the storage memory area of the proxy contract thereby allowing access to the persistent data without any additional overhead.

4.5 Possible Applications

The bytecode rewriter takes a patch template, which is specified as short snippet of EVM assembly language. This template is then specialized according to the patched contract and relocated to the end of the patched contract. This template-based approach to patch generation allows to specify multiple generic patches to address whole classes of vulnerabilities. In the following, we list possible vulnerability classes that can immediately benefit from our framework.

Improper access control to critical functions can be patched by just inserting a check at the beginning of a function to verify that the caller is a certain fixed address or equal to some address stored in the contract's state. Detection tools to handle this vulnerability have been investigated in prior work [20, 28].

Mishandled exceptions can occur when the contract uses a low-level call instruction, where the return value is not handled automatically, and the contract does not properly check the return value [23]. This issue can be patched by inserting a generic return-value check after such a call instructions.

Integer bugs are highly likely to occur when dealing with integer arithmetic since Solidity does not utilize checked arithmetic by default. This has resulted in many potentially vulnerable contracts being deployed and some being actively attacked [13, 30]. Given the prevalence of these vulnerabilities, we discuss in the next section how to automatically patch integer overflow bugs using EVMPATCH.

```

1 function initMultiowned(address[] _owners, uint _required)
2     ① internal {
3 // ...
4 function initDaylimit(uint _limit) ① internal {
5 // ...
6 // throw unless the contract is not yet initialized.
7 modifier only_uninitialized { if (m_numOwners > 0) throw; _;}
8
9 function initWallet(address[] _owners, uint _required,
10     uint _daylimit)
11     ② only_uninitialized {
12 // ...

```

Figure 3: Source code of patched Parity Multisig Wallet.

In what follows, we demonstrate the effectiveness of EVMPATCH by applying it to the two major bug classes of access control errors and integer bugs.

5 Evaluation of EVMPATCH

In this section, we report the evaluation results of EVMPATCH in patching two prominent types of bugs: (1) access control bugs, and (2) integer bugs (over-/underflow).

5.1 Patching Access Control Bugs

The Parity MultiSig Wallet is a prominent example for access control errors [3, 39]. This contract implements a wallet that is owned by multiple accounts. Any action taken by the wallet contract must be authorized by at least one of the owners. However, the contract suffered from a fatal bug that allowed anyone to become the sole owner because the corresponding functions *initWallet*, *initMultiowned*, and *initDayLimit* did not perform any access control checks.

Figure 3 shows the patched source code which adds the *internal* modifier to the functions *initMultiowned* and *initDayLimit* (marked with ① in Figure 3). This modifier makes these two functions inaccessible via the outside interface of the deployed contract. Furthermore, the patch adds the custom modifier *only_uninitialized*, which checks whether the contract was previously initialized (marked with ②).

The developers originally introduced a new vulnerability while deploying the patched the contract, which was actively exploited [38]. In contrast, because EVMPATCH performs bytecode rewriting, it would have immediately generated a securely patched version of the contract and would have deployed it automatically in a secure manner.

Consider Figure 4 which shows a customized patch in the domain-specific language employed by EVMPATCH to specify patches. As such, we insert a patch at the beginning of the *initWallet* function that checks whether the condition `sload(m_numOwners) == 0` holds, i.e., whether the contract is not yet initialized. If this does not hold, the contract execution will abort with a REVERT instruction. Note that here an explicit `sload` needs to be used to load variables from storage and the

```

1 add_require_patch:
2   initWallet:
3     - sload(m_numOwner) == 0
4
5 delete_public_function_patch:
6   - initDayLimit
7   - initMultiowned

```

Figure 4: Customized Patch for Parity Multisig Wallet.

expression is logically inverted from the patch in Figure 3, since this patch essentially inserts a Solidity `require` statement. Furthermore, two other publicly accessible functions need to be removed from the public function dispatcher. The patch shown in Figure 4 combines two existing patch templates provided by EVMPATCH. First, the *add require* patch template enforces a pre-condition before a function is entered. Second, the *delete public function* patch template removes a public function from the dispatcher, effectively marking the function as internal.

Evaluation Results. We verified that the patched contract is no longer exploitable by deploying a patched version of the *WalletLibrary* contract against the attack. Further, we compare a source-level patch with the patch applied by EVMPATCH. Table 2 shows an overview of the results. EVMPATCH only increases contract size by 25 B. The additional gas cost of the *initWallet* function is only 235 gas, i.e., 0.000,06 USD per transaction for 235.091 USD/ETH and a typical gas price of 1 Gwei. This demonstrates that EVMPATCH can efficiently and effectively insert patches for access control bugs.

5.2 Patching Integer Bugs

Typical integer types are bound to a minimum and/or maximum size due to the fixed bit-width of the integer type. However, programmers often do not pay sufficient attention to the size limitation of the actual integer type potentially causing integer bugs. Fortunately, several high-level programming languages (Python, Scheme) are able to avoid integer bugs since they leverage arbitrary precision integers with virtually unlimited size. However, the de-facto standard programming language for smart contracts, namely Solidity, does not embed such a mechanism. This leaves the burden of handling integer overflows completely on the developer who needs to either manually implement overflow checks or properly utilize the SafeMath library to safely perform numeric operations [33]. While common, the former is obviously error-prone. For instance, multiple vulnerabilities in ERC-20 token contracts were recently unveiled [1,

Table 2: Overhead of access control patch.

Version	Bytes	Size Increase	Gas Increase
Original	8290	0 %	0
Source-Patched	8201	-1.07 %	226
EVMPATCH'ed	8315	0.3 %	235

```

1 function batchTransfer(address[] _receivers, uint256 _value)
2     public whenNotPaused returns (bool) {
3     uint cnt = _receivers.length;
4     // OVERFLOW: 2 * ((INT_MAX / 2) + 1) == 0
5     uint256 amount = uint256(cnt) * _value;
6     require(cnt > 0 && cnt <= 20);
7     // BYPASSED CHECK: balances[msg.sender] >= 0
8     require(_value > 0 && balances[msg.sender] >= amount);
9     // RESULT: Transfer of ((INT_MAX / 2) + 1) tokens

```

Figure 5: Integer overflow bug reported by PeckShield [1].

26, 27]. These contracts manage subcurrencies, so-called tokens, on the Ethereum blockchain. Such tokens can deal with large amounts of currency since they track the token balance of every token owner and mediate the exchange of tokens and Ether. Figure 5 shows an excerpt of the *BEC* token contract’s code that exemplifies such integer overflow vulnerabilities. When computing the total amount in Line 6, an unchecked integer multiplication is used allowing an attacker to provide a very large *_value*. As a consequence, the *amount* variable will be set to a small amount. This effectively bypasses the balance check in Line 11 allowing the attacker to transfer a large amount of tokens to an attacker-controlled account. Recently, similar vulnerabilities have been discovered in over 42,000 contracts [13].

We developed patch templates for detecting integer overflows and underflows for the standard EVM integer width, i.e., unsigned 256 bit integers. For integer addition, subtraction, and multiplication, these templates add checks inspired by secure coding rules in the C programming language [34] and the *SafeMath* [33] Solidity library. When a violation is detected, EVMPATCH issues an exception to abort and roll back the current call to the contract.

5.2.1 Evaluation Results

To verify the correctness of the patches generated by our bytecode rewriter, we utilized the state-of-the-art integer detection tool Osiris [13] for vulnerability detection. After analyzing 50,535 unique contracts in the first 5,000,000 blocks of the Ethereum blockchain, Osiris detects at least one integer overflow vulnerability in 14,107 contracts. Using EVMPATCH, we were able to successfully patch almost all of these contracts automatically. More specifically, we could not patch 33 contracts amongst the 14107 investigated contracts because the basic block, where the detected vulnerability was located is too small for the trampoline code.

From those 14107 contracts, around 8000 involve transactions on the Ethereum network. To generate a large and representative evaluation data set, we extracted all transactions sent to these contracts up to block 7,755,100 (May 13 2019) from the Ethereum blockchain resulting in 26,385,532 transactions.

Replaying those transactions with our patch tester shows that for 95.5% of all vulnerable contracts, EVMPATCH’s generated patch was compliant to all of the prior transactions associated with those contracts. For the remaining 4.5% of

the investigated contracts, our patch rejected transactions for one of the following reasons: (1) we successfully stopped a malicious transaction, (2) the reported vulnerability was a false positive and should not have been patched, or (3) we unintentionally changed the contract’s functionality.

For close scrutiny, we selected ERC-20 token contracts from those contracts that could be successfully patched by EVMPATCH with confirmed integer overflow/underflow vulnerabilities that have been successfully attacked (see Table 3). For comparison purposes, we also manually patch these contracts on the Solidity source code level by replacing the vulnerable arithmetic operations with functions adapted from the *SafeMath* library [33]. The manually patched source code is then compiled with the exact same Solidity compiler version and optimization options used in the original contract (as reported on etherscan.io).

We applied the EVMPATCH patch tester to the generated patched contract versions and validated the reported outcome. This allows us to verify whether both patching approaches abort the same attack transactions. In addition, we can compare the overhead in gas consumption and the increase in code size. Note that in the manual patching method, we do not patch all potential vulnerabilities detected by Osiris as we skip adding checks on those arithmetic operations which cannot be exploited by an attacker, i.e., vulnerable arithmetic operations contained in functions that can only be called by the controller or owner of the contract. We verified the correctness of our patches using a total number of 506,607 real-world transactions associated with the ERC-20 token contracts listed in Table 3.

Table 3 shows the transaction execution results of the patch tester. We verified the aborted transactions and confirm that all of them correspond to genuine attacks except for one transaction⁷, which resembles a special case of token burning that we discuss in detail below. Apart from the valid attack transactions, the execution traces of the re-executed transactions match those of the original transactions, confirming that our patch does not break the contract’s functionality.

Out of the transactions identified as attacks, we found one particular transaction to the HXG token [17]. The transaction does indeed trigger an integer overflow but the HXG token rather burns some tokens by transferring them to a blackhole address 0x0. The burned tokens cannot be recovered and the balance of the blackhole address does not influence the behavior of the contract. When analyzing the contract, Osiris is not aware of the semantics of this blackhole address and reports a possible integer overflow. EVMPATCH then conservatively patches the integer overflow bugs reported by Osiris, which leads to one legitimate transaction failing. We argue that this pattern can be seen as bad coding practice as it wastes gas in unnecessarily storing the balance of the blackhole address.

Gas Overhead. The additional code introduced by the patching may potentially cause transactions to fail with an out-of-gas

⁷0x776da02ce8ce3cc882eb7f8104c31414f9fc756405745690bcf8df21e779e8a4

Contract	CVE	# Patches	# Transactions		Overhead (gas)		Code Size Increase (B)		Additional Cost RW (US\$)	
			Total	Attacks	RW	SM	RW	SM	per TX	per Upgrade
BEC [2]	2018-10299	1	424,229	1	83	164	117 (1.0%)	133 (1.1%)	< 0.01	0.01
SMT [36]	2018-10376	1	56,555	1	47	108	191 (0.8%)	97 (0.4%)	< 0.01	0.01
UET [43]	2018-10468	55	24,034	12	225	21	1,299 (18.2%)	541 (7.6%)	< 0.01	0.071
SCA [37]	2018-10706	1	292	10	47	0	3,811 (17.3%)	361 (1.6%)	< 0.01	0.189
HXG [17]	2018-11239	9	1497	5	120	541	997 (28.1%)	519 (14.6%)	< 0.01	0.057

Table 3: ERC-20 Token contracts investigated in depth with their respective CVE number, the number of patches introduced by EVMPATCH, and the number of transactions replayed by EVMPATCH’s patch tester and the number of attack transactions identified while testing the patches. We also give the average amount of overhead in gas consumption over all replayed transactions and overhead of contract size of the manual patched contracts (SM) and rewriter-generated patches (RW) and the overhead of the rewriter converted to US\$ (with a gas price of 1 Gwei and 235 US\$/eth; For readability we only show the exact US\$ figures only if they are more than one cent).

error. While the patches generally do not significantly increase gas consumption, such a behavior can nevertheless occur when the sender of the transactions provides a very tight gas budget. When the re-execution of a transaction with patched code fails early due to an out-of-gas exception, we could not accurately compare the behavior of the patched contract with the original contract. To remedy this, we disabled the gas-accounting in the EVM. We report the amount of additional gas consumption during transaction execution in Table 3. We excluded those transactions that do not execute functions which contain the vulnerable code, because they are not affected by the patches and therefore not relevant to our measurements.

Our results show that for contracts BEC, SMT, and HXG, those patched with EVMPATCH incur less gas overhead at runtime (83 gas, 47 gas and 120 gas) when compared to those patched on the source code level (164 gas, 108 gas and 541 gas). This is due to the fact that the Solidity compiler generates non-optimal code when only very few checks are added. In particular, Solidity utilizes internal function calls to invoke the SafeMath integer overflow checks. While this reduces code size (in case the check is needed at multiple places), it always requires executing additional instructions—thereby increasing gas overhead—to invoke and return from the internal function. In contrast, EVMPATCH inlines the safe numeric operations thereby introducing less gas overhead. One would need to instruct the Solidity compiler to selectively enable function inlining to yield similar gas costs as EVMPATCH.

Note that the average gas overhead is 0 gas for the manually patched SCA token. This is because only one transaction triggers the SafeMath integer overflow check. However, this is an attack transaction and it is aborted early, making gas overhead calculation not possible.

For UET and SCA, we identify higher gas overhead than for the manually patched version. In fact, UET requires on average 255 units of additional gas for every transaction in the patched version. In contrast, only 21 gas is added for manually patched version. This is due to the fact that our bytecode rewriter conservatively patches every potential vulnerability reported by Osiris in these two contracts (12 and 10 respectively).

However, not all of them are actually exploitable and as such we did not instrument them during manual patching.

Code Size Increase. Deploying contracts in the Ethereum blockchain also incurs costs proportionally to the size of the deployed contract. More specifically, Ethereum charges 200 gas per byte to store the contract code on the blockchain [44]. From Table 3, we recognize that the amount of extra code added by our rewriter is comparable to that of the SafeMath approach when a single vulnerability is patched. Since our approach duplicates the original basic blocks, the code size overhead depends on the specific location of the vulnerability. In the case of the BEC token contract, our rewriter increases the code size less than the source-level patches. The Solidity compiler generates more code for including the SafeMath library than is strictly necessary for the patch. Even considering the overhead of bytecode rewriting, we observe that EVMPATCH generates a smaller patch than the manual patching method for this contract.

However, in case many vulnerabilities are patched, EVMPATCH adds a slightly higher overhead. Naturally, the size of the upgraded contracts increases with the number of vulnerabilities to fix due to inlining. For instance, our bytecode rewriter generates 12 patches for UET contract and 10 patches for SCA contract resulting in 1299 B (18.2%) and 3811 B (17.3%) increase in code size. In the worst-case scenario in our dataset, this increase in code size induces negligible additional cost of US\$0.18 per deployment.

Our patch templates are currently optimized for patching a single vulnerable arithmetic. It is straightforward to adopt an approach akin to Solidity’s internal function calls when developing patch templates for our bytecode rewriter, which would reduce the code size overhead when patching many integer overflows.

EVMPATCH applies 3.9 patches on average to a contract in our data set of 14,107 contracts. The average code size of the original contracts is 8142.7 B (σ 5327.8 B). The average size increase after applying patches with EVMPATCH is 455.9 B (σ 333.5 B). This amounts to an average code size overhead of 5.6% after applying the patches. Given that Ethereum charges

200 gas per byte to the contract creation transaction, it incurs an average overhead of 91,180 gas or US\$0.02 at the time of writing. In the worst case that we observed, EVMPATCH incurs an overhead of 199,800 gas at deployment, which at the time of writing only amounts to about US\$0.04 additional deployment cost. This shows that the overhead of applying patches with bytecode rewriting is negligible for contract deployment, especially when compared to the number of Ether possibly at stake.

Costs of Deployment. The deployment cost of a newly patched contract dominates the costs of operating a smart contract with EVMPATCH. However, additionally there is a transaction needed to switch the address of the logic contract. Since the proxy pattern requires no state migration, this transaction requires a constant amount of gas. The proxy contract we utilize in EVMPATCH consumes 43.167 gas during a switchover transaction, i.e., about US\$0.01. Currently, state migration is the most viable contract upgrade strategy besides the proxy pattern. Prior work estimated that even with only 5000 ERC-20 holders, i.e., smart contract users, state migration will likely cost more than US\$100.00 in the best case [41]. Hence, compared to the cost of migrating all data to a new contract, the EVMPATCH’s additional cost of US\$0.01 is negligible.

Detecting Attacks. The patch tester of EVMPATCH allows us to also identify any prior attack transactions. In Figure 6, we additionally observe that while the vulnerabilities of the other token contracts have been reported within a fairly reasonable time after the first attack, UET has been exploited (5 months) long before the bug disclosure. More surprisingly, all contracts are still fairly active though they encountered a decrease of transaction volume after public disclosure of the vulnerabilities. Despite the fact that all of these vulnerabilities have been discovered around one year before the time of writing, there are still 23,630 transactions (4.66% of the evaluated transactions) issued to these vulnerable contracts after the public disclosure of the vulnerabilities, including successful attacks. This means that the owners of those contracts did not properly migrate to patched versions and users were not properly notified of the vulnerable state of these contracts.

5.2.2 Analysis of False Positives/Negatives

During our analysis of the vulnerable contracts, we identified false positives and false negatives caused by vulnerability reporting of Osiris [13]. This demonstrates that our patch testing is an important step in the process as many analysis tools are imprecise. We found that in the default configuration, Osiris often achieves limited code coverage. To this end, we utilized different timeout settings for both the whole analysis and for queries to the SMT solver and combined the results of multiple runs to achieve better code coverage. Furthermore, we found that—contrary to the claims in the original *Osiris* paper [13]—not all vulnerabilities are accurately detected by Osiris in two particular cases.

Hexagon (HXG) Token. This contract is vulnerable to an




Figure 6: Activity timeline of each contract. The grey shadow indicates the time window in which the vulnerabilities of these contracts are disclosed by Peckshield [30], and the big hollow points signify the occurrences of the attacks.

integer overflow, which allows an attacker to transfer very large amounts of ERC-20 tokens [26]. Osiris reports two false positives, which are caused by EVM code that is generated by the Solidity compiler. Even though all types are unsigned types in the Solidity source code, the compiler generates a signed addition. Here, Osiris reports a possible integer overflow, when -2 is added to the *balanceOf* mapping variable. When performing signed integer additions with negative values, the addition naturally overflows when the result moves from the negative value range into the positive value range and vice versa. As such, EVMPATCH patches a checked addition for an unsigned arithmetic operation which will always overflow. With our patch tester we observe all the failing transactions and perform manual analysis of the patched contract’s bytecode to determine that the root cause is an issue in the Solidity compiler, i.e., the generated code requires an additional instruction, when compared to a simple unsigned subtraction.

Social Chain (SCA). Our results also show a problem with Osiris when analyzing the SCA token. While Osiris does detect a possible overflow during multiplication in the problematic Solidity source code line, it does not detect the possible integer overflow for an addition in the same source code line. However, in the actual attack transaction, the integer overflow happens during the not-flagged addition operation. As such, this constitutes a false negative problem of Osiris. Since the vulnerable addition is not reported by Osiris, it is also not automatically patched by EVMPATCH. In contrast, for the manually patched version we took both arithmetic operations into account. The related attack transaction was previously reported as an attack transaction [27].

Summary of Evaluation. To summarize, our evaluation on integer overflow detection shows that EVMPATCH can correctly apply patches to smart contracts preventing *any* integer overflow attack. Furthermore, EVMPATCH incurs only a *negligible* gas overhead during deployment and runtime; especially compared to the Ether at stake. Our analysis shows that

the analyzed vulnerable smart contracts are *still in active use, even after being attacked* and the vulnerabilities being publicly disclosed. This motivates the need for a timely patching framework such as EVMPATCH. Lastly, based on an extensive and detailed analysis of 26,385,532 transactions, we demonstrate that EVMPATCH always preserves the contract’s original functionality except for a few cases, where the vulnerability report (generated by the third-party tool Osiris) was not accurate or bad coding practices were used (blackhole address).

5.3 Developer Study

Developer Background. To quantify the manual effort needed to patch smart contracts and evaluate the usefulness of EVMPATCH we conducted a thorough study with 6 professional developers with varying prior experience in using blockchain technologies and developing smart contracts. Our developers consider themselves familiar with blockchain technologies but not very familiar with developing Solidity code. None of the developers have developed an upgradable contract before. As such, we can quantify the effort needed for a smart contract developer to learn and apply an upgradable contract pattern.

Methodology. Throughout our study, we asked the developers to perform multiple tasks manually that are performed automatically by EVMPATCH: (1) manually patch three contracts vulnerable due to integer overflow bugs given the output of a static analyzer (OSIRIS [13]), (2) convert a contract to an upgradable contract manually and with EVMPATCH, and (3) patch an access control bug using EVMPATCH by writing a custom patch-template. The three tasks cover different scenarios, where EVMPATCH can be useful to a developer. The first two tasks cover the use of EVMPATCH to patch known bug classes with minimal human intervention. For these two tasks we assume no prior knowledge on patching smart contracts (see Table 5 how developers rated their prior experience with smart contracts). In contrast, the third task consists of extending EVMPATCH. This requires understanding a bug class and perform root cause analysis to properly patch the vulnerability. This is surely more challenging compared to the previous two tasks. Since the third task covers a different

Table 4: Timing results for the tasks as reported by the developers given in minutes and their reported confidence in the correctness of their results.

Task	Time (Minutes)			Confidence Median (1-7)
	Median	Min	Max	
Manual Integer Patches	47.50	35	78	6
Conversion	62.50	33	110	2.5
EVMPATCH Conversion	1.50	1	3	-
Patch Template	4.00	2	15	7

bug class, we believe there is no significant bias in the data due to the developers completing the other two tasks first.

For all tasks, we measured the time required by the developer to perform the task (excluding the time required for reading the tasks’ description). We asked the developers to rate their familiarity with relevant technologies, their confidence levels in their patches, and the difficulty of performing the tasks on a 7-point Likert scale. The full questionnaire and the answers of the developers are shown in Table 5, and the recorded time measurements are shown in Table 4. We provide the supporting files in a github repository.⁸

We then performed both a manual code review and a cross-check with EVMPATCH to analyze mistakes made by the developers. The results of our study show that significant effort is needed to correctly patch smart contracts manually, whereas EVMPATCH enables simple, user-friendly, and efficient patching. The time measurements show that the developers, who had no prior experience with EVMPATCH, were able to perform complex tasks utilizing EVMPATCH within minutes.

Patching Integer Overflow Bugs. We asked the developers to fix all integer overflow vulnerabilities in three contracts: 1 BEC [2] (CVE-2018-10299, 299 lines of code), and 2 HXG [17] (CVE-2018-11239, 102 lines of code) and 3 SCA [37] (CVE-2018-10706, 404 lines of code). To provide a representative set of contracts, we chose three ERC-20 contracts with varying complexity (in terms of lines of code) and where the static analysis also includes missed bugs and false alarms (see § 5.2.2). We ran OSIRIS on all three contracts and provided the developers the analysis output as well as a copy of the SafeMath Solidity library. This accurately resembles a real-world scenario, where a blockchain developer quickly needs to patch a smart contract based on the analysis results of recent state-of-the-art vulnerability analysis tools and can look-up manual patching tutorials available online. All developers manually and correctly patched the source code of all three contracts which demonstrates their expertise in blockchain development. However, on the downside, it took the developers on average 51.8 min ($\sigma = 16.6$ min) to create patched version for the three contracts. In contrast, EVMPATCH fully automates the patching process and is able to generate patches for the three contracts within a maximum of 10 s.

Converting to an Upgradable Contract. The developers had to convert a given smart contract into an upgradable smart contract. We provided the developers a short description of the delegatecall-proxy pattern and asked them to convert the given contract into two contracts: one proxy contract and a logic contract, which is based on the original contract. We provided no further information on how to handle the storage-layout problem, and we explicitly allowed using code found online. The developers required an average of 66.3 min⁹ to convert a contract into an upgradable contract. *None* of the developers performed a correct conversion into an upgradable contract,

⁸ github.com/uni-due-syssec/evmpatch-developer-study

⁹ $\sigma = 31.3$ min, fastest 33 min and slowest 110 min

which is also reflected in a median confidence of 2.5 in the correctness reported by the developers. We observed two major mistakes: (a) The proxy contract would only support a fixed set of functions, i.e., the proxy would not support adding functions to the contract, and (b) more importantly, only one out of six developers correctly handled storage collisions in the proxy and logic contract, i.e., five of the six converted contracts were broken by design. Hence, it remains open how long it would take developers to perform a correct conversion.

Next, we asked the developers to utilize EVMPATCH to create and deploy an upgradable contract. As EVMPATCH does not require *any* prior knowledge about upgradable contracts, the developers were able to deploy a correct upgradable contract within at most 3 min. In addition, patching with EVMPATCH inspires high confidence—a median of 7, the best rating on our scale—in the correctness of the patch. This gives a strong confirmation that deployment of a proxy with EVMPATCH is indeed superior to manual patching and upgrading.

Extending EVMPATCH. The developers had to write a custom patch template for EVMPATCH. We instructed the developers on how to use EVMPATCH and how patch templates are written with EVMPATCH’s patch template language (see Figure 4 for an example). Furthermore, we presented the developers an extended bug report that shows how an access control bug can be exploited. The developers leveraged the full EVMPATCH system, i.e., EVMPATCH applies the patch and validates the patch using the patch tester component which replays past transactions from the blockchain and notifies the developer whether: (a) the patch prevents a known attack, and (b) whether the patch broke functionality in other prior legitimate transactions. As such, EVMPATCH allowed the developers to create a fully functional and securely patched upgradable contract within a few minutes. On average, the developers only needed 5.5 min, and a maximum of 15 min, to create a custom patch template. As expected, all developers correctly patched the given contract using EVMPATCH, because a faulty patch would have been reported by EVMPATCH’s patch tester to the developer. EVMPATCH’s integrated patch tester gives the developers a high confidence into their patch. On average, the developers reported a confidence level of 6.6 ($\sigma = 0.4$), where 7 is the most confident. Furthermore, *none* of the developers considered writing such a custom patch template as particularly difficult.

Summary. Our study provides confirmation that EVMPATCH offers a high degree of automation, efficiency, and usability thereby freeing developers from manual and error-prone tasks. In particular, none of the six developers were able to produce a correct upgradable contract mainly due to the difficulty of preserving the storage-layout. Our study also confirms that extending EVMPATCH with custom patch templates is a feasible task, even for developers that are unaware of the inner workings of EVMPATCH.

6 Related Work

The infamous attack against “TheDAO” contract [7] received considerable attention from the community. Since then, many additional exploits and defenses, which mostly focus on discovering bugs before the contract is deployed, were revealed. Luu et al. presented the symbolic executor Oyente that explores a contracts code, while looking for possible vulnerabilities [23]. Since then many other symbolic execution tools with better precision, performance, and covering different vulnerabilities have been proposed [13, 20, 24, 25, 28]. Furthermore, static analyzers for both Solidity [12] and EVM bytecode have been proposed [42]. Information flow analysis and data sanitization in a multi-transaction setting is analyzed by Ethainter [4]. Furthermore, methods from formal verification and model checking have been applied to smart contracts [14, 19] and the semantics of the EVM and Solidity language have been formalized [15, 18]. However, only a small body of prior work has researched dynamic analysis and runtime protections. Tools such as Sereum [32] or ECFChecker [16] can detect live reentrancy attacks on vulnerable contracts. Recent work has further explored modular dynamic analysis frameworks for protecting smart contracts [6, 40]. Protection solutions that require modifications to the smart contract execution environment are unlikely to be integrated in to production blockchain systems.

Integer overflows have been widely studied in the context of Ethereum smart contracts. Osiris [13] is an extension to the symbolic execution tool Oyente [23] to accurately detect integer bugs. The improved symbolic execution engine first attempts to infer the integer type, i.e., signedness and bit width, from the specific instructions generated by Solidity compilers. Next, it checks for possible integer bugs, such as truncation, overflow, underflow, and wrong type casts. We leverage the detection capabilities of Osiris, because it pinpoints the exact location of the integer overflow bug. Other tools such as TeEther [20] and MAIAN [28] implicitly find integer bugs when they generate exploits for smart contracts. However, they do not report the exact location of the integer overflow, because they focus on exploit generation. ZEUS [19] utilizes abstract interpretation and symbolic model checking to verify safety properties of smart contracts. While ZEUS can detect potential integer overflow vulnerabilities, it does so at the LLVM intermediate level and cannot determine the exact location in the corresponding EVM bytecode.

Recently, bytecode rewriting for patching smart contracts has been explored with SMARTSHIELD [47]. SMARTSHIELD requires a complete control-flow graph (CFG) to update jump targets and data references. As discussed in § 4.1, generating a highly accurate CFG is highly challenging due to the EVM’s bytecode format. We believe that such a bytecode rewriting strategy does not scale to larger and more complicated contracts. In contrast, EVMPATCH’s trampoline-based rewriting strategy does not require an accurate CFG and is much more resilient when rewriting complex contracts.

SMARTSHIELD implements custom bytecode analysis to detect vulnerabilities, which may not be as accurate as specialized analyses. For example, SMARTSHIELD’s analysis does not infer whether an integer type is signed, which is important for accurate integer overflow detection [13]. EVMPATCH is a flexible framework that can integrate many static analysis tools for detecting vulnerabilities and can leverage analysis tool improvements with minimal effort. Last and most importantly, EVMPATCH automates the whole lifecycle of deploying and managing an upgradable contract, while SMARTSHIELD is designed to harden a contract pre-deployment. With EVMPATCH, a smart contract developer can also patch vulnerabilities that are discovered after deployment of the contract.

The Ethereum community explored several design patterns to allow upgradable smart contracts [11, 41, 45, 46] with manual migration to a new contract and the proxy pattern being the most popular (see § 2). The ZeppelinOS [46] framework supports upgradable contracts by implementing the delegatecall-proxy pattern. However, developers have to manually ensure compatibility of the legacy and patched contract on the Solidity level. This can be achieved using static analysis tools that perform “upgradeability” checks (e.g., Slither [35] checks for a compatible storage layout), which relies on accurate knowledge of compiler behavior with respect to storage allocations. On the other hand, EVMPATCH combines existing analysis tools and provides an automatic method to patch detected vulnerabilities while keeping storage layout consistent by design.

7 Conclusion

Updating erroneous smart contracts constitutes one of the major challenges in the field of blockchain technologies. The recent past has shown that attackers are fast in successfully abusing smart contract errors due to the natural design of the underlying technology: always online and available, one common and simple computing engine without any subtle software and configuration dependencies, and (often) high amount of cryptocurrency at disposal. While many proposals have introduced frameworks to aid developers in finding bugs, it remains open how developers and the community can quickly and automatically react to vulnerabilities on already deployed contracts. In this work, we developed a framework that supports automated and instant patching of smart contract errors based on bytecode rewriting. In terms of evaluation, we were able to demonstrate that real-world vulnerable contracts can be successfully patched without violating the functional correctness of the smart contract. Our developer study shows that an automated patching approach greatly reduces the time required for patching smart contracts and that our implementation, EVMPATCH, can be practically integrated into a smart contract developers workflow. We believe that automated patching will increase the trustworthiness and acceptance of smart contracts as it allows developers to quickly react on reported vulnerabilities.

Acknowledgment

The authors would like to thank the reviewers—and especially our shepherd Yinzhi Cao—for their valuable feedback, and the developers for taking the time to participate in our study. This work was partially funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany’s Excellence Strategy - EXC 2092 CASA - 390781972 and the DFG as part of project S2 within the CRC 1119 CROSSING. This work has been partially supported by the EU H2020-SU-ICT-03-2018 CyberSec4Europe project, funded by the European Commission under grant agreement no. 830929.

References

- [1] *ALERT: New batchOverflow Bug in Multiple ERC20 Smart Contracts (CVE-2018-10299)*. URL: <https://blog.peckshield.com/2018/04/22/batchOverflow/> (visited on 05/27/2019).
- [2] *BeautyChainToken*. URL: <https://etherscan.io/address/0xc5d105e63711398af9bbff092d4b6769c82f793d>.
- [3] Lorenz Breidenbach, Phil Daian, Ari Juels, and Emin Gün Sirer. *An In-Depth Look at the Parity Multisig Bug*. 2017. URL: <http://hackingdistributed.com/2017/07/22/deep-dive-parity-bug/> (visited on 04/20/2018).
- [4] Lexi Brent, Neville Grech, Sifis Lagouvardos, Bernhard Scholz, and Yannis Smaragdakis. “Ethainter: A Smart Contract Security Analyzer for Composite Vulnerabilities”. In: *Proceedings of the 41th ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI)*. 2020. DOI: [10.1145/3385412.3385990](https://doi.org/10.1145/3385412.3385990).
- [5] Bryan Roger Buck and Jeffrey K. Hollingsworth. “An API for Runtime Code Patching”. In: *Int. J. High Perform. Comput. Appl.* 14.4 (2000). DOI: [10.1177/109434200001400404](https://doi.org/10.1177/109434200001400404).
- [6] Ting Chen et al. “SODA: A Generic Online Detection Framework for Smart Contracts”. In: *Proceedings Network and Distributed System Security Symposium (NDSS)*. 2020. DOI: [10.14722/ndss.2020.24449](https://doi.org/10.14722/ndss.2020.24449).
- [7] Phil Daian. *Analysis of the DAO exploit*. 2016. URL: <http://hackingdistributed.com/2016/06/18/analysis-of-the-dao-exploit/> (visited on 10/10/2017).
- [8] Lucas Davi, Alexandra Dmitrienko, Manuel Egele, Thomas Fischer, Thorsten Holz, Ralf Hund, Stefan Nürnberger, and Ahmad-Reza Sadeghi. “MoCFI: A Framework to Mitigate Control-Flow Attacks on Smartphones”. In: *Proceedings Network and Distributed System Security Symposium (NDSS)*. 2012.
- [9] Solidity Documentation. *Layout of State Variables in Storage*. URL: <https://solidity.readthedocs.io/en/v0.5.10/miscellaneous.html> (visited on 05/27/2019).
- [10] Ethereum EIPs. *ERC930 - Eternal Storage Standard*. URL: <https://github.com/ethereum/EIPs/issues/930> (visited on 11/08/2019).
- [11] *Ethereum Smart Contract Best Practices: Upgrading Broken Contracts*. URL: https://consensys.github.io/smart-contract-best-practices/software_engineering/#upgrading-broken-contracts (visited on 09/28/2020).
- [12] Josselin Feist, Gustavo Grieco, and Alex Groce. “Slither: a static analysis framework for smart contracts”. In: *Proceedings of the 2nd International Workshop on Emerging Trends in Software Engineering for Blockchain (WETSEB@ICSE)*. 2019. DOI: [10.1109/WETSEB.2019.00008](https://doi.org/10.1109/WETSEB.2019.00008).

- [13] Christof Ferreira-Torres, Julian Schütte, and Radu State. “Osiris: Hunting for Integer Bugs in Ethereum Smart Contracts”. In: *Proceedings of the 34th Annual Computer Security Applications Conference (ACSAC)*. 2018. DOI: [10.1145/3274694.3274737](https://doi.org/10.1145/3274694.3274737).
- [14] Joel Frank, Cornelius Aschermann, and Thorsten Holz. “ETHBMC: A Bounded Model Checker for Smart Contracts”. In: *29th USENIX Security Symposium*. 2020.
- [15] Ilya Grishchenko, Matteo Maffei, and Clara Schneidewind. “Foundations and Tools for the Static Analysis of Ethereum Smart Contracts”. In: *Computer Aided Verification*. Springer International Publishing, 2018. DOI: [10.1007/978-3-319-96145-3_4](https://doi.org/10.1007/978-3-319-96145-3_4).
- [16] Shelly Grossman, Ittai Abraham, Guy Golan-Gueta, Yan Michalevsky, Noam Rinetzky, Mooly Sagiv, and Yoni Zohar. “Online detection of effectively callback free objects with applications to smart contracts”. In: *Proceedings of the ACM on Programming Languages* POPL (2018). DOI: [10.1145/3158136](https://doi.org/10.1145/3158136).
- [17] *HexagonToken*. URL: <https://etherscan.io/address/0xB5335e24d0aB29C190AB8C2B459238Da1153cEBA>.
- [18] Jiao Jiao, Shuanglong Kan, Shang-Wei Lin, David Sanan, Yang Liu, and Jun Sun. “Semantic Understanding of Smart Contracts: Executable Operational Semantics of Solidity”. In: *2020 IEEE Symposium on Security and Privacy*. 2020. DOI: [10.1109/SP40000.2020.00066](https://doi.org/10.1109/SP40000.2020.00066).
- [19] Sukrit Kalra, Seep Goel, Mohan Dhawan, and Subodh Sharma. “ZEUS: Analyzing Safety of Smart Contracts”. In: *Proceedings 2018 Network and Distributed System Security Symposium (NDSS)*. 2018. DOI: [10.14722/ndss.2018.23082](https://doi.org/10.14722/ndss.2018.23082).
- [20] Johannes Krupp and Christian Rossow. “teEther: Gnawing at Ethereum to Automatically Exploit Smart Contracts”. In: *27th USENIX Security Symposium*. USENIX Association, 2018. URL: <https://www.usenix.org/conference/usenixsecurity18/presentation/krupp>.
- [21] M A Laurenzano, M M Tikir, L Carrington, and A Snively. “PEBIL: Efficient static binary instrumentation for Linux”. In: *2010 IEEE International Symposium on Performance Analysis of Systems Software (ISPASS)*. 2010. DOI: [10.1109/ISPASS.2010.5452024](https://doi.org/10.1109/ISPASS.2010.5452024).
- [22] Chi-Keung Luk, Robert Cohn, Robert Muth, Harish Patil, Artur Klauser, Geoff Lowney, Steven Wallace, Vijay Janapa Reddi, and Kim Hazelwood. “Pin: building customized program analysis tools with dynamic instrumentation”. In: *Acm sigplan notices*. Vol. 40. ACM, 2005.
- [23] Loi Luu, Duc-Hiep Chu, Hrishi Olickel, Prateek Saxena, and Aquinas Hobor. “Making smart contracts smarter”. In: *Proceedings of the ACM SIGSAC Conference on Computer and Communications Security (CCS)*. ACM, 2016.
- [24] Mark Mossberg, Felipe Manzano, Eric Hennenfent, Alex Groce, Gustavo Grieco, Josselin Feist, Trent Brunson, and Artem Dinaburg. “Manticore: A User-Friendly Symbolic Execution Framework for Binaries and Smart Contracts”. In: *CoRR* abs/1907.03890 (2019). arXiv: [1907.03890](https://arxiv.org/abs/1907.03890).
- [25] ConsenSys. *Mythril v0.19.7 - Security analysis tool for EVM bytecode*. <https://github.com/ConsenSys/mythril>.
- [26] *New burnOverflow Bug Identified in Multiple ERC20 Smart Contracts (CVE-2018-11239)*. URL: <https://blog.peckshield.com/2018/05/18/burnOverflow> (visited on 05/27/2019).
- [27] *New multiOverflow Bug Identified in Multiple ERC20 Smart Contracts (CVE-2018-10706)*. URL: <https://blog.peckshield.com/2018/05/10/multiOverflow%7D> (visited on 05/27/2019).
- [28] Ivica Nikolic, Aashish Kolluri, Ilya Sergey, Prateek Saxena, and Aquinas Hobor. “Finding The Greedy, Prodigal, and Suicidal Contracts at Scale”. In: *34th Annual Computer Security Applications Conference (ACSAC)*. 2018. DOI: [10.1145/3274694.3274743](https://doi.org/10.1145/3274694.3274743).
- [29] Mathias Payer, Boris Bluntschli, Thomas R Gross, et al. “DynSec: On-the-fly Code Rewriting and Repair”. In: *HotSWUp*. 2013.
- [30] *PeckShield advisories*. URL: <https://blog.peckshield.com/advisories.html> (visited on 05/27/2019).
- [31] Rob Price. *Digital currency Ethereum is cratering because of a \$50 million hack*. June 2016. URL: <https://www.businessinsider.com/dao-hacked-ethereum-crashing-in-value-tens-of-millions-allegedly-stolen-2016-6> (visited on 04/20/2018).
- [32] Michael Rodler, Wenting Li, Ghassan O. Karamé, and Lucas Davi. “Sereum: Protecting Existing Smart Contracts Against Re-Entrancy Attacks”. In: *Proceedings of the Network and Distributed System Security Symposium (NDSS)*. 2019. DOI: [10.14722/ndss.2019.23413](https://doi.org/10.14722/ndss.2019.23413).
- [33] *SafeMath library - Wrappers over Solidity’s arithmetic operations with added overflow checks*. URL: <https://github.com/OpenZeppelin/openzeppelin-solidity/blob/master/contracts/math/SafeMath.sol> (visited on 06/03/2020).
- [34] *SEI CERT C Coding Standard: INT30-C. Ensure that unsigned integer operations do not wrap*. URL: <https://wiki.sei.cmu.edu/confluence/display/c/INT30-C.+Ensure+that+unsigned+integer+operations+do+not+wrap> (visited on 06/03/2020).
- [35] *Slither Wiki: Upgradeability Checks*. URL: <https://github.com/crytic/slither/wiki/Upgradeability-Checks> (visited on 06/02/2019).
- [36] *SmartMeshICO*. URL: <https://etherscan.io/address/0x55F93985431Fc9304077687a35A1BA103dC1e081>.
- [37] *Social Chain*. URL: <https://etherscan.io/address/0xb75a5e36cc668bc8fe468e8f272cd4a0fd0fd773>.
- [38] Parity Technologies. *A Postmortem on the Parity Multi-Sig Library Self-Destruct*. Nov. 2017. URL: <http://paritytech.io/a-postmortem-on-the-parity-multi-sig-library-self-destruct> (visited on 06/03/2020).
- [39] Parity Technologies. *Security Alert - Parity Wallet*. Nov. 2017. URL: <http://paritytech.io/security-alert> (visited on 04/06/2018).
- [40] Christof Ferreira Torres, Mathis Baden, Robert Norvill, Beltran Borja Fiz Pontiveros, Hugo Jonker, and Sjouke Mauw. “ÆGIS: Shielding Vulnerable Smart Contracts Against Attacks”. In: (Mar. 2020). arXiv: [2003.05987](https://arxiv.org/abs/2003.05987) [cs.CR].
- [41] *Trail of Bits Blog: How contract migration works*. URL: <https://blog.trailofbits.com/2018/10/29/how-contract-migration-works> (visited on 06/03/2020).
- [42] Petar Tsankov, Andrei Dan, Dana Drachler-Cohen, Arthur Gervais, Florian Buenzli, and Martin Vechev. “Security: Practical security analysis of smart contracts”. In: *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security*. 2018.
- [43] *UselessEthereumToken*. URL: <https://etherscan.io/address/0x27f706edde3aD952EF647Dd67E24e38CD0803DD6>.
- [44] Gavin Wood. *Ethereum: A secure decentralised generalised transaction ledger*. Specification. Version BYZANTIUM VERSION 7e819ec - 2019-10-20. Ethereum Foundation, Oct. 20, 2019. URL: <https://ethereum.github.io/yellowpaper/paper.pdf>.
- [45] *ZeppelinOS Blog: Proxy Patterns*. URL: <https://blog.zeppelinos.org/proxy-patterns> (visited on 06/03/2020).
- [46] *ZeppelinOS Documentation*. URL: <https://docs.zeppelinos.org> (visited on 06/03/2020).
- [47] Yuyao Zhang, Siqi Ma, Juanru Li, Kailai Li, Surya Nepal, and Dawu Gu. “SMARTSHIELD: Automatic Smart Contract Protection Made Easy”. In: *27th IEEE International Conference on Software Analysis, Evolution and Reengineering (SANER)*. 2020. DOI: [10.1109/SANER48275.2020.9054825](https://doi.org/10.1109/SANER48275.2020.9054825).

Table 5: Developer study questionnaire and answers by six developers (identified by the letters A to F).

Question		Answers						Median	Scale
		A	B	C	D	E	F		
Q1	Did you write Solidity code in the last two weeks?	no	no	no	no	yes	no		(yes/no)
Q2	Have you previously worked on a production-grade Solidity-based Ethereum contract?	yes	no	no	no	no	no		(yes/no)
Q3	Have you previously worked on a production-grade smart contract on another Blockchain Platform?	no	no	yes	no	yes	yes		(yes/no)
Q4	How familiar are you with Blockchain technologies in general?	6	5	7	6	6	6	6	(1 not familiar, 7 very familiar)
Q5	How familiar are you with the Ethereum Blockchain in particular?	6	5	4	2	6	2	4.5	(1 not familiar, 7 very familiar)
Q6	How familiar are you with the Solidity programming language?	6	3	2	1	5	1	2.5	(1 not familiar, 7 very familiar)
Q7	How familiar are you with upgradable contracts in Solidity?	5	3	1	1	4	1	2	(1 not familiar, 7 very familiar)
Task 1									
T1Q1	How confident are you in the correctness of your patch to contract 1?	5	7	7	6	7	6	6.5	(1 least confident, 7 most confident)
T1Q2	How confident are you in the correctness of your patch to contract 2?	6	7	7	4	7	6	6.5	(1 least confident, 7 most confident)
T1Q3	How confident are you in the correctness of your patch to contract 3?	3	5	6	5	2	4	4.5	(1 least confident, 7 most confident)
T1Q4	How much time did you need to patch all three contracts?	78	35	40	40	55	63	47.5	(Time in Minutes)
Task 2									
T2Q1	Have you previously used the delegatecall-proxy pattern in a Solidity contract?	no	no	no	no	no	no		(yes/no)
T2Q2	Have you previously used a different pattern to make a Solidity contract upgradable?	no	no	no	no	no	no		(yes/no)
T2Q3	Have you previously used a different upgradable smart contract?	no	no	no	no	no	no		(yes/no)
T2Q4	How confident are you in the correctness of your conversion?	5	3	1	1	5	2	2.5	(1 least confident, 7 most confident)
T2Q5	How difficult was the manual conversion?	4	5	5	6	4	6	5	(1 easy, 7 most difficult)
T2Q6	How difficult was the conversion using the evmpatch tool?	1	1	1	1	1	1	1	(1 easy, 7 most difficult)
T2Q8	How much time did you need to convert the contract to an upgradable contract (Step 1)?	110	80	45	90	40	33	62.5	(Time in Minutes)
T2Q8	How much time did you need to convert the contract using EVMPatch (Step 2)?	3	1	1	2	3	1	1.5	(Time in Minutes)
Task 3									
T3Q1	How confident are you in the correctness of your patch?	6	7	7	7	7	6	7	(1 least confident, 7 most confident)
T3Q2	How difficult was the conversion using the EVMPatch tool?	2	1	1	1	1	1	1	(1 easy, 7 most difficult)
T3Q3	How much time did you need to create and deploy the patch using EVMPatch?	15	2	5	2	6	3	4	(Time in Minutes)