

Finding smart contract vulnerabilities with ConCert’s property-based testing framework

Mikkel Milo  



Computer Science, Aarhus University

Eske Hoy Nielsen  

Computer Science, Aarhus University

Danil Annenkov  

Computer Science, Aarhus University

Bas Spitters  

Computer Science, Aarhus University

Abstract

We provide three detailed case studies of vulnerabilities in smart contracts, and show how property based testing would have found them: 1. the Dexter1 token exchange; 2. the iToken; 3. the ICO of Brave’s BAT token. The last example is, in fact, new, and was missed in the auditing process.

We have implemented this testing in ConCert, a general executable model/specification of smart contract execution in the Coq proof assistant. ConCert contracts can be used to generate verified smart contracts in Tezos’ LIGO and Concordium’s rust language. We thus show the effectiveness of combining formal verification and property-based testing of smart contracts.

2012 ACM Subject Classification Software and its engineering → Formal methods; Software and its engineering → Software verification and validation

Keywords and phrases Smart Contracts, Formal Verification, Property-Based Testing, Coq

Digital Object Identifier 10.4230/LIPIcs.CVIT.2016.23

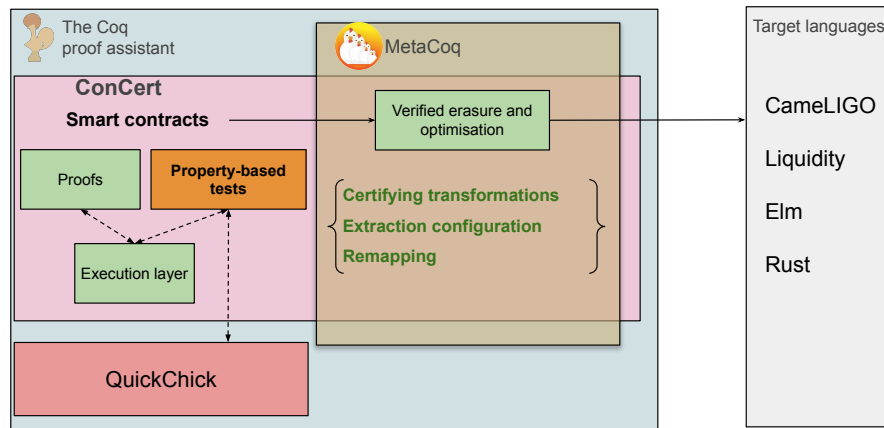
1 Introduction

Blockchain-based technologies have seen rising interest in recent years. This can be attributed to their ability to sustain a public distributed ledger with a high degree of reliability, integrity, and transparency, without requiring a trusted third party. Smart contracts are distributed applications deployed on a blockchain. They are typically used for sensitive transactions, for example, carrying large amounts of money or other valuable assets, but in principle, they can perform any computation. Once a smart contract is deployed on the blockchain, it is impossible to change its source code. The blockchain ensures that contracts are executed correctly according to the execution model. However, it gives no guarantee that the smart contract’s code is correct. Like other programs, smart contracts are susceptible to bugs.

Some attacks on smart contracts have resulted in substantial losses. For example, the “DAO attack” on Ethereum, where \$50 million worth of cryptocurrency was stolen due to a re-entrancy vulnerability¹. In April 2020, an attacker exploited a re-entrancy bug in the Lendf.me platform, resulting in a loss of about 99.5% of the platform’s funds (~\$25 million). In 2021 cryptocurrency-related crimes including smart contract attacks resulted in losses of approximately \$14 billion [6]. Hence, having a high assurance that a smart contract implementation is free of bugs is imperative. To address such issues, we are using the ConCert framework in the Coq proof assistant which facilitates formal verification and property-based testing of smart contracts.

¹ <https://www.wired.com/2016/06/50-million-hack-just-showed-dao-human/>





■ **Figure 1** The pipeline

42 Contributions.

43 In this paper, we present the details of the property-based testing functionality of the ConCert
 44 framework [3, 2]. We present three case studies demonstrating how ConCert can be used to
 45 find real-world bugs in smart contracts.

46 The first two case studies show how ConCert could have been used to find bugs that
 47 were found in smart contracts by auditors and attackers. The last case study shows how we
 48 used ConCert to find new bugs which could have led to upwards of \$8 million being stolen or
 49 frozen.

50 2 ConCert Overview

51 In this section, we give a brief overview of the ConCert framework, focusing on the smart
 52 contract execution layer and property-based testing. ConCert is open-source, and available
 53 at <https://github.com/AU-COBRA/ConCert/>.

54 2.1 Pipeline

55 The pipeline overview is presented in Figure 1. We start by developing a smart contract
 56 as a function in Coq using the ConCert infrastructure. We then can write a specification
 57 and test the smart contract function semi-automatically against it, using the integration with
 58 QuickChick [8]. With more effort, we can also prove the properties of smart contracts using
 59 the ConCert infrastructure. Proofs and tests crucially use the execution layer to reason about
 60 interacting contracts (see more details in Section 2.2), which enables us to capture properties
 61 beyond the mere functional correctness of a single contract invocation (see Section 3).

62 After testing and verification, one can obtain an executable implementation in one of
 63 the supported smart contract languages through *code extraction*. Our development uses
 64 the verified erasure procedure of MetaCoq [9] with verified optimisations and certifying
 65 pre-processing of ConCert. This gives us a code-generation procedure with strong correctness
 66 guarantees and a small trusted computing base of only MetaCoq and the pretty-printers into
 67 the target languages.

68 2.2 Smart Contract Execution Layer

69 The execution layer provides a model which facilitates reasoning about contract execution
 70 traces. This makes it possible to state and prove temporal properties of interacting smart
 71 contracts. Smart contracts in ConCert are modelled by abstracting a number of blockchains.²
 72 A contract consists of two functions:

73 ■ `init : Chain → ContractCallContext → Setup → option State`

74 The initialisation function is called after the contract is deployed on the blockchain. The
 75 first parameter of type `Chain` gives access to data about the blockchain (e.g. current chain
 76 height). The `ContractCallContext` parameter provides data about the current call (e.g.
 77 caller address, amount sent to the contract). `Setup` represents initialisation parameters.

78 ■ `receive : Chain → ContractCallContext → State → option Msg →`

79 `option (State * list ActionBody)` This function represents the main functionality of the
 80 contract that is executed for each call to the contract. `Chain` and `ContractCallContext` are
 81 the same as for `init`. The parameter of type `State` is the current state of the contract; `Msg`
 82 is a user-defined type of messages that contract accepts (the *entrypoints* of the contract).
 83 The result of a successful execution is a new state and a list of *actions* represented with
 84 `ActionBody`. The actions can be transfers, calls to other contracts (including itself), and
 85 contract deployments.

86 Both `receive` and `init` are ordinary Coq functions, making them convenient to reason
 87 about. However, reasoning about the contract functions in isolation is not sufficient, as
 88 many deployed contracts actually consist of a collection of interacting contracts, for example
 89 for the sake of modularity. One call to `receive` potentially emits more calls, which can
 90 create complex call graphs between deployed contracts. Therefore, it is necessary to consider
 91 execution traces to prove some safety properties of smart contracts. An execution trace
 92 `ChainTrace` is the reflexive, transitive closure of a proof-relevant `ChainStep` relation, which
 93 essentially captures the addition of a block to the blockchain. In this step, any *actions* (such
 94 as contract calls and transfers) coming from external users are executed.

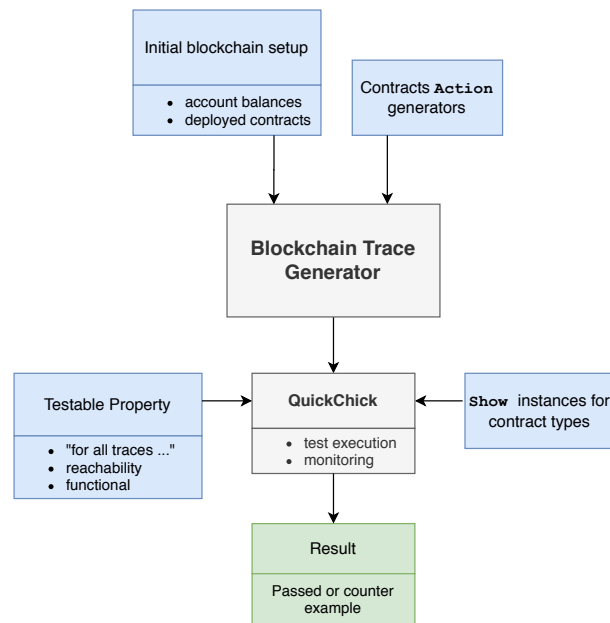
95 `ChainTrace` gives a relational operational semantics for the executions process. The
 96 semantics is non-deterministic since it allows for arbitrary execution order for the actions
 97 emitted by contract calls. Thus, ConCert provides two executable implementations: one
 98 follows depth-first and the other follows breadth-first order. It also provides proof that if
 99 running `add_block` succeeds, it results in a valid instance of `ChainTrace`. Having an executable
 100 implementation is crucial for property-based testing.

101 2.3 Property-based Testing framework

102 Property-based testing (henceforth abbreviated *PBT*), also known as *random-property testing*,
 103 is a technique for testing where test data is generated pseudo-randomly and tested in large
 104 quantities against some decidable property. We integrate the PBT library *QuickChick* [8] with
 105 the execution framework to obtain a method for testing contract executions. In particular, we
 106 support testing the functional correctness of contracts but also testing (decidable) properties
 107 of entire execution traces. The overview of the testing framework is given in Figure 2.

108 In brief, the PBT framework works by having the user provide *generators* for the `Msg`
 109 type of the contract(s) tested. In this context, generators are functions that produce pseudo-
 110 random values of the given type. These generators are used to populate randomly generated

² E.g. Concordium, Tezos, Dune, Æternity



■ **Figure 2** Property-based Testing in ConCert

111 execution traces with pseudo-random contract calls during testing with QuickChick. The
 112 user also configures the initial blockchain setup consisting of account balances and contracts
 113 that are currently available for interaction (deployed contracts). QuickChick also uses `Show`
 114 type class instances to print test results (e.g. counterexamples).

115 For example, consider how to test a token contract whose `Msg` type is

```

116
117 Inductive Msg :=
118   transfer of (address * address * nat)
119 | approve of (address * address * nat).
120
  
```

121 That is, it has two entrypoints: one for transferring tokens between the two given addresses
 122 and one for approving an address to spend a given number of tokens on behalf of another
 123 address. Generating pseudo-random values of `Msg` then amounts to either generating a
 124 `transfer` or an `approve`, and populating it with parameters by using the generators for
 125 `address` and `nat`. We can either implement this manually or have QuickChick automatically
 126 derive such a generator³. Note that we might prefer to implement this manually since
 127 we might want to ensure that the number of tokens to be transferred in `transfer` is never
 128 larger than the balance of the sender. We provide various combinators to make it easy and
 129 convenient to implement complex generators.

130 Suppose we want to test that `transfer` updates the internal balances correctly. In ConCert,
 131 this functional correctness property is specified by using pre- and post-conditions. Testing
 132 such a property with QuickChick could look like

```

133 QuickChick ({{msg_is_transfer}} Token.receive {{transfer_correct}}).
134
  
```

136 The code above states that if the incoming message is a transfer, then after executing the token

³ Due to limitations of QuickChick, the `Derive` command fails for some parameterised inductive types, e.g. `Msg` type in implicitly parameterised with some blockchain configuration. We have reported this issue: <https://github.com/QuickChick/QuickChick/issues/286>

137 contract's `receive` function, its state should be consistent with a predicate `transfer_correct`.
 138 By default, QuickChick will generate 10.000 inputs and test that the property is satisfied
 139 in all of them, or otherwise report a counterexample. The counterexamples reported are
 140 automatically minimized by the PBT framework to produce smaller counterexamples that
 141 are easy to understand. From our experience, these tests typically take less than a minute
 142 (see Section 7).

143 One can also test whether some state is reachable from the given state. For example, the
 144 following test

```
145 QuickChick (token_cb ~>> (person_has_tokens person_3 42)).
```

148 shows that from the state `token_cb` with three addresses participating in the token there is a
 149 state where `person_3` has 42 tokens. The corresponding trace is reported to the user.

150 3 Dexter decentralized exchange

151 In this section, we consider a bug in (an earlier version of) Dexter, a decentralized token
 152 exchange contract on the Tezos blockchain. The bug would have allowed an attacker to
 153 manipulate exchange rates to obtain unintended profit through a simple attack. The contract
 154 had previously been formally verified for functional correctness⁴. However, this bug can only
 155 be discovered when considering *execution traces* - that is, sequences of contract calls. We
 156 demonstrate how this bug can be found by testing a *natural* specification on traces. So, we
 157 argue that this bug would likely have been discovered when using ConCert as part of the
 158 specification process.

159 The Dexter exchange smart contract is used for exchanging tokens and `tez` (the on-chain
 160 currency of Tezos), it implements a so-called *constant-product market*, which means that the
 161 total value of the contract never decreases. A property of such markets is that the exchange
 162 rate cannot be significantly manipulated unless a party owns most of the market's assets [1].
 163 The rate at which tokens and `tez` can be exchanged is calculated dynamically at each trade
 164 according to the function

$$165 \text{getInputPrice}(Ts, Ts_{reserve}, Tez_{reserve}) = \frac{Ts \cdot 997 \cdot Tez_{reserve}}{Ts_{reserve} \cdot 1000 + Ts \cdot 997}$$

167 where Ts are the tokens being exchanged, $Ts_{reserve}$ is the reserve of tokens held by the
 168 Dexter contract, and $Tez_{reserve}$ is the contracts `tez` reserve.

169 One key property of constant-product markets, that cannot be verified from functional
 170 correctness alone, is that splitting trades is never profitable. Specifically, suppose a user
 171 trades N tokens for Z `tez`. Suppose this trade is split into $k > 1$ trades, totalling N tokens
 172 for a total of Z' `tez`. Then it should be the case that $Z' \leq Z$.

173 In ConCert, we can state this property by asserting that for each block added to generated
 174 traces, the total amount of `tez` gained from trades does not exceed what the user would have
 175 gained from trading the same amount of tokens in a single exchange. The full Coq definition
 176 can be found in `examples/dexter/DexterTests.v`

177 With this test, our PBT framework automatically finds a counterexample that violates
 178 the property. The counterexample show two consecutive exchanges; first trading 14 tokens
 179 for 5 `tez`, then 16 tokens for an additional 5 `tez`. However, the payout for a single trade of

⁴ <https://research-development.nomadic-labs.com/dexter-decentralized-exchange-for-tezos-formal-verification-work-by-nomadic-labs.html>

180 30 tokens would have been 9 tez, netting the user an extra one tez from splitting the trade.
 181 The vulnerability is due to a combination of Tezos' breadth-first execution model⁵ and the
 182 way the contract tracks its asset reserves. Concretely the problem is that in breadth-first
 183 both trades are executed before the actions emitted by the trades are executed, meaning
 184 that the second trade will start before the tez and tokens from the first trade have finished
 185 being transferred. The contract accounts for this by manually tracking the number of tokens,
 186 but fails to do the same for the tez reserve. Thus when the second trade starts the contract
 187 uses the wrong tez reserve for calculating the exchange rate. A strength of ConCert is that
 188 it allows testing in both depth-first and breadth-first execution order, running the same test
 189 with depth-first shows no vulnerability.

190 The bug was fixed prior to the deployment of Dexter.

191 4 iToken

192 In this section, we show how the bZx iToken smart contract was compromised and how
 193 ConCert could have discovered this vulnerability. The iToken smart contract is an interest
 194 accumulating ERC20 token used as part of the bZx decentralized finance platform. In Septem-
 195 ber 2020 an attacker stole \$8 million worth of cryptocurrency by exploiting a vulnerability
 196 in the iToken contract caused by a misplaced line of code⁶. This vulnerability was missed by
 197 two audits of the platform. The vulnerability was in the tokens `transferFrom`, which is used
 198 to transfer tokens between users. The transfer logic was implemented in the following way:

```
199
200 uint256 balanceFrom = balances[from];
201 uint256 balanceTo = balances[to];
202 balances[from] = balanceFrom.sub(amount);
203 balances[to] = balanceTo.add(amount);
```

205 This logic would have been safe had lines 2 and 3 been swapped. To see where this goes
 206 wrong, consider the case where `from = to`. In this case, the transferred amount would be
 207 subtracted from the sender's balance in line 3. However, in line 4 the original balance of the
 208 sender is used to add the transferred amount to the sender's balance, resulting in the sender
 209 ending gaining tokens through the self-transfer.

210 This bug could be found using the PBT framework by writing a test checking that the
 211 balance remains the same after a self-transfer. However, such a test would require knowledge
 212 of the possibility of a bug in this edge case. Instead, we formulate the property that *the*
 213 *sum of all balances should remain unchanged after a call*, with the exception of minting and
 214 burning calls. In ConCert testing such a property looks like:

```
215
216 Definition msg_is_not_mint_or_burn state msg :=
217   match msg with
218   | mint _ | burn _ => false
219   | _ => true
220   end.
221 Definition sum_balances_unchanged chain cctx (old_state : State) (msg : Msg)
222   (result : option (State * list ActionBody)) : bool :=
223   let balances_sum state := sum s.(balances) in
224   match result with
225   | Some (new_state, _) => balances_sum old_state =?= balances_sum new_state
226   | None => true (* Return true in the case that nothing changed *)
227   end.
```

⁵ Tezos moved to depth-first execution order after Dexter was developed

⁶ <https://bzx.network/blog/incident>

```
230 QuickChick ({{msg_is_not_mint_or_burn}} iTokenContract {{sum_balances_unchanged}})
```

```
🔗 examples/iTokenBuggy/iTokenBuggyTests.v:sum_balances_unchanged
```

231 By running the test, we indeed obtain a minimal counterexample showing that self-transfers
232 violate the property.

233 5 Basic Attention Token

234 In this section, we show how ConCert was used to find new bugs, that were missed by several
235 audits, in the Basic Attention Token (BAT) smart contract. BAT is an Ethereum initial coin
236 offering smart contract developed by Brave. It is a combination of an ERC-20 token and a
237 crowdsale contract, where users can fund ether to Braves' project in return for BAT tokens.
238 The crowdsale runs for a fixed amount of blocks, after which the funding either succeeds or
239 fails. If funding succeeds, Brave receives all the ether raised. If it fails, all users can claim
240 a refund of their ether by burning their tokens. As the contract owners, Brave get a fixed
241 amount of free tokens to spend.

242 We test functional correctness using a similar Hoare triple test as shown in Section 2.3.
243 In addition, we formulated five key safety properties.

- 244 1. **Funding is final:** Once the contract enters its funded state it cannot leave it again.
- 245 2. **Funding possible:** If there is enough ETH in the blockchain to reach the funding goal,
246 then it should be possible to reach a state in which the funding succeeded.
- 247 3. **No refunding for owners:** The free tokens given to the owners should not be refundable.
- 248 4. **Refund guarantee:** There should always be enough ETH in the contract balance to
249 refund all funded tokens. Unless funding succeeded.
- 250 5. **No frozen funds:** It should always be possible to completely drain the contract balance,
251 so no ETH gets permanently frozen.

252 Through testing, we found that only the first property holds. Most of the bugs occur from
253 combining token and crowdsale functionality and both parts behave safely on their own. *This*
254 *highlights that composing contracts is nontrivial and can easily introduce subtle bugs.*

255 5.1 Test Setup

256 In Sections 3 and 4 we showed that ConCert could find known bugs. For those, it was not so
257 important whether the generators would cover the entire input space. However, when testing
258 a complex contract with the purpose of finding potentially unknown bugs, it is crucial to
259 have good generators. A good quality generator should be able to cover the entire input
260 space of the smart contract and have a good balance between generating calls that succeed
261 and calls that fail. Using automatically derived generators will often result in too many
262 failing calls for complex smart contracts. For testing BAT we take the approach of combining
263 manually written generators designed to only produce valid calls with generators that are
264 likely to produce invalid calls. That is, for each entrypoint, we define two generators. This
265 is illustrated in Figure 3. The `finalize` entrypoint is an entrypoint that transitions the
266 contract from funding to the funded state. It can only be called by the owner after funding
267 succeeds. The first generator `gFinalize` only produces calls that we expect to succeed, while
268 the `gFinalizeinvalid` generator will generate calls with an arbitrary sender, which is unlikely
269 to be valid. The generators for potentially invalid calls can be automatically derived using
270 QuickChick. All the generators are combined into a single call generator.

271 This approach gives us a generator that can cover the entire input space while still
272 allowing us to tune the distribution of valid and invalid calls to different entrypoints. Using

```

Definition gFinalize env contract_state : G (option (Address * Msg)) :=
  if (isFullyFunded env contract_state) (* Check if funding succeeded *)
  then returnGen (Some (fund_addr, finalize)) (* Call finalize from owner address *)
  else returnGen None. (* Don't return call if not funded *)
Definition gFinalizeInvalid env contract_state : G (Address * Msg) :=
  sender ← gAddress ;; (* Generate arbitrary address *)
  returnGen (sender, finalize).

```

 `examples/bat/BATGens.v:gFinalize`

■ **Figure 3** Generators for the `finalize` endpoint

273 the PBT framework we can measure statistics about the generator and use that to tune the
274 distribution.

275 5.2 Finding Vulnerabilities

276 We test each of the five safety properties for the BAT contract defined in Section 5. Here we
277 detail a few of the tests.

278 A key property is that the contract doesn't deadlock, i.e. with enough user support it
279 should always be possible to reach the funded state. Since ConCert can test reachability
280 of states we can easily state this property by combining the reachability checker with a
281 deployment configuration generator. The following test states that for any BAT deployment
282 configuration there should exist a trace from the state where BAT is deployed with that
283 configuration to a state where the contract is funded.

```

284 QuickChick (forall gBATSetup (build_init_cb (fun cb => cb ~> is_finalized))).
285
286

```

 `examples/bat/BATTests.v`

287 Here `gBATSetup` is the configuration generator, `build_init_cb` builds an initial state with the
288 contract deployed, and `is_finalized` checks for a given blockchain state if the contract is
289 funded. By running the test, we obtain counterexamples showing four classes of configurations
290 where the contract cannot be fully funded. One of them is the case where the funding period
291 is empty or already over at the time of deployment. Ideally, the contract should have included
292 a check at deployment preventing such configurations.

A crucial safety property is that any user who donated should be guaranteed their money
back in case of failed funding. By testing the functional correctness of endpoints, we
already know that the contract will always refund the correct amount and will always succeed,
given that the contract has enough funds. Therefore, testing refund guarantee reduces to
checking that there is always enough funds to refund all tokens held by "real" users. Here we
distinguish between real users of the contract and the owner, because the owner's free tokens
should not be counted. That is, we want to test that the following is always true.

$$\text{contractBalance} \geq \frac{\text{totalTokenSupply} - \text{ownersTokens}}{\text{tokenExchangeRate}}$$

293 In ConCert a test of this looks like:

```

294
295 Definition contract_balance_lower_bound (cs : ChainState) :=
296   let contract_balance := env_account_balances cs contract_base_addr in
297   (* Get BAT contract state *)
298   match get_contract_state State cs contract_base_addr with
299   | Some cstate =>
300     (* Get token balance of owner *)

```



```

301   let bat_fund_balance := with_default 0 (FMap.find batFund (balances cstate)) in
302   if cstate.isFinalized
303   then checker true (* Case where refunds are not permitted *)
304   (* Assert that there is enough ETH to refund all tokens held by "real" users *)
305   else checker (Z.gcb contract_balance
306   (Z.of_N (((total_supply cstate) - bat_fund_balance) / cstate.tokenExchangeRate)))
307   | None => checker true (* Case where contract isn't deployed *)
308   end.
309 QuickChick (forallChainState contract_balance_lower_bound)

```

 examples/bat/BATTests.v:contract_balance_lower_bound

Running the test we get the following minimized counterexample from the testing framework.

```

312 Chain{
313   Block 1 [Action{act_from: 10, act_body: (act_deploy 0, Setup{...})}];
314   Block 2 [Action{act_from: 17, act_body: (act_call 128, 0, transfer 16 14)}]
315 }
316

```

This counterexample shows a trace where the BAT contract is deployed in the first block, after which the owner (address 17) immediately transfers some of its free tokens to another user. This is possible because the contract combines crowdsale and token contract behaviour. This violates two of the safety properties because nothing is preventing the second user from refunding the transferred tokens. Thus it is possible for the free tokens given to the owner to be refunded by first transferring them. This also breaks the property that all real users should be guaranteed a refund because if the owner refunds some of the free tokens then there is no longer enough ETH to refund all tokens held by real users.

The remaining safety properties were tested using similar methods.

6 Related Work

Various testing approaches have been applied to smart contracts. Tools like Truffle⁷ for Ethereum or SmartPy⁸ for Tezos mostly cover conventional unit testing that can be insufficient. The testing framework for LIGO⁹ supports unit testing and mutation testing. However, none of the conventional testing frameworks offers a possibility for generating random traces and testing properties of interacting contracts. We will now focus on works using fuzzing/property-based testing techniques.

The closest to our work is the property-based testing framework for the Tezos' Michelson language. The framework utilises QCheck, a QuickCheck-inspired property-based testing framework for OCaml. QCheck was extended by Nomadic Labs with the ability to generate arbitrary sequences of Liquidity Baking contract calls. The contract is manually reimplemented in OCaml and serves as a model for the original contract. The model implementation is then validated against the original contract through the actual Tezos execution model. The development is tailored to the Liquidity Baking contract and is not connected to the Michelson formalisation in Coq Mi-Cho-Coq [4]. We are currently collaborating with the Mi-Cho-Coq team on integrating ConCert with the formalisation of Michelson.

For the Ethereum blockchain, several works are using randomised testing techniques. Echidna [7] and Brownie¹⁰ use fuzzing-like techniques for testing smart contracts. The common challenge for this approach is that randomly generated transaction data might

⁷ <https://trufflesuite.com/>

⁸ <https://smartpy.io/docs/scenarios/testing/>

⁹ <https://ligolang.org/docs/advanced/testing>

¹⁰ Property-based testing framework for EVM: <https://github.com/eth-brownie/brownie>

not be enough to ensure good coverage. This is especially problematic in the case of smart contract interactions, since the whole sequence (trace) of actions must be generated. Echidna uses coverage-driven feedback to automatically tune the testing parameters. Brownie uses unit-test like tests with user-defined generators for randomising inputs to contract calls in the tests. Brownie does not generate calls or execution traces, which limits the types of bugs that it can find. In our approach, instead of tuning pre-defined parameters, we allow users to define generators that produce random data with fewer discarded tests. For simple cases, data generators can be derived automatically using the QuickChick infrastructure.

The EthPloit project [10] generates possible exploits using fuzzing techniques. The exploits are split into three categories. For each of these categories, a special exploit detector oracle is used to report an exploit. For example, the Balance Increment oracle compares the overall initial balance of attackers’ accounts with the current balance after a series of transfers and reports, if the balance of the attackers’ accounts increases. EthPloit utilises static analysis to focus attention on particular variables and functions. The input for selected functions is generated randomly, or chosen using a seed set. The seed sets are used to provide runtime feedback. This improves the fuzzing efficiency by exploiting the results of previous runs. In our approach, the users specify the properties to test, instead of searching for particular categories of exploits. Violation of such properties is reported as a counterexample, which points to vulnerabilities. The pure/functional nature of our smart contracts avoids many pitfalls and simplifies reasoning about smart contracts. When compared to effectfull languages, such as Solidity, static analysis is less urgent.

Finally, the `cooked-validators` library¹¹ for the Plutus smart contract language [5] facilitates property-based testing with arbitrary transaction sequences. Note, however, that the execution model for Plutus does not involve on-chain inter-contract communication.

7 Evaluation

We evaluate our framework in terms of usability, specifically regarding bug-finding capabilities. We demonstrated the testing framework on three concrete examples in the previous section, showing that it can find different types of real-world bugs. The vulnerabilities had a wide range of causes: the execution order, complex contract-to-contract interactions and the evolution of the contract state. Such bugs would not have been detected in other tools considering only functional correctness. This highlights ConCert’s unique capability of modelling and testing complex contract interactions. We have tested various other smart contracts, such as a reference implementation of the ERC-20 Token¹², and re-discovered known bugs, thus supporting the claim that our framework is effective at finding bugs.

Since we have the full power of Coq at our disposal, we can effectively test any *decidable* property on the `Chain` type. Hence, there are few limitations in terms of expressiveness. We also emphasise that once contracts are implemented (in ConCert) and the executable specifications are written (i.e. the decidable properties to be proven or tested), the only prerequisite for automatically testing the specifications is to implement the action generators and show instances, as discussed in Subsection 2.3. Implementing these requires only some expertise with Coq and QuickChick, and can in some cases be derived automatically. Hence, the setup is relatively simple, only requiring little extra effort compared to writing traditional tests.

¹¹ <https://iohk.io/en/blog/posts/2022/01/27/simple-property-based-tests-for-plutus-validators/>

¹² <https://github.com/AU-COBRA/ConCert/tree/master/examples/eip20>

389 Additionally, the feedback loop from executing tests is fast, making it easy to use
 390 during the contract development process. In our experience, QuickChick will usually report
 391 counterexamples, if they exist, within 1-2 seconds and otherwise report that all inputs (by
 392 default 10.000) passed — usually in than 5-10 seconds (for traces of 14 calls). Of course, the
 393 time depends on many factors, most importantly, the length of traces and the complexity of
 394 generators and contracts.

395 **8** Conclusions

396 We have presented the ConCert Coq framework for testing, verifying and extracting smart
 397 contracts. We have demonstrated the framework for property-based testing on three smart
 398 contracts using it to discover vulnerabilities used in previous attacks and new bugs that
 399 could have led to millions of dollars stolen or frozen. As stated in the previous section, the
 400 vulnerabilities had a wide range of causes covering the most common causes of flaws in smart
 401 contracts.

402 We have re-discovered several bugs in real-world contracts (not presented in this paper),
 403 such as the \$50 million “DAO attack” on Ethereum, and tested reference implementations of
 404 ERC-20 and FA2 Token Standards, common standards for tokens used in several blockchains¹³.

405 Hence, our approach to testing smart contracts scales to real-world contracts and is
 406 capable of finding significant bugs. Contracts in ConCert are extractable to Concordium’s
 407 Rust framework, Liquidity, and CameLIGO. Thus in total, we have a toolchain for producing
 408 executable code for smart contracts that are tested and verified. The importance of combined
 409 auditing, testing and verification is also starting to be recognized by the industry.¹⁴

410 Acknowledgements

411 We would like to thank the LIGO team and in particular Tom Jack, Raphaël Cauderlier,
 412 Exequiel Rivas, Rémi Lesénéchal, Gabriel Alfour, Thomas Letan and Arvid Jakobsson for
 413 the discussions about testing for LIGO. This research was partially supported by a grant
 414 from Nomadic Labs and by the Concordium Blockchain Research Center.

415 ——— References ———

- 416 1 Guillermo Angeris, Hsien-Tang Kao, Rei Chiang, Charlie Noyes, and Tarun Chitra.
 417 An Analysis of Uniswap markets. *Cryptoeconomic Systems*, 1(1), 2021. URL: [https://](https://cryptoeconomicsystems.pubpub.org/pub/angeris-uniswap-analysis)
 418 cryptoeconomicsystems.pubpub.org/pub/angeris-uniswap-analysis, doi:[https://doi.org/](https://doi.org/10.21428/58320208.c9738e64)
 419 [10.21428/58320208.c9738e64](https://doi.org/10.21428/58320208.c9738e64).
- 420 2 Danil Annenkov, Mikkel Milo, Jakob Botsch Nielsen, and Bas Spitters. Extracting Smart
 421 Contracts Tested and Verified in Coq. CPP 2021. Association for Computing Machinery, 2021.
 422 doi:10.1145/3437992.3439934.
- 423 3 Danil Annenkov, Jakob Botsch Nielsen, and Bas Spitters. ConCert: A Smart Contract
 424 Certification Framework in Coq. In *CPP’2020*, 2020. doi:10.1145/3372885.3373829.
- 425 4 Bruno Bernardo, Raphaël Cauderlier, Zhenlei Hu, Zhenlei Pesin, and Julien Tesson. Mi-Cho-
 426 Coq, a framework for certifying Tezos Smart Contracts. FMBC19, 2019.
- 427 5 James Chapman, Roman Kireev, Chad Nester, and Philip Wadler. System F in Agda, for fun
 428 and profit. In *MPC’19*, 2019. doi:10.1007/978-3-030-33636-3_10.

¹³ <https://github.com/AU-COBRA/ConCert>

¹⁴ e.g. <https://forum.cardano.org/t/cip-proposal-cardano-audit-best-practice-guidelines/100022>

- 429 6 Kim Grauer, Will Kueshner, and Henry Updegrave. Chainalysis 2022 Crypto Crime Report.
430 2022. URL: <https://go.chainalysis.com/2022-Crypto-Crime-Report.html>.
- 431 7 Gustavo Grieco, Will Song, Artur Cygan, Josselin Feist, and Alex Groce. Echidna: effective,
432 usable, and fast fuzzing for smart contracts. In *Proceedings of the 29th ACM SIGSOFT*
433 *International Symposium on Software Testing and Analysis*, pages 557–560, 2020. doi:
434 10.1145/3395363.3404366.
- 435 8 Zoe Paraskevopoulou, Catalin Hritcu, Maxime Dénès, Leonidas Lampropoulos, and Benjamin C.
436 Pierce. Foundational property-based testing. In Christian Urban and Xingyuan Zhang, editors,
437 *6th International Conference on Interactive Theorem Proving (ITP)*, volume 9236 of *Lecture*
438 *Notes in Computer Science*, pages 325–343. Springer, 2015. doi:10.1007/978-3-319-22102-
439 1_22.
- 440 9 Matthieu Sozeau, Abhishek Anand, Simon Boulter, Cyril Cohen, Yannick Forster, Fabian
441 Kunze, Gregory Malecha, Nicolas Tabareau, and Théo Winterhalter. The MetaCoq Project.
442 *Journal of Automated Reasoning*, Feb 2020. doi:10.1007/s10817-019-09540-0.
- 443 10 Qingzhao Zhang, Yizhuo Wang, Juanru Li, and Siqu Ma. EthPloit: From Fuzzing
444 to Efficient Exploit Generation against Smart Contracts. In *2020 IEEE 27th Interna-*
445 *tional Conference on Software Analysis, Evolution and Reengineering (SANER)*, 2020.
446 doi:10.1109/SANER48275.2020.9054822.