

1 Towards Formally Specifying and Verifying Smart 2 Contract Upgrades in Coq

3 Derek Sorensen   

4 University of Cambridge, *Department of Computer Science and Technology*, UK

5 — Abstract —

6 Smart contract upgrades are costly from a verification perspective and can be a meaningful source
7 of vulnerabilities when done incorrectly. Unfortunately, there is no established, formal framework
8 through which one can reason about contracts as they undergo upgrades, though much work has
9 been done to verify standalone smart contracts. Instead, one must repeat the full verification process
10 for each contract upgrade, something which relies heavily on fallible intuition, can lead to unexpected
11 vulnerabilities, and drives up the cost of formally verifying smart contracts. We propose a formal
12 framework for contract upgrades in ConCert, a Coq-based smart contract verification tool. Central
13 to this framework is our notion of a *contract morphism*, a theoretical tool which we introduce to
14 formally encode structural relationships between smart contracts, and with which we can formally
15 specify and verify an upgraded contract relative to its previous versions. We argue that ours is
16 a natural framework for specifying and verifying contract upgrades, and hope to offer a first step
17 towards rigorous, efficient specification and verification of contract upgrades.

18 **2012 ACM Subject Classification** Theory of computation → Program verification

19 **Keywords and phrases** smart contract verification, formal methods, interactive theorem prover,
20 smart contract upgrades

21 **Digital Object Identifier** 10.4230/OASlcs.FMBC.2024.1

22 **Supplementary Material** *Software*: <https://github.com/differentialsderek/FinCert>

23 **1** Introduction

24 Faulty upgrades are a meaningful source of smart contract vulnerabilities. Costly attacks
25 such as those on Uranium Finance (2021) [8], NowSwap (2021) [4], and Nomad (2022) [7, 9],
26 totaling 241 million USD in lost assets, are a few of many examples of contracts attacked
27 after an erroneous upgrade. Furthermore, because verifying software is time, labor, and
28 resource intensive, it can be difficult to justify formally verifying software which may be
29 upgraded quickly or frequently—a problem shared with other verified software, *e.g.* [16, 21].
30 Both of these factors limit the effectiveness of formal methods to address security issues in
31 real-world software, inhibiting verification as business and security propositions [18].

32 What is needed is a practical and formal framework through which to specify and verify
33 contract upgrades. As it stands we have no such framework apart from repeating the formal
34 specification and verification process on a new contract version. Not only are upgrades costly
35 from a verification perspective, as we have no good way of reusing much of the verification
36 work on previous contract versions, but incorrect specifications are themselves a meaningful
37 source of contract vulnerabilities [19]. Thus each time a specification is made from scratch
38 we risk introducing errors of incorrect specification.

39 To mitigate these issues we introduce a formal framework for specifying and verifying
40 contract upgrades, through which we can reuse formal specification and proof on previous
41 contract versions. This framework relies on the notion of a *contract morphism*, a theoretical
42 tool we introduce that formally encodes structural relationships between smart contracts,
43 and with which we can specify and reason about the structure and behavior of an upgraded
44 contract relative to its previous versions. We argue that this is a natural framework for



© Derek Sorensen;

licensed under Creative Commons License CC-BY 4.0

5th International Workshop on Formal Methods for Blockchains (FMBC 2024).

Editors: Bruno Bernardo and Diego Marmosler; Article No. 1; pp. 1:1–1:14

OpenAccess Series in Informatics



OASICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

45 specifying and verifying contract upgrades, one which could decrease the cost of formally
46 verifying contract upgrades as well as the risk of introducing vulnerabilities due to incorrect
47 specification.

48 We proceed as follows. In §2, we survey related work. In §3, we introduce *contract*
49 *morphisms* as a formal tool to specify and verify contract upgrades. In §4 we give two
50 examples of formally specifying a contract upgrade with contract morphisms. In §5 we
51 discuss formal verification with contract morphisms. We conclude in §6.

52 **2** Related Work

53 In the realm of smart contracts there is limited formal work on formal reasoning about
54 contract upgrades. Previous work [3, 6] proposes paradigm-shifting methods to either attach
55 formal proofs to smart contracts and their upgrades, which are verified by the chain, or to
56 trust a canonical third party to verify all contract upgrades before deployment. Unfortunately
57 this work is likely impractical, as both solutions require substantial paradigm shifts or re-
58 engineering of blockchain ecosystems. The latter also arguably contradicts the permissionless
59 ethos of blockchain ecosystems by mandating a trusted third party.

60 In the context of software more generally, much work has gone into ensuring that software
61 upgrades are carried out safely with formal methods [10, 12, 21]. Recent work has begun
62 to address the issue of adapting formal proofs in a proof assistant to changes in software in
63 order to lower the cost of formally verified software which may undergo regular upgrades [16].
64 This problem is complicated by the computable nature of proofs in proof assistants like Coq;
65 chosen data types strongly influence the structure of proofs, making adaptation difficult [11].
66 A notable contribution to this work is Ringer *et al.*'s work on *proof repair* [14, 15], which
67 seeks to relate a new program version to the old—by type equivalences or by comparing
68 inductive structures—and thereby reuse previously-completed proofs on the updated code.

69 Drawing on this previous work, particularly Ringer *et al.*'s idea of reusing formal proofs by
70 way of structural similarities between programs, our goal is to provide a framework for using
71 formal methods to formally specify and verify smart contract upgrades. Contract morphisms
72 (§3) will be our primary theoretical tool for specifying and verifying contract upgrades. Their
73 purpose is to formally encode a structural relationship between smart contracts which can
74 be used for both formal specification and proof reuse. With contract morphisms we address
75 the problem of formal reasoning about contract upgrades, but in contrast to previous work
76 on the subject our proposed framework does not require the paradigm-shifting reengineering
77 of blockchain systems in order to be used.

78 Finally, we note that for smart contracts there is a distinction between contract upgrades
79 and contract *upgradeability*. Some contracts come with a predefined logic to handle upgrades
80 and avoid hard forks, the most popular of these on Ethereum being the Diamond framework
81 [13]. However, they are complicated contracts as their specifications include the upgradeability
82 functionality and governance, as well as the functionality of a given version of the contract.
83 We will only consider upgrades via hard forks in this paper, leaving the question of rigorous
84 formal specification and verification of upgradeable contracts to future work.

85 **3** Contract Morphisms

86 In what follows we define *contract morphisms*, a theoretical tool which codifies formal
87 relationships between smart contracts. In later sections we use them to formally specify and
88 verify contract upgrades. We argue that this provides our desired formal framework.

89 3.1 Morphisms of Pure Functions

90 Before focusing on the specific case of smart contracts, we consider the more general case
 91 of programs formalized as pure functions. Take types A, A' and B, B' , and two functions
 92 $p : A \rightarrow B$ and $q : A' \rightarrow B'$. A *morphism* from p to q is a or a pair of functions f_i and f_o
 93 which form a commutative square,

$$94 \begin{array}{ccc} A & \xrightarrow{f_i} & A' \\ p \downarrow & \cong & \downarrow q \\ B & \xrightarrow{f_o} & B' \end{array}$$

i.e. for which

$$q \circ f_i = f_o \circ p.$$

Together, we call f_i and f_o the morphism

$$f : p \rightarrow q.$$

95 Via f_i and f_o , the commutative square like the above maps inputs and outputs of p to
 96 inputs and outputs of q . If p and q are programs (in particular, pure functions), we can also
 97 interpret this as execution traces of p to execution traces of q , such that transforming the
 98 inputs of p into those of q with f_i , and then applying q is the same as applying p first and
 99 then transforming the outputs over f_o .

We can define composition of morphisms easily as the composition of commutative squares.
 That is, given functions p, q , and r , and morphisms

$$f' : p \rightarrow q \text{ and } f'' : q \rightarrow r,$$

100 we can define a morphism $f := f'' \circ f' : p \rightarrow r$ by the outer square of the following diagram,

$$101 \begin{array}{ccccc} A & \xrightarrow{f'_i} & A' & \xrightarrow{f''_i} & A'' \\ p \downarrow & \cong & \downarrow q & \cong & \downarrow r \\ B & \xrightarrow{f'_o} & B' & \xrightarrow{f''_o} & B'' \end{array}$$

102 which is commutative if each of the inner squares are commutative. Note that composition is
 103 associative, assuming the underlying functions are associative, and that we have the obvious
 104 identity morphism $f_{\text{id}} : p \rightarrow p$ given by $f_i, f_o := \text{id}$,

$$105 \begin{array}{ccc} A & \xrightarrow{\text{id}} & A \\ p \downarrow & \cong & \downarrow p \\ B & \xrightarrow{\text{id}} & B \end{array}$$

106 which commutes trivially. Thus given a well-defined class of functions, which in our case will
 107 be smart contracts modeled in Coq by pure functions, we have a category on those functions
 108 with morphisms given by commutative squares on those pure functions.

109 In the coming sections, given a morphism $f : p \rightarrow q$, we might consider the case that q
 110 is an upgraded version of p . Because f relates execution traces of q to those of p , we will
 111 see this can be used to reason formally about q in terms of p , both in specification and
 112 verification.

113 **3.2 Contract Morphisms in ConCert**

114 In ConCert, a Coq-based tool for smart contract verification which models the execution
 115 semantics of third-generation blockchains [2] and features verified extraction to various
 116 blockchains [1], smart contracts are formalized with a `Contract` type as a pair of pure, stateful
 117 functions `init` and `receive`. The `init` function governs contract initialization and the `receive`
 118 function governs contract calls. The `Contract` type is polymorphic, parameterized by four
 119 types: `Setup`, `Msg`, `State`, and `Error` which, respectively, govern the data necessary for contract
 120 initialization, contract calls, contract storage, and contract errors.

121 For a contract

122
$$C : \text{Contract } \text{Setup } \text{Msg } \text{State } \text{Error}$$

123 the type signatures of each component function (`init C`) and (`receive C`) are given as follows,
 124 where the types `Chain` and `ContractCallContext` are ConCert-specific types used to model
 125 the underlying blockchain and context.

```
126 init C : Chain → ContractCallContext → Setup → result State Error.
127
128
129 receive C : Chain → ContractCallContext → State → option Msg →
130 result (State * list ActionBody) Error.
```

■ **Listing 1** Type signature of the `init` and `receive` functions, respectively, of a smart contract in ConCert.

132 Now consider contracts `C1` and `C2`,

133
$$C1 : \text{Contract } \text{Setup1 } \text{Msg1 } \text{State1 } \text{Error1}$$

134
$$C2 : \text{Contract } \text{Setup2 } \text{Msg2 } \text{State2 } \text{Error2}.$$

135 We define a data type of *morphisms* between contracts `C1` and `C2`,

136
$$\text{ContractMorphism } C1 \ C2.$$

137 This data type consists firstly of four *component functions* between the contract types of `C1`
 138 and `C2`—the `Setup`, `Msg`, `State`, and `Error` types respectively.

```
139 ■ setup_morph : Setup1 -> Setup2
140 ■ msg_morph : Msg1 -> Msg2
141 ■ state_morph : State1 -> State2
142 ■ error_morph : Error1 -> Error2.
```

143 We can use these component functions to make commutative squares like those we saw in §3.1
 144 for each of the `init` and `receive` functions. For `init`, the horizontal arrows of the squares are
 145 given by the functions `mA_init` and `mB_init`. For `receive`, the horizontal arrows are given
 146 by the functions `mA_recv` and `mB_recv`. See Listing 2 for the definition of these functions in
 147 terms of the four component functions given above.

$$\begin{array}{ccc}
 A_{\text{init}} & \xrightarrow{mA_init} & A'_{\text{init}} \\
 \text{init} \downarrow & \parallel & \downarrow \text{init}' \\
 B_{\text{init}} & \xrightarrow{mB_init} & B'_{\text{init}}
 \end{array}
 \qquad
 \begin{array}{ccc}
 A_{\text{recv}} & \xrightarrow{mA_recv} & A'_{\text{recv}} \\
 \text{receive} \downarrow & \parallel & \downarrow \text{receive}' \\
 B_{\text{recv}} & \xrightarrow{mB_recv} & B'_{\text{recv}}
 \end{array}$$

148

```

(* functions to form a commutative square on init *)
mA_init :=
  fun (c : Chain) (ctx : ContractCallContext) (s : Setup) =>
    (c, ctx, setup_morph s).
mB_init := fun (res : result State Error) =>
  match res with
  | Ok init_st => Ok (state_morph init_st)
  | Err e => Err (error_morph e)
  end.

(* functions to form a commutative square on receive *)
mA_recv := fun (c : Chain) (ctx : ContractCallContext)
  (st : State) (op_msg : option Msg) =>
  (c, ctx, state_morph st, option_map msg_morph op_msg).
mB_recv := fun (res : result (State * list ActionBody) Error) =>
  match res with
  | Ok (init_st, nacts) => Ok (state_morph init_st, nacts)
  | Err e => Err (error_morph e)
  end.

```

■ **Listing 2** The functions which we use for the horizontal arrows of a pair of commutative squares $f_init : \text{init } C1 \rightarrow \text{init } C2$ and $f_recv : \text{receive } C1 \rightarrow \text{receive } C2$, respectively, in the definition of a contract morphism.

149 The functions defined above give us squares, but to finish the definition of contract
 150 morphisms we need these squares to commute. Thus our definition includes two coherence
 151 conditions, one for the `init` square and one for the `receive` square, which are given as follows.

```

152
153 (* The coherence condition that makes the init square commute *)
154 init_coherence: forall c ctx s,
155 (match (init C1 c ctx s) with
156   | Ok init_st => Ok (state_morph init_st)
157   | Err e => Err (error_morph e)
158   end) =
159 (init C2 c ctx (setup_morph s)).
160
161 (* The coherence condition that makes the receive square commute *)
162 recv_coherence : forall c ctx st op_msg,
163 (match (receive C1 c ctx st op_msg) with
164   | Ok (new_st, new_acts) => Ok (state_morph new_st, new_acts)
165   | Err e => Err (error_morph e)
166   end) =
167 (receive C2 c ctx (state_morph st) (option_map msg_morph op_msg)).
168

```

169 Thus a contract morphism

```

170
171           m : ContractMorphism C1 C2

```

171 is defined as a pair of commutative squares, each of which are morphisms between the
 172 respective `init` and `receive` functions of each contract. We give the formal definition of a
 173 contract morphism in Listing 3.

174 As the name *morphism* suggests, we should expect contract morphisms to behave like
 175 morphisms in a well-defined category. That is, we should have an associative composition
 176 operation on morphisms, and for every contract C should have an identity morphism

1:6 Smart Contract Upgrades in Coq

```

Record ContractMorphism
  (C1 : Contract Setup1 Msg1 State1 Error1)
  (C2 : Contract Setup2 Msg2 State2 Error2) :=
  build_contract_morphism {
    (* the components of a morphism *)
    setup_morph : Setup1 → Setup2 ;
    msg_morph   : Msg1   → Msg2   ;
    state_morph : State1 → State2 ;
    error_morph : Error1 → Error2 ;
    (* coherence conditions *)
    init_coherence : forall c ctx s,
      result_functor state_morph error_morph (init C1 c ctx s) =
      init C2 c ctx (setup_morph s) ;
    recv_coherence : forall c ctx st op_msg,
      result_functor (fun '(st, l) => (state_morph st, l))
        error_morph
        (receive C1 c ctx st op_msg) =
        receive C2 c ctx (state_morph st)
        (option_map msg_morph op_msg) ;
  }.

```

■ **Listing 3** The formal definition of a contract morphism in ConCert, consisting of four component functions and two coherence conditions, which together give a pair of commutative squares.

177 `id_C : ContractMorphism C C`

178 with which composition is trivial.

179 Indeed, this is the case. We can compose morphisms by composing the morphism
 180 component functions. We have two results,

181 `compose_init_coh` and `compose_recv_coh`,

182 which show that coherence of the composed morphism follows from the coherence conditions
 183 of each individual morphism. These results simply show that commutative squares compose,
 184 as we saw in §3.1, giving us a well-defined composition function `compose_cm`.

```

185 compose_cm : forall C1 C2 C3  

186 (g : ContractMorphism C2 C3) (f : ContractMorphism C1 C2) : ContractMorphism C1 C3.  

187  

188

```

189 We also have a proof that composition is associative, drawing on the associativity of component
 190 functions, and we have the obvious identity morphism, given by four identity component
 191 functions, such that composition with the identity is trivial.

```

192 Definition id_cm (C : Contract Setup Msg State Error) :  

193 ContractMorphism C C := {  

194   (* components *)  

195   setup_morph := id ;  

196   msg_morph   := id ;  

197   state_morph := id ;  

198   error_morph := id ;  

199   (* coherence conditions *)  

200   init_coherence := init_coherence_id C ;  

201   recv_coherence := recv_coherence_id C ;  

202 }  

203

```

205 This gives us a well-defined category **Contracts** of smart contracts, with objects given by
 206 the `Contract` type and morphisms given by the `ContractMorphism` type.

207 Note that in many categories, *e.g.* the categories of sets, topological spaces, or groups,
 208 morphisms are structure-preserving functions. So too for us. The existence of a morphism

209 $f : \text{ContractMorphism } C1 \ C2$

210 indicates a structural and mathematical relationship between contracts `C1` and `C2`, in particular
 211 relating their execution traces via the four component morphisms. As we will see, this
 212 relationship can be exploited to prove theorems about one contract in terms of another
 213 contract, something which we will do here in the case of contract upgrades and upgradeability.

214 In many categories there are also different classes of morphisms, including injections
 215 (embeddings, monomorphisms), surjections (quotients, epimorphisms), and isomorphisms.
 216 Injections, or embeddings, typically preserve the structure of the domain faithfully within
 217 the codomain, essentially identifying a copy of the domain within the codomain. Surjections
 218 typically represent a compression of some kind, and the information lost in the compression
 219 can frequently be described by a kernel object. As we will see, we also have injective and
 220 surjective contract morphisms, which are given when the four component functions are,
 221 respectively, injective or surjective, and which follow analogous intuitions.

222 **4 Morphisms to Formally Specify and Verify Contract Upgrades**

223 Our goal now is to use contract morphisms as a tool to formally specify and verify con-
 224 tract upgrades in ConCert. Consider a contract upgrade from the perspective of a formal
 225 specification. Contracts are usually upgraded with a goal that relates the new to the previ-
 226 ous contract version, whether it be to patch a bug, add functionality, or improve contract
 227 features. Thus the new specification relates to the old—it should eliminate a vulnerability
 228 but preserve all other functionality, be backwards compatible while adding functionality, or
 229 make improvements such as greater gas-efficiency without deviating from the behavior of
 230 the previous contract version. Of course, in practice an upgraded contract is not formally
 231 specified in relation to an older version, but rather by altering the old specification into the
 232 new, or simply starting from scratch and writing a new specification by hand. As discussed
 233 in §1, this can be a source of vulnerabilities.

234 In this section, we will formally specify contract upgrades in two examples using contract
 235 morphisms. The advantage of using morphisms is that we are able to clearly articulate
 236 the intent of an upgrade in the formal specification by way of a morphism in such a way
 237 that formal verification consists of producing a morphism between the updated contract
 238 implementation and a previous version which meets the required specification.

239 ► **Example 1 (Swap Contract Upgrade).** Consider a smart contract `C1` that prices and executes
 240 trades, *e.g.* a decentralized exchange (DEX) or an automated market maker (AMM) [22].
 241 Suppose that we wish to upgrade `C1` to a contract `C2` so that it calculates trades at higher
 242 precision by a factor of ten, meaning that the internal token balances in storage have one
 243 more decimal place, and the trade calculation is able to calculate at one decimal place greater
 244 in precision. Then in ConCert our contract `C1` will have a storage type which keeps track of
 245 internal token balances, exposed by a function `get_bal`.

246 `Context { storage : Type } { get_bal : storage → N }.`
 247
 248

249 It will also have a `TRADE` entrypoint which accepts a payload of some type, `trade_data`,
 250 characterized by an entrypoint type, `entrypoint`, and an associated typeclass, `Msg_Spec`.

1:8 Smart Contract Upgrades in Coq

```
251
252 Class Msg_Spec (T : Type) := {
253   (* the trade entrypoint *)
254   trade : trade_data → T ;
255   (* for any other entrypoint types *)
256   other : other_entrypoint → option T ;
257 }.
258
259 (* We assume an entrypoint conforming to Msg_Spec *)
260 Context { entrypoint : Type } { e_msg : Msg_Spec entrypoint }.
```

■ **Listing 4** We assume an entrypoint type `entrypoint`, characterized by a typeclass `Msg_Spec`, which includes a trade function `trade`.

262 Now assume that `C1` has some internal function `calculate_trade` that it uses to calculate
263 how many tokens will be traded out for a given contract call to the `TRADE` entrypoint. The
264 trade quantity, internal token balances, and the `calculate_trade` function will all be accurate
265 up to some decimal place, commonly 9 in the wild, formalized in the following specification,
266 `spec_trade`, of `C1`.

```
267
268 (* the specification of C1's trading functionality with regards to the
269    calculate_trade function *)
270 Definition spec_trade : Prop :=
271   forall cstate chain ctx trade_data cstate' acts,
272   (* for any successful call to C1's trade entrypoint, *)
273   receive C1 chain ctx cstate (Some (trade trade_data)) =
274   Ok(cstate', acts) →
275   (* the balance in storage updates according to the
276      calculate_trade function *)
277   get_bal cstate' =
278   get_bal cstate + calculate_trade (trade_qty trade_data).
```

■ **Listing 5** The formalized proposition that `C1` uses `calculate_trade` to price trades.

280 The property of Listing 5, `spec_trade`, is a specification with regards to which `C1` is assumed
281 to be correct.

282 Now we wish to upgrade `C1` to a new contract `C2` such that `C2` calculates trades and keeps
283 balances at one decimal place higher of accuracy. We will first have a specification for `C2`
284 which is analogous to `spec_trade` in Listing 5, which says that `C2` uses some new function,
285 `calc_trade_precise`, to calculate its trades.

```
286
287 (* The specification of C2's trading functionality with regards to the
288    calculate_trade_precise function. This is analogous to spec_trade *)
289 Definition spec_trade_precise : Prop :=
290   forall cstate chain ctx trade_data cstate' acts,
291   (* for a successful call to C2's trade entrypoint, *)
292   receive C2 chain ctx cstate (Some (trade trade_data)) = Ok (cstate', acts) →
293   (* the balance in storage updates according to the
294      calculate_trade_precise function *)
295   get_bal cstate' =
296   get_bal cstate +
297   calculate_trade_precise (trade_qty trade_data).
```

■ **Listing 6** The formalized proposition that `C2` uses `calculate_trade_precise` to price trades.

299 Our goal now is to use a contract morphism to complete the formal specification of C_2 in
 300 terms of C_1 . Our specification is this: A correct implementation of the upgraded contract C_2
 301 must satisfy `spec_trade_precise` and be accompanied by a contract morphism

302 $f : \text{ContractMorphism } C_2 \ C_1$

303 with the following five properties, stated formally in Listing 7:

- 304 1. `msg_morph f` rounds down the precision of messages to `trade` by a factor of 10
- 305 2. `msg_morph f` is the identity morphism on all messages aside from messages to `trade`
- 306 3. `state_morph f` rounds down on the balances kept in storage exposed by `get_bal`
- 307 4. `error_morph f` and `setup_morph f` are the respective identity functions

```

308 (* FORMAL SPECIFICATION:
309    An upgrade C2 must admit a morphism
310    f : ContractMorphism C2 C1
311    with the following properties: *)
312
313
314 (* 1. msg_morph f rounds trades down when it maps inputs of the receive function *)
315 Definition f_recv_input_rounds_down
316   (f : ContractMorphism C2 C1) : Prop :=
317   forall t', exists t,
318   (msg_morph C2 C1 f) (trade t') = trade t ^
319   trade_qty t = (trade_qty t') / 10.
320
321 (* 2. msg_morph f only affects the trade entrypoint *)
322 Definition f_recv_input_other_equal
323   (f : ContractMorphism C2 C1) : Prop :=
324   forall msg o,
325   (* for calls to all other entrypoints, *)
326   msg = other o →
327   (* f is the identity *)
328   option_map (msg_morph C2 C1 f) (other o) = other o.
329
330 (* 3. state_morph f rounds down on the storage *)
331 Definition f_state_morph (f : ContractMorphism C2 C1) : Prop :=
332   forall st, get_bal (state_morph C2 C1 f st) = (get_bal st) / 10.
333
334 (* 4. error_morph f and setup_morph f are the identity functions *)
335 Definition f_recv_output_err (f : ContractMorphism C2 C1) : Prop :=
336   (error_morph C2 C1 f) = id.
337
338 Definition f_init_id (f : ContractMorphism C2 C1) : Prop :=
339   (setup_morph C2 C1 f) = id.
340

```

■ **Listing 7** The formal specification of the upgrade from C_1 to C_2 .

341 The meaning of a morphism f satisfying the above conditions, as a specification, is in
 342 the *coherence conditions* of f . We know that every possible execution trace of C_2 has a
 343 corresponding execution trace in C_1 , and we know that the input messages are identical
 344 except that C_2 accepts trades at a higher level of precision. The coherence conditions also
 345 tell us that the state of C_2 is always related to the analogous state of C_1 , expressed in the
 346 function `state_morph`. With regards to the trading functionality of our new contract C_2 , we
 347 know that the balance kept in the storage of C_2 , which is affected by trades, will always be
 348 identical to the analogous balance of C_1 after rounding down, which we can formally prove.

1:10 Smart Contract Upgrades in Coq

```

349
350 Theorem rounding_down_invariant bstate caddr
351   (trace : ChainTrace empty_state bstate):
352   (* Forall reachable states with contract at caddr, *)
353   env_contracts bstate caddr = Some (C2 : WeakContract) →
354   (* cstate is the state of the contract AND *)
355   exists (cstate' cstate : storage),
356   contract_state bstate caddr = Some cstate' ∧
357   (* cstate is contract-reachable for C1 AND *)
358   cstate_reachable C1 cstate ∧
359   (* such that for cstate, the state of C1 in bstate,
360     the balance in cstate is rounded-down from the
361     balance of cstate' *)
362   get_bal cstate = (get_bal cstate') / 10.
363

```

■ **Listing 8** All reachable states of C2 round down to their corresponding states in C1.

364 Most importantly, f guarantees a relationship between the trading functionality of C2 and
 365 that of C1: C2 emulates the exact same trading behavior as C1 after rounding down one
 366 decimal place in precision. This means that C2 does not introduce any novel vulnerabilities
 367 relating to trades and balances not extant to C1. In particular, a proof of this fact would
 368 have prevented the attacks on Uranium Finance [8], NowSwap [4], and Nomad [7].

369 Moving on, note that f of Example 1 was directed from C2 to C1. The coherence conditions
 370 of f forced all execution traces of C2 to conform to a pattern set by C1, which is precisely
 371 what lets us make the claim that we haven't introduced any new behaviors regarding trading
 372 functionality to C2 aside from the increase in precision. Morphisms directed in the opposite
 373 direction can also be used in specification. Rather than classifying all possible execution
 374 traces of the upgrade, in this case a morphism proves that certain desired behavior exists
 375 within the contract. We illustrate with an example of specifying backwards compatibility.

376 ► **Example 2 (Backwards Compatibility)**. Consider contracts C1 and C2, where C2 is again an
 377 upgrade of C1, and suppose that we wish to show that C2 is backwards compatible with C1.
 378 The intent of this upgrade is that the full functionality of C1 be present within C2. We show
 379 this by embedding C1 into C2 via an injective contract morphism.

380 We illustrate with a simple example of a counter contract C1 which keeps some $n : \mathbb{N}$ in
 381 storage and has one entrypoint `incr` that increments the natural number in storage by 1. C1
 382 is upgraded to C2, which in addition to an entrypoint to increment the natural number in
 383 storage also includes a `decr` entrypoint to decrement the natural number in storage by 1.

```

384
385 Inductive entrypoint1 := | incr (u : unit).
386 Inductive entrypoint2 := | incr' (u : unit) | decr (u : unit).
387

```

■ **Listing 9** The entrypoint types of C1 and C2, respectively.

388 We prove that C2 is backwards compatible with C1 by defining a contract morphism

```

389
390           f : ContractMorphism C1 C2

```

390 with the following component functions.

```

391
392 Definition msg_morph (e : entrypoint1) : entrypoint2 :=
393   match e with | incr _ => incr' tt end.
394 Definition setup_morph : setup → setup := id.
395 Definition state_morph : storage → storage := id.
396 Definition error_morph : error → error := id.
397

```

398 These component functions do the obvious thing—send calls to the increment entrypoint of
 399 C1 to the increment entrypoint of C2 with the same payload, and do nothing otherwise. And
 400 f is an embedding since each of its component functions are manifestly injective, which we
 401 can formally prove.

```
402 Lemma f_is_embedding : is_inj_cm f.
```

405 Again, the meaning of f as a specification is in its coherence conditions. Any reachable
 406 state of C1 necessarily has an analogous reachable state of C2 which is fully structure preserving:
 407 if we were to only use the functionality of C2 which it inherits from C1, we would get identical
 408 contract behavior to C1. We have a formal proof of this result.

```
409 Theorem injection_invariant bstate caddr
410   (trace : ChainTrace empty_state bstate):
411   env_contracts bstate caddr = Some (C1 : WeakContract) →
412   (* Forall reachable states cstate of C1,
413      there's a corresponding reachable state
414      cstate' of C2, related by the injection *)
415   exists (cstate' cstate : storage),
416   contract_state bstate caddr = Some cstate ∧
417   (* cstate' is a contract-reachable state of C2 *)
418   cstate_reachable C2 cstate' ∧
419   (* .. equal to cstate *)
420   cstate' = cstate.
```

■ **Listing 10** C2 is backwards compatible with C1 via the embedding f.

423 This is a toy example, but in practice specifying a new contract which is backwards compatible
 424 to the old in this strong sense may not be straightforward. Via contract embeddings, contract
 425 morphisms give us a way of formally specifying and verifying backwards compatibility.

5 Further Applications of Morphisms in Formal Verification

427 Contract morphisms establish a relationship between contracts which makes them suitable
 428 for specifying and verifying upgrades. For that same reason, contract morphisms may also
 429 have applications in proof reuse, or proof *transport*, more generally. The special case of
 430 contract *isomorphism* may also provide a stronger relationship between formal specification
 431 and proof on the associated contracts.

5.1 Hoare Properties and Contract Morphisms

433 First we consider properties that *transport* over a morphism, in particular those that we
 434 can pull back over a morphism. Hoare properties are a particularly strong example: they
 435 relate pre-conditions to post-conditions, which is relevant to morphisms because morphisms
 436 relate inputs and outputs of contract executions. As contracts are formalized in ConCert,
 437 constraints on on inputs amount to pre-conditions, and constraints on outputs amount to
 438 post-conditions. Thus for contracts C1 and C2 and a morphism $f : \text{ContractMorphism } C1 \ C2$,
 439 we might expect to be able to transport Hoare properties of one contract over f to the other.

440 Indeed, any Hoare property proved for C2 will always have an analogous result on C1,
 441 mediated by f. We proved this in two results which relate all reachable states of C1 to those
 442 of C2, and those of C2 to those of C1, via the `state_morph` component of f. These results,
 443 `left_cm_induction` and `right_cm_induction`, are collectively called morphism induction, as
 444 they allow us to induct along the execution trace of one contract in relation to that of another.

445 In particular, morphism induction says that properties of the state of $C2$ which are invariant
 446 over `state_morph` must hold for all states of $C1$.

447 As a toy example of this relationship, suppose that we can prove that if a certain boolean
 448 in the storage of $C2$ is set at `true`, a given entrypoint $e2$ of $C2$ can be successfully called, and
 449 that it fails otherwise. Suppose further that the `msg_morph` component of f sends all calls
 450 to an entrypoint $e1$ of $C1$ to calls to $e2$, and that the `state_morph` component of f sends a
 451 state of $C1$ with an analogous boolean set at `true` to one of $C2$ with the boolean set at `false`,
 452 and visa versa. Then by morphism induction on the trace of $C1$, we get for free that calls to
 453 $e1$ succeed only when the analogous boolean in the state of $C1$ is set at `false`, rather than
 454 `true`. The relationship encoded by f between contracts $C1$ and $C2$ shows that $C1$ and $C2$ use
 455 opposing, but predictably related, logic for execution, which allows us to reuse proofs on $C2$
 456 to prove analogous results on $C1$.

457 5.2 Isomorphisms and Propositional Indistinguishability

458 This relationship between contracts strengthens when we have a pair of morphisms

459 $f : \text{ContractMorphism } C1 \ C2$ and $g : \text{ContractMorphism } C2 \ C1$

460 such that `compose_cm g f = id_cm C1` and `compose_cm f g = id_cm C2`. This is an *isomorph-*
 461 *ism* of contracts. Isomorphisms of contracts are particularly strong; the component functions
 462 are equivalences of types and they induce a bisimulation of contracts in ConCert.

463 Since bisimulation is a strong and mathematically stable notion of equivalence [17], future
 464 work could investigate proof transport over contract isomorphisms, building on recent work
 465 in Coq-based formal methods. For example, we may wish to prove results on a contract
 466 optimized for formal reasoning, and transport those onto a bisimilar, performant contract,
 467 similar to the work of Cohen *et al.* [5]. This might include altering certain data types while
 468 maintaining an equivalence; chosen data types have a strong influence on the structure of
 469 proofs and can be nontrivial to transport [11, 15, 20].

470 6 Conclusion

471 Our goal in this paper was to provide a formal framework for formally specifying and verifying
 472 smart contract upgrades in Coq. To do so we introduced the notion of a contract morphism,
 473 which encodes a formal relationship between execution traces of two contracts. We argued
 474 that this was a suitable, formal notion with which to reason about contract upgrades and
 475 provided examples of contract upgrades which can be specified and verified with contract
 476 morphisms. To our knowledge, this is the first time that the intent of an upgrade has been
 477 articulated explicitly in formal specification, and is the first formal attempt at reasoning
 478 explicitly about contract upgrades in a formal setting.

479 This work is intended to be a preliminary framework for reasoning about contract upgrades
 480 in Coq. As such, there are practical questions to be asked, such as whether these tools
 481 are even feasible on gas-optimized code, which can be difficult to formally reason about.
 482 Even so we are optimistic, as the previously-mentioned work by Ringer *et al.* in proof
 483 repair is practically useful and resembles our framework from a theoretical standpoint. Since
 484 the status quo is to simply update the formal specification of a previous version into the
 485 specification of the new, we hope that contract morphisms will be a strong start to efficient
 486 and rigorous verification of contract upgrades.

487 — References —

- 488 1 Danil Annenkov, Mikkel Milo, Jakob Botsch Nielsen, and Bas Spitters. Extracting smart
 489 contracts tested and verified in Coq. In *Proceedings of the 10th ACM SIGPLAN International*
 490 *Conference on Certified Programs and Proofs*, CPP 2021, pages 105–121, New York, NY, USA,
 491 January 2021. Association for Computing Machinery. doi:10.1145/3437992.3439934.
- 492 2 Danil Annenkov, Jakob Botsch Nielsen, and Bas Spitters. ConCert: A smart contract
 493 certification framework in Coq. In *Proceedings of the 9th ACM SIGPLAN International*
 494 *Conference on Certified Programs and Proofs*, CPP 2020, pages 215–228, New York, NY, USA,
 495 January 2020. Association for Computing Machinery. doi:10.1145/3372885.3373829.
- 496 3 Pedro Antonino, Juliandson Ferreira, Augusto Sampaio, and A. W. Roscoe. Specification is
 497 Law: Safe Creation and Upgrade of Ethereum Smart Contracts. In *Software Engineering and*
 498 *Formal Methods: 20th International Conference, SEFM 2022, Berlin, Germany, September*
 499 *26–30, 2022, Proceedings*, pages 227–243. Springer, 2022.
- 500 4 Rob Behnke. Explained: The NowSwap Protocol Hack. [https://halborn.com/explained-the-](https://halborn.com/explained-the-nowswap-protocol-hack-september-2021/)
 501 [nowswap-protocol-hack-september-2021/](https://halborn.com/explained-the-nowswap-protocol-hack-september-2021/), September 2021. Accessed January 2024.
- 502 5 Cyril Cohen, Maxime Dénès, and Anders Mörtberg. Refinements for free! In *International*
 503 *Conference on Certified Programs and Proofs*, pages 147–162. Springer, 2013.
- 504 6 Thomas Dickerson, Paul Gazzillo, Maurice Herlihy, Vikram Saraph, and Eric Koskinen. Proof-
 505 Carrying Smart Contracts. In Aviv Zohar, Ittay Eyal, Vanessa Teague, Jeremy Clark, Andrea
 506 Bracciali, Federico Pintore, and Massimiliano Sala, editors, *Financial Cryptography and Data*
 507 *Security*, Lecture Notes in Computer Science, pages 325–338, Berlin, Heidelberg, 2019. Springer.
 508 doi:10.1007/978-3-662-58820-8_22.
- 509 7 etherscan.io. Nomad Bridge Exploit.
 510 Transaction 0xa5fe9d044e4f3e5aa5bc4c0709333cd2190cba0f4e7f16bcf73f49f83e4a5460, 2022.
- 511 8 Uranium Finance. Uranium Finance Exploit. [https://uraniumfinance.medium.com/exploit-](https://uraniumfinance.medium.com/exploit-d3a88921531c)
 512 [d3a88921531c](https://uraniumfinance.medium.com/exploit-d3a88921531c), April 2021. Accessed January 2024.
- 513 9 Immunefi. Hack Analysis: Nomad Bridge, August 2022. [https://medium.com/immunefi/hack-](https://medium.com/immunefi/hack-analysis-nomad-bridge-august-2022-5aa63d53814a)
 514 [analysis-nomad-bridge-august-2022-5aa63d53814a](https://medium.com/immunefi/hack-analysis-nomad-bridge-august-2022-5aa63d53814a), January 2023.
- 515 10 Oussama Jebbar, Ferhat Khendek, and Maria Toeroe. Upgrade of highly available systems:
 516 Formal methods at the rescue. In *2017 IEEE International Conference on Information Reuse*
 517 *and Integration (IRI)*, pages 270–274. IEEE, 2017.
- 518 11 Nicolas Magaud and Yves Bertot. Changing data structures in type theory: A study of
 519 natural numbers. In *International Workshop on Types for Proofs and Programs*, pages 181–196.
 520 Springer, 2000.
- 521 12 Stephen McCamant and Michael D Ernst. Predicting problems caused by component upgrades.
 522 In *Proceedings of the 9th European software engineering conference held jointly with 11th ACM*
 523 *SIGSOFT international symposium on Foundations of software engineering*, pages 287–296,
 524 2003.
- 525 13 Nick Mudge. EIP-2535: Diamonds, Multi-Facet Proxy. [https://eips.ethereum.org/EIPS/eip-](https://eips.ethereum.org/EIPS/eip-2535)
 526 [2535](https://eips.ethereum.org/EIPS/eip-2535). Accessed January 2024.
- 527 14 Talia Ringer. *Proof Repair*. University of Washington, 2021.
- 528 15 Talia Ringer, RanDair Porter, Nathaniel Yazdani, John Leo, and Dan Grossman. Proof repair
 529 across type equivalences. In *Proceedings of the 42nd ACM SIGPLAN International Conference*
 530 *on Programming Language Design and Implementation*, pages 112–127, 2021.
- 531 16 Talia Ringer, Nathaniel Yazdani, John Leo, and Dan Grossman. Adapting proof automation to
 532 adapt proofs. In *Proceedings of the 7th ACM SIGPLAN International Conference on Certified*
 533 *Programs and Proofs*, pages 115–129, 2018.
- 534 17 Davide Sangiorgi. On the bisimulation proof method. *Mathematical Structures in Computer*
 535 *Science*, 8(5):447–479, October 1998. doi:10.1017/S0960129598002527.
- 536 18 Amritraj Singh, Reza M Parizi, Qi Zhang, Kim-Kwang Raymond Choo, and Ali Dehghantanha.
 537 Blockchain smart contracts formalization: Approaches and challenges to address vulnerabilities.
 538 *Computers & Security*, 88:101654, 2020.

- 539 19 Derek Sorensen. (In)Correct Smart Contract Specifications. *IEEE International Conference*
540 *on Blockchain and Cryptocurrency (ICBC)*, 2024.
- 541 20 Nicolas Tabareau, Éric Tanter, and Matthieu Sozeau. Equivalences for free: univalent
542 parametricity for effective transport. *Proceedings of the ACM on Programming Languages*,
543 2(ICFP):1–29, 2018.
- 544 21 Doug Woos, James R. Wilcox, Steve Anton, Zachary Tatlock, Michael D. Ernst, and Thomas
545 Anderson. Planning for change in a formal verification of the raft consensus protocol. In
546 *Proceedings of the 5th ACM SIGPLAN Conference on Certified Programs and Proofs, CPP*
547 2016, pages 154–165, New York, NY, USA, 2016. Association for Computing Machinery.
548 doi:10.1145/2854065.2854081.
- 549 22 Jiahua Xu, Krzysztof Paruch, Simon Cousaert, and Yebo Feng. SoK: Decentralized Exchanges
550 (DEX) with Automated Market Maker (AMM) Protocols. *ACM Computing Surveys*, 55(11):1–
551 50, 2023.