



Meta Properties of Financial Smart Contracts

Derek Sorensen



Clare College

September 2023

This dissertation is submitted for the degree of Doctor of Philosophy

Declaration

This dissertation is the result of my own work and includes nothing which is the outcome of work done in collaboration except as declared in the Preface and specified in the text. It is not substantially the same as any that I have submitted, or am concurrently submitting, for a degree or diploma or other qualification at the University of Cambridge or any other University or similar institution except as declared in the Preface and specified in the text. I further state that no substantial part of my dissertation has already been submitted, or is being concurrently submitted, for any such degree, diploma or other qualification at the University of Cambridge or any other University or similar institution except as declared in the Preface and specified in the text. This dissertation does not exceed the prescribed limit of 60 000 words.

Derek Sorensen
September, 2023

Abstract

Meta Properties of Financial Smart Contracts

Derek Sorensen

Financial smart contracts routinely manage billions of US dollars worth of digital assets, making bugs in smart contracts extremely costly. Because of this, much work has been done in formal verification of smart contracts to prove a contract correct with regards to its specification. However, financial smart contracts have complicated specifications, and it is not all straightforward to write one which correctly captures all of its intended high-level behaviors. To mitigate this challenge, we develop formal tools to target *meta properties* of smart contracts, which are properties of a contract that are intended by, but out of scope of, its specification. The targeted properties include the economic behaviors of the contract, properties relating to its upgradeability features, and the intended behaviors of systems of contracts. The formal tools presented are written in Coq.

Acknowledgements

First to my primary supervisor Anil, for giving me a second chance at a PhD when my first topic did not work out as I'd hoped; for his unbounded enthusiasm in pursuing meaningful and impactful things; for his exceptional leadership; for his wise advice in business and research; for pressing me to find one single, coherent thread to my thinking that transformed this thesis it into what it is now.

To my secondary supervisor Keshav, for his extraordinary and unrelenting ability to find weaknesses in my thinking; for his instructive stories; for his groundedness that forces me to confront the world as it is and do research that is relevant to the world; for forcing me to justify all of my ideas with the utmost of rigor; and for taking the time to meticulously read my work.

To the Choir of Clare College, which gave me new life and love for the city of Cambridge and for my college after the profound isolation and suffering of successive lockdowns, having spent more than four months without getting within two meters of another human soul; for the music that inspires me to commune with the Divine and which has taught me to see research and mathematics in the transcendence that they embody; for the music, the parties, the late nights, the trips, and the unforgettable memories.

To the Rev'd Dr Mark Smith, for our many chats on life, faith, and theology; for our trips to the pub in between lockdowns; for our many Bible study sessions; for making the chapel a sacred space of worship—all of which ultimately culminated in my baptism. To the Chapel of Clare College, for our hilarious lunches, formals, and trivia nights; for the moral support to continue on with a PhD through the depths of lockdown; for making a community of faith in chapel. I could not have asked for a better College Dean or better chapel camaraderie.

To my good friend Anthony Davies, who kept me sane through lockdown by giving me an extensive and strict dietary and exercise routine, and who supported me to follow it diligently despite his own suffering through lockdown. You were the pin that held me together during those dark times. Thanks to you I am off anti-depressants. Thank you for giving me my life back.

And finally, to my mother for her quiet but vigorous encouragement to avoid abandoning my goals, and to my father whose unbridled confidence in my abilities has given me new (and perhaps unsubstantiated) confidence when I have found myself discouraged.

Thank you all.

Contents

1	Introduction	15
1.1	Formal Verification of Smart Contracts	15
1.1.1	Dexter2 Verification	16
1.2	Contract Vulnerabilities and (In)Correct Specifications	17
1.2.1	Economic Attacks	17
1.2.2	Unsafe Contract Upgrades	19
1.2.3	Complex System Behavior	19
1.3	Meta Properties of Financial Smart Contracts	20
2	Related Work	21
2.1	Smart Contract Verification	21
2.1.1	Smart Contract Language Embeddings	22
2.1.2	ConCert	23
2.2	Verifying Meta Properties	24
3	Background	25
3.1	Smart Contracts in ConCert	25
3.2	The Blockchain in ConCert	27
3.2.1	Blockchain Semantics in ConCert	29
3.3	Specification and Proof in ConCert	31

3.4	Contract Induction	34
4	A Contract's Economic Properties	35
4.1	Contract Specification With Economic Intent	35
4.2	Case Study: Structured Pools	37
4.2.1	The Issue of Fungibility in Tokenized Carbon Credits	37
4.2.2	Structured Pools	38
4.2.3	Properties of Structured Pools	43
4.3	Specification and Metaspecification	48
4.4	Formal Specification: A Contract Axiomatization	48
4.4.1	The Structured Pool Formal Specification	49
4.4.2	Typeclasses to Characterize Contract Types	49
4.4.3	Specifying Contract Initialization	52
4.4.4	Specifying Each Contract Entrypoint	52
4.4.5	The Formal Specification as a Predicate on Contracts	57
4.5	Formal Metaspecification	57
4.5.1	Formalizing the Metaspecification	58
4.5.2	Discussion	62
4.6	Conclusion	63
5	Contract Upgradeability	65
5.1	Contract Upgrades	65
5.1.1	Specifying Upgradeability: The Diamond Framework	66
5.1.2	Evolving Specifications Through Upgrades	67
5.1.3	Related Work	68
5.2	Contract Morphisms	68

5.2.1	Composition of Morphisms	74
5.3	Morphism Induction	76
5.3.1	Contract Trace and Reachability	76
5.3.2	Left Morphism Induction	77
5.3.3	Right Morphism Induction	78
5.4	Reasoning with Morphisms: Specification and Proof	79
5.4.1	Specifying a Contract Upgrade With Morphisms	79
5.4.2	Adding Features and Backwards Compatibility	83
5.4.3	Transporting Hoare-Like Properties Over a Morphism	84
5.4.4	Summary	88
5.5	A Mathematical Characterization of Contract Upgrades	89
5.5.1	The Varieties of Upgradeable Contracts	89
5.5.2	Isolating Mutable and Immutable Parts	90
5.5.3	Decomposing Upgradeability	93
5.5.4	Upgradeable Contracts are Fiber Bundles: A Digression	99
5.5.5	Summary	103
5.6	Conclusion	103
6	Contract Systems	105
6.1	Contracts Systems	105
6.1.1	Specifying Contract Systems	106
6.1.2	Related Work	107
6.2	Bisimulations of Contracts	108
6.2.1	Contract Trace Morphism	108
6.2.2	The Lifting Theorem	112
6.2.3	Contract Bisimulations	114

6.2.4	Discussion: Propositional Indistinguishability	115
6.3	Contract Systems as Bigraphs	117
6.3.1	Bigraphs	117
6.3.2	The Place Graph	117
6.3.2.1	Iteratively Building a Contract System	120
6.3.2.2	System Contracts, Morphisms, and Isomorphisms	123
6.3.3	The Link Graph	125
6.3.3.1	System Steps and System Trace	126
6.4	Bisimulations of Contract Systems	128
6.4.1	System Trace Morphisms and System Bisimulations	128
6.4.2	Lifting Theorems for Contract Systems	131
6.5	Conclusion	133
7	Conclusion	135
7.1	Limitations and Future Work	135
	References	139
A	Proofs and Definitions of Chapter 4	153
A.1	Formal Specification	153
A.1.1	Typeclass Specifications	153
A.1.2	The Formal Specification of a Structured Pool Contract	155
A.1.3	The Formal Specification Predicate	168
A.2	Formal Metaspecification	170
A.2.1	Demand Sensitivity	170
A.2.2	Nonpathological Prices	171
A.2.3	Swap Rate Consistency	171

A.2.4	Zero-Impact Liquidity Change	174
A.2.5	Arbitrage Sensitivity	175
A.2.6	Pooled Consistency	176
B	Proofs and Definitions of Chapter 5	177
B.1	Contract Morphisms	177
B.1.1	Composition of Morphisms	182
B.2	Morphism Induction	186
B.2.1	Contract Trace and Reachability	186
B.2.2	Right Morphism Induction	187
B.2.3	Right Morphism Induction	188
B.3	Reasoning with Morphisms: Specification and Proof	188
B.3.1	Specifying a Contract Upgrade with Morphisms	188
B.3.2	Adding Features and Backwards Compatibility	192
B.3.3	Transporting Hoare-Like Properties Over a Morphism	195
B.4	A Mathematical Characterization of Contract Upgrades	199
B.4.1	Isolating Mutable and Immutable Parts	202
B.4.2	Decomposing Upgradeability	204
B.4.3	Upgradeable Contracts are Fiber Bundles: A Digression	206
C	Proofs and Definitions of Chapter 6	209
C.1	Bisimulations of Contracts	209
C.1.1	Contract Trace Morphisms	209
C.1.2	The Lifting Theorem	213
C.1.3	Contract Bisimulations	215
C.1.4	Discussion: Propositional Indistinguishability	218

C.2	Contract Systems as Bigraphs	221
C.2.1	Bigraphs	221
C.2.2	The Place Graph	221
C.2.2.1	Iteratively Building a Contract System	223
C.2.2.2	System Contracts, Morphisms, and Isomorphisms	226
C.3	The Link Graph	227
C.3.1	System Steps and System Trace	227
C.4	Bisimulations of Contract Systems	229
C.4.1	System Trace Morphisms and System Bisimulations	229
C.4.2	Lifting Theorems for Contract Systems	233
	Glossary	239
	Acronyms	247

Chapter 1

Introduction

Smart contracts¹ are programs stored on a blockchain that automatically execute when certain predefined conditions are met. *Financial smart contracts* are broadly defined as smart contracts that serve as a digital intermediary between financial parties. These include contracts collectively referred to as decentralized finance (DeFi), and come in many forms, including decentralized exchanges (DEXs), automated market makers (AMMs), crypto lending, synthetics (including stablecoins), yield farming, crypto insurance protocols, and cross-chain bridges [135]. Financial smart contracts frequently manage huge quantities of money, making it essential for the underlying code to be rigorously tested and verified to ensure its correctness and security [144].

A defining characteristic of smart contracts is that once deployed, they are immutable. Thus if a contract has vulnerabilities, the victims of an attack are helpless to stop the attacker if the contract wasn't designed with the foresight sufficient to respond. Due to the high financial cost of exploits, it can be worth the large overhead cost to formally verify a smart contract before deployment.

1.1 Formal Verification of Smart Contracts

Much work has been done in formal verification of smart contracts [17, 38, 45, 63, 97, 121, 131]. Broadly speaking, the goal of formal verification is to formally prove that a contract is correct with regards to a specification. However, financial smart contracts have complicated specifications, and it is not at all straightforward to write one which has all of its intended behaviors. One reason for this is that specifications of financial contracts inevitably include behaviors which are expressed at a level of abstraction higher than a contract specification.

Let us illustrate with the formal verification work on Dexter2, an AMM on the Tezos blockchain.

¹This document contains a lot of blockchain jargon. These words are defined in the glossary, and each entry can be accessed from the body of the document by clicking on an instance of the corresponding word, *e.g.* on *smart contract*.




Figure 1.1: A trade of Δ_x for Δ_y along the indifference curve $xy = k$.

1.1.1 Dexter2 Verification

Dexter2 has been formally verified by three different groups using three different formal verification tools [34, 80, 102]. Each was based on the same informal specification [16].

The informal specification describes the contract interface, including its entrypoint functions, error messages, outgoing transactions, the contents of its storage, some invariants of the storage (including that its store of tokens never fully depletes), fees, and the logic of each of the entrypoint functions. This is a standard and detailed contract specification. Note, however, that while the specification is detailed on the contract design and interface, it doesn't include anything about expected or desired economic behavior.

This is not because the expected or desired economic behavior is unknown or uninteresting. It was clearly articulated by Vitalik Buterin, co-founder of Ethereum, in what he wrote in March 2018 about AMMs on the online forum Ethereum Research [33]. Buterin proposed that AMMs trade between a pair of tokens along an *indifference curve*

$$xy = k, \tag{1.1}$$

where x represents the quantity held by the contract of the token being traded in, y represents the quantity held by the contract of the token being traded out, and k is a constant. The tokens held by the contract come from liquidity providers, which are investors who deposit tokens into the AMM in exchange for a reward, most often a share of transaction fees. That k is constant means that a trade of Δ_x of one token yields Δ_y of another such that the product from (1.1) stays constant at k :

$$(x + \Delta_x)(y - \Delta_y) = k.$$

Buterin argued that an AMM that trades along (1.1) features efficient price discovery. He also argued that it can properly incentivize liquidity providers by charging a 0.3% fee on each trade to give to them.

We can probably convince ourselves that the informal specification of Dexter2 [16] features these economic qualities described by Buterin, including efficient price discovery and some suitable incentive mechanism so that investors deposit tokens into the AMM contract and provide liquidity to the market. However, concluding that the informal specification [16] or its formal counterparts [34, 80, 102] actually imply any of these economic behaviors is not a given fact. AMM fees and liquidity provision alone are highly complex topics, and are the subject of several economic studies, including: choosing optimal transaction fees [56, 62, 61], how liquidity providers react to market changes [69], and how all of that relates to the *curvature* of $xy = k$ [8]. It was also shown that front-running attack attacks can warp the incentive scheme of the blockchain itself in such a way that could compromise its underlying security [43].

To complicate matters, a brief study of the three formalizations [34, 80, 102] of the Dexter2 specification [16] reveals that each differs substantively both in how the properties of the informal specification are formalized, and in what assumptions are made in the process of verification.

We can see a general problem in specifying financial smart contracts. Financial smart contract specifications are designed to articulate properties, in this example of an economic nature, which are essential to the contract’s correct functionality but out of scope of said specification. Unfortunately, failing to correctly capture these intended properties in the specification routinely leads to extraordinarily expensive attacks. Furthermore, because these are vulnerabilities of the *specification*, as it stands formal verification is useless to mitigate them.

1.2 Contract Vulnerabilities and (In)Correct Specifications

Vulnerabilities of financial smart contracts are a serious cause of substantial financial loss. Blockchain-based applications lost 2.44 billion USD in 2021 and 3.6 billion USD in 2022 due to attacks [23, 146]. In 2022, financial smart contracts were the most attacked type of blockchain-based application, making up about two-thirds of all attacks [60, 23]. Importantly, attacks which exploit improper business logic or function design—those which concern us here—are in the top three causes of loss [23, p.10].

We can isolate three classes of vulnerabilities due to poor contract design and specification. These are: economic vulnerabilities, vulnerabilities introduced through poorly-specified upgrades or contract upgradeability, and vulnerabilities due to difficult-to-specify behaviors of a systems of interacting contracts. For each of these classes of vulnerabilities, we give examples of recent, successful attacks and discuss what is needed to target such vulnerabilities with formal verification.

1.2.1 Economic Attacks

Poorly specified financial smart contracts can be vulnerable to costly *economic attacks*, or attacks on the contract’s economic design.

Take for example Beanstalk, an Ethereum-based stablecoin protocol which uses a decentralized governance protocol. The governance protocol features an emergency commit function, which gives a supermajority of governance votes power to quickly respond to an emergency by approving and executing a proposal in one single vote. On April 17, 2022, an attacker used a *flash loan* to temporarily buy a supermajority of governance tokens and execute a proposal, draining the contract of approximately 77 million USD in lost contract assets [54].

The tool used in the attack, flash loans, are loans mediated by a smart contract, issued for the duration of a single transaction. Due to their atomicity, flash loans remove the creditor’s risk of debt default, and thus enable enormous, uncollateralized loans. For example, the Aave flash loan pool has had in excess of 1 billion USD which can be loaned out [112]. Flash loans can introduce unintuitive contract behavior, deviating from that of traditional markets, and have been extensively studied [65, 66, 112, 134]. Importantly, the Beanstalk attack leveraged the unexpected behavior due to the availability of flash loans, exploiting the faulty design of the governance mechanism rather than incorrect code [57].

More examples of successful flash loan attacks include attacks on the Spartan Protocol and Pancake Bunny, two DeFi contracts on the Binance Smart Chain (BSC). Attackers used flash loans to make huge trades and temporarily manipulate market prices of certain assets. Both of these contracts used these market prices in the contract logic, and in both cases this led to pathological—though, again, correct according to the specification—contract behavior. In May 2021, an attacker drained the Spartan Protocol contract for a profit of roughly the equivalent of 30 million USD [31, 79]. Another attacker drained Pancake Bunny of 114k WBNB and 697k BUNNY tokens, amounting to about 45 million USD at the time in lost funds [30, 67, 76, 108].

Finally, Mango Markets, a Solana-based DEX, was attacked in October 2022 for approximately 116 million USD in contract assets [123]. The attack consisted of a complicated and subtle trading strategy which only a sophisticated trader would be able to see and exploit [23]. CoinDesk reported that the attacker did everything within the parameters of the platform’s design [87]. Avraham Eisenberg, the attacker, wrote:

I believe all of our actions were legal open market actions, using the protocol as designed, even if the development team did not fully anticipate all the consequences of setting parameters the way they are. [15]

We would recognize each of these exploits as attacks, despite the fact that the contracts were functioning as specified. This tells us that the specifications did not accurately capture the intended economic behaviors, and were thus incorrect. To mitigate such attacks by ensuring a specification *does* accurately capture the intended economic behaviors, what is needed is a rigorous notion of a specification’s correctness as it relates to its economic properties, as well as formal tools to reason about a said correctness.

1.2.2 Unsafe Contract Upgrades

Poorly specified contract upgrades can introduce costly vulnerabilities.

Consider first Nomad, a cross-chain bridge protocol. In August 2022, more than 500 hacker addresses exploited a bug introduced by a faulty upgrade to one of the Nomad smart contracts [116]. The upgrade incorrectly added the null address ($0 \times 000 \dots 000$) as a trusted root, which turned off a key safety check, allowing anyone to withdraw arbitrary amounts of funds from the Nomad contract to their wallet by calling the contract with a particular payload. The attack resulted in 190 million USD in lost funds [55, 78].

Similarly, Uranium Finance, an AMM, suffered a costly exploit after a faulty contract upgrade. The original contract contained a constant, κ , equal to 1,000 in three different places, which was used to price trades. The update changed this value to 10,000 in two places but not the third, presumably to calculate trades with higher precision. The result of this was that the attacker could swap virtually nothing in for 98% of the total balance of any output token, which resulted in a loss of 50 million USD of contract funds [59]. NowSwap, a nearly identical application, upgraded with the same error and incurred a loss of 1 million USD [22].

It is clear that none of these contract upgrades captured the actual intent of the upgrade. Each introduced vulnerabilities, as small technical changes of the contract, which compromised the contract's fidelity to its intended design. Indeed, each upgrade was meant to preserve properties of the previous contract version, regarding pricing or permissions, while changing others. As each failed to do so, they were incorrect.

In order to safely specify contract upgrades and upgradeable contracts, what is needed is a rigorous notion of a correct specification of individual contract upgrades as well as contract upgradeability, as well as formal tools to reason about said correctness.

1.2.3 Complex System Behavior

Finally, poorly specified systems of contracts can be vulnerable to extremely costly attacks.

Consider Harvest Finance, a yield aggregator on Ethereum. On October 26, 2020, an attacker used a flash loan to trade about 17.2 million USDT for USDC on Curve, which temporarily increased the price of USDC in the Curve Y pool. Due to the fact that Harvest Finance uses the Curve Y pool as a price oracle in real-time to calculate the vault shares for a deposit, the attacker got into a Harvest vault at an advantageous rate. In the same transaction, the attacker reversed the trade on Curve, after which the price of USDC returned to normal in the Curve Y pool, which increased the value of the attacker's shares in terms of the now less expensive USDC. The hacker then exited the vault at this new exchange rate for a profit of 33 million USD in lost user funds [107].

Smart Contract	Vulnerability	Loss (USD)	dApp Type	Exploit	Year
Mango Markets	Contract Design [87]	115M	DEX	[123]	2022
Beanstalk	Contract Design [57]	77M	Stablecoin	[54]	2022
Pancake Bunny	Contract Design [67]	45M	Yield aggregator	[30]	2021
Spartan Protocol	Contract Design [79]	30M	AMM	[31]	2021
Nomad	Unsafe Upgrade [78]	190M	Cross-chain bridge	[55]	2022
Uranium	Unsafe Upgrade [59]	50M	AMM	[32]	2021
Cream Finance	Complex System [77]	130M	DeFi protocol	[53]	2021
Harvest Finance	Complex System [58]	34M	Yield aggregator	[52]	2020

Figure 1.2: A small sample of recent attacks, totalling to about 776M USD in lost funds.

Similarly, consider CREAM finance (short for *Crypto Rules Everything Around Me*), a multi-purpose DeFi protocol that brands itself as a one-stop shop for decentralized finance. It offers crypto lending, borrowing, yield farming, and trading services, and has several connected implementations across multiple chains. An October 2021 attack drained the pool of roughly 260 million USD in assets [53]. The attack was extremely complex, involving 68 assets and over 9 ETH in gas, roughly 36k USD at the time [53]. Immunefi, a bug bounty platform for smart contracts, diagnoses that the exploit was due to an easily manipulable price oracle, and uncapped supply of the token yUSD [77]. Even so, the attack is complex enough that only experts such as Immunefi can give a comprehensive diagnosis.

Aside from complex economic properties, the difficulty in specifying the above examples comes in the utter complexity of interacting systems of contracts. The intended behavior of a system of contracts, typically expressed and written as if it were a monolithic entity, can be difficult to preserve when modularizing the contract design into component pieces without introducing vulnerabilities. In order to safely specify such systems, what is needed is a rigorous notion of a specification’s correctness as it relates to how the system of contracts behaves when taken as a whole, and formal tools to reason about said correctness.

1.3 Meta Properties of Financial Smart Contracts

Our thesis is that we can mitigate costly vulnerabilities due to incorrect specifications by formally specifying and verifying a contract’s *meta properties*—or properties which are intended by, but out of scope of, a contract specification. To that end, we present formal tools targetting a contract’s economic properties (Chapter 4), upgradeability (Chapter 5), and compositionality (Chapter 6).

Chapter 2

Related Work

Here we survey previous work on formal verification of smart contracts with the goal of specifying and verifying meta properties. We find that, generally speaking, those verification tools which in principle are able to specify and verify meta properties are not able to do so on code which can be compiled and deployed, and those verification tools which reason about code that can be compiled and deployed are too low-level in scope to specify and verify the meta properties of interest here.

We are able to address this issue because we do our work in ConCert [10], a Coq-based verification tool with verified extraction [11, 13]. Because ConCert has a full embedding of the execution semantics of a blockchain in Coq, an interactive theorem prover [27], we are able to specify and verify arbitrary properties of smart contracts. Because it features verified extraction, we are able to do so on code which can be compiled and deployed.

2.1 Smart Contract Verification

We consider the landscape of formal verification with the goal of expressing and reasoning about meta properties of smart contracts. We will draw on several surveys of formal verification of smart contracts, including [17, 38, 45, 63, 97, 121, 131].

From among the plethora of tools available to formally verify smart contracts, we will mostly focus on those which use interactive theorem provers, or proof assistants. Verification tools based on proof assistants tend to use the proof assistant itself as a specification language, and so smart contracts are treated as mathematical objects, about which any mathematical statement can be made. In principle, this means that proof assistants can be used to verify any correct design [115]. As we will see, we use diverse mathematical tools to formalize meta properties, so the flexibility and robustness of proof assistants come particularly in handy. We also focus on proof assistants because they are recognized in the literature to be of the highest quality for formal verification [97].

2.1.1 Smart Contract Language Embeddings

Most verification tools which use proof assistants come in the form of an embedding of a smart contract language into the proof assistant. Smart contracts in that embedded language are then reasoned about through the embedding. We will survey these briefly, starting with embeddings of low-level languages and moving to those of intermediate- and high-level languages.

There are several examples of low-level language embeddings into proof assistants. For EVM bytecode alone we have embeddings in Isabelle [5, 74]; in Coq [139]; in the K Framework [73, 109]; in Why3 [143, 100]; and in Z3 [85]. There are similar embeddings for other smart contract languages, including for the Tezos smart contract language Michelson, which is a low-level, stack-based language. These include Mi-Cho-Coq, an embedding of Michelson into Coq [25]; K Michelson [81], an embedding of Michelson into the K Framework; and WhyISon, which transpiles Michelson into WhyML, the programming and specification language of the Why3 framework [42]. There are examples on other chains as well, including an embedding of the low-level, Bitcoin-based contract language Simplicity in Coq [105].

Due to their proximity in language and semantics of the lowest-level executing environment to each of these blockchains, low-level language embeddings are more likely to be faithful to the actual execution environment [84]. Equally, this low-level proximity can make it difficult to correctly specify high-level behavioral or economic properties of smart contracts without some form of abstraction.

On the other hand, higher-level languages tend to be more human-readable, and can be more straightforward to reason about [84], especially if they are statically typed functional programming languages [10]. Naively, this higher level of abstraction seems like it would make it more straightforward to reason about the correctness of a specification. However, the abstraction can come at the cost of rigor. As they sit at a higher level of abstraction, they require a rigorous language embedding as well as a correct compiler down to the low-level, executable code which preserves the proven properties [84].

We have many examples of intermediate- and high-level languages which have been embedded into proof assistants. For Ethereum smart contracts, we have Lolisa [138] and FEther [137], embeddings of Solidity into Coq; as well as embeddings of Solidity into Event-B [147]; into Isabelle [84]; into the K Framework [82]; and into F^* [28]. We also have custom implementations of finite state machines used by FSolidM [90, 91] and VeriSolid [92] that verify Solidity code. For Tezos, we have an embedding of Albert [26], an intermediate-level language which compiles down to Michelson, and which targets Mi-Cho-Coq; we also have ConCert [10], which has a certified extraction mechanism from Coq code into multiple smart contract languages [11], including into CameLIGO, an intermediate-level language which compiles down to Michelson [9]. On other chains we have Plutus in Agda [37], and verification for BNB in Coq [127].

At higher levels of abstraction, we also have some DSLs written in proof assistants which target various smart contract languages. Scilla is intermediate-level, and can be used to reason about temporal properties of smart contracts, which targets Solidity [118]. Archetype is a Tezos-based DSL which targets business

logic and uses Why3 [24]. Multi is a framework, written in Coq, which targets reasoning about smart contract interactions [36]. At an even higher level of abstraction, we have TLA+, a tool for reasoning about concurrent and distributed systems, which was used to verify the a cross-chain swap protocol [101].

For our purposes the key advantage to verification with intermediate- and high-level languages is that the abstraction makes it easier to write a correct specification, and reasoning about code also tends to be more straightforward [84]. However, there are disadvantages. Many of these are DSLs that specialize to target specific kinds of higher-level properties, such as a particular kind of business logic, and are thus limited in their scope. Others are specialized to reason about specifications or protocols, unattached to code which can be compiled and deployed, which then must be translated by hand down to a low-level specification. Furthermore, language embeddings tend to either omit a formalization of the semantics of the execution environment, or they make various assumptions which if inaccurate could systematically introduce unverifiable vulnerabilities.

In order to be fully rigorous, our work requires an intermediate- or high-level language embedding, which is unrestricted in the contracts it can write and properties it can specify, which includes an explicit model of the execution semantics of a blockchain, and which has a verified mechanism to extract, compile, and deploy code which has been reasoned about.

2.1.2 ConCert

We will use ConCert [10] to formally specify and verify meta properties for three reasons. Firstly, while ConCert has not yet been used to formalize meta properties of smart contracts, it is extremely well-suited to do so. The specification language is in Coq and so it is unrestricted, except by limitations of the blockchain model itself, in the kinds of mathematical statements that can be made and proved about smart contracts. It also formalizes blockchain execution semantics underlying the execution of a smart contract, which means that there is a well-defined smart contract type, `Contract`, in the context of the model, which has a specific semantics within the blockchain and which can thus be reasoned about abstractly. This, to our knowledge, is unique to ConCert. As we will see throughout this thesis, the combination of these characteristics allows us to reason unrestrictedly about smart contracts abstractly, as mathematical objects, which is key to the goal of this thesis.

Finally, ConCert’s extraction mechanism from Coq into its target smart contract languages is certified, so despite it not being a low-level language embedding like Mi-Cho-Coq, we have a high-fidelity translation into code which can be compiled and deployed. This is made possible by MetaCoq [11, 12, 126], which is a tool for reasoning about the extraction mechanism. It should be mentioned that we rely on the compilers of the target languages to be correct, though compilers can be certified and so there is no theoretical barrier which prevents ConCert’s certified extraction to extend to *e.g.* bytecode.

Thus ConCert gives us the advantages of verification at a higher level of abstraction without compromising the integrity of the verification results.

2.2 Verifying Meta Properties

Our work is set in the context of various types of contract vulnerabilities. While we argue for formally specifying and verifying meta properties to address these issues, there are alternative methods that one can take to address each of these classes of vulnerabilities. Rather than including a discussion of each of these here, in each of Chapters 4, 5, and 6, we discuss related work in formal methods which is more tailored to the particular class of vulnerability at hand.

Chapter 3

Background

Here we introduce the types and tactics of ConCert which are most relevant to the forthcoming work. We first look at the type of smart contracts in ConCert, the `Contract` type, and at the types which underlie the blockchain’s execution semantics. The latter abstracts the execution semantics at two levels: the `Environment` type, and the `ChainState` type, each of which can be acted on, respectively, by the `Action` and `ChainStep` types to model the progression of an executing blockchain.

We then discuss what contract specifications and proofs of contract invariants look like in ConCert, covering ConCert’s central custom Coq tactic, `contract_induction`. For any interested reader, the codebase and thorough documentation can be found at the ConCert GitHub repository [35].

3.1 Smart Contracts in ConCert

In ConCert, smart contracts are abstracted as a pair of functions: the initialization function, `init`, which governs how a contract initializes, and the receive function, `receive`, which governs how a contract handles calls to its entrypoints.

```
1 Record Contract (Setup Msg State Error : Type) :=
2   build_contract {
3     init :
4       Chain -> ContractCallContext -> Setup ->
5         result State Error;
6     receive :
7       Chain -> ContractCallContext -> State -> option Msg ->
8         result (State * list ActionBody) Error;
9   }.
```

Listing 3.1: The type of smart contracts in ConCert is a record type with two functions: `init`, which governs contract initialization, and `receive`, which governs contract calls.

To understand how smart contracts are modeled, let us briefly look at the `Chain`, `ContractCallContext`, `Setup`, `State`, `Msg`, `Error`, and `ActionBody` types. In brief,

- The `Chain` type carries data about the current state of the chain, such as the block height.
- The `ContractCallContext` type carries information about the context of a contract call, including the transaction sender, the transaction origin, the contract's balance, the amount of the native token (*e.g.* ETH or XTZ) sent in the transaction.
- The `Setup` type indicates what information is needed to deploy a contract.
- The `State` type is a contract's storage type.
- The `Msg` type is the type of messages a contract can receive.
- The `Error` type is the type of errors a contract can throw.
- Finally, the `ActionBody` type is ConCert's type of actions which can be emitted by a contract.

In ConCert, then, to define a smart contract one must define the `Setup`, `State`, `Msg`, and `Error` types and produce `init` and `receive` functions. As we will see, a call to a smart contract modifies the state of the blockchain by updating the contract state with the `receive` function and emits transactions of type `ActionBody`. If a call to a contract results in something of type `Error`, the execution rolls back and the `Environment` remains unchanged.

To deal with Coq's polymorphism, ConCert also features a serialized contract type `WeakContract`, though anyone doing contract verification work in ConCert should not ever encounter the `WeakContract` type explicitly. We will see the `WeakContract` type briefly in various definitions relevant to the chain's execution semantics later on. Note that, while we omitted it in Listing 3.1, because contracts need to be serialized, all four types parameterizing a contract must be serializable.

```

1 Inductive WeakContract :=
2   | build_weak_contract
3     (init :
4       Chain ->
5       ContractCallContext ->
6       SerializedValue (* setup *) ->
7       result SerializedValue SerializedValue)
8   (receive :
9     Chain ->
10    ContractCallContext ->
11    SerializedValue (* state *) ->
12    option SerializedValue (* message *) ->
13    result (SerializedValue * list ActionBody) SerializedValue).
```

Listing 3.2: The `WeakContract` type is a serialization of the `Contract` type used internally to ConCert to deal with contract polymorphism. It is defined coinductively with the `ActionBody` type.

3.2 The Blockchain in ConCert

In ConCert, the blockchain and its execution semantics are modeled at multiple levels of abstraction, which we go through here. Underlying everything is a typeclass, `ChainBase`, which represents basic assumptions made of any blockchain. This is almost always abstracted away when reasoning about smart contracts.

```
1  Class ChainBase :=
2    build_chain_base {
3      Address : Type;
4      address_eqb : Address -> Address -> bool;
5      address_eqb_spec :
6        forall (a b : Address), Bool.reflect (a = b) (address_eqb a b);
7      address_eqdec :> stdpp.base.EqDecision Address;
8      address_countable :> countable.Countable Address;
9      address_serializable :> Serializable Address;
10     address_is_contract : Address -> bool;
11   }.
12
```

Listing 3.3: The `ChainBase` typeclass, which represents basic assumptions made of any blockchain.

The basic assumptions of the `ChainBase` typeclass include an address type `Address`, which is countable and has decidable equality, and which has a distinction between wallet address and contract addresses. For example, on Tezos, this distinction can be seen in the format of the public keys, where contract addresses are of the form `KT...` and wallet addresses are of the form `tz...`

At the next level of abstraction, we have the record type `Chain`, which represents the view of the blockchain that a contract can access and interact with. The only information this type carries is the chain height, the current slot of a given block, and the finalized height.

```
1  Record Chain :=
2    build_chain {
3      chain_height : nat;
4      current_slot : nat;
5      finalized_height : nat;
6    }.
```

Listing 3.4: The `Chain` type, which represents the view of the blockchain that a contract can access and interact with.

From here, we have two types: the `Environment` type, which augments the `Chain` type to model the information that a realistic blockchain needs to implement operations, and the `ChainState` type, which augments the `Environment` type to include a queue of pending transactions that need to be executed.

The `Environment` type includes data about account balances, which contracts are at which addresses, and the states of deployed contracts.

```

1  Record Environment :=
2    build_env {
3      env_chain :> Chain;
4      env_account_balances : Address -> Amount;
5      env_contracts : Address -> option WeakContract;
6      env_contract_states : Address -> option SerializedValue;
7    }.

```

Listing 3.5: The `Environment` type augments the `Chain` type to model the information that a realistic blockchain needs to implement operations.

The `ChainState` type augments the `Environment` type to add a queue of outstanding transactions, shifting our view from the chain’s internal environment at any given block height to an external view of the chain itself, which executes transactions in a block.

```

1  Record ChainState :=
2    build_chain_state {
3      chain_state_env :> Environment;
4      chain_state_queue : list Action;
5    }.

```

Listing 3.6: the `ChainState` type augments the `Environment` type to include a queue of pending transactions that need to be executed.

Finally, we have `ChainBuilderType`, which is a typeclass representing implementations of blockchains. Part of the trust base of ConCert, then, is that the blockchain in question satisfies the semantics of the `ChainBuilderType`.

```

1  Class ChainBuilderType :=
2    build_builder {
3      builder_type : Type;
4      builder_initial : builder_type;
5      builder_env : builder_type -> Environment;
6      builder_add_block
7        (builder : builder_type)
8        (header : BlockHeader)
9        (actions : list Action) :
10      result builder_type AddBlockError;
11      builder_trace (builder : builder_type) :
12        ChainTrace empty_state (build_chain_state (builder_env builder) []);
13    }.

```

Listing 3.7: The `ChainBuilderType` typeclass characterizes implementations of blockchains.

3.2.1 Blockchain Semantics in ConCert

The `Environment` and `ChainState` types can be acted on by actions which represent the blockchain making progress by executing transactions in a block. Some of these can be initiated by users, and others relate to the blockchain’s execution semantics. The possible actions that a user can initiate are modeled by the `Action` and `ActionBody` types.

```
1 Record Action :=
2   build_act {
3     act_origin : Address;
4     act_from : Address;
5     act_body : ActionBody;
6   }.
```

Listing 3.8: The `Action` type, which includes the action’s origin, the sender, and the action’s body.

```
1 Inductive ActionBody :=
2   | act_transfer (to : Address) (amount : Amount)
3   | act_call (to : Address) (amount : Amount) (msg : SerializedValue)
4   | act_deploy (amount : Amount) (c : WeakContract) (setup : SerializedValue).
```

Listing 3.9: The `ActionBody` type, which specifies that a user can interact with the blockchain by transferring funds, calling a contract, or deploying a contract.

Every action carries with it the origin, `act_origin`, the sender, `act_from`, and what kind of action it is, whether it be a transfer, a contract call, or a contract deployment. From these we can build the types which act on the `Environment` and `ChainState` types to model the blockchain making progress.

First, let us look at the `ActionEvaluation` type, which acts on the `Environment` type. The definition of `ActionEvaluation` involves sixty-six lines of code, so we give a shortened version here.

```
1 Inductive ActionEvaluation
2   (prev_env : Environment) (act : Action)
3   (new_env : Environment) (new_acts : list Action) : Type :=
4   | eval_transfer :
5     forall (origin from_addr to_addr : Address)
6       (amount : Amount),
7     (* some omitted checks *)
8     ActionEvaluation prev_env act new_env new_acts
9   | eval_deploy :
10    forall (origin from_addr to_addr : Address)
11      (amount : Amount)
12      (wc : WeakContract)
13      (setup : SerializedValue)
14      (state : SerializedValue),
15    (* some omitted checks *)
16    ActionEvaluation prev_env act new_env new_acts
17   | eval_call :
```

```

18     forall (origin from_addr to_addr : Address)
19         (amount : Amount)
20         (wc : WeakContract)
21         (msg : option SerializedValue)
22         (prev_state : SerializedValue)
23         (new_state : SerializedValue)
24         (resp_acts : list ActionBody),
25     (* some omitted checks *)
26     ActionEvaluation prev_env act new_env new_acts.

```

Listing 3.10: The ActionEvaluation links two inhabitants of the Environment type to represent a blockchain making progress by evaluating an action.

The ActionEvaluation type is parameterized by a previous environment `prev_env`, and action `act`, a new environment `new_env`, and a list of actions `new_acts`. This models a blockchain making progress by evaluating an action, moving from the previous environment to a new environment.

Moving up to the ChainState type, we have the ChainStep type which acts on ChainState similar to how ActionEvaluation acts on Environment, forming a chain. As before, we give a shortened version of the type definition.

```

1 Inductive ChainStep (prev_bstate : ChainState) (next_bstate : ChainState) :=
2   | step_block :
3       forall (header : BlockHeader),
4           (* some omitted checks *)
5           ChainStep prev_bstate next_bstate
6   | step_action :
7       forall (act : Action)
8           (acts : list Action)
9           (new_acts : list Action),
10          ActionEvaluation prev_bstate act next_bstate new_acts ->
11          (* some omitted checks *)
12          ChainStep prev_bstate next_bstate
13   | step_action_invalid :
14       forall (act : Action)
15           (acts : list Action),
16          (* some omitted checks *)
17          ChainStep prev_bstate next_bstate
18   | step_permute :
19       EnvironmentEquiv next_bstate prev_bstate ->
20       Permutation (chain_state_queue prev_bstate) (chain_state_queue next_bstate) ->
21       ChainStep prev_bstate next_bstate.

```

Listing 3.11: The ChainStep type links two inhabitants of the ChainState type to represent a blockchain making progress.

The `ChainStep` type is parameterized by two chain states, the previous state `prev_bstate`, and the new state, `next_bstate`, and represents an update to the chain's state. The chain's state can be updated by: updating the environment with an inhabitant of an `ActionEvaluation` type, as given by `step_action`; adding a block, given by `step_block`; showing an action to be invalid, given by `setp_action_invalid`; or reordering the blockchain's transaction queue. Reordering the transaction queue is for the sake of generality, so that proofs are independent of depth-first or breadth-first transaction execution orderings, which can vary among chains.

Finally, the actual chained history of a blockchain is modeled through the `ChainTrace` type, which is a linked list of inhabitants of `ChainState`, linked by inhabitants of `ChainStep`.

```
1 Definition ChainTrace := ChainedList ChainState ChainStep.
```

Listing 3.12: The `ChainTrace` type, which models the chained history of a blockchain, and can be used to define the notion of a reachable chain state.

The `ChainedList` type models the chaining of points in some arbitrary type by a type of links, as follows.

```
1 Context {Point : Type} {Link : Point -> Point -> Type}.
2 Inductive ChainedList : Point -> Point -> Type :=
3   | clnil : forall {p}, ChainedList p p
4   | snoc : forall {from mid to},
5     ChainedList from mid -> Link mid to -> ChainedList from to.
```

Listing 3.13: The `ChainedList` type, described in the ConCert documentation as a proof-relevant transitive reflexive closure of a relation.

As we will see, the semantics of blockchain execution makes it possible for us to reason along execution traces of blockchains in a general way. In particular, the `ChainTrace` type gives us the notion of a reachable state of a blockchain, defined as a state to which there is a trace from the empty state, `empty_state`.

```
1 Definition reachable (state : ChainState) : Prop :=
2   inhabited (ChainTrace empty_state state).
```

Listing 3.14: The definition of a reachable state of a blockchain.

Many proofs of contract invariants begin by assuming a reachable chain state.

3.3 Specification and Proof in ConCert

A contract specification is simply a list of propositions, written in Coq, about a smart contract. For practical verification work, a specification typically references a specific smart contract. However, there is nothing stopping us from abstracting over smart contracts, which we will do in later chapters.

For now, let us look at a simple example of contract definition and specification. The contract in question will simply be a counter contract, which can increment and decrement a counter held in storage. We start by defining the Setup, Msg, State, and Error types.

```

1 Definition Setup := unit.
2
3 Inductive Msg :=
4 | incr (n : N)
5 | decr (n : N).
6
7 Record State :=
8   build_state { stor : Z }.
9
10 Definition Error : Type := N.

```

Listing 3.15: The counter contract’s four types Setup, Msg, State, and Error.

We then define the endpoint contracts and the contract’s main functionality.

```

1 (* endpoint functions *)
2 Definition incr_func (n : N) (st : State) :=
3   {| stor := st.stor + (Z.of_N n) |}.
4 Definition decr_func (n : N) (st : State) :=
5   {| stor := st.stor - (Z.of_N n) |}.
6
7 (* main contract functionality *)
8 Definition counter_func (st : State) (msg : Msg) : option State :=
9   match msg with
10  | incr n => Some (incr_func n st)
11  | decr n => Some (decr_func n st)
12  end.

```

Listing 3.16: The counter contract’s main functionality.

Finally, we can construct an inhabitant of Contract by defining init and receive functions.

```

1 Definition counter_init
2   (_ : Chain)
3   (_ : ContractCallContext)
4   (_ : Setup) :
5   option State :=
6     Some ({| stor := 0 |}).
7
8 Definition counter_rcv
9   (_ : Chain)
10  (_ : ContractCallContext)
11  (st : State)
12  (op_msg : option Msg) :
13  option (State * list ActionBody) :=

```



```

14     match op_msg with
15     | Some msg =>
16         match counter_func st msg with
17         | Some rslt => Some (rslt, [])
18         | None => None
19     end
20 | None => None
21 end.
22
23 Definition counter_contract : Contract Setup Msg State Error :=
24   build_contract counter_init counter_recv.

```

Listing 3.17: An inhabitant of the Contract type, defined by the init and receive functions.

Now that we have our contract `counter_contract` defined, we can prove invariants about it.

For example, we may wish to verify the property that at any given blockchain state, the value of `stor` in the state of `counter_contract` will equal the sum of the `incr` calls, minus the sum of the `decr` calls. In ConCert, we would write that statement like this:

```

1 Theorem counter_correct : forall bstate caddr (trace : ChainTrace empty_state bstate),
2   env_contracts caddr = Some (counter_contract : WeakContract) ->
3   exists cstate inc_calls,
4     contract_state bstate caddr = Some cstate /\
5     incoming_calls entrypoint trace caddr = Some inc_calls ->
6     (let sum_incr :=
7       sumN get_incr_qty inc_calls in
8     let sum_decr :=
9       sumN get_decr_qty inc_calls in
10    cstate.(stor) = sum_incr - sum_decr).

```

Listing 3.18: An invariant on `counter_contract`, which says the state of the counter is always the sum of all the `incr` calls minus the sum of the `decr` calls.

The theorem uses two functions, `get_incr_qty` and `get_decr_qty`, whose definitions we omit here but which extract from an incoming call the quantity to be incremented or decremented. Translating this theorem into prose, we would say something like:

Theorem 1 (`counter_correct`). *For all blockchain states `bstate`, contract addresses `caddr`, and chain traces `trace` from the genesis block to `bstate`, such that `caddr` is the contract address of `counter_contract`, there exists a contract state `cstate` and incoming calls `inc_calls`, such that `cstate` is the state of `counter_contract` at `bstate`, and `inc_calls` is all the incoming calls found in `trace`, such that: the value of `stor` in `cstate` is the sum of all the values of calls to the `incr` entrypoint, minus the sum of all the values of calls to the `decr` entrypoint.*

Shortened from there, this theorem states that at any reachable state, the value of `stor` in the storage of `counter_contract` is the sum of all the `incr` calls minus the sum of all the `decr` calls.

3.4 Contract Induction

An invariant like `counter_correct` is typically proved by a custom ConCert tactic `contract_induction`. As its name suggests, to prove an invariant by contract induction one proves it for a base case, contract deployment, and then for the inductive step, which consists of the various ways a blockchain can make progress.

The `contract_induction` tactic divides the proof of a contract invariant into seven subgoals. In the first six subgoals, one must (re)establish the invariant after:

1. deployment of the contract (the base case),
2. addition of a block,
3. an outgoing action,
4. a nonrecursive call,
5. a recursive call, and
6. permutation of the action queue.

Finally, in each of these steps, one can introduce facts about the contract to help with the proof. These must be proved in the seventh subgoal.

Chapter 4

A Contract's Economic Properties

The first class of meta properties that we target are the economic properties of financial smart contracts. As we saw in §1.2.1, poorly specified financial smart contracts can be vulnerable to costly economic attacks. In this chapter we rigorously develop the notion of correctness of the specification of a financial smart contract from the perspective of its economic properties.

To this end, we introduce a theoretical tool called a *metaspecification*, which is a specification of a specification, and which we will use throughout this thesis to specify and verify meta properties of financial smart contracts. In this chapter, we use it to verify that a contract specification adequately captures the economic properties which were the intent of its design. We illustrate with an example financial smart contract, called a *structured pool*, but the framework we introduce can be applied generally.

We proceed as follows. In §4.1, we discuss the problem of specifying financial smart contracts with economic intent. In §4.2, we specify the structured pool contract, introducing the economic problem it seeks to solve and mathematically proving that it addresses the problem. In §4.3, we formally introduce the notion of a metaspecification and discuss its relationship to the specification. In §4.4, we formalize the structured pool specification, treating it as an axiomatization of a contract, and separating the specification from the properties of the metaspecification. In §4.5, we formalize the structured pool metaspecification, formally proving that any contract conforming to the formalized specification exhibits the desirable properties given in §4.2. In §4.6 we conclude.

4.1 Contract Specification With Economic Intent

By definition, financial smart contracts are always specified with some economic intent, and that is manifest with varying degrees of rigor in the specification process.

The least rigorous of these communicate the contract's intended economic design through rhetorical

means such as diagrams, analogies, and imagery. Let us revisit Beanstalk, an Ethereum-based stablecoin protocol which was the victim of a 77 million USD attack in April 2022 [54]. Consider in particular the Beanstalk white paper [111]. Beanstalk is a complex protocol with many moving parts. Reading the white paper, one gets an intuitive notion for what each component of the protocol contributes to the stablecoin and its tokenomics through analogies to farming. The components of the protocol are referred to by a farm, the sun, silos, fields, barns, fertilizer, temperature, and humidity. There are stalks and seeds, which can be revitalized and fertilized. Finally, assets are referred to as ripe or unripe, and can be chopped. All of this imagery combines into that of a functioning farm, which gives an intuitive sense that the application design is correct in some sense. However, while each of the protocol components have a precise definition and technical notation, there is no codified, precise goal, and certainly no proofs that the definitions work together efficaciously.

Financial smart contracts which are specified with more rigor tend to define their goal in economic terms, and reason about how well their design satisfies those goals. Consider, for example, FairMM [40], an AMM designed to mitigate front-running attacks. It is specified with precise, mathematical definitions, and features proofs that the specification satisfies certain economic properties. Theorem 3, the main theorem which justifies the specification to be correct, shows that no rational, profit-seeking market maker has incentive to manipulate the price in advance of a trade. In principle, this addresses the problem of front-running; the reader might convince himself of this fact by reading the definition of a front-running attack in the paper’s abstract.

This leads us to a critical question. Proving theorems about a specification is undoubtedly an important step towards rigor, and indeed to ensuring that the specification captures certain economic meta properties. But *how do we choose which theorems to prove?*

In practice, there are various ways to answer this question. We can think of FairMM’s strategy as having defined a threat model, which is a common practice to evaluate software design [132]. Threat models can be designed systematically [89], and if they characterize the problem accurately then they can be useful in preventing vulnerabilities. The properties one proves in response to a threat model are those which, in the context of the model, are sufficient to neutralize or mitigate risk.

Another strategy is to explore the implications of a specification using tools like mutation and unit testing, and from there formulate propositions to be proved true of the specification. For example, Phipathananunth’s recent work uses mutation testing to identify any pathological, yet correct, behavior of a specification, and from these tests derive and verify various safety properties [110]. Similarly, the developers of Djed [140], a formally verified stablecoin, used mutation and unit testing to identify potentially pathological behavior of the specification and then targeted these behaviors with formal verification to justify the robustness of the contract specification. The method of using tools like unit and mutation testing is agnostic to any theory or model, but rather helps discover properties which, on their face, are evidently undesirable. They can be used in conjunction with a threat model or some other theory to discover where the model or theory breaks down.

In any case, the theorems one proves about a specification point to some fundamental notion of correct design, and the properties one chooses to prove should be oriented toward that notion of correctness. Thus the more one is able to couch theorems of correctness in a systematic understanding of the surrounding execution environment, corresponding threats, and economic goals of the contract, the more accurately one is able to approximate and prove correctness of a specification’s design.

In this spirit, substantial work has been done by Angeris *et al.* [6, 7], Bartoletti *et al.* [19, 20], and Xu *et al.* [136] to systematically understand market behavior of AMMs like Uniswap, and to create formal theories of DeFi and AMMs from which the properties constituting desirable behavior can be thoroughly reasoned through and derived. These studies range from highly theoretical to data-driven, and are a good start at a comprehensive understanding of the economic environment in which financial smart contracts operate.

In the coming sections, we will specify a novel smart contract designed to pool and trade tokenized carbon credits. We do so in the context of the cited studies and theories, and from them derive properties of well-behaved AMMs and other DeFi applications. While the context from which these properties are derived—the studies and formal models on which we draw—can always be improved, we argue that this is an example of rigorous contract design and specification, as there is a clear notion of correctness which is couched in substantial theoretical and practical work on the behavior of financial smart contracts.

4.2 Case Study: Structured Pools

We now move on to specify a financial smart contract, called a *structured pool*, which is designed to address issues of fungibility in blockchain-based carbon markets [125]. We will proceed by defining the economic problem that structured pools aim to solve (§4.2.1), specifying the contract (§4.2.2), and then justifying that the specification meets our economic goals (§4.2.3). Importantly, the theorems that we prove about structured pools are derived from the previously-mentioned work on theories of AMMs and DeFi [6, 7, 19, 20, 136].

Our goal will then be to formally verify this contract design in §4.4 and §4.5. We will show that the process of rigorous contract specification splits naturally into two components: the contract specification, which is a formal definition and axiomatization of the contract, and the metaspecification, which is a formal tool for reasoning about the implications and properties of the specification. Both are essential to ensure a contract’s correctness in the formal setting.

4.2.1 The Issue of Fungibility in Tokenized Carbon Credits

Tokenized carbon credits, which are of growing relevance to voluntary carbon markets [46, 48, 122], have unique metadata and are typically tokenized as non-fungible tokens (NFTs). However, this can lead to

low liquidity and high price volatility, so there is a push within the industry to make carbon credits as fungible as possible [49, 104, 124].

The current solution is to pool carbon credits which have similar features, such as a specific vintage or crediting methodology, and to value each pooled token equally within the pool [124]. Unfortunately, from a valuation perspective this discards the differences in constituent credits.

For example, Toucan, perhaps the most prominent provider of tokenized carbon credits [122], tokenizes carbon credits from the Verra registry as NFTs on Polygon [133]. Each NFT can then be fractionalized as an ERC20 token using a TCO2 token contract. Distinct TCO2 contracts are not mutually fungible because they carry the metadata of the carbon credit they fractionalize.

To achieve mutual fungibility, Toucan launched the Base Carbon Tonne (BCT) and the Nature Carbon Tonne (NCT) pools. Any TCO2 token which satisfies the acceptance criteria of one of these pools can be pooled, one-for-one, in exchange for BCT or NCT tokens, respectively. While these pools do increase token fungibility, they do so at the cost of valuing individual metadata.

Our goal then is to increase liquidity without ignoring individual token metadata, and we do so by removing the required one-for-one exchange rate. To that end we define a novel pooling mechanism, called a *structured pool*, which pools carbon credits without ignoring their differences by valuing pooled tokens relative to each other and facilitating trades between them. Thus our contract should act both as a pooling contract as well as an AMM, and any relevant properties of correctness should reflect that fact.

We now proceed to specify a structured pool contract.

4.2.2 Structured Pools

A structured pool contract has at least three entrypoints: `DEPOSIT`, `WITHDRAW`, and `TRADE`. The first two, `DEPOSIT` and `WITHDRAW`, are for pooling and unpooling constituent tokens, respectively, in exchange for a *pool token*. The exchange rate from a particular tokenized carbon credit to the pool token is called the *pooling exchange rate*, and is set individually for each carbon credit which can be pooled. These are also the entrypoints for, respectively, depositing and withdrawing liquidity used for trades, which are executed via the `TRADE` entrypoint. Each of these entrypoints is governed by equations which price trades and update the pooling exchange rates, which will see shortly.

The contract's storage must keep track of the family of tokens which can be pooled, each of which is called a *constituent token*, along with the pooling exchange rate of each token in the family. Each pooling exchange rate is assumed to be strictly positive when the contract is deployed. It must also keep track of the contract's balance of each constituent token, the address of the pool token contract, and the total number of outstanding pool tokens.

A brief comment on notation. We refer to our family of constituent tokens as T , where a token t_x in the family is described by its *token data*, which is a contract address and token ID pair. We will typically discuss trades from, *e.g.* t_x to t_y , where Δ_x refers to the quantity traded in, Δ_y refers to the quantity traded out, and x and y refer, respectively, to the quantity of each token held by the contract. We also write r_x as the pooling exchange rate in storage for token t_x .

Deposits

The `DEPOSIT` entrypoint accepts the token data of some t_x from the token family and a quantity q of tokens in t_x to be deposited. The pool contract checks that t_x is in the token family. It then transfers q tokens of t_x to itself, which is done by calling the `transfer` entrypoint of the token t_x , which is a standard entrypoint of token contracts. It simultaneously mints $q * r_x$ pool tokens and transfers them to the sender's wallet. This transaction is atomic, meaning that if any of the `TRANSFER` or `MINT` operations fail, the entire transaction fails.

Withdrawals

The `WITHDRAW` entrypoint accepts token data of some t_x from the token family and a quantity q of pool tokens the user wishes to burn in exchange for tokens in t_x . The pool contract checks that t_x is in the token family, and checks that it has sufficient tokens in t_x to execute the withdrawal transaction. The pool contract then transfers q pool tokens from the sender to itself and burns them by calling the `BURN` entrypoint, a standard entrypoint of token contracts. It simultaneously transfers $\frac{q}{r_x}$ tokens in t_x from itself to the sender's wallet. As before, the transaction is atomic, so if any of the `TRANSFER` or `BURN` operations fail, the entire transaction fails.

Trades

The `TRADE` entrypoint takes the token data of some token t_x in T to be traded in, the token data of some token t_y in T to be traded out, and the quantity Δ_x to be traded. It checks that both t_x and t_y are in the token family, that $k > 0$, and that $\Delta_x > 0$. It calculates Δ_y using formulae we will give below, and checks that it has a sufficient balance y in t_y to execute the trade action. Then in an atomic transaction, the contract updates the exchange rate r_x in response to the trade, transfers Δ_x of tokens t_x from the sender's wallet to itself, and transfers Δ_y of tokens t_y from itself to the sender's wallet. The specification is summarized in Figure 4.1.

The contract prices trades by simulating trading along the curve $xy = k$ (for some generic x and y), where k is the total number of outstanding pool tokens. A trade of Δ_x yields Δ_y tokens such that the following equation holds:

$$(x + \Delta_x)(y - \Delta_y) = k, \quad (4.1)$$

```

1 (* two auxiliary functions *)
2 fn CALCULATE_TRADE r_x r_y delta_x k =
3   let l = sqrt(k / (r_x r_y)) ;
4   l * r_x - k / (l * r_y + delta_x) ;
5
6 fn UPDATE_RATE x delta_x delta_y r_x r_y =
7   (r_x x + r_y * delta_y) / (x + delta_x);
8
9 (* pseudocode of the TRADE entrypoint *)
10 fn TRADE t_x t_y delta_x =
11   let delta_y = CALCULATE_TRADE
12     r_x r_y delta_x k ;
13   if (is_in_family t_x) &&
14     (is_in_family t_y) &&
15     (delta_x > 0) &&
16     (k > 0) &&
17     (self_balance t_y >= delta_y)
18   then
19     <atomic>
20       r_x <- UPDATE_RATE
21         x delta_x delta_y r_x r_y;
22       transfer (delta_x)
23         of (t_x)
24         from (sender)
25         to (self) ;
26       transfer (delta_y)
27         of (t_y)
28         from (self)
29         to (sender) ;
30     </atomic>
31   else
32     fail ;

```

Figure 4.1: Pseudocode of the TRADE entrypoint function.

giving

$$\Delta_y = y - \frac{k}{x + \Delta_x}. \quad (4.2)$$

This is how trades are priced in the wild for liquidity pools of fungible tokens [7]. We call $p_s = \frac{\Delta_y}{\Delta_x}$ the *swap price*.

An important consequence to (4.1) is that the smaller Δ_x is compared to k , the closer the exchange happens at a rate of $p_q = \frac{y}{x}$. This is because the derivative of $xy = k$, or $f(x) = \frac{k}{x}$, is

$$f'(x) = \frac{-k}{x^2} = \frac{-y}{x},$$

and the smaller Δ_x is relative to k , the more accurately the tangent line at some (x_0, y_0) approximates the convex curve $xy = k$. We call $p_q = \frac{y}{x}$ the *quoted price*.

The difference $p_q - p_s$ is called the *price slippage* [136, §3.2.4]. It is important to note that p_s is always less than p_q because p_q is calculated by moving Δ_x along the tangent line from a starting point (x_0, y_0) representing the current state of the contract's funds available for trading, and p_s is calculated by moving Δ_x along $xy = k$. Since $xy = k$ is convex, moving Δ_x along the tangent line always results in a larger Δ_y than moving along $xy = k$. See Figure 4.2 for a graphical illustration, where Δ_y^q is the output of a trade priced at p_q and Δ_y^s is the output of a trade priced at p_s .




Figure 4.2: A trade of $\Delta_x = 3$ for Δ_y^q and Δ_y^s , respectively, at $k = 50$. $\Delta_y^q = p_q \Delta_x$ is the trade priced at the *quoted price* p_q and $\Delta_y^s = p_s \Delta_x$ is the trade priced at the *swap price* p_s .

In particular, this means that

$$\Delta_y^s < p_q \Delta_x \quad (4.3)$$

always holds, since $\Delta_y^s = p_s \Delta_x$. This fact is crucial to the mechanics of how structured pools update relative prices in response to trading activity.

We use the pooling exchange rates to inform quoted prices between tokens, and then simulate trades along the curve $xy = k$ (for some generic x and y). If the token t_x pools at a rate of r_x , meaning r_x is the value of t_x in terms of pool tokens, and the token t_y pools at a rate of r_y , then t_x can be valued relative to t_y at a rate of

$$r_{x,y} := \frac{r_x}{r_y}. \quad (4.4)$$

It is perhaps counterintuitive that r_x is in the numerator and not the denominator of $r_{x,y}$, considering that $p_q = \frac{y}{x}$ in the generic case, but this is due to the fact that r_x indicates pool tokens per t_x , and we want $r_{x,y}$ to indicate t_y valued in terms of t_x .

To price a trade we begin by finding ℓ such that

$$(\ell r_y)(\ell r_x) = k,$$

where k is the total number of outstanding pool tokens in the contract's storage. The trade then yields Δ_y^s tokens such that

$$(\ell r_y + \Delta_x)(\ell r_x - \Delta_y^s) = k. \quad (4.5)$$

This formula yields the quoted price of this trade as

$$p_q = \frac{\ell r_x}{\ell r_y} = \frac{r_x}{r_y} = r_{x,y}, \quad (4.6)$$

as desired. The swap price, then, is

$$p_s = \frac{\Delta_y^s}{\Delta_x} \quad (4.7)$$

where

$$\Delta_y^s = \ell r_x - \frac{k}{\ell r_y + \Delta_x}. \quad (4.8)$$

For the rest of this document, we will write Δ_y^s simply as Δ_y unless explicitly stated otherwise.

After executing a trade, if we do not adjust pooling exchange rates, the pool token is now overcollateralized. We can see this because by (4.3),

$$r_y \Delta_y < r_x \Delta_x,$$

so a trade deposits *slightly more* in terms of pool tokens ($r_x \Delta_x$) than it removes ($r_y \Delta_y$). Thus the sum of the value of all the constituent tokens at their current valuation is now greater than the total number of outstanding pool tokens.

To avoid this, we adjust the values of the constituent tokens so that their sum at the new valuation is equal to the total number of outstanding pool tokens. In a trade t_x to t_y , because it is possible to deplete t_y from the pool, we cannot reliably regain pooled consistency by adjusting the value of the token being traded for in the pool. We know, however, that we have a supply of t_x because that was the deposited token. Thus to regain pooled consistency, we have to slightly devalue t_x in relation to the rest of the pool tokens. To do so, we divide the quantity of pool tokens by its collateral in t_x to get an updated exchange rate r'_x as follows:

$$r'_x := \frac{r_x x + r_y \Delta_y}{x + \Delta_x}. \quad (4.9)$$

Equation (4.9) updates the pooling exchange rate of t_x so that the pool token is neither under- nor over-collateralized.

Consider as an example a pool with three constituent tokens t_x , t_y , and t_z , and pooling exchange rates $r_x = 2$, $r_y = 1$, and $r_z = 1$. That is, t_x is valued at two pool tokens for one token, and each of t_y and t_z are valued at one pool token for one token. Suppose we have 10 t_x , 15 t_y , and 15 t_z pooled, thus having $20 + 15 + 15 = 50$ outstanding pool tokens.

Now suppose that we trade 1 t_x for slightly less than 2 t_y (the quoted price would be exactly 2). Using our formulae, $\ell = \sqrt{\frac{50}{2}} = 5$ and $\Delta_y \approx 1.67$ (slippage is high because of the small amount of liquidity). Post trade, we have in our pool 11 t_x , 13.33 t_y , and 15 t_z , giving us in the pool the equivalent in constituent tokens as

$$2 * 11 + 1 * 13.33 + 1 * 15 = 50.33$$

pool tokens with our unadjusted pooling exchange rates. To rectify this, bringing the pool back down to the value of 50 pool tokens, we slightly devalue t_x relative to the other tokens in the pool. We use the formula (4.9)

$$r'_x = \frac{\# \text{pool tokens}}{\# \text{tokens of } t_x} = \frac{r_x x + r_y \Delta_y}{x + \Delta_x} \approx 1.97,$$

which adjusts r_x so that the 11 t_x are now worth about 21.67 pool tokens instead of 22. This gives us our desired

$$1.97 * 11 + 1 * 13.33 + 1 * 15 = 50.$$

After this update, t_y is valued more in relation to t_x , which makes sense because t_x was sold to buy t_y . One t_y used to be worth half of t_x , and now it is valued at

$$\frac{r_y}{r'_x} \approx 0.508.$$

We need to make sure that the relative price of t_y didn't rise so much that if we trade back for t_x , we have more in t_x than we started with. If this were the case, we would have an opportunity for arbitrage within the structured pool, something we wish to avoid. The quoted price for trading our roughly 1.67 t_y back to t_x would give us about $0.508 * 1.67 \approx 0.848$, which is less than 1, as desired.

We end this section with a note that in these calculations, we implicitly assumed exchange rates r_x to be rational numbers by which we can multiply and divide freely so long as $r_x > 0$. Of course, implementations will include rounding error, and so we add to the specification that the `UPDATE_RATE` function return a positive number if the numerator and denominator of the quotient are positive. We also specify that the `CALCULATE_TRADE` function return a positive number if k , r_x , r_y , and Δ_x are positive, and that $\Delta_y < \frac{r_x}{r_y} \Delta_x$ always be true for successful trades.

4.2.3 Properties of Structured Pools

The structured pool contract is designed to imitate AMMs in how it prices trades and updates the pooling exchange rates. While AMMs such as Uniswap have been shown to exhibit desirable economic behaviors [7], it is not immediately obvious that the structured pool contract will do the same. To that end, we draw on work by Angeris *et al.* [6, 7], Bartoletti *et al.* [19, 20], and Xu *et al.* [136] on AMMs and DeFi, from which we derive six properties indicative of desirable market behavior from game-theoretic and economic perspectives. For each of the six properties, we give an informal definition, followed by a formal proposition and proof.

Demand Sensitivity

A trade for a given token increases its price relative to other constituent tokens, so that higher relative demand corresponds to a higher relative price. Likewise, trading one token in for another decreases the first's relative price in the pool, corresponding to slackened demand. This enforces the classical notion of supply and demand and is important to the proper functioning of an AMM, as we see in [20, §4.2].

Property 1 (Demand Sensitivity). *Let t_x and t_y be tokens in our family with nonzero pooled liquidity and exchange rates $r_x, r_y > 0$. In a trade t_x to t_y , as r_x is updated to r'_x , it decreases relative to r_z for*

all $z \neq x$, and r_y strictly increases relative to r_x .

Proof. First we prove that $r'_x < r_x$. We must prove:

$$r'_x = \frac{r_x x + r_y \Delta_y}{x + \Delta_x} < \frac{r_x x + r_x \Delta_x}{x + \Delta_x} = \frac{r_x(x + \Delta_x)}{x + \Delta_x} = r_x,$$

which holds if $r_y \Delta_y < r_x \Delta_x$. By (4.3) and (4.6):

$$\Delta_y < \frac{r_x}{r_y} \Delta_x = p_q \Delta_x,$$

so $r_y \Delta_y < r_x \Delta_x$ as desired. By the specification, r_z remains constant for all $t_z \neq t_x$ under `TRADE`, so as r_x is updated to r'_x it decreases relative to r_z . That r_y strictly increases relative to r_x is due to the fact that $r'_x < r_x$ and r_y stays constant. \square

Nonpathological prices

As relative prices shift through trades, a price that starts out nonzero never goes to zero or to a negative value. This is to avoid pathological behavior of zero or negative prices, and is true for standard AMMs like Uniswap [7, §2]. However, like most AMMs, prices can still get arbitrarily close to zero, so a constituent token which loses its value due to external factors can still become arbitrarily devalued within the pool. This is an important property so that the formulae which price trades never divide by zero.

Property 2 (Nonpathological Prices). *For a token t_x in T , if there is a contract state such that $r_x > 0$, then $r_x > 0$ holds for all future states of the contract.*

Proof. We only need to show that $r_x > 0$ implies $r'_x > 0$, since `TRADE` is the only entrypoint that updates exchange rates. Consider a contract state such that $r_x > 0$, and an incoming trade from t_x to some t_y of quantity $\Delta_x > 0$. Because Δ_y is calculated such that

$$(\ell r_y + \Delta_x)(\ell r_x - \Delta_y) = k,$$

and since r_x, r_y , and Δ_x are all positive, we know that Δ_y is positive so long as k is not zero. If k is zero, the transaction fails as we specified for the `TRADE` entrypoint, so we know that $\Delta_y > 0$. Since $r_y \Delta_y < r_x \Delta_x$ and x cannot be negative we have that

$$0 < r_y \Delta_y < r_x \Delta_x < r_x(x + \Delta_x),$$

rendering the numerator of r'_x ,

$$r_x x + r_y \Delta_y,$$

always positive. Since Δ_x is positive and x cannot be negative, the denominator of r'_x ,

$$x + \Delta_x,$$

is also positive, which gives our result. Our result holds, then, so long as the `UPDATE_RATE` function return a positive number if the numerator and denominator of the quotient are positive, which we specified for the `TRADE` entrypoint. \square

Swap Rate Consistency

For a token t_x in T and for any $\Delta_x > 0$, there is no sequence of trades, beginning and ending with t_x , such that $\Delta'_x > \Delta_x$, where Δ'_x is the output quantity of the sequence of trades. Swap rate consistency means that it is never profitable to trade in a loop, *e.g.* t_x to t_y , and back to t_x , which is important so that there are never any opportunities for arbitrage internal to the pool. This is similar to the assertion that trading cost be positive [7, §2], that trading from t_x to t_y , and back to t_x not be profitable in [6, §3, §4.1].

Property 3 (Swap Rate Consistency). *Let t_x be a token in our family with nonzero pooled liquidity and $r_x > 0$. Then for any $\Delta_x > 0$ there is no sequence of trades, beginning and ending with t_x , such that $\Delta'_x > \Delta_x$, where Δ'_x is the output quantity of the sequence of trades.*

Proof. Consider tokens t_x, t_y , and t_z with nonzero liquidity and with $r_x, r_y, r_z > 0$. First, we claim that the following inequality holds for all $x \geq 0$ and all trades from t_x to t_y :

$$r_y \Delta_y \leq r'_x \Delta_x. \tag{4.10}$$

Since

$$r'_x = \frac{r_x x + r_y \Delta_y}{x + \Delta_x}, \tag{4.9}$$

(4.10) simplifies to

$$r_y \Delta_y (x + \Delta_x) \leq \Delta_x (r_x x + r_y \Delta_y),$$

which in turn simplifies to

$$r_y \Delta_y x \leq r_x \Delta_x x.$$

Since we know that $r_y \Delta_y \leq r_x \Delta_x$ from (4.3), we can see that our inequality holds for all $x \geq 0$, as desired.

Now we consider sequences of trades beginning and ending with t_x . For a trade t_x to t_x , we have our result because

$$\Delta'_x < \frac{r_x}{r_x} \Delta_x = \Delta_x$$

by (4.3). Now consider a trading loop from t_x to t_y , and back to t_x , for $t_y \neq t_x$. We have our result if we

can show

$$\frac{r_y}{r'_x} \Delta_y \leq \Delta_x$$

is satisfied, because $\frac{r_y}{r'_x} \Delta_y$ is an upper bound on the quantity that Δ_y can be traded for as $p_s < p_q$. This, of course, is given by (4.10) and the fact that $r'_x > 0$ from Property 2.

Finally, consider a trade from t_x to t_y , to t_z , and back to t_x . Similar to before we need to show that

$$\frac{r_z}{r'_x} \Delta_z \leq \Delta_x$$

is satisfied. But we have from (4.10) that

$$r_z \Delta_z \leq r'_y \Delta_y \leq r_y \Delta_y \leq r'_x \Delta_x,$$

as desired. This proof can be easily seen to apply to trading loops of arbitrary length, which proves our result. \square

Zero-Impact Liquidity Change

The quoted price of trades is unaffected by depositing or withdrawing liquidity [136, §3.3.1]. Typically, for AMMs such as Uniswap [3] or Curve [2] this is implemented by requiring that liquidity providers provide liquidity in pairs such that the quoted price of the AMM does not change by depositing or withdrawing liquidity. Liquidity provision works differently for structured pools, but depositing or withdrawing liquidity still does not impact quoted prices.

Property 4 (Zero-Impact Liquidity Change). *The quoted price of trades is unaffected by calling `DEPOSIT` and `WITHDRAW`.*

Proof. We have this result because the quoted price depends only on the pooling exchange rates, as we saw in (4.4), and as per the specification, only the `TRADE` entrypoint alters pooling exchange rates. \square

Arbitrage sensitivity

If an external, demand-sensitive market prices a constituent token differently from the structured pool, a sufficiently large arbitrage transaction will equalize the prices of the external market and the structured pool, or deplete the pool. In our case, this happens because prices adapt through trades due to demand sensitivity or the pool depletes in that particular token. This is generally considered to be an important property so that prices adjust in line with supply and demand, see [20, §4.3] and [7].

Property 5 (Arbitrage sensitivity). *Let t_x be a token in our family with nonzero pooled liquidity and $r_x > 0$. If an external, demand-sensitive market prices t_x differently from the structured pool, then*

assuming sufficient liquidity, with a sufficiently large transaction either the price of t_x in the structured pool converges with the external market, or the trade depletes the pool of t_x .

Proof. Suppose the structured pool prices a constituent token t_x higher than an external market. Then an arbitrageur can buy t_x elsewhere and sell them into the structured pool. Doing so devalues t_x relative to the other tokens, as we have shown. Recall that $0 < r'_x < r_x$, so to prove our result we just need to show that 0 is the greatest lower bound of r'_x . Note that by definition, $\Delta_y = \Delta_y^s$, so substituting (4.8)

$$\Delta_y^s = \ell r_x - \frac{k}{\ell r_y + \Delta_x},$$

$$r'_x = \frac{r_x x + r_y \Delta_y}{x + \Delta_x} = \frac{r_x x + \ell r_x r_y - \frac{r_y k}{\ell r_y + \Delta_x}}{x + \Delta_x}.$$

Then

$$r'_x < \frac{r_x x + \ell r_x r_y}{x + \Delta_x}$$

and since x , r_x , r_y , and ℓ are constants for a trade, for any r , $0 < r < r_x$, by choosing a sufficiently large Δ_x we can make $r'_x < r$. Thus assuming sufficient external liquidity, we have our result.

Now suppose the structured pool prices a constituent token t_x lower than an external market. Then an arbitrageur can buy t_x from the structured pool and sell them elsewhere. Doing so does not change r_x , as per the specification. However, the external market is demand sensitive, so the price of t_x will decrease on that market. Then we know that after a trade of $\Delta_x = x$, either the external market now prices t_x lower than the structured pools contract, meaning there was some

$$\Delta'_x < \Delta_x$$

which gives our result, or the trade depletes the pool of t_x , giving our result. \square

Pooled Consistency

The number of outstanding pool tokens is equal to the value, in pool tokens, of all constituent tokens held by the contract. Mathematically, the sum of all the constituent, pooled tokens, multiplied by their value in terms of pooled tokens, always equals the total number of outstanding pool tokens. This means that the pool token is never under- or over-collateralized, and is similar to standard AMMs, where the LP token is always fully backed, representing a percentage of the liquidity pool, and is encoded in the literature as *preservation of net worth* [20, §3].

Property 6 (Pooled Consistency). *The following equation always holds:*

$$\sum_{t_x} r_x x = k \tag{4.11}$$

Proof. As a base case, by the specification, at the time of contract deployment $k = 0$ and we have no pooled liquidity, so (4.11) holds trivially because $x = 0$ for all t_x . For our inductive step, consider a contract state for which (4.11) holds. If we call `DEPOSIT`, (4.11) holds by definition because for a deposit of d_x of t_x , we mint $r_x d_x$ pool tokens. The same is true if we call `WITHDRAW`. Finally, if we call `TRADE` from tokens t_x to t_y , then there is an excess number of tokens in t_x , violating (4.11). This excess is quantified in (4.9) and remedied by adjusting r_x to r'_x as we saw before. \square

4.3 Specification and Metaspecification

The contents of the previous section show us that a correct specification has two components. The first is the contract specification, an axiomatization of the contract which defines the minimal structure needed for a contract implementation to be considered, in this case, a structured pool contract. The second is the *metaspecification*, which consists of properties that justify the specification to be correct. It proves theorems about an arbitrary contract satisfying that specification, which themselves are derived from some broader theory or understanding of the execution context.

These two components are complementary, but separate. The specification should be minimal, since in practice we wish to impose as few constraints as possible on the implementation of any given smart contract and minimize any verification work needed to prove the contract correct. Conversely, the metaspecification should be as comprehensive as possible in order to understand all the properties and behaviors associated with the specification that we possibly can.

In what follows, we formalize the contract specification in ConCert as a predicate on contracts (§4.4), and then formalize the metaspecification as a list of properties which can be proved by assuming a contract that conforms to the specification into the context (§4.5). As we will see, the metaspecification informs the specification, and vice versa, not only in the design of the specification but also in its formalization.

4.4 Formal Specification: A Contract Axiomatization

The formulation of a formal specification is not generally treated systematically in formal verification. Most formal verification is done on specific implementations of contracts, proving an instance of an informal specification. The specification is typically formalized ad hoc and not abstracted as a standalone, formal object. In particular, this means that there is no obvious way to reason abstractly and formally about one or various contracts which conform to a given specification.

We wish to be more systematic. We propose a generic reasoning technique which formalizes a specification as a predicate on smart contracts, axiomatizing that contract in a formal, mathematical way. Using ConCert, we formalize the specification of a structured pool contract [125], introducing a predicate on

the contract type.

```
is_structured_pool : forall (C : Contract Setup Msg State Error), Prop.
```

The predicate is a conjunction of each formal property of the specification. It delineates formally what it means to be a structured pool contract, and allows us to reason about the specification's consequences by assuming some contract `C` and a proof `(is_sp : is_structured_pool C)`. We will describe the formal properties of the specification here, and summarize them at the end of the section in Table 4.1.

4.4.1 The Structured Pool Formal Specification

We now outline the formal specification of structured pools. Recall from Chapter 3 that in ConCert, a smart contract is a record type of two functions: `init`, which describes how the contract initializes, and `receive`, which gives the semantics of a call to a contract entrypoint. Contracts are also parameterized by four types: `Setup`, the type to initialize, `Msg`, the entrypoint type, `State`, the storage type, and `Error`, the error type.

To formalize a contract specification, we first introduce typeclasses which characterize the contract types. We then specify how the contract must initialize. Finally, we specify contract calls by individually specifying each entrypoint.

4.4.2 Typeclasses to Characterize Contract Types

Of the four contract types, here we will look at the specification of `State`, the storage type, and `Msg`, the entrypoint type, of the structured pool contract. The formalized typeclasses of the remaining two contract types can be found in Appendix A.1.1.

First, the storage type. The informal specification states that the contract storage must contain:

- the exchange rates for each constituent token,
- the quantity of constituent tokens held in the pool,
- the address of the pool token, and
- the number of outstanding pool tokens.

We can specify this by using a typeclass, which simply requires that the storage type `T` of a structured pool contract have functions which reveal each of these data points.

- `stor_rates : T -> FMap token exchange_rate`
- `stor_tokens_held : T -> FMap token N`

- `stor_pool_token : T -> token`
- `stor_outstanding_tokens : T -> N`

```

1 Class State_Spec (T : Type) :=
2   build_state_spec {
3     (* the exchange rates *)
4     stor_rates : T -> FMap token exchange_rate ;
5     (* token balances *)
6     stor_tokens_held : T -> FMap token N ;
7     (* pool token data *)
8     stor_pool_token : T -> token ;
9     (* number of outstanding pool tokens *)
10    stor_outstanding_tokens : T -> N ;
11  }.

```

Listing 4.1: The typeclass characterizing the storage type of a structured pool contract.

Moving on, consider `Msg`, the entrypoint type. There must be at least three entrypoints: `POOL`, `UNPOOL`, and `TRADE`. Our specification first identifies the minimal amount of information needed to call each of these entrypoints by defining three types:

- `pool_data`, the payload type for the `POOL` entrypoint,
- `unpool_data`, the payload type for the `UNPOOL` entrypoint, and
- `trade_data`, the payload type for the `TRADE` entrypoint.

We then specify that the entrypoint type have at least three (but possibly more) entrypoints by a typeclass which requires the following functions into the entrypoint type `T`:

- `pool : pool_data -> T`
- `unpool : unpool_data -> T`
- `trade : trade_data -> T`
- `other : other_entrypoint -> option T`

The function `other`, going from some generic type `other_entrypoint` to `option T`, allows an implementation of the structured pool contract to optionally contain more entrypoints than the required three.

```

1 Class Msg_Spec (T : Type) :=
2   build_msg_spec {
3     pool : pool_data -> T ;
4     unpool : unpool_data -> T ;
5     trade : trade_data -> T ;
6     (* any other potential entrypoints *)
7     other : other_entrypoint -> option T ;
8   }.

```

Listing 4.2: The typeclass characterizing the entrypoint type of a structured pool contract.

When reasoning about a contract in ConCert, one must be able to reason about all entrypoints. Since we allow optionally for more than three entrypoints via the `Msg_Spec` typeclass, we need to include in the specification a stipulation that for any type T satisfying `Msg_Spec`, all inhabitants of T must be able to be written in terms of these four functions. That is, for any $m:T$, m is either:

- `pool p`, for some p ,
- `unpool u`, for some u ,
- `trade t`, for some t , or
- Some m equals `other o`, for some o .

In particular, this allows us to mimic induction over an arbitrary inhabitant of the message type, despite that type not being explicitly defined as an inductive type. We do so by formally encoding the following proposition into the contract specification:

Proposition 1. *For all $m:Msg$, one of the following holds:*

1. $m = \text{pool } p$, for some p
2. $m = \text{unpool } u$, for some u
3. $m = \text{trade } t$, for some t
4. Some $m = \text{other } o$, for some o .

```

1 Definition msg_destruct (contract:Contract Setup Msg State Error) :=
2   forall (m : Msg),
3     (exists p, m = pool p) \/
4     (exists u, m = unpool u) \/
5     (exists t, m = trade t) \/
6     (exists o, Some m = other o).

```

Listing 4.3: The formalization of Proposition 1, which is used to destruct inhabitants of the entrypoint type.

See Appendix A.1.1 for the formalized typeclasses characterizing `Setup`, the `setup` type, and `Error`, the error type.

4.4.3 Specifying Contract Initialization

We now specify the structured pool contract's `init` function, which governs how it initializes.

A structured pool contract must initialize with positive exchange rates, with no pooled tokens, and with no outstanding pool tokens. We encode each of these notions in three properties of the specification, formalized virtually identically to the informal statements made in the informal specification. They are as follows.

1. `initialized_with_positive_rates`,
2. `initialized_with_zero_balance`, and
3. `initialized_with_zero_outstanding`

Since rates are encoded in a map from tokens to exchange rates as we saw in 4.4.2, the first of these, `initialized_with_positive_rates`, stipulates that for all initialized contract states `cstate`, a rate `r` corresponding to a token `t` must satisfy $r > 0$. The second, `initialized_with_zero_balance`, gives us that the balance of any token `t` in storage initializes to 0. The third, `initialized_with_zero_outstanding`, stipulates that at the time of initialization, there be no outstanding pool tokens.

This covers what is explicitly mentioned in the informal specification, but also specifies that the data given in the `Setup` type is the same data used to initialize the rates and set the data for the pool token.

4.4.4 Specifying Each Contract Entrypoint

There are twenty-four properties of the full entrypoint specification, encoded as propositions. We will look at a few key properties here. For the full list of propositions, see Table 4.1.

Firstly, let us look at the specification of the `POOL` and `UNPOOL` entrypoints. One proposition in the specification of the `POOL` entrypoint is `pool_increases_tokens_held`, which specifies that when a token is pooled via a successful call to the `POOL` entrypoint, the balance of tokens held in the pool goes up in that token and the balance of all other tokens stays constant. This is given by the following proposition, formalized.

Proposition 2. *Consider a contract `contract`. Suppose that for some contract state `cstate`, chain, and contract call context, the `POOL` entrypoint of `contract` is called successfully with some payload `msg_payload`. Then in the updated contract state `cstate'`, the balance of the token pooled, given in the*

map (stor_tokens_held cstate'), is greater than its balance in the previous state, given in the map (stor_tokens_held cstate). The difference between the two is precisely the quantity pooled, and the balance for all other tokens stays constant between the two states.

```

1 Definition pool_increases_tokens_held
2   (contract : Contract Setup Msg State Error) : Prop :=
3     forall cstate chain ctx msg_payload cstate' acts,
4       (* If the call to POOL was successful, *)
5       receive contract chain ctx cstate (Some (pool msg_payload)) =
6       Ok(cstate', acts) ->
7       (* then in the new state cstate', tokens_held has
8         increased at the token pooled *)
9       let token := msg_payload.(token_pooled) in
10      let qty := msg_payload.(qty_pooled) in
11      let old_bal := get_bal token (stor_tokens_held cstate) in
12      let new_bal := get_bal token (stor_tokens_held cstate') in
13      new_bal = old_bal + qty /\
14      (* and tokens_held stays the same for all other tokens. *)
15      forall t,
16      t <> token ->
17      get_bal t (stor_tokens_held cstate) =
18      get_bal t (stor_tokens_held cstate').

```

Listing 4.4: The formalization of Proposition 2, which describes pre- and post-conditions of calling the POOL entrypoint.

The POOL and UNPOOL entrypoints are fully characterized as entrypoints by a set of propositions of a similar form to this one. These propositions describe what happens when tokens are (un)pooled, what transactions are emitted from those entrypoint calls, and how the storage gets updated.

Moving on to the specification of the TRADE entrypoint, we introduce two auxiliary functions: `calc_delta_y`, which calculates the output of a trade, typically written Δ_y in the literature, and `calc_rx'`, which calculates how exchange rates (and by implication, prices) update in response to trading activity. Structured pool contracts simulate trading along a convex curve, the most common of these being the Uniswap V1 curve $xy = k$ (see Figure 4.2).

Rather than require one implementation or curve, we simply specify that whatever functions are present in the implementation conform to a few requirements true of trades along any convex curve. Since in structured pools, the quoted price of any trade from a token t_x to a token t_y is given by the quotient of their rates

$$\frac{r_x}{r_y},$$

we require that the pricing and slippage (the difference between the quoted price of a trade and the price at which the trade is executed) of any implementation of `calc_delta_y` mimic the pricing and slippage

we get by trading along $xy = k$. For structured pools, this translates to a technical requirement that

$$r_y \Delta_y \leq r_x \Delta_x$$

for all trades from t_x to t_y of quantity Δ_x , where r_x is the exchange rate of t_x and r_y is the exchange rate of t_y . To express this in the specification, we formalize the following proposition.

Proposition 3. *Consider tokens t_x and t_y , with exchange rates r_x and r_y (resp.) and pooled balances x and y (resp.). For all trades from t_x to t_y of quantity Δ_x , the following inequality holds, where Δ_y is the quantity traded out which is calculated by `calc_delta_y`:*

$$r_y \Delta_y \leq r_x \Delta_x.$$

```
1 Definition trade_slippage :=
2   forall r_x r_y delta_x k x,
3   let delta_y := calc_delta_y r_x r_y delta_x k x in
4   r_y * delta_y <= r_x * delta_x.
```

Listing 4.5: The formalization of Proposition 3, which characterizes trade slippage.

We don't require any particular implementation of `calc_delta_y`, so long as it conforms to Proposition 3 (as well as the other relevant propositions of the specification).

Also included in the specification is a proposition that trades are calculated using the function `calc_delta_y`, formalized as follows.

Proposition 4. *Consider a contract `contract`. Suppose that for some contract state `cstate`, chain, and contract call context, the `TRADE` entrypoint is successfully called with some payload `msg_payload`, where `delta_x` is the quantity traded in. Then for the updated state `cstate'`:*

1. *The balance of the token traded in increases by `delta_x` from the previous state `cstate` to the updated state `cstate'`, and*
2. *The balance of the token traded out decreases by `delta_y` from the previous state `cstate` to the updated state `cstate'`, where `delta_y` is the quantity calculated by `calc_delta_y`.*

```
1 Definition trade_pricing
2   (contract : Contract Setup Msg State Error) : Prop :=
3   forall cstate chain ctx msg_payload cstate' acts,
4   (* the call to TRADE was successful *)
5   receive contract chain ctx cstate (Some (trade (msg_payload))) =
6   Ok(cstate', acts) ->
7   (* balances for t_x change appropriately *)
8   FMap.find (token_in_trade msg_payload) (stor_tokens_held cstate') =
9   Some (get_bal (token_in_trade msg_payload))
```

```

10  (stor_tokens_held cstate) + (qty_trade msg_payload)) /\
11  (* balances for t_y change appropriately *)
12  let t_x := token_in_trade msg_payload in
13  let t_y := token_out_trade msg_payload in
14  let delta_x := qty_trade msg_payload in
15  let rate_in := (get_rate t_x (stor_rates cstate)) in
16  let rate_out := (get_rate t_y (stor_rates cstate)) in
17  let k := (stor_outstanding_tokens cstate) in
18  let x := get_bal t_x (stor_tokens_held cstate) in
19  (* in the new state *)
20  FMap.find(token_out_trade msg_payload) (stor_tokens_held cstate') =
21  Some (get_bal(token_out_trade msg_payload) (stor_tokens_held cstate)
22  - (calc_delta_y rate_in rate_out delta_x k x)).

```

Listing 4.6: The formalization of Proposition 4, which requires that trades be priced with the function `calc_delta_y`.

The `TRADE` entrypoint is fully characterized as an entrypoint by a set of propositions similar to this one. These describe what happens when tokens are traded, what transactions are emitted from that entrypoint call, and how the storage updates.

Finally, to finish specifying the entrypoints we need to specify required behavior of any other entrypoint aside from `POOL`, `UNPOOL`, or `TRADE`. We simply specify that calling any other entrypoints does not affect exchange rates, token balances, or the outstanding tokens when called.

Proposition 5. *Consider a contract `contract`. Suppose that for some contract state `cstate`, chain, and contract call context, any entrypoint other than `POOL`, `UNPOOL`, or `TRADE` is successfully called with payload `o`. Then for the updated state `cstate'` the following hold:*

1. *the exchange rates remain constant, meaning*

$$(stor_rates\ cstate) = (stor_rates\ cstate')$$

2. *the token balances remain constant, meaning*

$$(stor_tokens_held\ cstate) = (stor_tokens_held\ cstate')$$

3. *the number of outstanding tokens remains constant, meaning*

$$(stor_outstanding_tokens\ cstate) = (stor_outstanding_tokens\ cstate').$$

	Propositions of the Specification	Summary
Preconditions	<code>none_fails, msg_destruct</code>	A contract call fails if the message payload is empty, and the message type must have at least the three entrypoints specified.
Specification of POOL	<code>pool_entrystate.check</code> <code>pool_emits_txns</code> <code>pool_increases_tokens_held</code> <code>pool_rates_unchanged</code> <code>pool_outstanding</code>	Calling POOL emits the correct transactions and alters storage correctly, and fails if someone attempts to pool a nonexistent token.
Specification of UNPOOL	<code>unpool_entrystate.check</code> <code>unpool_entrystate.check_2</code> <code>unpool_emits_txns</code> <code>unpool_decreases_tokens_held</code> <code>unpool_rates_unchanged</code> <code>unpool_outstanding</code>	Calling UNPOOL emits the correct transactions and alters storage correctly, and fails if someone attempts to unpool a nonexistent token.
Specification of TRADE	<code>trade_entrystate.check</code> <code>trade_entrystate.check_2</code> <code>trade_pricing_formula</code> <code>trade_update_rates</code> <code>trade_update_rates_formula</code> <code>trade_emits_transfers</code> <code>trade_tokens_held_update</code> <code>trade_outstanding_update</code> <code>trade_pricing</code> <code>trade_amounts_nonnegative</code>	Calling TRADE emits the correct transactions and updates the storage correctly; it prices trades and updates rates correctly; and it fails if someone attempts to trade a nonexistent token or a token with insufficient contract balance.
Specification of OTHER	<code>other_rates_unchanged</code> <code>other_balances_unchanged</code> <code>other_outstanding_unchanged</code>	Any other entrypoint must not alter exchange rates, token balances, or outstanding pool tokens.
<code>calc_delta_y</code> and <code>calc_rx'</code>	<code>rate_decrease</code> , <code>rates_balance</code> , <code>rates_balance_2</code> , <code>trade_slippage</code> , <code>trade_slippage_2</code> , <code>arbitrage_lt</code> , <code>arbitrage_gt</code>	Axiomatizes trading along a convex curve.
Contract Initialization	<code>initialized_with_positive_rates</code> <code>initialized_with_zero_balance</code> <code>initialized_with_zero_outstanding</code> <code>initialized_with_init_rates</code> <code>initialized_with_pool_token</code>	The contract initializes with positive rates, zero pooled balance, zero pool tokens outstanding, and using the initialization data.

Table 4.1: The propositions which constitute the formal specification of a structured pool contract.


```

1 Definition other_outstanding_unchanged
2   (contract : Contract Setup Msg State Error) : Prop :=
3     forall cstate cstate' chain ctx o acts,
4       (* the call to POOL was successful *)
5       receive contract chain ctx cstate (other o)
6       = Ok(cstate', acts) ->
7       (* balances all stay the same *)
8       (stor_outstanding_tokens cstate) =
9       (stor_outstanding_tokens cstate').

```

Listing 4.7: The formalization of Item 1 of Proposition 5, which requires that the number of outstanding tokens remain constant through a successful call to any entrypoint other than POOL, UNPOOL, or TRADE.

4.4.5 The Formal Specification as a Predicate on Contracts

Taken together, these propositions and typeclasses are the axiomatized definition of a structured pool contract, and we can reason about the specification if we assume the existence of some contract

$$\text{contract} : \text{Contract Setup Msg State Error}$$

such that each of Setup, Msg, State, and Error conform to their respective typeclasses Setup.Spec, Msg.Spec, State.Spec, and Error.Spec, and such that contract conforms to each of the properties defined in the specification.

To do so, we amalgamate all the properties into a predicate on contracts, which is a function

$$\text{is_structured_pool} : \text{forall } (C : \text{Contract Setup Msg State Error}), \text{Prop}.$$

See Table 4.1 for a summary. The full, formal statement of the `is_structured_pool` predicate can be found in Appendix A.1.3.

4.5 Formal Metaspecification

We now turn to give a formal treatment of the structured pool metaspecification. As we saw in §4.2.3, a metaspecification is critical to be sure that our specification is correct in the context of some underlying theory. Thus if we wish to formally verify a contract to be correct, we must verify its formal specification to also be correct. The formalization here has the same economic implications of correctness as we saw in the unformalized metaspecification of §4.2.3, but as we will see in the formal setting it has additional advantages that help us prove the *formalization* of the specification to be correct.

In particular, the reader may recall that the unformalized specification of §4.2.2 uses rational-number arithmetic in the properties of the specification and in the statements and proofs of these six properties

of the metaspecification. Using rationals or reals is common, especially with financial contracts like AMMs or lending pools, *e.g.* [18, 20]. However, we have the issue that smart contracts typically use only natural-number arithmetic, and so we need some way of ensuring that the transition from rational to natural-number arithmetic does not compromise the economic properties the contract is meant to satisfy.

4.5.1 Formalizing the Metaspecification

We begin by introducing an arbitrary structured pool contract into our Coq context, formalized as follows.

```
1 Context { contract : Contract Setup Msg State Error }
2         { is_sp : is_structured_pool contract }.
```

Listing 4.8: Introduce into the context an arbitrary contract `contract` which conforms to the structured pools specification.

We have formalized and proved all six of these properties in ConCert (see Appendix A.2), but will focus our discussion on those properties which illustrate how the formalization of the metaspecification helps to derive properties of the specification necessary for a safe transition from rational to natural-number arithmetic.

First, consider Demand Sensitivity (Property 1), which is the property that an individual token’s exchange rates decrease relative to other tokens with slackened demand, and increase with rising demand.

Property 1 (Demand Sensitivity). *Let t_x and t_y be tokens in our family with nonzero pooled liquidity and exchange rates $r_x, r_y > 0$. In a trade t_x to t_y , as r_x is updated to r'_x , it decreases relative to r_z for all $z \neq x$, and r_y strictly increases relative to r_x .*

We write the formalized theorem first in prose, and then give the formalized Coq code.

Theorem 1 (Demand Sensitivity, Formalized). *Consider a structured pool contract `contract` with state `cstate`. Furthermore, consider tokens t_x and t_y , rates r_x and r_y , and quantities x and y , where t_x is a token with nonzero pooled liquidity x and with rate $r_x > 0$, and t_y is a token with nonzero pooled liquidity y and with rate $r_y > 0$. In a trade t_x to t_y (a successful call to the contract’s `TRADE` entrypoint from t_x to t_y with $t_x <> t_y$), as r_x is updated to r'_x :*

1. r_x decreases relative to all rates r_z , for $t_z <> t_x$, and
2. r_y strictly increases relative to r_x .

```

1 Theorem demand_sensitivity cstate :
2   (* For all tokens t_x t_y, rates r_x r_y, and
3     quantities x and y, where *)
4   forall t_x r_x x t_y r_y y,
5     (* t_x is a token with nonzero pooled liquidity and
6       with rate r_x > 0, and *)
7     FMap.find t_x (stor_tokens_held cstate) = Some x /\ x > 0 /\
8     FMap.find t_x (stor_rates cstate) = Some r_x /\ r_x > 0 ->
9     (* t_y is a token with nonzero pooled liquidity and
10       with rate r_y > 0 *)
11     FMap.find t_y (stor_tokens_held cstate) = Some y /\ y > 0 /\
12     FMap.find t_y (stor_rates cstate) = Some r_y /\ r_y > 0 ->
13     (* In a trade t_x to t_y ... *)
14     forall chain ctx msg msg_payload acts cstate',
15       (* i.e.: a successful call to the contract *)
16       receive contract chain ctx cstate (Some msg) =
17       Ok(cstate', acts) ->
18       (* which is a trade *)
19       msg = trade msg_payload ->
20       (* from t_x to t_y *)
21       msg_payload.(token_in_trade) = t_x ->
22       msg_payload.(token_out_trade) = t_y ->
23       (* with t_x <> t_y *)
24       t_x <> t_y ->
25       (* ... as r_x is updated to r_x': ... *)
26       let r_x' := get_rate t_x (stor_rates cstate') in
27       (* (1) r_x decreases relative to all rates r_z,
28         for t_z <> t_x, and *)
29       (forall t_z,
30         t_z <> t_x ->
31         let r_z := get_rate t_z (stor_rates cstate) in
32         let r_z' := get_rate t_z (stor_rates cstate') in
33         rel_decr r_x r_z r_x' r_z') /\
34       (* (2) r_y strictly increases relative to r_x *)
35       let t_y := msg_payload.(token_out_trade) in
36       let r_y := get_rate t_y (stor_rates cstate) in
37       let r_y' := get_rate t_y (stor_rates cstate') in
38       rel_incr r_y r_x r_y' r_x'.

```

Listing 4.9: Theorem 1, the formalized statement of Demand Sensitivity (Property 1).

To see the first issue in the transition from rational to natural-number arithmetic, note first that the unformalized version of Demand Sensitivity (Property 1) uses a notion of “relative increase” and “relative decrease,” which can be understood easily in mathematical terms, but which needs to be encoded somehow formally. These two notions of natural numbers are defined by the following functions, used in the formalization in Listing 4.9. They state that as variables x, y , and z change to x', y' , and z' respectively, x increases relative to y if $y - x \leq y' - x'$, and x decreases relative to z if $z - x \leq z' - x'$ (see Listing 4.10).

```

1 Definition rel_incr (y x y' x' : N) :=
2   ((Z.of_N y) - (Z.of_N x) <= (Z.of_N y') - (Z.of_N x')) % Z.
3
4 Definition rel_decr (x z x' z' : N) :=
5   ((Z.of_N z) - (Z.of_N x) <= (Z.of_N z') - (Z.of_N x')) % Z.

```

Listing 4.10: The formal notion of relative increase and decrease between natural numbers.

The careful reader will notice that Property 1 states that the relative decrease is strict: For an exchange rate r_x updated to r'_x , $r'_x < r_x$. In contrast, the inequality in Theorem 1 is not strict: $r_{-x'} \leq r_{-x}$. This is because the informal specification uses rational arithmetic in all calculations, but in smart contracts do arithmetic with natural numbers. Thus if $r_{-x'} < r_{-x}$, meaning the rate update is strict for all trades, then after a finite number of trades the rate r_{-x} could update to 0, contradicting Nonpathological Prices (Property 2).

Let us look at the formalization of Pooled Consistency (Property 6), which states that the total value of pooled tokens, calculated in terms of pool tokens, is always the same as the total number of outstanding pool tokens.

Property 6 (Pooled Consistency). *The following equation always holds:*

$$\sum_{t_x} r_x x = k$$

In what follows, the function `tokens_to_values` helps to formalize the sum in Property 6 by taking all tokens with a nontrivial exchange rate and multiplies the contract’s pooled balance in that token by its exchange rate. We then fold over this list to take the sum in the formal statement.

Theorem 6 (Pooled Consistency, Formalized). *Consider a structured pool contract `contract` with state `cstate`. Then the the sum of all the constituent, pooled tokens, multiplied by their value in terms of pooled tokens, always equals the total number of outstanding pool tokens.*

```

1 Theorem pooled_consistency bstate caddr :
2   (* Consider a reachable bstate, with contract deployed at caddr ... *)
3   reachable bstate ->
4   env_contracts bstate caddr = Some (contract : WeakContract) ->
5   exists (cstate : State),
6   (* ... with state cstate. *)
7   contract_state bstate caddr = Some cstate /\
8   (* Then the sum over all tokens of rates * (qty held) equals
9     the total number of outstanding tokens. *)
10  suml (tokens_to_values
11        (stor_rates cstate)
12        (stor_tokens_held cstate)) =
13  (stor_outstanding_tokens cstate).

```

Listing 4.11: The formalization of Property 6, which requires that the sum of all the constituent, pooled tokens, multiplied by their value in terms of pooled tokens, always equals the total number of outstanding pool tokens.

```

1 Definition tokens_to_values
2   (rates : FMap token exchange_rate)
3   (tokens_held : FMap token N) : list N :=
4   List.map
5     (fun k =>
6       let rate := get_rate k rates in
7       let qty_held := get_bal k tokens_held in
8       rate * qty_held)
9   (FMap.keys rates).

```

Listing 4.12: The `tokens_to_values` function which produces a list of token balances multiplied by their exchange rates.

From the formal proof of Pooled Consistency (Theorem 6) we derived the property `rates_balance`, which stipulates that the calculations for pooling and unpooling have to be inverses. That is, if one pools tokens and then unpool the output of that transaction, they should end up with the same amount of tokens as they started with. This is an implicit property of the informal specification because it uses rational exchange rates and their inverses to pool and unpool tokens.

Aside from Pooled Consistency, two other formal proofs helped characterize the change from rational to natural-number arithmetic. The formal proof of Swap Rate Consistency (Property 3) showed that another strict inequality had to be relaxed from the informal specification. The formal proof of Arbitrage Sensitivity (Property 5) showed that prices must be able to range in the open interval $(0, \infty)$. This is a property of trading along a convex curve, and is encoded as `arbitrage_lt` and `arbitrage_gt` in the formal specification.

Formalizing the metaspecification also forced us to define the invariants on all entrypoints other than `POOL`, `UNPOOL`, or `TRADE` that we saw in §4.4.4. Specifically, these are Nonpathological Prices (Property 2)

and Pooled Consistency (Property 6), the results which required reasoning about arbitrary contract calls. Like other financial contract specifications we saw earlier [145, 40], the structured pool specification is designed to be minimal and explicitly leaves out other entrypoints. Now, because of the metaspecification, we have a specification that includes arbitrary entrypoints aside from the three explicitly specified.

4.5.2 Discussion

The complexity and nuance of proving these six results is an indication of how nontrivial the problem of correct economic specification is.

We can observe a few notable benefits of having formalized the metaspecification.

1. The formalization made it clear which properties of rational arithmetic must continue to hold for an implementation using natural-number arithmetic so that key economic behaviors of the specified contract remain intact. Specifically, these are the conditions such that Demand Sensitivity (Property 1), Swap Rate Consistency (Property 3), Arbitrage Sensitivity (Property 5), and Pooled Consistency (Property 6) hold, which are the properties of `calc_delta_y` and `calc_rx'` we see in Table 4.1. Because every instance of rational or natural-number arithmetic in the specification is designed to satisfy the economic properties of the metaspecification, these are only discoverable through the formal metaspecification.
2. Similarly, by formalizing the metaspecification we discovered assumptions about the specified and unspecified entrypoints that were implicit to the informal specification but not obvious. Now, the specification is fully precise and unambiguous, and justified to be correct by the metaspecification.

Think back to the costly attacks on BeanStalk and Mango Markets from §1.2.1. The formalization presented here shows that reasoning about a specification’s economic properties, including the pathological contract behavior exploited by the attackers, is highly nontrivial. The fact that virtually all contracts are specified and deployed, only reasoning informally at best about a specification’s economic properties, points to the potential source of many of these vulnerabilities.

Furthermore, economic bugs can be subtle. We just saw that three of the properties of the metaspecification—designed to prevent pathological economic behavior—would not hold if the transition from rational to natural-number arithmetic wasn’t done correctly. Furthermore, the conditions for a correct transition were nontrivial and discovered only by verifying the specification correct with regards to the metaspecification.

4.6 Conclusion

The goal of this chapter was to formally develop the notion of correctness of a specification of a financial smart contract from the perspective of its economic properties. We argued that rigorous specification of any financial smart contract, because it must have economic intent, must be defined using precise definitions and accompanied by theorems which prove it to be correct. Furthermore, these theorems should be couched in substantial theoretical and practical work on the behavior of financial smart contracts.

This gave rise naturally to the separation of a contract specification from its metaspecification. A metaspecification is a specification of a specification, and consists of properties which prove the specification to be correct within the context of a theory. It does so by proving that any contract conforming to the specification exhibits properties which characterize correct economic behavior.

Introducing the specification and metaspecification into the formal setting required defining a predicate on smart contracts, which is the contract specification, and then formally reasoning about an arbitrary contract which has a proof of the predicate corresponding to the specification. This allowed us to keep the specification as concise as possible so as to minimize the work required to formally verify a contract to be correct with regards to the specification, while still inheriting all the desirable economic behavior proved about in the metaspecification.

Finally, we showed that the metaspecification not only improves the rigor of the specification itself in terms of its ability to successfully capture economic meta properties, but also to prove the formalization itself to be correct. As there are always choices that have to be made when both designing and formalizing a contract specification, the metaspecification can help ensure that pathological contract behavior is not introduced because of a poorly formed or formalized specification.

We conclude with a note that the theories from which we derived the properties of our metaspecification [6, 7, 19, 20, 136] were not themselves formalized, which meant that the properties of the metaspecification were derived from the theories by hand and not formally. Future work on this topic could include formalizing said theories and using them to rigorously and systematically derive properties of financial contracts metaspecifications.

Chapter 5

Contract Upgradeability

The meta properties that we target in this chapter are upgradeability properties of financial smart contracts. As we saw in §1.2.2, poorly specified contract upgrades can introduce costly vulnerabilities. In this chapter we rigorously develop the notion of correctness of a specification of both individual contract upgrades, as well as contract upgradeability, articulating meta properties as properties of one contract in relation to those of another.

To this end, we introduce a theoretical tool called a *contract morphism*, which is a formal mechanism to structurally relate smart contracts in ConCert. We show how morphisms can be used in proof and specification, targeting upgradeable contracts in particular, and how they can be used to mathematically characterize the bounds of a contract’s upgradeability.

This chapter is organized as follows. In §5.1, we motivate contract morphisms by discussing the problem of formally specifying and verifying contract upgrades. In §5.2, we introduce contract morphisms. In §5.3, we show how contract morphisms can be used with ConCert’s contract induction tactic. In §5.4, we discuss strategies in proof and specification using contract morphisms. In §5.5 we mathematically characterize contract upgrades using contract morphisms. In §5.6 we conclude. Each section contains various code snippets; Appendix B mirrors section headings and gives a more complete version of the code from which the snippets are taken.

5.1 Contract Upgrades

Like the economic properties of a contract specification, contract upgradeability involves complex contract behavior which can be difficult to specify correctly. Blockchains do not generally feature built-in methods for upgrading a smart contract once it has been deployed. Instead, if one wishes to upgrade a smart contract, one has to encode an upgradeability framework into the contract before deployment. As we saw in §1.2.2, this is hard to do well and can result in costly bugs.

The alternative to encoding upgradeability into a smart contract is to upgrade via a hard fork. This involves deploying a new contract and convincing users to migrate to the new one, for example with each of the Uniswap upgrades [4, 70]. Especially for financial smart contracts such as AMMs which rely on liquidity providers, this transition can be expensive and difficult.

It is perhaps an indication of the complexity of deploying safe, upgradeable contracts that many financial smart contracts use hard forks. Best practices [142, 106] and established upgrade frameworks [95, 141] are not enough to prevent vulnerabilities. Like the economic properties of Chapter 4, contract upgradeability features nontrivial meta properties, so writing a correct specification is also nontrivial.

5.1.1 Specifying Upgradeability: The Diamond Framework

For example, consider the EIP-2535 Diamond upgrade framework [95], which is a specification of a popular, generic, and flexible upgrade protocol for Ethereum smart contracts. The specification describes contract storage and entrypoints, defining a system of contracts that all interact with each other. Reading the specification, one can convince one’s self that it does in fact specify an upgradeable contract. However, there is no property of the specification which precisely characterizes what it means to be upgradeable, much less *how* upgradeable or mutable this defined structure is.

Rather than giving a precise notion of upgradeability, the intuition behind what it means to be upgradeable, including the bounds of a contract’s upgradeability, are communicated in the Diamond framework through pictures, diagrams, the motivation section of the specification, and the analogy of a diamond in naming conventions for different parts of the specification. Each of these are rhetorical tools to communicate what it means to be “upgradeable” but they are not mathematically or technically rigorous. In particular, if we were to formally verify this standard, we would have no independent, mathematical notion to verify the claim that this is, indeed, upgradeable in some precise sense.

Thus, just as with the economic properties implied by a specification that we saw in Chapter 4, we see that the specification of an upgradeable contract tries to target some notion of upgradeability which we can understand intuitively; however, whether or not the specification succeeds in doing so is a matter of intuition at best. If we can articulate a precise notion of upgradeability, then we can verify a contract specification correct with regards to a metaspecification that includes these upgradeability properties.

How then do we mathematically capture the notion of upgradeability? We will approach the problem as follows. An upgradeable contract is made up of two constituent parts: the part governing contract upgrades, which is a skeleton that remains constant no matter the state of the contract, and the part corresponding to a specific contract version, which can change through upgrades. We might describe the logic governing upgrades as a contract in its own right, a metacontract of sorts which governs the data and versioning of the upgradeable contract. The contract version corresponding to any given state is one of many possibilities, which we might describe with a family of contracts which parameterizes all possible forms the upgradeable contract can take.

We will make this fully precise using contract morphisms, giving a formal decomposition of an upgradeable smart contract into a *base contract*—its immutable part—and a family of *version contracts*—its mutable part. The base contract contains the upgradeability framework, and the family of *version contracts* contains the contract functionality which can be upgraded.

5.1.2 Evolving Specifications Through Upgrades

Another aspect which merits our attention is that upgrading a contract, whether by hard fork or through built-in upgradeability, involves specifying the new contract version. How do we know the specification of the new contract accurately reflects what we wish to happen in the upgrade?

Consider a contract upgrade from the perspective of a formal specification. Generally speaking, we upgrade with a goal which relates to the previous contract version, whether it be to patch a bug, add functionality, or improve contract features. Indeed, the new specification likely relates to the old: the new should eliminate a vulnerability of the old, be backwards compatible while adding functionality to the old, or make improvements on the old, for example to be more gas-efficient.

The actual intention of the upgrade, and therefore its specification, might then be best written in relation to the old contract. Informally, if the upgrade is to patch a bug, we might specify that the new contract behave identically to the old, except that it patches the buggy functionality. If the upgrade adds functionality, we might specify that new functionality as normal, and then specify that all the functionality of the old contract still hold for the new—or specify exactly how it changes in relation to the old. Finally, if we are optimizing, we might specify that the new contract behave exactly like the old, except that it improve in some metric like gas efficiency or precision of some calculation.

Of course, in practice upgrades are not specified in relation to an older version, but rather by altering the old specification into the new, or simply starting from scratch and writing a new specification by hand. However, we run into the same issue that we saw with the economic properties implied by a specification in Chapter 4: such a specification tries to target what is meant by an upgrade, which is typically described informally by the difference between old and new contract versions; whether or not the specification succeeds is at best a matter of intuition. In the case of a financial contract, how can we be sure that small changes to the specification do not corrupt its correctness with regards to a metaspecification? As we saw in §1.2.2, it is not straightforward to get these right and can lead to costly exploits.

If we are able to formally compare the old and new contracts and their specifications, we can verify the specification of a contract upgrade correct via a metaspecification that articulates how the new contract should relate to the old, just as we have done informally above. We will show in §5.4 that a contract upgrade can be specified by formally relating the new contract version to the previous version by using contract morphisms.

5.1.3 Related Work

Little work has gone into understanding contract upgradeability from a formal perspective. There are two notable exceptions, both of which have similarities to what we propose here.

The first is work by Antonino *et al.* [14], which seeks to address vulnerabilities in contract upgrades by proposing a novel *systematic deployment framework* that requires contracts to be formally verified before they are deployed or upgraded. The framework relies on a *trusted deployer*, which is an off-chain service that vets contract creations and updates. Under this framework, at deployment of an upgradeable contract one can specify the bounds of upgradeability by requiring from the trusted deployer that certain invariants hold through all upgrades.

The second is work by Dickerson *et al.* [47], which proposes a paradigm shift for blockchains and smart contracts so that smart contracts can carry with them proofs of correctness. These are called *proof-carrying smart contracts*. These contracts can be upgradeable, so long as the upgraded contract carries a proof that it conforms to a specification by the original deployed contract. This work relates to previous work outside the context of blockchains on dynamic software updating [72] and proof-carrying code [99].

Both of these seek to limit the bounds of upgradeability by requiring that upgrades conform to some kind of a specification, where Antonino *et al.* rely on a trusted deployer to verify that an upgrade meets a specification before deploying the upgrade, and Dickerson *et al.* require a new paradigm of proof-carrying smart contracts so that the blockchain itself can verify a proof that a contract upgrade conforms to certain specified standards. In both cases, there is a desire to be able to specify certain invariants at the time of deployment that cannot be altered through upgrades.

Our approach is distinct. It requires no trusted third party, nor a fundamental paradigm shift into smart contracts that carry proofs. Rather, using contract morphisms we can mathematically characterize the bounds of a contract's upgradeability, and rigorously specify upgraded contracts in relation to older versions. If one wishes to impose invariants that cannot be changed through upgrades, one can do so by reasoning about an upgradeable contract's decomposition into its base contract and version contracts.

5.2 Contract Morphisms

We now move on to introduce contract morphisms, which are a formal mechanism to structurally compare smart contracts in ConCert. The goal of contract morphisms is to be able to reason about a contract and its specification in relation to another contract and its specification. As we will see, contract morphisms can be used to prove and to specify properties of one contract in terms of another (§5.4), as well as to decompose an upgradeable contract into its mutable and immutable parts (§5.5).

Recall that the `Contract` type, parameterized by types `Setup`, `Msg`, `State`, and `Error`, is a record type with two constructors: the `init` function, which dictates how a contract is initialized, and the `receive` function, which dictates the semantics of calling a contract entrypoint. In the context of a given state of the blockchain, the `init` function takes something of type `Setup` and, if successful, deploys the contract with an initial storage of type `State`.

```
1 init : Chain -> ContractCallContext -> Setup -> result State Error.
```

Listing 5.1: The `init` function of a smart contract.

The `receive` function then takes the current state of the contract and something of type `Msg`, and, if successful, produces an updated storage of type `State` and a list of emitted transactions.

```
1 receive : Chain -> ContractCallContext -> State -> option Msg ->
2         result (State * list ActionBody) Error.
```

Listing 5.2: The `receive` function of a smart contract.

Now consider contracts

```
C1 : Contract Setup1 Msg1 State1 Error1
C2 : Contract Setup2 Msg2 State2 Error2
```

with respective `init` and `receive` functions `init1`, `init2` and `receive1`, `receive2`. We define a morphism of contracts, which we write `f : ContractMorphism C1 C2`, as a natural transformation of `init` and `receive` functions,

```
f_init : init1 -> init2 and f_recv : receive1 -> receive2.
```

In other words, we define a function from inputs to `init1` to inputs to `init2`, and from outputs of `init1` to outputs of `init2`, such that if we transform inputs to `init1` and then take their image under `init2`, we get the same result as if we had first taken the image of the inputs under `init1` and then transformed the outputs to those of `init2`. We do the same for `receive1` and `receive2`. See Figure 5.1 for a graphical representation.




Figure 5.1: The `init` and `receive` components of a contract morphism, where I_i^1 and I_o^1 (resp. I_i^2 and I_o^2) are the input and output types of `C1.(init)` (resp. `C2.(init)`), and R_i^1 and R_o^1 (resp. R_i^2 and R_o^2) are the input and output types of `C1.(receive)` (resp. `C2.(receive)`). Both diagrams *commute*, meaning that starting from the upper-left corner, moving horizontally first and then vertically yields the same result as moving vertically first and then horizontally.

We will derive the natural transformations between `init` and `receive` functions from functions on the contract's parameterizing types, requiring four *component functions*:

- `setup_morph : Setup1 -> Setup2`
- `msg_morph : Msg1 -> Msg2`
- `state_morph : State1 -> State2`
- `error_morph : Error1 -> Error2,`

such that the following two coherence conditions hold:

- For all `c`, `ctx`, and `s`, transforming `(C1.(init) c ctx s)` using `state_morph` and `error_morph` gives us `(C2.(init) c ctx (setup_morph s))`; and
- For all `c`, `ctx`, `st`, and `op_msg`, transforming `(C1.(receive) c ctx st op_msg)` with `state_morph` and `error_morph` gives us

`C2.(receive) c ctx (state_morph st) (option_map msg_morph op_msg).`

See the definition of the type `ContractMorphism C1 C2` in Coq below.

```

1 Record ContractMorphism
2   (C1 : Contract Setup1 Msg1 State1 Error1)
3   (C2 : Contract Setup2 Msg2 State2 Error2) :=
4   build_contract_morphism {
5     (* the components of a morphism f *)
6     setup_morph : Setup1 -> Setup2 ;
7     msg_morph   : Msg1   -> Msg2   ;
8     state_morph : State1 -> State2 ;
9     error_morph : Error1 -> Error2 ;
10    (* coherence conditions *)
11    init_coherence : forall c ctx s,
12      result_functor state_morph error_morph
13        (init C1 c ctx s) =
14        init C2 c ctx (setup_morph s) ;
15    recv_coherence : forall c ctx st op_msg,
16      result_functor (fun '(st, l) => (state_morph st, l)) error_morph
17        (receive C1 c ctx st op_msg) =
18        receive C2 c ctx (state_morph st) (option_map msg_morph op_msg) ;
19  }.

```

Listing 5.3: The definition of contract morphisms in ConCert.

Here, `result_functor` is the following function, which for types `T` and `E`, simply takes functions of type `T -> T'` and `E -> E'` and returns a function of type `result T E -> result T' E'`.

```

1 Definition result_functor {T T' E E' : Type} :
2   (T -> T') -> (E -> E') -> result T E -> result T' E' :=
3   fun (f_t : T -> T') (f_e : E -> E') (res : result T E) =>
4   match res with | Ok t => Ok (f_t t) | Err e => Err (f_e e) end.

```

Listing 5.4: The result functor to transform results of init and receive.

Example 5.2.1 (Hello, world!). Consider a contract C_1 which keeps a `nat` in storage, which can be incremented by calling the contract’s unique entrypoint `incr`. Consider another contract C_2 which has the same storage type as C_1 and two entrypoints: the first, `incr'` which behaves identically to `incr`, and the second, `reset`, which can be called to reset the `nat` in storage to 0.

We can construct a contract morphism $f : \text{ContractMorphism } C_1 \ C_2$ by simply sending messages to the `incr` entrypoint of C_1 to the `incr'` entrypoint of C_2 , and using the identity function for all other component functions.

```

1 Definition msg_morph (e : entrypoint) : entrypoint' :=
2   match e with | incr _ => incr' tt end.
3 Definition setup_morph : setup -> setup := id.
4 Definition state_morph : storage -> storage := id.
5 Definition error_morph : error -> error := id.

```

Listing 5.5: The component functions of a morphism from C_1 to C_2 which sends the `incr` entrypoint of C_1 to the `incr'` entrypoint of C_2 .

After proving the coherence results, `init_coherence` and `recv_coherence`, we can construct a morphism.

```

1 Definition f : ContractMorphism C1 C2 :=
2   build_contract_morphism C1 C2 setup_morph msg_morph state_morph error_morph
3   init_coherence recv_coherence.

```

Listing 5.6: A morphism from C_1 to C_2 which sends messages to the `incr` entrypoint of C_1 to messages to the `incr'` entrypoint of C_2 .

Example 5.2.2 (Identity Morphism). One important contract morphism is the identity morphism `id_cm`, which is defined for any contract C and inhabits `ContractMorphism C C`.

```

1 Definition id_cm (C : Contract Setup Msg State Error) : ContractMorphism C C := { |
2   (* components *)
3   setup_morph := id ;
4   msg_morph   := id ;
5   state_morph := id ;
6   error_morph := id ;
7   (* coherence conditions *)
8   init_coherence := init_coherence_id C ;
9   recv_coherence := recv_coherence_id C ;
10  | }.

```

Listing 5.7: The identity contract morphism `id_cm C` defined for any contract C .

The associated coherence results are proved trivially by reflexivity.

```
1 Lemma init_coherence_id (C : Contract Setup Msg State Error) :
2   forall c ctx s,
3     result_functor id id (init C c ctx s) =
4     init C c ctx s.
```

Listing 5.8: The init coherence lemma for the identity morphism which is proved by reflexivity.

```
1 Lemma recv_coherence_id (C : Contract Setup Msg State Error) :
2   forall c ctx st op_msg,
3     result_functor
4       (fun ' (st, l) => (id st, l)) id
5       (receive C c ctx st op_msg) =
6       receive C c ctx (id st) (option_map id op_msg).
```

Listing 5.9: The receive coherence lemma for the identity morphism which is proved by reflexivity.

Example 5.2.3 (Injective, surjective morphisms). Injective and surjective functions are ubiquitous in mathematics. Injective functions, also called embeddings, are those which are fully structure-preserving. We can give a formal definition for injectivity here:

```
1 Definition is_inj {A B : Type} (f : A -> B) : Prop :=
2   forall (a a' : A), f a = f a' -> a = a'.
```

Listing 5.10: A formal definition of injectivity.

In other words, injective functions send distinct terms to distinct terms.

Surjective functions, also called quotients, are those for which every term in the function’s codomain has a preimage. Rather than being structure-preserving, surjective functions are often (but not always) *structure-compressing*. We can give a formal definition for surjectivity here:

```
1 Definition is_surj {A B : Type} (f : A -> B) : Prop :=
2   forall (b : B), exists (a : A), f a = b.
```

Listing 5.11: A formal definition of surjectivity.

We wish to define the analogues for contract morphisms. Contract embeddings, or injections, will be those which are fully structure-preserving: if $f : \text{ContractMorphism } C1 \ C2$ is an embedding, then we can think of $C2$ as having a copy of $C1$ in it. This will become relevant later, *e.g.* in Example 5.4.2.

A contract morphism f , then, is an *embedding* if all of its component functions are injective. Formalized in ConCert, we have a predicate `is_inj_cm` which is defined as follows in Listing 5.12.


```

1 Definition is_inj_cm (f : ContractMorphism C1 C2) : Prop :=
2   is_inj (setup_morph C1 C2 f) /\
3   is_inj (msg_morph C1 C2 f) /\
4   is_inj (state_morph C1 C2 f) /\
5   is_inj (error_morph C1 C2 f).

```

Listing 5.12: An embedding of contracts is a contract morphism whose component morphisms are injective.

Likewise, we can define contract quotients, or surjections. These will be contract morphisms that compress structure, and will become relevant *e.g.* in Example 5.4.1 as a tool to categorize contract behavior.

A contract morphism f , then, is a *quotient of contracts* if all of its component functions are surjective. Formalized in ConCert, we have a predicate `is_surj_cm` which is defined as follows in Listing 5.13.

```

1 Definition is_surj_cm (f : ContractMorphism C1 C2) : Prop :=
2   is_surj (setup_morph C1 C2 f) /\
3   is_surj (msg_morph C1 C2 f) /\
4   is_surj (state_morph C1 C2 f) /\
5   is_surj (error_morph C1 C2 f).

```

Listing 5.13: A quotient of contracts is a contract morphism whose component morphisms are surjective.

Example 5.2.4 (Equality of Morphisms). Given two morphisms

$$f \ g : \text{ContractMorphism } C1 \ C2,$$

we might ask ourselves whether or not they are equal. This will be relevant *e.g.* in Example 5.2.7 and Chapter 6 when we introduce contract isomorphisms and equivalences.

By assuming proof irrelevance, we get $f = g$ if and only if each of the component functions are equal via function extensionality. This is because the proofs of coherence for f have the same type as those for g if their component functions are equal.

```

1 Lemma eq_cm_iff :
2   forall (f g : ContractMorphism C1 C2),
3     (* the components are equal ... *)
4     (setup_morph C1 C2 f) = (setup_morph C1 C2 g) /\
5     (msg_morph C1 C2 f) = (msg_morph C1 C2 g) /\
6     (state_morph C1 C2 f) = (state_morph C1 C2 g) /\
7     (error_morph C1 C2 f) = (error_morph C1 C2 g) <->
8     (* ... iff the morphisms are equal *)
9     f = g.

```

Listing 5.14: Equality of contract morphisms.

5.2.1 Composition of Morphisms

Contract morphisms can be composed. We define composition via a function `compose_cm`, which takes morphisms

`f : ContractMorphism C1 C2 and g : ContractMorphism C2 C3`

and returns a morphism

`compose_cm g f : ContractMorphism C1 C3.`

To compose contract morphisms, we simply compose their component functions.

```

1 Definition compose_cm (g : ContractMorphism C2 C3) (f : ContractMorphism C1 C2) :
2   ContractMorphism C1 C3 := { |
3     (* the components *)
4     setup_morph := compose (setup_morph C2 C3 g) (setup_morph C1 C2 f) ;
5     msg_morph   := compose (msg_morph   C2 C3 g) (msg_morph   C1 C2 f) ;
6     state_morph := compose (state_morph C2 C3 g) (state_morph C1 C2 f) ;
7     error_morph := compose (error_morph C2 C3 g) (error_morph C1 C2 f) ;
8     (* the coherence results *)
9     init_coherence := compose_init_coh g f ;
10    recv_coherence := compose_recv_coh g f ;
11  }.
```

Listing 5.15: Composition of contract morphisms in ConCert.

The function `compose_cm` relies on two lemmas, `compose_init_coh g f` and `compose_recv_coh g f`, which prove the coherence results for the composed natural transformations of Figure 5.1. These lemmas simply show that commuting diagrams compose. That is, if we have the diagram below such that each of the left and right squares commute, then the outer rectangle also commutes.

$$\begin{array}{ccccc}
 A & \xrightarrow{i_1} & C & \xrightarrow{j_1} & E \\
 \downarrow f & & \downarrow g & & \downarrow h \\
 B & \xrightarrow{i_2} & D & \xrightarrow{j_2} & F
 \end{array}$$

```

1 Lemma compose_init_coh (g : ContractMorphism C2 C3) (f : ContractMorphism C1 C2) :
2   let setup_morph' := (compose (setup_morph C2 C3 g) (setup_morph C1 C2 f)) in
3   let state_morph' := (compose (state_morph C2 C3 g) (state_morph C1 C2 f)) in
4   let error_morph' := (compose (error_morph C2 C3 g) (error_morph C1 C2 f)) in
5   forall c ctx s,
6     result_functor state_morph' error_morph'
7       (init C1 c ctx s) =
8       init C3 c ctx (setup_morph' s).
```

Listing 5.16: The init coherence lemma for morphism composition.

```

1 Lemma compose_recv_coh (g : ContractMorphism C2 C3) (f : ContractMorphism C1 C2) :
2   let msg_morph' := (compose (msg_morph C2 C3 g) (msg_morph C1 C2 f)) in
3   let state_morph' := (compose (state_morph C2 C3 g) (state_morph C1 C2 f)) in
4   let error_morph' := (compose (error_morph C2 C3 g) (error_morph C1 C2 f)) in
5   forall c ctx st op_msg,
6     result_functor
7       (fun ' (st, l) => (state_morph' st, l)) error_morph'
8       (receive C1 c ctx st op_msg) =
9       receive C3 c ctx (state_morph' st) (option_map msg_morph' op_msg).

```

Listing 5.17: The receive coherence lemma for morphism composition.

Example 5.2.5 (Composition with the identity morphism). Composing any morphism with the identity morphism does nothing, and the proof is trivial using the equality lemma `eq_cm_iff` from Example 5.2.4.

```

1 Lemma compose_id_cm_left (f : ContractMorphism C1 C2) :
2   compose_cm (id_cm C2) f = f.
3
4 Lemma compose_id_cm_right (f : ContractMorphism C1 C2) :
5   compose_cm f (id_cm C1) = f.

```

Listing 5.18: Left and right composition of the identity morphism is trivial.

Example 5.2.6 (Morphism Composition is Associative). Composition is associative. This is a trivial proof using the equality lemma `eq_cm_ff` from Example 5.2.4, since composition of each component function of a contract morphism is also associative.

```

1 Lemma compose_cm_assoc
2   (f : ContractMorphism C1 C2)
3   (g : ContractMorphism C2 C3)
4   (h : ContractMorphism C3 C4) :
5   compose_cm h (compose_cm g f) =
6   compose_cm (compose_cm h g) f.

```

Listing 5.19: Composition of contract morphisms is associative

Example 5.2.7 (Contract Isomorphism). Finally, using notions of composition, identity, and equality, we can define contract *isomorphisms*. As we will see in Chapter 6, these are morphisms for which the formal, structural relationship expressed between the contracts is an equivalence. We define these via the following predicate.

```

1 Definition is_iso_cm (f : ContractMorphism C1 C2) (g : ContractMorphism C2 C1) : Prop :=
2   compose_cm g f = id_cm C1 /\
3   compose_cm f g = id_cm C2.

```

Listing 5.20: The formal definition of an isomorphism of contracts.

That is, morphisms $f : \text{ContractMorphism } C1 \ C2$ and $g : \text{ContractMorphism } C2 \ C1$ are together an isomorphism of contracts if, when composed in either order, their composition is equal to the identity morphism. As we will see in Chapter 6, a contract isomorphism results in a bisimulation of contracts.

5.3 Morphism Induction

In sections 5.4 and 5.5 we will explore specification and proof using contract morphisms in depth. Before doing so, we must introduce a proof technique which incorporates contract morphisms with ConCert’s contract induction tactic in order to prove contract invariants.

Consider contracts $C1$ and $C2$ and a contract morphism

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

We prove two theorems, addressing two cases. The first, which we call left morphism induction, can be used to prove an invariant of $C1$ using known properties of $C2$. The second, which we call right morphism induction, can be used to prove an invariant of $C2$ using known properties of $C1$.

5.3.1 Contract Trace and Reachability

Before introducing left and right morphism induction, we need the notions of contract trace and contract reachability. These are analogous to ConCert’s notions of chain trace and reachable chain states. The two are related in that the state of a contract deployed in a reachable chain state is always a contract-reachable state, but contract-reachable states could in principle include some states which do not correspond to any reachable chain state.

First, we define contract steps as successful calls to the `receive` function.

```

1 Record ContractStep (C : Contract Setup Msg State Error)
2   (prev_cstate : State) (next_cstate : State) :=
3   build_contract_step {
4     seq_chain : Chain ;
5     seq_ctx : ContractCallContext ;
6     seq_msg : option Msg ;
7     seq_new_acts : list ActionBody ;
8     (* we can call receive successfully *)
9     recv_some_step :
10       receive C seq_chain seq_ctx prev_cstate seq_msg = Ok (next_cstate, seq_new_acts) ;
11   }.

```

Listing 5.21: Contract steps are successful calls to the `receive` function.

Then a contract’s trace is a chained list of contract states, linked together by contract steps.

```

1 Definition ContractTrace (C : Contract Setup Msg State Error) :=
2   ChainedList State (ContractStep C).

```

Listing 5.22: A contract’s trace is a chained list of contract states, linked together by contract steps.

This gives us a natural notion of a contract-genesis state, which is defined as a state into which it is possible for the contract to initialize.

```

1 Definition is_genesis_state (C : Contract Setup Msg State Error) (init_cstate : State) :=
2   exists init_chain init_ctx init_setup,
3   init C init_chain init_ctx init_setup = Ok init_cstate.

```

Listing 5.23: A contract genesis state is one into which a contract can initialize.

It also gives us a natural notion of a reachable contract state, which is a state with a trace from some contract genesis state.

```

1 Definition cstate_reachable (C : Contract Setup Msg State Error) (cstate : State) :=
2   exists init_cstate,
3   (* init_cstate is a valid initial cstate *)
4   is_genesis_state C init_cstate /\
5   (* with a trace to cstate *)
6   inhabited (ContractTrace C init_cstate cstate).

```

Listing 5.24: A contract-reachable state is one for which there exists a trace from a contract genesis state.

We will use these definitions in the remainder of this chapter, as well as in Chapter 6.

5.3.2 Left Morphism Induction

Recall from §3.4 that to prove a contract invariant with contract induction, one proves the invariant on the base case (contract deployment), and the inductive step consists of all the ways that the blockchain can make progress. A contract morphism $f : \text{ContractMorphism } C1 \ C2$ indicates that there is a structural relationship between $C1$ and $C2$ at each of the relevant steps of induction: first at contract deployment, the base case, via the transformation of the `init` function; and then at each inductive step via the transformation of the `receive` function.

This means that when performing contract induction on $C1$, we should have access to the corresponding relational structure of $C2$ —and the same going the other way if we’re inducting on $C2$. As we will see, we might wish to prove an invariant of $C1$ in terms of an invariant on $C2$.

Left morphism induction facilitates such a goal, stating that for any blockchain state `bstate` where $C1$ is deployed at some contract address `caddr` with state `cstate1`, there exists a corresponding state `cstate2` of $C2$, which is a contract-reachable state of $C2$, and which is the image of `cstate1` under the `state_morph` component of f .

```

1 (* f : C1 -> C2, inducting on C1 *)
2 Theorem left_cm_induction :
3   (* forall simple morphism and reachable bstate, *)
4   forall (f : ContractMorphism C1 C2) bstate caddr
5     (trace : ChainTrace empty_state bstate),
6   (* where C is at caddr with state cstate, *)
7   env_contracts bstate caddr = Some (C1 : WeakContract) ->
8   exists (cstate1 : State1),
9   contract_state bstate caddr = Some cstate1 /\
10  (* every reachable cstate1 of C1 corresponds to a contract-reachable cstate2 of C2: *)
11  exists (cstate2 : State2),
12  (* 1. init_cstate2 is a valid initial cstate of C'  *)
13  cstate_reachable C2 cstate2 /\
14  (* 2. cstate and cstate' are related by state_morph. *)
15  cstate2 = state_morph C1 C2 f cstate1.

```

Listing 5.25: Left contract morphism induction.

In particular, if we can deduce an invariant of `cstate1` from a known invariant of `cstate2` by their relationship defined by `state_morph`, then we can prove our invariant of `C1` via an invariant of `C2`.

5.3.3 Right Morphism Induction

Right morphism induction takes the opposite case: we prove an invariant of `C2` via known properties of `C1` using a morphism `f : ContractMorphism C1 C2`. Because of the direction of `f`, we do not know that every reachable state of `C2` has a corresponding reachable state of `C1` related to that of `C2` via `state_morph`. Rather, we prove that if we can find a contract trace of `C1`, it can be transformed into one of `C2`.

```

1 (* f : C1 -> C2, inducting on C2 *)
2 Theorem right_cm_induction:
3   forall (from to : State1) (f : ContractMorphism C1 C2),
4   (* every contract trace, and thus reachable state, of C1 ... *)
5   ContractTrace C1 from to ->
6   (* has a corresponding contract trace of C2 *)
7   ContractTrace C2 (state_morph C1 C2 f from) (state_morph C1 C2 f to).

```

Listing 5.26: Right contract morphism induction.

Right morphism induction, then, can be used to prove the existence of some contract-reachable state of `C2` via one of `C1`. For example, if `C2` is an upgraded version of `C1`, one way to specify backwards compatibility is that every contract-reachable state of `C1` has a corresponding contract-reachable state of `C2` which preserves essential data. We prove this by right morphism induction in Example 5.4.2.

5.4 Reasoning with Morphisms: Specification and Proof

Morphisms can be used in a variety of ways in proof and specification of smart contracts. We will demonstrate with three examples: specifying a contract upgrade in relation to the old (§5.4.1), proving backwards compatibility (§5.4.2), and proving Hoare-like properties of partial correctness (§5.4.3).

5.4.1 Specifying a Contract Upgrade With Morphisms

Let us revisit the Uranium Finance exploit from §1.2.2.

Example 5.4.1 (Uranium Finance Upgrade Specification). Recall from §1.2.2 Uranium Finance, a contract which was exploited because developers replaced a constant `k` set at 1,000 with 10,000 in all but one of its instances during an upgrade. The result was wildly incorrect pricing, which rapidly drained their liquidity pools.

Suppose `C1` is the Uranium Finance contract pre-upgrade, and `C2` is the contract post-upgrade, and suppose the upgrade was only to adjust how the contract prices trades. Suppose further that the upgrade was to increase the decimal precision of this calculation by a factor of ten, meaning that the internal token balances in storage have one more decimal place, and the trade calculation is able to calculate at one decimal place greater in precision.

The original contract `C1` will have a storage type, then, which keeps track of internal token balances.

```
1 Context { storage : Type } { get_bal : storage -> N }.
```

Listing 5.27: We assume a storage type and a function which calculates balances.

It will also have a `TRADE` entrypoint which accepts a message whose payload includes a desired trade quantity, which we formalize with a typeclass as we did in Chapter 4.

```
1 (* A typeclass which characterizes the entrypoint type *)
2 Class Msg_Spec (T : Type) := {
3   trade : trade_data -> T ;
4   (* for any other entrypoint types *)
5   other : other_entrypoint -> option T ;
6 }.
7
8 (* We assume an entrypoint conforming to Msg_Spec *)
9 Context { entrypoint : Type } { e_msg : Msg_Spec entrypoint }.
```

Listing 5.28: We assume an entrypoint type, characterized by `Msg_Spec`, which includes a trade function, and introduce into the context an entrypoint type and an instance of `Msg_Spec entrypoint`.

We now assume that `C1` has some function `calculate_trade` which calculates how many tokens will be traded in for a given contract call to the `TRADE` entrypoint. The trade quantity, internal token balances,

and the `calculate_trade` function will all be accurate up to some decimal place, commonly 6 or 9 in the wild. This is formalized in the following property, which we assume into our context.

```

1 (* get_bal changes according to calculate_trade, meaning that: *)
2 Definition spec_trade : Prop :=
3   forall cstate chain ctx trade_data cstate' acts,
4     (* for any successful call to the trade entrypoint of C1, *)
5     receive C1 chain ctx cstate (Some (trade trade_data)) = Ok(cstate', acts) ->
6     (* the balance in storage updates as follows. *)
7     get_bal cstate' =
8     get_bal cstate + calculate_trade (trade_qty trade_data).

```

Listing 5.29: The formalized proposition that C1 uses `calculate_trade` to price trades.

Now when we upgrade C1 to C2 with the exclusive purpose of increasing the accuracy by one decimal place for internal token balances and trades, we update `calculate_trade` to `calculate_trade_precise` which calculates at one higher decimal place of accuracy. This is formalized as follows.

```

1 (* an auxiliary function which rounds down by one decimal place *)
2 Definition round_down (n : N) := n / 10.
3
4 (* we assume that calculate_trade_precise is related to calculate_trade by round_down *)
5 Context { calculate_trade_precise : N -> N }
6   (* (i.e. calculate_trade_precise rounds down to calculate_trade) *)
7   { calc_trade_coherence : forall n,
8     round_down (calculate_trade_precise n) =
9     calculate_trade (round_down n) }.

```

Listing 5.30: We assume a function `calculate_trade_precise`, which always rounds down to the `calculate_trade` function.

Then C2 is assumed to use `calculate_trade_precise` to calculate its trades, formalized in this assumed proposition which is the analogue of `spec_trade` but for C2.

```

1 (* Now trades are calculated in line with calculate_trade_precise. *)
2 Definition spec_trade_precise : Prop :=
3   forall cstate chain ctx trade_data cstate' acts,
4     (* ... meaning that for a successful call to the trade entrypoint of C2, *)
5     receive C2 chain ctx cstate (Some (trade trade_data)) = Ok(cstate', acts) ->
6     (* the balance held in storage goes up by calculate_trade_precise. *)
7     get_bal cstate' =
8     get_bal cstate + calculate_trade_precise (trade_qty trade_data).

```

Listing 5.31: The formalized proposition that C2 uses `calculate_trade_precise` to price trades.

Assuming that we have already formally verified C1 correct with respect to some formal specification (and metaspecification), and wished to do the same for C2, we might do so by altering the specification of C1 until it correctly specifies the contract C2, and then verify C2 from the ground up. As nearly all of the

contract functionality remained unchanged in the upgrade, this would require re-proving many results already proved about `C1`, where the slight changes we made to `C1` would likely make it impossible to simply copy/paste the proofs. Furthermore, as we saw in the previous chapter, specifications are difficult to write correctly, and small changes to the specification may have unintended consequences. And in altering the specification we might make the same error that the Uranium Finance engineers did and unintentionally create a vulnerability.

To ensure safety, our metaspecification specifies a relationship between `C2` and `C1`: `C2` should “do everything that `C1` does,” except that internal balances and trades should be calculated at one decimal place higher in precision. We can articulate this rigorously through a contract morphism f from the contract upgrade `C2` to the original contract `C1` that has the following properties.

First, that when f sends inputs of `C2.receive` to `C1.receive`, it rounds down the precision of the requested trades using our rounding function `round_down`.

```
1 (* 1. f rounds trades down when it maps inputs of the receive function *)
2 Definition f_recv_input_rounds_down (f : ContractMorphism C2 C1) : Prop :=
3   forall t', exists t,
4     (msg_morph C2 C1 f) (trade t') = trade t /\
5     trade_qty t = round_down (trade_qty t').
```

Listing 5.32: f rounds down the inputs of the receive function.

Second, that f is the identity morphism on all entrypoints aside from the `trade` entrypoint.

```
1 (* 2. aside from trade, f doesn't touch the other entrypoints *)
2 Definition f_recv_input_other_equal (f : ContractMorphism C2 C1) : Prop :=
3   forall msg o,
4     (* for calls to all other entrypoints, *)
5     msg = other o ->
6     (* f is the identity *)
7     option_map (msg_morph C2 C1 f) (other o) = other o.
```

Listing 5.33: f is the identity morphism on all but the `trade` entrypoint.

Third, that f rounds down on the balances kept in storage exposed by `get_bal`, but is the identity on all other aspects of the storage.

```
1 (* 3. f rounds down on the storage, but doesn't touch anything else. *)
2 Context {st_morph : storage -> storage}
3   {state_rounds_down : forall st, get_bal (st_morph st) = round_down (get_bal st)}.
4
5 Definition f_state_morph (f : ContractMorphism C2 C1) : Prop :=
6   (state_morph C2 C1 f) = st_morph.
```

Listing 5.34: f rounds down balances kept in storage.

Fourth, that f is the identity function on error values.

```

1 (* 4. f is the identity on error values *)
2 Definition f_recv_output_err (f : ContractMorphism C2 C1) : Prop :=
3   (error_morph C2 C1 f) = id.

```

Listing 5.35: f is the identity on error values.

Fifth and finally, that f is the identity on setup values.

```

1 (* 5. f is the identity on contract initialization *)
2 Definition f_init_id (f : ContractMorphism C2 C1) : Prop :=
3   (setup_morph C2 C1 f) = id.

```

Listing 5.36: f is the identity on setup values.

The two coherence results of a contract morphism f satisfying these properties show that the upgrade from $C1$ to $C2$ was done as intended: we increased precision successfully without changing how trades were priced. We know this because we have that, after rounding, a trade on $C2$ is the same as the analogous trade on $C1$. We also know that all other functionality remains the same. In particular, this is where the Uranium Finance engineers would have realized the pricing mechanism of the upgraded contract did not conform to that of the old, avoiding the catastrophic mistake and subsequent exploitation.

We show this via the following formalized contract invariant on $C2$, which is proved by left morphism induction. It states that for any reachable state of $C2$, the analogue state of $C1$ given by the contract morphism $f : \text{ContractMorphism } C2 \ C1$ simply rounds the balances down.

```

1 Theorem rounding_down_invariant bstate caddr (trace : ChainTrace empty_state bstate):
2   (* Forall reachable states with contract at caddr, *)
3   env_contracts bstate caddr = Some (C2 : WeakContract) ->
4   (* cstate is the state of the contract AND *)
5   exists (cstate' cstate : storage),
6   contract_state bstate caddr = Some cstate' /\
7   (* cstate is contract-reachable for C1 AND *)
8   cstate_reachable C1 cstate /\
9   (* such that for cstate, the state of C1 in bstate,
10    the balance in cstate is rounded-down from the balance of cstate' *)
11   get_bal cstate = round_down (get_bal cstate').

```

Listing 5.37: All reachable states of $C2$ round down to their corresponding states in $C1$.

In particular, any contract-reachable state of $C2$ has an analogous contract-reachable state of $C1$, so any invariants of $C1$ which are independent of the precision of balances kept in storage and of trades still hold for $C2$. Those which depend on the precision may have an analogous form which holds for $C2$; with the contract morphism in place, this would be the only verification work required to formally verify $C2$ to be correct.

5.4.2 Adding Features and Backwards Compatibility

In a similar spirit to Example 5.4.1, let us consider how we might formally specify backwards compatibility using a contract morphism.

Example 5.4.2 (Backwards Compatibility). Consider contracts C_1 and C_2 , where C_2 is an upgrade of C_1 , and suppose that we wish to show that C_2 is backwards compatible with C_1 .

This can be expressed via a contract morphism in a very precise way: one can not only prove that C_2 is backwards compatible with C_1 , but also indicate exactly *how*—we can indicate which entrypoints and actions in the new contract correspond to which functionality of the old through the component functions of the contract morphism.

We illustrate with an example of a counter contract C_1 which keeps $n : \mathbb{N}$ in storage and has one entrypoint `incr` that increments the natural number in storage by 1. C_1 is upgraded to C_2 , which in addition to an entrypoint to increment the natural number in storage also includes a `decr` entrypoint to decrement the natural number in storage by 1.

```
1 Inductive entrypoint1 := | incr (u : unit).
2 Inductive entrypoint2 := | incr' (u : unit) | decr (u : unit).
```

Listing 5.38: The entrypoint types of C_1 and C_2 , respectively.

We will prove that C_2 is backwards compatible with C_1 by defining a contract morphism

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

with the following component functions.

```
1 Definition msg_morph (e : entrypoint1) : entrypoint2 :=
2   match e with | incr _ => incr' tt end.
3 Definition setup_morph : setup -> setup := id.
4 Definition state_morph : storage -> storage := id.
5 Definition error_morph : error -> error := id.
```

Listing 5.39: The component functions of a morphism defining backwards compatibility.

Note first that f is an embedding since each of its component functions are injections.

```
1 Lemma embedding : is_inj_cm f.
```

Listing 5.40: The resulting morphism f is an embedding of contracts.

In particular, this means that any reachable state of C_1 has an analogous reachable state of C_2 which is fully structure preserving: if we were to only use the functionality of C_2 which it inherits from C_1 , we would get identical contract behavior.

We prove this result using morphism induction as follows. The theorem states that for all reachable chain states with C_1 deployed, there is a corresponding contract-reachable state of C_2 whose state is equal to that of C_1 .

```

1 Theorem injection_invariant bstate caddr (trace : ChainTrace empty_state bstate):
2   (* Forall reachable states with contract C1 at caddr, *)
3   env_contracts bstate caddr = Some (C1 : WeakContract) ->
4   (* forall reachable states of C1 cstate, there's a corresponding reachable state
5     cstate' of C2, related by the injection *)
6   exists (cstate' cstate : storage),
7   contract_state bstate caddr = Some cstate /\
8   (* cstate' is a contract-reachable state of C2 *)
9   cstate_reachable C2 cstate' /\
10  (* .. equal to cstate *)
11  cstate' = cstate.

```

Listing 5.41: C_2 is backwards compatible with C_1 via the contract embedding f .

That the corresponding, contract-reachable state of C_2 is equal to the state of C_1 is precisely what is meant by backwards-compatibility: were we to use only the entrypoints of C_2 inherited from C_1 , we would get identical contract behavior.

Like in Example 5.4.1, which specified meta properties of an upgrade, backwards compatibility can be articulated in a metaspecification for an arbitrary pair of contracts and their specifications, that requires a contract embedding of one contract into another.

5.4.3 Transporting Hoare-Like Properties Over a Morphism

Departing slightly from the previous two examples, we introduce a generic proof technique which uses the coherence proofs of a contract morphism to transport Hoare-like properties over a contract morphism.

Consider the property `unpool_emits_txns` from the structured pools specification of §4.4, which is formalized as follows (see line 219 of Listing A.5 in Appendix A.1.2 for the full context).

```

1 (* When the UNPOOL entrypoint is successfully called, it emits a BURN call to the
2   pool_token, with q in the payload *)
3 Definition unpool_emits_txns (contract : Contract Setup Msg State Error) : Prop :=
4   forall cstate chain ctx msg_payload cstate' acts,
5   (* the call to UNPOOL was successful *)
6   receive contract chain ctx cstate (Some (unpool msg_payload)) = Ok(cstate', acts) ->
7   (* in the acts list there are burn and transfer transactions *)
8   exists burn_data burn_payload transfer_to transfer_data transfer_payload,
9   (* there is a burn call in acts *)
10  let burn_call := (act_call
11    (* calls the pool token address *)
12    (stor_pool_token cstate).(token_address)

```

```

13      (* with amount 0 *)
14      0
15      (* with payload burn_payload *)
16      (serialize (FA2Spec.Retire burn_payload))) in
17  (* with has burn_data in the payload *)
18  In burn_data burn_payload /\
19  (* and burn_data has these properties: *)
20  burn_data.(FA2Spec.retire_amount) = msg_payload.(qty_unpooled) /\
21  (* the burned tokens go from the unpooler *)
22  burn_data.(FA2Spec.retiring_party) = ctx.(ctx_from) /\
23  (* there is a transfer call *)
24  let transfer_call := (act_call
25    (* call to the token address *)
26    (msg_payload.(token_unpooled).(token_address))
27    (* with amount = 0 *)
28    0
29    (* with payload transfer_payload *)
30    (serialize (FA2Spec.Transfer transfer_payload))) in
31  (* with a transfer in it *)
32  In transfer_data transfer_payload /\
33  (* which itself has transfer data *)
34  In transfer_to transfer_data.(FA2Spec.txs) /\
35  (* whose quantity is the quantity pooled *)
36  let r_x := get_rate msg_payload.(token_unpooled) (stor_rates cstate) in
37  transfer_to.(FA2Spec.amount) = msg_payload.(qty_unpooled) / r_x /\
38  (* and these are the only emitted transactions *)
39  (acts = [ burn_call ; transfer_call ] \/
40   acts = [ transfer_call ; burn_call ]).

```

Listing 5.42: The `unpool_emits_txns` property from the formal specification of §4.4, which is a Hoare-like property of partial correctness.

Note that on line 6 of Listing 5.42, we assume that the contract executes without error, assuming

$$\text{receive contract chain ctx state (Some msg) = Ok(cstate', acts),}$$

where `msg` is of the form `(unpool msg_payload)` and `cstate'` is the updated state after the transaction. The remainder of `unpool_emits_txns` is various properties of the list of emitted transactions, `acts`, stating that there has to be a transfer transaction in the list and specifying what the payload of that transaction needs to look like. A proof of `unpool_emits_txns` consists of constructing an outgoing transaction which is emitted by any successful call under those conditions.

More fundamentally, `unpool_emits_txns` is a Hoare-like assertion of partial correctness, reasoning about a successful call to the structured pool contract, where the preconditions are that the state of the chain be reachable with contract address `caddr` and state `cstate`, and the postconditions are statements about `acts` and `cstate'` including that `acts` contain an appropriate `TRANSFER` transaction. Indeed, a reflection on the formal specification of §4.4 reveals that most of the properties of the specification are of this form.

Consider contracts $C1$ and $C2$ and morphism $f : \text{ContractMorphism } C1 \ C2$, and suppose that $C2$ satisfies `unpool_emits_txns`. Suppose further that $C1$ also has an `UNPOOL` entrypoint. If we wish to prove `unpool_emits_txns` for $C1$, we must reason about successful calls to the `UNPOOL` entrypoint. As we will see, we can take advantage of the coherence results which come with f , namely that a successful call to $C1$

```
receive C1 chain ctx state (Some msg) = Ok(cstate', acts)
```

results in a successful call to $C2$

```
receive C2 chain ctx (f.(state_morph) state) (Some (f.(msg_morph) msg)) =
  (Ok (f.(state_morph) cstate', acts)).
```

Using this we can reason about successful calls to the `UNPOOL` entrypoint of $C1$ via the corresponding call under f to the `UNPOOL` entrypoint of $C2$.

Example 5.4.3 (Transporting a Property From the Specification). Consider contracts

```
1 C1 : Contract setup entrypoint storage error
2 C2 : Contract setup entrypoint' storage error
```

and assume `(is_sp : is_structured_pool C2)`, meaning that $C2$ is a structured pool, satisfying the specification of §4.4. Suppose further that the types `entrypoint` and `entrypoint'` are defined as follows:

```
1 (* entrypoint type *)
2 Inductive entrypoint :=
3 | Pool : pool_data -> entrypoint
4 | Unpool : unpool_data -> entrypoint.
5
6 (* entrypoint' type *)
7 Inductive entrypoint' :=
8 | Pool' : pool_data -> entrypoint'
9 | Unpool' : unpool_data -> entrypoint'
10 | Trade' : trade_data -> entrypoint'.
```

Listing 5.43: The entrypoint types of $C1$ and $C2$, respectively.

This gives us an embedding between the entrypoint types of $C1$ and $C2$.

```
1 Definition embed_entrypoint (e : entrypoint) : entrypoint' :=
2 match e with
3 | Pool p => Pool' p
4 | Unpool p => Unpool' p
5 end.
```

Listing 5.44: An embedding of `entrypoint` into `entrypoint'`.

Finally, assume that the functionality of $C1$ regarding its `POOL` and `UNPOOL` entrypoints is identical to that of $C2$ regarding its `POOL'` and `UNPOOL'` entrypoints via the embedding.

```

1 Definition init_coherence_prop : Prop :=
2   forall (c : Chain) (ctx : ContractCallContext) (s : setup),
3     init C c ctx s = init C' c ctx s.
4 Axiom init_coherence_pf : init_coherence_prop.
5
6 Definition recv_coherence_prop : Prop :=
7   forall (c : Chain) (ctx : ContractCallContext) (st : storage) (op_msg : option
8     entrypoint),
9     receive C c ctx st op_msg =
10    receive C' c ctx st (option_map embed_entrypoint op_msg).
11 Axiom recv_coherence_pf : recv_coherence_prop.

```

Listing 5.45: The init and receive functions behave identically under the embedding.

Then we can construct a morphism $f : \text{ContractMorphism } C1 \ C2$ using the coherence proofs and the `embed_entrypoint` function that we’ve just seen:

```

1 (* construct a contract morphism *)
2 Definition f : ContractMorphism C1 C2 := { |
3   setup_morph := id ;
4   msg_morph   := embed_entrypoint ;
5   state_morph := id ;
6   error_morph := id ;
7   (* coherence *)
8   init_coherence := init_coherence_pf ;
9   recv_coherence := recv_coherence_pf ;
10 | }.

```

Listing 5.46: A contract morphism $f : \text{ContractMorphism } C1 \ C2$.

Now suppose that we wish to prove `unpool_emits_txns C` from the structured pools specification.

```

1 Theorem pullback_unpool_emits_txns : unpool_emits_txns C1.

```

We already have a proof of `unpool_emits_txns C2`, and we know that for all messages `msg` into `C1`,

```
receive C1 c ctx st (Some msg) = receive C2 c ctx st (Some (embed_entrypoint msg)).
```

Since in proving `unpool_emits_txns C1` we assume a successful execution,

```
receive C1 c ctx st (Some msg) = Ok(cstate', acts),
```

via f we get a successful execution of `C2`,

```
receive C2 c ctx st (Some embed_entrypoint msg) = Ok(cstate', acts).
```

Our assumed proof of `unpool_emits_txns C2` gives us a proof of the postconditions of `unpool_emits_txns` for `C2`, so we can derive the fact that the relevant postconditions of `unpool_emits_txns` also hold for this call of `C1`. This proves `unpool_emits_txns C1`.

We note that this proof relies on C_1 and C_2 being very similar to each other. The more C_1 and C_2 are similar to each other—for example, in a contract fork or upgrade—the easier it is to use the coherence results of the contract morphism to prove results about C_1 in terms of C_2 . In particular, if a contract is only altered slightly to be upgraded, we could save tedious work of re-proving specified properties or re-specifying by simply using what is already known about the old contract to prove things about the new via a contract morphism.

5.4.4 Summary

Let us reflect on our goal from §5.1.2 of developing formal language to specify contract upgrades. Our main observation was that we need language to formally compare old and new contract versions and their specifications in order to be able to verify that the specification of a new contract version is correct with regards to the intent of the upgrade.

The examples of this section have shown us that contract morphisms provide such a formal language. In Example 5.4.1, we articulated the goal of an upgrade which isolated one piece of the contract functionality, altered it, and left everything else untouched. The intent of the upgrade was expressed through a morphism which asserted that calculations in the upgrade, when rounded down, be equivalent to calculations in the previous version. In Example 5.4.2, we precisely specified our goal that a contract upgrade be backward compatible with a previous version. Again the intent of this upgrade, that we add functionality while preserving the old, was expressed by a contract embedding which identified precisely which pieces of the new contract version corresponded to the old. Finally, Example 5.4.3 showed us that properties of one contract can be used to prove properties of another by way of a contract morphism, which can help us prove facts about an upgrade in terms of its previous versions.

This gives us our desired formal language for individual upgrades, and it is now our task to treat generic upgradeable contracts.

5.5 A Mathematical Characterization of Contract Upgrades

In what follows we will mathematically characterize contract upgradeability using contract morphisms. Before doing so, we go into greater detail on upgradeable contracts.

5.5.1 The Varieties of Upgradeable Contracts

Upgradeable contracts vary, ranging from limited upgradeability to highly flexible frameworks.

Starting from the limited end, take for example MakerDAO, the contract that governs the stablecoin DAI. Certain contract parameters can be updated through a vote by MKR token holders, such as the so-called *stability rate*, which is an interest rate that affects the price of DAI [83]. While MakerDAO cannot radically change the contract features and functionality, we still consider it upgradeable as its functionality can be updated by governance token holders in these limited ways.

Contracts like MakerDAO are an important class of upgradeable contracts called decentralized autonomous organizations (DAOs). These are smart contracts that are governed collectively by holders of a governance token [51]. In some DAOs, governance token holders can also vote to alter, and in some cases fully upgrade, contract functionality. These contracts are governed through a complex web of tokens and incentives which is a fully-fledged research topic in its own right [29].

For example, take Murmuration, a generic DAO template built on Tezos [128]. Murmuration includes a governance token contract and a DAO, the latter of which is governed by those who hold governance tokens. A user submits a proposal which consists of an anonymous lambda function and various metadata. If the proposal passes and executes, which is determined through a voting procedure by governance token holders, the lambda of the proposal becomes the new contract functionality, replacing an entrypoint function. By definition, a lambda has few restrictions, if any, so the upgradeability afforded by Murmuration is substantially greater than that of MakerDAO.

Moving beyond DAOs, the Diamond upgrade framework [95] from §5.1.1 is an extremely flexible system of proxy contracts, which essentially allows for fully altering contract functionality. In addition to changing entrypoint functions like Murmuration, contracts conforming to the Diamond standard can also modify entrypoints to the contract and add to the contract storage.

Each of these examples lend themselves to an intuitive notion of “how upgradeable” the contracts are, and we can say with some confidence that the Diamond framework affords greater upgradeability features than Murmuration, which in turn is more flexible than MakerDAO. But what does this mean, formally?

5.5.2 Isolating Mutable and Immutable Parts

We wish to formally understand contract upgradeability in order to adequately form metaspecifications for upgradeable contracts. To do so, we need language which can characterize its upgradeability properties—by precisely identifying its immutable and mutable parts—as well as language for the governance process and how the upgrade process relates to the actual contract functionality. To that end, we now consider an upgradeable contract in relation to its mutable and immutable parts.

First, the immutable part, which we call the *base contract*. This part of an upgradeable contract is what governs the upgradeability framework and the corresponding incentive and game-theoretic structure of the upgrade process. It indicates what about the contract can change, and through what process it can change, similar to constitutional rules which govern legislative rule-making.

To express this mathematically, we isolate the immutable parts of a contract C into a base contract C_b , and construct a morphism

$$f_p : \text{ContractMorphism } C \ C_b$$

which forgets anything other than the contract structure that governs upgrades and contract versions. This will be a quotient of contracts.

Example 5.5.1 (Quotient Onto the Base Contract). Consider a contract C whose storage contains some natural number $n : \mathbb{N}$ and a function $s : \mathbb{N} \rightarrow \mathbb{N}$, and which has two entrypoints: `next`, which applies s to n , and updates the number in storage with $(s \ n)$; and `upgrade_fun`, which accepts a parameter $s' : \mathbb{N} \rightarrow \mathbb{N}$, and replaces s in storage with s' . This contract is upgradeable in the sense that the functionality of `next`, the primary way in which a user can act on the number in contract storage, can be changed by calling `upgrade_fun`.

```
1 Inductive entrypoint :=
2 | next (u : unit)
3 | upgrade_fun (s' : N -> N) .
4 Record storage := { n : N ; s : N -> N ; }.
```

Listing 5.47: The entrypoint and storage types of our contract C .

Now let us consider what the immutable part of C is. Because the upgradeability refers to the functionality of the `upgrade_fun` entrypoint rather than `next`, the base contract of upgradeable C forgets any natural number in the storage of C and remembers only the function $s : \mathbb{N} \rightarrow \mathbb{N}$, which indicates the version of C . It only has one entrypoint relevant to the structure of C , `upgrade_fun`, which is the upgrade functionality.

To express this formally, we define the base contract C_b , whose entrypoint type, `entrypoint_b`, and storage type, `storage_b`, are defined as follows in Listing 5.48.

```

1 Inductive entrypoint_b' := | upgrade_fun_b (s' : N -> N).
2 Definition entrypoint_b := (entrypoint_b' + unit)%type.
3 Record storage_b := { s_b : N -> N ; }.

```

Listing 5.48: The entrypoint and storage types of the base contract C_b .

The entrypoint type of C_b , entrypoint_b , is defined as a sum type of $\text{entrypoint}_b'$, a type which isolates the upgradeability entrypoint of C , and the unit type. This is so that we can define a function

$$\text{msg_morph} : \text{entrypoint} \rightarrow \text{entrypoint}_b.$$

In particular, we need to be able to send messages of type entrypoint which do not correspond to upgradeability somewhere other than an entrypoint—namely, into the summed unit.

We now define a morphism $f : \text{ContractMorphism } C \ C_b$ with the following component functions.

```

1 Definition msg_morph_p (e : entrypoint) : entrypoint_b :=
2   match e with
3   | next _ => inr tt (* not upgrade functionality *)
4   | upgrade_fun s' => inl (upgrade_fun_b s') (* corresponds to an upgrade *)
5   end.
6 Definition state_morph_p : storage -> storage_b :=
7   (fun (x : storage) => {| s_b := x.(s) ; |}).

```

Listing 5.49: Two component functions of a morphism $f : \text{ContractMorphism } C \ C_b$.

The component function state_morph isolates the function s in storage , which is the data corresponding to a contract’s version; msg_morph sends a message to next to the unit, as it does not correspond to an upgrade, and forwards a message to upgrade_fun . The coherence results come easily by design— C_b is designed to simply be a compressed version of C . We will see later on that C_b does in fact characterize the immutable parts of C .

Let us now move to the mutable part of C , which we call the *version contracts*. This is the contract’s functionality which stands apart from the upgradeability framework, which changes through upgrades. Were it not implemented as an upgradeable contract, it could in principle be implemented as a standalone, non-upgradeable contract. Indeed, any specific version of the upgradeable contract emulates a version contract from within the upgradeability framework.

Example 5.5.2 (Family of Contract Embeddings). We continue with C from Example 5.5.1 which always has some $n : N$ and $s_next : N \rightarrow N$ in storage . Now we wish to define a family of contracts which describes the functionality that is specific to any particular version of the contract.

Opposite to distilling the upgradeability skeleton as we did in Example 5.5.1, we isolate and embed the functionality of a specific version of C . That is, we remove upgrade_fun from the entrypoint type and s from storage , using a (fixed) $s_next : N \rightarrow N$ when the next entrypoint is called.

```

1 Inductive entrypoint_version := | next_f (u : unit).
2 Record storage_version := { n_f : N }.

```

Listing 5.50: The entrypoint and storage types of a version contract $C.f$.

Any contract $C.f'$ with these entrypoint and storage types can be called to update its internal natural number in storage, but it is unable to upgrade in the sense that s_next is fixed. Importantly, were we to initialize C with s_next and never call `upgrade_fun`, then C and $C.f'$ would be identical behaviorally.

Note further that the particular version that $C.f'$ corresponds to is precisely the information in the storage of our base contract $C.b$. In fact, we can define a family of contracts which is parameterized by inhabitants of the storage type of $C.b$, consisting of contracts whose types are defined as follows.

```

1 Definition entrypoint_f : storage_b -> Type := fun v => entrypoint_version.
2 Definition storage_f    : storage_b -> Type := fun v => storage_version.
3 Definition setup_f      : storage_b -> Type := fun v => N.
4 Definition error_f      : storage_b -> Type := fun v => N.
5 Definition result_f     : storage_b -> Type :=
6   fun v => ResultMonad.result ((storage_f v) * list ActionBody) (error_f v).

```

Listing 5.51: A family of contract types parameterized by `storage_b`.

In particular, for a version $v : \text{storage_b}$, the receive function of each version contract uses the function $v.(s_b)$ to execute the next entrypoint. The `init` function is similarly parameterized.

```

1 Definition receive_f (v : storage_b)
2   (_ : Chain)
3   (_ : ContractCallContext)
4   (storage_f : storage_f v)
5   (msg : option (entrypoint_f v))
6   : result_f v :=
7   match msg with
8   | Some (next_f _) =>
9     let st := {| n_f := v.(s_b) storage_f.(n_f) ; |} in
10    Ok (st, [])
11   | None => Err 0
12   end.

```

Listing 5.52: The receive function of $C.f$ version, parameterized by version : `storage_b`

This gives us our family of version contracts as follows.

```

1 Definition C_f (v : storage_b) :
2   Contract (setup_f v) (entrypoint_f v) (storage_f v) (error_f v) :=
3   build_contract (init_f v) (receive_f v).

```

Listing 5.53: A family of version contracts, parameterized by `storage_b`.

We now construct a family of embeddings

```
fi_param (v : storage_b) : ContractMorphism (C_f v)
```

whose component functions are defined as follows.

```
1 Definition msg_morph_i (v : storage_b) (e : entrypoint_f v) : entrypoint :=
2   match e with
3   | next_f _ => next tt
4   end.
5 Definition setup_morph_i (v : storage_b) (st_f : setup_f v) : setup := {|
6   n := st_f ;
7   s := s_b v ; |}.
8 Definition state_morph_i (v : storage_b) (st_f : storage_f v) : storage :=
9   {| n := st_f.(n_f) ; s := s_b v ; |}.
10 Definition error_morph_i (v : storage_b) : error_f v -> error := id.
```

Listing 5.54: The component functions of the family of morphisms `fi_param`.

Intuitively, our family of morphisms shows us that any contract-reachable state of C can be characterized as having a copy of some contract in the family C_f , which stays constant until an upgrade is called.

By studying the structure of C , C_b , and C_f from Examples 5.5.1 and 5.5.2, one might be able to convince one's self that there is some sort of decomposition of C , characterized by C_b and C_f . Indeed, this is the case. In the following section we present this first as an abstract theory, and then prove the decomposition results for our contract C .

5.5.3 Decomposing Upgradeability

Generalizing from Examples 5.5.1 and 5.5.2, consider a contract C ,

```
1 C : Contract Setup Msg State Error.
```

types `Setup_b`, `Msg_b`, `State_b`, and `Error_b` and a contract `C_b`,

```
1 C_b : Contract Setup_b (Msg_b + unit) State_b Error_b.
```

and a family of contracts and types parameterized by `State_b`.

```
1 setup_f : State_b -> Type.
2 msg_f   : State_b -> Type.
3 state_f : State_b -> Type.
4 error_f : State_b -> Type.
5 C_f : forall (v : State_b), Contract (setup_f v) (msg_f v) (state_f v) (error_f v).
```

Listing 5.55: A family of contracts parameterized by `State_b`.

As before, we call C_b the base contract and C_f the family of version contracts. Furthermore, suppose the following conditions hold.

First, that if C receives an empty message, it fails.

```
1 Definition msg_required := forall chain ctx prev_state,
2   exists e, receive C chain ctx prev_state None = Err e.
```

Listing 5.56: Condition 1: C returns an error if called with no message.

Second, that the `init` function gives versioned states according to the following predicate `is_versioned`.

```
1 Definition is_versioned
2   (i_param : forall c_version, ContractMorphism (C_f c_version) C)
3   cstate :=
4   exists c_version cstate_f,
5   cstate = state_morph (C_f c_version) C (i_param c_version) cstate_f.
```

Listing 5.57: The predicate `is_versioned`.

```
1 Definition init_versioned
2   (i_param : forall c_version, ContractMorphism (C_f c_version) C) :=
3   forall init_state chain ctx setup,
4   init C chain ctx setup = Ok init_state ->
5   is_versioned i_param init_state.
```

Listing 5.58: Condition 2: All initial states are versioned.

Third, that messages into C can be categorized as either a message to the base contract portion of C , or to the current version contract portion of C . This is defined by the condition that a message is sent to the unit under p if and only if it has a preimage under i .

```
1 Definition msg_decomposable
2   c_version
3   (i : ContractMorphism (C_f c_version) C)
4   (p : ContractMorphism C C_b) :=
5   forall m,
6   msg_morph C C_b p m = inr tt <->
7   (exists m', m = msg_morph (C_f c_version) C i m').
```

Listing 5.59: Condition 3: A message is sent to the unit under p if and only if it has a preimage under i .

Fourth, that all possible states of C can be categorized by the version that they belong to. This is defined by the condition that a contract state has a preimage under i if and only if the version contract C_f has the version corresponding to the image under p .

```

1 Definition states_categorized
2   c_version
3   (i : ContractMorphism (C_f c_version) C)
4   (p : ContractMorphism C C_b) :=
5   forall st,
6   (exists st_f, st = state_morph (C_f c_version) C i st_f) <->
7   state_morph C C_b p st = c_version.

```

Listing 5.60: Condition 4: A contract state has a preimage under i if and only if the version contract C_f has the version corresponding to the image under p .

Fifth and finally, that there are functions `extract_version` and `new_version_state` which describe how upgrades move between contract versions. The `extract_version` function takes a message to upgrade C (in other words, an incoming message to C_b under p) and extracts the version into which the message will upgrade C . The `new_version_state` function takes the previous version `old_v` and the upgrading message `msg`, and sends an inhabitant of $(\text{state_f } \text{old_v})$, the storage type of the previous version contract, to an inhabitant of $(\text{state_f } (\text{extract_version } \text{msg}))$, the storage type of the next version contract. As we will see, in the contracts from Examples 5.5.1 and 5.5.2, `extract_version` simply uses the payload of `upgrade_fun`, a function, to get the version of the new contract, and since the storage type of all version contracts is the same, the `new_version_state` function is simply the identity function.

```

1 Definition version_transition
2   old_v
3   (i_param : forall c_version, ContractMorphism (C_f c_version) C)
4   (p : ContractMorphism C C_b)
5   (extract_version : Msg_b -> State_b)
6   (new_version_state : forall old_v msg,
7     state_f old_v -> state_f (extract_version msg)) :=
8   forall cstate cstate_f,
9   (* forall states of version old_v *)
10  cstate = state_morph (C_f old_v) C (i_param old_v) cstate_f ->
11  (* and forall successful calls ... *)
12  forall chain ctx msg new_state new_acts msg',
13  receive C chain ctx cstate (Some msg) = Ok (new_state, new_acts) ->
14  (* to upgrade the contract C ... *)
15  msg_morph C C_b p msg = inl msg' ->
16  (* then the new state is the state given by new_version_state *)
17  let new_v := extract_version msg' in
18  new_state =
19    state_morph (C_f new_v) C (i_param new_v) (new_version_state old_v msg' cstate_f).

```

Listing 5.61: Condition 5: The functions `extract_version` and `new_version_state` which describe how upgrades move between contract versions.

Then given contracts C and C_b , a family of contracts C_f , a family of embeddings i_{param} of contracts from C_f into C , and a quotient p of C onto C_b , C has an upgradeability decomposition into C_b and C_f

if there are functions `extract_version` and `new_version_state` such that the above-listed conditions hold.

```

1 Definition upgradeability_decomposition
2   (i_param : forall c_version, ContractMorphism (C_f c_version) C)
3   (p : ContractMorphism C C_b)
4   (extract_version : Msg_b -> State_b)
5   (new_version_state : forall old_v msg,
6     state_f old_v -> state_f (extract_version msg)) :=
7   (* Forall versions of a contract C, *)
8   msg_required /\
9   init_versioned i_param /\
10  forall c_version,
11  let i := i_param c_version in
12  msg_decomposable c_version i p /\
13  states_categorized c_version i p /\
14  version_transition c_version i_param p extract_version new_version_state.

```

Listing 5.62: The definition of an upgradeability decomposition, which is a conjunction of all five conditions listed above.

A proof of `upgradeability_decomposition C` gives us two results which characterize the contract trace of `C` in terms of its base and version contracts.

First, that all reachable states have a version, defined by the `is_versioned` predicate from Listing 5.58.

```

1 Theorem versioned_invariant
2   (* Consider family of embeddings, and *)
3   (i_param : forall c_version, ContractMorphism (C_f c_version) C)
4   (* a projection onto the skeleton C_b. *)
5   (p : ContractMorphism C C_b)
6   (extract_version : Msg_b -> State_b)
7   (new_version_state : forall old_v msg,
8     state_f old_v -> state_f (extract_version msg)) :
9   (* Then forall reachable states ... *)
10  forall bstate caddr (trace : ChainTrace empty_state bstate),
11  (* where C is at caddr with state cstate, *)
12  env_contracts bstate caddr = Some (C : WeakContract) ->
13  exists (cstate : State),
14  contract_state bstate caddr = Some cstate /\
15  (* if the contract's upgradeability can be decomposed *)
16  (upgradeability_decomposition i_param p extract_version new_version_state ->
17  (* then every contract state cstate is versioned *)
18  is_versioned i_param cstate).

```

Listing 5.63: All reachable states of an upgradeable contract are versioned.

Second, that if we have a contract `C` with upgradeability decomposition `upgradeability_decomposition`, then all incoming calls are either upgrades—a call to contract upgradeability, resulting in a new version—or

correspond are calls to the current version contract, the version staying constant. Transitions from one version to another result in a new state that is the image of the state of a new version contract.

```

1 Theorem upgradeability_decomposed
2   (* Consider family of embeddings, and *)
3   (i_param : forall c_version, ContractMorphism (C_f c_version) C)
4   (* a projection onto the base contract C_b. *)
5   (p : ContractMorphism C C_b)
6   (extract_version : Msg_b -> State_b)
7   (new_version_state : forall old_v msg,
8     state_f old_v -> state_f (extract_version msg)) :
9   (* forall contract states and corresponding contract versions, *)
10  forall cstate c_version cstate_f,
11  (* with i, the embedding for version c_version, ... *)
12  let i := i_param c_version in
13  (* if C_f -> C ->> C_b is the decomposition of a contract's upgradeability ... *)
14  upgradeability_decomposition i_param p extract_version new_version_state ->
15  (* and cstate is in the image of cstate_f under the embedding i
16     (meaning that cstate has version c_version) ... *)
17  cstate = state_morph (C_f c_version) C i cstate_f ->
18  (* Then forall calls to the versioned contract *)
19  forall chain ctx m new_state new_acts,
20  receive C chain ctx cstate (Some m) = Ok (new_state, new_acts) ->
21  (* it either stays within this version *)
22  (exists cstate_f', new_state = state_morph (C_f c_version) C i cstate_f') \/
23  (* it moves onto a new version *)
24  (exists c_version' cstate_f',
25  new_state = state_morph (C_f c_version') C (i_param c_version') cstate_f').

```

Listing 5.64: All contract calls to an upgradeable contract are either upgrades (to a new version) or stay in the same version; transitions between versions behave as expected.

These two results show that all contract states are versioned by C_f and C_b , and then give precise semantics of how the contract moves between version contracts in C_f by calling the base contract C_b .

Let us reflect on what these mean for our goal of formally characterizing contract upgradeability. We've shown that our decomposition isolates the functionality relating to contract upgrades from the functionality relating to each particular contract version, and that it accurately describes all contract states and transitions between them. The definition of the family of version contracts is a precise definition of the bounds of a contract's upgradeability. And by isolating the governance features of contract upgrades into the base contract, we are able to rigorously reason about any incentive, economic, or game-theoretic aspects of the contract's upgradeability.

Of course, upgradeable contracts are not typically specified in such a modular way, but as with the Diamond framework [95], upgradeable contracts are typically specified as one, unified contract. We can thus formulate the metaspecification of an upgradeable contract in terms of its base contract, with its

associated governance features, and its version contracts, defining the structure common to every possible contract version. The specification of an upgradeable contract is correct if it specifies a contract which can be decomposed into the base and version contracts of the metaspecification.

Example 5.5.3 (Decomposing C). We conclude this section by continuing with Examples 5.5.1 and 5.5.2, showing that C can be decomposed into the base contract C_b and the family of version contracts C_f . To do so, we need functions `extract_version` and `new_version_state` which describe how upgrades move between version contracts.

Since versions of C are given by the function s in storage, the `extract_version` function simply takes the function from the message into the `upgrade_fun` entrypoint.

```
1 Definition extract_version (m : entrypoint_b') : storage_b :=
2   match m with | upgrade_fun_b s_b => {| s_b := s_b |} end.
```

Listing 5.65: The `extract_version` function, which isolates the contract version resulting from a contract call by taking the new function from the upgrade message payload.

The `new_version_state` function takes the state of a previous version contract to the state of a new version contract. Since the storage types of all version contracts are the same for this particular family, the function is constant.

```
1 Definition new_version_state old_v msg (st : storage_f old_v) :
2   storage_f (extract_version msg) := st.
```

Listing 5.66: The `new_version_state` function which takes the state of a previous version contract to the state of a new version contract.

We have the following theorem, which says that C is decomposable with regards to C_b and the quotient f_p , and the family of version contracts C_f and embeddings fi_param .

```
1 Theorem upgradeability_decomposition :
2   upgradeability_decomposition fi_param f_p extract_version new_version_state.
```

Listing 5.67: C is decomposable with regards to C_b and the family of version contracts C_f .

What does this tell us about C ? We know that all reachable states of C are versioned, *i.e.* have an embedding of a contract in the family C_f which indicates its version. We also have precise semantics of how the upgrades move between version contracts, via the function `new_version_state`, and we can characterize all incoming messages as either upgrades, or calls to a particular version contract.

5.5.4 Upgradeable Contracts are Fiber Bundles: A Digression

The decomposition of an upgradeable contract resembles a topological phenomenon called a *fiber bundle*. Just as an upgradeable contract knits together its base contract with a family of version contracts, a fiber bundle knits a topological space in with a family of topological spaces into one so-called *total space* [75].

Fiber bundles are hugely influential within mathematics, particularly in geometry and topology, and have an extremely rich theory surrounding them. They are a tool used to understand topological spaces, which are famously complex and mysterious mathematical objects. We point out an analogous structure between topological spaces and smart contracts because smart contracts (and software more broadly) are also complex mathematical objects, and understanding them is precisely the goal of formal verification.

For what follows, we assume basic knowledge about topological spaces and continuous maps; for an introduction to topology, see *e.g.* Munkres' introductory course [96].

First, recall that for a map of topological spaces $f : X \rightarrow Y$, the *fiber* of f over $y \in Y$ is the preimage $f^{-1}(y)$, where $f^{-1}(y) := \{x \in X \mid f(x) = y\}$.

Definition 5.5.1 (Fiber bundle). Consider topological spaces E and B . A continuous surjection $\pi : E \rightarrow B$ is a fiber bundle if the following conditions hold:

1. For all $b \in B$, $\pi^{-1}(b)$ is homeomorphic to a fixed topological space F .
2. There is an open cover $\{U_\alpha\}_{\alpha \in I}$ of B with isomorphisms

$$\phi_\alpha : \pi^{-1}(U_\alpha) \xrightarrow{\sim} U_\alpha \times F$$

which restricts on fibers to a homeomorphism.

3. For $\alpha, \beta \in I$, the composition $\phi_\beta^{-1} \circ \phi_\alpha : (U_\alpha \cap U_\beta) \times F \rightarrow (U_\alpha \cap U_\beta) \times F$ is well-defined and satisfies

$$\phi_\beta^{-1} \circ \phi_\alpha(x, v) = (x, g_{\alpha\beta}(x)v)$$

for some $\text{Aut}(F)$ -valued function $g_{\alpha\beta} : U \cup V \rightarrow \text{Aut}(F)$.

4. These maps satisfy

$$g_{\alpha\alpha} = \text{Id} \quad \text{and} \quad g_{\alpha\beta}(x)g_{\beta\gamma}(x)g_{\gamma\alpha}(x) = \text{Id}.$$

The functions $\phi_\beta^{-1} \circ \phi_\alpha$ are called transition functions, E is called the total space, B is called the base space, and F is called the fiber. Diagrammatically, a fiber bundle is often drawn as follows.

$$\begin{array}{ccc} F & \hookrightarrow & E \\ & & \downarrow \\ & & B \end{array}$$

One can think of fiber bundles intuitively as a quotient, similar to a group quotient, where F takes on the analogous role played by a normal subgroup in a group quotient.

Fiber bundles form a category, where maps of fiber bundles, which we simply call *bundle maps*, are commuting squares

$$\begin{array}{ccc} E & \longrightarrow & E' \\ \downarrow & & \downarrow \\ B & \longrightarrow & B', \end{array} \quad (5.1)$$

where $E \rightarrow B$ and $E' \rightarrow B'$ are fiber bundles and all the maps are continuous. If we fix a base space B , we can define a category of fiber bundles over B by defining morphisms to be commuting squares such as (5.1) with the condition that the map on the bottom row be the identity map.

One important feature of fiber bundles is that they have the homotopy lifting property for all CW-complexes. The function $\pi : E \rightarrow B$ has the homotopy lifting property with respect to a space X if, for all homotopies

$$h : [0, 1] \times X \rightarrow B,$$

if there exists a map $f_0 : \{0\} \times X \rightarrow E$ such that $\pi \circ f_0 = h|_{\{0\} \times X}$, then there exists a homotopy $f : [0, 1] \times X \rightarrow E$ such that $\pi \circ f = h$ and $f|_{\{0\} \times X} = f_0$. A *fibration* is a surjection $\pi : E \rightarrow B$ which satisfies the homotopy lifting property with respect to any topological space.

If the lift is unique, one important corollary is that for points $b, b' \in B$, a path from b to b' induces a function from the fiber of b to the fiber of b' : given $x \in \pi^{-1}(b)$, use the homotopy lifting property to lift the path; the image of our function is the endpoint of that path in $\pi^{-1}(b')$.

Example 5.5.4 (Trivial Bundle). For all topological spaces B and F , $E := B \times F$ is a fiber bundle over B given by the obvious inclusion and projection functions.

Example 5.5.5 (Covering spaces). Let X be a topological space. A covering of X is a fiber bundle, with X as the base space, such that the fibers of X are discrete topological spaces. Paths in a covering space lift uniquely to paths in the covering space, yielding a function on fibers.

Example 5.5.6 (Möbius Bundle). Consider a Möbius band M defined by taking the space $[0, 1] \times [-1, 1]$ and identifying $(0, x)$ with $(1, -x)$ for all $x \in [-1, 1]$. Similarly, construct a copy of S^1 by taking $[0, 1]$ and identifying 0 with 1. We can define a function $\pi : M \rightarrow S^1$ given by $(x, y) \mapsto x$. Then π is a fiber bundle with fibers $[-1, 1]$.

Some fiber bundles admit a *section*, which is a function $s : B \rightarrow E$ such that $\pi \circ s = \text{id}$. Sections are useful tools in analysis, topology, and differential geometry for a variety of reasons. Topologically, one way to conceptualize the existence of a section is as an indicator of how “twisted” (or rather, untwisted) a fiber bundle is. The trivial bundle, for example, has a section for every $x \in F$ given by $b \mapsto (b, x)$, and is considered to be the “least twisted” type of fiber bundle.

Another example is the Möbius bundle M from Example 5.5.6. M has a section $s : S^1 \rightarrow M$ given by $x \mapsto (x, 0)$. However if we remove $(x, 0)$ from every fiber, this gives us a fiber bundle M' with disconnected fibers $[-1, 0) \cup (0, 1]$, which now admits no sections. This is because the Möbius band has a twist: if we try to construct a section of S^1 into M' , *e.g.* by drawing a loop in M' , no matter what we choose, after one revolution around M' we've swapped sides in the fiber.

The analogy between fiber bundles and upgradeable contracts is as follows. We consider contracts as analogous to topological spaces, where the points of the space are inhabitants of a contract's storage. A successful contract call alters the storage of the contract, and is thus a continuous path—which indicates movement—within the contract.

Our decomposition of a contract into its base and version contracts

$$\begin{array}{ccc} \mathsf{C_fv} & \hookrightarrow & \mathsf{C} \\ & & \downarrow \\ & & \mathsf{C_b} \end{array}$$

follows the analogy of a fiber bundle, where the version contracts are the fibers and the base contract is the base space of the fiber bundle. That $\mathsf{C_f}$ parameterizes all the fibers over the storage type of $\mathsf{C_b}$ is the analogue of a continuous parameterization of the fibers by the base space in E .

Some of the conditions of §5.5.3 also have an interpretation in the context of fiber bundles. Condition 2, that all initial states are versioned, indicates that every initialization corresponds to the point in a fiber of p . Condition 3, that messages are decomposable, means that a path in C is either a path that stays within a fiber, meaning it is a call to a version contract $\mathsf{C_f_v}$, or a path moves *between* fibers, meaning it is a call to the base contract $\mathsf{C_b}$. Condition 4 is that $\mathsf{C_f_v}$ is precisely the fiber over v under p .

Condition 5 gives us the analogy with the path lifting property: a call to upgrade a contract, which is a path within the base contract $\mathsf{C_b}$, lifts to a path in C which starts in the fiber $\mathsf{C_f_prev_v}$ of the previous version $\mathsf{prev_v}$ and ends in the new fiber $\mathsf{C_f_new_v}$ of the new contract version $\mathsf{new_v}$. The path of an upgrade in $\mathsf{C_b}$ is given by the `extract_version` function, which isolates the new version from an incoming upgrade call. Then `new_version_state` is our function between fibers, taking an inhabitant of the state of the previous fiber contract to an inhabitant of the state of the new fiber contract.

Considering this analogy, we might ask if $\mathsf{f_p} : \text{ContractMorphism } \mathsf{C} \ \mathsf{C_b}$ admits a section. Recall that in defining $\mathsf{C_b}$, we had to sum the `entrypoint_type` `entrypoint_b'` with the unit in order to be able to define the morphism from C to $\mathsf{C_b}$. This construction, which we call the `pointed_contract` construction, is a general construction on contracts: given a contract

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

we have a so-called *pointed contract*

`pointed_contract C : Contract Setup (Msg + unit) State Error`

which admits an extra entrypoint which succeeds when called, but does nothing to the state and emits no actions. In particular, $C.b$ from our example is `pointed_contract C.b'` for a contract $C.b'$.

There is also a canonical embedding of C into `pointed_contract C` for all C ,

```
pointed_include : ContractMorphism C (pointed_contract C),
```

whose component functions are the identity apart from on the message type, where the component function is `(fun m => inl m)`.

We will then define a section of our fiber bundle analogue to be a morphism

```
fp_rinv : ContractMorphism C.b' C
```

such that

```
compose_cm f_p (fp_rinv n) = pointed_include C.b'.
```

As it turns out, we have such a section for the upgradeable contract of Example 5.5.3, defined by the following component morphisms.

```
1 Definition setup_morph_s n : setup_b -> setup :=
2   (fun (x : setup_b) => {| n := n ; s := x.(s_b) |}).
3 Definition msg_morph_s (e : entrypoint_b') : entrypoint :=
4   match e with | upgrade_fun_b s' => upgrade_fun s' end.
5 Definition state_morph_s n : storage_b -> storage :=
6   (fun (x : storage_b) => {| n := n ; s := x.(s_b) ; |}).
7 Definition error_morph_s : error_b -> error := (fun (x : error_b) => x).
```

Listing 5.68: A right inverse of p from Example 5.5.3

Note in particular that `setup_morph_s` and `state_morph_s` are parameterized by a natural number n . This is precisely what parameterizes p' : every inhabitant of the storage of $C.b$ can be canonically lifted to an inhabitant of its fiber by simply inserting n as the natural number in storage. Thus we have a family of sections parameterized by $n : \mathbb{N}$.

```
1 Definition fp_rinv (n : N) : ContractMorphism C.b' C.
```

And we can prove that for all n , `fp_rinv n` is a section of our fiber bundle.

```
1 Theorem p_rinv_section (n : N) : compose_cm f_p (fp_rinv n) = pointed_include C.b'.
```

This family of sections is reminiscent of the sections of the trivial bundle we saw before, which were given by $b \mapsto (b, x)$ for $x \in F$. Because our fibers here are not isomorphic contracts, we do not have a decomposition of C as a product of $C.b$ and some fixed $C.f'$. Even so, the family of sections indicates that C is a simple analogue of a product of contracts, which seems appropriate on reflection as C clearly separates the upgrade functionality from that corresponding to each contract version.

5.5.5 Summary

Let us reflect on the motivating example for this section, the Diamond upgrade framework of §5.1.1. Our main observation was that the specification uses diagrams and analogies to a diamond in order to communicate the different components of the upgradeable contract’s functionality. The parts of the specification which deal with the process of adding or removing a facet are a description of the Diamond standard’s base contract—the functionality governing how the contract can be changed. The parts of the specification which describe the structure the contract as a diamond, having many facets, describes the version contracts—the structure common to all possible versions of a Diamond contract.

Of course, producing an actual decomposition of the Diamond standard would likely be highly nontrivial. Instead, we would take the approach of writing the specification of an upgradeable contract explicitly in terms of its base and version contracts. Any implementation conscious of the required decomposition can then be crafted in such a way that the decomposition, using contract morphisms, yields itself easily.

5.6 Conclusion

Our goal in this chapter was to develop formal tools to evaluate the correctness of a specification of an individual contract upgrade or an upgradeable contract. This is because contract upgrades have nontrivial meta properties which can make correct specification of individual upgrades, as well as upgradeable contracts, a nontrivial endeavor. We argued that correct specification of contract upgradeability has at least two components, both of which we addressed with contract morphisms.

The first component is that in an upgrade, whether by a hard fork or within an upgradeable contract, the specification of the upgraded contract is typically written with the intention to relate to the specification of a previous version in some way. One can thus evaluate the correctness of the specification of an upgrade if one is able to formally express that intended relationship.

To address this component, we illustrated how contract morphisms can be used to clearly specify the intention of upgrades, by specifying properties of upgrades such as backwards compatibility. In particular, we revisited a major attack on a faulty contract upgrade, giving a simple metaspecification of the intended upgrade which clearly exposes the vulnerability. We also showed how one can prove Hoare-like properties of one contract using those of another.

The second component is that upgradeable contracts are difficult to specify with a prose specification because an upgradeable contract encodes the logic of two separate contracts—one that governs upgrades, and one that represents the contract’s version at any given state. The relationship between these two contracts is complex, so to evaluate the correctness of a specification of an upgradeable contract we need some formal way to decompose it into its mutable and immutable parts.

Using contract morphisms we gave a formal and general decomposition of an upgradeable contract into a base contract and a family of version contracts. This rigorously and precisely articulates what it means for the contract to be upgradeable, including what the semantics of upgrades are, and it completely characterizes the forms that contract upgrades can take. Furthermore, there is a strong mathematical analogue to fiber bundles, a tool used in topology and geometry to understand the structure of a topological space in terms of a decomposition into a *base space* and a family of *fibers*.

By addressing both of these components with contract morphisms, we have a rigorous, mathematical treatment of upgradeability. This can help address vulnerabilities in contract upgrades by evaluating the correctness of specifications of individual contract upgrades, as well as upgradeable contracts generally.

Chapter 6

Contract Systems

The meta properties that we target in this chapter are those relating to the behavior of a system of contracts, taken as a whole. As we saw in §1.2.3, poorly specified systems of contracts can be vulnerable to extremely costly attacks. In this chapter we rigorously develop the notion of a specification’s correctness as it relates to how the system of contracts behaves when taken as a whole.

To this end, we develop various notions of equivalences of contracts so that we can understand the specification of a system of contracts in terms of a monolithic contract, bisimilar to the system. This will allow us to distinguish the specification of a system’s infrastructure—a description of when and with what payload contracts in the system call each other—from the specification of its core, intended functionality when considering the contract system as a whole.

We proceed as follows. In §6.1, we motivate contract and system bisimulations by discussing the problem of specifying contract systems. In §6.2, we introduce the notion of a contract bisimulation as a precursor to a bisimulation of contract systems, and show that isomorphic contracts are also bisimilar. In §6.3, we formally define the notion of contract systems, using Milner’s *bigraphs* [94] as our data type, to isolate the specification of the system infrastructure from its core functionality. In §6.4, we introduce bisimulations of contract systems. In §6.5, we conclude. Each section contains various code snippets; Appendix C mirrors section headings and gives a more complete version of the code from which the snippets are taken.

6.1 Contracts Systems

Financial smart contracts are ubiquitously implemented as modular systems of contracts rather than as monolithic contracts. One reason for this is purely practical. A blockchain is a highly resource-constrained platform on which to write programs due to gas fees, which are paid by the user [148]. Depending on the contract, it can be more gas-efficient to implement modularly, for example to prevent the storage of

a contract from getting too large and thus expensive to access. In this case, implementing a contract modularly is an optimization technique.

Other reasons are more endemic. Systems of contracts are ubiquitous in financial smart contracts because new applications frequently build off of existing ones [130]. Indeed, one of the strengths of smart contracts is that they are *composable*—so much so that decentralized finance is colloquially referred to as “money legos” [130], which refers to building new DeFi protocols by composing existing ones as if with lego bricks.

Recall for example that the implementation of Dexter2—indeed, that of any AMM—features a main, trading contract, and a secondary LP token contract [102]. The functionality of the main contract depends on that of the token contract, and conversely the token contract depends on the main trading contract. Furthermore, Dexter2 facilitates trading between a diversity of tokens, each of which has its own tokenomics (rules governing minting, burning, and distribution) and are themselves contracts independent from the AMM. One previously mentioned example of the ensuing complexity is that *governance tokens*—tokens that give the holder governance rights over an upgradeable contract or a DAO—can be traded on AMMs. In the exploit mentioned in §1.2.1, the Beanstalk attacker first traded for governance tokens on an AMM before executing flash loan attack (via yet another contract), winning a majority governance vote and draining the Beanstalk contract of its funds.

The examples continue. Yield aggregators like Harvest or Yearn [41] optimize over a set of existing DeFi protocols so that users can maximize yield farming. DEX aggregators like 1inch or Matcha minimize trading fees over various DEXs. Synthetics, including stablecoins like DAI, rely on price oracles to manage their tokenomics and maintain their peg [129]. CREAM, the multi-purpose DeFi protocol of §1.2.3, relies intimately on a whole web of contracts across multiple chains to offer its financial services.

In each of these examples, contract security and even correctness depends intimately on the accumulated and interconnected behavior of many interacting contracts [130]. Because of the overwhelming prevalence of contract systems, the very endeavor of smart contract verification can be rendered inept if we don’t have robust language to specify and reason about contract systems. Famously, however, contract composability is a source of extraordinary complexity in formal reasoning [103, 130].

6.1.1 Specifying Contract Systems

We focus on the complexity introduced to formal specification by contract composability. Consider what is needed to specify of a system of contracts. As well as specifying each individual contract in the system, we must inevitably specify when and with what payload contracts in the system call each other. From these specifications one must deduce, either by intuition or by other means, the intended behavior of the system of contracts when taken as a whole, or considering it as a single process.

We argue that specifying a system of contracts in this way is difficult to do correctly because the details of inter-contract communication can obfuscate the actual intended behavior of the system specification.

Recall the specification of the Diamond standard [95] from Chapter 5. The core, intended functionality is illustrated by pictures and naming conventions (*e.g.* facets, diamonds, loupes), precisely because the specification alone has many moving parts and is difficult to understand. If it were possible to somehow separate the specification of system infrastructure, including inter-contract communication, from that of the contract’s core, desired functionality, the specification would likely be much easier to understand.

There is more at work here. One can see conceptually that the intended behavior of a contract system is generally articulated in terms of one, coherent process, which is agnostic to whether and how the system is modularized, or whether it is monolithic. Indeed one can imagine various implementations of a given specification, where some are modular and some are monolithic, but one has an intuitive notion that as long as the contracts of a given implementation behave together, coherently, according to the specification, that implementation is correct.

We make this rigorous by introducing formal tools to reduce the modular system of contracts to a monolithic contract, by way of a process algebra, such that the modular and monolithic contracts are bisimilar. We can then reason about our system in terms of that single contract. By using Milner’s bigraphs [94] as our data type for contract systems, we can isolate the system infrastructure, allowing us to separate the specification of the interacting system from the specification of the core, desired contract behavior. From there, we have tools already at our disposal from previous chapters to rigorously specify and reason about the desired core contract behavior.

There are other reasons for which we may wish to consider contract bisimilarity. Consider the problem of specifying and verifying a contract which is highly optimized for performance. Performant code is frequently difficult to read, much less reason about—and conversely, contracts optimized for formal reasoning are seldom optimized for performance. Thus a contract on which it is feasible to articulate and verify meta properties might be too inefficient (and thus expensive) to feasibly deploy. In this scenario, we may wish to verify a contract optimized for formal reasoning, and then perform bisimulation-preserving optimizations on the verified contract before deploying it.

6.1.2 Related Work

While applying process-algebraic techniques to formal reasoning about smart contracts is still a relatively new practice, there is related work that reasons about systems of contracts and that draws in various ways on process-algebraic techniques.

One example is Tolmach *et al.*’s work on the formal analysis of composable DeFi protocols [130]. Tolmach *et al.* formulate a process-algebraic model of DEXs and tokens, abstracting each as primitives. These are used to analyze the behavior of contract systems by formally abstracting contract interactions and then using a model-checker to verify correctness properties. All of the correctness properties are properties of one or more components of the system.

Another example is 20squares, a smart contract auditing firm that uses compositional game theory [64] to analyze the economic properties of the specification of a composable set of Ethereum-based smart contracts. They are able to ask certain game-theoretic questions about contract specifications, such as how much certain agents have to be bribed to deviate from intended behavior [1]. In particular, they address system complexity by composing games which represent individual contracts, and studying the resulting incentives.

Finally, Madl *et al.* [86] reason about contract systems using interface automata, which is a formal way of specifying interactions and synchronization of input and output actions within a system.

Each of these examples takes a specification of a modular system and uses tools to reason in a composable manner over that system. Our contrasting approach is to separate the specification of the system infrastructure from that of the core functionality, which removes the complexity of the contract system before we formally reason about it. We do so by formally developing the notion of contract and system bisimulations, allowing us to consider contract specifications that are agnostic to whether or not, or how, the implementation is modularized.

6.2 Bisimulations of Contracts

The essential goal of this chapter is to formally define bisimulations of contract systems, but as a precursor our first task is to formally introduce bisimulations of individual contracts. As we will see, contract bisimulations extend naturally into bisimulations of contract systems. They are also a generalization of contract isomorphisms, which we saw briefly in Chapter 5.

6.2.1 Contract Trace Morphism

Recall from §5.3.1 that a contract’s trace is a chained list of contract steps, where contract steps are the data for a successful call to the `receive` function. We restate the formal definitions of each here.

```

1 Record ContractStep (C : Contract Setup Msg State Error)
2   (prev_cstate : State) (next_cstate : State) :=
3   build_contract_step {
4     seq_chain : Chain ;
5     seq_ctx : ContractCallContext ;
6     seq_msg : option Msg ;
7     seq_new_acts : list ActionBody ;
8     (* we can call receive successfully *)
9     recv_some_step :
10       receive C seq_chain seq_ctx prev_cstate seq_msg = Ok (next_cstate, seq_new_acts) ;
11   }.

```

Listing 6.1: Contract steps are successful calls to the `receive` function.

```

1 Definition ContractTrace (C : Contract Setup Msg State Error) :=
2   ChainedList State (ContractStep C).

```

Listing 6.2: A contract’s trace is a chained list of contract states, linked together by contract steps.

A *contract trace morphism*, similar to a contract morphism, encodes a formal, structural relationship between the traces of two contracts. Like contract morphisms, a contract trace morphism features a function between contract storage types. However, unlike contract morphisms, which also features functions between message, error, and setup types, a contract trace morphism simply requires that the function between state types send initial states to initial states, and that for any two states connected by a step of the first contract, the corresponding states be also connected by a step in the second contract.

More concretely, for contracts

`C1:Contract Setup1 Msg1 State1 Error1` and `C2:Contract Setup2 Msg2 State2 Error2`,

a morphism of contract traces includes the following data:

- A function `ct_state_morph : State1 -> State2`
- A proof that if there are inputs to the `init` function of `C1` that result in an initialized state `init_state`, then there are inputs to the `init` function of `C2` that result in a corresponding initialized state `(ct_state_morph init_state)`
- For states `state1` and `state2`, and any step forward of `C1`,

`step1 : ContractStep C1 state1 state2,`

we have a corresponding step forward of `C2` between the analogous states

`step2 : ContractStep C2 (ct_state_morph state1) (ct_state_morph state2)`

```

1 Record ContractTraceMorphism
2   (C1 : Contract Setup1 Msg1 State1 Error1)
3   (C2 : Contract Setup2 Msg2 State2 Error2) :=
4   build_ct_morph {
5     (* a function of state types *)
6     ct_state_morph : State1 -> State2 ;
7     (* init state C1 -> init state C2 *)
8     genesis_fixpoint : forall init_cstate,
9       is_genesis_cstate C1 init_cstate ->
10      is_genesis_cstate C2 (ct_state_morph init_cstate) ;
11     (* coherence *)
12     cstep_morph : forall state1 state2,
13       ContractStep C1 state1 state2 ->
14       ContractStep C2 (ct_state_morph state1) (ct_state_morph state2) ;
15   }.

```

Listing 6.3: The formal definition of a morphism of contract traces.

Inductively, this gives us a relationship between reachable states: initial states of each contract are related via the function between state types, and from there, for any step of the first contract there is a related step of the second contract that respects the function on states.

Example 6.2.1 (The Identity Contract Trace Morphism). One important contract trace morphism is the identity morphism `id_ctm`, which is defined for every contract `C` and inhabits the type

$$\text{ContractTraceMorphism } C \ C.$$

This is analogous to the identity morphism of Example 5.2.2. We give the formal definition here.

```
1 Definition id_ctm (C : Contract Setup1 Msg1 State1 Error1) :
2   ContractTraceMorphism C C := {|
3   ct_state_morph := id ;
4   genesis_fixpoint := id_genesis_fixpoint C ;
5   cstep_morph := id_cstep_morph C ;
6 |}.
```

Listing 6.4: The identity contract trace morphism `id_ctm C` defined for any contract `C`.

The definition relies on a lemma, `id_genesis_fixpoint C`, and a function, `id_cstep_morph C`, which are both defined trivially.

```
1 Definition id_genesis_fixpoint (C : Contract Setup1 Msg1 State1 Error1)
2   init_cstate (gen_C : is_genesis_cstate C init_cstate) :
3   is_genesis_cstate C (id init_cstate) :=
4   gen_C.
```

Listing 6.5: The genesis fixpoint result for the identity contract trace morphism.

```
1 Definition id_cstep_morph (C : Contract Setup1 Msg1 State1 Error1)
2   state1 state2 (step : ContractStep C state1 state2) :
3   ContractStep C (id state1) (id state2) :=
4   step.
```

Listing 6.6: The step result for the identity contract trace morphism.

Example 6.2.2 (Equality of Contract Trace Morphisms). Given two contract trace morphisms

$$f \ g : \text{ContractTraceMorphism } C_1 \ C_2,$$

like with contract morphisms we may ask ourselves whether or not they are equal.

Because of the dependent nature of the definition of contract trace morphisms, this definition is not quite as straightforward as equality of contract morphisms (Example 5.2.4). We must first assume equality of each state morphisms, and then we can state the (weaker) equality result, which states that we have equality of contract trace morphisms if the two functions between contract steps are equal.

```

1 Lemma eq_ctm_dep
2   (C1 : Contract Setup1 Msg1 State1 Error1)
3   (C2 : Contract Setup2 Msg2 State2 Error2)
4   (* one single state morphism *)
5   (ct_st_m : State1 -> State2)
6   (gen_fix1 gen_fix2 : forall init_cstate,
7     is_genesis_cstate C1 init_cstate ->
8     is_genesis_cstate C2 (ct_st_m init_cstate))
9   (cstep_m1 cstep_m2 : forall state1 state2,
10    ContractStep C1 state1 state2 ->
11    ContractStep C2 (ct_st_m state1) (ct_st_m state2)) :
12   (* if the step morphisms are equal ... *)
13   cstep_m1 = cstep_m2 ->
14   (* then the contract trace morphisms are equal *)
15   {| ct_state_morph := ct_st_m ;
16     genesis_fixpoint := gen_fix1 ;
17     cstep_morph := cstep_m1 ; |}
18   =
19   {| ct_state_morph := ct_st_m ;
20     genesis_fixpoint := gen_fix2 ;
21     cstep_morph := cstep_m2 ; |}.

```

Listing 6.7: Equality of contract trace morphisms.

Of course, a more sophisticated formulation of equality would transport over an equality of state morphisms, though we leave that for future work.

Example 6.2.3 (Contract Trace Morphism Composition). Like contract morphisms, contract trace morphisms can be composed. Similar to the composition of contract morphisms (§5.2.1), we define composition via a function `compose_ctm`, which takes morphisms

```
f : ContractTraceMorphism C1 C2 and g : ContractTraceMorphism C2 C3
```

and returns a morphism

```
compose_ctm g f : ContractTraceMorphism C1 C3.
```

To compose contract morphisms, we simply compose their component functions.

```

1 Definition compose_ctm
2   (m2 : ContractTraceMorphism C2 C3)
3   (m1 : ContractTraceMorphism C1 C2) :
4   ContractTraceMorphism C1 C3 :=
5   {|
6     ct_state_morph := compose (ct_state_morph C2 C3 m2) (ct_state_morph C1 C2 m1) ;
7     genesis_fixpoint := genesis_compose m2 m1 ;
8     cstep_morph := cstep_compose m2 m1 ;
9   |}.

```

Listing 6.8: Composition of CT Morphisms.

This relies on two functions: `genesis_compose` and `cstep_compose`. However, since `genesis_fixpoint` is a simple implication and `cstep_morph` is a function, these simply compose each component.

```

1 Definition genesis_compose
2   (m2 : ContractTraceMorphism C2 C3) (m1 : ContractTraceMorphism C1 C2) :
3   forall init_cstate,
4   is_genesis_state C1 init_cstate ->
5   is_genesis_state C3
6   (compose (ct_state_morph C2 C3 m2) (ct_state_morph C1 C2 m1) init_cstate).

```

Listing 6.9: The function `genesis_compose`

```

1 Definition cstep_compose
2   (m2 : ContractTraceMorphism C2 C3) (m1 : ContractTraceMorphism C1 C2) :
3   forall state1 state2,
4   ContractStep C1 state1 state2 ->
5   ContractStep C3
6   (compose (ct_state_morph C2 C3 m2) (ct_state_morph C1 C2 m1) state1)
7   (compose (ct_state_morph C2 C3 m2) (ct_state_morph C1 C2 m1) state2).

```

Listing 6.10: The function `cstep_compose`.

That composition is associative comes trivially.

```

1 Lemma compose_ctm_assoc
2   (f : ContractTraceMorphism C1 C2)
3   (g : ContractTraceMorphism C2 C3)
4   (h : ContractTraceMorphism C3 C4) :
5   compose_ctm h (compose_ctm g f) =
6   compose_ctm (compose_ctm h g) f.

```

Listing 6.11: Composition of CT morphisms is associative.

Similarly, it comes immediately that composition on either side with the identity is a trivial operation.

```

1 Lemma compose_id_ctm_left (f : ContractTraceMorphism C1 C2) :
2   compose_ctm (id_ctm C2) f = f.
3
4 Lemma compose_id_ctm_right (f : ContractTraceMorphism C1 C2) :
5   compose_ctm f (id_ctm C1) = f.

```

Listing 6.12: Composition with the identity does nothing.

6.2.2 The Lifting Theorem

Our first result about contract trace morphisms is that contract morphisms carry all the information needed to define a contract trace morphism. Indeed, the coherence components of a contract morphism

are simply stronger versions of the components of a contract trace morphism. We say that a contract morphism *lifts* to a contract trace morphism, and we prove this via a *lifting theorem*.

We prove this result by defining a function `cm_lift_ctm`, which takes a contract morphism

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

and returns a contract trace morphism

$$\text{cm_lift_ctm } f : \text{ContractTraceMorphism } C1 \ C2,$$

which we define formally here.

```
1 Definition cm_lift_ctm (f : ContractMorphism C1 C2) : ContractTraceMorphism C1 C2 :=
2   {}
3   ct_state_morph := state_morph _ _ f ; (* use the state morph of f *)
4   genesis_fixpoint := lift_genesis f ;  (* from f's init coherence *)
5   cstep_morph := lift_cstep_morph f ;   (* from f's recv coherence *)
6   {}.
```

Listing 6.13: Contract morphisms lift to contract trace morphisms.

As we can see in the definition, we use the function of contract states

$$\text{state_morph } _ _ f$$

from the contract morphism f to define the contract trace morphism.

We then use the `init` and `receive` coherence results of f to, respectively, prove the `genesis fixpoint` result and the `contract step` result. These are each proved and stated, respectively, as functions `lift_genesis` and `lift_cstep_morph`.

```
1 Definition lift_genesis (f : ContractMorphism C1 C2) :
2   forall init_cstate,
3     is_genesis_state C1 init_cstate ->
4     is_genesis_state C2 (state_morph C1 C2 f init_cstate).
```

Listing 6.14: Using `init` coherence from f , we prove the `genesis fixpoint` result.

```
1 Definition lift_cstep_morph (f : ContractMorphism C1 C2) :
2   forall state1 state2,
3     ContractStep C1 state1 state2 ->
4     ContractStep C2
5       (state_morph C1 C2 f state1)
6       (state_morph C1 C2 f state2).
```

Listing 6.15: Using `receive` coherence from f , we prove the `contract step` result.

Importantly, the identity contract morphism lifts to the identity contract trace morphism, and compositions of contract morphisms lift to compositions of contract trace morphisms. This will give us that isomorphic contracts are also bisimilar.

```

1 Theorem cm_lift_ctm_id :
2   cm_lift_ctm (id_cm C1) = id_ctm C1.

```

Listing 6.16: Identity lifts to identity.

```

1 Theorem cm_lift_ctm_compose
2   (g : ContractMorphism C2 C3) (f : ContractMorphism C1 C2) :
3   cm_lift_ctm (compose_cm g f) = (* lifting the composition = ... *)
4   compose_ctm (cm_lift_ctm g) (cm_lift_ctm f). (* composing the lifts *)

```

Listing 6.17: Compositions lift to compositions.

6.2.3 Contract Bisimulations

Our analogue to a contract isomorphism (Example 5.2.7) is now a contract *bisimulation*, or a contract trace isomorphism. A pair of contract trace morphisms form an isomorphism if they compose each way to the identity.

```

1 (** A bisimulation of contracts, or an isomorphism of contract traces *)
2 Definition is_iso_ctm
3   (m1 : ContractTraceMorphism C1 C2) (m2 : ContractTraceMorphism C2 C1) :=
4   compose_ctm m2 m1 = id_ctm C1 /\
5   compose_ctm m1 m2 = id_ctm C2.

```

Listing 6.18: Formal definition of a bisimulation of contracts.

Let us briefly reflect on this as a definition of a bisimulation. Bisimilarity is a stable and mathematically natural concept formulated to capture the notion of equivalence between processes [71, 117].

Consider a labelled transition system $(S, \Lambda, \rightarrow)$, where S is a set of states, Λ is a set of labels, and \rightarrow is a set of labelled transitions (a subset of $S \times \Lambda \times S$). Recall that a bisimulation is a binary relation $R \subseteq S \times S$ such that for every pair of states $(p, q) \in R$ and labels $\alpha, \beta \in \Lambda$,

- if $p \xrightarrow{\alpha} p'$, then there is $q \xrightarrow{\beta} q'$ such that $(p', q') \in R$
- if $q \xrightarrow{\beta} q'$, then there is $p \xrightarrow{\alpha} p'$ such that $(p', q') \in R$.

We can see that our formal definition of a bisimulation in Listing 6.18 achieves precisely this correspondence, as it gives a one-to-one correspondence on contract states and transitions such that corresponding transitions move between corresponding states.

With this definition in hand we can show an important corollary of the lifting theorem. Contract isomorphisms lift to contract bisimulations, due simply to the fact that the identity lifts to the identity, and compositions lift to compositions.

```

1 Corollary ciso_to_ctiso (f : ContractMorphism C1 C2) (g : ContractMorphism C2 C1) :
2   is_iso_cm f g -> is_iso_ctm (cm_lift_ctm f) (cm_lift_ctm g).

```

Listing 6.19: Contract isomorphisms induce bisimulations of contracts.

We summarize with the bisimilarity predicate on contract pairs.

```

1 Definition contracts_bisimilar
2   (C1 : Contract Setup1 Msg1 State1 Error1)
3   (C2 : Contract Setup2 Msg2 State2 Error2) :=
4   exists (f : ContractTraceMorphism C1 C2) (g : ContractTraceMorphism C2 C1),
5   is_iso_ctm f g.

```

Listing 6.20: The bisimilarity predicate on contract pairs.

Finally, we point out that bisimilarity is an equivalence relation. We prove this with three results:

Reflexivity,

```

1 Lemma bisim_refl C : contracts_bisimilar C C.

```

symmetry,

```

1 Lemma bisim_sym C1 C2 :
2   contracts_bisimilar C1 C2 ->
3   contracts_bisimilar C2 C1.

```

and transitivity.

```

1 Lemma bisim_trans C1 C2 C3 :
2   contracts_bisimilar C1 C2 /\ contracts_bisimilar C2 C3 ->
3   contracts_bisimilar C1 C3.

```

6.2.4 Discussion: Propositional Indistinguishability

Contract bisimulations indicate a strong structural correspondence between bisimilar contracts. In particular, they require an isomorphism of contract states which is preserved under contract steps. This tells us that state invariants which are also invariant under the state isomorphism of the bisimulation carry over that bisimulation.

Example 6.2.4 (Propositional Indistinguishability). Consider contracts $C1$ and $C2$, and suppose that we have morphisms

$$m1 : \text{ContractTraceMorphism } C1 \ C2 \text{ and } m2 : \text{ContractTraceMorphism } C2 \ C1$$

which form a bisimulation of contracts, *i.e.* we have a witness

```
C1-C2.bisim : is_iso_ctm m1 m2.
```

Assume that both state types, `State1` and `State2`, have a constant in storage, which we have access to via functions

```
const_in_stor_C1 : State1 -> nat and const_in_stor_C2 : State2 -> nat.
```

Furthermore, suppose that constant is invariant over the state morphism components of `m1`, meaning that for all `st : State1`,

```
const_in_stor_C1 st = const_in_stor_C2 (ct_state_morph C1 C2 m1 st).
```

Finally, suppose that `C1` has an invariant relating to this constant, *e.g.* that for all contract-reachable states,

```
const_in_stor_C1 st > 0.
```

Using the bisimulation of contracts, we can prove that the same invariant holds with the following theorem:

Theorem 6.2.1. *For all reachable chain states `bstate` with `C2` deployed at an address `caddr` with state `cstate`, the inequality `const_in_stor_C2 cstate > 0` holds.*

```
1 Theorem invariant_C2 bstate caddr (trace : ChainTrace empty_state bstate):
2   (* Forall reachable states with contract at caddr, *)
3   env_contracts bstate caddr = Some (C2 : WeakContract) ->
4   (* such that cstate is the state of the contract, *)
5   exists (cstate : State2),
6   contract_state bstate caddr = Some cstate /\
7   (* the constant in storage is > 0 *)
8   (const_in_stor_C2 cstate > 0)%nat.
```

Listing 6.21: caption text

We prove this by contract induction—since every initial state of `C2` corresponds to an initial state of `C1` under the state isomorphism of the bisimulation, and all steps of `C1` correspond to isomorphic steps of `C2`, at each step of contract induction we can take advantage of the bisimulation and the state isomorphism to prove our invariant.

Of course, contracts can be bisimilar but not identical, so it will not likely be true that *any* proposition holding for one holds for the other. Even so, this example is an illustration toward understanding the degree to which bisimilar contracts are propositionally indistinguishable. We leave any further, formal analysis of the degree to which bisimilar contracts are also propositionally indistinguishable to future work.

6.3 Contract Systems as Bigraphs

With bisimulations of contracts formally defined, we now move on to consider contract systems. Our goal will be to represent a system of interacting contracts as a single process in ConCert, which we can then specify and reason about as a single entity. The corresponding notions of bisimilarity will build off of what we have just seen.

To do so, we require a data type in which to represent systems of contracts, as well as a formal way to separate the specification of a contract system’s infrastructure—a specification of which contracts pass messages to which other contracts, and when—from the specification of its core functionality when considering the contracts as a whole. The data type that we will use is a *bigraph*, which we define in the following section.

6.3.1 Bigraphs

We give a very brief introduction to bigraphs here and direct the reader to some of Milner’s earlier writings [93, 94] for a full exposition. A bigraph is a universal mathematical model for representing the spatial configuration of physical or virtual objects, their interaction capabilities, and temporal evolution [120]. A bigraph consists of a set of nodes, denoting processes, and defines two independent structures on those nodes, which are:

- The *place graph*, a directed forest representing a spacial or nesting relationship between nodes
- The *link graph*, a hypergraph representing interactions between nodes.

Bigraphs have been the subject of both theoretical and practical work to reason about systems of processes [119, 120, 88]. Because they facilitate generic, process-algebraic approaches to reasoning about smart contracts, we will use bigraphs as a data type here in which to encode contract systems, treating individual contracts as nodes. We proceed to formalize the place graph and discuss the link graph.

6.3.2 The Place Graph

The place graph of a contract system indicates a nesting or hierarchical structure within the contract system. We argue that this is a natural way to conceptualize contract systems, and thus a natural data type for contract systems as opposed to *e.g.* lists.

Consider an AMM, one of our most basic and important financial smart contracts, which facilitates trades between token contracts. Conceptually, we might think of token contracts being nested inside of the process of the AMM. From the AMM’s perspective, the interface to interacting with the tokens is the trading contract; the trading contract then propogates transactions into its nested processes when it




Figure 6.1: A visualization of the place graph Node t $[t1 ; t2 ; t3]$. Nodes $t1$ and $t3$ have further child n-trees, while $t2$ is a leaf.

executes trades. This lends intuition to understanding how the system works, which is a good thing, but it may have practical applications as well. If one models multiple AMMs on the same blockchain, one can see visually how these AMMs interact with each other by where their place graphs intersect—along token or other contracts.

Other contract systems that we might consider include yield aggregators, like Harvest Finance or other one-stop-shop DeFi protocols like CREAM, both from the exploits of §1.2.3. Yield and DEX aggregators have an interface contract, and then try to optimize whatever operation they do—whether that is looking for highest yield pools or best price for trades. In both cases, the interface contract can be thought of as the root node of the place graph, and the processes across which it tries to optimize can be thought of as nested within that node. A similar argument goes for one-stop-shop DeFi protocols such as CREAM, which provide a single, united interface that builds off of other DeFi protocols by forwarding messages. The interface contract has nested within it the contracts that form its backend.

Finally, thinking of DeFi more broadly as a set of highly interconnected financial contracts, dubbed “money legos,” a new DeFi contract that builds off of previous contracts can be visualized as having those contracts on which it builds nested within the new contract.

With all this in mind, we proceed to define an inductive data structure for contract systems, where contracts are nodes. We first give a formal definition of an n-ary tree, or n-tree, over an arbitrary type T .

```
1 Inductive ntree (T : Type) : Type :=
2 | Node : T -> list (ntree T) -> ntree T.
```

Listing 6.22: The formal definition of an n-tree.

For example, an n-tree of the form

Node t $[t1 ; t2 ; t3]$

is a node inhabited by $t : T$, and has three child trees $t1$, $t2$, and $t3$, each of which can have any number of their own child trees. We call the node t the *root*. See Figure 6.1.

```
1 Definition ContractPlaceGraph (Setup Msg State Error : Type) :=
2   ntree (Contract Setup Msg State Error).
```

Listing 6.23: The formal definition of a contract system via its place graph.

A contract system's place graph is then simply defined to be an n-tree whose nodes are contracts. Of course, this formal definition is of a homogeneous tree, only allowing for nodes to be contracts parameterized by the same four types. In practice, systems of contracts are not homogeneous, but we will show presently how to construct a homogeneous system from a heterogeneous one.

Once defined, a contract system

```
sys : ContractPlaceGraph Setup Msg State Error
```

can be initialized and called, like a contract, but with custom `init` and `receive` functions `sys.init` and `sys.receive`. To initialize a system, we initialize the root contract with some data `s : Setup`. If successful, this returns a state `st : State`, which we call the state of the system.

```
1 Definition sys_init
2   (sys : ContractPlaceGraph Setup Msg State Error)
3   (c : Chain)
4   (ctx : ContractCallContext)
5   (s : Setup) : result State Error :=
6   match sys with
7   | Node _ root_contract _ =>
8     init root_contract c ctx s
9   end.
```

Listing 6.24: System initialization.

Then to call a system, we give it a state `st : State` and an optional message `op_msg : option Msg` and, if successful, it returns an updated state and a list of emitted actions.

```
1 Definition sys_receive
2   (sys : ContractPlaceGraph Setup Msg State Error)
3   (c : Chain)
4   (ctx : ContractCallContext)
5   (st : State)
6   (op_msg : option Msg) : result (State * list ActionBody) Error :=
7   ntree_fold_left
8   (fun (recv_propagate : result (State * list ActionBody) Error)
9     (contr : Contract Setup Msg State Error) =>
10     match recv_propagate with
11     | Ok (st0, lacts0) =>
12       match receive contr c ctx st0 op_msg with
13       | Ok (st1, lacts1) => Ok (st1, lacts0 ++ lacts1)
14       | Err e => Err e
15     end
16     | Err e => Err e
17   end)
18   sys
19   (Ok (st, nil)).
```

Listing 6.25: Calls to a contract system are governed by the `sys.receive` function.

The `sys_receive` function iteratively calls *every contract* in the system with the given message. Each call either returns an updated state or an error, and we either fold over all of those updates to the state, or propagate the error. We also accumulate all emitted actions. This is done with the n-tree fold function.

```

1 Fixpoint ntree_fold_left {A T}
2   (f : A -> T -> A)
3   (sys : ntree T)
4   (a0 : A) : A :=
5   match sys with
6   | Node _ t list_child_trees =>
7     List.fold_left
8       (fun (a0' : A) (sys' : ntree T) =>
9         ntree_fold_left f sys' a0')
10      list_child_trees
11      (f a0 t)
12   end.

```

Listing 6.26: Folding over an n-tree.

These semantics for a contract system may not seem obviously intuitive; that they accurately describe the behavior of an actual system is made more clear when we define the iterative process of building an n-tree of contracts.

6.3.2.1 Iteratively Building a Contract System

As we mentioned previously, a contract system is parameterized by the same four types that parameterize a contract, and the definition requires that all the contracts in a system be parameterized by the same four types. Of course, this does not reflect actual systems of contracts, so we define an iterative building technique to take heterogeneous contracts and make them homogeneous.

We define two functions, `c_sum_l` and `c_sum_r`. The aim of each is to produce a contract which preserves the essential functionality of, respectively, the left (`c1`) and right (`c2`) contracts of the pair of contracts (`c1 c2`) given it as an input.

Each takes a pair of contracts `c1` and `c2` of distinct types, and returns a contract whose type is the *product* of the setup and state types, and the *sum* of the message and error types

$$\text{Contract } (\text{Setup1} * \text{Setup2}) (\text{Msg1} + \text{Msg2}) (\text{State1} * \text{State2}) (\text{Error1} + \text{Error2}).$$

The intuition behind this is that the state of a system requires state values for each constituent contract; likewise, to initialize a system we need setup data for each constituent contract. But to call the contract, we only need a message for the specific contract we wish to call. Likewise if a call results in an error it only need be an error the constituent contract that was called.

We give the formal definition of each function below.


```

1 Definition c_sum_l
2   (C1 : Contract Setup1 Msg1 State1 Error1)
3   (C2 : Contract Setup2 Msg2 State2 Error2) :
4   Contract (Setup1 * Setup2) (Msg1 + Msg2) (State1 * State2) (Error1 + Error2).

```

Listing 6.27: `c_sum_l` produces a contract with the essential functionality of the left contract, `C1`.

```

1 Definition c_sum_r
2   (C1 : Contract Setup1 Msg1 State1 Error1)
3   (C2 : Contract Setup2 Msg2 State2 Error2) :
4   Contract (Setup1 * Setup2) (Msg1 + Msg2) (State1 * State2) (Error1 + Error2).

```

Listing 6.28: `c_sum_r` produces a contract with the essential functionality of the right contract, `C2`.

A call to `c_sum_l C1 C2` takes a state $(st1, st2) : State1 * State2$. Given a message `inl msg`—that is, a message `msg : Msg1` to `C1`—it calls the receive function of `C1` on `st1` and `msg`. If the call to `(receive C1)` is successful, returning an updated state `st1'`, we return $(st1', st2)$ and any emitted actions. If given a message `inr msg`—that is, a message `msg : Msg2` to `C2`—the call does nothing, returning $(st1, st2)$ and no emitted actions.

By executing along the semantics of `C1` when called with a message to `C1`, and doing nothing otherwise, the function `c_sum_l` is meant to give us a contract that represents `C1` within the context of `C1` and `C2` together. A call to `c_sum_r C1 C2` executes in an analogous fashion, but calling `C2` with a message `inr msg` and doing nothing otherwise.

Using `c_sum_l` and `c_sum_r`, we can iteratively build systems. First, we define the simple example of nesting a contract `C2` within another contract `C1`. We might do this, for example, if `C1` is an interface contract and `C2` is a backend contract.

```

1 Definition nest
2   (C1 : Contract Setup1 Msg1 State1 Error1)
3   (C2 : Contract Setup2 Msg2 State2 Error2) :
4   ContractPlaceGraph (Setup1 * Setup2) (Msg1 + Msg2) (State1 * State2) (Error1 + Error2)
5   := let T :=
6   Contract (Setup1 * Setup2) (Msg1 + Msg2) (State1 * State2) (Error1 + Error2) in
7   Node T (c_sum_l C1 C2) [Node T (c_sum_r C1 C2) nil].

```

Listing 6.29: A function to nest a contract `C2` within another contract `C1`.

The function `nest` takes two contracts `C1` and `C2`, and takes their image under `c_sum_l` and `c_sum_r`, producing a place graph

$$\text{nest } C1 \ C2 := \text{Node } T \ (\text{c_sum_l } C1 \ C2) \ [\text{Node } T \ (\text{c_sum_r } C1 \ C2) \ \text{nil}].$$

The root contract, `c_sum_l C1 C2`, represents `C1`, and the child contract, `c_sum_r C1 C2`, represents `C2`. The place graph `nest C1 C2` can receive messages of the form `inl msg`, targeting `C1`, or `inr msg`,




Figure 6.2: A contract C_2 nested within another contract C_1 .

targeting C_2 , and the state of $\text{nest } C_1 \ C_2$ is a witness of the state of C_1 as well as the state of C_2 . The system $\text{nest } C_1 \ C_2$ makes progress as its constituents, C_1 and C_2 , are called successfully and update the system's state. See Figure 6.2.

Moving to a slightly more complicated example, we can iteratively add children to an existing contract system. We start with the singleton system, which is simply a singleton n-tree, or an n-tree with only one node, of the form $\text{Node } _ \ C \ \text{nil}$.

```

1 Definition singleton_place_graph
2   (C : Contract Setup Msg State Error)
3   : ContractPlaceGraph Setup Msg State Error :=
4   Node _ C nil.

```

From here, we can add children to the singleton system via a function `sys_add_child_r`.

```

1 Definition sys_add_child_r
2   (sys : ContractPlaceGraph Setup1 Msg1 State1 Error1)
3   (C : Contract Setup2 Msg2 State2 Error2) :
4   ContractPlaceGraph (Setup1 * Setup2) (Msg1 + Msg2) (State1 * State2) (Error1 + Error2)
5   :=
6   let T := Contract (Setup1 * Setup2) (Msg1 + Msg2) (State1 * State2) (Error1 + Error2)
7   in
8   match sys with
9   | Node _ root_contract _ =>
10      match (ntree_map (fun C1 => c_sum_l C1 C) sys) with
11      | Node _ root_contract' children =>
12          Node T root_contract' (children ++ [Node T (c_sum_r root_contract C) nil])
13      end
14   end.

```

Listing 6.30: Iteratively add a child C to a contract system sys .

Given a contract system sys , we iteratively add a child C to the system by first mapping over sys with $\text{c_sum_l } _ \ C$, or rather

(fun C1 => c_sum_l C1 C),

and then appending the right sum of the root contract `root_contract` and C ,

c_sum_r root_contract C,




Figure 6.3: Appending a child C to a system Node $\text{root } [\text{child}_1 ; \text{child}_2 ; \dots]$.

to the list of children of the root contract. This maintains the same functionality of each node inherited from sys , while adding C to the system, by iteratively applying c_sum_l and c_sum_r . See Figure 6.3.

6.3.2.2 System Contracts, Morphisms, and Isomorphisms

System place graphs are inductive structures built on contracts, so naturally we might consider morphisms of contract systems, building off of contract morphisms. In this context, our main goal is to establish notions of equivalences of system place graphs analogous to what we saw in §6.2.3.

Like contract morphisms, a morphism of system place graphs consists of four component morphisms, along with coherence conditions for sys_init and sys_receive .

```

1  (** A morphism of system place graphs *)
2  Record SystemMorphism
3    (sys1 : ContractPlaceGraph Setup1 Msg1 State1 Error1)
4    (sys2 : ContractPlaceGraph Setup2 Msg2 State2 Error2) :=
5    build_system_morphism {
6      (* the components of a morphism *)
7      sys_setup_morph : Setup1 -> Setup2 ;
8      sys_msg_morph   : Msg1   -> Msg2   ;
9      sys_state_morph : State1 -> State2 ;
10     sys_error_morph : Error1 -> Error2 ;
11     (* coherence conditions *)
12     sys_init_coherence : forall c ctx s,
13       result_functor sys_state_morph sys_error_morph
14         (sys_init sys1 c ctx s) =
15         sys_init sys2 c ctx (sys_setup_morph s) ;
16     sys_recv_coherence : forall c ctx st op_msg,
17       result_functor (fun ' (st, l) => (sys_state_morph st, l)) sys_error_morph
18         (sys_receive sys1 c ctx st op_msg) =
19         sys_receive sys2 c ctx (sys_state_morph st) (option_map sys_msg_morph op_msg)
20   ;

```

Listing 6.31: The formal definition of a morphism of system place graphs.

As before, composition is given by composition of component morphisms, which we formalize as follows, and composition is associative.

```

1 Definition compose_sm (g : SystemMorphism sys2 sys3) (f : SystemMorphism sys1 sys2) :
2   SystemMorphism sys1 sys3 := { |
3     (* the components *)
4     sys_setup_morph := compose (sys_setup_morph sys2 sys3 g) (sys_setup_morph sys1 sys2 f) ;
5     sys_msg_morph    := compose (sys_msg_morph sys2 sys3 g) (sys_msg_morph sys1 sys2 f) ;
6     sys_state_morph := compose (sys_state_morph sys2 sys3 g) (sys_state_morph sys1 sys2 f) ;
7     sys_error_morph := compose (sys_error_morph sys2 sys3 g) (sys_error_morph sys1 sys2 f) ;
8     (* the coherence results *)
9     sys_init_coherence := sys_compose_init_coh g f ;
10    sys_recv_coherence := sys_compose_recv_coh g f ;
11  | }.

```

Listing 6.32: Composition of morphisms of system place graphs.

The identity system morphism is given by identity component functions, given as follows.

```

1 Definition id_sm (sys : ContractPlaceGraph Setup Msg State Error) :
2   SystemMorphism sys sys := { |
3     (* components *)
4     sys_setup_morph := id ;
5     sys_msg_morph    := id ;
6     sys_state_morph := id ;
7     sys_error_morph := id ;
8     (* coherence conditions *)
9     sys_init_coherence := sys_init_coherence_id sys ;
10    sys_recv_coherence := sys_recv_coherence_id sys ;
11  | }.

```

Listing 6.33: The identity morphism of system place graphs.

Similar to before we define an isomorphism of systems as a pair of morphisms that compose each way to the identity.

```

1 Definition is_iso_sm (m1 : SystemMorphism sys1 sys2) (m2 : SystemMorphism sys2 sys1) :=
2   compose_sm m2 m1 = id_sm sys1 /\
3   compose_sm m1 m2 = id_sm sys2.

```

Listing 6.34: An isomorphism of contract systems.

The reader may notice that system morphisms look almost identical to contract morphisms. Because system place graphs have the `sys.init` and `sys.receive` functions, which mimic the `init` and `receive` of contracts, we might ask whether systems are themselves contracts.

Indeed, this is the case. We can define the *system contract*, which is the contract in ConCert that represents a contract system.

```

1 Definition sys_contract (sys : ContractPlaceGraph Setup Msg State Error) :=
2   build_contract (sys_init sys) (sys_receive sys).

```

Listing 6.35: The system contract.

Furthermore, system morphisms are in one-to-one correspondence with the morphisms of the corresponding contracts. We have two functions—one that takes system morphisms to contract morphisms,

```

1 Definition sysm_to_cm
2   (sys1 : ContractPlaceGraph Setup1 Msg1 State1 Error1)
3   (sys2 : ContractPlaceGraph Setup2 Msg2 State2 Error2)
4   (f : SystemMorphism sys1 sys2) : ContractMorphism (sys_contract sys1) (sys_contract
   sys2).

```

and one that takes contract morphisms to system morphisms—

```

1 Definition cm_to_sysm
2   (sys1 : ContractPlaceGraph Setup1 Msg1 State1 Error1)
3   (sys2 : ContractPlaceGraph Setup2 Msg2 State2 Error2)
4   (f : ContractMorphism (sys_contract sys1) (sys_contract sys2)) : SystemMorphism sys1
   sys2.

```

which compose each way to the identity. In this correspondence, the identity corresponds to the identity, and compositions to compositions.

6.3.3 The Link Graph

Now that we have the data type for contract systems as its place graph, let us turn our attention to a contract system’s link graph. In a bigraph, the link graph represents the interactions of processes in a system, as opposed to the place graph which only represents a spacial, or hierarchical relationship between processes.

For contract systems, a link graph has an obvious definition, at least intuitively. That is, that for contracts $C1$ and $C2$ in a contract system, there is an edge

$$C1 \longrightarrow C2$$

when there is a contract step of $C1$ which emits a call to $C2$. Even more, because of the semantics of a blockchain, that call to $C1$ succeeds only if the emitted call to $C2$ succeeds. Indeed, a contract call succeeds if and only if it doesn’t throw an error and all the calls which it emits succeed.

Furthermore, the place graph is essential to the semantics of a system of contracts. Going back to the example of an interface and backend contract forming a system, that the link graph connect calls to the

interface to (permissioned) calls to the backend is essential to the system functioning as intended. Indeed, the link graph dictates how systems move forward, grouping and ordering contract calls that constitute a genuine step forward for the system. Depending on the link graph structure, which itself depends on contract addresses and message-passing, a contract system can behave in very different ways.

A genuine step forward for a system, then, is a full traversal of the link graph from the point of entry of an initiating call. Such a traversal is a sequence of recursive calls to `sys.receive`, each of which is a step forward for the system, such that in the end the accumulated, emitted transactions do not include any system-recursive calls.

6.3.3.1 System Steps and System Trace

To encode the semantics of a link graph into a contract system, we need first to formalize the semantics of a system stepping forward, and its trace.

Consider a system of interacting contracts that has a frontend, or interface, contract, and a backend contract. Messages come in through the frontend, whose storage has no meaningful data to the system, and are simply forwarded to the backend, where the meaningful part of the contract system’s storage is. A transaction coming in through the interface only really has full meaning for the system when the ensuing chain of transactions is completed.

Now compare such a system to its monolithic counterpart—where the interface and backend are housed in the same contract and there is no need for any message-passing—then the contract call corresponding to the call to the interface and message to the backend would be a single call and execution to the monolithic contract which updates the storage.

Because of this, we distinguish between incremental, or single, steps, and system steps. The distinction is that a system step can be one, many, or no incremental steps. An incremental step is defined to be a single, successful call to the contract system.

```

1 Record SingleSystemStep (sys : ContractPlaceGraph Setup Msg State Error)
2   (prev_sys_state next_sys_state : State) :=
3   build_sys_single_step {
4     sys_step_chain : Chain ;
5     sys_step_ctx : ContractCallContext ;
6     sys_step_msg : option Msg ;
7     sys_step_nacts : list ActionBody ;
8     (* we can call receive successfully *)
9     sys_recv_ok_step :
10       sys_receive sys sys_step_chain sys_step_ctx prev_sys_state sys_step_msg =
11       Ok (next_sys_state, sys_step_nacts) ;
12 }.
```

Listing 6.36: An incremental, or single, step of a contract system.

Irrespective of the contract system in question, any meaningful step forward for a system must be expressible as a chained list of single system steps, and thus has semantics in the following type.

```
1 Definition ChainedSingleSteps (sys : ContractPlaceGraph Setup Msg State Error) :=
2   ChainedList State (SingleSystemStep sys).
```

Listing 6.37: A system step can be one or several incremental steps.

On the other hand, it is not true that *any* chained list of system steps constitutes a meaningful step forward. We must instead define what it means for a contract system to meaningfully step forward by defining a function,

$$\text{sys_link} : \text{State} \rightarrow \text{State} \rightarrow \text{Type},$$

which links two system states by a system step forward. Those steps must have semantics as chained single steps via a function of the form

$$\text{sys_link_semantics } st1 \ st2 : \text{sys_link } st1 \ st2 \rightarrow \text{ChainedSingleSteps } st1 \ st2.$$

Such a definition of system steps is in fact a *specification* because whether or not a system actually satisfies these semantics for stepping forward depends on conditions of a given state of the chain, including contract addresses and contract storage. This is our definition of a system's link graph.

From here, we can give a full definition of a contract system as a bigraph: a set of contracts, encoded in a place graph, with link graph semantics.

```
1 Record ContractSystem (Setup Msg State Error : Type) :=
2   build_contract_system {
3     (* the place and link graphs *)
4     sys_place : ContractPlaceGraph Setup Msg State Error ;
5     sys_link   : State -> State -> Type ;
6     (* the link graph has semantics in ChainedSingleSteps *)
7     sys_link_semantics : forall st1 st2,
8       sys_link st1 st2 ->
9       ChainedSingleSteps sys_place st1 st2 ;
10  }.
```

Given this definition of a contract system, we can define a type of system steps for any system *sys*, which are steps given by the link graph semantics.

```
1 Definition SystemStep (sys : ContractSystem Setup Msg State Error) :=
2   sys_link _ _ _ sys.
```

Listing 6.38: The type of system steps.

Finally, a system trace is a chained list of system steps.

```
1 Definition SystemTrace (sys : ContractSystem Setup Msg State) :=
2   ChainedList (SystemState State) (SystemStep sys).
```

Listing 6.39: A contract system's trace is a chained list of system steps.

As we did for contracts in §6.2.3, we use the definitions of system steps and system trace to reason about bisimulations of systems.

6.4 Bisimulations of Contract Systems

With the data type for contract systems in place, we are able to now define bisimulations of contract systems, building off of §6.2.1. Like bisimulations of contracts, a bisimulation of contract systems is a correspondence between system states, and steps forward in a system state. Because system steps are defined by the link graph, a bisimulation of systems indicates a correspondence between the link graph structure of two contract systems. As before, we will consider a system's trace, and system trace morphisms, to establish bisimulations.

6.4.1 System Trace Morphisms and System Bisimulations

A system trace morphism between systems `sys1` and `sys2` consists of a function from the state type of `sys1` to that of `sys2` which sends initial states to initial states, and steps defined by the link graph of `sys1` to corresponding steps defined by the link graph of `sys2`.

```
1 Record SystemTraceMorphism
2   (sys1 : ContractSystem Setup1 Msg1 State1 Error1)
3   (sys2 : ContractSystem Setup2 Msg2 State2 Error2) :=
4   build_st_morph {
5     (* a function *)
6     st_state_morph : State1 -> State2 ;
7     (* init state sys1 -> init state sys2 *)
8     sys_genesis_fixpoint : forall init_sys_state,
9       is_genesis_sys_state sys1 init_sys_state ->
10      is_genesis_sys_state sys2 (st_state_morph init_sys_state) ;
11     (* step morphism *)
12     sys_step_morph : forall sys_state1 sys_state2,
13       SystemStep sys1 sys_state1 sys_state2 ->
14       SystemStep sys2 (st_state_morph sys_state1) (st_state_morph sys_state2) ;
15   }.
```

Listing 6.40: A morphism of system traces.

We can define composition analogously to contract trace morphisms, and as before, composition is associative. Composition relies on two functions, `sys_genesis_compose` and `sys_step_compose`, which are also defined analogously to their counterparts in contract trace morphisms.

```

1 Definition compose_stm
2   (m2 : SystemTraceMorphism sys2 sys3)
3   (m1 : SystemTraceMorphism sys1 sys2) : SystemTraceMorphism sys1 sys3 := { |
4     st_state_morph := compose (st_state_morph _ _ m2) (st_state_morph _ _ m1) ;
5     sys_genesis_fixpoint := sys_genesis_compose m2 m1 ;
6     sys_step_morph := sys_step_compose m2 m1 ;
7   | }.

```

Listing 6.41: Composition of system trace morphisms.

The dependent notion of equality of system trace morphisms also mirrors that of contract trace morphisms.

```

1 Lemma eq_stm_dep
2   (sys1 : ContractSystem Setup1 Msg1 State1 Error1)
3   (sys2 : ContractSystem Setup2 Msg2 State2 Error2)
4   (st_st_m : State1 -> State2)
5   sys_gen_fix1 sys_gen_fix2
6   (sys_step_m1 sys_step_m2 : forall sys_state1 sys_state2,
7     SystemStep sys1 sys_state1 sys_state2 ->
8     SystemStep sys2 (st_st_m sys_state1) (st_st_m sys_state2)) :
9   sys_step_m1 = sys_step_m2 ->
10  { | st_state_morph := st_st_m ;
11     sys_genesis_fixpoint := sys_gen_fix1 ;
12     sys_step_morph := sys_step_m1 ; | }
13  =
14  { | st_state_morph := st_st_m ;
15     sys_genesis_fixpoint := sys_gen_fix2 ;
16     sys_step_morph := sys_step_m2 ; | }.

```

Listing 6.42: Equality of system trace morphisms.

Finally, the identity morphism is defined similarly to that of contract trace morphisms.

```

1 Definition id_stm
2   (sys : ContractSystem Setup Msg State Error) : SystemTraceMorphism sys sys :=
3   { |
4     st_state_morph := id ;
5     sys_genesis_fixpoint := id_sys_genesis_fixpoint sys ;
6     sys_step_morph := id_sys_step_morph sys ;
7   | }.

```

Listing 6.43: The identity system trace morphism.

It relies on two functions, which are both defined trivially.

```

1 Definition id_sys_genesis_fixpoint (sys : ContractSystem Setup Msg State Error)
2   init_sys_state
3   (gen_sys : is_genesis_sys_state sys init_sys_state) :
4   is_genesis_sys_state sys (id init_sys_state) :=
5   gen_sys.

```

Listing 6.44: The genesis fixpoint result for the identity system trace morphism.

```

1 Definition id_sys_step_morph (sys : ContractSystem Setup Msg State Error)
2   sys_state1 sys_state2 (step : SystemStep sys sys_state1 sys_state2) :
3   SystemStep sys (id sys_state1) (id sys_state2) :=
4   step.

```

Listing 6.45: The step result for the identity system trace morphism.

We then define system trace isomorphisms as morphisms which compose each way to the identity.

```

1 Definition is_iso_stm
2   (m1 : SystemTraceMorphism sys1 sys2) (m2 : SystemTraceMorphism sys2 sys1) :=
3   compose_stm m2 m1 = id_stm sys1 /\
4   compose_stm m1 m2 = id_stm sys2.

```

Listing 6.46: An isomorphism of system trace morphisms is a pair of system trace morphisms which compose each way to the identity.

And finally, a bisimulation of systems is an isomorphism of system traces.

```

1 Definition systems_bisimilar
2   (sys1 : ContractSystem Setup1 Msg1 State1 Error1)
3   (sys2 : ContractSystem Setup2 Msg2 State2 Error2) :=
4   exists (f : SystemTraceMorphism sys1 sys2) (g : SystemTraceMorphism sys2 sys1),
5   is_iso_stm f g.

```

Listing 6.47: The formal definition of a bisimulation of contract systems.

A note on system bisimulations. Because system traces are defined by a system’s link graph, in order for a bisimulation of contract systems to have legitimate semantic meaning, the system must satisfy the specification given by its link graph. This will typically mean that constituent contracts of the system are deployed at certain addresses, and that contracts are able to call each other, *e.g.* by having the appropriate addresses in storage so that they can emit transactions which call other contracts in the system. Crucially, a bisimulation of systems only reflects an equivalence for the behavior of a system which is deployed in such a way that it satisfies the place graph.

6.4.2 Lifting Theorems for Contract Systems

We have various lifting theorems for system morphisms to system trace morphisms which generalize how contract morphisms lift to contract trace morphisms.

System morphisms indicate a relationship between single system steps of two systems, and so from a system morphism we can derive a correspondence between all single system steps. Thus system morphisms lift to system trace morphisms if the link graph structure on each system is compatible with regards to individual system steps. In particular, if both system steps have the *discrete* link graph, for which all system steps are single system steps, then system morphisms lift to system trace morphisms.

More formally, given a system place graph, we can construct the discrete link graph, which is a link graph where all system steps are given by single system steps.

```

1 Definition discrete_link (sys : ContractPlaceGraph Setup Msg State Error) st1 st2 :=
2   SingleSystemStep sys st1 st2.
3
4 Definition discrete_link_semantics (sys : ContractPlaceGraph Setup Msg State Error)
5   st1 st2 (step : discrete_link sys st1 st2) :
6   ChainedSingleSteps sys st1 st2 :=
7   (snoc clnil step).
8
9 Definition discrete_sys (sys : ContractPlaceGraph Setup Msg State Error) := { |
10   sys_place := sys ;
11   sys_link := discrete_link sys ;
12   sys_link_semantics := discrete_link_semantics sys ;
13 | }.

```

Listing 6.48: The discrete link graph construction on any contract place graph.

The lifting theorem for system trace morphisms is that a system morphism lifts to a system trace morphism of discrete systems.

```

1 Definition sm_lift_stm (f : SystemMorphism sys1 sys2) :
2   SystemTraceMorphism (discrete_sys sys1) (discrete_sys sys2).

```

Listing 6.49: A function that lifts system morphisms to system trace morphisms.

We have that the identity lifts to the identity,

```

1 Theorem sm_lift_stm_id :
2   sm_lift_stm (id_sm sys1) = id_stm (discrete_sys sys1).

```

Listing 6.50: The identity system morphism lifts to the identity system trace morphism under the discrete link graph.

and compositions to compositions.

```

1 Lemma sm_lift_stm_compose (g : SystemMorphism sys2 sys3) (f : SystemMorphism sys1 sys2) :
2   sm_lift_stm (compose_sm g f) =
3   compose_stm (sm_lift_stm g) (sm_lift_stm f).

```

Listing 6.51: Compositions of system morphisms lift to compositions of system trace morphisms.

Thus isomorphic systems are also bisimilar, under the discrete link graph.

```

1 Corollary sys_iso_to_bisim
2   (sys1 : ContractPlaceGraph Setup1 Msg1 State1 Error1)
3   (sys2 : ContractPlaceGraph Setup2 Msg2 State2 Error2) :
4   systems_isomorphic sys1 sys2 ->
5   systems_bisimilar (discrete_sys sys1) (discrete_sys sys2).

```

Listing 6.52: Isomorphism contract systems are bisimilar under the discrete link graph.

Furthermore, because system morphisms correspond to contract morphisms of the system contract, a contract morphism lifts to a system morphism of singleton systems (systems with only one constituent contract). We define this with a function `lift_cm_to_sm`.

```

1 Definition lift_cm_to_sm (f : ContractMorphism C1 C2) :
2   SystemMorphism (singleton_place_graph C1) (singleton_place_graph C2).

```

Listing 6.53: A contract morphism lifts to a system morphism of singleton systems.

Similarly, contract morphisms lift to system trace morphism on the singleton system.

```

1 Definition lift_ctm_to_stm (f : ContractTraceMorphism C1 C2) :
2   SystemTraceMorphism
3     (discrete_sys (singleton_place_graph C1))
4     (discrete_sys (singleton_place_graph C2)).

```

The identity contract morphism lifts to the identity system morphism on the singleton place graph,

```

1 Lemma lift_id_cm_to_id_sm :
2   lift_cm_to_sm (id_cm C) = id_sm (singleton_place_graph C).

```

Listing 6.54: The identity contract morphism lifts to the identity system morphism of singleton systems.

and compositions lift to compositions.

```

1 Lemma lift_cm_to_sm_comp
2   (f : ContractMorphism C1 C2) (g : ContractMorphism C2 C3) :
3   lift_cm_to_sm (compose_cm g f) = compose_sm (lift_cm_to_sm g) (lift_cm_to_sm f).

```

Thus isomorphic contracts are isomorphic (and thus bisimilar) as singleton systems with the discrete link graph.

```

1 Theorem c_iso_csys_iso :
2   contracts_isomorphic C1 C2 ->
3   systems_isomorphic (singleton_place_graph C1) (singleton_place_graph C2).

```

Listing 6.55: The singleton systems of isomorphic contracts are isomorphic.

These lifting theorems show the details of the generalization from contracts, contract trace morphisms, and contract bisimulations to contract systems, system trace morphisms, and system bisimulations.

6.5 Conclusion

Contract system specifications obfuscate the intent of the core functionality of the system with details about how different processes pass messages between each other and how their executions relate to each other. Our core thesis of this chapter was that this can be mitigated by separating the specification into that of system infrastructure and that of the core, intended behavior, the latter being agnostic to a contract’s modular structure. In contrast with other efforts to address system complexity, which use tools like model checkers to reason in a composable manner *over* the system, we develop process-algebraic formalisms with the intent of separating the specification of the system’s infrastructure from its core functionality.

In particular, we formalized contract systems as bigraphs, which describe the spacial hierarchy of a contract system as well as how the constituent contracts are linked through contract calls. A contract system’s link graph is a specification of system infrastructure. Furthermore, contract systems defined with a place and link graph take a computational form which acts coherently, as one single process. Our formalisms are designed so that a specified contract can be implemented as a single contract, or proved to be bisimilar to a contract whose specification isolates that of the desired, core functionality.

Though we only used bigraphs in this chapter as a data type, the process algebraic properties of bigraphs have applications further than what we have mentioned. By embedding the process-algebraic semantics of bigraphs into ConCert, it may be possible to reason about contract systems in more sophisticated ways that leverage the process algebra, building on previous work relating to bigraphs [120] and formally reasoning about composable DeFi protocols [130].

Future work includes formalizing theories to support such work, as well as using the formalisms of this chapter to attempt verification of a practical, deployed system of contracts in terms of a single process bisimilar to the system.

Chapter 7

Conclusion

Financial smart contracts have complicated specifications. Meta properties, by definition, are difficult to accurately capture with a contract specification, and the failure to do so routinely exposes fatal smart contract vulnerabilities. However, formal verification tools based in interactive theorem provers, like ConCert, offer us a chance to introduce mathematical techniques to reason about contracts and their specifications with greater precision and mathematical maturity.

We targeted three classes of meta properties common to financial smart contracts: the intended economic properties (§4), the intended upgradeability properties (§5), and the properties intended of a system of interacting contracts, taken as a whole (§6). For each we introduced formal and theoretical tools into ConCert designed to rigorously specify each class of meta properties, including contract metaspecifications, contract morphisms, and bisimulations of both individual contracts and contract systems. In each of these cases, we showed examples to illustrate how these tools can be used in specification and proof.

The goal of this work has been to add mathematical precision and maturity to the art of contract specification and verification, in the hopes that critical financial infrastructure can be more carefully designed, rigorously specified, formally verified, and safely deployed to the blockchain.

7.1 Limitations and Future Work

There are limitations to our work which we hope to address in future work.

In Chapter 4, we stated and verified economic properties derived from a theoretical formulation of AMMs and DeFi which modeled a blockchain as a state machine [19, 20]. Crucially, these formulations are not formalized. We could more rigorously state and prove economic properties of financial smart contracts if we formalized a theory of AMMs and DeFi in ConCert, for which the bigraph embedding of Chapter 6 lays a foundation.

To do so, we must first formally define DeFi *primitives*, including swaps, swap rate, exchange rate, liquidity provider, and liquidity. From these we can give a formalized derivation of key properties of financial smart contracts, including those that we stated in the structured pools metaspecification of §4.5: demand sensitivity, incentive consistency, positive trading cost, and zero-impact liquidity change. Such an embedding makes *all* of the formal reasoning relevant to the contract specification’s economic properties fully and clearly encoded in ConCert. The result improves on the rigor of previously cited work on theories of DeFi and AMMs [19, 20] because the DeFi primitives and subsequently derived properties have explicit, formalized semantics in a process algebra, and can be stated of contracts with verified extraction.

The ultimate goal of embedding a theory of DeFi and AMMs into ConCert is to produce a comprehensive, usable, Coq-based toolkit for precisely stating and verifying economic properties of a contract specification. From there it could be made adaptable and automated with specialized tactics and a broad library of verified economic properties, usable at least in principle in contract specifications by engineers in the wild. While the foundations will be process-algebraic (using bigraphs), carefully formalized DeFi primitives could make this scalable by largely abstracting complex process-algebraic reasoning away. Ideally, future researchers could build off of these foundations to reason, in an ever more sophisticated way, about behaviors and vulnerabilities of financial smart contracts.

Another direction of future work builds off the observation of §5.5.4 that general upgradeability frameworks exist in analogy to *fiber bundles*, a geometric construction which is used in mathematics and physics to separate the interacting structure of various components of complicated mathematical objects. These results suggest that programs formalized within a proof assistant exhibit properties commonly articulated by mathematicians—in this case, of geometric objects and topological spaces.

There is good reason to suspect that these results are indicative of something deeper, that techniques from topology have real application to formal verification. Topology, by way of homotopy theory, is known to have deep connections to computation by way of *computational trinitarianism* [68, 98]. In this theoretical context, computation, propositions, and proofs all have a geometric, or topological, interpretation. Geometries also emerge in various type theories under active study, including Homotopy Type Theory (HoTT) [21] and others [39, 50]. Each of these type theories are logics as well as computable foundations of geometry, so propositions and proofs have an associated geometry. From this it is natural to hypothesize that the formal structure and behavior of programs may have geometric properties, and so might the propositions and proofs of a formal specification—and that these geometries may interact.

An explicit study of the geometry of programs, propositions, and proofs may be of interest to computer scientists and mathematicians alike. To computer scientists, in formal methods. We have already shown that programs themselves can have geometric properties which can be leveraged in formal specification and verification. It may also be worthwhile to understand the relationship of a proposition’s formal proof to its geometry to support work in *proof repair* [113, 114], by understanding geometrically how alterations to a proposition correspond to alterations to its proof. To mathematicians, an understanding of the

geometry of propositions and their proofs gives a formal and geometric way to study the (non-)equivalence of proofs and proof spaces of theorems, adding theoretical richness to the common practice of searching for multiple proofs of a single theorem [44].

The overarching goal, both of this thesis and any future work, is to make the practice of formal verification—stating propositions and supplying proofs—more effectual by adding to its mathematical maturity. This is done by treating formalized programs as well-defined mathematical objects, and introducing formal, mathematical techniques to state and prove propositions about programs which are difficult to state correctly in prose. Because programs are vulnerable to poor specifications as much as they are to incorrect code, doing so could make formally verified software more secure by grounding the process of formal verification deeper in mathematical theory.

References

- [1] 20squares. <https://20squares.xyz/>. Accessed July 2023.
- [2] Curve finance. <https://curve.fi/>. Accessed July 2023.
- [3] Uniswap Protocol. <https://uniswap.org/>. Accessed July 2023.
- [4] Hayden Adams, Noah Zinsmeister, Moody Salem, River Keefer, and Dan Robinson. Uniswap v3 core. <https://rholang.io/community-site/whitepaper-v3.pdf>. Accessed July 2023.
- [5] Sidney Amani, Myriam Bégel, Maksym Bortin, and Mark Staples. Towards verifying ethereum smart contract bytecode in Isabelle/HOL. In *Proceedings of the 7th ACM SIGPLAN International Conference on Certified Programs and Proofs*, CPP 2018, pages 66–77, New York, NY, USA, January 2018. Association for Computing Machinery. ISBN 978-1-4503-5586-5. doi: 10.1145/3167084.
- [6] Guillermo Angeris, Akshay Agrawal, Alex Evans, Tarun Chitra, and Stephen Boyd. Constant Function Market Makers: Multi-Asset Trades via Convex Optimization. arXiv:2107.12484, July 2021.
- [7] Guillermo Angeris, Hsien-Tang Kao, Rei Chiang, Charlie Noyes, and Tarun Chitra. An Analysis of Uniswap markets. *Cryptoeconomic Systems*, 0(1), April 2021. ISSN 2767-4207,. doi: 10.21428/58320208.c9738e64.
- [8] Guillermo Angeris, Tarun Chitra, and Alex Evans. When Does The Tail Wag The Dog? Curvature and Market Making. *Cryptoeconomic Systems*, 2(1), June 2022. ISSN 2767-4207,. doi: 10.21428/58320208.e9e6b7ce.
- [9] Danil Annenkov and Bas Spitters. Deep and shallow embeddings in Coq. TYPES, 2019.
- [10] Danil Annenkov, Jakob Botsch Nielsen, and Bas Spitters. ConCert: A smart contract certification framework in Coq. In *Proceedings of the 9th ACM SIGPLAN International Conference on Certified Programs and Proofs*, CPP 2020, pages 215–228, New York, NY, USA, January 2020. Association for Computing Machinery. ISBN 978-1-4503-7097-4. doi: 10.1145/3372885.3373829.
- [11] Danil Annenkov, Mikkel Milo, Jakob Botsch Nielsen, and Bas Spitters. Extracting smart contracts tested and verified in Coq. In *Proceedings of the 10th ACM SIGPLAN International Conference*

- on *Certified Programs and Proofs*, CPP 2021, pages 105–121, New York, NY, USA, January 2021. Association for Computing Machinery. ISBN 978-1-4503-8299-1. doi: 10.1145/3437992.3439934.
- [12] Danil Annenkov, Mikkel Milo, and Bas Spitters. Code Extraction from Coq to ML-like languages. 2021. URL <https://icfp21.sigplan.org/details/mlfamilyworkshop-2021-papers/8/Code-Extraction-from-Coq-to-ML-like-languages>.
 - [13] Danil Annenkov, Mikkel Milo, Jakob Botsch Nielsen, and Bas Spitters. Extracting functional programs from Coq, in Coq. *Journal of Functional Programming*, 32:e11, 2022/ed. ISSN 0956-7968, 1469-7653. doi: 10.1017/S0956796822000077.
 - [14] Pedro Antonino, Juliandson Ferreira, Augusto Sampaio, and A. W. Roscoe. Specification is Law: Safe Creation and Upgrade of Ethereum Smart Contracts. In *Software Engineering and Formal Methods: 20th International Conference, SEFM 2022, Berlin, Germany, September 26–30, 2022, Proceedings*, pages 227–243. Springer, 2022.
 - [15] Avraham Eisenberg (@avi_eisen). Mango Markets Exploit. https://twitter.com/avi_eisen/status/1581326199682265088, October 2022. Accessed July 2023.
 - [16] b. Dexter2 Specification. <https://gitlab.com/dexter2tz/dexter2tz/-/blob/master/docs/informal-spec/dexter2-cpmm.md>. Accessed July 2023.
 - [17] Massimo Bartoletti and Roberto Zunino. Formal models of bitcoin contracts: A survey. *Frontiers in Blockchain*, 2:8, 2019.
 - [18] Massimo Bartoletti, James Hsin-yu Chiang, and Alberto Lluch Lafuente. SoK: Lending Pools in Decentralized Finance. In Matthew Bernhard, Andrea Bracciali, Lewis Gudgeon, Thomas Haines, Arian Klages-Mundt, Shin’ichiro Matsuo, Daniel Perez, Massimiliano Sala, and Sam Werner, editors, *Financial Cryptography and Data Security. FC 2021 International Workshops*, Lecture Notes in Computer Science, pages 553–578, Berlin, Heidelberg, 2021. Springer. ISBN 978-3-662-63958-0. doi: 10.1007/978-3-662-63958-0_40.
 - [19] Massimo Bartoletti, James Hsin-yu Chiang, and Alberto Lluch Lafuente. Towards a Theory of Decentralized Finance. In Matthew Bernhard, Andrea Bracciali, Lewis Gudgeon, Thomas Haines, Arian Klages-Mundt, Shin’ichiro Matsuo, Daniel Perez, Massimiliano Sala, and Sam Werner, editors, *Financial Cryptography and Data Security. FC 2021 International Workshops*, Lecture Notes in Computer Science, pages 227–232, Berlin, Heidelberg, 2021. Springer. ISBN 978-3-662-63958-0. doi: 10.1007/978-3-662-63958-0_20.
 - [20] Massimo Bartoletti, James Hsin-yu Chiang, and Alberto Lluch-Lafuente. A Theory of Automated Market Makers in DeFi. In Ferruccio Damiani and Ornella Dardha, editors, *Coordination Models and Languages*, Lecture Notes in Computer Science, pages 168–187, Cham, 2021. Springer International Publishing. ISBN 978-3-030-78142-2. doi: 10.1007/978-3-030-78142-2_11.

- [21] Andrej Bauer, Jason Gross, Peter LeFanu Lumsdaine, Michael Shulman, Matthieu Sozeau, and Bas Spitters. The HoTT library: A formalization of homotopy type theory in Coq. In *Proceedings of the 6th ACM SIGPLAN Conference on Certified Programs and Proofs*, pages 164–172, Paris France, January 2017. ACM. ISBN 978-1-4503-4705-1. doi: 10.1145/3018610.3018615. URL <https://dl.acm.org/doi/10.1145/3018610.3018615>.
- [22] Rob Behnke. Explained: The NowSwap Protocol Hack. <https://halborn.com/explained-the-nowswap-protocol-hack-september-2021/>, September 2021. Accessed July 2023.
- [23] Beosin. Beosin — Global Web3 Security Report 2022. https://medium.com/Beosin_com/beosin-global-web3-security-report-2022-7aa2e4bb13. Accessed July 2023.
- [24] Bruno Bernardo, Raphaël Cauderlier, Guillaume Claret, Arvid Jakobsson, Basile Pesin, and Julien Tesson. Making Tezos Smart Contracts More Reliable with Coq. In Tiziana Margaria and Bernhard Steffen, editors, *Leveraging Applications of Formal Methods, Verification and Validation: Applications*, Lecture Notes in Computer Science, pages 60–72, Cham, 2020. Springer International Publishing. ISBN 978-3-030-61467-6. doi: 10.1007/978-3-030-61467-6_5.
- [25] Bruno Bernardo, Raphaël Cauderlier, Zhenlei Hu, Basile Pesin, and Julien Tesson. Mi-Cho-Coq, a Framework for Certifying Tezos Smart Contracts. In Emil Sekerinski, Nelma Moreira, José N. Oliveira, Daniel Ratiu, Riccardo Guidotti, Marie Farrell, Matt Luckcuck, Diego Marmsoler, José Campos, Troy Astarte, Laure Gonnord, Antonio Cerone, Luis Couto, Brijesh Dongol, Martin Kutrib, Pedro Monteiro, and David Delmas, editors, *Formal Methods. FM 2019 International Workshops*, Lecture Notes in Computer Science, pages 368–379, Cham, 2020. Springer International Publishing. ISBN 978-3-030-54994-7. doi: 10.1007/978-3-030-54994-7_28.
- [26] Bruno Bernardo, Raphaël Cauderlier, Basile Pesin, and Julien Tesson. Albert, an intermediate smart-contract language for the Tezos blockchain. arXiv:2001.02630, January 2020.
- [27] Yves Bertot and Pierre Castéran. *Interactive Theorem Proving and Program Development: Coq’Art: The Calculus of Inductive Constructions*. Springer Science & Business Media, 2013.
- [28] Karthikeyan Bhargavan, Antoine Delignat-Lavaud, Cédric Fournet, Anitha Gollamudi, Georges Gonthier, Nadim Kobeissi, Natalia Kulatova, Aseem Rastogi, Thomas Sibut-Pinote, Nikhil Swamy, and Santiago Zanella-Béguelin. Formal Verification of Smart Contracts: Short Paper. In *Proceedings of the 2016 ACM Workshop on Programming Languages and Analysis for Security*, PLAS ’16, pages 91–96, New York, NY, USA, October 2016. Association for Computing Machinery. ISBN 978-1-4503-4574-3. doi: 10.1145/2993600.2993611.
- [29] Giancarlo Bigi, Andrea Bracciali, Giovanni Meacci, and Emilio Tuosto. Validation of Decentralised Smart Contracts Through Game Theory and Formal Methods. In Chiara Bodei, Gianluigi Ferrari, and Corrado Priami, editors, *Programming Languages with Applications to Biology and Security: Essays Dedicated to Pierpaolo Degano on the Occasion of His 65th Birthday*, Lecture Notes in

- Computer Science, pages 142–161. Springer International Publishing, Cham, 2015. ISBN 978-3-319-25527-9. doi: 10.1007/978-3-319-25527-9_11.
- [30] BscScan.com. Pancake Bunny Exploiter. Address 0x158c244b62058330f2c328c720b072d8db2c612f, 2021.
 - [31] BscScan.com. Spartan Protocol Exploit. Transaction 0xb64ae25b0d836c25d115a9368319902c972a0215bd108ae17b1b9617dfb93af8, 2021.
 - [32] BscScan.com. Uranium Finance Exploiter. Address 0x2b528a28451e9853f51616f3b0f6d82af8bea6ae, 2021.
 - [33] Vitalik Buterin. Improving front running resistance of $x*y=k$ market makers - Decentralized exchanges. <https://ethresear.ch/t/improving-front-running-resistance-of-x-y-k-market-makers/1281>, March 2018. Accessed July 2023.
 - [34] Raphael Cauderlier. Dexter2 Specification (Mi-Cho-Coq). https://gitlab.com/nomadic-labs/mi-cho-coq/-/blob/dexter-verification/src/contracts.coq/dexter_spec.v. Accessed July 2023.
 - [35] COBRA Research Center. ConCert. <https://github.com/AU-COBRA/ConCert>. Accessed July 2023.
 - [36] Martán Ceresa and César Sánchez. Multi: A Formal Playground for Multi-Smart Contract Interaction. arXiv:2207.06681, July 2022.
 - [37] James Chapman, Roman Kireev, Chad Nester, and Philip Wadler. System F in Agda, for Fun and Profit. In Graham Hutton, editor, *Mathematics of Program Construction*, Lecture Notes in Computer Science, pages 255–297, Cham, 2019. Springer International Publishing. ISBN 978-3-030-33636-3. doi: 10.1007/978-3-030-33636-3_10.
 - [38] Huashan Chen, Marcus Pendleton, Laurent Njilla, and Shouhuai Xu. A survey on ethereum systems security: Vulnerabilities, attacks, and defenses. *ACM Computing Surveys (CSUR)*, 53(3):1–43, 2020.
 - [39] Felix Cherubini, Thierry Coquand, and Matthias Hutzler. A Foundation for Synthetic Algebraic Geometry.
 - [40] Michele Ciampi, Muhammad Ishaq, Malik Magdon-Ismail, Rafail Ostrovsky, and Vassilis Zikas. FairMM: A Fast and Frontrunning-Resistant Crypto Market-Maker. In Shlomi Dolev, Jonathan Katz, and Amnon Meisels, editors, *Cyber Security, Cryptology, and Machine Learning*, Lecture Notes in Computer Science, pages 428–446, Cham, 2022. Springer International Publishing. ISBN 978-3-031-07689-3. doi: 10.1007/978-3-031-07689-3_31.
 - [41] Simon Cousaert, Jiahua Xu, and Toshiko Matsui. SoK: Yield Aggregators in DeFi. In *2022 IEEE International Conference on Blockchain and Cryptocurrency (ICBC)*, pages 1–14, May 2022. doi: 10.1109/ICBC54727.2022.9805523.

- [42] Luís Pedro Arrojado da Horta, João Santos Reis, Simão Melo de Sousa, and Mário Pereira. A tool for proving michelson smart contracts in why3. In *2020 IEEE International Conference on Blockchain (Blockchain)*, pages 409–414. IEEE, 2020.
- [43] Philip Daian, Steven Goldfeder, Tyler Kell, Yunqi Li, Xueyuan Zhao, Iddo Bentov, Lorenz Breidenbach, and Ari Juels. Flash Boys 2.0: Frontrunning in Decentralized Exchanges, Miner Extractable Value, and Consensus Instability. In *2020 IEEE Symposium on Security and Privacy (SP)*, pages 910–927, May 2020. doi: 10.1109/SP40000.2020.00040.
- [44] John W. Dawson. Why do mathematicians re-prove theorems? *Philosophia Mathematica*, 14(3): 269–286, 2006. doi: 10.1093/phimat/nkl009.
- [45] Monika Di Angelo and Gernot Salzer. A survey of tools for analyzing ethereum smart contracts. In *2019 IEEE International Conference on Decentralized Applications and Infrastructures (DAPPCON)*, pages 69–78. IEEE, 2019.
- [46] Andres Diaz-Valdivia and Marta Poblet. Governance of ReFi Ecosystem and the Integrity of Voluntary Carbon Markets as a Common Resource. doi: 10.2139/ssrn.4286167, November 2022.
- [47] Thomas Dickerson, Paul Gazzillo, Maurice Herlihy, Vikram Saraph, and Eric Koskinen. Proof-Carrying Smart Contracts. In Aviv Zohar, Ittay Eyal, Vanessa Teague, Jeremy Clark, Andrea Bracciali, Federico Pintore, and Massimiliano Sala, editors, *Financial Cryptography and Data Security*, Lecture Notes in Computer Science, pages 325–338, Berlin, Heidelberg, 2019. Springer. ISBN 978-3-662-58820-8. doi: 10.1007/978-3-662-58820-8_22.
- [48] Xiaoqun Dong, Rachel Chi Kiu Mok, Durreh Tabassum, Pierre Guigon, Eduardo Ferreira, Chandra Shekhar Sinha, Neeraj Prasad, Joe Madden, Tom Baumann, Jason Libersky, Eamonn McCormick, and Jefferson Cohen. Blockchain and emerging digital technologies for enhancing post-2020 climate markets. <https://tinyurl.com/bdz5wczb>.
- [49] Gregor Dorfleitner, Franziska Muck, and Isabel Scheckenbach. Blockchain applications for climate protection: A global empirical investigation. *Renewable and Sustainable Energy Reviews*, 149: 111378, October 2021. ISSN 1364-0321. doi: 10.1016/j.rser.2021.111378.
- [50] Christoph Dorn and Christopher L. Douglas. Framed combinatorial topology, December 2021. URL <http://arxiv.org/abs/2112.14700>.
- [51] Youssef El Faqir, Javier Arroyo, and Samer Hassan. An overview of decentralized autonomous organizations on the blockchain. In *Proceedings of the 16th International Symposium on Open Collaboration*, OpenSym 2020, pages 1–8, New York, NY, USA, August 2020. Association for Computing Machinery. ISBN 978-1-4503-8779-8. doi: 10.1145/3412569.3412579.
- [52] etherscan.io. Harvest Finance Hacker. Address 0xf224ab004461540778a914ea397c589b677e27bb, 2020.

- [53] etherscan.io. CREAM Finance Exploit.
Transaction 0x0fe2542079644e107cbf13690eb9c2c65963ccb79089ff96bfaf8dced2331c92, 2021.
- [54] etherscan.io. Beanstalk Exploit.
Transaction 0xcd314668aaa9bbfebaf1a0bd2b6553d01dd58899c508d4729fa7311dc5d33ad7, 2022.
- [55] etherscan.io. Nomad Bridge Exploit.
Transaction 0xa5fe9d044e4f3e5aa5bc4c0709333cd2190cba0f4e7f16bcf73f49f83e4a5460, 2022.
- [56] Alex Evans, Guillermo Angeris, and Tarun Chitra. Optimal Fees for Geometric Mean Market Makers. In Matthew Bernhard, Andrea Bracciali, Lewis Gudgeon, Thomas Haines, Arian Klages-Mundt, Shin’ichiro Matsuo, Daniel Perez, Massimiliano Sala, and Sam Werner, editors, *Financial Cryptography and Data Security. FC 2021 International Workshops*, Lecture Notes in Computer Science, pages 65–79, Berlin, Heidelberg, 2021. Springer. ISBN 978-3-662-63958-0. doi: 10.1007/978-3-662-63958-0_6.
- [57] Beanstalk Farms. Beanstalk Governance Exploit. <https://bean.money/blog/beanstalk-governance-exploit>. Accessed July 2023.
- [58] Harvest Finance. Harvest Flashloan Economic Attack Post-Mortem. <https://medium.com/harvest-finance/harvest-flashloan-economic-attack-post-mortem-3cf900d65217>, October 2020.
- [59] Uranium Finance. Uranium Finance Exploit. <https://uraniumfinance.medium.com/exploit-d3a88921531c>, April 2021. Accessed July 2023.
- [60] Savannah Fortis. DeFi-type projects received the highest number of attacks in 2022: Report. <https://cointelegraph.com/news/defi-type-projects-received-the-highest-number-of-attacks-in-2022-report>. Accessed July 2023.
- [61] Robin Fritsch and Roger Wattenhofer. A Note on Optimal Fees for Constant Function Market Makers. *DeFi@CCS*, 2021. doi: 10.1145/3464967.3488589.
- [62] Robin Fritsch, Samuel Käser, and Roger Wattenhofer. The economics of automated market makers. In *Proceedings of the 4th ACM Conference on Advances in Financial Technologies*, pages 102–110, 2022.
- [63] Ikram Garfatta, Kais Klai, Walid Gaaloul, and Mohamed Graiet. A Survey on Formal Verification for Solidity Smart Contracts. In *2021 Australasian Computer Science Week Multiconference, ACSW ’21*, pages 1–10, New York, NY, USA, February 2021. Association for Computing Machinery. ISBN 978-1-4503-8956-3. doi: 10.1145/3437378.3437879.
- [64] Neil Ghani, Jules Hedges, Viktor Winschel, and Philipp Zahn. Compositional game theory. arXiv:1603.04641, February 2018.
- [65] Florian Gronde. Flash loans and decentralized lending protocols: An in-depth analysis. Master’s thesis, Center for Innovative Finance, University of Basel Basel, Switzerland, 2020.

- [66] Lewis Gudgeon, Daniel Perez, Dominik Harz, Benjamin Livshits, and Arthur Gervais. The Decentralized Financial Crisis. In *2020 Crypto Valley Conference on Blockchain Technology (CVCBT)*, pages 1–15, June 2020. doi: 10.1109/CVCBT50464.2020.00005.
- [67] Samuel Haig. PancakeBunny tanks 96% following \$200M flash loan exploit. <https://cointelegraph.com/news/pancakebunny-tanks-96-following-200m-flash-loan-exploit>, May 2021.
- [68] Robert Harper. The Holy Trinity (2011). <https://existentialtype.wordpress.com/2011/03/27/the-holy-trinity/>. Accessed August 2023.
- [69] Lioba Heimbach, Ye Wang, and Roger Wattenhofer. Behavior of Liquidity Providers in Decentralized Exchanges. arXiv:2105.13822, October 2021.
- [70] Lioba Heimbach, Eric Schertenleib, and Roger Wattenhofer. Risks and Returns of Uniswap V3 Liquidity Providers. In *Proceedings of the 4th ACM Conference on Advances in Financial Technologies*, pages 89–101, September 2022. doi: 10.1145/3558535.3559772.
- [71] Matthew Hennessy and Robin Milner. Algebraic laws for nondeterminism and concurrency. *J. ACM*, 32(1):137–161, January 1985. ISSN 0004-5411. doi: 10.1145/2455.2460.
- [72] Michael Hicks and Scott Nettles. Dynamic software updating. *ACM Trans. Program. Lang. Syst.*, 27(6):1049–1096, November 2005. ISSN 0164-0925. doi: 10.1145/1108970.1108971.
- [73] Everett Hildenbrandt, Manasvi Saxena, Nishant Rodrigues, Xiaoran Zhu, Philip Daian, Dwight Guth, Brandon Moore, Daejun Park, Yi Zhang, Andrei Stefanescu, and Grigore Rosu. KEVM: A Complete Formal Semantics of the Ethereum Virtual Machine. In *2018 IEEE 31st Computer Security Foundations Symposium (CSF)*, pages 204–217, July 2018. doi: 10.1109/CSF.2018.00022.
- [74] Yoichi Hirai. Defining the Ethereum Virtual Machine for Interactive Theorem Provers. In Michael Brenner, Kurt Rohloff, Joseph Bonneau, Andrew Miller, Peter Y.A. Ryan, Vanessa Teague, Andrea Bracciali, Massimiliano Sala, Federico Pintore, and Markus Jakobsson, editors, *Financial Cryptography and Data Security*, volume 10323, pages 520–535. Springer International Publishing, Cham, 2017. ISBN 978-3-319-70277-3 978-3-319-70278-0. doi: 10.1007/978-3-319-70278-0_33.
- [75] Dale Husemoller. *Fibre Bundles*, volume 20 of *Graduate Texts in Mathematics*. Springer, New York, NY, 1994. ISBN 978-1-4757-2263-5 978-1-4757-2261-1. doi: 10.1007/978-1-4757-2261-1.
- [76] Igor Igamberdiev (@FrankResearcher). BUNNY Exploit Report. <https://twitter.com/FrankResearcher/status/1395196961108774915>, May 2021. Accessed July 2023.
- [77] Immunefi. Hack Analysis: Cream Finance Oct 2021. <https://medium.com/immunefi/hack-analysis-cream-finance-oct-2021-fc222d913fc5>, November 2022.

- [78] ImmuneFi. Hack Analysis: Nomad Bridge, August 2022. <https://medium.com/immuneFi/hack-analysis-nomad-bridge-august-2022-5aa63d53814a>, January 2023.
- [79] PeckShield Inc. The Spartan Incident: Root Cause Analysis. <https://peckshield-94632.medium.com/the-spartan-incident-root-cause-analysis-b14135d3415f>, May 2021. Accessed July 2023.
- [80] Runtime Verification Inc. Dexter2 Specification (K Framework). <https://github.com/runtime-verification/michelson-semantics/blob/a46be4a542e01b17a93134395c889df1468a067b/tests/proofs/dexter/dexter-spec.md>. Accessed July 2023.
- [81] Runtime Verification Inc. K-Michelson: A Michelson Semantics. <https://github.com/runtime-verification/michelson-semantics>, August 2022. Accessed July 2023.
- [82] 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 (SP)*, pages 1695–1712. IEEE, 2020.
- [83] Roman Kozhan and Ganesh Viswanath-Natraj. Fundamentals of the MakerDAO Governance Token. In *3rd International Conference on Blockchain Economics, Security and Protocols (Tokenomics 2021)*. Schloss Dagstuhl-Leibniz-Zentrum für Informatik, 2022.
- [84] Ximeng Li, Zhiping Shi, Qianying Zhang, Guohui Wang, Yong Guan, and Ning Han. Towards Verifying Ethereum Smart Contracts at Intermediate Language Level. In Yamine Ait-Ameur and Shengchao Qin, editors, *Formal Methods and Software Engineering*, Lecture Notes in Computer Science, pages 121–137, Cham, 2019. Springer International Publishing. ISBN 978-3-030-32409-4. doi: 10.1007/978-3-030-32409-4_8.
- [85] Loi Luu, Duc-Hiep Chu, Hrishi Olickel, Prateek Saxena, and Aquinas Hobor. Making Smart Contracts Smarter. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security, CCS '16*, pages 254–269, New York, NY, USA, October 2016. Association for Computing Machinery. ISBN 978-1-4503-4139-4. doi: 10.1145/2976749.2978309.
- [86] Gabor Madl, Luis Bathen, German Flores, and Divyesh Jadav. Formal Verification of Smart Contracts Using Interface Automata. In *2019 IEEE International Conference on Blockchain (Blockchain)*, pages 556–563, July 2019. doi: 10.1109/Blockchain.2019.00081.
- [87] Shaurya Malwa. How Market Manipulation Led to a \$100M Exploit on Solana DeFi Exchange Mango. <https://www.coindesk.com/markets/2022/10/12/how-market-manipulation-led-to-a-100m-exploit-on-solana-defi-exchange-mango/>, October 2022.
- [88] Alessio Mansutti, Marino Miculan, and Marco Peressotti. Multi-agent Systems Design and Prototyping with Bigraphical Reactive Systems. In Kostas Magoutis and Peter Pietzuch, editors, *Distributed Applications and Interoperable Systems*, volume 8460, pages 201–208. Springer

- Berlin Heidelberg, Berlin, Heidelberg, 2014. ISBN 978-3-662-43351-5 978-3-662-43352-2. doi: 10.1007/978-3-662-43352-2_16.
- [89] Aaron Marback, Hyunsook Do, Ke He, Samuel Kondamarri, and Dianxiang Xu. A threat model-based approach to security testing. *Software: Practice and Experience*, 43(2):241–258, 2013. ISSN 1097-024X. doi: 10.1002/spe.2111.
 - [90] Anastasia Mavridou and Aron Laszka. Designing secure ethereum smart contracts: A finite state machine based approach. In *Financial Cryptography and Data Security: 22nd International Conference, FC 2018, Nieuwpoort, Curaçao, February 26–March 2, 2018, Revised Selected Papers 22*, pages 523–540. Springer, 2018.
 - [91] Anastasia Mavridou and Aron Laszka. Tool Demonstration: FSolidM for Designing Secure Ethereum Smart Contracts. In Lujo Bauer and Ralf Küsters, editors, *Principles of Security and Trust*, Lecture Notes in Computer Science, pages 270–277, Cham, 2018. Springer International Publishing. ISBN 978-3-319-89722-6. doi: 10.1007/978-3-319-89722-6_11.
 - [92] Anastasia Mavridou, Aron Laszka, Emmanouela Stachtari, and Abhishek Dubey. VeriSolid: Correct-by-Design Smart Contracts for Ethereum. In Ian Goldberg and Tyler Moore, editors, *Financial Cryptography and Data Security*, Lecture Notes in Computer Science, pages 446–465, Cham, 2019. Springer International Publishing. ISBN 978-3-030-32101-7. doi: 10.1007/978-3-030-32101-7_27.
 - [93] Robin Milner. The Bigraphical Model. <https://www.cl.cam.ac.uk/archive/rm135/uam-theme.html>. Accessed July 2023.
 - [94] Robin Milner. Bigraphs and Their Algebra. *Electronic Notes in Theoretical Computer Science*, 209: 5–19, April 2008. ISSN 1571-0661. doi: 10.1016/j.entcs.2008.04.002.
 - [95] Nick Mudge. EIP-2535: Diamonds, Multi-Facet Proxy. <https://eips.ethereum.org/EIPS/eip-2535>. Accessed July 2023.
 - [96] James R. Munkres. *Topology*. Prentice Hall, Incorporated, 2000. ISBN 978-0-13-181629-9.
 - [97] Yvonne Murray and David A. Anisi. Survey of Formal Verification Methods for Smart Contracts on Blockchain. In *2019 10th IFIP International Conference on New Technologies, Mobility and Security (NTMS)*, pages 1–6, June 2019. doi: 10.1109/NTMS.2019.8763832.
 - [98] ncatlab.org. Computational trinity. <https://ncatlab.org/nlab/show/computational+trilogy>. Accessed August 2023.
 - [99] George C. Necula. Proof-carrying code. In *Proceedings of the 24th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, POPL ’97, pages 106–119, New York, NY, USA, January 1997. Association for Computing Machinery. ISBN 978-0-89791-853-4. doi: 10.1145/263699.263712.

- [100] Zeinab Nehai and François Bobot. Deductive Proof of Ethereum Smart Contracts Using Why3. arXiv:1904.11281, August 2019.
- [101] Zeinab Nehai, François Bobot, Sara Tucci-Piergiovanni, Carole Delporte-Gallet, and Hugues Fauconnier. A TLA+ Formal Proof of a Cross-Chain Swap. In *23rd International Conference on Distributed Computing and Networking, ICDCN 2022*, pages 148–159, New York, NY, USA, January 2022. Association for Computing Machinery. ISBN 978-1-4503-9560-1. doi: 10.1145/3491003.3491006.
- [102] Eske Hoy Nielsen, Danil Annenkov, and Bas Spitters. Formalising Decentralised Exchanges in Coq. In *Proceedings of the 12th ACM SIGPLAN International Conference on Certified Programs and Proofs*, pages 290–302, 2023.
- [103] Jakob Botsch Nielsen and Bas Spitters. Smart Contract Interactions in Coq. In *FM Workshops (1)*, volume 12232 of *Lecture Notes in Computer Science*, pages 380–391. Springer, 2019.
- [104] Eric Nowak. Voluntary Carbon Markets. <https://tinyurl.com/k8sbhemf>, March 2022.
- [105] Russell O’Connor. Simplicity: A new language for blockchains. In *Proceedings of the 2017 Workshop on Programming Languages and Analysis for Security*, pages 107–120, 2017.
- [106] Trail of Bits. Contract Upgrade Anti-Patterns. <https://blog.trailofbits.com/2018/09/05/contract-upgrade-anti-patterns>, September 2018. Accessed July 2023.
- [107] Kris Oosthoek. Flash Crash for Cash: Cyber Threats in Decentralized Finance. arXiv:2106.10740, June 2021.
- [108] pancakebunny.finance (@PancakeBunnyFin). BUNNY Exploit Report. <https://twitter.com/PancakeBunnyFin/status/1395173389208334342>, May 2021. Accessed July 2023.
- [109] Daejun Park, Yi Zhang, Manasvi Saxena, Philip Daian, and Grigore Roşu. A formal verification tool for Ethereum VM bytecode. In *Proceedings of the 2018 26th ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering*, pages 912–915, Lake Buena Vista FL USA, October 2018. ACM. ISBN 978-1-4503-5573-5. doi: 10.1145/3236024.3264591.
- [110] Siraphob Phipathananunth. Using Mutations to Analyze Formal Specifications. In *Companion Proceedings of the 2022 ACM SIGPLAN International Conference on Systems, Programming, Languages, and Applications: Software for Humanity, SPLASH Companion 2022*, pages 81–83, New York, NY, USA, December 2022. Association for Computing Machinery. ISBN 978-1-4503-9901-2. doi: 10.1145/3563768.3563960.
- [111] Publius. Beanstalk: A Permissionless Fiat Stablecoin Protocol. <https://bean.money/beanstalk.pdf>. Accessed July 2023.

- [112] Kaihua Qin, Liyi Zhou, Benjamin Livshits, and Arthur Gervais. Attacking the DeFi Ecosystem with Flash Loans for Fun and Profit. In Nikita Borisov and Claudia Diaz, editors, *Financial Cryptography and Data Security*, Lecture Notes in Computer Science, pages 3–32, Berlin, Heidelberg, 2021. Springer. ISBN 978-3-662-64322-8. doi: 10.1007/978-3-662-64322-8_1.
- [113] Talia Ringer. *Proof Repair*. University of Washington, 2021.
- [114] Talia Ringer, Randair Porter, Nathaniel Yazdani, John Leo, and Dan Grossman. Proof repair across type equivalences. In *Proceedings of the 42nd ACM SIGPLAN International Conference on Programming Language Design and Implementation*, pages 112–127, 2021.
- [115] John Rushby. Theorem Proving for Verification. In Franck Cassez, Claude Jard, Brigitte Rozoy, and Mark Dermot Ryan, editors, *Modeling and Verification of Parallel Processes: 4th Summer School, MOVEP 2000 Nantes, France, June 19–23, 2000 Revised Tutorial Lectures*, Lecture Notes in Computer Science, pages 39–57. Springer, Berlin, Heidelberg, 2001. ISBN 978-3-540-45510-3. doi: 10.1007/3-540-45510-8_2.
- [116] @samczsun. Nomad Tweet Thread. <https://twitter.com/samczsun/status/1554252024723546112>, August 2022. Accessed July 2023.
- [117] Davide Sangiorgi. On the bisimulation proof method. *Mathematical Structures in Computer Science*, 8(5):447–479, October 1998. ISSN 1469-8072, 0960-1295. doi: 10.1017/S0960129598002527.
- [118] Ilya Sergey, Amrit Kumar, and Aquinas Hobor. Temporal Properties of Smart Contracts. In Tiziana Margaria and Bernhard Steffen, editors, *Leveraging Applications of Formal Methods, Verification and Validation. Industrial Practice*, Lecture Notes in Computer Science, pages 323–338, Cham, 2018. Springer International Publishing. ISBN 978-3-030-03427-6. doi: 10.1007/978-3-030-03427-6_25.
- [119] Michele Sevegnani and Muffy Calder. Bigraphs with sharing. *Theoretical Computer Science*, 577: 43–73, April 2015. ISSN 0304-3975. doi: 10.1016/j.tcs.2015.02.011.
- [120] Michele Sevegnani and Muffy Calder. BigraphER: Rewriting and analysis engine for bigraphs. In *Computer Aided Verification: 28th International Conference, CAV 2016, Toronto, ON, Canada, July 17-23, 2016, Proceedings, Part II* 28, pages 494–501. Springer, 2016.
- [121] Amritraj Singh, Reza M. Parizi, Qi Zhang, Kim-Kwang Raymond Choo, and Ali Dehghantanha. Blockchain smart contracts formalization: Approaches and challenges to address vulnerabilities. *Computers & Security*, 88:101654, January 2020. ISSN 0167-4048. doi: 10.1016/j.cose.2019.101654.
- [122] Adam Siphthorpe, Sabine Brink, Tyler Van Leeuwen, and Iain Staffell. Blockchain solutions for carbon markets are nearing maturity. *One Earth*, 5(7):779–791, July 2022. ISSN 2590-3330, 2590-3322. doi: 10.1016/j.oneear.2022.06.004.
- [123] solscan.io. Mango Markets Exploiter. Address CQvKSNnYtPTZfQRQ5jkHq8q2swJyRsdQLcFcj3EmKFfX, 2022.

- [124] Derek Sorensen. Tokenized carbon credits. *Preprint*, 2023. URL <https://derekhsorensen.com/docs/sorensen-tokenized-carbon-credits.pdf>.
- [125] Derek Sorensen. Structured Pools for Tokenized Carbon Credits. In *2023 IEEE International Conference on Blockchain and Cryptocurrency (ICBC)*, pages 1–6, May 2023. doi: 10.1109/ICBC56567.2023.10174866.
- [126] Matthieu Sozeau, Abhishek Anand, Simon Boulier, Cyril Cohen, Yannick Forster, Fabian Kunze, Gregory Malecha, Nicolas Tabareau, and Théo Winterhalter. The MetaCoq Project. *J Autom Reasoning*, 64(5):947–999, June 2020. ISSN 1573-0670. doi: 10.1007/s10817-019-09540-0.
- [127] Tianyu Sun and Wensheng Yu. A Formal Verification Framework for Security Issues of Blockchain Smart Contracts. *Electronics*, 9(2):255, February 2020. ISSN 2079-9292. doi: 10.3390/electronics9020255.
- [128] Keefer Taylor. Murmuration: A Generalizable DAO for Tezos. <https://github.com/Hover-Labs/murmuration>. Accessed March 2023.
- [129] Huang Teng, Wayne Tian, Haocheng Wang, and Zhiyuan Yang. Applications of the Decentralized Finance (DeFi) on the Ethereum. In *2022 IEEE Asia-Pacific Conference on Image Processing, Electronics and Computers (IPEC)*, pages 573–578. IEEE, 2022.
- [130] Palina Tolmach, Yi Li, Shang-Wei Lin, and Yang Liu. Formal Analysis of Composable DeFi Protocols. In *Financial Cryptography and Data Security. FC 2021 International Workshops*, Lecture Notes in Computer Science, pages 149–161, Berlin, Heidelberg, 2021. Springer. ISBN 978-3-662-63958-0. doi: 10.1007/978-3-662-63958-0_13.
- [131] Palina Tolmach, Yi Li, Shang-Wei Lin, Yang Liu, and Zengxiang Li. A Survey of Smart Contract Formal Specification and Verification. *ACM Comput. Surv.*, 54(7):148:1–148:38, July 2021. ISSN 0360-0300. doi: 10.1145/3464421.
- [132] P. Torr. Demystifying the threat modeling process. *IEEE Security & Privacy*, 3(5):66–70, September 2005. ISSN 1558-4046. doi: 10.1109/MSP.2005.119.
- [133] Toucan. Toucan Whitepaper. <https://docs.toucan.earth/>. Accessed March 2023.
- [134] Dabao Wang, Siwei Wu, Ziling Lin, Lei Wu, Xingliang Yuan, Yajin Zhou, Haoyu Wang, and Kui Ren. Towards A First Step to Understand Flash Loan and Its Applications in DeFi Ecosystem. In *Proceedings of the Ninth International Workshop on Security in Blockchain and Cloud Computing*, pages 23–28, May 2021. doi: 10.1145/3457977.3460301.
- [135] Sam Werner, Daniel Perez, Lewis Gudgeon, Aariah Klages-Mundt, Dominik Harz, and William Knottenbelt. SoK: Decentralized Finance (DeFi). In *Proceedings of the 4th ACM Conference on Advances in Financial Technologies, AFT ’22*, pages 30–46, New York, NY, USA, July 2023. Association for Computing Machinery. ISBN 978-1-4503-9861-9. doi: 10.1145/3558535.3559780.

- [136] Jiahua Xu, Krzysztof Paruch, Simon Cousaert, and Yebo Feng. SoK: Decentralized Exchanges (DEX) with Automated Market Maker (AMM) Protocols. *ACM Comput. Surv.*, 55(11):238:1–238:50, February 2023. ISSN 0360-0300. doi: 10.1145/3570639.
- [137] Zheng Yang and Hang Lei. Fether: An extensible definitional interpreter for smart-contract verifications in coq. *IEEE Access*, 7:37770–37791, 2019.
- [138] Zheng Yang and Hang Lei. Lolisa: Formal Syntax and Semantics for a Subset of the Solidity Programming Language in Mathematical Tool Coq. *Mathematical Problems in Engineering*, 2020: e6191537, December 2020. ISSN 1024-123X. doi: 10.1155/2020/6191537.
- [139] Zheng Yang, Hang Lei, and Weizhong Qian. A Hybrid Formal Verification System in Coq for Ensuring the Reliability and Security of Ethereum-Based Service Smart Contracts. *IEEE Access*, 8: 21411–21436, 2020. ISSN 2169-3536. doi: 10.1109/ACCESS.2020.2969437.
- [140] Joachim Zahnentferner, Dmytro Kaidalov, Jean-Frédéric Etienne, and Javier Díaz. Djed: A Formally Verified Crypto-Backed Autonomous Stablecoin Protocol. In *2023 IEEE International Conference on Blockchain and Cryptocurrency (ICBC)*, pages 1–9, May 2023. doi: 10.1109/ICBC56567.2023.10174901.
- [141] Open Zeppelin. Proxy Upgrade Pattern. <https://docs.openzeppelin.com/upgrades-plugins/1.x/proxies>, Accessed July 2023.
- [142] Open Zeppelin. Writing Upgradeable Contracts. <https://docs.openzeppelin.com/upgrades-plugins/1.x/writing-upgradeable>, Accessed July 2023.
- [143] Xiyue Zhang, Yi Li, and Meng Sun. Towards a Formally Verified EVM in Production Environment. *Coordination Models and Languages*, 12134:341–349, May 2020. doi: 10.1007/978-3-030-50029-0_21.
- [144] Zibin Zheng, Shaoan Xie, Hong-Ning Dai, Weili Chen, Xiangping Chen, Jian Weng, and Muhammad Imran. An overview on smart contracts: Challenges, advances and platforms. *Future Generation Computer Systems*, 105:475–491, April 2020. ISSN 0167-739X. doi: 10.1016/j.future.2019.12.019.
- [145] Liyi Zhou, Kaihua Qin, and Arthur Gervais. A2mm: Mitigating frontrunning, transaction reordering and consensus instability in decentralized exchanges. *arXiv preprint arXiv:2106.07371*, 2021.
- [146] Liyi Zhou, Xihan Xiong, Jens Ernstberger, Stefanos Chaliasos, Zhipeng Wang, Ye Wang, Kaihua Qin, Roger Wattenhofer, Dawn Song, and Arthur Gervais. SoK: Decentralized Finance (DeFi) Attacks, April 2023.
- [147] Jian Zhu, Kai Hu, Mamoun Filali, Jean-Paul Bodeveix, and Jean-Pierre Talpin. Formal Verification of Solidity contracts in Event-B, May 2020.
- [148] Weiqin Zou, David Lo, Pavneet Singh Kochhar, Xuan-Bach Dinh Le, Xin Xia, Yang Feng, Zhenyu Chen, and Baowen Xu. Smart Contract Development: Challenges and Opportunities. *IEEE Transactions on Software Engineering*, 47(10):2084–2106, October 2021. ISSN 1939-3520. doi: 10.1109/TSE.2019.2942301.

Appendix A

Proofs and Definitions of Chapter 4

A.1 Formal Specification

We include the typeclass specifications for each contract type and the specification predicate.

A.1.1 Typeclass Specifications

Listing A.1: The formalization of the typeclass `Setup_Spec`, which characterizes the storage type `Setup` of a structured pool contract.

```
1 Class Setup_Spec (T : Type) :=  
2   build_setup_spec {  
3     init_rates : T -> FMap token exchange_rate ;  
4     init_pool_token : T -> token ;  
5   }.
```

Listing A.2: The formalization of the typeclass `Msg_Spec` and its associated types, which characterize the storage type `Msg` of a structured pool contract.

```
1 Record pool_data := {  
2   token_pooled : token ;  
3   qty_pooled : N ; (* the qty of tokens to be pooled *)  
4 }.  
5  
6 Record unpool_data := {  
7   token_unpooled : token ;  
8   qty_unpooled : N ; (* the qty of pool tokens being turned in *)  
9 }.  
10  
11 Record trade_data := {  
12   token_in_trade : token ;
```

```

13   token_out_trade : token ;
14   qty_trade : N ; (* the qty of token_in going in *)
15 }.
16
17 Class Msg_Spec (T : Type) :=
18   build_msg_spec {
19     pool : pool_data -> T ;
20     unpool : unpool_data -> T ;
21     trade : trade_data -> T ;
22     (* any other potential entrypoints *)
23     other : other_entrypoint -> option T ;
24   }.

```

Listing A.3: The formalization of the typeclass `State_Spec` and its associated types, which characterize the storage type `State` of a structured pool contract.

```

1 Context { other_entrypoint : Type }.
2 Class State_Spec (T : Type) :=
3   build_state_spec {
4     (* the exchange rates *)
5     stor_rates : T -> FMap token exchange_rate ;
6     (* token balances *)
7     stor_tokens_held : T -> FMap token N ;
8     (* pool token data *)
9     stor_pool_token : T -> token ;
10    (* number of outstanding pool tokens *)
11    stor_outstanding_tokens : T -> N ;
12  }.

```

Listing A.4: The formalization of the typeclass `Error_Spec`, which characterizes the storage type `Error` of a structured pool contract.

```

1 Definition error : Type := N.
2 Class Error_Spec (T : Type) :=
3   build_error_type {
4     error_to_Error : error -> T ;
5   }.

```

A.1.2 The Formal Specification of a Structured Pool Contract

Listing A.5: The full, formal sepcification of a structured pool contract.

```
1 (* =====
2  * The Contract Specification 'is_structured_pool' :
3      We detail a list of propositions of a contract's behavior which must be
4      proven true of a given contract for it to be a correct structured pool contract.
5  * ===== *)
6
7  { Setup Msg State Error : Type }
8  { other_entrypoint : Type }
9  `{Serializable Setup} `{Serializable Msg} `{Serializable State} `{Serializable Error}.
10
11 (** Specification of the Msg type:
12  - A Pool entrypoint, whose interface is defined by the pool_data type
13  - An Unpool entrypoint, whose interface is defined by the unpool_data type
14  - A Trade entrypoint, whose interface is defined by the trade_data type
15  - Possibly other types
16 *)
17
18 Record pool_data := {
19   token_pooled : token ;
20   qty_pooled : N ; (* the qty of tokens to be pooled *)
21 }.
22
23 Record unpool_data := {
24   token_unpooled : token ;
25   qty_unpooled : N ; (* the qty of pool tokens being turned in *)
26 }.
27
28 Record trade_data := {
29   token_in_trade : token ;
30   token_out_trade : token ;
31   qty_trade : N ; (* the qty of token_in going in *)
32 }.
33
34 (* The Msg typeclass *)
35 Class Msg_Spec (T : Type) :=
36   build_msg_spec {
37     pool : pool_data -> T ;
38     unpool : unpool_data -> T ;
39     trade : trade_data -> T ;
40     (* any other potential entrypoints *)
41     other : other_entrypoint -> option T ;
42   }.
43
44
45
46
```

```

47 (** Specification of the State type:
48     The contract state keeps track of:
49     - the exchange rates
50     - tokens held
51     - pool token address
52     - number of outstanding pool tokens
53 *)
54 Class State_Spec (T : Type) :=
55   build_state_spec {
56     (* the exchange rates *)
57     stor_rates : T -> FMap token exchange_rate ;
58     (* token balances *)
59     stor_tokens_held : T -> FMap token N ;
60     (* pool token data *)
61     stor_pool_token : T -> token ;
62     (* number of outstanding pool tokens *)
63     stor_outstanding_tokens : T -> N ;
64   }.
65
66 (** Specification of the Setup type:
67     To initialize the contract, we need:
68     - the initial rates
69     - the pool token
70 *)
71 Class Setup_Spec (T : Type) :=
72   build_setup_spec {
73     init_rates : T -> FMap token exchange_rate ;
74     init_pool_token : T -> token ;
75   }.
76
77 (* specification of the Error type *)
78 Class Error_Spec (T : Type) :=
79   build_error_type {
80     error_to_Error : error -> T ;
81   }.
82
83 (* we assume that our contract types satisfy the type specifications *)
84 Context `{Msg_Spec Msg}  `{Setup_Spec Setup}  `{State_Spec State}  `{Error_Spec Error}.
85
86 (* First, we assume that all successful calls require a message *)
87 Definition none_fails (contract : Contract Setup Msg State Error) : Prop :=
88   forall cstate chain ctx,
89     (* the receive function returns an error if the token to be pooled is not in the
90        rates map held in the storage (=> is not in the semi-fungible family) *)
91     exists err : Error,
92     receive contract chain ctx cstate None = Err err.
93
94
95

```

```

96 (* We also specify that the Msg type is fully characterized by its typeclass *)
97 Definition msg_destruct (contract : Contract Setup Msg State Error) : Prop :=
98   forall (m : Msg),
99     (exists p, m = pool p) /\
100     (exists u, m = unpool u) /\
101     (exists t, m = trade t) /\
102     (exists o, Some m = other o).
103
104
105 (** Specification of the POOL entrypoint *)
106
107 (* A successful call to POOL means that token_pooled has an exchange rate (=> is in T) *)
108 Definition pool_entrypoint_check (contract : Contract Setup Msg State Error) : Prop :=
109   forall cstate cstate' chain ctx msg_payload acts,
110     (* a successful call *)
111     receive contract chain ctx cstate (Some (pool (msg_payload))) = Ok(cstate', acts) ->
112     (* an exchange rate exists *)
113     exists r_x,
114     FMap.find msg_payload.(token_pooled) (stor_rates cstate) = Some r_x.
115
116 (* When the POOL entrypoint is successfully called, it emits a TRANSFER call to the
117    token in storage, with q tokens in the payload of the call *)
118 Definition pool_emits_txns (contract : Contract Setup Msg State Error) : Prop :=
119   forall cstate chain ctx msg_payload cstate' acts,
120     (* the call to POOL was successful *)
121     receive contract chain ctx cstate (Some (pool (msg_payload))) = Ok(cstate', acts) ->
122     (* in the acts list there is a transfer call with q tokens as the payload *)
123     exists transfer_to transfer_data transfer_payload mint_data mint_payload,
124     (* there is a transfer call *)
125     let transfer_call := (act_call
126       (* calls the pooled token address *)
127       (msg_payload.(token_pooled).(token_address))
128       (* with amount 0 *)
129       0
130       (* and payload transfer_payload *)
131       (serialize (FA2Spec.Transfer transfer_payload))) in
132     (* with a transfer in it *)
133     In transfer_data transfer_payload /\
134     (* which itself has transfer data *)
135     In transfer_to transfer_data.(FA2Spec.txs) /\
136     (* whose quantity is the quantity pooled *)
137     transfer_to.(FA2Spec.amount) = msg_payload.(qty_pooled) /\
138     (* there is a mint call in acts *)
139     let mint_call := (act_call
140       (* calls the pool token contract *)
141       (stor_pool_token cstate).(token_address)
142       (* with amount 0 *)
143       0
144       (* and payload mint_payload *)

```

```

145     (serialize (FA2Spec.Mint mint_payload))) in
146   (* with has mint_data in the payload *)
147   In mint_data mint_payload /\
148   (* and the mint data has these properties: *)
149   let r_x := get_rate msg_payload.(token_pooled) (stor_rates cstate) in
150   mint_data.(FA2Spec.qty) = msg_payload.(qty_pooled) * r_x /\
151   mint_data.(FA2Spec.mint_owner) = ctx.(ctx_from) /\
152   (* these are the only emitted transactions *)
153   (acts = [ transfer_call ; mint_call ] \/
154   acts = [ mint_call ; transfer_call ]).
155
156 (* When the POOL entrypoint is successfully called, tokens_held goes up appropriately *)
157 Definition pool_increases_tokens_held
158   (contract : Contract Setup Msg State Error) : Prop :=
159   forall cstate chain ctx msg_payload cstate' acts,
160   (* the call to POOL was successful *)
161   receive contract chain ctx cstate (Some (pool msg_payload)) = Ok(cstate', acts) ->
162   (* in cstate', tokens_held has increased at token *)
163   let token := msg_payload.(token_pooled) in
164   let qty := msg_payload.(qty_pooled) in
165   let old_bal := get_bal token (stor_tokens_held cstate) in
166   let new_bal := get_bal token (stor_tokens_held cstate') in
167   new_bal = old_bal + qty /\
168   forall t,
169   t <> token ->
170   get_bal t (stor_tokens_held cstate) =
171   get_bal t (stor_tokens_held cstate').
172
173 (* And the rates don't change *)
174 Definition pool_rates_unchanged (contract : Contract Setup Msg State Error) : Prop :=
175   forall cstate cstate' chain ctx msg_payload acts,
176   (* the call to POOL was successful *)
177   receive contract chain ctx cstate (Some (pool msg_payload)) = Ok(cstate', acts) ->
178   (* rates all stay the same *)
179   forall t,
180   FMap.find t (stor_rates cstate) = FMap.find t (stor_rates cstate').
181
182 (* The outstanding tokens change appropriately *)
183 Definition pool_outstanding (contract : Contract Setup Msg State Error) : Prop :=
184   forall cstate cstate' chain ctx msg_payload acts,
185   (* the call to POOL was successful *)
186   receive contract chain ctx cstate (Some (pool msg_payload)) = Ok(cstate', acts) ->
187   (* rates all stay the same *)
188   let rate_in := get_rate msg_payload.(token_pooled) (stor_rates cstate) in
189   let qty := msg_payload.(qty_pooled) in
190   stor_outstanding_tokens cstate' =
191   stor_outstanding_tokens cstate + rate_in * qty.
192
193

```

```

194 (** Specification of the UNPOOL entrypoint *)
195
196 (* We assume an inverse rate function *)
197 Context { calc_rx_inv : forall (r_x : N) (q : N), N }.
198
199 (*A successful call to UNPOOL means that token_pooled has an exchange rate (=> is in T)*)
200 Definition unpool_entrypoint_check (contract : Contract Setup Msg State Error) : Prop :=
201   forall cstate cstate' chain ctx msg_payload acts,
202     (* a successful call *)
203     receive contract chain ctx cstate (Some (unpool (msg_payload))) = Ok(cstate', acts) ->
204     (* an exchange rate exists *)
205     exists r_x,
206     FMap.find msg_payload.(token_unpooled) (stor_rates cstate) = Some r_x.
207
208 Definition unpool_entrypoint_check_2 (contract : Contract Setup Msg State Error) : Prop :=
209   forall cstate cstate' chain ctx msg_payload acts,
210     (* a successful call *)
211     receive contract chain ctx cstate (Some (unpool (msg_payload))) = Ok(cstate', acts) ->
212     (* we don't unpool more than we have in reserves *)
213     qty_unpooled msg_payload <=
214     get_rate (token_unpooled msg_payload) (stor_rates cstate) *
215     get_bal (token_unpooled msg_payload) (stor_tokens_held cstate).
216
217 (* When the UNPOOL entrypoint is successfully called, it emits a BURN call to the
218    pool_token, with q in the payload *)
219 Definition unpool_emits_txns (contract : Contract Setup Msg State Error) : Prop :=
220   forall cstate chain ctx msg_payload cstate' acts,
221     (* the call to UNPOOL was successful *)
222     receive contract chain ctx cstate (Some (unpool msg_payload)) = Ok(cstate', acts) ->
223     (* in the acts list there are burn and transfer transactions *)
224     exists burn_data burn_payload transfer_to transfer_data transfer_payload,
225     (* there is a burn call in acts *)
226     let burn_call := (act_call
227       (* calls the pool token address *)
228       (stor_pool_token cstate).(token_address)
229       (* with amount 0 *)
230       0
231       (* with payload burn_payload *)
232       (serialize (FA2Spec.Retire burn_payload))) in
233     (* with has burn_data in the payload *)
234     In burn_data burn_payload /\
235     (* and burn_data has these properties: *)
236     burn_data.(FA2Spec.retire_amount) = msg_payload.(qty_unpooled) /\
237     (* the burned tokens go from the unpooler *)
238     burn_data.(FA2Spec.retiring_party) = ctx.(ctx_from) /\
239     (* there is a transfer call *)
240     let transfer_call := (act_call
241       (* call to the token address *)
242       (msg_payload.(token_unpooled).(token_address))

```

```

243     (* with amount = 0 *)
244     0
245     (* with payload transfer_payload *)
246     (serialize (FA2Spec.Transfer transfer_payload))) in
247   (* with a transfer in it *)
248   In transfer_data transfer_payload /\
249   (* which itself has transfer data *)
250   In transfer_to transfer_data.(FA2Spec.txs) /\
251   (* whose quantity is the quantity pooled *)
252   let r_x := get_rate msg_payload.(token_unpooled) (stor_rates cstate) in
253   transfer_to.(FA2Spec.amount) = msg_payload.(qty_unpooled) / r_x /\
254   (* and these are the only emitted transactions *)
255   (acts = [ burn_call ; transfer_call ] \/
256    acts = [ transfer_call ; burn_call ]).
257
258
259 (*When the UNPOOL entrypoint is successfully called, tokens_held goes down appropriately*)
260 Definition unpool_decreases_tokens_held
261   (contract : Contract Setup Msg State Error) : Prop :=
262   forall cstate chain ctx msg_payload cstate' acts,
263   (* the call to POOL was successful *)
264   receive contract chain ctx cstate (Some (unpool msg_payload)) = Ok(cstate', acts) ->
265   (* in cstate', tokens_held has increased at token *)
266   let token := msg_payload.(token_unpooled) in
267   let r_x := get_rate token (stor_rates cstate) in
268   let qty := calc_rx_inv r_x msg_payload.(qty_unpooled) in
269   let old_bal := get_bal token (stor_tokens_held cstate) in
270   let new_bal := get_bal token (stor_tokens_held cstate') in
271   new_bal = old_bal - qty /\
272   forall t,
273   t <> token ->
274   get_bal t (stor_tokens_held cstate) =
275   get_bal t (stor_tokens_held cstate').
276
277 (*When the UNPOOL entrypoint is successfully called, tokens_held goes down appropriately*)
278 Definition unpool_rates_unchanged (contract : Contract Setup Msg State Error) : Prop :=
279   forall cstate cstate' chain ctx msg_payload acts,
280   (* the call to POOL was successful *)
281   receive contract chain ctx cstate (Some (unpool msg_payload)) = Ok(cstate', acts) ->
282   (* rates all stay the same *)
283   forall t,
284   FMap.find t (stor_rates cstate) = FMap.find t (stor_rates cstate').
285
286 (* Defines how the UNPOOL entrypoint updates outstanding tokens *)
287 Definition unpool_outstanding (contract : Contract Setup Msg State Error) : Prop :=
288   forall cstate cstate' chain ctx msg_payload acts,
289   (* the call to POOL was successful *)
290   receive contract chain ctx cstate (Some (unpool msg_payload)) = Ok(cstate', acts) ->
291   (* rates all stay the same *)

```



```

292   let rate_in := get_rate msg_payload.(token_unpooled) (stor_rates cstate) in
293   let qty := msg_payload.(qty_unpooled) in
294   stor_outstanding_tokens cstate' =
295     stor_outstanding_tokens cstate - qty.
296
297 (* A successful call to TRADE means that token_in_trade and token_out_trade have exchange
298    rates (=> are in T) *)
299 Definition trade_entrypoint_check (contract : Contract Setup Msg State Error) : Prop :=
300   forall cstate chain ctx msg_payload cstate' acts,
301     (* a successful call *)
302     receive contract chain ctx cstate (Some (trade (msg_payload))) = Ok(cstate', acts) ->
303     (* exchange rates exist *)
304     exists y r_x r_y,
305       (FMap.find msg_payload.(token_out_trade) (stor_tokens_held cstate) = Some y) /\
306       (FMap.find msg_payload.(token_in_trade) (stor_rates cstate) = Some r_x) /\
307       (FMap.find msg_payload.(token_out_trade) (stor_rates cstate) = Some r_y)).
308
309 (* A successful call to TRADE means that the inputs to the trade calculation are
310    all positive *)
311 Definition trade_entrypoint_check_2 (contract : Contract Setup Msg State Error) : Prop :=
312   forall cstate chain ctx msg_payload cstate' acts,
313     (* a successful call *)
314     receive contract chain ctx cstate (Some (trade (msg_payload))) = Ok(cstate', acts) ->
315     (* exchange rates exist *)
316     exists x r_x r_y k,
317       (FMap.find msg_payload.(token_in_trade) (stor_tokens_held cstate) = Some x) /\
318       (FMap.find msg_payload.(token_in_trade) (stor_rates cstate) = Some r_x) /\
319       (FMap.find msg_payload.(token_out_trade) (stor_rates cstate) = Some r_y) /\
320       (stor_outstanding_tokens cstate = k) /\
321       r_x > 0 /\ r_y > 0 /\ x > 0 /\ k > 0).
322
323 (* Specification of the TRADE entrypoint *)
324 (* We assume the existence of two functions *)
325 Context
326   { calc_delta_y : forall (rate_in : N) (rate_out : N) (qty_trade : N) (k : N) (x : N),
327     N }
328   { calc_rx' : forall (rate_in : N) (rate_out : N) (qty_trade : N) (k : N) (x : N), N }.
329
330 (* When TRADE is successfully called, the trade is priced using the correct formula
331    given by calculate_trade. The updated rate is also priced using the formula from
332    calculate_trade. *)
333 Definition trade_pricing_formula (contract : Contract Setup Msg State Error) : Prop :=
334   forall cstate chain ctx msg_payload t_x t_y q cstate' acts,
335     (* the TRADE entrypoint was called succesfully *)
336     t_x = msg_payload.(token_in_trade) ->
337     t_y = msg_payload.(token_out_trade) ->
338     t_x <> t_y ->
339     q = msg_payload.(qty_trade) ->
340     receive contract chain ctx cstate (Some (trade msg_payload)) =

```

```

341     Ok(cstate', acts) ->
342     (* calculate the diffs delta_x and delta_y *)
343     let delta_x :=
344         (get_bal t_x (stor_tokens_held cstate')) -
345         (get_bal t_x (stor_tokens_held cstate)) in
346     let delta_y :=
347         (get_bal t_y (stor_tokens_held cstate)) -
348         (get_bal t_y (stor_tokens_held cstate')) in
349     let rate_in := (get_rate t_x (stor_rates cstate)) in
350     let rate_out := (get_rate t_y (stor_rates cstate)) in
351     let k := (stor_outstanding_tokens cstate) in
352     let x := get_bal t_x (stor_tokens_held cstate) in
353     (* the diff delta_x and delta_y are correct *)
354     delta_x = q /\
355     delta_y = calc_delta_y rate_in rate_out q k x.
356
357
358 Definition trade_update_rates (contract : Contract Setup Msg State Error) : Prop :=
359     forall cstate chain ctx msg_payload cstate' acts,
360     (* the TRADE entrypoint was called succesfully *)
361     receive contract chain ctx cstate (Some (trade msg_payload)) =
362     Ok(cstate', acts) ->
363     let t_x := msg_payload.(token_in_trade) in
364     let t_y := msg_payload.(token_out_trade) in
365     t_x <> t_y /\
366     (* calculate the diffs delta_x and delta_y *)
367     let rate_in := (get_rate t_x (stor_rates cstate)) in
368     let rate_out := (get_rate t_y (stor_rates cstate)) in
369     let q := msg_payload.(qty_trade) in
370     let k := (stor_outstanding_tokens cstate) in
371     let x := get_bal t_x (stor_tokens_held cstate) in
372     (* the new rate of t_x is correct *)
373     let r_x' := calc_rx' rate_in rate_out q k x in
374     (stor_rates cstate') =
375     (FMap.add (token_in_trade msg_payload)
376     (calc_rx'
377         (get_rate (token_in_trade msg_payload) (stor_rates cstate))
378         (get_rate (token_out_trade msg_payload) (stor_rates cstate))
379         (qty_trade msg_payload)
380         (stor_outstanding_tokens cstate)
381         (get_bal (token_in_trade msg_payload) (stor_tokens_held cstate)))
382     (stor_rates cstate))).
383
384
385 Definition trade_update_rates_formula
386     (contract : Contract Setup Msg State Error) : Prop :=
387     forall cstate chain ctx msg_payload cstate' acts,
388     (* the TRADE entrypoint was called succesfully *)
389     receive contract chain ctx cstate (Some (trade msg_payload)) =

```

```

390     Ok(cstate', acts) ->
391     let t_x := msg_payload.(token_in_trade) in
392     let t_y := msg_payload.(token_out_trade) in
393     t_x <> t_y /\
394     (* calculate the diffs delta_x and delta_y *)
395     let rate_in := (get_rate t_x (stor_rates cstate)) in
396     let rate_out := (get_rate t_y (stor_rates cstate)) in
397     let q := msg_payload.(qty_trade) in
398     let k := (stor_outstanding_tokens cstate) in
399     let x := get_bal t_x (stor_tokens_held cstate) in
400     (* the new rate of t_x is correct *)
401     let r_x' := calc_rx' rate_in rate_out q k x in
402     FMap.find t_x (stor_rates cstate') = Some r_x' /\
403     (forall t, t <> t_x ->
404         FMap.find t (stor_rates cstate') =
405         FMap.find t (stor_rates cstate)).
406
407
408 (* When TRADE is successfully called, it emits two TRANSFER actions *)
409 Definition trade_emits_transfers (contract : Contract Setup Msg State Error) : Prop :=
410     forall cstate cstate' chain ctx msg_payload acts,
411     (* the call to TRADE was successful *)
412     receive contract chain ctx cstate (Some (trade (msg_payload))) = Ok(cstate', acts) ->
413     (* the acts list consists of two transfer actions, specified as follows: *)
414     exists transfer_to_x transfer_data_x transfer_payload_x
415         transfer_to_y transfer_data_y transfer_payload_y,
416     (* there is a transfer call for t_x *)
417     let transfer_tx := (act_call
418         (* call to the correct token address *)
419         (msg_payload.(token_in_trade).(token_address))
420         (* with amount = 0 *)
421         0
422         (* and payload transfer_payload_x *)
423         (serialize (FA2Spec.Transfer transfer_payload_x))) in
424     (* with a transfer in it *)
425     In transfer_data_x transfer_payload_x /\
426     (* which itself has transfer data *)
427     In transfer_to_x transfer_data_x.(FA2Spec.txs) /\
428     (* whose quantity is the quantity traded, transferred to the contract *)
429     transfer_to_x.(FA2Spec.amount) = msg_payload.(qty_trade) /\
430     transfer_to_x.(FA2Spec.to_) = ctx.(ctx_contract_address) /\
431     (* there is a transfer call for t_y *)
432     let transfer_ty := (act_call
433         (* call to the correct token address *)
434         (msg_payload.(token_out_trade).(token_address))
435         (* with amount = 0 *)
436         0
437         (* and payload transfer_payload_y *)
438         (serialize (FA2Spec.Transfer transfer_payload_y))) in

```

```

439   (* with a transfer in it *)
440   In transfer_data_y transfer_payload_y /\
441   (* which itself has transfer data *)
442   In transfer_to_y transfer_data_y.(FA2Spec.txs) /\
443   (* whose quantity is the quantity traded, transferred to the contract *)
444   let t_x := msg_payload.(token_in_trade) in
445   let t_y := msg_payload.(token_out_trade) in
446   let rate_in := (get_rate t_x (stor_rates cstate)) in
447   let rate_out := (get_rate t_y (stor_rates cstate)) in
448   let k := (stor_outstanding_tokens cstate) in
449   let x := get_bal t_x (stor_tokens_held cstate) in
450   let q := msg_payload.(qty_trade) in
451   transfer_to_y.(FA2Spec.amount) = calc_delta_y rate_in rate_out q k x /\
452   transfer_to_y.(FA2Spec.to_) = ctx.(ctx_from) /\
453   (* acts is only these two transfers *)
454   (acts = [ transfer_tx ; transfer_ty ] /\
455    acts = [ transfer_ty ; transfer_tx ]).
456
457
458 Definition trade_tokens_held_update (contract : Contract Setup Msg State Error) : Prop :=
459   forall cstate chain ctx msg_payload cstate' acts,
460   (* the call to TRADE was successful *)
461   receive contract chain ctx cstate (Some (trade (msg_payload))) = Ok(cstate', acts) ->
462   (* in the new state *)
463   let t_x := msg_payload.(token_in_trade) in
464   let t_y := msg_payload.(token_out_trade) in
465   let rate_in := (get_rate t_x (stor_rates cstate)) in
466   let rate_out := (get_rate t_y (stor_rates cstate)) in
467   let k := (stor_outstanding_tokens cstate) in
468   let x := get_bal t_x (stor_tokens_held cstate) in
469   let delta_x := msg_payload.(qty_trade) in
470   let delta_y := calc_delta_y rate_in rate_out delta_x k x in
471   let prev_bal_y := get_bal t_y (stor_tokens_held cstate) in
472   let prev_bal_x := get_bal t_x (stor_tokens_held cstate) in
473   (* balances update appropriately *)
474   get_bal t_y (stor_tokens_held cstate') = (prev_bal_y - delta_y) /\
475   get_bal t_x (stor_tokens_held cstate') = (prev_bal_x + delta_x) /\
476   forall t_z,
477     t_z <> t_x ->
478     t_z <> t_y ->
479     get_bal t_z (stor_tokens_held cstate') =
480     get_bal t_z (stor_tokens_held cstate).
481
482 Definition trade_outstanding_update (contract : Contract Setup Msg State Error) : Prop :=
483   forall cstate chain ctx msg_payload cstate' acts,
484   (* the call to TRADE was successful *)
485   receive contract chain ctx cstate (Some (trade (msg_payload))) = Ok(cstate', acts) ->
486   (* in the new state *)
487   (stor_outstanding_tokens cstate') = (stor_outstanding_tokens cstate).

```

```

488
489 Definition trade_pricing (contract : Contract Setup Msg State Error) : Prop :=
490   forall cstate chain ctx msg_payload cstate' acts,
491     (* the call to TRADE was successful *)
492     receive contract chain ctx cstate (Some (trade (msg_payload))) = Ok(cstate', acts) ->
493     (* balances for t_x change appropriately *)
494     FMap.find (token_in_trade msg_payload) (stor_tokens_held cstate') =
495     Some (get_bal (token_in_trade msg_payload) (stor_tokens_held cstate) + (qty_trade
496       msg_payload)) /\
497     (* balances for t_y change appropriately *)
498     let t_x := token_in_trade msg_payload in
499     let t_y := token_out_trade msg_payload in
500     let delta_x := qty_trade msg_payload in
501     let rate_in := (get_rate t_x (stor_rates cstate)) in
502     let rate_out := (get_rate t_y (stor_rates cstate)) in
503     let k := (stor_outstanding_tokens cstate) in
504     let x := get_bal t_x (stor_tokens_held cstate) in
505     (* in the new state *)
506     FMap.find (token_out_trade msg_payload) (stor_tokens_held cstate') =
507     Some (get_bal (token_out_trade msg_payload) (stor_tokens_held cstate)
508       - (calc_delta_y rate_in rate_out delta_x k x)).
509
510 Definition trade_amounts_nonnegative (contract : Contract Setup Msg State Error) : Prop :=
511   forall cstate chain ctx msg_payload cstate' acts,
512     (* the call to TRADE was successful *)
513     receive contract chain ctx cstate (Some (trade (msg_payload))) = Ok(cstate', acts) ->
514     (* delta_x and delta_y are nonnegative *)
515     let t_x := msg_payload.(token_in_trade) in
516     let t_y := msg_payload.(token_out_trade) in
517     let rate_in := (get_rate t_x (stor_rates cstate)) in
518     let rate_out := (get_rate t_y (stor_rates cstate)) in
519     let k := (stor_outstanding_tokens cstate) in
520     let x := get_bal t_x (stor_tokens_held cstate) in
521     let delta_x := msg_payload.(qty_trade) in
522     let delta_y := calc_delta_y rate_in rate_out delta_x k x in
523     0 <= delta_x /\
524     0 <= delta_y.
525
526 (* Specification of all other entrypoints *)
527 Definition other_rates_unchanged (contract : Contract Setup Msg State Error) : Prop :=
528   forall cstate cstate' chain ctx o acts,
529     (* the call to POOL was successful *)
530     receive contract chain ctx cstate (other o) = Ok(cstate', acts) ->
531     (* rates all stay the same *)
532     forall t,
533     FMap.find t (stor_rates cstate) = FMap.find t (stor_rates cstate').
534
535 Definition other_balances_unchanged (contract : Contract Setup Msg State Error) : Prop :=
536   forall cstate cstate' chain ctx o acts,

```

```

536 (* the call to POOL was successful *)
537 receive contract chain ctx cstate (other o) = Ok(cstate', acts) ->
538 (* balances all stay the same *)
539 forall t,
540 FMap.find t (stor_tokens_held cstate) = FMap.find t (stor_tokens_held cstate').
541
542 Definition other_outstanding_unchanged (contract : Contract Setup Msg State Error) : Prop
    :=
543 forall cstate cstate' chain ctx o acts,
544 (* the call to POOL was successful *)
545 receive contract chain ctx cstate (other o) = Ok(cstate', acts) ->
546 (* balances all stay the same *)
547 (stor_outstanding_tokens cstate) = (stor_outstanding_tokens cstate').
548
549 (** Specification of the functions calc_rx' and calc_delta_y
550 - This specification distils the features of trading along a convex curve which are
551 relevant to the correct functioning of structured pools.
552 - While these are supposed to be nontrivial functions that simulate trading along
553 a convex curve, 'calc_rx'', 'calc_delta_y', and 'calc_rx_inv' can be made trivial
554 by setting all rates to 1, setting the first function to the identity function,
555 the second to multiplication by r_x = 1, and the third to division by r_x = 1
556 *)
557
558 (* update_rate function returns positive number if the num, denom are positive *)
559 Definition update_rate_stays_positive :=
560 forall r_x r_y delta_x k x,
561 let r_x' := calc_rx' r_x r_y delta_x k x in
562 r_x > 0 ->
563 r_x' > 0.
564
565 (* r_x' <= r_x *)
566 Definition rate_decrease :=
567 forall r_x r_y delta_x k x,
568 let r_x' := calc_rx' r_x r_y delta_x k x in
569 r_x' <= r_x.
570
571 (* the inverse rate function is a right inverse of r_x *)
572 Definition rates_balance :=
573 forall q t rates prev_state,
574 let r_x := get_rate t rates in
575 let x := get_bal t (stor_tokens_held prev_state) in
576 r_x * calc_rx_inv r_x q = q.
577
578 Definition rates_balance_2 :=
579 forall t prev_state,
580 let r_x' := calc_rx' (get_rate (token_in_trade t) (stor_rates prev_state))
581 (get_rate (token_out_trade t) (stor_rates prev_state)) (qty_trade t) (
582 stor_outstanding_tokens prev_state)
583 (get_bal (token_in_trade t) (stor_tokens_held prev_state)) in

```

```

583   let delta_y := calc_delta_y (get_rate (token_in_trade t) (stor_rates prev_state))
584         (get_rate (token_out_trade t) (stor_rates prev_state)) (qty_trade t) (
585         stor_outstanding_tokens prev_state)
586         (get_bal (token_in_trade t) (stor_tokens_held prev_state)) in
587   let r_x := get_rate (token_in_trade t) (stor_rates prev_state) in
588   let x := get_bal (token_in_trade t) (stor_tokens_held prev_state) in
589   let y := get_bal (token_out_trade t) (stor_tokens_held prev_state) in
590   let r_y := get_rate (token_out_trade t) (stor_rates prev_state) in
591   r_x' * (x + qty_trade t) + r_y * (y - delta_y) =
592   r_x * x + r_y * y.
593
594 (* the trade always does slightly worse than predicted *)
595 Definition trade_slippage :=
596   forall r_x r_y delta_x k x,
597   let delta_y := calc_delta_y r_x r_y delta_x k x in
598   r_y * delta_y <= r_x * delta_x.
599
600 Definition trade_slippage_2 :=
601   forall r_x r_y delta_x k x,
602   let delta_y := calc_delta_y r_x r_y delta_x k x in
603   let r_x' := calc_rx' r_x r_y delta_x k x in
604   r_y * delta_y <= r_x' * delta_x.
605
606 (* rates have no positive lower bound *)
607 Definition arbitrage_lt :=
608   forall rate_x rate_y ext k x,
609   0 < ext ->
610   ext < rate_x ->
611   exists delta_x,
612   calc_rx' rate_x rate_y delta_x k x <= ext.
613
614 (* calc_delta_y has no positive upper bound *)
615 Definition arbitrage_gt :=
616   forall rate_x rate_y ext_goal k x,
617   rate_x > 0 /\
618   rate_y > 0 /\
619   x > 0 /\
620   k > 0 ->
621   exists delta_x,
622   ext_goal <= calc_delta_y rate_x rate_y delta_x k x.
623
624 (* Initialization specification *)
625 Definition initialized_with_positive_rates (contract : Contract Setup Msg State Error) :=
626   forall chain ctx setup cstate,
627   (* If the contract initializes successfully *)
628   init contract chain ctx setup = Ok cstate ->
629   (* then all rates are nonzero *)
630   forall t r,
631   FMap.find t (stor_rates cstate) = Some r ->
632   r > 0.

```

```

631 Definition initialized_with_zero_balance (contract : Contract Setup Msg State Error) :=
632   forall chain ctx setup cstate,
633   (* If the contract initializes successfully *)
634   init contract chain ctx setup = Ok cstate ->
635   (* then all token balances initialize to zero *)
636   forall t,
637   get_bal t (stor_tokens_held cstate) = 0.
638
639 Definition initialized_with_zero_outstanding (contract : Contract Setup Msg State Error)
640   :=
641   forall chain ctx setup cstate,
642   (* If the contract initializes successfully *)
643   init contract chain ctx setup = Ok cstate ->
644   (* then there are no outstanding tokens *)
645   stor_outstanding_tokens cstate = 0.
646
647 Definition initialized_with_init_rates (contract : Contract Setup Msg State Error) :=
648   forall chain ctx setup cstate,
649   (* If the contract initializes successfully *)
650   init contract chain ctx setup = Ok cstate ->
651   (* then the init rates map is the same as given in the setup *)
652   (stor_rates cstate) = (init_rates setup).
653
654 Definition initialized_with_pool_token (contract : Contract Setup Msg State Error) :=
655   forall chain ctx setup cstate,
656   (* If the contract initializes successfully *)
657   init contract chain ctx setup = Ok cstate ->
658   (* then the pool token is the same as given in the setup *)
659   (stor_pool_token cstate) = (init_pool_token setup).

```

A.1.3 The Formal Specification Predicate

Listing A.6: The amalgamation of the propositions of the specification of Listing ?? into a single predicate on smart contracts, which is the formal specification of a structured pool contract.

```

1 Definition is_structured_pool
2   (C : Contract Setup Msg State Error) : Prop :=
3   none_fails C /\
4   msg_destruct C /\
5   (* pool entrypoint specification *)
6   pool_entrypoint_check C /\
7   pool_emits_txns C /\
8   pool_increases_tokens_held C /\
9   pool_rates_unchanged C /\
10  pool_outstanding C /\
11  (* unpool entrypoint specification *)
12  unpool_entrypoint_check C /\

```



```

13  unpool_entrystate_check_2 C /\
14  unpool_emits_txns C /\
15  unpool_decreases_tokens_held C /\
16  unpool_rates_unchanged C /\
17  unpool_outstanding C /\
18  (* trade entrystate specification *)
19  trade_entrystate_check C /\
20  trade_entrystate_check_2 C /\
21  trade_pricing_formula C /\
22  trade_update_rates C /\
23  trade_update_rates_formula C /\
24  trade_emits_transfers C /\
25  trade_tokens_held_update C /\
26  trade_outstanding_update C /\
27  trade_pricing C /\
28  trade_amounts_nonnegative C /\
29  (* specification of all other entrystates *)
30  other_rates_unchanged C /\
31  other_balances_unchanged C /\
32  other_outstanding_unchanged C /\
33  (* specification of calc_rx' and calc_delta_y *)
34  update_rate_stays_positive /\
35  rate_decrease /\
36  rates_balance /\
37  rates_balance_2 /\
38  trade_slippage /\
39  trade_slippage_2 /\
40  arbitrage_lt /\
41  arbitrage_gt /\
42  (* initialization specification *)
43  initialized_with_positive_rates C /\
44  initialized_with_zero_balance C /\
45  initialized_with_zero_outstanding C /\
46  initialized_with_init_rates C /\
47  initialized_with_pool_token C.

```

A.2 Formal Metaspecification

Here we include the formalization and proofs of each property of the metaspecification in Chapter 4.

A.2.1 Demand Sensitivity

Property 1 (Demand Sensitivity). *Let t_x and t_y be tokens in our family with nonzero pooled liquidity and exchange rates $r_x, r_y > 0$. In a trade t_x to t_y , as r_x is updated to r'_x , it decreases relative to r_z for all $z \neq x$, and r_y strictly increases relative to r_x .*

Listing A.7: The formalization and proof of Demand Sensitivity, Property 1.

```
1 Theorem demand_sensitivity cstate :
2   (* For all tokens t_x t_y, rates r_x r_y, and quantities x and y, where *)
3   forall t_x r_x x t_y r_y y,
4   (* t_x is a token with nonzero pooled liquidity and with rate r_x > 0, and *)
5   FMap.find t_x (stor_tokens_held cstate) = Some x /\ x > 0 /\
6   FMap.find t_x (stor_rates cstate) = Some r_x /\ r_x > 0 ->
7   (* t_y is a token with nonzero pooled liquidity and with rate r_y > 0 *)
8   FMap.find t_y (stor_tokens_held cstate) = Some y /\ y > 0 /\
9   FMap.find t_y (stor_rates cstate) = Some r_y /\ r_y > 0 ->
10  (* In a trade t_x to t_y ... *)
11  forall chain ctx msg msg_payload acts cstate',
12    (* i.e.: a successful call to the contract *)
13    receive contract chain ctx cstate (Some msg) = Ok(cstate', acts) ->
14    (* which is a trade *)
15    msg = trade msg_payload ->
16    (* from t_x to t_y *)
17    msg_payload.(token_in_trade) = t_x ->
18    msg_payload.(token_out_trade) = t_y ->
19    (* with t_x <> t_y *)
20    t_x <> t_y ->
21    (* ... as r_x is updated to r_x': ... *)
22    let r_x' := get_rate t_x (stor_rates cstate') in
23    (* (1) r_x decreases relative to all rates r_z, for t_z <> t_x, and *)
24    (forall t_z,
25      t_z <> t_x ->
26      let r_z := get_rate t_z (stor_rates cstate) in
27      let r_z' := get_rate t_z (stor_rates cstate') in
28      rel_decr r_x r_z r_x' r_z') /\
29    (* (2) r_y strictly increases relative to r_x *)
30    let t_y := msg_payload.(token_out_trade) in
31    let r_y := get_rate t_y (stor_rates cstate) in
32    let r_y' := get_rate t_y (stor_rates cstate') in
33    rel_incr r_y r_x r_y' r_x'.
```

Listing A.8: Definitions and lemmas supporting the formal definition and proof of Demand Sensitivity.

```

1 (* x decreases relative to z as x => x', z => z' : z - x <= z' - x' *)
2 Definition rel_decr (x z x' z' : N) :=
3   ((Z.of_N z) - (Z.of_N x) <= (Z.of_N z') - (Z.of_N x'))%Z.
4
5 Lemma rel_decr_lem : forall x x' z : N, x' <= x -> rel_decr x z x' z.
6
7 (* y increases relative to x as y => y', x => x' : y - x <= y' - x' *)
8 Definition rel_incr (y x y' x' : N) :=
9   ((Z.of_N y) - (Z.of_N x) <= (Z.of_N y') - (Z.of_N x'))%Z.
10
11 Lemma rel_incr_lem : forall x x' y : N, x' <= x -> rel_incr y x y x'.

```

A.2.2 Nonpathological Prices

Property 2 (Nonpathological Prices). *For a token t_x in T , if there is a contract state such that $r_x > 0$, then $r_x > 0$ holds for all future states of the contract.*

Listing A.9: The formalization and proof of Nonpathological Prices, Property 2.

```

1 Theorem nonpathological_prices bstate caddr :
2   (* reachable state with contract at caddr *)
3   reachable bstate ->
4   env_contracts bstate caddr = Some (contract : WeakContract) ->
5   (* the statement *)
6   exists (cstate : State),
7   contract_state bstate caddr = Some cstate /\
8   (* For a token t_x in T and rate r_x, *)
9   forall t_x r_x,
10  (* if r_x is the exchange rate of t_x, then r_x > 0 *)
11  FMap.find t_x (stor_rates cstate) = Some r_x -> r_x > 0.

```

A.2.3 Swap Rate Consistency

Property 3 (Swap Rate Consistency). *Let t_x be a token in our family with nonzero pooled liquidity and $r_x > 0$. Then for any $\Delta_x > 0$ there is no sequence of trades, beginning and ending with t_x , such that $\Delta'_x > \Delta_x$, where Δ'_x is the output quantity of the sequence of trades.*

Listing A.10: The formalization and proof of Swap Rate Consistency, Property 3.

```

1 Theorem swap_rate_consistency bstate cstate :
2   (* Let t_x be a token with nonzero pooled liquidity and rate r_x > 0 *)
3   forall t_x r_x x,
4     FMap.find t_x (stor_rates cstate) = Some r_x /\ r_x > 0 ->
5     FMap.find t_x (stor_tokens_held cstate) = Some x /\ x > 0 ->
6     (*then for any delta_x > 0 and any sequence of trades, beginning and ending with t_x*)
7     forall delta_x (trade_sequence : list trade_sequence_type) t_fst t_last,
8       delta_x > 0 ->
9       (* trade_sequence is a list of successive trades *)
10      are_successive_trades trade_sequence ->
11      (* with a first and last trade, t_fst and t_last respectively, *)
12      (hd_error trade_sequence) = Some t_fst ->
13      (hd_error (rev trade_sequence)) = Some t_last ->
14      (* starting from our current bstate and cstate *)
15      seq_chain t_fst = bstate ->
16      seq_cstate t_fst = cstate ->
17      (* the first trade is from t_x *)
18      token_in_trade (seq_trade_data t_fst) = t_x ->
19      qty_trade (seq_trade_data t_fst) = delta_x ->
20      (* the last trade is to t_x *)
21      token_out_trade (seq_trade_data t_last) = t_x ->
22      FMap.find t_x (stor_rates cstate) = FMap.find t_x (stor_rates (seq_cstate t_last)) ->
23      (* delta_x', the output of the last trade, is never larger than delta_x. *)
24      let delta_x' := trade_to_delta_y t_last in
25      delta_x' <= delta_x.
26 Proof.
27   intros * H_rate H_held * dx_geq_0 trades_successive fst_txn lst_txn start_bstate
28     start_cstate from_tx trade_delta_x to_tx one_tx_txn *.
29   unfold delta_x'.
30   rewrite <- trade_delta_x.
31   apply (geq_list_is_sufficient trade_sequence t_x t_fst t_last cstate r_x); auto.
32   now apply swap_rate_lemma.
33 Qed.

```

Listing A.11: Definitions and lemmas relevant to the formal definition of Swap Rate Consistency

```

1 (* first a type to describe successive trades *)
2 Record trade_sequence_type := build_trade_sequence_type {
3   seq_chain : ChainState ;
4   seq_ctx : ContractCallContext ;
5   seq_cstate : State ;
6   seq_trade_data : trade_data ;
7   seq_res_acts : list ActionBody ;
8 }.
9
10 (* a function to calculate the the trade output of the final trade, delta_x' *)
11 Definition trade_to_delta_y (t : trade_sequence_type) :=
12   let cstate := seq_cstate t in

```

```

13   let token_in := token_in_trade (seq_trade_data t) in
14   let token_out := token_out_trade (seq_trade_data t) in
15   let rate_in := get_rate token_in (stor_rates cstate) in
16   let rate_out := get_rate token_out (stor_rates cstate) in
17   let delta_x := qty_trade (seq_trade_data t) in
18   let k := stor_outstanding_tokens cstate in
19   let x := get_bal token_in (stor_tokens_held cstate) in
20   (* the calculation *)
21   calc_delta_y rate_in rate_out delta_x k x.
22
23   (* a proposition that indicates a list of trades are successive, successful trades *)
24   Fixpoint are_successive_trades (trade_sequence : list trade_sequence_type) : Prop :=
25     match trade_sequence with
26     | [] => True
27     | t1 :: l =>
28       match l with
29       | [] =>
30         (* if the list has one element, it just has to succeed *)
31         exists cstate' acts,
32         receive contract
33           (seq_chain t1)
34           (seq_ctx t1)
35           (seq_cstate t1)
36           (Some (trade (seq_trade_data t1)))
37         = Ok(cstate', acts)
38       | t2 :: l' =>
39         (* the trade t1 goes through, connecting t1 and t2 *)
40         receive contract
41           (seq_chain t1)
42           (seq_ctx t1)
43           (seq_cstate t1)
44           (Some (trade (seq_trade_data t1)))
45         = Ok(seq_cstate t2, seq_res_acts t2) /\
46           (qty_trade (seq_trade_data t2) = trade_to_delta_y t1) /\
47           (token_in_trade (seq_trade_data t2) = token_out_trade (seq_trade_data t1)) /\
48           (are_successive_trades l)
49       end
50     end.
51
52   Fixpoint geq_list (l : list N) : Prop :=
53     match l with
54     | [] => True
55     | h :: t1 =>
56       match t1 with
57       | [] => True
58       | h' :: t1' => (h >= h') /\ geq_list t1
59       end
60     end.

```

Listing A.12: The two lemmas that constitute the essence of the formal proof of Swap Rate Consistency.

```

1 Lemma geq_list_is_sufficient : forall trade_sequence t_x t_fst t_last cstate r_x,
2   (* more assumptions *)
3   (hd_error trade_sequence) = Some t_fst ->
4   (hd_error (rev trade_sequence)) = Some t_last ->
5   token_in_trade (seq_trade_data t_fst) = t_x ->
6   token_out_trade (seq_trade_data t_last) = t_x ->
7   seq_cstate t_fst = cstate ->
8   FMap.find t_x (stor_rates cstate) = Some r_x /\ r_x > 0 ->
9   FMap.find t_x (stor_rates cstate) = FMap.find t_x (stor_rates (seq_cstate t_last)) ->
10  (* the statement *)
11  geq_list (map trade_to_ry_delta_y trade_sequence) ->
12  let delta_x := qty_trade (seq_trade_data t_fst) in
13  let delta_x' := trade_to_delta_y t_last in
14  delta_x' <= delta_x.
15
16 Lemma swap_rate_lemma : forall trade_sequence,
17  (* if this is a list of successive trades *)
18  are_successive_trades trade_sequence ->
19  (* then *)
20  geq_list (map trade_to_ry_delta_y trade_sequence).

```

Listing A.13: Definitions relevant to the formal definition and proof of the two auxiliary lemmas of the proof of Swap Rate Consistency.

```

1 Definition trade_to_ry_delta_y (t : trade_sequence_type) :=
2   let delta_y := trade_to_delta_y t in
3   let rate_y := get_rate (token_out_trade (seq_trade_data t)) (stor_rates (seq_cstate t))
4   ) in
5   rate_y * delta_y.

```

A.2.4 Zero-Impact Liquidity Change

Property 4 (Zero-Impact Liquidity Change). *The quoted price of trades is unaffected by calling DEPOSIT and WITHDRAW.*

Listing A.14: The formalization and proof of Zero-Impact Liquidity Change, Property 4.

```

1 Theorem zero_impact_liquidity_change :
2   (* Consider the quoted price of a trade t_x to t_y at cstate, *)
3   forall cstate t_x t_y r_x r_y,
4   FMap.find t_x (stor_rates cstate) = Some r_x ->
5   FMap.find t_y (stor_rates cstate) = Some r_y ->
6   let quoted_price := r_x / r_y in
7   (* and a successful POOL or UNPOOL action. *)
8   forall chain ctx msg payload_pool payload_unpool acts cstate' r_x' r_y',

```

```

9      receive contract chain ctx cstate (Some msg) = Ok(cstate', acts) ->
10      msg = pool payload_pool \/
11      msg = unpool payload_unpool ->
12      (* Then take the (new) quoted price of a trade t_x to t_y at cstate'. *)
13      FMap.find t_x (stor_rates cstate') = Some r_x' ->
14      FMap.find t_y (stor_rates cstate') = Some r_y' ->
15      let quoted_price' := r_x' / r_y' in
16      (* The quoted price is unchanged. *)
17      quoted_price = quoted_price'.

```

A.2.5 Arbitrage Sensitivity

Property 5 (Arbitrage sensitivity). *Let t_x be a token in our family with nonzero pooled liquidity and $r_x > 0$. If an external, demand-sensitive market prices t_x differently from the structured pool, then assuming sufficient liquidity, with a sufficiently large transaction either the price of t_x in the structured pool converges with the external market, or the trade depletes the pool of t_x .*

```

1 Theorem arbitrage_sensitivity :
2   forall cstate t_x r_x x,
3     (* t_x is a token with nonzero pooled liquidity *)
4     FMap.find t_x (stor_rates cstate) = Some r_x /\ r_x > 0 /\
5     FMap.find t_x (stor_tokens_held cstate) = Some x /\ x > 0 ->
6     (* we consider some external price *)
7     forall external_price,
8       0 < external_price ->
9       (* and a trade of trade_qty succeeds *)
10      forall chain ctx msg msg_payload cstate' acts,
11        receive contract chain ctx cstate msg = Ok(cstate', acts) ->
12        msg = Some(trade msg_payload) ->
13        t_x = (token_in_trade msg_payload) ->
14        (* the arbitrage opportunity is resolved *)
15        let r_x' := get_rate t_x (stor_rates cstate') in
16        (* first the case that the external price was lower *)
17        (external_price < r_x ->
18          exists trade_qty,
19            msg_payload.(qty_trade) = trade_qty ->
20            external_price >= r_x') /\
21        (* second the case that the external price is higher *)
22        (external_price > r_x ->
23          exists trade_qty,
24            msg_payload.(qty_trade) = trade_qty ->
25            external_price <= r_x' /\
26        let t_y := token_out_trade msg_payload in
27        let r_y := get_rate t_y (stor_rates cstate) in
28        let x := get_bal t_x (stor_tokens_held cstate) in
29        let y := get_bal t_y (stor_tokens_held cstate) in

```

```

30   let balances := (stor_tokens_held cstate) in
31   let k := (stor_outstanding_tokens cstate) in
32   get_bal t_y balances <= calc_delta_y r_x r_y trade_qty k x).

```

Listing A.15: The formalization and proof of Arbitrage Sensitivity, Property 5.

A.2.6 Pooled Consistency

Property 6 (Pooled Consistency). *The following equation always holds:*

$$\sum_{t_x} r_x x = k \quad (\text{A.1})$$

Listing A.16: The formalization and proof of Pooled Consistency, Property 6.

```

1 Theorem pooled_consistency bstate caddr :
2   reachable bstate ->
3   env_contracts bstate caddr = Some (contract : WeakContract) ->
4   exists (cstate : State),
5   contract_state bstate caddr = Some cstate /\
6   (* The sum of all the constituent, pooled tokens, multiplied by their value in terms
7   of pooled tokens, always equals the total number of outstanding pool tokens. *)
8   suml (tokens_to_values (stor_rates cstate) (stor_tokens_held cstate)) =
9   (stor_outstanding_tokens cstate).

```

Listing A.17: Definitions and lemmas supporting the formal definition and proof of Pooled Consistency.

```

1 (* map over the keys *)
2 Definition tokens_to_values
3   (rates : FMap token exchange_rate) (tokens_held : FMap token N) : list N :=
4   List.map
5     (fun k =>
6       let rate := get_rate k rates in
7       let qty_held := get_bal k tokens_held in
8       rate * qty_held)
9     (FMap.keys rates).
10
11 (* take the sum of a list *)
12 Definition suml l := fold_right N.add 0 l.

```


Appendix B

Proofs and Definitions of Chapter 5

B.1 Contract Morphisms

Listing B.1: The formal definition of a contract morphism.

```
1 Section MorphismDefinition.
2 Context `{Serializable Setup1} `{Serializable Msg1} `{Serializable State1} `{Serializable
   Error1}
3     `{Serializable Setup2} `{Serializable Msg2} `{Serializable State2} `{Serializable
   Error2}.
4
5 (** The definition *)
6 Record ContractMorphism
7   (C1 : Contract Setup1 Msg1 State1 Error1)
8   (C2 : Contract Setup2 Msg2 State2 Error2) :=
9   build_contract_morphism {
10     (* the components of a morphism f *)
11     setup_morph : Setup1 -> Setup2 ;
12     msg_morph   : Msg1   -> Msg2   ;
13     state_morph : State1 -> State2 ;
14     error_morph : Error1 -> Error2 ;
15     (* coherence conditions *)
16     init_coherence : forall c ctx s,
17       result_functor state_morph error_morph
18         (init C1 c ctx s) =
19       init C2 c ctx (setup_morph s) ;
20     recv_coherence : forall c ctx st op_msg,
21       result_functor (fun ' (st, l) => (state_morph st, l)) error_morph
22         (receive C1 c ctx st op_msg) =
23       receive C2 c ctx (state_morph st) (option_map msg_morph op_msg) ;
24   }.
25
26 End MorphismDefinition.
```

Listing B.2: The definition of result_functor.

```

1 Definition result_functor {T T' E E' : Type} : (T -> T') -> (E -> E') -> result T E ->
  result T' E' :=
2   fun (f_t : T -> T') (f_e : E -> E') (res : result T E) =>
3   match res with | Ok t => Ok (f_t t) | Err e => Err (f_e e) end.

```

Listing B.3: The identity contract morphism.

```

1 Section IdentityMorphism.
2 Context `{Serializable Msg} `{Serializable Setup} `{Serializable State} `{Serializable
  Error}.
3
4 Lemma init_coherence_id (C : Contract Setup Msg State Error) :
5   forall c ctx s,
6   result_functor id id (init C c ctx s) =
7   init C c ctx s.
8 Proof.
9   intros.
10  unfold result_functor.
11  now destruct (init C c ctx s).
12 Qed.
13
14 Lemma recv_coherence_id (C : Contract Setup Msg State Error) :
15   forall c ctx st op_msg,
16   result_functor
17     (fun '(st, l) => (id st, l)) id
18     (receive C c ctx st op_msg) =
19   receive C c ctx (id st) (option_map id op_msg).
20 Proof.
21   intros.
22   unfold result_functor, option_map. cbn.
23   destruct op_msg.
24   - now destruct (receive C c ctx st (Some m)); try destruct t.
25   - now destruct (receive C c ctx st None); try destruct t.
26 Qed.
27
28 (** The identity morphism *)
29 Definition id_cm (C : Contract Setup Msg State Error) : ContractMorphism C C := {|
30   (* components *)
31   setup_morph := id ;
32   msg_morph   := id ;
33   state_morph := id ;
34   error_morph := id ;
35   (* coherence conditions *)
36   init_coherence := init_coherence_id C ;
37   recv_coherence := recv_coherence_id C ;
38   |}.

```

39

40 End IdentityMorphism.

Listing B.4: The formal definition of a contract injection and surjection.

```

1 Section Injections.
2 Context `{Serializable Setup1} `{Serializable Msg1} `{Serializable State1} `{Serializable
   Error1}
3     `{Serializable Setup2} `{Serializable Msg2} `{Serializable State2} `{Serializable
   Error2}
4     {C1 : Contract Setup1 Msg1 State1 Error1}
5     {C2 : Contract Setup2 Msg2 State2 Error2}.
6
7 Definition is_inj {A B : Type} (f : A -> B) : Prop :=
8     forall (a a' : A), f a = f a' -> a = a'.
9
10 Lemma is_inj_compose {A B C : Type} :
11     forall (f1 : A -> B) (f2 : B -> C),
12         is_inj f1 ->
13         is_inj f2 ->
14         is_inj (compose f2 f1).
15 Proof.
16     intros * f1_inj f2_inj.
17     unfold is_inj in *.
18     intros * H_compose.
19     simpl in H_compose.
20     apply f2_inj in H_compose.
21     now apply f1_inj in H_compose.
22 Qed.
23
24 (* A morphism is a *weak embedding* if it embeds the message and storage types *)
25 Definition is_weak_inj_cm (f : ContractMorphism C1 C2) : Prop :=
26     is_inj (msg_morph C1 C2 f) /\
27     is_inj (state_morph C1 C2 f).
28
29 Definition contract_weakly_embeds : Prop :=
30     exists (f : ContractMorphism C1 C2), is_weak_inj_cm f.
31
32 (* A morphism is an embedding if it embeds on all contract types *)
33 Definition is_inj_cm (f : ContractMorphism C1 C2) : Prop :=
34     is_inj (setup_morph C1 C2 f) /\
35     is_inj (msg_morph C1 C2 f) /\
36     is_inj (state_morph C1 C2 f) /\
37     is_inj (error_morph C1 C2 f).
38
39 Definition contract_embeds : Prop :=
40     exists (f : ContractMorphism C1 C2), is_inj_cm f.
41
42 End Injections.

```

```

43
44
45 (** Surjective contract morphisms *)
46 Section Surjections.
47 Context `{Serializable Setup1} `{Serializable Msg1} `{Serializable State1} `{Serializable
    Error1}
48     `{Serializable Setup2} `{Serializable Msg2} `{Serializable State2} `{Serializable
    Error2}
49     {C1 : Contract Setup1 Msg1 State1 Error1}
50     {C2 : Contract Setup2 Msg2 State2 Error2}.
51
52 Definition is_surj {A B : Type} (f : A -> B) : Prop :=
53     forall (b : B), exists (a : A), f a = b.
54
55 (* A morphism is a *weak quotient* if it embeds on all contract types *)
56 Definition is_weak_surj_cm (f : ContractMorphism C1 C2) : Prop :=
57     is_surj (msg_morph C1 C2 f) /\
58     is_surj (state_morph C1 C2 f).
59
60 Definition contract_weakly_surjects : Prop :=
61     exists (f : ContractMorphism C1 C2), is_weak_surj_cm f.
62
63 (* A morphism is a surjection if it surjects on all contract types *)
64 Definition is_surj_cm (f : ContractMorphism C1 C2) : Prop :=
65     is_surj (setup_morph C1 C2 f) /\
66     is_surj (msg_morph C1 C2 f) /\
67     is_surj (state_morph C1 C2 f) /\
68     is_surj (error_morph C1 C2 f).
69
70 Definition contract_surjects : Prop :=
71     exists (f : ContractMorphism C1 C2), is_surj_cm f.
72
73 End Surjections.

```

Listing B.5: Equality of contract morphisms.

```

1 Section EqualityOfMorphisms.
2 Context `{Serializable Setup1} `{Serializable Msg1} `{Serializable State1} `{Serializable
    Error1}
3     `{Serializable Setup2} `{Serializable Msg2} `{Serializable State2} `{Serializable
    Error2}
4     {C1 : Contract Setup1 Msg1 State1 Error1}
5     {C2 : Contract Setup2 Msg2 State2 Error2}.
6
7 Lemma eq_cm :
8     forall (f g : ContractMorphism C1 C2),
9     (* if the components are equal ... *)
10     (setup_morph C1 C2 f) = (setup_morph C1 C2 g) ->
11     (msg_morph C1 C2 f) = (msg_morph C1 C2 g) ->

```

```

12   (state_morph C1 C2 f) = (state_morph C1 C2 g) ->
13   (error_morph C1 C2 f) = (error_morph C1 C2 g) ->
14   (* ... then the morphisms are equal *)
15   f = g.
16 Proof.
17   intros f g.
18   destruct f, g.
19   cbn.
20   intros msg_eq set_eq st_eq err_eq.
21   subst.
22   f_equal;
23   apply proof_irrelevance.
24 Qed.
25
26 Lemma eq_cm_rev :
27   forall (f g : ContractMorphism C1 C2),
28   (* if the morphisms are equal ... *)
29   f = g ->
30   (* ... then the components are equal *)
31   (setup_morph C1 C2 f) = (setup_morph C1 C2 g) /\
32   (msg_morph C1 C2 f) = (msg_morph C1 C2 g) /\
33   (state_morph C1 C2 f) = (state_morph C1 C2 g) /\
34   (error_morph C1 C2 f) = (error_morph C1 C2 g).
35 Proof.
36   intros f g fg_eq.
37   now inversion fg_eq.
38 Qed.
39
40 Lemma eq_cm_iff :
41   forall (f g : ContractMorphism C1 C2),
42   (* the components are equal ... *)
43   (setup_morph C1 C2 f) = (setup_morph C1 C2 g) /\
44   (msg_morph C1 C2 f) = (msg_morph C1 C2 g) /\
45   (state_morph C1 C2 f) = (state_morph C1 C2 g) /\
46   (error_morph C1 C2 f) = (error_morph C1 C2 g) <->
47   (* ... iff the morphisms are equal *)
48   f = g.
49 Proof.
50   intros.
51   split.
52   - intro H_components.
53     destruct H_components as [H_set [H_msg [H_st H_err]]].
54     now apply eq_cm.
55   - now apply eq_cm_rev.
56 Qed.
57
58 End EqualityOfMorphisms.

```

B.1.1 Composition of Morphisms

Listing B.6: Composition of contract morphisms.

```

1 Section MorphismComposition.
2 Context `{Serializable Setup1} `{Serializable Msg1} `{Serializable State1} `{Serializable
   Error1}
3     `{Serializable Setup2} `{Serializable Msg2} `{Serializable State2} `{Serializable
   Error2}
4     `{Serializable Setup3} `{Serializable Msg3} `{Serializable State3} `{Serializable
   Error3}
5     { C1 : Contract Setup1 Msg1 State1 Error1 }
6     { C2 : Contract Setup2 Msg2 State2 Error2 }
7     { C3 : Contract Setup3 Msg3 State3 Error3 }.
8
9 Lemma compose_init_coh (g : ContractMorphism C2 C3) (f : ContractMorphism C1 C2) :
10   let setup_morph' := (compose (setup_morph C2 C3 g) (setup_morph C1 C2 f)) in
11   let state_morph' := (compose (state_morph C2 C3 g) (state_morph C1 C2 f)) in
12   let error_morph' := (compose (error_morph C2 C3 g) (error_morph C1 C2 f)) in
13   forall c ctx s,
14     result_functor state_morph' error_morph'
15       (init C1 c ctx s) =
16       init C3 c ctx (setup_morph' s).
17 Proof.
18   intros.
19   unfold setup_morph'.
20   cbn.
21   rewrite <- (init_coherence C2 C3 g).
22   rewrite <- (init_coherence C1 C2 f).
23   unfold result_functor.
24   now destruct (init C1 c ctx s).
25 Qed.
26
27 Lemma compose_recv_coh (g : ContractMorphism C2 C3) (f : ContractMorphism C1 C2) :
28   let msg_morph' := (compose (msg_morph C2 C3 g) (msg_morph C1 C2 f)) in
29   let state_morph' := (compose (state_morph C2 C3 g) (state_morph C1 C2 f)) in
30   let error_morph' := (compose (error_morph C2 C3 g) (error_morph C1 C2 f)) in
31   forall c ctx st op_msg,
32     result_functor
33       (fun ' (st, l) => (state_morph' st, l)) error_morph'
34       (receive C1 c ctx st op_msg) =
35       receive C3 c ctx (state_morph' st) (option_map msg_morph' op_msg).
36 Proof.
37   intros.
38   pose proof (recv_coherence C2 C3 g).
39   pose proof (recv_coherence C1 C2 f).
40   unfold state_morph', msg_morph'.
41   cbn.
42   replace (option_map (compose (msg_morph C2 C3 g) (msg_morph C1 C2 f)) op_msg)
43   with (option_map (msg_morph C2 C3 g) (option_map (msg_morph C1 C2 f) op_msg)).

```

```

44   2:{ now destruct op_msg. }
45   rewrite <- H11.
46   rewrite <- H12.
47   unfold result_functor.
48   now destruct (receive C1 c ctx st op_msg).
49   Qed.
50
51   (** Composition *)
52   Definition compose_cm (g : ContractMorphism C2 C3) (f : ContractMorphism C1 C2) :
53     ContractMorphism C1 C3 := { |
54       (* the components *)
55       setup_morph := compose (setup_morph C2 C3 g) (setup_morph C1 C2 f) ;
56       msg_morph    := compose (msg_morph C2 C3 g) (msg_morph C1 C2 f) ;
57       state_morph  := compose (state_morph C2 C3 g) (state_morph C1 C2 f) ;
58       error_morph  := compose (error_morph C2 C3 g) (error_morph C1 C2 f) ;
59       (* the coherence results *)
60       init_coherence := compose_init_coh g f ;
61       recv_coherence := compose_recv_coh g f ;
62     |}.
63
64   End MorphismComposition.

```

Listing B.7: Some results on contract morphisms.

```

1   Section MorphismCompositionResults.
2   Context `{Serializable Setup1} `{Serializable Msg1} `{Serializable State1} `{Serializable
      Error1}
3     `{Serializable Setup2} `{Serializable Msg2} `{Serializable State2} `{Serializable
      Error2}
4     `{Serializable Setup3} `{Serializable Msg3} `{Serializable State3} `{Serializable
      Error3}
5     `{Serializable Setup4} `{Serializable Msg4} `{Serializable State4} `{Serializable
      Error4}
6     { C1 : Contract Setup1 Msg1 State1 Error1 }
7     { C2 : Contract Setup2 Msg2 State2 Error2 }
8     { C3 : Contract Setup3 Msg3 State3 Error3 }
9     { C4 : Contract Setup4 Msg4 State4 Error4 }.
10
11   (** Composition with the Identity morphism is trivial *)
12   Lemma compose_id_cm_left (f : ContractMorphism C1 C2) :
13     compose_cm (id_cm C2) f = f.
14   Proof.
15     now apply eq_cm.
16   Qed.
17
18   Lemma compose_id_cm_right (f : ContractMorphism C1 C2) :
19     compose_cm f (id_cm C1) = f.
20   Proof.
21     now apply eq_cm.

```

```

22 Qed.
23
24 (** Composition is associative *)
25 Lemma compose_cm_assoc
26   (f : ContractMorphism C1 C2)
27   (g : ContractMorphism C2 C3)
28   (h : ContractMorphism C3 C4) :
29   compose_cm h (compose_cm g f) =
30   compose_cm (compose_cm h g) f.
31 Proof.
32   now apply eq_cm.
33 Qed.
34
35 End MorphismCompositionResults.

```

Listing B.8: Definition of a contract isomorphism.

```

1 Section IsomorphismDefinition.
2 Context `{Serializable Setup1} `{Serializable Msg1} `{Serializable State1} `{Serializable
   Error1}
3       `{Serializable Setup2} `{Serializable Msg2} `{Serializable State2} `{Serializable
   Error2}
4       {C1 : Contract Setup1 Msg1 State1 Error1}
5       {C2 : Contract Setup2 Msg2 State2 Error2}.
6
7 Definition is_iso {A B : Type} (f : A -> B) (g : B -> A) : Prop :=
8   compose g f = @id A /\ compose f g = @id B.
9
10 Lemma is_iso_reflexive {A B : Type} (f : A -> B) (g : B -> A) :
11   is_iso f g -> is_iso g f.
12 Proof.
13   unfold is_iso.
14   intro H_is_iso.
15   now destruct H_is_iso.
16 Qed.
17
18 Definition is_iso_cm (f : ContractMorphism C1 C2) (g : ContractMorphism C2 C1) : Prop :=
19   compose_cm g f = id_cm C1 /\
20   compose_cm f g = id_cm C2.
21
22 Lemma iso_cm_components :
23   forall (f : ContractMorphism C1 C2) (g : ContractMorphism C2 C1),
24   is_iso (msg_morph C1 C2 f) (msg_morph C2 C1 g) /\
25   is_iso (setup_morph C1 C2 f) (setup_morph C2 C1 g) /\
26   is_iso (state_morph C1 C2 f) (state_morph C2 C1 g) /\
27   is_iso (error_morph C1 C2 f) (error_morph C2 C1 g)
28   <->
29   is_iso_cm f g.
30 Proof.

```



```

31   intros f g.
32   unfold is_iso.
33   unfold iff.
34   split.
35   (* -> *)
36   -   intro H_iso.
37       unfold is_iso_cm.
38       split; now apply eq_cm.
39   (* <- *)
40   -   unfold is_iso_cm, compose_cm, id_cm.
41       now intro H_iso.
42   Qed.
43
44 End IsomorphismDefinition.

```

Listing B.9: Some results on contract isomorphisms.

```

1 Section IsomorphismsResults.
2 Context `{Serializable Setup1} `{Serializable Msg1} `{Serializable State1} `{Serializable
   Error1}
3     `{Serializable Setup2} `{Serializable Msg2} `{Serializable State2} `{Serializable
   Error2}.
4
5 (** An equality predicate on contracts *)
6 Definition contracts_isomorphic
7   (C1 : Contract Setup1 Msg1 State1 Error1)
8   (C2 : Contract Setup2 Msg2 State2 Error2) : Prop :=
9   exists (f : ContractMorphism C1 C2) (g : ContractMorphism C2 C1),
10  is_iso_cm f g.
11
12 Context `{Serializable Setup3} `{Serializable Msg3} `{Serializable State3} `{Serializable
   Error3}
13     `{Serializable Setup4} `{Serializable Msg4} `{Serializable State4} `{Serializable
   Error4}
14     { C1 : Contract Setup1 Msg1 State1 Error1 }
15     { C2 : Contract Setup2 Msg2 State2 Error2 }
16     { C3 : Contract Setup3 Msg3 State3 Error3 }
17     { C4 : Contract Setup4 Msg4 State4 Error4 }.
18
19 Lemma iso_cm_reflexive (f : ContractMorphism C1 C2) (g : ContractMorphism C2 C1) :
20   is_iso_cm f g ->
21   is_iso_cm g f.
22 Proof.
23   intro H_is_iso.
24   apply iso_cm_components in H_is_iso.
25   apply iso_cm_components.
26   destruct H_is_iso as [H_ind_iso H_is_iso].
27   do 2 (apply is_iso_reflexive in H_ind_iso;
28   split;

```

```

29   try exact H_ind_iso;
30   clear H_ind_iso;
31   destruct H_is_iso as [H_ind_iso H_is_iso]).
32   split; now apply is_iso_reflexive.
33   Qed.
34
35   Lemma composition_iso_cm
36     (f1 : ContractMorphism C1 C2)
37     (g1 : ContractMorphism C2 C3)
38     (f2 : ContractMorphism C2 C1)
39     (g2 : ContractMorphism C3 C2) :
40     is_iso_cm f1 f2 ->
41     is_iso_cm g1 g2 ->
42     is_iso_cm (compose_cm g1 f1) (compose_cm f2 g2).
43   Proof.
44     unfold is_iso_cm.
45     intros iso_f iso_g.
46     destruct iso_f as [iso_f1 iso_f2].
47     destruct iso_g as [iso_g1 iso_g2].
48     split; rewrite compose_cm_assoc.
49     - rewrite <- (compose_cm_assoc g1 g2 f2).
50       rewrite iso_g1.
51       now rewrite (compose_id_cm_right f2).
52     - rewrite <- (compose_cm_assoc f2 f1 g1).
53       rewrite iso_f2.
54       now rewrite (compose_id_cm_right g1).
55   Qed.
56
57 End IsomorphismsResults.

```

Listing B.10: An isomorphism of contracts is both injective and surjective.

```

1 Theorem inj_surj_iso_cm (f : ContractMorphism C1 C2) :
2   (exists (g : ContractMorphism C2 C1), is_iso_cm f g) ->
3   is_inj_cm f /\ is_surj_cm f.

```

B.2 Morphism Induction

B.2.1 Contract Trace and Reachability

Listing B.11: The definition of contract trace and contract reachability.

```

1 Section ContractTrace.
2 Context { Setup Msg State Error : Type }
3   '{Serializable Msg} '{Serializable Setup} '{Serializable State} '{Serializable
   Error}.

```

```

4
5 (* Notions of contract state stepping forward *)
6 Record ContractStep (C : Contract Setup Msg State Error)
7   (prev_cstate : State) (next_cstate : State) :=
8   build_contract_step {
9     seq_chain : Chain ;
10    seq_ctx : ContractCallContext ;
11    seq_msg : option Msg ;
12    seq_new_acts : list ActionBody ;
13    (* we can call receive successfully *)
14    recv_some_step :
15      receive C seq_chain seq_ctx prev_cstate seq_msg = Ok (next_cstate, seq_new_acts) ;
16  }.
17
18 Definition ContractTrace (C : Contract Setup Msg State Error) :=
19   ChainedList State (ContractStep C).
20
21 Definition is_genesis_state (C : Contract Setup Msg State Error) (init_cstate : State) :
22   Prop :=
23   exists init_chain init_ctx init_setup,
24   init C init_chain init_ctx init_setup = Ok init_cstate.
25
26 Definition cstate_reachable (C : Contract Setup Msg State Error) (cstate : State) : Prop
27   :=
28   exists init_cstate,
29   (* init_cstate is a valid initial cstate *)
30   is_genesis_state C init_cstate /\
31   (* with a trace to cstate *)
32   inhabited (ContractTrace C init_cstate cstate).
33
34 Lemma inhab_trace_trans (C : Contract Setup Msg State Error) :
35   forall from mid to,
36     (ContractStep C mid to) ->
37     inhabited (ContractTrace C from mid) ->
38     inhabited (ContractTrace C from to).
39
40 Proof.
41   intros from mid to step.
42   apply inhabited_covariant.
43   intro mid_t.
44   apply (snoc mid_t step).
45
46 Qed.
47
48 End ContractTrace.

```

B.2.2 Right Morphism Induction

Listing B.12: The theorem of left morphism induction.

```

1 (* f : C1 -> C2, inducting on C1 *)
2 Theorem left_cm_induction :
3   (* forall simple morphism and reachable bstate, *)
4   forall (f : ContractMorphism C1 C2) bstate caddr (trace : ChainTrace empty_state
5     bstate),
6   (* where C is at caddr with state cstate, *)
7   env_contracts bstate caddr = Some (C1 : WeakContract) ->
8   exists (cstate1 : State1),
9   contract_state bstate caddr = Some cstate1 /\
10  (* every reachable cstate1 of C1 corresponds to a contract-reachable cstate2 of C2 ...
11    *)
12  exists (cstate2 : State2),
13  (* 1. init_cstate2 is a valid initial cstate of C'  *)
14  cstate_reachable C2 cstate2 /\
15  (* 2. cstate and cstate' are related by state_morph. *)
16  cstate2 = state_morph C1 C2 f cstate1.

```

B.2.3 Right Morphism Induction

Listing B.13: The theorem of right morphism induction.

```

1 (* f : C1 -> C2, inducting on C2 *)
2 Theorem right_cm_induction:
3   forall (from to : State1) (f : ContractMorphism C1 C2),
4   ContractTrace C1 from to ->
5   ContractTrace C2 (state_morph C1 C2 f from) (state_morph C1 C2 f to).

```

B.3 Reasoning with Morphisms: Specification and Proof

B.3.1 Specifying a Contract Upgrade with Morphisms

Listing B.14: The Uranium Finance example.

```

1 (** Example 5.3.1:
2   This example recalls the Uranium Finance hack of 2021 due to an incorrect upgrade:
3   A constant 'k' was changed from 1_000 to 10_000 in all but one of its instances
4   in the contract.
5
6   This example illustrates how a contract upgrade can be *specified* using contract
7   morphisms, and uses that example.
8   We have formulated this example to be as general as possible.
9 *)
10

```

```

11 Section UpgradeSpec.
12 Context { Base : ChainBase }.
13
14 (** Assume we have a calculate_trade function which is used to calculate trades in
15     a smart contract. It takes some input in N and returns the output of the trade, also
16     in N. *)
16 Context { calculate_trade : N -> N }.
17
18 (** Assume that we have some storage type which keeps track of balances via
19     a function 'get_bal' *)
20 Context { storage : Type } '{ storage_ser : Serializable storage }
21     { get_bal : storage -> N }
22     (* Also assume some other relevant types of the contract. *)
23     { setup : Type } '{ setup_ser : Serializable setup }
24     { other_entrpoint : Type }.
25
26 (** We assume that entrpoint type includes a trade function, among other entrpoints. *)
27 Context { trade_data : Type } { trade_qty : trade_data -> N }.
28
29 Class Msg_Spec (T : Type) := {
30     trade : trade_data -> T ;
31     (* for any other entrpoint types *)
32     other : other_entrpoint -> option T ;
33 }.
34
35 Context { entrpoint : Type } '{ e_ser : Serializable entrpoint } '{ e_msg : Msg_Spec
36     entrpoint }.
37
38 (** And we assume anything in the entrpoint type is of the form 'trade n' or (roughly) '
39     other o'. *)
38 Definition msg_destruct : Prop :=
39     forall e,
40     (exists n, e = trade n) \/
41     (exists o, Some e = other o).
42 Context { e_msg_destruct : msg_destruct }.
43
44 (** Thus, the entrpoint type has this form:
45
46     Inductive entrpoint :=
47         | trade (qty : N)
48         | ... .
49 *)
50
51 (* final definitions of contract types *)
52 Definition error := N.
53 Definition result : Type := ResultMonad.result (storage * list ActionBody) error.
54
55 (***) Now suppose that we have a contract with those types ... *)
56 Context { C1 : Contract setup entrpoint storage error }.

```

```

57
58 (** such that get_bal changes according to calculate_trade, meaning that: *)
59 Definition spec_trade : Prop :=
60   forall cstate chain ctx trade_data cstate' acts,
61     (** for any successful call to the trade entrypoint of C1, *)
62     receive C1 chain ctx cstate (Some (trade trade_data)) = Ok(cstate', acts) ->
63     (** the balance in storage updates as follows. *)
64     get_bal cstate' =
65     get_bal cstate + calculate_trade (trade_qty trade_data).
66
67 Context { spec_trade : Prop }.
68
69 (** Now suppose that we have another calculate_trade function, this time which calculates
70     trades at one more digit of precision. *)
71 Definition round_down (n : N) := n / 10.
72
73 Context { calculate_trade_precise : N -> N }
74     (** (i.e. calculate_trade_precise rounds down to calculate_trade) *)
75     { calc_trade_coherence : forall n,
76       round_down (calculate_trade_precise n) =
77       calculate_trade (round_down n) }.
78
79 (** Suppose also that we have a round-down function on the storage type. *)
80
81 (** And that we have another contract, C2, ... *)
82 Context { C2 : Contract setup entrypoint storage error }.
83
84 (** but now trades are calculated in line with calculate_trade_precise. *)
85 Definition spec_trade_precise : Prop :=
86   forall cstate chain ctx trade_data cstate' acts,
87     (** ... meaning that for a successful call to the trade entrypoint of C2, *)
88     receive C2 chain ctx cstate (Some (trade trade_data)) = Ok(cstate', acts) ->
89     (** the balance held in storage goes up by calculate_trade_precise. *)
90     get_bal cstate' =
91     get_bal cstate + calculate_trade_precise (trade_qty trade_data).
92
93 Context { spec_trade_precise : Prop }.
94
95
96 (** Now, to specify the *upgrade* from C1 to C2, we specify that there exist some morphism
97     f : C1 -> C2 which satisfies the following conditions: *)
98 Context { st_morph : storage -> storage }
99     { state_rounds_down : forall st, get_bal (st_morph st) = round_down (get_bal st)
100     }.
101
102 (** 1. f rounds trades down when it maps inputs *)
103 Definition f_rcv_input_rounds_down (f : ContractMorphism C2 C1) : Prop :=
104   forall t', exists t,
105     (msg_morph C2 C1 f) (trade t') = trade t /\

```

```

105     trade_qty t = round_down (trade_qty t').
106
107 (** 2. aside from trade, f doesn't touch the other entypoints *)
108 Definition f_recv_input_other_equal (f : ContractMorphism C2 C1) : Prop :=
109     forall msg o,
110     (* for calls to all other entypoints, *)
111     msg = other o ->
112     (* f is the identity *)
113     option_map (msg_morph C2 C1 f) (other o) = other o.
114
115 (** 3. f rounds down on the storage, but doesn't touch anything else. *)
116 Definition f_state_morph (f : ContractMorphism C2 C1) : Prop :=
117     (state_morph C2 C1 f) = st_morph.
118
119 (** 4. f is the identity on error values *)
120 Definition f_recv_output_err (f : ContractMorphism C2 C1) : Prop :=
121     (error_morph C2 C1 f) = id.
122
123 (** 5. f is the identity on the setup *)
124 Definition f_init_id (f : ContractMorphism C2 C1) : Prop :=
125     (setup_morph C2 C1 f) = id.
126
127 (** Now we have a specification of the correct upgrade in terms of the existence of
128     a contract morphism. *)
129 Definition upgrade_spec (f : ContractMorphism C2 C1) : Prop :=
130     f_recv_input_rounds_down f /\
131     f_recv_input_other_equal f /\
132     f_state_morph f /\
133     f_recv_output_err f /\
134     f_init_id f.
135
136
137 (** The Upgrade Metaspecification.
138     To justify that upgrade_spec actually specifies a correct upgrade, we prove
139     the following result(s). *)
140
141 (***) Suppose there exists some f : C2 -> C1 satisfying upgrade_spec. *)
142 Context { f : ContractMorphism C2 C1 }
143     { is_upgrade_morph : upgrade_spec f }.
144
145 (* All states of C2 relate to equivalent states of C1 by rounding down *)
146 Theorem rounding_down_invariant bstate caddr (trace : ChainTrace empty_state bstate):
147     (* Forall reachable states with contract at caddr, *)
148     env_contracts bstate caddr = Some (C2 : WeakContract) ->
149     (* cstate is the state of the contract AND *)
150     exists (cstate' cstate : storage),
151     contract_state bstate caddr = Some cstate' /\
152     (* cstate is contract-reachable for C1 AND *)
153     cstate_reachable C1 cstate /\

```

```

154   (* such that for cstate, the state of C1 in bstate,
155       the balance in cstate is rounded-down from the balance of cstate' *)
156   get_bal cstate = round_down (get_bal cstate').
157 Proof.
158   intros c_at_caddr.
159   pose proof (left_cm_induction f bstate caddr trace c_at_caddr)
160   as H_cm_ind.
161   destruct H_cm_ind as [cstate' [contr_cstate' [cstate [reach H_cm_ind]]]].
162   exists cstate', cstate.
163   repeat split; auto.
164   cbn in H_cm_ind.
165   rewrite H_cm_ind.
166   destruct is_upgrade_morph as [_ [_ [f_state [_ _]]]].
167   now rewrite f_state.
168 Qed.
169
170 End UpgradeSpec.

```

B.3.2 Adding Features and Backwards Compatibility

Listing B.15: An example of backwards compatibility.

```

1 Section BackwardsCompatible.
2 Context { Base : ChainBase }.
3 Set Primitive Projections.
4 Set Nonrecursive Elimination Schemes.
5
6 (** The initial contract C *)
7 (* contract types definition *)
8 Inductive entrypoint1 := | incr (u : unit).
9 Definition storage := N.
10 Definition setup := N.
11 Definition error := N.
12 Definition result : Type := ResultMonad.result (storage * list ActionBody) error.
13
14 Section Serialization.
15   Global Instance entrypoint1_serializable : Serializable entrypoint1 :=
16     Derive Serializable entrypoint1_rect<incr>.
17 End Serialization.
18
19 (* init function definition *)
20 Definition init (_ : Chain) (_ : ContractCallContext) (n : setup) :
21   ResultMonad.result storage N :=
22   Ok (n).
23
24 (* receive function definition *)
25 Definition receive (_ : Chain) (_ : ContractCallContext) (n : storage)
26   (msg : option entrypoint1) : result :=

```



```

27     match msg with
28     | Some (incr _) => Ok (n + 1, [])
29     | None => Err 0
30     end.
31
32 (* construct the contract *)
33 Definition C1 : Contract setup entrypoint1 storage error :=
34     build_contract init receive.
35
36
37 (** The updated contract C' *)
38 (* contract types definition *)
39 Inductive entrypoint2 := | incr' (u : unit) | decr (u : unit).
40
41 Section Serialization.
42     Global Instance entrypoint2_serializable : Serializable entrypoint2 :=
43         Derive Serializable entrypoint2_rect<incr',decr>.
44 End Serialization.
45
46 (* receive function definition *)
47 Definition receive' (_ : Chain) (_ : ContractCallContext) (n : storage)
48     (msg : option entrypoint2) : result :=
49     match msg with
50     | Some (incr' _) => Ok (n + 1, [])
51     | Some (decr _) => Ok (n - 1, [])
52     | None => Err 0
53     end.
54
55 (* construct the contract *)
56 Definition C2 : Contract setup entrypoint2 storage error :=
57     build_contract init receive'.
58
59
60 (** The contract morphism confirming backwards compatibility *)
61 Definition msg_morph (e : entrypoint1) : entrypoint2 :=
62     match e with | incr _ => incr' tt end.
63 Definition setup_morph : setup -> setup := id.
64 Definition state_morph : storage -> storage := id.
65 Definition error_morph : error -> error := id.
66
67 (* the coherence results *)
68 Lemma init_coherence : forall c ctx s,
69     result_functor state_morph error_morph (init c ctx s) =
70     init c ctx (setup_morph s).
71 Proof. auto. Qed.
72
73 Lemma recv_coherence : forall c ctx st op_msg,
74     result_functor (fun' (st, l) => (state_morph st, l)) error_morph (receive c ctx st
75     op_msg) =

```

```

75   receive' c ctx (state_morph st) (option_map msg_morph op_msg).
76 Proof.
77   intros.
78   unfold result_functor, msg_morph, state_morph, error_morph.
79   induction op_msg; auto.
80   now destruct a.
81 Qed.
82
83 (* construct the morphism *)
84 Definition f : ContractMorphism C1 C2 :=
85   build_contract_morphism C1 C2 setup_morph msg_morph state_morph error_morph
86     init_coherence recv_coherence.
87
88
89 (* This theorem shows a strong notion of backwards compatibility because there is
90    an embedding of the old contract into the new *)
91 Lemma embedding : is_inj_cm f.
92 Proof.
93   unfold is_inj_cm; unfold is_inj.
94   repeat split; intros.
95   - cbn in H.
96     now unfold setup_morph in H.
97   - now destruct a, a', u, u0.
98   - cbn in H.
99     now unfold state_morph in H.
100  - cbn in H.
101    now unfold error_morph in H.
102 Qed.
103
104 (** Theorem:
105     All reachable states have a corresponding reachable state, related by the
106     *embedding* f. *)
107 Theorem injection_invariant bstate caddr (trace : ChainTrace empty_state bstate):
108   (* Forall reachable states with contract C1 at caddr, *)
109   env_contracts bstate caddr = Some (C1 : WeakContract) ->
110   (* forall reachable states of C1 cstate, there's a corresponding reachable state
111      cstate' of C2, related by the injection *)
112   exists (cstate' cstate : storage),
113   contract_state bstate caddr = Some cstate /\
114   (* cstate' is a contract-reachable state of C2 *)
115   cstate_reachable C2 cstate' /\
116   (* .. equal to cstate *)
117   cstate' = cstate.
118 Proof.
119   intros c_at_caddr.
120   pose proof (left_cm_induction f bstate caddr trace c_at_caddr)
121   as H_cm_ind.
122   destruct H_cm_ind as [cstate [cstate_c [cstate' [reach H_cm_ind]]]].
123   now exists cstate', cstate.

```

```

124 Qed.
125
126 End BackwardsCompatible.

```

B.3.3 Transporting Hoare-Like Properties Over a Morphism

Listing B.16: Transporting a Hoare-like property of partial correctness over a contract morphism.

```

1 Section TransportHoare.
2 Context { Base : ChainBase }.
3
4 Context {setup storage error : Type}
5   '{setup_ser : Serializable setup}  '{stor_ser : Serializable storage}
6   '{err_ser : Serializable error}  '{storage_state : @State_Spec Base storage}
7   '{err_err : Error_Spec error}.
8
9 Set Primitive Projections.
10 Set Nonrecursive Elimination Schemes.
11
12 Inductive entrypoint :=
13 | Pool : pool_data -> entrypoint
14 | Unpool : unpool_data -> entrypoint
15 | Null : trade_data -> entrypoint.
16
17 Inductive entrypoint' :=
18 | Pool' : pool_data -> entrypoint'
19 | Unpool' : unpool_data -> entrypoint'
20 | Trade' : trade_data -> entrypoint'.
21
22 Context { other_entrypoint : Type }.
23
24 Definition e_pool (p : pool_data) : entrypoint := Pool p.
25 Definition e_unpool (u : unpool_data) : entrypoint := Unpool u.
26 Definition e_trade (t : trade_data) : entrypoint := Null t.
27 Definition e_other (o : other_entrypoint) : option entrypoint := None.
28 Global Instance entrypoint_msg_spec : Msg_Spec entrypoint := {
29   pool := e_pool ;
30   unpool := e_unpool ;
31   trade := e_trade ;
32   other := e_other ;
33 }.
34
35 Definition e'_pool (p : pool_data) : entrypoint' := Pool' p.
36 Definition e'_unpool (u : unpool_data) : entrypoint' := Unpool' u.
37 Definition e'_trade (t : trade_data) : entrypoint' := Trade' t.
38 Definition e'_other (o : other_entrypoint) : option entrypoint' := None.
39 Global Instance entrypoint'_msg_spec : @Msg_Spec Base other_entrypoint entrypoint' := {
40   pool := e'_pool ;

```



```

87 Axiom init_coherence_pf : init_coherence_prop.
88
89 Definition recv_coherence_prop : Prop := forall c ctx st op_msg,
90   result_functor (fun ' (st, l) => (id st, l)) id
91     (receive C1 c ctx st op_msg) =
92     receive C2 c ctx (id st) (option_map embed_entrypoint op_msg).
93 Axiom recv_coherence_pf : recv_coherence_prop.
94
95
96 (* construct a contract morphism *)
97 Definition f : ContractMorphism C1 C2 := { |
98   setup_morph := id ;
99   msg_morph := embed_entrypoint ;
100   state_morph := id ;
101   error_morph := id ;
102   (* coherence *)
103   init_coherence := init_coherence_pf ;
104   recv_coherence := recv_coherence_pf ;
105 |}.
106
107
108 (* TODO get rid of this *)
109 Tactic Notation "is_sp_destruct" :=
110   match goal with
111   | is_sp : is_structured_pool _ |- _ =>
112     unfold is_structured_pool in is_sp;
113     destruct is_sp as [none_fails_pf is_sp'];
114     destruct is_sp' as [msg_destruct_pf is_sp'];
115     (* isolate the pool entrypoint specification *)
116     destruct is_sp' as [pool_entrypoint_check_pf is_sp'];
117     destruct is_sp' as [pool_emits_txns_pf is_sp'];
118     destruct is_sp' as [pool_increases_tokens_held_pf is_sp'];
119     destruct is_sp' as [pool_rates_unchanged_pf is_sp'];
120     destruct is_sp' as [pool_outstanding_pf is_sp'];
121     (* isolate the unpool entrypoint specification *)
122     destruct is_sp' as [unpool_entrypoint_check_pf is_sp'];
123     destruct is_sp' as [unpool_entrypoint_check_2_pf is_sp'];
124     destruct is_sp' as [unpool_emits_txns_pf is_sp'];
125     destruct is_sp' as [unpool_decreases_tokens_held_pf is_sp'];
126     destruct is_sp' as [unpool_rates_unchanged_pf is_sp'];
127     destruct is_sp' as [unpool_outstanding_pf is_sp'];
128     (* isolate the trade entrypoint specification *)
129     destruct is_sp' as [trade_entrypoint_check_pf is_sp'];
130     destruct is_sp' as [trade_entrypoint_check_2_pf is_sp'];
131     destruct is_sp' as [trade_pricing_formula_pf is_sp'];
132     destruct is_sp' as [trade_update_rates_pf is_sp'];
133     destruct is_sp' as [trade_update_rates_formula_pf is_sp'];
134     destruct is_sp' as [trade_emits_transfers_pf is_sp'];
135     destruct is_sp' as [trade_tokens_held_update_pf is_sp'];

```

```

136     destruct is_sp' as [trade_outstanding_update_pf is_sp'];
137     destruct is_sp' as [trade_pricing_pf is_sp'];
138     destruct is_sp' as [trade_amounts_nonnegative_pf is_sp'];
139     (* isolate the specification of all other entrypoints *)
140     destruct is_sp' as [other_rates_unchanged_pf is_sp'];
141     destruct is_sp' as [other_balances_unchanged_pf is_sp'];
142     destruct is_sp' as [other_outstanding_unchanged_pf is_sp'];
143     (* isolate the specification of calc_rx' and calc_delta_y *)
144     destruct is_sp' as [update_rate_stays_positive_pf is_sp'];
145     destruct is_sp' as [rate_decrease_pf is_sp'];
146     destruct is_sp' as [rates_balance_pf is_sp'];
147     destruct is_sp' as [rates_balance_2_pf is_sp'];
148     destruct is_sp' as [trade_slippage_pf is_sp'];
149     destruct is_sp' as [trade_slippage_2_pf is_sp'];
150     destruct is_sp' as [arbitrage_lt_pf is_sp'];
151     destruct is_sp' as [arbitrage_gt_pf is_sp'];
152     (* isolate the initialization specification *)
153     destruct is_sp' as [initialized_with_positive_rates_pf is_sp'];
154     destruct is_sp' as [initialized_with_zero_balance_pf is_sp'];
155     destruct is_sp' as [initialized_with_zero_outstanding_pf is_sp'];
156     destruct is_sp' as [initialized_with_init_rates_pf initialized_with_pool_token_pf]
157 end.
158
159
160 (* Prove the following about C1 using C2 and f *)
161 Theorem pullback_unpool_emits_txns : unpool_emits_txns C1.
162 Proof.
163   pose proof is_sp as is_sp.
164   is_sp_destruct.
165   pose proof recv_coherence_pf as recv_eq.
166   unfold recv_coherence_prop in recv_eq.
167   unfold unpool_emits_txns.
168   intros * recv_some.
169   pose proof (unpool_emits_txns_pf cstate chain ctx msg_payload cstate' acts )
170   as sp_spec_result.
171   pose proof (recv_eq chain ctx cstate (Some (unpool msg_payload))) as recv_eq.
172   unfold result_functor in recv_eq.
173   rewrite recv_some in recv_eq.
174   unfold embed_entrypoint in recv_eq.
175   cbn in recv_eq.
176   pose proof (eq_sym recv_eq) as recv_eq'.
177   now apply sp_spec_result in recv_eq'.
178 Qed.
179
180 End TransportHoare.

```

B.4 A Mathematical Characterization of Contract Upgrades

Listing B.17: The definitions of the various contracts in the upgradeability decomposition.

```
1 Section ContractDefinitions.
2 (** Contract types definition *)
3 (** Main contract C *)
4 Inductive entrypoint :=
5 | next (u : unit)
6 | upgrade_fun (s' : N -> N).
7 Record storage := { n : N ; s : N -> N ; }.
8 Definition setup := storage.
9 Definition error := N.
10 Definition result : Type := ResultMonad.result (storage * list ActionBody) error.
11
12 (** Base' contract C_b' (to be C_b) *)
13 Inductive entrypoint_b' :=
14 | upgrade_fun_b (s' : N -> N).
15 Record storage_b := { s_b : N -> N ; }.
16 Definition setup_b := storage_b.
17 Definition error_b := N.
18 Definition result_b : Type := ResultMonad.result (storage_b * list ActionBody) error_b.
19
20 (** Version contracts C_f v, for v : storage_b *)
21 Inductive entrypoint_version := | next_f (u : unit).
22 Record storage_version := { n_f : N }.
23 Definition entrypoint_f : storage_b -> Type := fun v => entrypoint_version.
24 Definition storage_f : storage_b -> Type := fun v => storage_version.
25 Definition setup_f : storage_b -> Type := fun v => N.
26 Definition error_f : storage_b -> Type := fun v => N.
27 Definition result_f : storage_b -> Type :=
28   fun v => ResultMonad.result ((storage_f v) * list ActionBody) (error_f v).
29
30 Section Serialization.
31   Section SerializeFunctionType.
32     Context `{Serializable A} `{Serializable B}.
33
34     (* for simplicity, we assume that function types are serializable *)
35     Definition serialize_nn (f : A -> B) : SerializedValue. Admitted.
36     Definition deserialize_nn (val : SerializedValue) : option (A -> B). Admitted.
37     Lemma deserialize_serialize_nn (f : A -> B) :
38       deserialize_nn (serialize_nn f) = Some f.
39     Admitted.
40
41     Global Instance nn_serializable : Serializable (A -> B) :=
42       { | serialize := serialize_nn ;
43         deserialize := deserialize_nn ;
44         deserialize_serialize := deserialize_serialize_nn ; | }.
45   End SerializeFunctionType.
```

```

46
47   (* assuming function types are serializable, ... *)
48   Global Instance storage_serializable : Serializable storage :=
49     Derive Serializable storage_rect<Build_storage>.
50   Global Instance entrypoint_serializable : Serializable entrypoint :=
51     Derive Serializable entrypoint_rect<next,upgrade_fun>.
52   Global Instance storage_b_serializable : Serializable storage_b :=
53     Derive Serializable storage_b_rect<Build_storage_b>.
54   Global Instance entrypoint_s_serializable : Serializable entrypoint_b' :=
55     Derive Serializable entrypoint_b'_rect<upgrade_fun_b>.
56   Global Instance storage_f_serializable v : Serializable (storage_f v) :=
57     Derive Serializable storage_version_rect<Build_storage_version>.
58   Global Instance entrypoint_f_serializable v : Serializable (entrypoint_f v) :=
59     Derive Serializable entrypoint_version_rect<next_f>.
60 End Serialization.
61
62 (** Contract, init, and receive definitions *)
63
64 (** Main contract C *)
65 (* init *)
66 Definition init (_ : Chain)
67   (_ : ContractCallContext)
68   (init_state : setup)
69   : ResultMonad.result storage N :=
70   Ok (init_state).
71
72 (* receive *)
73 Definition receive (_ : Chain)
74   (_ : ContractCallContext)
75   (storage : storage)
76   (msg : option entrypoint)
77   : result :=
78   match msg with
79   | Some (next _) =>
80     let st := {| n := storage.(s) storage.(n) ; s := storage.(s) ; |} in
81     Ok (st, [])
82   | Some (upgrade_fun s') =>
83     let st := {| n := storage.(n) ; s := s' ; |} in
84     Ok (st, [])
85   | None => Err 0
86   end.
87
88 (* the contract C *)
89 Definition C : Contract setup entrypoint storage error :=
90   build_contract init receive.
91
92 (** Base' Contract (to be the base contract) *)
93 (* init *)
94 Definition init_b' (_ : Chain)

```



```

95         (_ : ContractCallContext)
96         (init_state_b : setup_b)
97         : ResultMonad.result storage_b N :=
98     Ok (init_state_b).
99
100 (* receive *)
101 Definition receive_b' (_ : Chain)
102     (_ : ContractCallContext)
103     (_ : storage_b)
104     (msg : option entrypoint_b')
105     : result_b :=
106     match msg with
107     | Some (upgrade_fun_b s_new) =>
108         let st := {| s_b := s_new ; |} in
109         Ok (st, [])
110     | None => Err 0
111     end.
112
113 (* the contract C_b' *)
114 Definition C_b' : Contract setup_b entrypoint_b' storage_b error_b :=
115     build_contract init_b' receive_b'.
116
117 (** For the morphisms, we define the base contract *)
118 Definition C_b := pointed_contract C_b'.
119 Definition entrypoint_b := (entrypoint_b' + unit)%type.
120
121 (** The family of version contracts C_f *)
122 (* init *)
123 Definition init_f (version : storage_b)
124     (_ : Chain)
125     (_ : ContractCallContext)
126     (n : setup_f version)
127     : ResultMonad.result (storage_f version) (error_f version) :=
128     let storage_f := {| n_f := n ; |} in
129     Ok (storage_f).
130
131 (* receive *)
132 Definition receive_f (version : storage_b)
133     (_ : Chain)
134     (_ : ContractCallContext)
135     (storage_f : storage_f version)
136     (msg : option (entrypoint_f version))
137     : result_f version :=
138     match msg with
139     | Some (next_f _) =>
140         let st := {| n_f := version.(s_b) storage_f.(n_f) ; |} in
141         Ok (st, [])
142     | None => Err 0
143     end.

```

```

144
145 (* the contract C_f *)
146 Definition C_f (version : storage_b) : Contract (setup_f version) (entrypoint_f version) (
    storage_f version) (error_f version) :=
147   build_contract (init_f version) (receive_f version).
148
149 End ContractDefinitions.

```

B.4.1 Isolating Mutable and Immutable Parts

Listing B.18: The quotient onto the base contract.

```

1  (** f_p : C ->> C_b *)
2  Section Quotient.
3  Definition zero_fn : N -> N := (fun (x : N) => 0).
4  Definition null_storage_b : storage_b := {| s_b := zero_fn |}.
5  Definition null_setup_b : setup_b := {| s_b := zero_fn |}.
6
7  (* component morphisms *)
8  Definition msg_morph_p (e : entrypoint) : entrypoint_b :=
9    match e with
10   | next _ => inr tt (* not upgrade functionality *)
11   | upgrade_fun s' => inl (upgrade_fun_b s') (* corresponds to an upgrade *)
12   end.
13 Definition state_morph_p : storage -> storage_b := (fun (x : storage) => {| s_b := x.(s) ;
    |}).
14 Definition setup_morph_p : setup -> setup_b := (fun (x : setup) => {| s_b := x.(s) |}).
15 Definition error_morph_p : error -> error_b := (fun (x : error) => x).
16
17 (* the coherence results *)
18 Lemma init_coherence_p :
19   init_coherence_prop C C_b
20     setup_morph_p state_morph_p error_morph_p.
21 Proof. unfold init_coherence_prop. auto. Qed.
22
23 Lemma recv_coherence_p :
24   recv_coherence_prop C C_b
25     msg_morph_p state_morph_p error_morph_p.
26 Proof.
27   unfold recv_coherence_prop.
28   intros.
29   unfold result_functor.
30   cbn.
31   destruct op_msg; cbn.
32   - unfold msg_morph_p.
33     destruct e eqn:H_e.
34     + now destruct st.
35     + now unfold state_morph_p.

```

```

36   - now unfold error_morph_p.
37 Qed.
38
39 (* construct the morphism *)
40 Definition f_p : ContractMorphism C C_b :=
41   build_contract_morphism C C_b
42     (* the morphisms *)
43     setup_morph_p msg_morph_p state_morph_p error_morph_p
44     (* coherence *)
45     init_coherence_p recv_coherence_p.
46
47 End Quotient.

```

Listing B.19: The family of contract embeddings.

```

1  (** f_i : forall v, C_f v -> C *)
2  Section Embedding.
3
4  Definition msg_morph_i (v : storage_b) (e : entrypoint_f v) : entrypoint :=
5    match e with
6    | next_f _ => next tt
7    end.
8
9  Definition setup_morph_i (v : storage_b) (st_f : setup_f v) : setup := {|
10    n := st_f ;
11    s := s_b v ; |}.
12
13  Definition state_morph_i (v : storage_b) (st_f : storage_f v) : storage :=
14    {| n := st_f.(n_f) ; s := s_b v ; |}.
15
16  Definition error_morph_i (v : storage_b) : error_f v -> error := id.
17
18  (* the coherence results *)
19
20  Lemma init_coherence_i (v : storage_b) :
21    init_coherence_prop (C_f v) C (setup_morph_i v) (state_morph_i v) (error_morph_i v).
22  Proof. unfold init_coherence_prop. auto. Qed.
23
24  Lemma recv_coherence_i (v : storage_b) :
25    recv_coherence_prop (C_f v) C (msg_morph_i v) (state_morph_i v) (error_morph_i v).
26  Proof.
27    unfold recv_coherence_prop.
28    intros.
29    destruct op_msg; auto.
30    now destruct e.
31  Qed.
32
33  (* construct the morphism *)
34
35  Definition fi_param (v : storage_b) : ContractMorphism (C_f v) C :=
36    build_contract_morphism (C_f v) C
37      (* the morphisms *)
38      (setup_morph_i v) (msg_morph_i v) (state_morph_i v) (error_morph_i v)
39      (* coherence results *)

```

```

35      (init_coherence_i v) (recv_coherence_i v).
36
37 End Embedding.

```

B.4.2 Decomposing Upgradeability

Listing B.20: A proof of the upgradeability decomposition.

```

1  (** Here we prove the core result that the family of versioned contracts and the
2     base contract given above provide an upgradeability decomposition of our contract C.
3     *)
4  Section UpgradeabilityDecomposition.
5  Definition extract_version (m : entrypoint_b') : storage_b :=
6     match m with | upgrade_fun_b s_b => {| s_b := s_b |} end.
7  Definition new_version_state old_v msg (st : storage_f old_v) : storage_f (extract_version
8     msg) := st.
9
10 Theorem upgradeability_decomposition :
11     upgradeability_decomposition fi_param f_p extract_version new_version_state.
12 Proof.
13   unfold upgradeability_decomposition.
14   repeat split.
15   (* msg_required *)
16   - unfold ContractMorphisms.msg_required.
17     intros.
18     now exists 0.
19   (* init_versioned *)
20   - unfold init_versioned.
21     intros ? ? ? s init_ok.
22     destruct init_state as [n_i s_i].
23     unfold is_versioned.
24     now exists {| s_b := s_i |}, {| n_f := n_i |}.
25   (* msg_decomposable *)
26   (* -> *)
27   - simpl.
28     intro H_null.
29     unfold msg_morph_p in H_null.
30     destruct m; try inversion H_null.
31     destruct u.
32     now exists (next_f tt).
33   (* <- *)
34   - simpl.
35     intro H_preim.
36     unfold msg_morph_i in H_preim.
37     destruct H_preim as [m' H_preim].
38     destruct m'.
39     unfold msg_morph_p.
40     now destruct m; try inversion H_preim.

```

```

39   (* states categorized *)
40   (* -> *)
41   -   simpl.
42       intro H_preim.
43       destruct H_preim as [st_f H_preim].
44       destruct st_f as [nf].
45       unfold state_morph_i in H_preim.
46       simpl in H_preim.
47       destruct st as [n s].
48       now inversion H_preim.
49   (* <- *)
50   -   simpl.
51       unfold state_morph_p, state_morph_i.
52       destruct c_version as [sb], st as [n s].
53       cbn.
54       intro H.
55       inversion H.
56       now exists {| n_f := n |}.
57   (* version transition *)
58   -   unfold version_transition.
59       simpl.
60       unfold msg_morph_p, state_morph_i.
61       intros * H_cstate_preim ? ? ? ? ? rcv_some is_upgrade_msg.
62       destruct msg; try inversion is_upgrade_msg.
63       inversion rcv_some.
64       f_equal.
65       unfold new_version_state.
66       now rewrite H_cstate_preim.
67   Qed.
68
69   End UpgradeabilityDecomposition.

```

Listing B.21: Some results provable due to the decomposition.

```

1   Section Decomposition.
2
3   (** Theorem: all reachable contract states are versioned according to this indexing *)
4   Theorem all_states_versioned :
5     forall bstate caddr (trace : ChainTrace empty_state bstate),
6     (* where C is at caddr with state cstate, *)
7     env_contracts bstate caddr = Some (C : WeakContract) ->
8     exists (cstate : storage),
9     contract_state bstate caddr = Some cstate /\
10    (* then every contract state cstate is versioned *)
11    is_versioned fi_param cstate.
12   Proof.
13     intros * ? c_at_caddr.
14     pose proof (versioned_invariant fi_param f_p extract_version new_version_state bstate
15       caddr trace c_at_caddr).

```

```

15   destruct H as [cstate [cstate_st versioned]].
16   exists cstate.
17   split; auto.
18   apply (versioned upgradeability_decomposition).
19   Qed.
20
21   (** Theorem: Versioning moves along as we've described it. *)
22   Theorem upgrade_decomposed :
23     (* Forall versioned contract states (incl. all reachable states), *)
24     forall cstate c_version cstate_f,
25     cstate = state_morph (C_f c_version) C (fi_param c_version) cstate_f ->
26     (* And forall calls to the versioned contract *)
27     forall chain ctx m new_state new_acts,
28     receive chain ctx cstate (Some m) = Ok (new_state, new_acts) ->
29     (* the contract state either stays within this version *)
30     (exists cstate_f', new_state = state_morph (C_f c_version) C (fi_param c_version)
31      cstate_f') \/
32     (* or it moves onto a new version as we've described it. *)
33     (exists c_version' cstate_f',
34      new_state = state_morph (C_f c_version') C (fi_param c_version') cstate_f').
35   Proof.
36     intros * cstate_preim * recv_some.
37     apply (upgradeability_decomposed fi_param f_p extract_version new_version_state cstate
38           c_version
39           cstate_f upgradeability_decomposition cstate_preim chain ctx m new_state new_acts
40           recv_some).
41   Qed.
42
43   End Decomposition.

```

B.4.3 Upgradeable Contracts are Fiber Bundles: A Digression

Listing B.22: The fiber bundle of our upgradeability decomposition admits a section.

```

1 Section FiberBundle.
2
3   (* A section of p is a morphism C_b' -> C such that C_b' -> C -> C_b = C_b' -> C_b *)
4   Definition setup_morph_s n : setup_b -> setup := (fun (x : setup_b) => {| n := n ; s := x
5     .(s_b) |}).
6   Definition msg_morph_s (e : entrypoint_b') : entrypoint :=
7     match e with | upgrade_fun_b s' => upgrade_fun s' end.
8   Definition state_morph_s n : storage_b -> storage :=
9     (fun (x : storage_b) => {| n := n ; s := x.(s_b) ; |}).
10  Definition error_morph_s : error_b -> error := (fun (x : error_b) => x).
11
12  Definition fp_rinv (n : N) : ContractMorphism C_b' C.
13  Proof.
14    apply (build_contract_morphism C_b' C

```

```

14      (setup_morph_s n) msg_morph_s (state_morph_s n) error_morph_s).
15  -   intros.
16      simpl.
17      now unfold state_morph_s, setup_morph_s, init.
18  -   intros.
19      simpl.
20      unfold result_functor, state_morph_s, error_morph_s, receive.
21      simpl.
22      destruct op_msg; cbn; auto.
23      now destruct e.
24 Defined.
25
26 Theorem p_rinv_section (n : N) :
27   compose_cm f_p (fp_rinv n) = pointed_include C_b'.
28 Proof.
29   unfold compose_cm, pointed_include.
30   apply eq_cm; cbn.
31   -   now unfold setup_morph_p, setup_morph_s.
32   -   unfold msg_morph_p, msg_morph_s.
33       apply functional_extensionality.
34       intro e.
35       now destruct e.
36   -   now unfold state_morph_p, state_morph_s.
37   -   now unfold error_morph_p, error_morph_s.
38 Qed.
39
40 End FiberBundle.

```


Appendix C

Proofs and Definitions of Chapter 6

C.1 Bisimulations of Contracts

C.1.1 Contract Trace Morphisms

Listing C.1: The formal definition of a contract trace morphism.

```
1 Section ContractTraceMorphism.
2 Context `{Serializable Setup1} `{Serializable Msg1} `{Serializable State1} `{Serializable
   Error1}
3     `{Serializable Setup2} `{Serializable Msg2} `{Serializable State2} `{Serializable
   Error2}.
4
5 Record ContractTraceMorphism
6   (C1 : Contract Setup1 Msg1 State1 Error1)
7   (C2 : Contract Setup2 Msg2 State2 Error2) :=
8   build_ct_morph {
9     (* a function of state types *)
10    ct_state_morph : State1 -> State2 ;
11    (* init state C1 -> init state C2 *)
12    genesis_fixpoint : forall init_cstate,
13      is_genesis_cstate C1 init_cstate ->
14      is_genesis_cstate C2 (ct_state_morph init_cstate) ;
15    (* coherence *)
16    cstep_morph : forall state1 state2,
17      ContractStep C1 state1 state2 ->
18      ContractStep C2 (ct_state_morph state1) (ct_state_morph state2) ;
19  }.
20
21 End ContractTraceMorphism.
```

Listing C.2: The identity contract trace morphism.

```

1 Section IdentityCTMorphism.
2 Context `{Serializable Setup1} `{Serializable Msg1} `{Serializable State1} `{Serializable
   Error1}.
3
4 Definition id_genesis_fixpoint (C : Contract Setup1 Msg1 State1 Error1)
5   init_cstate
6   (gen_C : is_genesis_cstate C init_cstate) :
7   is_genesis_cstate C (id init_cstate) :=
8   gen_C.
9
10 Definition id_cstep_morph (C : Contract Setup1 Msg1 State1 Error1)
11   state1 state2
12   (step : ContractStep C state1 state2) :
13   ContractStep C (id state1) (id state2) :=
14   step.
15
16 Definition id_ctm (C : Contract Setup1 Msg1 State1 Error1) : ContractTraceMorphism C C :=
17   { |
18     ct_state_morph := id ;
19     genesis_fixpoint := id_genesis_fixpoint C ;
20     cstep_morph := id_cstep_morph C ;
21   |}.
22 End IdentityCTMorphism.

```

Listing C.3: Equality of contract trace morphisms.

```

1 Section EqualityOfCTMorphisms.
2 Context `{Serializable Setup1} `{Serializable Msg1} `{Serializable State1} `{Serializable
   Error1}
3   `{Serializable Setup2} `{Serializable Msg2} `{Serializable State2} `{Serializable
   Error2}.
4
5 Lemma eq_ctm_dep
6   (C1 : Contract Setup1 Msg1 State1 Error1)
7   (C2 : Contract Setup2 Msg2 State2 Error2)
8   (ct_st_m : State1 -> State2)
9   (gen_fix1 gen_fix2 : forall init_cstate,
10    is_genesis_cstate C1 init_cstate ->
11    is_genesis_cstate C2 (ct_st_m init_cstate))
12   (cstep_m1 cstep_m2 : forall state1 state2,
13    ContractStep C1 state1 state2 ->
14    ContractStep C2 (ct_st_m state1) (ct_st_m state2)) :
15   cstep_m1 = cstep_m2 ->
16   { | ct_state_morph := ct_st_m ;
17     genesis_fixpoint := gen_fix1 ;
18     cstep_morph := cstep_m1 ; |}
19   =

```

```

20   {| ct_state_morph := ct_st_m ;
21       genesis_fixpoint := gen_fix2 ;
22       cstep_morph := cstep_m2 ; |}.
23 Proof.
24   intro cstep_equiv.
25   subst.
26   f_equal.
27   apply proof_irrelevance.
28 Qed.
29
30 End EqualityOfCTMorphisms.

```

Listing C.4: Composition of contract trace morphisms.

```

1 Section CTMorphismComposition.
2 Context '{Serializable Setup1} '{Serializable Msg1} '{Serializable State1} '{Serializable
   Error1}
3       '{Serializable Setup2} '{Serializable Msg2} '{Serializable State2} '{Serializable
   Error2}
4       '{Serializable Setup3} '{Serializable Msg3} '{Serializable State3} '{Serializable
   Error3}
5   {C1 : Contract Setup1 Msg1 State1 Error1}
6   {C2 : Contract Setup2 Msg2 State2 Error2}
7   {C3 : Contract Setup3 Msg3 State3 Error3}.
8
9 Definition genesis_compose (m2 : ContractTraceMorphism C2 C3) (m1 : ContractTraceMorphism
   C1 C2)
10   init_cstate (gen_C1 : is_genesis_cstate C1 init_cstate) :
11   is_genesis_cstate C3 (compose (ct_state_morph C2 C3 m2) (ct_state_morph C1 C2 m1)
   init_cstate) :=
12   match m2 with
13   | build_ct_morph _ _ _ gen_fix2 step2 =>
14       match m1 with
15       | build_ct_morph _ _ _ gen_fix1 step1 =>
16           gen_fix2 _ (gen_fix1 _ gen_C1)
17       end
18   end.
19
20 Definition cstep_compose (m2 : ContractTraceMorphism C2 C3) (m1 : ContractTraceMorphism C1
   C2)
21   state1 state2 (step : ContractStep C1 state1 state2) :
22   ContractStep C3
23       (compose (ct_state_morph C2 C3 m2) (ct_state_morph C1 C2 m1) state1)
24       (compose (ct_state_morph C2 C3 m2) (ct_state_morph C1 C2 m1) state2) :=
25   match m2, m1 with
26   | build_ct_morph _ _ _ _ step2, build_ct_morph _ _ _ _ step1 =>
27       step2 _ _ (step1 _ _ step)
28   end.
29

```

```

30 Definition compose_ctm
31   (m2 : ContractTraceMorphism C2 C3)
32   (m1 : ContractTraceMorphism C1 C2) : ContractTraceMorphism C1 C3 :=
33   { |
34     ct_state_morph := compose (ct_state_morph _ _ m2) (ct_state_morph _ _ m1) ;
35     genesis_fixpoint := genesis_compose m2 m1 ;
36     cstep_morph := cstep_compose m2 m1 ;
37   |}.
38
39 End CTMorphismComposition.

```

Listing C.5: Some results about contract trace morphism composition, including that composition is associative, and that left and right composition with the identity is trivial.

```

1 Section CTMorphismCompositionResults.
2 Context `{Serializable Setup1} `{Serializable Msg1} `{Serializable State1} `{Serializable
   Error1}
3   `{Serializable Setup2} `{Serializable Msg2} `{Serializable State2} `{Serializable
   Error2}
4   `{Serializable Setup3} `{Serializable Msg3} `{Serializable State3} `{Serializable
   Error3}
5   `{Serializable Setup4} `{Serializable Msg4} `{Serializable State4} `{Serializable
   Error4}
6   { C1 : Contract Setup1 Msg1 State1 Error1 }
7   { C2 : Contract Setup2 Msg2 State2 Error2 }
8   { C3 : Contract Setup3 Msg3 State3 Error3 }
9   { C4 : Contract Setup4 Msg4 State4 Error4 }.
10
11 (* Composition is associative *)
12 Lemma compose_ctm_assoc
13   (f : ContractTraceMorphism C1 C2)
14   (g : ContractTraceMorphism C2 C3)
15   (h : ContractTraceMorphism C3 C4) :
16   compose_ctm h (compose_ctm g f) =
17   compose_ctm (compose_ctm h g) f.
18 Proof. now destruct f, g, h. Qed.
19
20 (* Composition with the identity is trivial *)
21 Lemma compose_id_ctm_left (f : ContractTraceMorphism C1 C2) :
22   compose_ctm (id_ctm C2) f = f.
23 Proof. now destruct f. Qed.
24
25 Lemma compose_id_ctm_right (f : ContractTraceMorphism C1 C2) :
26   compose_ctm f (id_ctm C1) = f.
27 Proof. now destruct f. Qed.
28
29 End CTMorphismCompositionResults.

```

C.1.2 The Lifting Theorem

Listing C.6: The lifting theorem.

```
1 Section LiftingTheorem.
2 Context `{Serializable Setup1} `{Serializable Msg1} `{Serializable State1} `{Serializable
   Error1}
3     `{Serializable Setup2} `{Serializable Msg2} `{Serializable State2} `{Serializable
   Error2}
4     {C1 : Contract Setup1 Msg1 State1 Error1}
5     {C2 : Contract Setup2 Msg2 State2 Error2}.
6
7 Definition lift_genesis (f : ContractMorphism C1 C2) :
8   forall init_cstate,
9     is_genesis_cstate C1 init_cstate ->
10    is_genesis_cstate C2 (state_morph C1 C2 f init_cstate).
11 Proof.
12   destruct f as [setup_morph msg_morph state_morph error_morph i_coh r_coh].
13   cbn.
14   intros * genesis.
15   unfold is_genesis_cstate in *.
16   destruct genesis as [c [ctx [s init_coh]]].
17   exists c, ctx, (setup_morph s).
18   rewrite <- i_coh.
19   unfold result_functor.
20   now destruct (init C1 c ctx s).
21 Defined.
22
23 Definition lift_cstep_morph (f : ContractMorphism C1 C2) :
24   forall state1 state2,
25     ContractStep C1 state1 state2 ->
26     ContractStep C2
27       (state_morph C1 C2 f state1)
28       (state_morph C1 C2 f state2).
29 Proof.
30   destruct f as [setup_morph msg_morph state_morph error_morph i_coh r_coh].
31   cbn.
32   intros * step.
33   destruct step as [seq_chain seq_ctx seq_msg seq_new_acts recv_step].
34   apply (build_contract_step C2 (state_morph state1) (state_morph state2) seq_chain
   seq_ctx
35     (option_map msg_morph seq_msg) seq_new_acts).
36   rewrite <- r_coh.
37   unfold result_functor.
38   destruct (receive C1 seq_chain seq_ctx state1 seq_msg);
39   try destruct t;
40   now inversion recv_step.
41 Defined.
42
43 (** Lifting Theorem *)
```

```

44 Definition cm_lift_ctm (f : ContractMorphism C1 C2) : ContractTraceMorphism C1 C2 :=
45   build_ct_morph _ _ (state_morph _ _ f) (lift_genesis f) (lift_cstep_morph f).
46
47 End LiftingTheorem.

```

Listing C.7: Some results on the lifting theorem, including that the identity lifts to the identity, and compositions to compositions.

```

1 Section LiftingTheoremResults.
2 Context `{Serializable Setup1} `{Serializable Msg1} `{Serializable State1} `{Serializable
   Error1}
3   `{Serializable Setup2} `{Serializable Msg2} `{Serializable State2} `{Serializable
   Error2}
4   `{Serializable Setup3} `{Serializable Msg3} `{Serializable State3} `{Serializable
   Error3}
5   {C1 : Contract Setup1 Msg1 State1 Error1}
6   {C2 : Contract Setup2 Msg2 State2 Error2}
7   {C3 : Contract Setup3 Msg3 State3 Error3}.
8
9 (* id lifts to id *)
10 Theorem cm_lift_ctm_id :
11   cm_lift_ctm (id_cm C1) = id_ctm C1.
12 Proof.
13   unfold cm_lift_ctm, id_ctm.
14   simpl.
15   apply (eq_ctm_dep C1 C1 (@id State1)).
16   apply functional_extensionality_dep.
17   intro st1.
18   apply functional_extensionality_dep.
19   intro st1'.
20   apply functional_extensionality_dep.
21   intro cstep.
22   destruct cstep as [s_chn s_ctx s_msg s_nacts s_recv].
23   unfold id_cstep_morph.
24   cbn.
25   unfold option_map.
26   destruct s_msg;
27   cbn;
28   f_equal;
29   apply proof_irrelevance.
30 Qed.
31
32 (* compositions lift to compositions *)
33 Theorem cm_lift_ctm_compose
34   (g : ContractMorphism C2 C3) (f : ContractMorphism C1 C2) :
35   cm_lift_ctm (compose_cm g f) =
36   compose_ctm (cm_lift_ctm g) (cm_lift_ctm f).
37 Proof.
38   unfold cm_lift_ctm, compose_ctm.

```

```

39   cbn.
40   apply (eq_ctm_dep C1 C3 (compose (state_morph C2 C3 g) (state_morph C1 C2 f))).
41   apply functional_extensionality_dep.
42   intro st1.
43   apply functional_extensionality_dep.
44   intro st1'.
45   apply functional_extensionality_dep.
46   intro cstep.
47   destruct cstep as [s_chn s_ctx s_msg s_nacts s_recv].
48   unfold lift_cstep_morph.
49   destruct g as [smorph_g msgmorph_g stmorph_g errmorph_g initcoh_g recvcoh_g].
50   destruct f as [smorph_f msgmorph_f stmorph_f errmorph_f initcoh_f recvcoh_f].
51   cbn.
52   destruct s_msg;
53   cbn;
54   f_equal;
55   apply proof_irrelevance.
56 Qed.
57
58 End LiftingTheoremResults.

```

C.1.3 Contract Bisimulations

Listing C.8: The definition and results about contract bisimulations.

```

1 Section ContractBisimulation.
2
3 Section ContractTraceIsomorphism.
4 Context '{Serializable Setup1} '{Serializable Msg1} '{Serializable State1} '{Serializable
   Error1}
5       '{Serializable Setup2} '{Serializable Msg2} '{Serializable State2} '{Serializable
   Error2}
6       {C1 : Contract Setup1 Msg1 State1 Error1}
7       {C2 : Contract Setup2 Msg2 State2 Error2}.
8
9 (* a bisimulation of contracts, or an isomorphism of contract traces *)
10 Definition is_iso_ctm
11   (m1 : ContractTraceMorphism C1 C2) (m2 : ContractTraceMorphism C2 C1) :=
12   compose_ctm m2 m1 = id_ctm C1 /\
13   compose_ctm m1 m2 = id_ctm C2.
14
15 (* contract isomorphism -> contract trace isomorphism *)
16 Corollary ciso_to_ctiso (f : ContractMorphism C1 C2) (g : ContractMorphism C2 C1) :
17   is_iso_cm f g -> is_iso_ctm (cm_lift_ctm f) (cm_lift_ctm g).
18 Proof.
19   unfold is_iso_cm, is_iso_ctm.
20   intro iso_cm.
21   destruct iso_cm as [iso_cm_l iso_cm_r].

```

```

22   rewrite <- (cm_lift_ctm_compose g f).
23   rewrite <- (cm_lift_ctm_compose f g).
24   rewrite iso_cm_l.
25   rewrite iso_cm_r.
26   now repeat rewrite cm_lift_ctm_id.
27   Qed.
28
29 End ContractTraceIsomorphism.
30
31 (* the definition of bisimilar contracts *)
32 Definition contracts_bisimilar
33   `{Serializable Setup1} `{Serializable Msg1} `{Serializable State1} `{Serializable
   Error1}
34   `{Serializable Setup2} `{Serializable Msg2} `{Serializable State2} `{Serializable
   Error2}
35   (C1 : Contract Setup1 Msg1 State1 Error1)
36   (C2 : Contract Setup2 Msg2 State2 Error2) :=
37   exists (f : ContractTraceMorphism C1 C2) (g : ContractTraceMorphism C2 C1),
38   is_iso_ctm f g.
39
40 (* bisimilarity is an equivalence relation *)
41 Lemma bisim_refl
42   `{Serializable Setup} `{Serializable Msg} `{Serializable State} `{Serializable Error}
43   (C : Contract Setup Msg State Error) :
44   contracts_bisimilar C C.
45 Proof.
46   exists (id_ctm C), (id_ctm C).
47   unfold is_iso_ctm.
48   split; apply compose_id_ctm_left.
49   Qed.
50
51 Lemma bisim_sym
52   `{Serializable Setup1} `{Serializable Msg1} `{Serializable State1} `{Serializable
   Error1}
53   `{Serializable Setup2} `{Serializable Msg2} `{Serializable State2} `{Serializable
   Error2}
54   (C1 : Contract Setup1 Msg1 State1 Error1)
55   (C2 : Contract Setup2 Msg2 State2 Error2) :
56   contracts_bisimilar C1 C2 ->
57   contracts_bisimilar C2 C1.
58 Proof.
59   intro c_bisim.
60   unfold contracts_bisimilar in *.
61   destruct c_bisim as [f [f' iso_f_g]].
62   exists f', f.
63   unfold is_iso_ctm in *.
64   destruct iso_f_g as [f_id1 f_id2].
65   split.
66   - apply f_id2.

```



```

67   -   apply f_id1.
68 Qed.
69
70 Lemma bisim_trans
71   `{Serializable Setup1} `{Serializable Msg1} `{Serializable State1} `{Serializable
    Error1}
72   `{Serializable Setup2} `{Serializable Msg2} `{Serializable State2} `{Serializable
    Error2}
73   `{Serializable Setup3} `{Serializable Msg3} `{Serializable State3} `{Serializable
    Error3}
74   {C1 : Contract Setup1 Msg1 State1 Error1}
75   {C2 : Contract Setup2 Msg2 State2 Error2}
76   {C3 : Contract Setup3 Msg3 State3 Error3} :
77   contracts_bisimilar C1 C2 /\ contracts_bisimilar C2 C3 ->
78   contracts_bisimilar C1 C3.
79 Proof.
80   intros c_bisims.
81   destruct c_bisims as [[f [f' iso_f]] [g [g' iso_g]]].
82   unfold contracts_bisimilar in *.
83   exists (compose_ctm g f), (compose_ctm f' g').
84   destruct iso_g as [g_id1 g_id2].
85   destruct iso_f as [f_id1 f_id2].
86   unfold is_iso_ctm.
87   split.
88   -   rewrite <- compose_ctm_assoc.
89       replace (compose_ctm g' (compose_ctm g f)) with (compose_ctm (compose_ctm g' g) f)
90       .
91       2:{ now rewrite <- compose_ctm_assoc. }
92       rewrite g_id1.
93       rewrite compose_id_ctm_left.
94       apply f_id1.
95   -   rewrite <- compose_ctm_assoc.
96       replace (compose_ctm f (compose_ctm f' g')) with (compose_ctm (compose_ctm f f') g
97       ' ).
98       2:{ now rewrite <- compose_ctm_assoc. }
99       rewrite f_id2.
100      rewrite compose_id_ctm_left.
101      apply g_id2.
102 Qed.
103
104 (* an isomorphism of contracts lifts to a bisimulation *)
105 Theorem c_iso_to_bisim
106   `{Serializable Setup1} `{Serializable Msg1} `{Serializable State1} `{Serializable
    Error1}
107   `{Serializable Setup2} `{Serializable Msg2} `{Serializable State2} `{Serializable
    Error2}
108   {C1 : Contract Setup1 Msg1 State1 Error1}
109   {C2 : Contract Setup2 Msg2 State2 Error2} :
110   contracts_isomorphic C1 C2 -> contracts_bisimilar C1 C2.

```

```

109 Proof.
110   intro c_iso.
111   destruct c_iso as [f [g [is_iso_1 is_iso_2]]].
112   unfold contracts_bisimilar.
113   exists (cm_lift_ctm f), (cm_lift_ctm g).
114   unfold is_iso_ctm.
115   split;
116   rewrite <- cm_lift_ctm_compose;
117   try rewrite is_iso_1;
118   try rewrite is_iso_2;
119   now rewrite cm_lift_ctm_id.
120 Qed.
121
122 End ContractBisimulation.

```

C.1.4 Discussion: Propositional Indistinguishability

Listing C.9: An example illustrating the degree to which bisimilar contracts are propositionally indistinguishable.

```

1 (* Bisimilar contracts have very similar invariants, depending on the state
2    isomorphism.
3
4    In this example, we show that, for bisimilar contracts C1 and C2, if there
5    is a state invariant of C1 which is preserved by the state isomorphism of
6    the bisimulation, then that same state invariant holds for C2.
7 *)
8
9 Section PropositionalIndistinguishability.
10 Context {Base : ChainBase}.
11
12 Context
13   '{Serializable Setup1} '{Serializable Msg1} '{Serializable State1} '{Serializable
   Error1}
14   '{Serializable Setup2} '{Serializable Msg2} '{Serializable State2} '{Serializable
   Error2}
15 (* Consider contracts C1 and C2 ... *)
16   {C1 : Contract Setup1 Msg1 State1 Error1}
17   {C2 : Contract Setup2 Msg2 State2 Error2}.
18
19 (* such that C1 and C2 are bisimilar. *)
20 Context {m1 : ContractTraceMorphism C1 C2} {m2 : ContractTraceMorphism C2 C1}
21         {C1_C2_bisim : is_iso_ctm m1 m2}.
22
23 (* Assume that both State1 and State2 have a constant in storage, given by a function
24    const_in_stor *)
25 Context { const_in_stor_C1 : State1 -> nat } { const_in_stor_C2 : State2 -> nat }

```

```

26 (* and assume that this constant is invariant under the ct_state isomorphism *)
27     { const_pres : forall st,
28         const_in_stor_C2 (ct_state_morph C1 C2 m1 st) = const_in_stor_C1 st }.
29
30 (* Assume that an invariant holds for C1 ... *)
31 Axiom invariant_C1_init : forall c ctx s cstate,
32     init C1 c ctx s = Ok cstate ->
33     (const_in_stor_C1 cstate > 0)%nat.
34
35 Axiom invariant_C1_recv : forall c ctx cstate op_msg new_st nacts,
36     receive C1 c ctx cstate op_msg = Ok (new_st, nacts) ->
37     (const_in_stor_C1 new_st > 0)%nat.
38
39 (* Because the invariant is preserved under the state isomorphism of the bisimulation,
40     we can prove that the same invariant of C2 *)
41 Theorem invariant_C2 bstate caddr (trace : ChainTrace empty_state bstate):
42     (* Forall reachable states with contract at caddr, *)
43     env_contracts bstate caddr = Some (C2 : WeakContract) ->
44     (* such that cstate is the state of the contract, *)
45     exists (cstate : State2),
46     contract_state bstate caddr = Some cstate /\
47     (* the constant in storage is > 0 *)
48     (const_in_stor_C2 cstate > 0)%nat.
49 Proof.
50     intros.
51     contract_induction; auto; intros.
52     (* deployment *)
53     - assert (is_genesis_cstate C2 result)
54       as is_gen2.
55       { unfold is_genesis_cstate.
56         now exists chain, ctx, setup. }
57     pose proof (genesis_fixpoint C2 C1 m2 result is_gen2)
58     as is_gen1.
59     destruct is_gen1 as [c2 [ctx2 [s2 init_ok2]]].
60     pose proof (invariant_C1_init _ _ _ _ init_ok2)
61     as invar_C1.
62     rewrite <- const_pres in invar_C1.
63     assert ((ct_state_morph C1 C2 m1 (ct_state_morph C2 C1 m2 result)) = result)
64     as state_id.
65     { unfold is_iso_ctm in C1_C2_bisim.
66       destruct C1_C2_bisim as [iso_l iso_r].
67       unfold compose_ctm, id_ctm in *.
68       replace (ct_state_morph C1 C2 m1 (ct_state_morph C2 C1 m2 result))
69       with (Basics.compose (ct_state_morph C1 C2 m1) (ct_state_morph C2 C1 m2)
70 result).
71       2:{ auto. }
72       inversion iso_r.
73       now rewrite H8. }
74     now rewrite state_id in invar_C1.

```

```

74   (* nonrecursive call *)
75   - assert (ContractStep C2 prev_state new_state)
76   as step_C2.
77   {   apply (build_contract_step C2 prev_state new_state chain ctx msg
78         new_acts receive_some). }
79   pose proof (cstep_morph C2 C1 m2 prev_state new_state step_C2)
80   as step_morph.
81   destruct step_morph.
82   pose proof (invariant_C1_rcv seq_chain seq_ctx
83         (ct_state_morph C2 C1 m2 prev_state) seq_msg
84         (ct_state_morph C2 C1 m2 new_state) seq_new_acts rcv_ok_step)
85   as invar_C2.
86   rewrite <- const_pres in invar_C2.
87   replace (ct_state_morph C1 C2 m1 (ct_state_morph C2 C1 m2 new_state))
88   with (Basics.compose (ct_state_morph C1 C2 m1) (ct_state_morph C2 C1 m2) new_state
89   )
89   in invar_C2.
90   2:{ auto. }
91   assert (Basics.compose (ct_state_morph C1 C2 m1) (ct_state_morph C2 C1 m2) = id)
92   as state_id.
93   {   unfold is_iso_ctm in C1_C2_bisim.
94         destruct C1_C2_bisim as [iso_l iso_r].
95         unfold compose_ctm, id_ctm in *.
96         now inversion iso_r. }
97   now rewrite state_id in invar_C2.
98   (* recursive call *)
99   - assert (ContractStep C2 prev_state new_state)
100  as step_C2.
101  {   apply (build_contract_step C2 prev_state new_state chain ctx msg
102        new_acts receive_some). }
103  pose proof (cstep_morph C2 C1 m2 prev_state new_state step_C2)
104  as step_morph.
105  destruct step_morph.
106  pose proof (invariant_C1_rcv seq_chain seq_ctx
107        (ct_state_morph C2 C1 m2 prev_state) seq_msg
108        (ct_state_morph C2 C1 m2 new_state) seq_new_acts rcv_ok_step)
109  as invar_C2.
110  rewrite <- const_pres in invar_C2.
111  replace (ct_state_morph C1 C2 m1 (ct_state_morph C2 C1 m2 new_state))
112  with (Basics.compose (ct_state_morph C1 C2 m1) (ct_state_morph C2 C1 m2) new_state
113  )
113  in invar_C2.
114  2:{ auto. }
115  assert (Basics.compose (ct_state_morph C1 C2 m1) (ct_state_morph C2 C1 m2) = id)
116  as state_id.
117  {   unfold is_iso_ctm in C1_C2_bisim.
118        destruct C1_C2_bisim as [iso_l iso_r].
119        unfold compose_ctm, id_ctm in *.
120        now inversion iso_r. }

```

```

121     now rewrite state_id in invar_C2.
122     (* prove facts *)
123     - solve_facts.
124 Qed.
125
126 End PropositionalIndistinguishability.

```

C.2 Contract Systems as Bigraphs

C.2.1 Bigraphs

C.2.2 The Place Graph

Listing C.10: The formalization of an n-ary tree, or n-tree.

```

1 Section ntree.
2
3 Inductive ntree (T : Type) : Type :=
4 | Node : T -> list (ntree T) -> ntree T.
5
6 Definition singleton_ntree {T} (t : T) := Node T t nil.
7
8 (* fold/traversal for ntrees *)
9 Fixpoint ntree_fold_left {A T}
10   (f : A -> T -> A)
11   (sys : ntree T)
12   (a0 : A) : A :=
13   match sys with
14   | Node _ t list_child_trees =>
15     List.fold_left
16       (fun (a0' : A) (sys' : ntree T) =>
17         ntree_fold_left f sys' a0')
18       list_child_trees
19       (f a0 t)
20   end.
21
22 (* ntree map : the functoriality of ntrees *)
23 Fixpoint ntree_map {T T'} (f : T -> T') (tree : ntree T) : ntree T' :=
24   match tree with
25   | Node _ v children =>
26     Node T' (f v) (List.map (fun child => ntree_map f child) children)
27   end.
28
29 Fixpoint replace_at_index {T : Type} (n : nat) (new_elem : T) (l : list T) : list T :=
30   match l, n with
31   | nil, _ => nil

```

```

32 | _ :: tl, 0 => new_elem :: tl
33 | hd :: tl, S n' => hd :: replace_at_index n' new_elem tl
34 end.
35
36 Fixpoint add_tree_at_leaf {T} (orig append : ntree T) (leaf_index : list nat) : ntree T :=
37   match orig, leaf_index with
38   | Node _ v children, nil => Node T v (append :: children)
39   | Node _ v children, i :: is =>
40     match List.nth_error children i with
41     | Some child => Node T v (replace_at_index i (add_tree_at_leaf child append is)
42                               children)
43     | None => orig
44   end
45 end.
46 End ntree.

```

Listing C.11: The formal definition of a contract's place graph.

```

1 Definition ContractPlaceGraph
2   (Setup Msg State Error : Type)
3   `{sys_set : Serializable Setup}
4   `{sys_msg : Serializable Msg}
5   `{sys_st  : Serializable State}
6   `{sys_err : Serializable Error} :=
7   ntree (Contract Setup Msg State Error).

```

Listing C.12: The formal definition of a contract place graph's interface.

```

1 Section SystemInterface.
2 Context `{Serializable Setup} `{Serializable Msg} `{Serializable State} `{Serializable
   Error}.
3
4 (* system init : just initialize the root, since all contract init behaves identically *)
5 Definition sys_init
6   (sys : ContractPlaceGraph Setup Msg State Error)
7   (c : Chain)
8   (ctx : ContractCallContext)
9   (s : Setup) : result State Error :=
10  match sys with
11  | Node _ root_contract _ =>
12    init root_contract c ctx s
13  end.
14
15 (* system receive: take the message and state and run through the entire system.
16    since systems are iteratively built so that a message not intended for a given
17    contract
18    returns the identity, this targets the contract in question and leaves the rest
19    untouched. *)

```

```

18 Definition sys_receive
19   (sys : ContractPlaceGraph Setup Msg State Error)
20   (c : Chain)
21   (ctx : ContractCallContext)
22   (st : State)
23   (op_msg : option Msg) : result (State * list ActionBody) Error :=
24   ntree_fold_left
25   (fun (recv_propagate : result (State * list ActionBody) Error)
26     (contr : Contract Setup Msg State Error) =>
27     match recv_propagate with
28     | Ok (st0, lacts0) =>
29       match receive contr c ctx st0 op_msg with
30       | Ok (st1, lacts1) => Ok (st1, lacts0 ++ lacts1)
31       | Err e => Err e
32     end
33     | Err e => Err e
34   end)
35   sys
36   (Ok (st, nil)).
37
38 (* thes two functions give us a contract *)
39 Definition sys_contract (sys : ContractPlaceGraph Setup Msg State Error) :=
40   build_contract (sys_init sys) (sys_receive sys).
41
42 End SystemInterface.

```

C.2.2.1 Iteratively Building a Contract System

Listing C.13: Some definitions and functions for iteratively constructing the place graph of a contract system.

```

1 Section IterativePlaceGraphBuild.
2 (* the definition of a singleton system *)
3 Definition singleton_place_graph
4   `{Serializable Setup} `{Serializable Msg} `{Serializable State} `{Serializable Error}
5   (C : Contract Setup Msg State Error)
6   : ContractPlaceGraph Setup Msg State Error := singleton_ntree C.
7
8 Section IterativeSum.
9 Context `{Serializable Setup1} `{Serializable Msg1} `{Serializable State1} `{Serializable
   Error1}
10   `{Serializable Setup2} `{Serializable Msg2} `{Serializable State2} `{Serializable
   Error2}.
11
12 (* an iterative add to contract systems s.t. type goals are satisfied *)
13
14 (* accepts messages on the left *)
15 Definition c_sum_l

```

```

16   (C1 : Contract Setup1 Msg1 State1 Error1)
17   (C2 : Contract Setup2 Msg2 State2 Error2) :
18   Contract (Setup1 * Setup2) (Msg1 + Msg2) (State1 * State2) (Error1 + Error2).
19 Proof.
20   destruct C1 as [init1 recv1].
21   destruct C2 as [init2 recv2].
22   apply build_contract.
23   (* each setup must succeed, providing the new system state *)
24   - apply (fun c ctx s' =>
25       let '(s1, s2) := s' in
26       match init1 c ctx s1, init2 c ctx s2 with
27       | Ok st1, Ok st2 => Ok (st1, st2)
28       | Err e, _ => Err (inl e) (* the left error is first *)
29       | _, Err e => Err (inr e) (* followed by the right *)
30       end).
31   - apply (fun c ctx st' op_msg =>
32       let '(st1, st2) := st' in
33       match op_msg with
34       | Some msg =>
35         match msg with
36         (* the message was intended for this contract,
37          so we attempt to update the state *)
38         | inl msg =>
39           match recv1 c ctx st1 (Some msg) with
40           | Ok (new_st1, nacts) => Ok ((new_st1, st2), nacts)
41           | Err e => Err (inl e)
42           end
43         (* the message was not intended for this contract, so we do nothing *)
44         | inr msg => Ok (st', nil)
45         end
46       | None => (* if there is no message, we call the contract with None *)
47         match recv1 c ctx st1 None with
48         | Ok (new_st1, nacts) => Ok ((new_st1, st2), nacts)
49         | Err e => Err (inl e)
50         end
51       end).
52 Defined.
53
54 (* same as before, but accepts messages on the right now *)
55 Definition c_sum_r
56   (C1 : Contract Setup1 Msg1 State1 Error1)
57   (C2 : Contract Setup2 Msg2 State2 Error2) :
58   Contract (Setup1 * Setup2) (Msg1 + Msg2) (State1 * State2) (Error1 + Error2).
59 Proof.
60   destruct C1 as [init1 recv1].
61   destruct C2 as [init2 recv2].
62   apply build_contract.
63   (* each setup must succeed, providing the new system state *)
64   - apply (fun c ctx s' =>

```



```

65         let '(s1, s2) := s' in
66         match init1 c ctx s1, init2 c ctx s2 with
67         | Ok st1, Ok st2 => Ok (st1, st2)
68         | Err e, _ => Err (inl e) (* the left error is first *)
69         | _, Err e => Err (inr e) (* followed by the right *)
70         end).
71 -   apply (fun c ctx st' op_msg =>
72         let '(st1, st2) := st' in
73         match op_msg with
74         | Some msg =>
75             match msg with
76             (* the message was not intended for this contract, so we do nothing *)
77             | inl msg => Ok (st', nil)
78             (* the message was intended for this contract,
79              so we attempt to update the state *)
80             | inr msg =>
81                 match recv2 c ctx st2 (Some msg) with
82                 | Ok (new_st2, nacts) => Ok ((st1, new_st2), nacts)
83                 | Err e => Err (inr e)
84                 end
85             end
86         | None => (* if there is no message, we call the contract with None *)
87             match recv2 c ctx st2 None with
88             | Ok (new_st2, nacts) => Ok ((st1, new_st2), nacts)
89             | Err e => Err (inr e)
90             end
91         end).
92 Defined.
93
94 End IterativeSum.
95
96
97 Section IterativeChild.
98 Context '{Serializable Setup1} '{Serializable Msg1} '{Serializable State1} '{Serializable
   Error1}
99         '{Serializable Setup2} '{Serializable Msg2} '{Serializable State2} '{Serializable
   Error2}.
100
101 (* add a contract as a child to a system(/nest contracts) *)
102 Definition sys_add_child_r
103   (sys : ContractPlaceGraph Setup1 Msg1 State1 Error1)
104   (C : Contract Setup2 Msg2 State2 Error2) :
105   ContractPlaceGraph (Setup1 * Setup2) (Msg1 + Msg2) (State1 * State2) (Error1 + Error2)
   :=
106   let T := (Contract (Setup1 * Setup2) (Msg1 + Msg2) (State1 * State2) (Error1 + Error2)
   ) in
107   match sys with
108   | Node _ root_contract _ =>
109       match (ntree_map (fun C1 => c_sum_l C1 C) sys) with

```

```

110         | Node _ root_contract' children =>
111             Node T root_contract' (children ++ [Node T (c_sum_r root_contract C) nil])
112     end
113 end.
114
115 (* nest C1 C2 indicates that C2 is nested within C1 *)
116 Definition nest
117   (C1 : Contract Setup1 Msg1 State1 Error1)
118   (C2 : Contract Setup2 Msg2 State2 Error2) :
119   ContractPlaceGraph (Setup1 * Setup2) (Msg1 + Msg2) (State1 * State2) (Error1 + Error2)
120   :=
121   let T := (Contract (Setup1 * Setup2) (Msg1 + Msg2) (State1 * State2) (Error1 + Error2)
122     ) in
123   Node T (c_sum_l C1 C2) [Node T (c_sum_r C1 C2) nil].
124
125 End IterativeChild.
126
127 End IterativePlaceGraphBuild.

```

C.2.2.2 System Contracts, Morphisms, and Isomorphisms

Listing C.14: The formal definition of a system morphism, or a morphism of system place graphs.

```

1 Record SystemMorphism
2   (sys1 : ContractPlaceGraph Setup1 Msg1 State1 Error1)
3   (sys2 : ContractPlaceGraph Setup2 Msg2 State2 Error2) :=
4   build_system_morphism {
5     (* the components of a morphism *)
6     sys_setup_morph : Setup1 -> Setup2 ;
7     sys_msg_morph   : Msg1   -> Msg2   ;
8     sys_state_morph : State1 -> State2 ;
9     sys_error_morph : Error1 -> Error2 ;
10    (* coherence conditions *)
11    sys_init_coherence : forall c ctx s,
12      result_functor sys_state_morph sys_error_morph
13        (sys_init sys1 c ctx s) =
14        sys_init sys2 c ctx (sys_setup_morph s) ;
15    sys_recv_coherence : forall c ctx st op_msg,
16      result_functor (fun ' (st, l) => (sys_state_morph st, l)) sys_error_morph
17        (sys_receive sys1 c ctx st op_msg) =
18        sys_receive sys2 c ctx (sys_state_morph st) (option_map sys_msg_morph op_msg)
19    ;
20  }.

```

Listing C.15: A system place graph can be defined as a contract, and system morphisms are in one-to-one correspondence with contract morphisms of that contract.

```

1 Definition sys_contract (sys : ContractPlaceGraph Setup Msg State Error) :=

```

```

2   build_contract (sys_init sys) (sys_receive sys).
3
4   (* a system morphism is in one-to-one correspondence with a morphism of contracts,
5      when we consider a system as its own contract *)
6   Definition cm_to_sysm
7     (sys1 : ContractPlaceGraph Setup1 Msg1 State1 Error1)
8     (sys2 : ContractPlaceGraph Setup2 Msg2 State2 Error2)
9     (f : ContractMorphism (sys_contract sys1) (sys_contract sys2)) : SystemMorphism sys1
10    sys2.
11   Proof.
12     destruct f.
13     apply (build_system_morphism sys1 sys2 setup_morph msg_morph state_morph error_morph
14           init_coherence recv_coherence).
15   Defined.
16
17   Definition sysm_to_cm
18     (sys1 : ContractPlaceGraph Setup1 Msg1 State1 Error1)
19     (sys2 : ContractPlaceGraph Setup2 Msg2 State2 Error2)
20     (f : SystemMorphism sys1 sys2) : ContractMorphism (sys_contract sys1) (sys_contract
21     sys2).
22   Proof.
23     destruct f as [sys_setup_morph sys_msg_morph sys_state_morph sys_error_morph
24     sys_init_coh sys_recv_coh].
25     apply (build_contract_morphism (sys_contract sys1) (sys_contract sys2)
26           sys_setup_morph sys_msg_morph sys_state_morph sys_error_morph
27           sys_init_coh sys_recv_coh).
28   Defined.
29
30   Lemma cm_sysm_one_to_one
31     (sys1 : ContractPlaceGraph Setup1 Msg1 State1 Error1)
32     (sys2 : ContractPlaceGraph Setup2 Msg2 State2 Error2) :
33     compose (cm_to_sysm sys1 sys2) (sysm_to_cm sys1 sys2) = id /\
34     compose (sysm_to_cm sys1 sys2) (cm_to_sysm sys1 sys2) = id.
35   Proof.
36     split;
37     unfold sysm_to_cm, cm_to_sysm;
38     apply functional_extensionality;
39     intro;
40     now destruct x.
41   Qed.

```

C.3 The Link Graph

C.3.1 System Steps and System Trace

Listing C.16: The formal definition of a single system step, and chained single steps.

```

1 Section LinkGraph.
2 Context `{Serializable Setup} `{Serializable Msg} `{Serializable State} `{Serializable
   Error}.
3
4 (* system state stepping forward *)
5 Record SingleSystemStep (sys : ContractPlaceGraph Setup Msg State Error)
6   (prev_sys_state next_sys_state : State) :=
7   build_sys_single_step {
8     sys_step_chain : Chain ;
9     sys_step_ctx : ContractCallContext ;
10    sys_step_msg : option Msg ;
11    sys_step_nacts : list ActionBody ;
12    (* we can call receive successfully *)
13    sys_rcv_ok_step :
14      sys_receive sys sys_step_chain sys_step_ctx prev_sys_state sys_step_msg =
15      Ok (next_sys_state, sys_step_nacts) ;
16  }.
17
18 Definition ChainedSingleSteps (sys : ContractPlaceGraph Setup Msg State Error) :=
19   ChainedList State (SingleSystemStep sys).
20
21 End LinkGraph.

```

Listing C.17: The formal definition of a contract system consists of a definition of its place and link graphs. The link graph must have semantics in chained single steps.

```

1 Record ContractSystem
2   (Setup Msg State Error : Type)
3   `{sys_set : Serializable Setup}
4   `{sys_msg : Serializable Msg}
5   `{sys_st : Serializable State}
6   `{sys_err : Serializable Error} :=
7   build_contract_system {
8     (* the place and link graphs *)
9     sys_place : ContractPlaceGraph Setup Msg State Error ;
10    sys_link : State -> State -> Type ;
11    (* the link graph has semantics in ChanedSingleSteps *)
12    sys_link_semantics : forall st1 st2,
13      sys_link st1 st2 ->
14      ChainedSingleSteps sys_place st1 st2 ;
15  }.

```

Listing C.18: The formal definition of the steps of a contract system.

```

1 Definition SystemStep (sys : ContractSystem Setup Msg State Error) :=
2   sys_link' sys.

```

Listing C.19: The formal definition of the a contract system's trace.

```
1 Definition SystemTrace (sys : ContractSystem Setup Msg State Error) :=
2   ChainedList State (SystemStep sys).
```

C.4 Bisimulations of Contract Systems

C.4.1 System Trace Morphisms and System Bisimulations

Listing C.20: The formal definition of a system trace morphism.

```
1 Record SystemTraceMorphism
2   (sys1 : ContractSystem Setup1 Msg1 State1 Error1)
3   (sys2 : ContractSystem Setup2 Msg2 State2 Error2) :=
4   build_st_morph {
5     (* a function *)
6     st_state_morph : State1 -> State2 ;
7     (* init state sys1 -> init state sys2 *)
8     sys_genesis_fixpoint : forall init_sys_state,
9       is_genesis_sys_state sys1 init_sys_state ->
10      is_genesis_sys_state sys2 (st_state_morph init_sys_state) ;
11     (* step morphism *)
12     sys_step_morph : forall sys_state1 sys_state2,
13       SystemStep sys1 sys_state1 sys_state2 ->
14       SystemStep sys2 (st_state_morph sys_state1) (st_state_morph sys_state2) ;
15   }.
```

Listing C.21: The identity system trace morphism.

```
1 Section IdentitySTMorphism.
2 Context `{Serializable Setup} `{Serializable Msg} `{Serializable State} `{Serializable
   Error}.
3
4 Definition id_sys_genesis_fixpoint (sys : ContractSystem Setup Msg State Error)
5   init_sys_state
6   (gen_sys : is_genesis_sys_state sys init_sys_state) :
7   is_genesis_sys_state sys (id init_sys_state) :=
8   gen_sys.
9
10 Definition id_sys_step_morph (sys : ContractSystem Setup Msg State Error)
11   sys_state1 sys_state2 (step : SystemStep sys sys_state1 sys_state2) :
12   SystemStep sys (id sys_state1) (id sys_state2) :=
13   step.
14
15 Definition id_stm (sys : ContractSystem Setup Msg State Error) : SystemTraceMorphism sys
   sys :=
16 { |
```

```

17   st_state_morph := id ;
18   sys_genesis_fixpoint := id_sys_genesis_fixpoint sys ;
19   sys_step_morph := id_sys_step_morph sys ;
20 |}.
21
22 End IdentitySTMorphism.

```

Listing C.22: A lemma to assert equality of system trace morphisms.

```

1 Section EqualityOfSTMorphisms.
2 Context '{Serializable Setup1} '{Serializable Msg1} '{Serializable State1} '{Serializable
   Error1}
3       '{Serializable Setup2} '{Serializable Msg2} '{Serializable State2} '{Serializable
   Error2}.
4
5 Lemma eq_stm_dep
6   (sys1 : ContractSystem Setup1 Msg1 State1 Error1)
7   (sys2 : ContractSystem Setup2 Msg2 State2 Error2)
8   (st_st_m : State1 -> State2)
9   sys_gen_fix1 sys_gen_fix2
10  (sys_step_m1 sys_step_m2 : forall sys_state1 sys_state2,
11    SystemStep sys1 sys_state1 sys_state2 ->
12    SystemStep sys2 (st_st_m sys_state1) (st_st_m sys_state2)) :
13  sys_step_m1 = sys_step_m2 ->
14  {| st_state_morph := st_st_m ;
15    sys_genesis_fixpoint := sys_gen_fix1 ;
16    sys_step_morph := sys_step_m1 ; |}
17  =
18  {| st_state_morph := st_st_m ;
19    sys_genesis_fixpoint := sys_gen_fix2 ;
20    sys_step_morph := sys_step_m2 ; |}.
21 Proof.
22   intro cstep_equiv.
23   subst.
24   f_equal.
25   apply proof_irrelevance.
26 Qed.
27
28 End EqualityOfSTMorphisms.

```

Listing C.23: The formal definition of system trace morphism composition.

```

1 Section STMorphismComposition.
2 Context '{Serializable Setup1} '{Serializable Msg1} '{Serializable State1} '{Serializable
   Error1}
3       '{Serializable Setup2} '{Serializable Msg2} '{Serializable State2} '{Serializable
   Error2}
4       '{Serializable Setup3} '{Serializable Msg3} '{Serializable State3} '{Serializable
   Error3}

```

```

5      {sys1 : ContractSystem Setup1 Msg1 State1 Error1}
6      {sys2 : ContractSystem Setup2 Msg2 State2 Error2}
7      {sys3 : ContractSystem Setup3 Msg3 State3 Error3}.
8
9 Definition sys_genesis_compose
10   (m2 : SystemTraceMorphism sys2 sys3) (m1 : SystemTraceMorphism sys1 sys2)
11   init_sys_state (gen_s1 : is_genesis_sys_state sys1 init_sys_state) :
12   is_genesis_sys_state sys3
13     (compose (st_state_morph sys2 sys3 m2) (st_state_morph sys1 sys2 m1)
14     init_sys_state) :=
15   match m2, m1 with
16   | build_st_morph _ _ _ gen_fix2 step2, build_st_morph _ _ _ gen_fix1 step1 =>
17     gen_fix2 _ (gen_fix1 _ gen_s1)
18   end.
19
20 Definition sys_step_compose
21   (m2 : SystemTraceMorphism sys2 sys3) (m1 : SystemTraceMorphism sys1 sys2)
22   sys_state1 sys_state2
23   (step : SystemStep sys1 sys_state1 sys_state2) :
24   SystemStep sys3
25     (compose (st_state_morph sys2 sys3 m2) (st_state_morph sys1 sys2 m1) sys_state1)
26     (compose (st_state_morph sys2 sys3 m2) (st_state_morph sys1 sys2 m1) sys_state2)
27   :=
28   match m2, m1 with
29   | build_st_morph _ _ _ _ step2, build_st_morph _ _ _ _ step1 =>
30     step2 _ _ (step1 _ _ step)
31   end.
32
33 Definition compose_stm
34   (m2 : SystemTraceMorphism sys2 sys3)
35   (m1 : SystemTraceMorphism sys1 sys2) : SystemTraceMorphism sys1 sys3 :=
36   { |
37     st_state_morph := compose (st_state_morph _ _ m2) (st_state_morph _ _ m1) ;
38     sys_genesis_fixpoint := sys_genesis_compose m2 m1 ;
39     sys_step_morph := sys_step_compose m2 m1 ;
40   | }.
41
42 End STMorphismComposition.

```

Listing C.24: Some results about system trace morphism composition.

```

1 Section STMorphismComposition.
2 Context `{Serializable Setup1} `{Serializable Msg1} `{Serializable State1} `{Serializable
   Error1}
3   `{Serializable Setup2} `{Serializable Msg2} `{Serializable State2} `{Serializable
   Error2}
4   `{Serializable Setup3} `{Serializable Msg3} `{Serializable State3} `{Serializable
   Error3}
5   {sys1 : ContractSystem Setup1 Msg1 State1 Error1}

```

```

6      {sys2 : ContractSystem Setup2 Msg2 State2 Error2}
7      {sys3 : ContractSystem Setup3 Msg3 State3 Error3}.
8
9 Definition sys_genesis_compose
10   (m2 : SystemTraceMorphism sys2 sys3) (m1 : SystemTraceMorphism sys1 sys2)
11   init_sys_state (gen_s1 : is_genesis_sys_state sys1 init_sys_state) :
12   is_genesis_sys_state sys3
13     (compose (st_state_morph sys2 sys3 m2) (st_state_morph sys1 sys2 m1)
14      init_sys_state) :=
15   match m2, m1 with
16   | build_st_morph _ _ _ gen_fix2 step2, build_st_morph _ _ _ gen_fix1 step1 =>
17     gen_fix2 _ (gen_fix1 _ gen_s1)
18   end.
19
20 Definition sys_step_compose
21   (m2 : SystemTraceMorphism sys2 sys3) (m1 : SystemTraceMorphism sys1 sys2)
22   sys_state1 sys_state2
23   (step : SystemStep sys1 sys_state1 sys_state2) :
24   SystemStep sys3
25     (compose (st_state_morph sys2 sys3 m2) (st_state_morph sys1 sys2 m1) sys_state1)
26     (compose (st_state_morph sys2 sys3 m2) (st_state_morph sys1 sys2 m1) sys_state2)
27   :=
28   match m2, m1 with
29   | build_st_morph _ _ _ _ step2, build_st_morph _ _ _ _ step1 =>
30     step2 _ _ (step1 _ _ step)
31   end.
32
33 Definition compose_stm
34   (m2 : SystemTraceMorphism sys2 sys3)
35   (m1 : SystemTraceMorphism sys1 sys2) : SystemTraceMorphism sys1 sys3 :=
36   { |
37     st_state_morph := compose (st_state_morph _ _ m2) (st_state_morph _ _ m1) ;
38     sys_genesis_fixpoint := sys_genesis_compose m2 m1 ;
39     sys_step_morph := sys_step_compose m2 m1 ;
40   }|.
41
42 End STMorphismComposition.

```

Listing C.25: A system trace isomorphism.

```

1 Definition is_iso_stm
2   (m1 : SystemTraceMorphism sys1 sys2) (m2 : SystemTraceMorphism sys2 sys1) :=
3   compose_stm m2 m1 = id_stm sys1 /\
4   compose_stm m1 m2 = id_stm sys2.

```

Listing C.26: The formal definition of a bisimulation of contract systems.

```

1 Definition systems_bisimilar
2   `{Serializable Setup1} `{Serializable Msg1} `{Serializable State1} `{Serializable
3   Error1}

```



```

3   '{Serializable Setup2} '{Serializable Msg2} '{Serializable State2} '{Serializable
   Error2}
4   (sys1 : ContractSystem Setup1 Msg1 State1 Error1)
5   (sys2 : ContractSystem Setup2 Msg2 State2 Error2) :=
6   exists (f : SystemTraceMorphism sys1 sys2) (g : SystemTraceMorphism sys2 sys1),
7   is_iso_stm f g.

```

C.4.2 Lifting Theorems for Contract Systems

Listing C.27: The discrete link graph construction on a place graph.

```

1 Section DiscreteLinkSys.
2 Context '{Serializable Setup} '{Serializable Msg} '{Serializable State} '{Serializable
   Error}.
3
4 Definition discrete_link (sys : ContractPlaceGraph Setup Msg State Error) st1 st2 :=
5   SingleSystemStep sys st1 st2.
6
7 Definition discrete_link_semantics (sys : ContractPlaceGraph Setup Msg State Error)
8   st1 st2 (step : discrete_link sys st1 st2) :
9   ChainedSingleSteps sys st1 st2 :=
10   (snoc clnil step).
11
12 Definition discrete_sys (sys : ContractPlaceGraph Setup Msg State Error) := { |
13   sys_place := sys ;
14   sys_link := discrete_link sys ;
15   sys_link_semantics := discrete_link_semantics sys ;
16 | }.
17
18 End DiscreteLinkSys.

```

Listing C.28: The lifting theorem for system trace morphisms.

```

1 Definition sm_lift_stm (f : SystemMorphism sys1 sys2) :
2   SystemTraceMorphism (discrete_sys sys1) (discrete_sys sys2) :=
3   build_st_morph _ _ (sys_state_morph _ _ f) (lift_sys_genesis f) (lift_sys_step_morph f
   ).

```

Listing C.29: Some results on the lifting theorem for contract trace morphisms.

```

1 Section LiftingTheoremResults.
2 Context '{Serializable Setup1} '{Serializable Msg1} '{Serializable State1} '{Serializable
   Error1}
3   '{Serializable Setup2} '{Serializable Msg2} '{Serializable State2} '{Serializable
   Error2}
4   '{Serializable Setup3} '{Serializable Msg3} '{Serializable State3} '{Serializable
   Error3}

```

```

5      {sys1 : ContractPlaceGraph Setup1 Msg1 State1 Error1}
6      {sys2 : ContractPlaceGraph Setup2 Msg2 State2 Error2}
7      {sys3 : ContractPlaceGraph Setup3 Msg3 State3 Error3}.
8
9  (* id lifts to id *)
10 Lemma sm_lift_stm_id :
11    sm_lift_stm (id_sm sys1) = id_stm (discrete_sys sys1).
12 Proof.
13   apply (eq_stm_dep (discrete_sys sys1) (discrete_sys sys1) (@id State1)).
14   apply functional_extensionality_dep.
15   intro st1.
16   apply functional_extensionality_dep.
17   intro st1'.
18   apply functional_extensionality_dep.
19   intro sys_step.
20   destruct sys_step.
21   unfold lift_sys_step_morph, id_sm, discrete_sys, option_map, id_sys_step_morph.
22   cbn.
23   do 2 f_equal; auto.
24   destruct sys_step_msg;
25   apply f_equal;
26   apply proof_irrelevance.
27 Qed.
28
29 (* compositions lift to compositions *)
30 Lemma sm_lift_stm_compose
31   (g : SystemMorphism sys2 sys3) (f : SystemMorphism sys1 sys2) :
32   sm_lift_stm (compose_sm g f) =
33   compose_stm (sm_lift_stm g) (sm_lift_stm f).
34 Proof.
35   apply (eq_stm_dep (discrete_sys sys1) (discrete_sys sys3)
36     (compose (sys_state_morph sys2 sys3 g) (sys_state_morph sys1 sys2 f))).
37   apply functional_extensionality_dep.
38   intro st1.
39   apply functional_extensionality_dep.
40   intro st1'.
41   apply functional_extensionality_dep.
42   intro sys_step.
43   induction sys_step.
44   destruct g as [smorph_g msgmorph_g stmorph_g errmorph_g initcoh_g recvcoh_g].
45   destruct f as [smorph_f msgmorph_f stmorph_f errmorph_f initcoh_f recvcoh_f].
46   unfold lift_sys_step_morph, sys_step_compose, compose_sm.
47   destruct sys_step_msg;
48   cbn;
49   f_equal;
50   apply proof_irrelevance.
51 Qed.
52
53 End LiftingTheoremResults.

```

Listing C.30: Isomorphic contract systems are bisimilar under the discrete link graph.

```

1 Corollary sys_iso_to_bisim
2   '{Serializable Setup1} '{Serializable Msg1} '{Serializable State1} '{Serializable
   Error1}
3   '{Serializable Setup2} '{Serializable Msg2} '{Serializable State2} '{Serializable
   Error2}
4   (sys1 : ContractPlaceGraph Setup1 Msg1 State1 Error1)
5   (sys2 : ContractPlaceGraph Setup2 Msg2 State2 Error2) :
6   systems_isomorphic sys1 sys2 -> systems_bisimilar (discrete_sys sys1) (discrete_sys
   sys2).
7 Proof.
8   intro sys_iso.
9   destruct sys_iso as [f [g [is_iso_1 is_iso_2]]].
10  unfold systems_bisimilar.
11  exists (sm_lift_stm f), (sm_lift_stm g).
12  unfold is_iso_stm.
13  split;
14  rewrite <- sm_lift_stm_compose;
15  try rewrite is_iso_1;
16  try rewrite is_iso_2;
17  now rewrite sm_lift_stm_id.
18 Qed.

```

Listing C.31: Contract morphisms lift to system morphisms, and system morphisms lift to system trace morphisms.

```

1 Section LiftCMtoSM.
2 Context '{Serializable Setup1} '{Serializable Msg1} '{Serializable State1} '{Serializable
   Error1}
3   '{Serializable Setup2} '{Serializable Msg2} '{Serializable State2} '{Serializable
   Error2}
4   {C1 : Contract Setup1 Msg1 State1 Error1}
5   {C2 : Contract Setup2 Msg2 State2 Error2}.
6
7 Definition lift_cm_to_sm (f : ContractMorphism C1 C2) :
8   SystemMorphism (singleton_place_graph C1) (singleton_place_graph C2).
9 Proof.
10  destruct f as [setup_morph msg_morph state_morph error_morph init_coherence
   recv_coherence].
11  apply (build_system_morphism (singleton_place_graph C1) (singleton_place_graph C2)
   setup_morph msg_morph state_morph error_morph);
12  unfold singleton_place_graph, singleton_ntree, sys_init, sys_receive, ntree_fold_left
   in *.
13  - apply init_coherence.
14  - intros.
15  - rewrite <- recv_coherence.
16  - cbn.

```

```

18     now destruct (receive C1 c ctx st op_msg).
19 Defined.
20
21 Definition lift_ctm_to_stm (f : ContractTraceMorphism C1 C2) :
22   SystemTraceMorphism
23     (discrete_sys (singleton_place_graph C1))
24     (discrete_sys (singleton_place_graph C2)).
25 Proof.
26   destruct f as [ct_st_morph gen_fixp cstep_morph].
27   apply (build_st_morph
28     (discrete_sys (singleton_place_graph C1)) (discrete_sys (singleton_place_graph C2))
29     ) ct_st_morph);
30   unfold singleton_place_graph, singleton_ntree, sys_init, sys_receive, ntree_fold_left
31   in *.
32   - apply gen_fixp.
33   - intros * step.
34     assert (ContractStep C2 (ct_st_morph sys_state1) (ct_st_morph sys_state2)
35       -> SingleSystemStep (Node (Contract Setup2 Msg2 State2 Error2) C2 [])
36       (ct_st_morph sys_state1) (ct_st_morph sys_state2))
37     as H_step.
38     {
39       intro cstep.
40       destruct cstep as [c ctx msg nacts rcv_ok].
41       apply (build_sys_single_step _ _ _ c ctx msg nacts).
42       unfold sys_receive.
43       cbn.
44       destruct (receive C2 c ctx (ct_st_morph sys_state1) msg); auto.
45       now destruct t. }
46   apply H_step, cstep_morph.
47   clear H_step.
48   destruct step as [c ctx msg nacts rcv_ok].
49   apply (build_contract_step C1 sys_state1 sys_state2 c ctx msg nacts).
50   unfold sys_receive in rcv_ok.
51   cbn in *.
52   destruct (receive C1 c ctx sys_state1 msg); auto.
53   destruct t.
54   now inversion rcv_ok.
55 Defined.
56
57 End LiftCMtoSM.

```

Listing C.32: The identity contract morphism lifts to the identity system morphism, and compositions lift to compositions. Thus isomorphic contracts lift to isomorphic singleton contract systems under the discrete link graph.

```

1 (* id lifts to id *)
2 Lemma lift_id_cm_to_id_sm
3   '{Serializable Setup} '{Serializable Msg} '{Serializable State} '{Serializable Error}
4   {C : Contract Setup Msg State Error} :
5   lift_cm_to_sm (id_cm C) = id_sm (singleton_place_graph C).

```

```

6 Proof.
7   unfold lift_cm_to_sm, id_cm, id_sm, singleton_place_graph.
8   cbn.
9   f_equal;
10  apply proof_irrelevance.
11 Qed.
12
13 (* compositions lift to compositions *)
14 Lemma lift_cm_to_sm_comp
15   `{Serializable Setup1} `{Serializable Msg1} `{Serializable State1} `{Serializable
16   Error1}
17   `{Serializable Setup2} `{Serializable Msg2} `{Serializable State2} `{Serializable
18   Error2}
19   `{Serializable Setup3} `{Serializable Msg3} `{Serializable State3} `{Serializable
20   Error3}
21   {C1 : Contract Setup1 Msg1 State1 Error1}
22   {C2 : Contract Setup2 Msg2 State2 Error2}
23   {C3 : Contract Setup3 Msg3 State3 Error3}
24   (f : ContractMorphism C1 C2) (g : ContractMorphism C2 C3) :
25   lift_cm_to_sm (compose_cm g f) = compose_sm (lift_cm_to_sm g) (lift_cm_to_sm f).
26 Proof.
27   destruct g as [smorph_g msgmorph_g stmorph_g errmorph_g initcoh_g recvcoh_g].
28   destruct f as [smorph_f msgmorph_f stmorph_f errmorph_f initcoh_f recvcoh_f].
29   unfold compose_cm, compose_sm, lift_cm_to_sm.
30   cbn.
31   f_equal;
32   apply proof_irrelevance.
33 Qed.
34
35 (* isomorphic contracts => isomorphic singleton systems *)
36 Theorem c_iso_csys_iso
37   `{Serializable Setup1} `{Serializable Msg1} `{Serializable State1} `{Serializable
38   Error1}
39   `{Serializable Setup2} `{Serializable Msg2} `{Serializable State2} `{Serializable
40   Error2}
41   {C1 : Contract Setup1 Msg1 State1 Error1}
42   {C2 : Contract Setup2 Msg2 State2 Error2} :
43   contracts_isomorphic C1 C2 ->
44   systems_isomorphic (singleton_place_graph C1) (singleton_place_graph C2).
45 Proof.
46   intro c_iso.
47   destruct c_iso as [f [g [is_iso_1 is_iso_2]]].
48   unfold systems_isomorphic.
49   exists (lift_cm_to_sm f), (lift_cm_to_sm g).
50   unfold is_iso_sm.
51   split;
52   rewrite <- lift_cm_to_sm_comp;
53   try rewrite is_iso_1;
54   try rewrite is_iso_2;

```

```
50     now rewrite lift_id_cm_to_id_sm.  
51 Qed.
```

Glossary

address An address, or wallet address, is a public key used to represent a destination for transactions on a blockchain network. Addresses are used to send and receive digital assets, and they are generated using cryptographic algorithms to ensure security. See also wallet. 19

arbitrage In the context of cryptocurrency, arbitrage refers to the practice of taking advantage of price differences between different decentralized exchanges (DEXs) or different trading pairs within the same DEX to make a profit. If a token is priced differently on different DEXs, an arbitrageur can buy the token on the DEX where it is underpriced, and then sell it on the DEX where it is overpriced, making a profit. 45, 46

automated market maker An automated market maker (AMM) is a type of decentralized exchange (DEX) that uses a mathematical formula to determine the price of assets being traded on the platform. This in contrast to traditional centralized exchanges, which match buyers and sellers and take a cut of the transaction as a fee. AMMs are commonly used in decentralized finance (DeFi) applications, and they play an important role in providing liquidity and enabling the trading of digital assets. 15, 244, 245, 247

Binance Smart Chain Binance Smart Chain (BSC) is a blockchain developed by the Binance exchange. It is built on top of the Ethereum Virtual Machine (EVM) and uses a similar smart contract language to Ethereum. 18, 247

blockchain A blockchain is a distributed and decentralized digital ledger that records a secure and immutable ledger of transactions across a network of computers. Each transaction, bundled into a block, is cryptographically linked to the previous one, forming a chain of blocks, hence the name “blockchain.” Blockchains are commonly used for cryptocurrencies like Bitcoin or Ethereum, but have applications beyond digital currencies including supply chain management, voting systems, and smart contracts. 15, 17

collateralized In DeFi, an asset is collateralized if it is backed by collateral greater than or equal to its value, often as part of a crypto lending scheme in DeFi applications. The borrower puts up the collateral as a guarantee that they will be able to repay the loan. If the borrower defaults, the lender can seize the collateral to recover their funds, often by auctioning the seized collateral. This

arrangement is typically mediated by a smart contract and executed automatically. An asset can be over- or under-collateralized, indicating that the value of collateral exceeds or subceeds, respectively, the value of the collateralized asset. 47

constituent token A constituent token of a pool is a (typically non-fungible) token which can be pooled in exchange for a (typically fungible) token, called a pool token. Whether and under what conditions the pool token can be redeemed for underlying constituent tokens will vary contract by contract, depending on the tokenomics. 44, 47, 243

cross-chain bridge A cross-chain bridge is an application, not necessarily decentralized, that allows for the transfer of assets or information between different blockchains. Possible assets that can be exchanged or transferred include tokens, coins, and other digital assets, even if the blockchains being bridged have different consensus algorithms, security models, and underlying infrastructure. 15, 19, 20

crypto insurance protocols Crypto insurance protocols are DeFi applications that provide insurance coverage for assets in the crypto space. These protocols are designed to mitigate the risk of loss for investors in the event of unexpected events such as hacking, smart contract vulnerabilities, or market crashes. They typically work by pooling funds from multiple investors, who then share the risk of loss. Claims are processed and approved by a variety of different mechanisms, depending on the nature of the insured event and whether or not human intervention is required to assess claims. 15, 241

crypto lending Crypto lending, or decentralized lending, refers to the practice of borrowing and lending digital assets within the context of blockchain-based financial systems, such as DeFi platforms. Users can lend their cryptocurrency holdings to others in exchange for earning interest on their loans, while borrowers can access funds by providing collateral in the form of other cryptocurrencies. These transactions are facilitated through smart contracts on blockchain networks, instead of traditional intermediaries like banks. 15, 20, 239, 241

Curve Curve is a decentralized exchange (DEX) on the Ethereum blockchain that specializes in stablecoins such as USDC, DAI, and USDT. Its key value proposition is low slippage on trades. 19, 46

DAI DAI is a stablecoin on the Ethereum blockchain, pegged to the value of the US dollar. It is minted as a receipt of debt, where the ETH is held in a smart contract as collateral to a loan. Interest rates on the loan fluctuate in response to market prices in order to stabilize the price of DAI. 89, 106, 240, 242

decentralized application A decentralized application (dApp) is a software application that runs on a decentralized network like a blockchain, and is not controlled by any central authority. Decentralized applications can be used for a variety of purposes, ranging from financial applications like exchanges, to gaming platforms, and social media sites. 241, 247

decentralized autonomous organization A decentralized autonomous organization (DAO) is a type of organization that is run using rules encoded in a smart contract, on a blockchain. DAOs are designed to operate without the need for intermediaries or a central authority. Members of a DAO participate by holding and voting with governance tokens, or tokens that represent rights or ownership in the organization, and decisions are made through a consensus mechanism such as voting or staking. DAOs are often used in DeFi as a way to create and manage decentralized investment funds, decentralized exchanges, or other types of decentralized organizations. 89, 247

decentralized exchange A decentralized exchange (DEX) is a type of decentralized marketplace that operates on a blockchain, allowing users to trade cryptocurrencies and digital assets directly with each other without the need for intermediaries. Examples of DEXs include, but are not limited to, automated market makers AMMs and decentralized auctions. 15, 239, 240, 244, 245, 247

decentralized finance Decentralized finance (DeFi) refers to a financial ecosystem built on blockchain technology that largely operates without traditional intermediaries. It enables peer-to-peer transactions and interactions with digital assets through smart contracts and decentralized protocols. DeFi includes various services like lending, borrowing and trading. DeFi applications include, but are not limited to, DEXs, AMMs, crypto lending, and crypto insurance protocols. 15, 20, 106, 247

decentralized governance Decentralized governance refers to a system of decision-making and administration in which power is distributed among a network of individuals or entities, rather than being centralized in a single governing body. In the context of blockchain technology and cryptocurrencies, decentralized governance typically refers to a system in which stakeholders collectively make decisions about the direction and management of a particular network or protocol, often through the use of decentralized voting mechanisms or on-chain proposals. 18

DEX aggregator A DEX aggregator is a platform that enables users to trade cryptocurrencies across multiple decentralized exchanges (DEXs) at once, using complex algorithms to find the best prices across all supported exchanges. 106

entrypoint function A contract entrypoint function is a public-facing function that allow users to interact with the smart contract and make changes to its state. Examples of entrypoint functions in a smart contract might include functions to transfer funds, mint new tokens, vote on proposals, or access information about the state of the contract. In general, the entrypoint functions are defined by the contract developer and are intended to provide a way for external users to interact with the contract in a meaningful way. 16

Ethereum Ethereum is a decentralized, open-source blockchain platform that enables the creation and execution of smart contracts and decentralized applications (dApps). It was created in 2015 by Vitalik Buterin and has since become one of the largest and most widely used blockchain platforms in the world. Its native token, ETH, is used to pay for gas fees. It also supports Solidity, a Turing-complete programming language, which allows developers to build a wide variety of applications,

from decentralized exchanges and prediction markets to games and social networks. 16, 19, 22, 108, 239, 240, 242–245, 247

Ethereum Virtual Machine The Ethereum Virtual Machine (EVM) is a decentralized, Turing-complete virtual machine that serves as the runtime environment for executing smart contracts on the Ethereum blockchain. Developers can write and deploy smart contract code in high-level programming languages like Solidity, which is then compiled into bytecode that can be understood and executed by the EVM. 239, 244, 247

flash loan In DeFi, a flash loan is a type of decentralized loan that allows a user to borrow funds for a single, atomic transaction, without collateral and with a very low interest rate. The must loan is repaid automatically at the end of the transaction. Any transaction that takes out a flash loan which does not end in repayment is invalid. Flash loans are used for a variety of purposes in DeFi, including exploiting market inefficiencies and executing arbitrage strategies. 18, 19, 242

flash loan attack A flash loan attack is a type of exploit in DeFi in which an attacker uses funds borrowed with a flash loan to execute a series of manipulative trades or arbitrage opportunities. The attacker profits from these trades at the expense of other users, such as liquidity providers, taking advantage of the borrowed funds without bearing any risk or requiring collateral. 18, 106

fractionalize In the context of non-fungible tokens (NFTs), “fractionalizing” refers to the process of dividing ownership of a single NFT into smaller, tradeable fractions or shares, typically as fungible tokens. 38

front-running attack In a front-running attack, an actor inserts certain transactions in a block ahead of others which, by their order, are profitable by gaining an unfair advantage and causing financial losses to victims. This manipulation involves monitoring pending transactions, identifying lucrative opportunities, and quickly submitting their own transaction with higher fees to exploit market inefficiencies. 17, 36

fungible See fungible token. 38, 240, 243

fungible token A fungible token is a type of digital asset that represents a unit of value that is interchangeable with other units of the same value. This means that each unit of the token is identical and interchangeable with any other unit. Examples of fungible tokens include the cryptocurrencies BTC or ETH, stablecoins USDC or DAI, or other tokens conforming to the ERC20 token standard. 242, 247

gas Gas is a term used to describe the fee required to process a transaction on Ethereum and other blockchains, paid in the blockchain’s native token (*e.g.* ETH). The amount of gas required for a transaction depends on the computational complexity of the operation being performed and the demand for block space. 20, 67, 105, 241

governance token A governance token is a type of token used to facilitate decision-making in a decentralized organization, such as a DAO. Holders of governance tokens can be given voting rights that allow them to participate in decisions relevant to the organization, such as changes to the organization's protocol or the allocation of funds. Governance tokens can also be used to propose and vote on changes to the organization's governance structure, as well as to elect or recall members of the organization's leadership. 89, 106, 241

liquidity provider In DeFi, a liquidity provider is an entity who supplies their digital assets to liquidity pools, enabling trading and financial activities on the platform and earning rewards in return. 16, 66, 242, 243

LP token An LP token, short for liquidity provider token, is a digital asset issued to liquidity providers in DeFi platforms, representing their share in a liquidity pool and enabling them to manage and withdraw their deposited assets and rewards. 47, 106, 243

non-fungible See non-fungible token. 240, 243

non-fungible token A non-fungible token (NFT) is a unique and indivisible digital asset built on blockchain, often representing ownership of distinct items like digital art, collectibles, or virtual real estate. NFTs are associated with the ERC721 token standard on Ethereum. 37, 242, 243, 247

peg A peg in the context of digital assets refers to a fixed exchange rate between a cryptocurrency and a real-world asset, such as the US dollar or gold. A pegged asset or cryptocurrency is one whose value maintains a stable value at some peg. Some pegged assets are backed by a reserve of the asset they are pegged to, which can help ensure the stability of the peg even during market fluctuations. Others maintain their peg algorithmically through various debt mechanisms. 106, 240, 244

pool In DeFi, a pool is a shared aggregation of assets that provides liquidity for a specific asset or market. Users can add their own assets to the pool and receive a proportional share of the pool's tokens, called LP tokens or pool tokens. These tokens give users a share of the fees generated by the trades that occur in the pool. 240, 243

pool token A pool token is a (usually fungible) token which can be minted in exchange for pooling a (usually non-fungible) constituent token into a pool. The terms governing when and under what conditions the pool token can be redeemed for underlying constituent tokens are set by the contract governing the pool. 38, 47, 240, 243

price oracle In DeFi, a price oracle is a source of truth for the current price of a cryptocurrency or other asset. A price oracle is typically a smart contract to which information about the price of an asset from an external data source, such as an exchange, is pushed, and provides it to other smart contracts in a standardized format. In order for a price oracle to be effective, it must be reliable and resistant to tampering and manipulation. 19, 106

slippage In the context of decentralized exchanges (DEXs) and automated market makers (AMMs), slippage refers to the difference between the expected price of an asset and the actual price at which it is traded, which can differ due to the changing market conditions such as an large trades. This is particularly relevant in DeFi, where price slippage can have a significant impact on the profitability of trading strategies such as arbitrage. 240

smart contract A smart contract is a computer program, stored on a blockchain, that automatically executes when certain conditions are met. Smart contracts do not require intermediaries to enforce their terms. They can facilitate exchange of assets, such as cryptocurrencies, and have a wide variety of use cases. 15, 18, 239

Solidity Solidity is a high-level programming language for writing smart contracts that can be compiled into bytecode be executed on the Ethereum Virtual Machine (EVM), and stored on the Ethereum blockchain. 22, 241, 242

stablecoin A stablecoin is a type of cryptocurrency designed to maintain a stable value, typically pegged to a stable asset, such as a fiat currency (e.g., USD, EUR) or a commodity (e.g., gold). The main goal of stablecoins is to reduce price volatility commonly associated with traditional cryptocurrencies like Bitcoin or Ethereum, providing a more reliable medium of exchange and store of value. Stablecoins achieve stability through various mechanisms, such as collateralization, algorithmic control, or a combination of both, which are designed to ensure that the value of the stablecoin remains relatively constant over time. 15, 18, 20, 36, 89, 106, 240, 244, 245

synthetics Synthetics are digital assets that are designed to track the price of another asset, such as a traditional financial instrument, a commodity, or another cryptocurrency. They include, but are not limited to, stablecoins, and like stablecoins maintain their peg through various mechanisms such as collateralization, algorithmic control, or a combination of both. 15, 106

Tezos Tezos is a third-generation, proof-of-stake blockchain which supports smart contracts written in Michelson. Its native token is XTZ. 15, 22, 244, 247

token A token is a digital asset that represents a unit of value and can be traded on a blockchain platform. Tokens can serve a variety of purposes, such as representing ownership in a company, access to a particular product or service, or as a medium of exchange. Tokens can be created using smart contracts on blockchain platforms, such as Ethereum, and they are typically stored and traded using a digital wallet. Tokens are standardized on most blockchains, for example the ERC20 and ERC721 standards on the Ethereum blockchain, and the FA2 standard on the Tezos blockchain. 89, 106, 107, 242–245, 247

tokenization Tokenization is the process of converting ownership rights or assets, external to the blockchain, into a blockchain-based digital representation such as a token. 244

tokenize See tokenization. 245

tokenized carbon credit Tokenized carbon credits are tokens, stored on a blockchain, that represent a carbon credit, typically a unit of carbon emissions reduction. Most are tokenized from legacy organizations that verify carbon emissions reductions such as Verra. 37, 38

tokenomics Tokenomics, a portmanteau of “token” and “economics,” refers to the study and design of the economic system and incentives behind a cryptocurrency or digital token. It encompasses the rules, policies, and mechanisms that determine the creation, distribution, usage, and value of the tokens within a blockchain or decentralized ecosystem. Tokenomics plays a crucial role in determining the success and sustainability of a cryptocurrency or digital token project. 36, 106, 240

transaction A transaction is an action or operation that alters the state of the blockchain network. It typically involves the transfer of digital assets, such as cryptocurrencies or tokens, from one user to another, but can also encompass various other types of interactions, such as smart contract executions, data storage, or authentication processes. Transactions are recorded in blocks and cryptographically linked together. Each transaction must be verified and validated by network participants, known as miners or validators, to ensure its authenticity and compliance with the network’s consensus rules before being added to the blockchain. 18, 245

Uniswap Uniswap is a decentralized exchange (DEX) built on the Ethereum blockchain. It was the first, and is arguably the most popular, automated market maker (AMM). 43, 44, 46, 66

USDC USDC (USD Coin) is a stablecoin pegged to the value of the US dollar. 19, 240, 242

USDT Tether, or USDT, is a stablecoin pegged to the value of the US dollar. 19, 240

vault A vault is a secure digital container or smart contract function designed to hold and manage digital assets, such as cryptocurrencies or tokens, in a decentralized and tamper-resistant manner. Vaults are often utilized in DeFi protocols to facilitate various financial operations like lending, borrowing, and yield farming. They are programmed with specific rules and conditions governing the access, withdrawal, and management of the assets they hold, and they aim to provide a high level of security and transparency. Vaults can have different strategies and mechanisms to optimize asset utilization, protect against risks, and maximize returns for users within the blockchain ecosystem. 19

wallet In the context of a blockchain, a wallet refers to a public, private key pair which can be used to execute transactions on a blockchain and hold digital assets. Transactions originating from a wallet must be signed by the private key. The term wallet also refers to software or a device that stores and manages blockchain-based digital assets. 19, 239

yield aggregator A yield aggregator is a type of smart contract that allows users to automatically aggregate their crypto assets into various yield-generating DeFi protocols. The purpose of a yield aggregator is to maximize the yield received by the user on their cryptocurrency investments. Yield aggregators typically work by automatically re-allocating funds to the highest yielding protocols, taking into account various factors such as fees, liquidity, and performance. 19, 20, 106

yield farming Yield farming is the practice of providing liquidity to DeFi protocols in exchange for rewards in the form of interest or tokens. The rewards can come from the fees generated by the protocols, or through inflation, and are typically distributed to liquidity providers as a way of incentivizing them to provide liquidity and ensure the stability of the platform. 15, 20, 106, 245

Acronyms

AMM Automated market maker. 15, 16, 19, 20, 36–38, 43, 44, 46, 47, 66, 106, 117, 118, 239, 241

BSC Binance Smart Chain. 18

DAO decentralized autonomous organization. 89, 106, 243

dApp decentralized application. 241

DeFi decentralized finance. 15, 18, 20, 37, 43, 106, 107, 118, 133, 239–246

DEX Decentralized exchange. 15, 20, 106, 107, 118, 239–241

ERC20 A token standard for fungible tokens on the Ethereum blockchain. 38, 242, 244

ERC721 A token standard for non-fungible tokens on the Ethereum blockchain. 243, 244

ETH The native token of the Ethereum blockchain. 20, 240–242

EVM Ethereum Virtual Machine. 22, 239

FA2 A generic standard on the Tezos blockchain for fungible and non-fungible tokens. 244

NFT See non-fungible token. 37, 38, 242

XTZ The native token of the Tezos blockchain. 244