



NOVA
NOVA SCHOOL OF
SCIENCE & TECHNOLOGY

DEPARTMENT OF
COMPUTER SCIENCE

JOÃO CARLOS RAPOSO DOS REIS
Bachelor in Computer Science

TOWARDS A SOLIDER SOLIDITY

MASTER IN COMPUTER SCIENCE
NOVA University Lisbon
December, 2023



TOWARDS A SOLIDER SOLIDITY

JOÃO CARLOS RAPOSO DOS REIS

Bachelor in Computer Science

Adviser: António Ravara
Associate Professor, NOVA University Lisbon

Co-adviser: Mário Pereira
Assistant Professor, NOVA University Lisbon

Examination Committee

Chair: Nuno Preguiça
Full Professor, NOVA University Lisbon

Rapporteur: Simão Melo de Sousa
Full Professor, University of Algarve

Adviser: António Ravara
Associate Professor, NOVA University Lisbon

Towards a Solider Solidity

Copyright © João Carlos Raposo dos Reis, NOVA School of Science and Technology, NOVA University Lisbon.

The NOVA School of Science and Technology and the NOVA University Lisbon have the right, perpetual and without geographical boundaries, to file and publish this dissertation through printed copies reproduced on paper or on digital form, or by any other means known or that may be invented, and to disseminate through scientific repositories and admit its copying and distribution for non-commercial, educational or research purposes, as long as credit is given to the author and editor.

This document was created with the (pdf/Xe/Lua)LaTeX processor and the [NOVAtesis](#) template (v7.0.3) [[novathesis-manual](#)].

ABSTRACT

Blockchains are being widely adopted since its inception. Since Ethereum popularized smart contracts and introduced its platform for developing and deploying these programs on the blockchain, there has been an unprecedented boom in decentralized applications (dApps). Smart contracts bring a new flexibility which we didn't see in the beginning with the first generation blockchains. These programs rule important agreements between two parties involving transacting valuable assets between them. Recently we have seen a record of stealing assets from the blockchain, as they exploded and became more popular in 2020. Many known vulnerabilities include: Re-entrancy, Gasless Send, Phishing with tx.origin and Type Casts.

This reinforces the importance to ensure that these programs are correct. However this is still a hard task, as the programming languages used to develop them often do not really help and smart contracts are still being exploited. Static analysis tools can prevent these bugs, if they are robust enough to detect most of the vulnerabilities which are known and not known, and ensure that these contracts are developed with the best design practices. Many of those existing tools and programming languages are still in a very early stage and not yet prepared to be used in production environments.

The development of smart contracts can be done in two types of languages: domain-specific and general-purpose languages. Domain-specific languages, specifically designed for writing smart contracts, are generally considered safer than their general-purpose counterparts. Solidity, created by the Ethereum Foundation, is a DSL and stands out as the most popular language to date. However, it is not without its flaws, primarily stemming from its complexity and Turing completeness. While recent academic efforts have introduced languages aimed at addressing known vulnerabilities in Solidity, the language's widespread adoption needs continuous improvements to improve its safety.

Our contribution targets a vulnerability in Solidity related to Type Casts. We introduce a proof-of-concept language inspired by Featherweight Solidity, a formalization of Solidity as a subset with an improved type system. Specifically, our language features an extended typing address to mitigate this vulnerability.

RESUMO

As blockchains ganharam ampla adoção desde sua criação. Desde que a Ethereum popularizou os smart contracts e introduziu sua plataforma para desenvolver e implantar esses programas na blockchain, houve um crescimento sem precedentes nas aplicações descentralizadas (dApps). Smart contracts trazem uma nova flexibilidade que não vimos no início com as blockchains de primeira geração. Esses programas regem acordos importantes entre duas partes envolvendo a transação de ativos valiosos entre elas. Recentemente, observamos um aumento no roubo de ativos da blockchain. Muitas vulnerabilidades conhecidas incluem Re-entrancy, Gasless Send, Phishing com tx.origin e Type Casts.

Isso reforça a importância de garantir que esses programas estejam corretos. No entanto, ainda é uma tarefa difícil, pois as linguagens de programação usadas para desenvolvê-los muitas vezes não ajudam realmente, e smart contracts continuam sendo explorados. Ferramentas de análise estática podem prevenir esses bugs se forem robustas o suficiente para detectar a maioria das vulnerabilidades conhecidas e desconhecidas, e garantir que esses contratos sejam desenvolvidos com as melhores práticas de design. Muitas dessas ferramentas e linguagens de programação existentes ainda estão numa fase inicial e ainda não estão preparadas para serem usadas em ambientes de produção.

O desenvolvimento de smart contracts pode ser feito em dois tipos de linguagens: linguagens de domínio específico e linguagens de propósito geral. As primeiras são desenvolvidas especificamente para escrever smart contracts, sendo geralmente consideradas mais seguras do que as segundas. O Solidity, criado pela Ethereum Foundation, destaca-se como a linguagem mais popular até o momento. No entanto, não está isento de falhas, principalmente devido à sua complexidade e ser Turing complete. Recentemente têm surgido linguagens acadêmicas a abordar vulnerabilidades conhecidas no Solidity. Contudo a já ampla adoção da linguagem aprimorar sua segurança nas áreas em que atualmente apresenta deficiências.

A nossa contribuição visa uma vulnerabilidade no Solidity: Type Casts. Introduzimos uma linguagem de prova de conceito inspirada no Featherweight Solidity, uma formalização do Solidity como um subconjunto com um sistema de tipos estendido. Especificamente, a nossa linguagem apresenta um tipo address estendido mitigando essa vulnerabilidade.

CONTENTS

1	Introduction	1
1.1	Motivation	1
1.2	The Problem	1
1.3	Objectives	2
1.4	Contributions	2
1.5	Organization of this document	3
2	Background	4
2.1	General Context	4
2.2	Distributed System	4
2.2.1	General Context	4
2.2.2	Consensus Protocols	5
2.3	Blockchain	5
2.3.1	Block	6
2.3.2	Chain	6
2.3.3	Ledger	6
2.3.4	Consensus Protocols	7
2.3.5	Blockchain Types	8
2.3.6	Examples of Blockchain	9
2.4	Smart Contracts	10
2.4.1	General Context	10
2.4.2	Brief Presentation of Industry Languages	12
3	State of the Art	19
3.1	Academic Languages	19
3.1.1	Vyper	19
3.1.2	Scilla	20
3.1.3	Bamboo	20
3.1.4	Flint	20

3.1.5	Obsidian	20
3.2	Known Vulnerabilities in Solidity	21
3.2.1	Type Casts	21
3.2.2	Reentrancy	22
3.2.3	Phishing waith tx.origin	23
3.2.4	Delegatecall Injection	24
3.2.5	Unprotected Selfdestruct	25
3.2.6	Gasless Send/Transfer	26
3.2.7	Unprotected ether withdrawal	27
3.3	Tools To Verify Solidity Smart Contracts	28
3.3.1	Slither	28
3.3.2	Mythril	29
3.3.3	SMTChecker	29
3.3.4	Solidifier	30
3.3.5	Solc-verify	30
3.3.6	Other Tools	31
4	Featherweight Solidity	33
4.1	General Context	33
4.2	Meta Variables	33
4.3	Grammar	34
4.3.1	Contract Definition	35
4.3.2	Constructor Declaration	35
4.3.3	Expression	36
4.3.4	Values	36
4.3.5	Types	36
4.4	Configuration	36
4.5	Lookup Functions	37
4.6	Auxiliary Predicates	38
4.7	Operational Semantics	39
4.7.1	Conditional and Sequential Statements	39
4.7.2	Variable Operations	40
4.7.3	Mappings	41
4.7.4	Contract Instantiation	41
4.7.5	Balance and Address	42
4.7.6	Cast	42
4.7.7	Money Transfer	42
4.7.8	Function Calls	43
4.8	Type System	44
4.8.1	Axioms	44
4.8.2	Standard Rules	45

4.8.3	Mappings	46
4.8.4	Contract Instantiation and Access	46
4.8.5	Casts and Money Transfers	47
4.8.6	Functions	47
5	FS 2.0: Revised Operational Semantics	49
5.1	Grammar	49
5.1.1	Contract Definition	50
5.1.2	Constructor Declaration	50
5.2	Configuration	50
5.3	Contract Example in FS 2.0	52
5.4	Auxiliary Predicates	53
5.4.1	Default Values	53
5.4.2	Multiple Inheritance	53
5.4.3	Functions	55
5.4.4	State Variables	57
5.4.5	Update Contract Balance	58
5.4.6	Constructors	58
5.5	Operational Semantics Rules	59
5.5.1	Arithmetic and Boolean Operations	59
5.5.2	Variable Operations	60
5.5.3	Constructor Rule	61
5.5.4	Conditional and Sequential Statements	62
5.5.5	Mappings	63
5.5.6	Functions	64
5.5.7	Computation Rules	67
5.5.8	Operational Semantics by Example	69
5.5.9	Implementation in OCaml	72
6	Property-based Testing	75
6.1	General Context	75
6.2	Arithmetic and Boolean Operations	75
6.3	Conditional Expressions	76
6.4	Future Work	77
7	FS 2.0: Type System	80
7.1	General Context	80
7.2	Type Environment	80
7.3	Rules	81
7.3.1	Rules Without Inheritance	81
7.3.2	Rules With Inheritance	86
7.3.3	Type System by Example	90

7.3.4	Implementation in OCaml	91
8	Case Studies	95
8.1	General Context	95
8.2	Contract Hierarchy and C3 Linearization Results	99
8.3	Execution Example	100
9	Conclusions and Future Work	103
9.1	Conclusion	103
9.2	Future Work	103
	Bibliography	105

INTRODUCTION

1.1 Motivation

Blockchains are being widely adopted since its inception. Since Ethereum popularized smart contracts and introduced its platform for developing and deploying these programs on the blockchain, there has been an unprecedented boom in decentralized applications (dApps). Smart contracts bring a new flexibility which we did not see in the beginning with the first generation blockchains. These programs rule important agreements between two parties involving transacting valuable assets between them. It is therefore important to ensure that these programs are correct. However this is still a hard task, as the programming languages used to developed them often do not really help and smart contracts are still being exploited. Static analysis tools can prevent these bugs, if they are robust enough to detect most of the vulnerabilities which are known and not known, and ensure that these contracts are developed with the best design practices. Many of those existing tools and programming languages are still in a very early stage and not yet prepared to be used in production environments.

1.2 The Problem

Since smart contracts once deployed cannot be modified, unless they use proxy contracts, which are contracts that can be upgraded, there is an urge that these programs are auditable. Being auditable, means that they are certified to do whichever features and only these features they are intended to. The main job of an auditor is to find any errors or vulnerabilities that could cause a loss of funds and/or control of the project. Static analysis can then help them to do this important task.

Recently we have seen a record of stealing assets from the blockchain , as they exploded and became more popular in 2020. One example occurred in March 2022, when [Ronin said](#) that hackers stole cryptocurrency worth almost 615 milion of dollars from its systems. At the time, this was a record-making in smart contracts history. There is a [good article](#) explaining what was the purposes of this project and how it got exploited by the hackers.

Understanding the underlying language we are studying is crucial when building tools to deal with vulnerabilities. To improve the safety of a programming language, there are two possible approaches: designing a language to be proof against vulnerabilities or developing static analysis tools to detect them. We have encountered both approaches in either academic or industrial contexts and we will delve deeper into them in Chapter 3.

1.3 Objectives

The goal of the dissertation is, in a nutshell, to develop solutions to deal with presently undetected vulnerabilities in Solidity. To achieve this, our initial step involved an in-depth analysis of the state-of-the-art landscape, encompassing programming languages designed for smart contract development, existing vulnerabilities within Solidity, and the corresponding static analysis tools employed for their detection. One of our objectives is to discern the efficacy of domain-specific languages versus general-purpose languages for smart contract development. Additionally, we seek to identify vulnerabilities within Solidity that may not have received adequate attention. Our investigation revealed that Type Casts is a vulnerability that has been relatively overlooked. This vulnerability is related to the weakness in the typing of contract addresses, which is the most common way to refer to them. When casting an address to a contract instance and calling a function of it, Solidity's compiler does not verify whether such a function is defined in the target contract. When we try to execute it and it does not exist, an exception is thrown, resulting in users losing funds without the possibility of reimbursement. However, with a more refined address type, we could prevent the function from being called, thus safeguarding the funds being sent. Consequently, we have undertaken the development of an environment where this particular vulnerability is mitigated.

1.4 Contributions

There have been some formalizations on Solidity semantics and type system [29, 14]. However, being a relatively new language, formalizations are currently dispersed and focus on different aspects. In this master thesis we gathered all together, introducing many concepts on the one we are extending: Featherweight Solidity [21, 6].

We present herein a proof-of-concept: a refined implementation and formalization of Solidity. We introduce Featherweight Solidity 2.0 as a Solidity subset with a new type system. This new type system addresses what it was called as Type Casts vulnerability, where there is currently no way to have subtyping in address type. Therefore our language supports an extended typing address, aiming to fix this problem. Beyond formalizing it, we also leveraged the expressiveness power of OCaml, a functional programming language known for its strong type inference and formal verification capabilities, to implement Featherweight Solidity as a proof-of-concept compiler.¹

¹<https://github.com/jcrrreis/featherweight-solidity>

1.5 Organization of this document

The rest of this document is structured as follows:

- Chapter 2 - Foundational concepts of blockchain and smart contracts.
- Chapter 3 - A study of the state of the art in smart contracts, covering academic languages, vulnerabilities and static analysis tools in Solidity.
- Chapter 4 - Introduction to Featherweight Solidity, a calculus language that models the core features of Solidity.
- Chapter 5 - Presentation of our revised operational semantics for FS 2.0.
- Chapter 6 - Demonstration of how we use property-based testing to prove properties of our operational semantics.
- Chapter 7 - Presentation of our revised type system rules for FS 2.0.
- Chapter 8 - Illustration of our work using a complex case study.
- Chapter 9 - Summary of our work and suggestions for future development.

BACKGROUND

2.1 General Context

Blockchain has been one of the hottest and promising technologies of the decade. In 2008, Satoshi Nakamoto (an alias used by the creator/creators) released a paper in which he explained a new form of digital payment without relying in a trusted third party, using this technology [19]. This was a big event as this was the first time someone did conceptualize a **decentralized blockchain** effectively, a new kind of distributed system. Later on, in 2014 a new decentralized blockchain surged: Ethereum [2]. This new architecture popularized an old concept in computer science: **smart contracts**. They were first proposed by Nick Szabo in 1996 who coined the term, using it to refer to "a set of promises, specified in digital form, including protocols within which the parties perform on these promises" [24]. A smart contract is simply a computer program that run on a blockchain and allow to automatically execute, control or document legally relevant events and actions between two parties.

2.2 Distributed System

2.2.1 General Context

A distributed system is a set of hardware and software components interconnected through a communications infrastructure, which cooperate and coordinate with each other only by exchanging messages, for the execution of distributed applications [25]. In a distributed system, data is replicated through many equivalent nodes that can be replaced by another if one fails, improving performance and reliability of the system. Nevertheless, having multiple copies of the same data in many nodes can lead to consistency problems and this systems need to guarantee that its users always view a single coherent system, requiring that all nodes always have consistent replica of the system.

Consistency in distributed systems means that every node has the same view of data at a given point in time regardless of whichever client has updated the data [25]. When you make a request to any node, you receive the exact same response so that it looks

like there is only a single node performing all the operations. There are many types of consistency models in distributed systems, where the extremes are **strong consistency** and **weak consistency** [4]. Blockchains by their nature, being an immutable distributed ledger that often stores and transacts valuable assets, need to have a strong consistency model. This consistency model requires that all accesses are seen by all nodes in the same order and there are one and only one consistent state. On the other hand, weak consistency does not require that all accesses are seen in the same order and there are no guarantees that all nodes have the same data at any times. Strong consistency requires a much stronger synchronization between nodes and for this reason it is usually slower to accomplish than weak consistency.

2.2.2 Consensus Protocols

To achieve consistency, distributed systems need a **consensus protocols**, which are the basis for the state machine replication approach to distributed computing [17]. However, reaching consensus on distributed networks, in a safe and efficient way, is far from being an easy task. If some nodes fail or act maliciously the whole distributed system is compromised. This is known as the Byzantine Generals problem (or Byzantine fault). The Byzantine Generals problem¹ was conceived in 1982 and is a condition of a distributed computer system, where components may fail and there is imperfect information on whether a component has failed [16]. A system that is able to resist the class of failures derived from the Byzantine Generals' Problem, has a property called **Byzantine Fault Tolerant** (BFT). A BFT system is able to continue operating even if some of the nodes fail or act maliciously.

One of the first algorithms proved to be correct to achieve consensus in a distributed system was Paxos [15]. This algorithm is used to achieve consensus among a distributed set of computers that communicate via an asynchronous network. One or more clients proposes a value to Paxos and we have consensus when a majority of systems running Paxos agrees on one of the proposed values. Even today Paxos is widely used by many big cloud computing companies such as Amazon, Google and Microsoft.

2.3 Blockchain

Blockchain is a technology that aims to allow digital information to be recorded and distributed, but not edited. As a distributed ledger technology, the blockchain is intentionally designed to be highly resistant to modification and frauds (such as double-spending) [27].

¹This analogy behind the Byzantine Generals Problem is that several generals are besieging Byzantium and they have surrounded the city, but they must collectively decide when to attack. If all generals attack at the same time, they will win, but if they attack at different times, they will lose. The generals have no secure communication channels with one another because any messages they send or receive may have been intercepted or deceptively sent by Byzantium's defenders [28].

A distributed ledger is a database that is consensually and synchronized shared between multiple nodes, any changes made to the ledger are reflected and copied to all participants, which check the validity of the operation. Every node in a distributed ledger processes and validates every record.

To understand better how does blockchains work we need to know what is a block, a chain and a consensus protocol. Next, I will present these concepts and explain what role each one have in this architecture.

2.3.1 Block

Blocks are data structures that store immutable data. Each block must store a pointer to the **hash of the previous block**, a **timestamp** representing the time when a block has been validated by the network and more data that varies depending on the blockchain [5]. This blocks are permanently stored in a distributed ledger, meaning that a blockchain only grows and there are no delete operations. Although this is important to secure blockchains, this also is a problem because blockchain grow unboundly over time. To overcome this problem, blockchains in general use Merkle trees [18]. A Merkle tree is a data structure, that is a generalisation of hash lists and hash chains, where every leaf node is labelled with the hash of a data block and every non-leaf node is labelled with the cryptographic hashes of the labels of its child nodes. This data structures allow efficient and secure verification of the contents of large data structures, without relying in the actual data.

2.3.2 Chain

A chain is a set of interconnected ordered blocks, formed through the pointer to the hash of the previous block that each block has. This mechanism turns out to be similar to linked lists but instead of pointing to the next block, each block points to the previous one [5]. Adding a new block to a chain is a transactional operation: either a block is committed, if they are validated by nodes participating in the consensus protocol or the transaction is aborted if a block is malicious. As hashes are cryptographically derived from the block data, collisions are very unlikely and one change in any block will invalidate all the subsequent blocks in the chain, because all their hashes will also change.

2.3.3 Ledger

A ledger is a database with the purpose of recording contracts and transactions. Centralized ledgers are used by companies to contain all the accounts for recording transactions relating to a company's assets, liabilities, owners' equity, revenue, and expenses. A distributed ledger is a database that is synchronized and accessible across different sites and geographies by multiple participants. This type of ledgers are used by blockchains and it is a decentralized form of a ledger. Distributed ledgers are more secure and reduce cyber

attacks and financial fraud because it does not have a central point of failure and are better able to withstand malicious attacks.

2.3.4 Consensus Protocols

Blockchains are intentionally designed to be decentralized, working as a digital ledger that is maintained by a distributed network of computer nodes. For this reason, blockchain technology allowed the creation of trustless economic systems, where transparent and reliable financial transactions could be executed without the need for intermediaries. Just as most distributed systems, the participants of a cryptocurrency network need to regularly agree on the current state of the blockchain, known as **consensus achievement**.

As blockchains are decentralized distributed systems used to transfer valuable digital assets, they need to be **BFT**. To be a BFT, a system needs to embed a **consensus protocol**. However if a node or group of nodes control over 50% of a blockchain hashing power they can change the order of blocks, reverse past approved blocks (which would then lead to a double-spending problem) and prevent other nodes from mining new blocks. This is known as the 51% attack [9]. Blockchain's consensus protocol discourage attackers from doing this attack by making it cost inefficient.

Although there are many consensus protocols currently used in blockchains, the two most popular ones are **Proof-of-Work** (PoW) and **Proof-of-Stake** (PoS), which I will explain next.

2.3.4.1 Proof-of-Work

This concept was invented in 1993 by computer scientists Cynthia Dwork and Moni Naor to combat e-mail spamming [8]. Eventually in 2009 when Bitcoin was released it was adopted as its consensus protocol. Proof-of-Work is a system where many miners compete with each other to solve a very complex puzzle, but easy to verify. The first miner solving this puzzle will then broadcast this block to the network and other miners will verify that the solution is correct. If the solution is approved by other nodes, the miner which solved the puzzle will earn a reward. Although this ensures that it is extremely costly to control the majority of the network, preventing a 51% attack, as it would require a lot of computational power to do it and therefore making the attack inefficient, this could also have an impact on environment, if the energy used by miners are from fossil fuels. This protocol is currently used in blockchains like Bitcoin and Ethereum.

2.3.4.2 Proof-of-Stake

This protocol surged in 2012, when [Peercoin](#) was launched and used both Proof-of-Work and Proof-of-Stake for its consensus protocol [23]. Therefore it was adopted by many other blockchains as their consensus protocol. [Cardano](#), [Solana](#) and [Polkadot](#) are examples of 3rd generation blockchains that adopted a fully PoS consensus algorithm and Ethereum

is planning switching from PoW to PoS in the first quarter of 2022. This algorithm uses validators instead of miners. Validators are coin owners that simply deposit (or stake) an amount of coins into the network. They are selected randomly by the algorithm to verify new block. This action is called minting, which is equivalent to mining in PoW. This is one of the differences from PoW, instead of being a competition-based algorithm, it chooses validators randomly, with validators with higher amount of coins staked having an higher probability of being choose. In order to perform a 51% attack a validator needs to own more than 50% of the token supply, meaning that if Bitcoin used PoS an attacker would need 383 billion euros to perform the attack. Also, Proof-of-Stake requires much less computational power and consequently have a much lesser environment footprint than Proof-of-Work.

2.3.5 Blockchain Types

Blockchain types can be characterized as permissionless, permissioned or both. A permissionless blockchain allow any user to join the network and don't restrict the rights of the nodes. Conversely, a permissioned blockchain can restrict access to the network to certain nodes and can also restrict the rights of those nodes on the network. There are 4 types of blockchain: **Public Blockchains**, **Private Blockchains**, **Consortium Blockchains** and **Hybrid blockchains** [20].

2.3.5.1 Public Blockchains

Public blockchains are permissionless, allowing anyone to join and are completely decentralized. They allow all nodes of the blockchain to have equal rights to access blockchain, create and validate new blocks. Some examples of public Blockchains are Bitcoin and Ethereum.

2.3.5.2 Private Blockchains

Private blockchains are permissioned, controlled by a single organization. The central authority controls the rights of each node to perform function. This type of blockchains are only partially decentralized, because public access to these networks are restricted. The most prominent private blockchain is [HyperLedger Fabric](#).

Both private and public blockchains have drawbacks, with public blockchains tending to have longer validation times and private blockchains being more vulnerable to fraud and bad actors. The next two types of blockchain, address these drawbacks.

2.3.5.3 Consortium Blockchains

Consortium blockchains are permissioned blockchains governed by a group of organizations rather than a single one. Therefore this blockchains are more decentralized than private ones, resulting in higher levels of security. [Global Shipping Business Network](#)

(GSBN), a blockchain built by [CargoSmart](#), is an example of a consortium blockchain, aiming to digitalize the shipping industry and allow maritime industry operators to work more collaboratively.

2.3.5.4 Hybrid blockchains

An Hybrid blockchain is a blockchain that is controlled by a single organization, but with a level of supervision performed by the public blockchain, which is requires to perform certain transaction validations. An example of a hybrid blockchain is [IBM Food Trust](#), which aims to improve efficiency throughout the whole food supply chain

2.3.6 Examples of Blockchain

2.3.6.1 Bitcoin

Bitcoin, as I already mentioned in this article, was the first decentralized blockchain released. Launched in early 2009, Bitcoin bring to the internet a new revolutionary payment process, fully decentralized and without a central authority, meaning that it does not rely in a trusted third party, like Visa, MasterCard or PayPal. Everyone can create a bitcoin address (equivalent to a bank account in traditional financial system) without needing an approve, send transactions without approve (only miners need to confirm validity of transactions) and become a miner. However bitcoin transactions still take 10 or more minutes to be processed, making them much slower than transactions made by other centralized protocols. Because of this many question if it is a viable alternative to them and many believe that Bitcoin could just become a digital store of value, a digital alternative to gold.

2.3.6.2 Ethereum

Ethereum was launched in 2015 and brought with it the ability to create applications that run on the blockchain. The main difference from Bitcoin is that Ethereum is more than just a cryptocurrency, it is a platform. It allows developers to deploy smart contracts and decentralized apps (DApps), these run on a fully decentralized blockchain, without any downtime or interference from a third party and is powered by its own native cryptocurrency: Ether (ETH). Nevertheless, as each block in a blockchain is immutable, unlike traditional computer programs, these applications once deployed cannot be modified and redeployed, meaning that if a security or safety breach is detected it cannot be corrected, unless owners deploy it to a new address.

2.3.6.3 HyperLedger Fabric

HyperLedger Fabric is an open-source platform managed by Linux Foundation designed to create business-to-business (B2B) businesses and cross-industry applications. As we already have seen, it is a private blockchain with a modular architecture that delivers high

degrees of confidentiality, flexibility, resiliency, and scalability. Unlike Ethereum, where all transactions are public, in HyperLedger Fabric access is limited. This platform also supports smart contracts, they are packaged in a [chaincode](#) and then can be deployed into blockchain. It also has a flexible consensus protocol, that can be configurable by developers using well-established toolkits for Crash Fault-Tolerant (CFT) or Byzantine Fault-Tolerant (BFT) ordering.

2.3.6.4 Ripple

Ripple is a blockchain created by three former bitcoin developers. Their goal was to create a more sustainable and faster consensus algorithm than PoW, to make transactions faster and more efficient than Bitcoin. This protocol uses [XRP Ledger \(XRPL\)](#) as its consensus protocol, in which designated servers (or validators) come to an agreement on the order and outcome of transactions made in Ripple network. All validators process transactions following the same rules and transactions are public, with strong cryptography to guarantee the integrity of the system. At the time of writing there are [156 validators](#) in the network but the consensus protocol ensures more decentralization over time as the validator pool grows. Unlike bitcoin which can only process 7 per second, Ripple can process more than 1500 transactions per second and its goal is to create a [SWIFT 2.0](#).

2.4 Smart Contracts

2.4.1 General Context

Smart contracts aim to automatize the execution of business processes, removing arbitrariness that humans can introduce. Typically every blockchain node have a virtual machine, which is a computation engine which acts like a decentralized computer that runs programs, smart contracts, just like Java Virtual Machine (JVM) does and maintain consensus across the network. This virtual machine is very import to support smart contracts as they interpret these programs on a lower level basis, often an assembly-like language or a bytecode.

There are many examples of these type of virtual machines being Ethereum Virtual Machine (EVM) the most famous one. This virtual machine executes as a stack machine with a depth of 1024 items, where each item is a 256-bit word, which was chosen for the ease of use with 256-bit cryptography, normally [Keccak-256](#). Smart contract's bytecode are then executed as a number of EVM operation codes (opcodes), which perform standard stack operations like XOR, AND, ADD, SUB, etc. or blockchain-specific stack operations, such as ADDRESS, BALANCE, BLOCKHASH, etc. On other hand Algorand Virtual Machine (AVM) interprets an assembly-like language, TEAL. TEAL programs are comprised of a set of opcodes, which are used to implement the logic of smart contracts and smart signatures. Examples of opcodes interpreted by AVM can be found [here](#).

Although bitcoin support smart contracts, through [BitML](#), they are less flexible than in newer architectures like Ethereum, Tezos or Algorand. In Bitcoin transactions can specify simple conditions on how to redeem them, using a limited set of logic, arithmetic and cryptographic operators. Despite the limited expressiveness of these conditions, it is possible to encode a variety of smart contracts (e.g. gambling games, escrow services, crowdfunding systems), by suitably chaining transactions.

The most popular programming languages are general-purpose programming languages, because they are used for many areas of research in computer science. If we look into [Tiobe](#), an index where ratings are based on the number of skilled engineers world-wide, courses and third party vendors, that uses popular search engines such as Google, Bing, Yahoo!, Wikipedia, Amazon, YouTube and Baidu to calculate the ratings, we can see that Python, C and Java are the most famous programming languages by these metrics. This suggests that these languages boast a thriving developer community with abundant online resources, facilitating newer developers with a lot of free and paid content to learn more about them. This can lead to an increase in popularity of a blockchain that chooses to use one of these languages for smart contract programming.

When you think of languages to write smart contracts you can either choose general-purpose programming languages or domain-specific programming languages. The former are designed to be used for building software in a wide variety of application domains. The latter are only designed to be used for building software in a specific application domain, in this case they are designed to write smart contracts. However digital contracts need to protect valuable assets and general-purpose programming languages don't provide any means to help developers making them more secure and safer. In 2018 Vitalik Buterin, co-founder of Ethereum blockchain, published a tweet where he shares an opinion that unlike regular coding smart contracts need to focus on 4 properties [1]:

- Very small code size - This makes smart contracts less prone to bugs and more intelligible.
- Much higher focus on safety - Smart contracts receive, store and transfers valuable user assets and they need to have safe operations.
- Much higher focus on auditability (misleading code very bad) - In smart contracts code dictates what a participant can do or cannot do ("CODE IS LAW").
- Perfect determinism - Code in smart contracts cannot have an unpredictable behaviour, like we have in C/C++ when for instance you try to access a position out of an array length, called [undefined behaviour](#).

For these reasons, many organisations managing blockchains have developed domain-specific programming languages. However these languages are still very new, many times suffering from several vulnerabilities and need to be developed and improved. To summarize and answer our initial question, **Domain-specific programming languages**

are **better to write smart contracts** but, as they are new programming languages, they are still susceptible to many vulnerabilities and need to be improved.

Even though Vitalik Buterin stated that general-purpose programming languages are worse than domain-specific languages to write smart contracts, there are a lot of blockchains which allow developers to write smart contracts in one of these languages.

Next I will present the most popular languages in the industry to write smart contracts, including some general-purpose ones.

2.4.2 Brief Presentation of Industry Languages

2.4.2.1 Solidity

Ethereum's official programming language to build and deploy smart contracts is [Solidity](#). This programming language is compatible with every blockchain protocol that runs on Ethereum Virtual Machine (EVM). [Polygon](#) and [Binance Smart Chain](#) are two examples of protocols that support EVM and consequently Solidity. Solidity is an object-oriented, high-level language for implementing smart contracts, created and designed by the Ethereum Foundation to target the EVM. It is a strongly typed and Turing-complete language, inspired by C++, Python and JavaScript, supporting inheritance, libraries and complex user-defined types (structs). It also is the most popular programming language to write smart contracts. This popularity combined with its expressiveness made the language a target for exploits, with several vulnerabilities and bugs discovered over the last years.

2.4.2.2 C++

C++ is a low level language, that delegates into developers memory management. it is supported and the native programming language in [EOSIO](#) blockchain to write smart contracts. Solana also supports C++ for smart contract development, but they recommend developers using Rust instead to have a better development experience. As we have already seen C++ is not a perfect deterministic programming language, many operations are prescribed to be unpredictable and as one of the lowest level languages, contracts written in this language are much more prone to have errors, bugs or vulnerabilities.

If you have some curiosity to see how smart contracts are written in C++, [here](#) is a repository with some examples of smart contracts for EOSIO blockchain.

```
1 void token::transfer( const name&    from,
2                     const name&    to,
3                     const asset&    quantity,
4                     const string&    memo )
5 {
6     check( from != to, "cannot transfer to self" );
7     require_auth( from );
8     check( is_account( to ), "to account does not exist");
9     auto sym = quantity.symbol.code();
10    stats statstable( get_self(), sym.raw() );
```

```

11     const auto& st = statstable.get( sym.raw() );
12
13     require_recipient( from );
14     require_recipient( to );
15
16     check( quantity.is_valid(), "invalid quantity" );
17     check( quantity.amount > 0, "must transfer positive quantity" );
18     check( quantity.symbol == st.supply.symbol, "symbol precision mismatch" );
19     check( memo.size() <= 256, "memo has more than 256 bytes" );
20
21     auto payer = has_auth( to ) ? to : from;
22
23     sub_balance( from, quantity );
24     add_balance( to, quantity, payer );
25 }

```

This is a function implementing a token transfer in C++ for EOSIO blockchain. As we can see, this approach brings a more confusing and harder to interpret code, making it less auditable. The **name** data type, represents an address and this data type wraps an `uint64_t` to ensure it is only passed to methods that expect a name.

2.4.2.3 Python

[Algorand](#) is the most popular blockchain project that uses Python as the primary language to write smart contracts. To write smart contracts for Algorand you will need to use [PyTeal](#) library that binds python code into [TEAL](#) code, an assembly like language used to express contracts in Algorand. This blockchain uses pure Proof-of-Stake (PPoS) consensus protocol. The main edge over Ethereum is the incredible transaction processing, it can process 1000 transactions per second. Another way to write smart contracts in Python is using [SmartPY](#), an intuitive platform to write smart contracts for [Tezos](#). However, Tezos uses Michelson as their native supported language for smart contract development.

When you look at PyTeal documentation, they note that **PyTeal hasn't been security audited**.

In the following [url](#) you can see some examples of smart contracts written using PyTeal for Algorand Virtual Machine.

```

1  alice = Addr("6ZHGH5Z5CTPCF5WCESXMGRSVK7QJETR63M3NY5FJCUYDH057VTCMJOBGY")
2  bob = Addr("7Z5PW02C6LFNQFGHWKSK5H47IQP50JW2M3HA2QPXTY3WTNP5NU2MHBW27M")
3  secret = Bytes("base32", "2323232323232323")
4  timeout = 3000
5
6  def htlc(
7      tpl_seller=alice,
8      tpl_buyer=bob,
9      tpl_fee=1000,
10     tpl_secret=secret,
11     tpl_hash_fn=Sha256,

```

```
12     tmpl_timeout=timeout,
13 ):
14
15     fee_cond = Txn.fee() < Int(tmpl_fee)
16     safety_cond = And(
17         Txn.type_enum() == TxnType.Payment,
18         Txn.close_remainder_to() == Global.zero_address(),
19         Txn.rekey_to() == Global.zero_address(),
20     )
21
22     recv_cond = And(Txn.receiver() == tmpl_seller, tmpl_hash_fn(Arg(0)) ==
23         tmpl_secret)
24
25     esc_cond = And(Txn.receiver() == tmpl_buyer, Txn.first_valid() >
26         Int(tmpl_timeout))
27
28     return And(fee_cond, safety_cond, Or(recv_cond, esc_cond))
```

Python brings a more readable approach, as its syntax is more concise than C++. However, it is a dynamic data type language, which makes it prone to errors and not type safe for this purpose. Addresses are treated with a special class **Addr** offered by this library.

2.4.2.4 Plutus

Cardano uses **Plutus** as the native language for smart contract development. Plutus is a Turing-complete language and smart contracts written in Plutus are effectively Haskell programs. Haskell is a statically-typed, purely functional programming language with type inference and lazy evaluation, very popular in academic world.

You can find examples of smart contracts written with Plutus for the Cardano blockchain in the following [repository](#).

```
1 (tid1, tid2) <- startExampleCampaign 700 700
2 tick 20
3 void $ addTx $ Tx
4     { _txInputs    = [Input (optr tid1 0) unit]
5       , _txOutputs = [Output "Bob" (fromAda 700) unit]
6       , _txSignees = ["Bob"]
7       , _txSlotRange = SlotRange 20 Forever
8       , _txForge    = mempty
9     }
10 void $ addTx $ Tx
11     { _txInputs    = [Input (optr tid2 0) unit]
12       , _txOutputs = [Output "Charlie" (fromAda 700) unit]
13       , _txSignees = ["Charlie"]
14       , _txSlotRange = SlotRange 20 Forever
15       , _txForge    = mempty
16     }
```

In the snippet above, these two procedures transfer 700 ADA from a certain address to another target address, using Plutus a subset of Haskell. This target address receives a total amount of 1400 ADA. As we can see Haskell brings simplicity and an easy to read code. However developers coming from Orient-Object programming languages or non-developers might struggle to understand what each line of code does. It is required that auditors have some previous Haskell knowledge, to have a more accurate audit.

2.4.2.5 Rust

Rust is a low-level statically-typed programming language that is fast and memory-efficient. Solana's smart contracts are fully written in Rust, while Polkadot has a framework built on top of Rust called [Substrate](#). [Near protocol](#) also uses Rust for the purpose of smart contract development.

[Here](#) is a tutorial teaching how to handle the creation and ownership of non-fungible tokens (NFTs) in the Substrate framework.

```
1 fn transfer_from_to(&mut self, from: AccountId, to: AccountId, value:
  Balance) -> bool {
2   let from_balance = self.balance_of_or_zero(&from);
3   if from_balance < value {
4     return false
5   }
6
7   // Update the sender's balance.
8   self.balances.insert(from, from_balance - value);
9
10  // Update the receiver's balance.
11  let to_balance = self.balance_of_or_zero(&to);
12  self.balances.insert(to, to_balance + value);
13
14  self.env().emit_event(Transfer {
15    from: Some(from),
16    to: Some(to),
17    value,
18  });
19
20  true
21 }
```

The snippet above is a function that transfers, in Polkadot network, some value of a specific ERC-20 token, from an account to another, using Substrate framework. As we can see Rust syntax is clear, easy to audit and supports either static or dynamic typing. Although it also has support for dynamic typing, Rust's compiler infers types by determining the type of a value based on the context in which it is used, making it a type safe language. Addresses in Substrate are handled with a specific data type created by Polkadot, named **AccountId**.

2.4.2.6 Go

Go is a statically typed, compiled programming language designed by Google very similar to C. This language offers many features like memory safety, garbage collection and structural typing. Hyperledger fabric's chaincodes can be written in Go.

If you want to have know how chaincodes written in Go look like, you can see some examples [here](#).

```
1 func (s *SmartContract) TransferAsset(ctx
    contractapi.TransactionContextInterface, id string, newOwner string) error
    {
2
3     asset, err := s.ReadAsset(ctx, id)
4     if err != nil {
5         return err
6     }
7
8     clientID, err := s.GetSubmittingClientIdentity(ctx)
9     if err != nil {
10        return err
11    }
12
13    if clientID != asset.Owner {
14        return fmt.Errorf("submitting client not authorized to update asset, does
            not own asset")
15    }
16
17    asset.Owner = newOwner
18    assetJSON, err := json.Marshal(asset)
19    if err != nil {
20        return err
21    }
22
23    return ctx.GetStub().PutState(id, assetJSON)
24 }
```

2.4.2.7 C#

NEO is a blockchain protocol fully built in C# and consequently it is also the platform native language to write smart contracts. However, this protocol supports many high level languages, providing compilers and plug-ins for these languages, which are used to compile high-level languages into instruction sets supported by NEO virtual machines. Apart from C# you can write smart contracts for NEO virtual machines in: Python, Java/Kotlin, Go and JavaScript.

[Here](#) are some examples of smart contracts written in C# for NEO blockchain.

```

1 private static bool Transfer(byte[] from, byte[] to, BigInteger amount,
   byte[] callscript)
2 {
3     //Check parameters
4     if (from.Length != 20 || to.Length != 20)
5         throw new InvalidOperationException("The parameters from and to
   SHOULD be 20-byte addresses.");
6     if (amount <= 0)
7         throw new InvalidOperationException("The parameter amount MUST be
   greater than 0.");
8     if (!Runtime.CheckWitness(from))
9         return false;
10    StorageMap asset = Storage.CurrentContext.CreateMap(nameof(asset));
11    var fromAmount = asset.Get(from).AsBigInteger();
12    if (fromAmount < amount)
13        return false;
14    if (from == to)
15        return true;
16
17    //Reduce payer balances
18    if (fromAmount == amount)
19        asset.Delete(from);
20    else
21        asset.Put(from, fromAmount - amount);
22
23    //Increase the payee balance
24    var toAmount = asset.Get(to).AsBigInteger();
25    asset.Put(to, toAmount + amount);
26
27    Transferred(from, to, amount);
28    return true;
29 }

```

2.4.2.8 Move

[Move](#) is a safe and flexible programming language for the [Diem Blockchain](#), originally developed by Facebook and tailored as a secure and verified programming language. It allows developers to define types to represent and control assets, but has no native support for contract protocols or gas usage. Move is designed to be a platform-agnostic language to enable common libraries, tooling, and developer communities across diverse blockchains with vastly different data and execution models. This language aims to be the "JavaScript of web3" in terms of ubiquity, when developers want to quickly write safe code involving assets, it should be written in Move.

In the following [repository](#), you can find an example of a contract implementing ERC-20 standard, that models the standard for fungible tokens, in Move.

2.4.2.9 Michelson

[Michelson](#) is an open-source function programming language designed for Tezos smart contracts development, introduced in 2014, when Tezos whitepaper [10] was released. The language is stack-based, with high level data types and primitives, strict static type checking and was designed to facilitate formal verification, allowing users to prove the properties of their contracts. It is inspired in not so popular languages like Forth, Scheme, ML and Cat.

2.4.2.10 TEAL

TEAL is an assembly-like language, processed by the AVM. The language is Turing-complete that supports looping and subroutines, but limits the amount of time the contract has to execute using a dynamic opcode cost evaluation algorithm. TEAL programs are processed one line at a time pushing and popping values on and off the stack. These stack values are either unsigned 64 bit integers or byte strings. TEAL provides a set of operators that operate on the values within the stack.

2.4.2.11 Marlowe

[Marlowe](#) is a programming language designed by Cardano Foundation for the eponymous blockchain. It is written as a Haskell data type, enabling users to create smart contracts in a low code platform called Blockly. This platform makes it easier for non-developers to build and deploy smart contracts. You can also write smart contracts in Haskell or JavaScript, which are more familiar to programmers and then convert into Marlowe code.

STATE OF THE ART

In this chapter, we offer an overview of smart contract languages developed through academic research. These languages are carefully crafted to address specific functionalities while emphasizing safety and security—a crucial consideration given the foundational role of smart contracts in decentralized applications. While we will not delve deeply into technical details, we will explore the key features and advantages of each language, shedding light on their contributions to secure and reliable smart contract execution. Familiarizing ourselves with these academic developments provides valuable insights into the broader landscape of smart contract development and the ongoing pursuit of safer blockchain technologies. Additionally, we will discuss well-known vulnerabilities in Solidity and the static analysis tools—both academic and industrial—that can be employed to detect them. Finally, we introduce the most important features of Featherweight Solidity.

3.1 Academic Languages

In the following table there is a summary of the different domain-specific languages to develop smart contracts.

	Solidity	Vyper	Scilla	Move	Michelson	Bamboo	Flint	Obsidian	TEAL
Turing Completeness	X				X				X
Type safe		X	X	X	X	X	X	X	X
Recursive calls	X				X				X
Function Modifiers	X								
Production ready	X				X				X
Assembly Language					X				X
Platform Agnostic			X	X				X	

3.1.1 Vyper

[Vyper](#) is a pythonic programming language for EVM which addresses resource usage and overflow checking. This programming language dropped some features that Solidity have in order to make contracts more secure and easier to audit. However it still suffers from

re-entrancy vulnerabilities and it is not intended to be a full replacement for everything that can be done in Solidity. It will deliberately forbid things or make things harder if it deems fit to do so for the goal of increasing security.

3.1.2 Scilla

[Scilla](#) is a language being developed for the [Zilliqa blockchain](#). It imposes a structure on smart contracts that makes applications less vulnerable to attacks by eliminating known vulnerabilities directly at the language-level. Zilliqa has been designed to be scalable, employing the idea of sharding¹ to validate transactions in parallel thus making transaction processing faster [22]. Although Scilla was designed in the context of Zilliqa, it is platform-agnostic and hence can be used with any other blockchain.

3.1.3 Bamboo

[Bamboo](#) is a programming language designed to write smart contracts for EVM. Bamboo makes state transition explicit and avoids reentrancy problems by default. However this language does not support features the following features:

- Loop constructs (for, while, ...) - Due to the constant block gas limit, loops should be avoided and each iteration should be done in separate transactions.
- Assignments into storage variables, except array elements - Instead of assigning a new value to a storage variable, the new value can be given as an argument in the continuation

[Here](#) are a lot of smart contract examples written in Bamboo.

3.1.4 Flint

[Flint](#) is a new type-safe contract-oriented programming language to write smart contracts for EVM, that has a variety of novel contract-oriented features, such as caller capabilities, immutability by default, and asset types. This programming language also provides safer operations, meaning that arithmetic operations on integers are safe by default: an overflow/underflow causes the Ethereum transaction to be reverted. However, Flint is still in alpha development, is an academic language developed in Imperial College London thus is still not ready to be used in production yet, lacking several important aspects like asset protection or gas consumption control.

3.1.5 Obsidian

[Obsidian](#) is another academic programming language that enables developing smart contracts for EVM. The authors of this project claim that Obsidian includes two innovations

¹Sharding is a process by which processing of the network can be made less intensive and allow for greater transaction volume[13].

relative to current languages for these platforms: state-oriented programming which lets developers declare and transition among states explicitly and linear types which ensure that important resources managed by your programs are managed correctly. Like Flint, this programming language is an academic language currently under development and is not production ready yet.

3.2 Known Vulnerabilities in Solidity

For the purpose of this dissertation, we have chosen to focus on studying Solidity, as it is currently the most popular language for writing smart contracts and has the most extensive resources available on the internet. Next, we will address known vulnerabilities in Solidity and provide examples of contracts that have the vulnerability and are not safe.

3.2.1 Type Casts

The Solidity compiler can detect some standard type errors (e.g., assigning an integer value to a variable of type string), as any strongly typed programming language. However, as a contract written in the Solidity language can call another contract by directly referencing the callee contract's instance, when a contract calls another contract's function, it only checks if the interface matches; it does not check in particular if the contract passed as an argument is from the same type as it is declared in a function. Therefore, a developer should be careful whenever a public function in a contract calls another contract interface.

```
1 contract CounterLibrary {
2     uint256 public counter;
3
4     function add() public returns (uint256) {
5         counter++;
6         return counter;
7     }
8 }
9
10 contract FakeCounter {
11     uint256 public counter;
12
13     function add() public returns (uint256) {
14         return counter;
15     }
16 }
17
18 contract Game {
19     uint256 public counter;
20
21     function play(CounterLibrary c1) public returns (uint256) {
22         uint256 res = c1.add();
23         return res;
24     }
25 }
```

Figure 3.1: An example of Type Casts vulnerability

Figure 3.1 shows an example of this vulnerability. Assuming these three contracts, if you deploy them into the blockchain, you can then call the play function from Game contract, which receives a CounterLibrary type contract. Nevertheless, you can either call this function with CounterLibrary contract address or FakeCounter contract address, even if only the first contract match the type declared, both match the interface required by play function. Obviously they will have different execution and output (CounterLibrary will return a true counter, but FakeCounter will always return 0).

There is still no way to prevent this vulnerability.

3.2.2 Reentrancy

This vulnerability was first seen in the DAO attack [7]. The DAO was a digital decentralized autonomous organization and a form of investor-directed venture capital fund. This idea was inspired by crowdfunding and was meant to work as a big decentralized [hedge fund](#). At the time, the DAO was the largest crowdfunding project ever and raised 150 million dollars, approximately 15% of Ethereum market cap. It was simply too big to fail and Ethereum Foundation decided controversially to hard fork [3] their own blockchain. However, some miners refused to fork because the DAO incident was not a defect in the protocol, arguing that "Code is Law" and formed Ethereum Classic, the original unaltered Ethereum blockchain. The vulnerability occurs when an external callee contract calls back to a function in the caller contract before the caller contract finishes, allowing an attack to

bypass the due validity check until the caller contract is drained of Ether or the transaction runs out of gas. Reentrancy can be prevented using one of the following methods:

- Assuring that a contract's state variables are updated before calling another contract.
- Introducing a mutex lock on the contract state to assure that only the lock owner can change the state(see [ReentrancyGuard](#) created by OpenZeppelin).
- Using the transfer or send (instead of call) methods to send money to other contracts, because this methods only forwards 2,300 gas to the callee contract.

```
1 function withdraw() public {
2     uint bal = balances[msg.sender];
3     require(bal > 0);
4
5     (bool sent, ) = msg.sender.call{value: bal}("");
6     require(sent, "Failed to send ether");
7     balances[msg.sender] = 0;
8 }
9
10 function nonReentrantWithdraw() public {
11     uint bal = balances[msg.sender];
12     require(bal > 0);
13
14     balances[msg.sender] = 0;
15     (bool sent, ) = msg.sender.call{value: bal}("");
16     require(sent, "Failed to send ether");
17 }
```

Figure 3.2: An example of a function vulnerable to reentrant attacks

3.2.3 Phishing waith tx.origin

In Solidity, tx.origin is a global variable that points to the initial externally owned account (EOA) responsible for initiating the current transaction. This vulnerability occurs when a contract uses tx.origin for authorization, which can be compromised by a phishing attack. A malicious contract can deceive the owner of a contract into calling a function that only the owner should be able to call.

This vulnerability can be prevented by using msg.sender, instead of tx.origin, for authentication, because if a contract A calls B, and B calls C, in C msg.sender will return the address of contract B and tx.origin will return the address of contract A.

```

1  contract Wallet {
2      address payable owner;
3
4      constructor() payable {
5          owner = msg.sender;
6      }
7
8      function transfer(address payable _to, uint _amount) public payable {
9          require(tx.origin == owner, "Not owner");
10         // require(msg.sender == owner, "Not owner"); correct way to
11         // authenticate!
12
13         (bool sent, ) = _to.call{value: amount}("");
14         require(sent, "Failed to send ether");
15     }
16 }
17 contract Attack {
18     address payable public owner;
19     Wallet wallet;
20
21     constructor(Wallet _wallet) {
22         wallet = Wallet(_wallet);
23         owner = payable(msg.sender);
24     }
25
26     function attack() public {
27         wallet.transfer(owner, address(wallet).balance);
28     }
29 }

```

Figure 3.3: An example of a vulnerable contract to phishing

3.2.4 Delegatecall Injection

Delegatecall is a solidity function to facilitate code-reuse inserting a callee contract's bytecode into the bytecode of the caller contract: it simply let's a contract call another contract's external function. However using delegatecall is tricky to use and wrong usage or incorrect understanding can lead to devastating results. As a consequence, a malicious callee contract can directly modify (or manipulate) the state variables of the caller contract.

```

1  function delegateCallToAnotherContract(address _contract) public {
2      (bool success, bytes memory data) =
3          _contract.delegatecall(abi.encodeWithSignature("setVars(uint256)"))
4  }

```

Figure 3.4: An example of a dangerous delegatecall usage

3.2.5 Unprotected Selfdestruct

```

1  contract EtherGame {
2      uint public targetAmount = 7 ether;
3      address public winner;
4
5      function deposit() public payable {
6          require(msg.value == 1 ether, "You can only send 1 ether");
7          uint balance = address(this).balance;
8          require(balance <= targetAmount, "Game is over");
9
10         if(balance == targetAmount) {
11             winner = msg.sender;
12         }
13     }
14     function claimReward public {
15         require(msg.sender == winner, "Not winner");
16         (bool sent, ) = msg.sender.call{value: address(this).balance}("");
17         require(sent, "Failed to send ether");
18     }
19 }
20
21 contract Attack {
22     EtherGame etherGame;
23
24     constructor(EtherGame _etherGame) {
25         etherGame = EtherGame(_etherGame);
26     }
27
28     public attack() public payable {
29         //You can simply break the game forcing sending ether so that
30         // the game balance >= 7
31         address payable addr = payable(address(etherGame));
32         selfdestruct(etherGame);
33     }
34 }

```

Figure 3.5: An example of a malicious use of selfdestruct

Selfdestruct is a built-in Solidity function that destroys the current contract, sending its funds to the given address and end execution. This function has some peculiarities inherited from the EVM:

- The receiving contract's receive function is not executed.
- The contract is only really destroyed at the end of the transaction and revert s might "undo" the destruction.

This vulnerability was first seen in Parity wallet [26]. The Parity wallet was designed to integrate seamlessly with all standard tokens as well as manage Ether transfers. However, there was a vulnerability with this contract reported on 6th November 2017 by an anonymous user. The user decided to exploit this vulnerability and made himself the "owner" of

the library contract. Subsequently, the user could access a special function which only the owner of the Parity wallet should have been able to access. This function allowed a user to destroy this component affecting all subsequent contracts that relied on this contract. This action blocked funds in 587 wallets holding a total amount of 513,774.16 Ether as well as additional tokens. To prevent this vulnerability a developer must be careful when implementing a function that call the built-in selfdestruct function, making sure that an arbitrary user can never get access to it.

An attacker can also create a contract with a selfdestruct() function, send ether to it, call selfdestruct(target) and force ether to be sent to a target. This vulnerability arises from the misuse of the globally available variable address(this).balance. For this reason your contract should never rely on this variable, as it can be artificially manipulated.

3.2.6 Gasless Send/Transfer

```
1 contract C {
2
3     constructor() payable {
4
5     }
6
7     function pay(uint n, D1 d) public payable {
8         address d1 = address(d);
9         payable(d1).transfer(n);
10    }
11    function balance() public view returns (uint) {
12        return address(this).balance;
13    }
14 }
15
16 contract D1 {
17     uint public counter = 0;
18
19     //this fallback function costs more than 2300 gas
20     //so the transfer/send method will fail, use call instead
21     fallback() external payable {
22         counter++;
23     }
24
25     function balance() public view returns (uint) {
26         return address(this).balance;
27     }
28 }
```

Figure 3.6: An example of a transfer that will fail due to not having enough gas

When using the functions send() or transfer() to transfer ether to another contract, it is possible to incur in an out-of-gas exception, as these functions only forward 2300 gas to the callee contract. This may be quite unexpected by programmers, because transferring ether is not generally associated to executing code. The reason behind this exception is

subtle. Note that `c.send(amount)` or `c.transfer(amount)` are compiled in the same way of a call with empty signature, but remember that the actual number of gas units available to the callee is always bound by 2300, to prevent reentrancy attacks. Now, since the call has no signature, it will invoke the callee's receive or fallback function. However, 2300 units of gas only allow to execute a limited set of bytecode instructions, e.g. those which do not alter the state of the contract. In any other case, the call will end up in an out-of-gas exception.

Summing up, sending ether via `send` or `transfer` functions succeed in two cases: when the recipient is a contract with an inexpensive fallback, or when the recipient is a user. If a developer wants to execute complex computations and change a contract state, after a contract receives ether, he should use the `call` function to transfer ether since this function does not impose a forward gas limit when sending funds to the callee contract.

3.2.7 Unprotected ether withdrawal

```
1
2 function owner() public view virtual returns (address) {
3     // returns owner
4     ....
5 }
6
7 modifier onlyOwner() {
8     require(owner() == msg.sender, "Not owner");
9 }
10
11 function protectedWithdraw(uint amount) public onlyOwner {
12     require(amount > 0);
13     require(amount <= balance, "No funds to withdraw the amount requested.");
14
15     balance -= amount;
16     payable(msg.sender).transfer(amount);
17 }
18
19 function unprotectedWithdraw(uint amount) public {
20     require(amount > 0);
21     require(amount <= balance, "No funds to withdraw the amount requested.");
22
23     balance -= amount;
24     payable(msg.sender).transfer(amount);
25 }
```

Figure 3.7: An example of a protected and an unprotected ether withdraw function

This vulnerability occurs when bad actors without adequate access controls, are able to withdraw some or all ether from a contract. To avoid this vulnerability, only allow withdrawals to be triggered by those authorized, or as intended, normally only the owner of the contract should be able to withdraw ether. This vulnerability can be avoided extending `Ownable` class offered by OpenZeppelin.

3.3 Tools To Verify Solidity Smart Contracts

In our pursuit of enhancing the security of Solidity smart contracts, it is essential to understand how different static analysis tools can help us mitigate vulnerabilities. Before we delve into the specifics of each tool, let's take a moment to explore a comprehensive summary table. This table will provide a quick reference, showcasing which vulnerabilities each tool excels at identifying. Understanding these strengths will help us make informed choices when selecting the most suitable tool for our particular needs.

	Slither	Mythril	SMTChecker	Solc-verify	Solidifier
Type Casts					
Reentrancy	X	X		X	
Phishing with tx.origin	X	X	X		X
Delegatecall Injection	X	X			
Unprotected Selfdestruct	X	X			
Gasless send/transfer					
Unprotected Ether Withdrawal	X	X*	X	X*	X*

3.3.1 Slither

[Slither](#) is a Solidity static analysis framework that checks for specific vulnerabilities. Moreover it also lets developers extend this tool through an API that enables one to write custom analyses. At the time of its launch, it was the first open-source static analysis framework for Solidity. The main features provided by this tool are:

- Very fast and precise, averaging less than 1 second of execution time per contract and detects vulnerable Solidity code with low false positives.
- Identifies where the error condition occurs in the source code.
- Easily integrates into continuous integration and Truffle builds.
- Built-in 'printers' quickly report crucial contract information.
- Detector API to write custom analyses in Python.
- Correctly parses 99.9% of all public Solidity code
- Ability to analyze contracts written with Solidity ≥ 0.4 and so compatible with the most recent Solidity versions.
- Intermediate representation (SlithIR) enables simple, high-precision analyses.

This tool was the best performer in both execution speed time and accuracy in all these 7 vulnerabilities. It worked very well to detect 5 of those, only missing Type Casts and Gasless send/transfer vulnerabilities.

3.3.2 Mythril

Mythril is a security analysis tool for EVM bytecode. It offers a command that checks a contract or all the contracts inside a directory, similarly to `slither`. However I could note that this tool is less precise and slower than `slither`, as this tool checks Solidity compiled code (binaries), giving more false-positives and more complex to analyze. It is also not possible to customize this tool, because it does not offer any API to do so, like `slither` does. It is also worth noting that this tool could not be accurate when checking a contract protected against Unprotected Ether Withdrawal when using modifiers, like "onlyOwner" implemented by OpenZeppelin. It only stopped giving this vulnerability if we explicit put a pre-condition "**require(owner == msg.sender)**" before executing a contract code. This tool was not as good as `Slither` when treating reentrancy vulnerability as it warned reentrancy when we transferred ether using `call()` method, even if the function was protected against reentrancy with `ReentrancyGuard` or using good development practices.

3.3.3 SMTChecker

SMTChecker is a Solidity implementation of a formal verification approach based on Satisfiability Modulo Theories (SMT) and Horn solving. This module automatically tries to prove that the code satisfies the specification given by `require` (pre-conditions) and `assert` (post-conditions) statements. That is, it considers `require` statements as assumptions and tries to prove that the conditions in `assert` statements are always true. If an assertion fails, a counterexample may be given to the user showing how the assertion can be violated. If no warning is given by the **SMTChecker** for a property, it means that the property is safe. This module also targets other verifications that are checked at compile time:

- Arithmetic underflow and overflow.
- Division by zero.
- Trivial conditions and unreachable code.
- Popping an empty array.
- Out of bounds index access.
- Insufficient funds for a transfer.

All the targets above are automatically checked by default if all engines are enabled, except underflow and overflow for Solidity versions equal or greater than 0.8.7.

I have to note that I talked with one of the authors and did not get any response whether it is possible or not, to detect Reentrancy vulnerability with this tool. The way I tried was not working, as the tool was stuck in a infinite loop, producing no output.

3.3.4 Solidifier

Solidifier is a bounded model checker for Solidity. This tool translates Solidity code into the [Boogie language](#) which is later verified by [Corral](#). Unlike many of the tools developed to analyse smart contracts that check for known possibly problematic/vulnerable behaviours, Solidifier looks for violation to assertions. The semantic properties described in this way should better describe the intent of developers and so they more finely capture mistakes and possibly vulnerabilities of smart contracts. This tool works very similar like VeriSol and Solc-verify. Both try to falsify/prove semantic properties of smart contracts.

This tool requires a contract to declare the following library in order to get access to verification primitives:

```
1 library Verification {
2     function Assume (bool b);
3     function Assert (bool b);
4     function CexPrintui (uint ui);
5 }
```

Verification.Assume(condition) primitive fails silently (stopping the contract's execution) when the condition is not met and **Verification.Assert(condition)** fails with an error that is reported by Solidifier. Both this functions are identical to SMTChecker primitives **require** and **assert**. The Verification library can also declare functions to print the value of expressions of basic types to the counterexample trace. These are prefixed by **CexPrint**. In this example, CexPrintui (uint ui) declares a function that prints a unsigned integer. On other hand we could have declared a function CexPrintad (address ad) that print addresses. It is important noting that printing only happens if the counterexample reaches a given printing statement.

Using this tool we were only able to detect Phishing with tx.origin vulnerability. This tool does not support function modifiers, delegatecall or external contract calls, thus making impossible to detect other vulnerabilities. However it was possible to kinda detect and prevent Unprotected Ether Withdrawal using Assume function as a pre-condition to ensure that only the owner of the contract is able to withdraw ether from a contract and using Assert to prove that pre-condition.

3.3.5 Solc-verify

Solc-verify [11] is an extended version of the Solidity compiler that is able to perform automated formal verification on Solidity smart contracts using specification annotations and modular program verification. They must be side-effect free Solidity expressions and can refer to variables within the scope of the annotated element. This tool allows access to loop and contract invariants, function pre and post conditions, functional modifiers, specific functions referring to a sum of a collection or old variable values and some more that you can find [here](#).

However this tool [does not support built-in Solidity functions or state variables](#) like `tx.origin`, `delegatecall` and `selfdestruct` making it impossible to detect vulnerabilities involving these. It also has false positives for reentrancy when using `send()` and `transfer()` functions which are indeed reentrant proof.

3.3.6 Other Tools

In this section I will talk about tools that I was not able to try because they either are discontinued or I could not install and the authors did not help me with my installation errors.

3.3.6.1 VeriSmart

[VeriSmart](#) is a safety analyzer for Solidity smart contracts written in OCaml. This tool apparently is still under development, since the last commit on GitHub was in January 2023. It supports two types of analysis modes:

- **Safety verification** - In this mode, VeriSmart can be used to prove the absence of bugs or detect bugs. The key feature in this mode is, VeriSmart automatically infers transaction invariants, which are conditions that hold under arbitrary interleaving of transactions, of smart contracts. With this ability, VeriSmart can precisely analyze safety properties in smart contracts.
- **Vulnerable transaction sequence generation** - In this mode, VeriSmart can be used to generate vulnerable transaction sequences (with concrete arguments for each transaction) that demonstrate the flaws. They claim that the key feature in this mode is, they guide a symbolic execution procedure with language models to find vulnerable sequences effectively.

Unfortunately I was not able to experiment this tool because when running a specific contract I got an error and even though [I opened an issue on GitHub](#), the owners did not reply me in time.

3.3.6.2 SmartAce

[SmartACE](#) is a tool which extends the Solidity compiler to support bounded and parameterized smart contract verification. This tool is built upon a novel model for smart contracts, in which users are processes and communication is explicit. In this model, communication is over-approximated by static analysis, and the results are sufficient to find all local neighbourhoods. This is achieved by using [SeaHorn](#), an automated analysis framework for LLVM-based languages. It takes as input a smart contract annotated with assertions, then it checks that all assertions hold.

I did not have time to explore more this tool and had some problems running it. It is worth noting that this tool is only compatible with contracts written in Solidity version 0.5.17 and we need to port them back if we want to test contracts from version above.

3.3.6.3 VeriSol

[VeriSol](#) is a Microsoft Research project for prototyping a formal verification and analysis system for smart contracts developed in the popular Solidity programming language. It is based on translating programs in Solidity language to programs in Boogie intermediate verification language, and then leveraging and extending the verification toolchain for Boogie programs. It is not compatible with Solidity versions 0.8 authors are not actively maintaining the tool since 2021.

3.3.6.4 SmartCheck

[SmartCheck](#) is an extensible static analysis tool for discovering vulnerabilities and other code issues in Ethereum smart contracts that are written in the Solidity programming language. However this tool is deprecated since 2020 and I could not run it properly on my machine.

FEATHERWEIGHT SOLIDITY

4.1 General Context

As previously stated, our main focus during this master thesis was studying, and redefining Featherweight Solidity and then implementing it in OCaml. Featherweight Solidity is a minimal core calculus for Solidity, inspired by Featherweight Java [12], focused on the core of Solidity and dropping complex features, to study its type system. Therefore, it enables the study of various aspects of smart contract programming, such as new contracts deployment, interaction among deployed contracts, and money transfers.

It includes both the semantics and the type system of Solidity, providing a comprehensive understanding of the language's behavior and its potential vulnerabilities. It also allows a new subtype for the type address, making this new redefining system able to detect which contract an address refers to.

The formalization of Solidity in FS was a significant step towards understanding the language's features and vulnerabilities, marking the first formalization of Solidity. While FS has its limitations and some aspects of Solidity have not been fully modeled yet, it provides a foundation for further research into new language features and contract vulnerabilities. Our implementation keeps all of essential feature embedded in the original formalization, but we extended it to meet some of the most important feature from Solidity yet to be modeled: multiple inheritance, flexibility on constructor bodies, default values assigned to variables and another adjustments to be in-line with the most recent Solidity version to date (0.8.*).

Next, we will present the original formalization.

4.2 Meta Variables

Let's begin by introducing some meta variables that will be used in the subsequent rules:

- C : The set of contract names, and C, D identifies each element in the set.
- S : The set of state variable identifiers, and s identifies each element.

- \mathcal{F} : The set of function names, and f identifies each element.
- \mathcal{X} : The set of argument names, and x, y identifies each element.
- \mathcal{V} : The set of values, and v identifies each element.
- \mathcal{T} : The set containing all possible types of our language, and T identifies each element.

it is important to note that we consider null as a concrete value for a contract name. Therefore, every reference to c represents a value within an infinite set that also includes null.

Additionally, the variable f represents all possible values for a function identifier, including the unknown. However, this special value cannot be used in a function defined by the user.

4.3 Grammar

(Contract decl.)	SC	$::=$	<code>contract C is $D\{\widetilde{T} s\} K \widetilde{F}$</code>
(Constructor decl.)	K	$::=$	<code>constructor $((\widetilde{T}_1 y, \widetilde{T}_2 x))\{\text{super}(\widetilde{y}); \text{this}.\widetilde{s} = \widetilde{x}\}$</code>
(Function decl.)	F	$::=$	<code>$T f \langle C \rangle ((\widetilde{T} x))\{\text{return } e\}$ <code>unit $fb \langle C \rangle () \{\text{return } e\}$</code></code>
(Expression)	e	$::=$	<code>$v x b \text{this}.s \text{this}.f(\widetilde{e})$ <code>msg.sender msg.value </code> <code>address(e) $e.s$ $e.transfer(e)$ </code> <code>new $C.value(e)(\widetilde{e})$ <code>$e; e T x = e x = e e.s = e$ <code>$e[e] e[e \rightarrow e] e.value(e)(\widetilde{e})$ <code>$e.f.value(e)(\widetilde{e}) \text{revert}$ <code>$e.f.value(e).sender(e)(\widetilde{e})$ <code>if e then e else e</code></code></code></code></code></code></code>
(Values)	v	$::=$	<code>true false n a u M </code> c
(Types)	T	$::=$	<code>$\widetilde{T} \rightarrow T \text{bool} \text{uint} \text{address}\langle C \rangle$ <code>unit mapping($T \Rightarrow T$) C</code></code>

Figure 4.1: Syntax of Featherweight Solidity+

This is the grammar for the original extended Featherweight Solidity formalization. Please note that it only allows for simple inheritance. Although mutual recursion is permitted according to this grammar definition, we assume non-mutually recursive contract definitions. It is important to highlight that their grammar enforces the requirement for every contract to inherit from another. **However, we disagree with this choice because we**

believe it should be an axiom that every contract always inherits from a top-level class or contract and thus there should be the possibility for a contract to not inherit from any class but the top class.

Additionally, there are four important properties of subtyping that we'd like to emphasize:

$$\begin{array}{c}
 \text{(TOP)} \\
 \hline
 C <: \text{Top}
 \end{array}
 \qquad
 \begin{array}{c}
 \text{(REFLEXIVITY)} \\
 \hline
 C <: C
 \end{array}$$

$$\begin{array}{c}
 \text{(CONTRACT)} \\
 \text{contract } C \text{ is } D\{\widetilde{T}s\} K \widetilde{F}\} \\
 \hline
 C <: D
 \end{array}
 \qquad
 \begin{array}{c}
 \text{(TRANSITIVITY)} \\
 C <: D \quad D <: E \\
 \hline
 C <: E
 \end{array}$$

These rules also apply to subtyping in address subtyping. In short, for $\text{address}\langle C \rangle <: \text{address}\langle D \rangle$ to be true, C must be a subtype of D , that is $C <: D$. This is known as co-variance and its rule is listed below:

$$\begin{array}{c}
 \text{(ADDRESS-COVARIANCE)} \\
 C <: D \\
 \hline
 \text{address}\langle C \rangle <: \text{address}\langle D \rangle
 \end{array}$$

We can also see that constructor declarations are more strict than Solidity, enforcing one to initialize all state variables.

Functions are identical to Solidity, except in FS+ we can restrict what type `msg.sender` can be. This states the required maximum supertype allowed for the implicit variable `msg.sender`. It is validated at compile-time, supposing `msg.sender` as an expression of type $\text{address}\langle C \rangle$ and checking the subtyping for each call of f . **We also disagree with making it a mandatory annotation because this approach is not retro-compatible with Solidity. By default, it should expect all types of addresses.**

4.3.1 Contract Definition

We write $\text{contract } C \text{ is } D\{\widetilde{T}s\} K \widetilde{F}\}$ to denote a smart contract declaration. It indicates that contract C depends on a contract D , a sequence possibly empty of typed variables, denoted by $\widetilde{T}s$, a mandatory constructor declaration and a set of functions (\widetilde{F}), also can be empty. However if a contract does not have at least a fallback function (`unit fb(){return e}`) it will not be able to receive ether.

4.3.2 Constructor Declaration

In this version, it is imperative that all state variables of the contract be meticulously initialized. The set of typed parameters, denoted as \widetilde{T}_1y , must align in length with the

state variables of the super contract. Similarly, the set \widetilde{T}_2x must conform to the contract's own variables, ensuring a precise match.

4.3.3 Expression

This represents our set of possible expressions. Our keyword **this** allow to referring to the current contract state variables and functions, unlike Solidity. In Solidity, it is only used to reference function, we do not need to use the keyword to access contract's state variables.

4.3.4 Values

As values we accept boolean values (true or false) and an unsigned integer, denoted by n . There is also a possibility to write a specific address, if one uses the constructor `address(0)` it will refer to the special empty address (default value for all addresses). The rest of the values are only allowed as a return type of a function or being the value of specific keywords, for example the keyword **this**.

4.3.5 Types

We have all types matching each possible value for our expressions, plus a type for typing functions $\widetilde{T} \rightarrow T$.

4.4 Configuration

$$\begin{array}{ll}
 \text{(Blockchain)} & \beta ::= \emptyset \mid \beta \cdot [x \mapsto v] \mid \\
 & \beta \cdot [(c, a) \mapsto (C, \widetilde{s} : \widetilde{v}, n)] \mid \\
 \text{(Call Stack)} & \sigma ::= \beta \mid \sigma \cdot a \\
 \text{(Contract Table)} & CT ::= \emptyset \mid CT \cdot [C \mapsto D]
 \end{array}$$

Figure 4.2: Environments of Featherweight Solidity+

The call stack σ is responsible for tracking the nesting of function calls, specifically the addresses of their enclosing contracts, during the execution of a transaction. To simplify the operational semantics during transaction execution, they mandate that the initial element of the call stack must represent a copy of the blockchain as it existed at the commencement of the top-level transaction.

This initial blockchain copy, situated at the base of the call stack, serves a critical purpose in the event of an abort, where the blockchain state must be reverted to the beginning of the transaction. Conversely, if the transaction successfully commits, the modified blockchain will be duplicated to replace the copy held within σ .

Their configuration differs slightly from ours. As previously mentioned, they store a copy of the blockchain within the call stack, denoting their configuration as $\langle \beta, \sigma, e \rangle$, with its initial value set as $\langle \beta_0, \beta_0, e \rangle$.

On the other hand, our approach utilizes a quadruple configuration, and we maintain a separate entry for the copy of the blockchain, represented as $\langle \beta, \beta_0, \sigma, e \rangle$.

In summary, a successfully evaluated transaction concludes with a configuration in the form of $\langle \beta, \beta_0, v \rangle$, leading to the commitment of the updated blockchain β . Conversely, when an error occurs, the configuration is represented as $\langle \beta, \beta_0, \text{revert} \rangle$. In this scenario, the modified blockchain β is discarded, returning to the initial state β_0 and effectively simulating the rollback of the aborted transaction.

4.5 Lookup Functions

These lookup functions are used in operational semantics and type system rules and are formally defined in this section.

The first function accepts a contract as input and retrieves the contract's state variables. The second function takes a contract, a function name, and a set of values, returning a tuple containing the function's parameters and its body. The third function, given a contract and a function name, provides the type associated with the function's parameters and body

$$\text{sv}(\text{Top}) = \emptyset$$

$$\frac{\text{(STATE VARIABLE LOOKUP)} \quad \text{contract } C \text{ is } D\{\widetilde{T} s\} K \widetilde{F}\} \quad \text{sv}(D) = \widetilde{T}_1 r}{\text{sv}(C) = \widetilde{T} s; \widetilde{T}_1 r}$$

$$\text{fbody}(\text{Top}, fb, \widetilde{v}) = (\{\}, \text{revert})$$

$$\frac{\text{(FUNCTION BODY LOOKUP - 1)} \quad \text{CT}(C) = \text{contract } C \text{ is } D\{\widetilde{T} s\} K \widetilde{F}\} \quad T f \langle C' \rangle ((\widetilde{T} x))\{\text{return } e\} \in \widetilde{F} \quad \widetilde{x} = \widetilde{v}}{\text{fbody}(C, f, \widetilde{v}) = (\widetilde{x}, \text{return } e)}$$

$$\frac{\text{(FUNCTION BODY LOOKUP - 2)} \quad \text{CT}(C) = \text{contract } C \text{ is } D\{\widetilde{T} s\} K \widetilde{F}\} \quad T f \langle C' \rangle ((\widetilde{T} x))\{\text{return } e\} \notin \widetilde{F} \vee \widetilde{x} \neq \widetilde{v}}{\text{fbody}(C, f, \widetilde{v}) = \text{fbody}(SC, f, \widetilde{v})}$$

$$\frac{\text{(FUNCTION SIGNATURE LOOKUP - 1)} \\ CT(C) = \text{contract } C \text{ is } D\{\widetilde{T} s\} K \widetilde{F} \quad T f \langle C' \rangle ((\widetilde{T} x))\{\text{return } e\} \in \widetilde{F}}{\text{ftype}(C, f) = \widetilde{T} \rightarrow T}$$

$$\frac{\text{(FUNCTION SIGNATURE LOOKUP - 2)} \\ CT(C) = \text{contract } C \text{ is } D\{\widetilde{T} s\} K \widetilde{F} \quad T f \langle C' \rangle ((\widetilde{T} x))\{\text{return } e\} \notin \widetilde{F}}{\text{ftype}(C, f) = \text{ftype}(SC, f)}$$

Lastly, we introduce `fsender` function. It retrieves the type of `msg.sender` inside a function.

$$\frac{\text{(SENDER MAXIMUM TYPE LOOKUP - 1)} \\ CT(C) = \text{contract } C \text{ is } D\{\widetilde{T} s\} K \widetilde{F} \quad T f \langle C' \rangle ((\widetilde{T} x))\{\text{return } e\} \in \widetilde{F}}{\text{fsender}(C, f) = C'}$$

$$\frac{\text{(SENDER MAXIMUM TYPE LOOKUP - 2)} \\ CT(C) = \text{contract } C \text{ is } D\{\widetilde{T} s\} K \widetilde{F} \quad T f \langle C' \rangle ((\widetilde{T} x))\{\text{return } e\} \notin \widetilde{F}}{\text{fsender}(C, f) = \text{fsender}(SC, f)}$$

Note that all these functions are recursive. If the contract does not provide either the function of state variable we go up on the hierarchy and search for it again, until we reach the Top contract and return the base case value for each one.

4.6 Auxiliary Predicates

$$\text{uptbal} : \text{BLOCKCHAIN} \times \text{ADDRESS} \times \text{BALANCE} \rightarrow \text{BLOCKCHAIN}$$

$$\text{uptbal}(\beta, a, n) = \begin{cases} \beta[(c, a) \mapsto (C, \widetilde{s} : \widetilde{v}, n' + n)] & \beta((c, a)) = (C, \widetilde{s} : \widetilde{v}, n') \wedge n' + n \geq 0 \\ \perp & \text{otherwise} \end{cases}$$

We have simplified the original definition, making it equivalent. This predicate succeeds if and only if there exists a valid tuple (c, a) in β , and the sum of $n' + n$ is greater than or equal to 0. it is crucial to emphasize that n can be positive, negative, or zero.

$$\text{Top}(\sigma) = \begin{cases} a & \text{if } \sigma = \sigma' \cdot a \\ \emptyset & \text{if } \sigma = \beta \end{cases}$$

This predicate simply returns the top of the call stack, if there is one or \emptyset , if the call stack is empty.

4.7 Operational Semantics

4.7.1 Conditional and Sequential Statements

(REVERT)

$$\frac{}{\langle \beta, \beta_i, \sigma, E[\text{revert}] \rangle \rightarrow \langle \beta, \beta_i, \sigma, \text{revert} \rangle}$$

Every expression can eventually be evaluated to a revert. This rule is for cases that our expression violates a pre-condition and can't be evaluated by our rules, it should fallback to this rule. In the end, if we do not have explicitly made a rule for a failure, by omission, that rule should always be treated as a revert.

(IF-TRUE)

$$\frac{}{\langle \beta, \sigma, \text{if true then } e1 \text{ else } e2 \rangle \rightarrow \langle \beta, \sigma, e1 \rangle}$$

(IF-FALSE)

$$\frac{}{\langle \beta, \sigma, \text{if false then } e1 \text{ else } e2 \rangle \rightarrow \langle \beta, \sigma, e2 \rangle}$$

These rules have the standard meaning. Important to note that if we want to omit either branch the omitted branch should take the value unit internally.

(SEQ-C)

$$\frac{\sigma = \beta_i}{\langle \beta, \beta_i, \sigma, v; e \rangle \rightarrow \langle \beta, \beta, e \rangle}$$

(SEQ-R)

$$\frac{\sigma = \beta_i}{\langle \beta, \sigma, \text{revert}; e \rangle \rightarrow \langle \beta_i, \sigma, \text{revert} \rangle}$$

(SEQ)

$$\frac{\text{Top}(\sigma) = a}{\langle \beta, \sigma, v; e \rangle \rightarrow \langle \beta, \sigma, e \rangle}$$

These rules have to take into account different scenarios. First, when an expression like $v; e$ is found at the top level ($\text{Top}(\sigma) = \emptyset$) then the last transaction was successful. In fact, it evaluated to a value v , which can be discarded to proceed with the next transaction e . Furthermore, since evaluating to a value is equivalent to a commit, we apply the changes and proceed with β . This scenario is modeled by SEQ-C. Secondly, $v; e$ might be at an inner

level ($\text{Top}(\sigma) = a$). In this case (SEQ) we cannot apply any changes, since the transaction is not over yet. Hence, we just proceed with the evaluation of e , discarding v . Lastly, a revert may have been thrown during the evaluation, so instead of v ; e the expression involving sequential composition is $\text{revert}; e$ (SEQ-R). In this case we do not further proceed in evaluating the other transactions and we replace β with β_i .

4.7.2 Variable Operations

(VAR)

$$\frac{}{\langle \beta, \sigma, x \rangle \rightarrow \langle \beta, \sigma, \beta(x) \rangle}$$

Accesses a variable and returns its value.

(ASSIGN)

$$\frac{x \in \text{dom}(\beta)}{\langle \beta, \sigma, \sigma, x = v \rangle \rightarrow \langle \beta \cdot [x \mapsto v], \sigma, v \rangle}$$

Assigns a value to an existing variable.

(LET)

$$\frac{x \notin \text{dom}(\beta)}{\langle \beta, \sigma, T x = v; e \rangle \rightarrow \langle \beta \cdot [x \mapsto v], \sigma, v; e \rangle}$$

When a variable x is declared, it is added to the β . This behavior is represented and modeled by the LET function. VAR and ASSIGN functions, on the other hand, solely access β to either read or modify the value associated with a variable x .

The rules governing the reading and mutation of state variables are the following:

(STATEREAD)

$$\frac{\exists a. \beta((c, a)) = (C, \tilde{s}v, n) \quad s \in \tilde{s}}{\langle \beta, \sigma, c.s \rangle \rightarrow \langle \beta, \sigma, v_s \rangle}$$

(STATEASSIGN)

$$\frac{\exists a. \beta((c, a)) = (C, \tilde{s}v, n) \quad s \in \tilde{s}}{\langle \beta, \sigma, c.s = v' \rangle \rightarrow \langle \beta \cdot [c.s \mapsto v'], \sigma, v' \rangle}$$

When provided with a contract reference and a state variable identifier, STATEREAD returns the value of the specified variable. Similarly, STATEASSIGN facilitates the modification of this value.

4.7.3 Mappings

(MAPREAD)

$$\frac{}{\langle \beta, \sigma, M[v1] \rangle \rightarrow \langle \beta, \sigma, M(v1) \rangle}$$

(MAPASSIGN)

$$M' = M \setminus \{(v1, M(v1))\} \cup \{(v1, v2)\}$$

$$\frac{}{\langle \beta, \sigma, M[v1 \mapsto v2] \rangle \rightarrow \langle \beta, \sigma, M' \rangle}$$

Although Di Pirro claims that their MAPREAD rule is a total function, we disagree because if $v1$ has not been previously assigned a value, there is no insight on how to handle the situation.

4.7.4 Contract Instantiation

(NEW-1)

$$\frac{(c, a) \notin \text{Dom}(\beta) \quad \text{sv}(C) = \widetilde{T}s \quad |\widetilde{v}| = |\widetilde{s}| \quad \text{Top}(\sigma) \neq \emptyset}{\langle \beta, \sigma, \text{new } C.\text{value}(n)(\widetilde{v}) \rangle \rightarrow \langle \text{uptbal}(\beta, \text{Top}(\sigma), -n).[(c, a) \mapsto (C, \widetilde{s} : \widetilde{v}, n)], \sigma, c \rangle}$$

(NEW-2)

$$\frac{(c, a) \notin \text{Dom}(\beta) \quad \text{sv}(C) = \widetilde{T}s \quad |\widetilde{v}| = |\widetilde{s}| \quad \text{Top}(\sigma) = \emptyset}{\langle \beta, \sigma, \text{new } C.\text{value}(n)(\widetilde{v}) \rangle \rightarrow \langle \beta.[(c, a) \mapsto (C, \widetilde{s} : \widetilde{v}, n)], \sigma, c \rangle}$$

(NEW-R)

$$\frac{\text{uptbal}(\beta, \text{Top}(\sigma), -n) = \perp \quad \text{Top}(\sigma) \neq \emptyset}{\langle \beta, \sigma, \text{new } C.\text{value}(n)(\widetilde{v}) \rangle \rightarrow \langle \beta, \sigma, \text{revert} \rangle}$$

NEW rules deploy a new contract on the blockchain. To ensure consistency with the definition of K , where the body consists solely of a sequence of assignments, the number of typed values provided as input must match the number of state variables in contract C . This differs slightly from Solidity's capabilities, as it restricts explicit references to `msg.sender` or `msg.value` and initializing state variables with different types. In FS+, a strict one-to-one correspondence between parameters and state variables is required.

The NEW-2 rule increases liquidity by allowing external sources to provide Wei without checking balances. In Ethereum, the balance of the contract posing as a sender is decreased instead. This rule also enables the creation of externally-owned accounts on their implementation.

4.7.5 Balance and Address

$$\begin{array}{c}
 \text{(BALANCE)} \\
 \beta(a) = (C, \tilde{sv}, n) \\
 \hline
 \langle \beta, \sigma, \text{balance}(a) \rangle \rightarrow \langle \beta, \sigma, n \rangle
 \end{array}$$

$$\begin{array}{c}
 \text{(ADDRESS)} \\
 \beta(c) = a \\
 \hline
 \langle \beta, \sigma, \text{address}(c) \rangle \rightarrow \langle \beta, \sigma, a \rangle
 \end{array}$$

In Solidity, balance is a property of address values that returns the balance of the account referred to by each address. Here they claim that to avoid confusion with contract state variables, they chose to define balance as an expression of the language. Furthermore, in Solidity, an implicit cast is applied whenever a contract reference is used instead of an address value. To model such behavior in FS, which does not have any implicit casts, they defined an explicit global function called address that extracts the address corresponding to a given reference.

4.7.6 Cast

$$\begin{array}{c}
 \text{(CONTRRETR)} \\
 \exists c. \beta(c, a) = C \quad C \leq: D \\
 \hline
 \langle \beta, \sigma, D(a) \rangle \rightarrow \langle \beta, \sigma, c \rangle
 \end{array}$$

$$\begin{array}{c}
 \text{(CONTRRETR-R)} \\
 \exists c. \beta(c, a) = C \quad C \not\leq: D \\
 \hline
 \langle \beta, \sigma, D(a) \rangle \rightarrow \langle \beta, \sigma, \text{revert} \rangle
 \end{array}$$

These rules apply to the operation of casting an address to a contract. If there is a contract c with the associated address a on the blockchain, we check if the contract type C is equal to or inherits from D at some point.

These rules are applied to the operation which we can cast an address to a contract. If there is in the blockchain a contract c with the associated address a we simply check if the contract type C is equal or inherits from D at some point.

4.7.7 Money Transfer

$$\begin{array}{c}
 \text{(TRANSFER)} \\
 \beta(a) = (C, sv, n) \\
 \text{fbody}(C, fb, \{\}) = (\{\}, e) \quad \beta' = \text{uptbal}(\text{uptbal}(\beta, \text{Top}(\sigma), -n), a, n) \\
 \hline
 \langle \beta, \sigma, a.\text{transfer}(n) \rangle \rightarrow \langle \beta', \sigma * a, e\{\text{this} := c, \text{msg.sender} := \text{Top}(\sigma), \text{msg.value} := n\} \rangle
 \end{array}$$

$$\begin{array}{c}
 \text{(TRANSFER-R)} \\
 \text{uptbal}(\beta, \text{Top}(\sigma), -n) = \perp \\
 \hline
 \langle \beta, \sigma, a.\text{transfer}(n) \rangle \rightarrow \langle \beta, \sigma, \text{revert} \rangle
 \end{array}$$

These rules model the transfer function in Solidity as it is. If contract executing this function has sufficient balance to transfer n to a , then it is executed successfully and the contract with associated with address a executes its fallback function. Otherwise the transaction is reverted.

4.7.8 Function Calls

$$\begin{array}{c}
 \text{(CALL)} \\
 \exists a.\beta((c, a)) = (C, \tilde{sv}, n) \quad \text{fbody}(C, f, \tilde{v}) = (\tilde{x}, e) \\
 \tilde{x} \notin \text{Dom}(\beta) \quad \beta' = \text{uptbal}(\text{uptbal}(\beta, \text{Top}(\sigma), -n), a, n) \cdot [\tilde{x} \mapsto \tilde{v}] \\
 e_s = e\{\text{this} := c, \text{msg.sender} := \text{Top}(\sigma), \text{msg.value} := n\} \\
 \hline
 \langle \beta, \sigma, c.f.\text{value}(n)(\tilde{v}) \rangle \rightarrow \langle \beta', \sigma \cdot a, e_s \rangle
 \end{array}$$

$$\begin{array}{c}
 \text{(CALL-R)} \\
 \text{uptbal}(\beta, \text{Top}(\sigma), -n) = \perp \\
 \hline
 \langle \beta, \sigma, c.f.\text{value}(n)(\tilde{v}) \rangle \rightarrow \langle \beta, \sigma, \text{revert} \rangle
 \end{array}$$

$$\begin{array}{c}
 \text{(CALLTOPLEVEL)} \\
 \exists a.\beta((c, a)) = (C, \tilde{sv}, n) \quad \text{fbody}(C, f, \tilde{v}) = (\tilde{x}, e) \\
 \tilde{x} \notin \text{Dom}(\beta) \quad \beta' = \text{uptbal}(\text{uptbal}(\beta_w, a', -n), a, n) \cdot [\tilde{x} \mapsto \tilde{v}] \\
 e_s = e\{\text{this} := c, \text{msg.sender} := a', \text{msg.value} := n\} \quad \text{Top}(\sigma) = \emptyset \\
 \hline
 \langle \beta, \sigma, c.f.\text{value}(n).\text{sender}(a')(\tilde{v}); e' \rangle \rightarrow \langle \beta', \sigma \cdot a, e_s; e' \rangle
 \end{array}$$

$$\begin{array}{c}
 \text{(CALLTOPLEVEL-R)} \\
 \text{uptbal}(\beta, a', -n) = \perp \quad \text{Top}(\sigma) = \emptyset \\
 \hline
 \langle \beta, \sigma, c.f.\text{value}(n).\text{sender}(a')(\tilde{v}); e' \rangle \rightarrow \langle \beta', \sigma \cdot a, \text{revert}; e' \rangle
 \end{array}$$

In our interpretation and comparison with Solidity's state of the art, CALLTOPLEVEL is used to initiate a transaction that is always triggered by an externally-owned account, while CALL is used when a contract wants to invoke a function of another contract. Both are quite similar, and you can pass a number n posing as the value of Wei that the other party will receive. However, if the initiator does not have enough funds to cover the amount, the transaction will fail.

(RETURN)

$$\frac{}{\langle \beta, \sigma \cdot a, \text{return } v \rangle \rightarrow \langle \beta, \sigma, v \rangle}$$

(RETURN-R)

$$\frac{}{\langle \beta, \sigma \cdot a, \text{return revert} \rangle \rightarrow \langle \beta, \sigma, \text{revert} \rangle}$$

4.8 Type System

4.8.1 Axioms

(REF)

$$\frac{\Gamma, c : C, \Gamma' \vdash \langle \rangle}{\Gamma, c : C, \Gamma' \vdash c : C}$$

(VAR)

$$\frac{\Gamma, x : T, \Gamma' \vdash \langle \rangle}{\Gamma, x : T, \Gamma' \vdash x : T}$$

(TRUE)

$$\frac{\Gamma \vdash \langle \rangle}{\Gamma \vdash \text{true} : \text{bool}}$$

(FALSE)

$$\frac{\Gamma \vdash \langle \rangle}{\Gamma \vdash \text{false} : \text{bool}}$$

(UNIT)

$$\frac{\Gamma \vdash \langle \rangle}{\Gamma \vdash u : \text{unit}}$$

(REVERT)

$$\frac{\Gamma \vdash \langle \rangle}{\Gamma \vdash \text{revert} : T}$$

$$\begin{array}{c}
 \text{(UINT)} \\
 n \in \mathbb{N}_0 \quad \Gamma \vdash \langle \rangle \\
 \hline
 \Gamma \vdash n : \text{uint}
 \end{array}$$

$$\begin{array}{c}
 \text{(ADDRESS)} \\
 \Gamma, a : \text{address}\langle C \rangle, \Gamma' \vdash \langle \rangle \\
 \hline
 \Gamma, a : \text{address}\langle C \rangle, \Gamma' \vdash a : \text{address}\langle C \rangle
 \end{array}$$

The axioms have all a standard meaning. Here it is important to note that $\Gamma \vdash \text{revert} : T$ is well-typed regardless of T . We do not give a unique type to this expression, since it may be used in any context whatsoever. For instance, $\Gamma \vdash \text{if true then } 10 \text{ else revert} : \text{uint}$ and $\Gamma \vdash \text{if true then } u \text{ else revert} : \text{unit}$ are both valid expressions in FS, but the former has type `uint` and the latter has type `unit`. If we want to omit the `else` branch then its value should always be `u`.

4.8.2 Standard Rules

These rules have a standard meaning are listed below. They are composed by the expression, variable assignment and declaration as well as sequential composition and special operations `balance` and `address`.

$$\begin{array}{c}
 \text{(BALANCE)} \\
 \Gamma \vdash e : \text{address} \\
 \hline
 \Gamma \vdash \text{balance}(e) : \text{uint}
 \end{array}$$

$$\begin{array}{c}
 \text{(ADDR)} \\
 \Gamma \vdash e : C \\
 \hline
 \Gamma \vdash \text{address}(e) : \text{address}
 \end{array}$$

$$\begin{array}{c}
 \text{(RETURN)} \\
 \Gamma \vdash e : T \\
 \hline
 \Gamma \vdash \text{return } e : T
 \end{array}$$

$$\begin{array}{c}
 \text{(LET)} \\
 \Gamma \vdash e_1 : T'_1 \quad T'_1 <: T_1 \quad \Gamma, x : T_1 \vdash e_2 : T_2 \\
 \hline
 \Gamma \vdash T_1 \ x = e_1; e_2 : T_2
 \end{array}$$

(ASSIGN)

$$\frac{\Gamma \vdash x : T \quad \Gamma \vdash e : T_e \quad T_e <: T}{\Gamma \vdash x = e : T_e}$$

(IF)

$$\frac{\Gamma \vdash e_1 : \text{bool} \quad \Gamma \vdash e_2 : T_1 \quad \Gamma \vdash e_3 : T_2 \quad T_1 <: T \quad T_2 <: T}{\Gamma \vdash \text{if } e_1 \text{ then } e_2 \text{ else } e_3 : T}$$

(SEQ)

$$\frac{\Gamma \vdash e_1 : T_1 \quad \Gamma \vdash e_2 : T_2}{\Gamma \vdash e_1; e_2 : T_2}$$

4.8.3 Mappings

The following rules provide a type to mapping values, as well as to read and write accesses on them. To read these rules it is important to mention that they consider $\Gamma \vdash \tilde{v} : T$, represents $n \in \mathbb{N}$ distinct judgments where the type T is fixed and does not vary: $\Gamma \vdash v_i : T, 1 \leq i \leq n$

(MAPPING)

$$\frac{M = \{\widetilde{(k, v)}\} \quad \Gamma \vdash \tilde{k} : T_k \quad \Gamma \vdash \tilde{v} : T_v \quad T_k <: T_1 \quad T_v <: T_2}{\Gamma \vdash M : \text{mapping}(T_1 \Rightarrow T_2)}$$

(MAPASSIGN)

$$\frac{\Gamma \vdash e1 : \text{mapping}(T_1 \rightarrow T_2) \quad \Gamma \vdash e2 : T'_1 \quad \Gamma \vdash e3 : T'_2 \quad T'_1 <: T_1 \quad T'_2 <: T_2}{\Gamma \vdash e1[e2 \rightarrow e3] : \text{mapping}(T_1 \Rightarrow T_2)}$$

(MAPREAD)

$$\frac{\Gamma \vdash e1 : \text{mapping}(T_1 \rightarrow T_2) \quad \Gamma \vdash e2 : T'_1 \quad T'_1 <: T_1}{\Gamma \vdash e1[e2] : T_2}$$

4.8.4 Contract Instantiation and Access

The following rules are about the deploy of a new contract as well as the access (read/write) to its state variables.

$$\frac{\text{(NEW)} \quad \text{sv}(C) = \widetilde{T}s \quad \Gamma \vdash \widetilde{e} : \widetilde{T}' \quad \widetilde{T}' <: \widetilde{T} \quad |\widetilde{e}| = |\widetilde{s}| \quad \Gamma \vdash e' : \text{uint}}{\Gamma \vdash \text{new } C.\text{value}(e')(\widetilde{e}) : C}$$

$$\frac{\text{(STATEASSIGN)} \quad \Gamma \vdash e_1.s : T \quad \Gamma \vdash e_2 : T' \quad T' <: T}{\Gamma \vdash e_1.s = e_2 : T}$$

$$\frac{\text{(STATEREAD)} \quad \Gamma \vdash e : C \quad \text{sv}(C) = \widetilde{T}s \quad s_i \in s}{\Gamma \vdash e.s_i : T_i}$$

4.8.5 Casts and Money Transfers

$$\frac{\text{(CONTRRETR)} \quad \Gamma \vdash e : \text{address}\langle C \rangle \quad C <: D}{\Gamma \vdash D(e) : D}$$

$$\frac{\text{(TRANSFER)} \quad \Gamma \vdash e_1 : \text{address}\langle C \rangle \quad \text{ftype}(C, \text{fb}) = \{\} \rightarrow \text{unit} \quad \Gamma \vdash e_2 : \text{uint} \quad \Gamma \vdash \text{this} : C' \quad C' <: \text{fsender}(C, \text{fb})}{\Gamma \vdash e_1.\text{transfer}(e_2) : \text{unit}}$$

Starting by the first rule, they claim we now we have precise information about address e and the operation $D(e)$ is well-typed if and only if C is either equals D or a subclass of it. In other words, we can cast an address pointing to a subcontract of D to a reference of type D

In the second rule, we can notice that it is mandatory for the contract invoking the transfer function to inherit from a contract defining a fb function that has to be annotated with a contract C that is a supertype of our current contract. We disagree with this implementation because it is mandatory on our language to have a fb function without any annotation, i.e it can be called by every address in our environment including externally-owned accounts.

4.8.6 Functions

$$\frac{\text{(CALLTOPLEVEL)} \quad \Gamma \vdash e_1 : C \quad \Gamma \vdash e_2 : \text{uint} \quad \text{ftype}(C, f) = \widetilde{T}_1 \rightarrow T_2 \quad |\widetilde{e}| = |\widetilde{T}_1| \quad \Gamma \vdash e_3 : \text{address}\langle C' \rangle \quad C' <: \text{fsender}(C, f)}{\Gamma \vdash e_1.f.\text{value}(e_2).\text{sender}(e_3)(\widetilde{e}) : T_2}$$

(CALL)

$$\frac{\begin{array}{l} \Gamma \vdash \mathbf{e}_1 : \mathbf{C} \quad \Gamma \vdash e_2 : \mathbf{uint} \\ \text{ftype}(C, f) = \tilde{T}_1 \rightarrow T_2 \quad |\tilde{e}| = |\tilde{T}_1| \quad \Gamma \vdash \mathbf{this} : C' \quad C' <: \text{fsender}(C, f) \end{array}}{\Gamma \vdash e_1.f.\text{value}(e_2)(\tilde{e}) : T_2}$$

(FUN)

$$\frac{\Gamma \vdash \mathbf{e}_1 : \mathbf{C} \quad \text{ftype}(C, f) = \tilde{T}_1 \rightarrow T_2 \quad \Gamma \vdash \mathbf{this} : C' \quad C' <: \text{fsender}(C, f)}{\Gamma \vdash c.f : \tilde{T}_1 \rightarrow T_2}$$

(CALLVALUE)

$$\frac{\Gamma \vdash \tilde{T}_1 \rightarrow T_2 \quad \Gamma \vdash e_2 : \mathbf{uint} \quad \Gamma \vdash \tilde{e}_1 : \tilde{T}'_1 \quad \tilde{T}'_1 <: T_1 \quad |\tilde{e}| = \tilde{T}_1}{\Gamma \vdash e_1.\text{value}(e_2)(\tilde{e}) : T_2}$$

In these rules, it is worth mentioning that since every function has a restricted `msg.sender` type value, we need to ensure that the contract calling the function f is equal to or a subclass of the value we obtain from the annotation. However, we disagree with restricting it on `CALLTOPLEVEL`. Our argument is that this constructor is only intended to be called by externally-owned accounts (EOAs), i.e., users initiating a transaction and not by contracts. Therefore, restricting it is pointless.

FS 2.0: REVISED OPERATIONAL SEMANTICS

In this chapter, we introduce the revised operational semantics for the programming language, building upon the original formalization by Di Pirro and Silvia Crafa. Key modifications include: integration of multiple inheritance, clarification of ambiguous definitions identified during our study and align more closely with the actual behavior of the current Solidity version (0.8.*).

In our notation going forward, we will employ \tilde{X} to denote a sequence of values, and $\tilde{\tilde{X}}$ to signify a set comprising sequences of values, encapsulating the concept of a sequence of sequences.

5.1 Grammar

(Contract decl.)	SC	$::=$	<code>contract C is \tilde{D} {(\tilde{T} s) K \tilde{F}}</code>
(Constructor decl.)	K	$::=$	<code>constructor $\tilde{C}(\tilde{e})$ ((\tilde{T} x))(e)</code>
(Function decl.)	F	$::=$	<code>T f ((\tilde{T} x)){return e} unit fb() {return e}</code>
(Expression)	e	$::=$	<code>v x b this.s this.f(\tilde{e}) msg.sender msg.value address(e) e.s e.transfer(e) new C.value(e)((\tilde{e})) e;e T x = e x = e e.s = e e[e] e[e → e] e.value(e)(\tilde{e}) e.f.value(e)(\tilde{e}) revert e.f.value(e).sender(e)(\tilde{e}) if e then e else e</code>
(Values)	v	$::=$	<code>true false n a u M c</code>
(Types)	T	$::=$	<code>$\tilde{T} \rightarrow T$ bool uint address unit mapping($T \Rightarrow T$) C</code>

Figure 5.1: Syntax of Featherweight Solidity 2.0

Like the original formalization, the grammar of Featherweight Solidity 2.0 only contains a subset of Solidity. Looping primitives, arrays, structs, modifiers and events are removed from our syntax to simplify the formal system and associated proofs and remove extra complexity from a proof-of-concept. Nevertheless, conditional expressions are included, we allow multiple inheritance and one can also simulate loops using recursion. We also do not take into account the different types of visibility (represented in Solidity with external, internal, public, and private) and we removed pure and view from function declaration. Lastly, we consider every function as payable, and we do not model gas. Next we explain in detail the user syntax.

5.1.1 Contract Definition

We write contract C is $\widetilde{D} \{(\widetilde{T} s) K \widetilde{F}\}$ to denote a smart contract declaration. It indicates that contract C contains a sequence, possibly empty, $D_1 \dots D_n$, representing a possible contract extension, another sequence possibly empty of typed variables, denoted by $\widetilde{T} s$, a mandatory constructor declaration and a set of functions (\widetilde{F}), also can be empty. However if a contract does not have at least a fallback function (`unit fb(){return e}`) it will not be able to receive ether.

5.1.2 Constructor Declaration

Our implementation doesn't enforce to initialize all state variables here, bringing more flexibility when declaring constructors. Thereafter we are much more liberal here and treat constructor essentially as a special function, just like it currently works in Solidity. Note that if a contract has super contracts inheriting from, their constructors should be also called here in place of $\widetilde{C}(\overline{e})$.

5.2 Configuration

$$\begin{array}{lll}
 (\textit{Blockchain}) & \beta & ::= \emptyset \mid \beta \cdot [x \mapsto v] \mid \\
 & & \beta \cdot [(c, a) \mapsto (C, \widetilde{s} : \overline{v}, n)] \mid \\
 & & \beta \cdot [a \mapsto n] \\
 (\textit{Call Stack}) & \sigma & ::= \varepsilon \mid \sigma \cdot a \\
 (\textit{Contract Table}) & CT & ::= \emptyset \mid CT \cdot [C \mapsto SC]
 \end{array}$$

Figure 5.2: Environments of Featherweight Solidity 2.0

Here we make two changes to Di Pirro's original formalization. One of those being, instead of having a call stack that contains an initial copy of the blockchain and a stack of addresses, we separate these two concepts and our call stack becomes just the stack of addresses. In our definition whenever we refer to β_i in the rules, we are referring to a copy

of the blockchain before the first call. Another change we make is having a blockchain composed by three productions. The first two are identical, a map between variables and their values and a mapping between a key comprised with contract reference and an address to a tuple with contract information (name, state variables and balance). The third formation is a mapping between addresses and balances, representing externally-owned accounts.

5.3 Contract Example in FS 2.0

Before introducing our rules, we would like to introduce a smart contract example that we will use to some of the rules we will be presenting. This contract will also be used to exemplify type system rules and in chapter 8.

```

1 import "./NFTStorage.sol";
2 contract Game is Context{
3     mapping(address => NFTStorage) stores;
4     constructor() Context(){}
5     function fb() {}
6     function createStore() returns (address) {
7         NFTStorage store = new NFTStorage{value: 0}();
8         this.stores = (this.stores[address(store)] = store);
9         return address(store);
10    }
11    function addExternalStore(address<NFTStorage> store) {
12        this.stores = (this.stores[store] = NFTStorage(store));
13    }
14    function setNFTPrice(address<NFTStorage> store, uint price) {
15        NFTStorage storeInstance = this.stores[store];
16        storeInstance.setNFTPrice{value: 0}(price);
17    }
18    function createNFT(address<NFTStorage> store, address to) {
19        NFTStorage storeInstance = this.stores[store];
20        storeInstance.createNFT{value: 0}(to);
21    }
22    function buyNFT(address<NFTStorage> store) {
23        NFTStorage storeInstance = this.stores[store];
24        storeInstance.buyNFT{value: msgvalue}();
25    }
26    function transferNFT(address<NFTStorage> store, uint tokenId, address
27        from, address to) {
28        address sndr = this.msgSender();
29        NFTStorage storeInstance = this.stores[store];
30        storeInstance.transferNFT{value: 0}(tokenId, sndr, to);
31    }
32    function destroyNFT(address<NFTStorage> store, uint tokenId) {
33        address sndr = this.msgSender();
34        NFTStorage storeInstance = this.stores[store];
35        storeInstance.destroyNFT{value: 0}(sndr, tokenId);
36    }

```

Figure 5.3: A smart contract written in our Solidity refined grammar.

This is a simple contract that serves as a proxy to interact with a set of NFT stores, being a high-level API to interact them. Each externally-owned account can invoke this transactions, to create a new store, add an existing external store, set NFT prices for a store, create NFTs and so one.

We can see that our syntax is very similar with the one Solidity has, except we allow declaring a subset of type address, as we can see in function signatures. With this refined syntax, we can detect before executing each function, if the address passed is indeed

pointing to the `NFTStorage` contract specification. Thus passing an address belonging to an externally-owned account, will result in an error, before executing the function.

5.4 Auxiliary Predicates

During this chapter, we will be referencing some auxiliary methods and predicates, to simplify our rules presentation. In this section we will be providing and explaining each one.

5.4.1 Default Values

$$\begin{aligned} & \text{defaultvalue} : \mathbf{T} \rightarrow \mathbf{V} \\ \text{defaultvalue}(T) = & \begin{cases} 0 & T = \text{uint} \\ \text{false} & T = \text{bool} \\ \mathbf{0x0} & T = \text{address} \\ \text{Null} & T = \mathbf{C} \\ [] & T = \text{mapping}(T \Rightarrow T) \\ \text{unknown} & T = \widetilde{T} \rightarrow T \end{cases} \\ & \text{init} : \mathbf{T} \times \mathbf{S} \rightarrow \mathbf{S} \times \mathbf{V} \end{aligned}$$

$$\text{init}(Ts) = [s \mapsto \text{defaultvalue}(T)]$$

Our formalization does not enforce one to initialize state variables. Like Solidity if the user does not initialize it through constructor, we will be assuming a default value for each one. This auxiliary function is intended to assign default values to the state variables of a contract when it is initialized, ensuring that they always have values even when none are assigned in the constructor's body.

The function $\text{init}(\widetilde{Ts})$ is a sequence of functions $\text{init}(Ts)$, where \widetilde{Ts} represents a list of tuples with the type and the name of a certain state variable of a contract and Ts represents each tuple. Moreover $\widetilde{SC} \mapsto \widetilde{e}$ denotes a mapping like $[C1 \mapsto \widetilde{e1}; \dots; Cn \mapsto \widetilde{en}]$

For the concrete contract example, this is applied to the state variable `stores`. This variable is not initialized in the constructor with a value and so will fallback to the value defined by this function, which is simply the empty map.

5.4.2 Multiple Inheritance

The implementation of multiple inheritance in our system adheres to a set of fundamental rules. As previously mentioned, we have adopted the C3 linearization algorithm as the cornerstone of our approach to determine the hierarchy order for each contract and effectively resolve the diamond problem.

The diamond problem is an ambiguity that arises when two classes B and C inherit from A, and class D inherits from both B and C.




Figure 5.4: The diamond problem

As illustrated in the figure above, the graph have two distinct topological sort orders: [D;C;B;A] and [D;B;C;A]. In the context of programming languages and software design, it is imperative to establish consistency in choosing a specific order to resolve such ambiguities.

To incorporate C3 Linearization, as an algorithm to solve contract's dependencies we adapted an existing generic implementation to suit our project [0].

$$A : C \times \tilde{A}$$

$$\text{remove} : \tilde{A} \times C \rightarrow \tilde{A}$$

(REMOVE)

$$\tilde{L} = [L_1; \dots; L_n]$$

$$SC = \text{contract } C \text{ is } \tilde{D} \{(\tilde{T} s); K \tilde{F}\} \quad \tilde{L}_1 = [L_1 \setminus \{SC\}; \dots; L_n \setminus \{SC\}]$$

$$\text{remove}(\tilde{L}, SC) = \tilde{L}_1$$

$$\text{merge} : \tilde{A} \rightarrow \tilde{C}$$

$$\text{Let } SC = \text{contract } C \text{ is } \tilde{D} \{(\tilde{T} s); K \tilde{F}\}$$

$$\text{merge}(\tilde{L}) = \begin{cases} SC \cdot (\text{merge}(\text{remove}(\tilde{L}, SC))) & \text{if } SC \in L_k, L_k \in \tilde{L}, SC = \text{head}(L_k), SC \notin \text{tail}(\tilde{L}) \\ \text{fail} & \text{otherwise} \end{cases}$$

$$\text{c3linearization} : SC \rightarrow \tilde{C}$$

$$\text{Let } SC = \text{contract } C \text{ is } \tilde{D} \{(\tilde{T} s); K \tilde{F}\}$$

$$\text{c3linearization}(SC) = \begin{cases} [C] & \tilde{D} = \emptyset \\ [C] \cdot \text{merge}(\text{c3linearization}(\tilde{D})) & \text{otherwise} \end{cases}$$

The c3linearization(SC) function computes a contract's hierarchy based on its super-class relationships. The functions head and tail are elementary operations employed

to extract information from a set. Specifically, `head` retrieves the first element of the set, while `tail` gathers all elements except the first one.

In Chapter 8, we provide an illustrative example of its execution with a set of contracts.

5.4.3 Functions

This section presents the auxiliary methods used to decide which functions our contract will execute, based on its hierarchy.

$$\frac{\text{name} : \mathcal{F} \rightarrow f \quad (\text{GET FUNCTION NAME}) \quad F = T f ((\widetilde{T} x))\{e\}}{\text{name}(F) = f}$$

This is a trivial function to retrieve a function name. For our illustrative contract possible return values are: `fb`, `createStore`, `addExternalStore`, `setNFTPrice`, `createNFT`, `buyNFT`, `transferNFT` and `destroyNFT`.

$$\frac{\text{functions} : C \rightarrow \widetilde{\mathcal{F}} \quad (\text{RETRIEVE CONTRACT FUNCTIONS MAP}) \quad CT(C) = \text{contract } C \text{ is } \widetilde{D} [D_1 \mapsto \widetilde{e}_1; \dots; D_n \mapsto \widetilde{e}_n]\{(\widetilde{T} s); K F_1; \dots; F_n\} \quad n \geq 0}{\text{functions}(C) = [\text{name}(F_1) \mapsto F_1; \dots; \text{name}(F_n) \mapsto F_n]}$$

This function computes a map for each contract, where the key is the name of the function and the value is its declaration. In the illustrative example we would get the following map:

$$[\text{fb} \mapsto \text{unit}; \text{createStore} \mapsto e_1; \dots; \text{destroyNFT} \mapsto e_n]$$

where e_1 and e_n represent the expressions of each function body.

Now for this example suppose that the Game contract inherits from a contract Context that follows this specification:

```

1  contract Context{
2      uint balance;
3      constructor() Context() {}
4      function fb() {}
5      function createStore() returns (address) {
6          ....
7      }
8      function newFunction() {
9          ....
10     }
11 }
```

Figure 5.5: A dummy example for Context contract.

$$\text{cmap} : (\tilde{\mathcal{F}} \times \tilde{\mathcal{F}})^2 \rightarrow \tilde{\mathcal{F}} \times \tilde{\mathcal{F}}$$

(CONCAT FUNCTIONS MAP)

$$\frac{M = [f_{name_1} \mapsto f_1; \dots; f_{name_n} \mapsto f_n] \quad fm = [fn_1 \mapsto fm_1; \dots; fn_{n_1} \mapsto fm_{n_1}] \quad n \geq 0 \quad n_1 \geq 0 \quad r = M \cup (fm \setminus \{\forall n_1 \in fm : fn_{n_1} \in M\})}{\text{cmap}(M, fm) = r}$$

This method concatenates a function to a map M if and only if there is not a function with the same name (key). This is useful when calculating our table lookup of each contract, that we will represent next.

$$\text{tlookup}(\tilde{C}_s, \text{tbl}) : (\tilde{S}\tilde{C}, [f \mapsto F]) \rightarrow [f \mapsto F] = \begin{cases} \text{tbl} & \tilde{C}_s = [] \\ \text{tlookup}(\tilde{C}_s', \text{cmap}(\text{tbl}, \text{functions}(C_i))) & \tilde{C}_s = C_i \cdot \tilde{C}_s' \end{cases}$$

This function computes the table lookup for a contract. It is simply a recursive function using the previous cmap function and it starts from the head of the computed linearization of classes. If we assume Context as illustrated in 5.5 and we try to apply this function to our Game contract we get the following result:

$$\text{tlookup}([\text{Game}; \text{Context}], []) =$$

$$[\text{fb} \mapsto \text{unit}; \dots; \text{destroyNFT} \mapsto e_n; \text{newFunction} \mapsto e_{n2}]$$

Note that the only function we concatenate from Context contract is **newFunction**, because it does not exist in the previous iterated contracts.

(RETRIEVE FUNCTION BODY)

$$\frac{\text{tlookup}(\tilde{C}_s, []) [f] = T f ((\tilde{T} x)) \{e\} \quad \tilde{C}_s = \text{c3linearization}(CT(C)) \quad |\tilde{v}| = |\tilde{x}|}{\text{fbody}(C, \tilde{v}, f) = (\tilde{x}, e)}$$

(RETRIEVE FUNCTION BODY)

$$\frac{\tilde{C}_s = \text{c3linearization}(CT(C)) \quad \text{tlookup}(\tilde{C}_s, []) [f] = \perp \vee (\text{tlookup}(\tilde{C}_s, []) [f] = T f ((\tilde{T} x)) \{e\} \wedge |\tilde{v}| \neq |\tilde{x}|)}{\text{fbody}(C, \tilde{v}, f) = (\{\}, \text{return revert})}$$

This function retrieves a function body and its arguments from a function table. It receives a contract, a set of values and a function name. Imagine that we apply this function to Game contract, with $f = \text{addExternalStore}$, $\tilde{v} = [\text{address}(0)]$ we get the following tuple as return: $([\text{store}], e_b)$ where e_b represents the body of our function.

$$\begin{array}{c}
 \text{(RETRIEVE FUNCTION SIGNATURE)} \\
 \hline
 CT(C) = \text{contract } C \text{ is } \widetilde{D} [\widetilde{SC} \mapsto \widetilde{e}] \{(\widetilde{T} s); K \widetilde{F}\} \quad B f ((\widetilde{A} x))\{e\} \in \widetilde{F} \\
 \hline
 \text{ftype}(C, f) = \widetilde{A} \rightarrow B
 \end{array}$$

$$\begin{array}{c}
 \text{(SENDER MAXIMUM TYPE LOOKUP - 1)} \\
 \hline
 CT(C) = \text{contract } C \text{ is } \widetilde{D}\{(\widetilde{T} s) K \widetilde{F}\} \quad T f \langle C' \rangle ((\widetilde{T} x))\{\text{return } e\} \in \widetilde{F} \\
 \hline
 \text{fsender}(C, f) = C'
 \end{array}$$

$$\begin{array}{c}
 \text{(SENDER MAXIMUM TYPE LOOKUP - 2)} \\
 \hline
 CT(C) = \text{contract } C \text{ is } \widetilde{D}\{(\widetilde{T} s) K \widetilde{F}\} \quad T f ((\widetilde{T} x))\{\text{return } e\} \in \widetilde{F} \\
 \hline
 \text{fsender}(C, f) = \emptyset
 \end{array}$$

$$\begin{array}{c}
 \text{(SENDER MAXIMUM TYPE LOOKUP - 3)} \\
 \hline
 CT(C) = \text{contract } C \text{ is } \widetilde{D}\{(\widetilde{T} s) K \widetilde{F}\} \quad T f \langle C' \rangle ((\widetilde{T} x))\{\text{return } e\} \notin \widetilde{F} \\
 \hline
 \text{fsender}(C, f) = \perp
 \end{array}$$

These two functions will be used later on when we present our typing rules for expressions. The first one retrieves a tuple with a set of types for arguments and a type for a function body, while the latter retrieves for a function f , its required type for the implicit variable msg.sender . When the annotation is omitted from the function declaration, we return an empty set, as there will be no restrictions on the type.

5.4.4 State Variables

$$\begin{array}{c}
 \text{(STATE VARIABLE LOOKUP)} \\
 \hline
 CT(C) = \text{contract } C \text{ is } \widetilde{D} [\widetilde{SC} \mapsto \widetilde{e}] \{(\widetilde{T} s); K \widetilde{F}\} \\
 \hline
 \text{sv}(C) = \widetilde{T} s
 \end{array}$$

This is a state variable lookup, retrieving a contract's state variables.

$$\text{mkstate}(\widetilde{Cs}, \text{state}) = \begin{cases} \text{state} & \widetilde{Cs} = [] \\ \text{mkstate}(\widetilde{Cs}', \text{state} \cdot \text{sv}(C_i)) & \widetilde{Cs} = C_i \cdot \widetilde{Cs}' \end{cases}$$

This function computes the state variables of a contract based on its linearization. It receives a list of contracts representing the contract's linearization and a state which will be an empty map in the beginning. Then it recursively pop the head of the list and concatenates each contract's state to our state. When the list of contracts is empty it returns our final state. Supposing that Context has the implementation we seen in 5.5, computing the state of contract Game we will get:

$$\text{mkstate}([\text{Game}; \text{Context}], []) = \\ [(\text{mapping}(\text{address} \Rightarrow \text{NFTStorage}), \text{stores}); (\text{uint}, \text{balance})]$$

5.4.5 Update Contract Balance

$$\text{uptbal} : \text{BLOCKCHAIN} \times \text{ADDRESS} \times \text{BALANCE} \rightarrow \text{BLOCKCHAIN}$$

$$\text{uptbal}(\beta, a, n) = \begin{cases} \beta[(c, a) \mapsto (C, \widetilde{s} : \widetilde{v}, n' + n)] & \beta((c, a)) = (C, \widetilde{s} : \widetilde{v}, n') \wedge n' + n \geq 0 \\ \beta[a \mapsto (n' + n)] & \text{if } \beta(a) = n' \wedge n' + n \geq 0 \\ \perp & \text{otherwise} \end{cases}$$

This function updates the balance of both EOAs or contracts. It receives a blockchain β , an address a and a number n . Note that n can either be negative or positive. Also this is a function because in β both c and a are unique and so is the pair formed by both, c representing a reference to a contract. Obviously an EOA does not have c , so it is only referenced by an address.

$$\text{Top} : \widetilde{\text{ADDRESS}} \rightarrow \text{ADDRESS} \cup \{\emptyset\}$$

$$\text{Top}(\sigma) = \begin{cases} a & \text{if } \sigma = \sigma' \cdot a \\ \emptyset & \text{if } \sigma = \epsilon \end{cases}$$

Returns the last element added to our stack of addresses, identifying the last address which initiated a transaction.

5.4.6 Constructors

(RETRIEVE CONSTRUCTOR PARAMETERS)

$$\text{CT}(C) = \text{contract } C \text{ is } \widetilde{D} [D_1 \mapsto \widetilde{e}_1; \dots; D_n \mapsto \widetilde{e}_n] \{(\widetilde{T} s); K F_1; \dots; F_n\} \\ K = \text{constructor } ((\widetilde{T} x))(e)$$

$$\text{cparams}(C) = \widetilde{x}$$

Auxiliary function that receives a contract and returns its constructor parameters.

(RETRIEVE CONSTRUCTOR BODY)

$$\text{CT}(C) = \text{contract } C \text{ is } \widetilde{D} [D_1 \mapsto \widetilde{e}_1; \dots; D_n \mapsto \widetilde{e}_n] \{(\widetilde{T} s); K F_1; \dots; F_n\} \\ K = \text{constructor } ((\widetilde{T} x))(e)$$

$$\text{cbody}(C) = e$$

Same as the previous one, however this function returns the body of the constructor.

$$\begin{array}{c}
 \text{(RETRIEVE SUPER CONTRACTS ARGUMENTS MAP)} \\
 \hline
 CT(C) = \text{contract } C \text{ is } \tilde{D} [D_1 \mapsto \tilde{e}_1; \dots; D_n \mapsto \tilde{e}_n] \{(\tilde{T} s); K F_1; \dots; F_n\} \quad n \geq 1 \\
 \hline
 \text{cargs}(C) = [D_1 \mapsto [\text{cparams}(D_1) \mapsto \tilde{e}_1]; \dots; D_n \mapsto [\text{cparams}(D_n) \mapsto \tilde{e}_n]]
 \end{array}$$

This function computes a map where it maps each super contract to a map between its constructor parameters and values each one is going to get. It receives a contract definition as input.

$$\text{mkcvars}(\tilde{C}s, \tilde{v}s, \tilde{e}) = \begin{cases} \tilde{v}s & \tilde{C}s = [] \\ \text{mkcvars}(\tilde{C}s', (\tilde{v}s \cdot [C_i \mapsto [\text{cparams}(C_i) \mapsto \tilde{e}]])) \cdot \text{cargs}(C_i), \tilde{e}) & \tilde{C}s = C_i \cdot \tilde{C}s' \\ & \wedge i = 1 \\ \text{mkcvars}(\tilde{C}s', \tilde{v}s \cdot \text{cargs}(C_i), \tilde{e}) & \tilde{C}s = C_i \cdot \tilde{C}s' \\ & \wedge i \neq 1 \end{cases}$$

This auxiliary function takes as input a set of contract definitions, a mapping between constructor parameters and their corresponding values, and a set of values. It then computes a mapping between constructor parameters and values required for the initialization of each contract and its supercontracts.

$$\begin{array}{c}
 \text{(EXECUTE CONSTRUCTOR)} \\
 \hline
 \text{cparams}(C) = \tilde{x} \quad \text{cbody}(C) = e \quad e_s = \langle \beta', e\{\tilde{v}s[C]\} \rangle \\
 \hline
 \text{exec}(C, \beta, \sigma, \tilde{v}s) = e_s
 \end{array}$$

Simple function where constructor is executed and we get a tuple with a modified blockchain and the expression executed with the arguments passed as parameters. It receives a contract name, a blockchain, a sigma and a set of values (arguments).

5.5 Operational Semantics Rules

Now we will present our operational semantics rules. These rules are used by our PoC interpreter and its goal is to be as close as possible to Solidity.

5.5.1 Arithmetic and Boolean Operations

$$\begin{array}{c}
 \text{(ARIT/BOOL)} \\
 v = a \diamond b \\
 \hline
 \langle \beta, \beta_i, \sigma, a \diamond b \rangle \rightarrow \langle \beta, \beta_i, \sigma, v \rangle
 \end{array}$$

All arithmetic and boolean operators follow the consensual trivial rule. it is important to note that we support the following arithmetic operations: plus, minus, division, multiplication, exponential and modulo. For boolean operations we support: negation, conjunction, disjunction, lesser (or equal), greater (or equal), equals and inequals.

Next, one can see how these rules should be applied. In this case, we demonstrate a simple addition operation, where we sum two values to obtain a new value, which represents the result.

$$\langle \beta, \beta, \emptyset, 1 + 1 \rangle \rightarrow \langle \beta, \beta, \emptyset, 2 \rangle$$

5.5.2 Variable Operations

(VAR)

$$\frac{}{\langle \beta, \beta_i, \sigma, x \rangle \rightarrow \langle \beta, \beta_i, \sigma, \beta(x) \rangle}$$

Accesses a variable and returns its value.

(ASSIGN)

$$\frac{}{\langle \beta, \beta_i, \sigma, x = v \rangle \rightarrow \langle \beta \cdot [x \mapsto v], \beta_i, \sigma, v \rangle}$$

Assigns a value to an existing variable.

(LET)

$$\frac{x \notin \text{dom}(\beta)}{\langle \beta, \beta_i, \sigma, T x = v; e \rangle \rightarrow \langle \beta \cdot [x \mapsto v], \beta_i, \sigma, v; e \rangle}$$

This rule allows developers to declare a new variable. It is important to highlight that, unlike Solidity, which allows for collision names between state variables and local variables, our approach is designed with a distinct advantage. We do allow for collisions, but the state variable is never shadowed. To access its value, users must employ the keyword **this**.

Next, we illustrate these rules with an example. Initially, we declare a variable x as a uint with a value of 3, using LET rule. Subsequently, we modify this value using the ASSIGN rule, and finally, we read the variable x from the blockchain, using VAR rule.

$$\langle \beta, \beta, \emptyset, \text{uint } x = 3; x = 2 \rangle \rightarrow \langle \beta \cdot [x \mapsto 3], \beta, \emptyset, x = 2 \rangle \rightarrow$$

$$\langle \beta \cdot [x \mapsto 2], \beta, \emptyset, 2 \rangle$$

$$\langle \beta \cdot [x \mapsto 2], \beta, \emptyset, x \rangle \rightarrow \langle \beta \cdot [x \mapsto 2], \beta, \emptyset, 2 \rangle$$

Now, we present the rules for reading and writing state variables.

(STATEREAD)

$$\frac{\exists a. \beta((c, a)) = (C, \tilde{s}v, n) \quad s \in \tilde{s}}{\langle \beta, \beta_i, \sigma, c.s \rangle \rightarrow \langle \beta, \beta_i, \sigma, v_s \rangle}$$

$$\begin{array}{c}
 \text{(STATEASSIGN)} \\
 \frac{\exists a. \beta((c, a)) = (C, \tilde{s}v, n) \quad s \in \tilde{s}}{\langle \beta, \beta_i, \sigma, c.s = v' \rangle \rightarrow \langle \beta \cdot [c.s \mapsto v'], \beta_i, \sigma, v' \rangle}
 \end{array}$$

These rules lets us access a state variable or write a new value for it. Obviously if there is no such state variable s this should fail.

These rules allow us to access a state variable or write a new value to it. However, if there is no such state variable s , this operation should fail.

To illustrate these rules, we assume that we already have a contract in the blockchain. In this case, we use generic values to identify the contract, denoted by the pair containing values c and a , and C for the contract name. The examples are as follows:

$$\langle [(c, a) \mapsto (C, [x \mapsto 0], 0) \cdot [\mathbf{this} \mapsto c], [(c, a) \mapsto (C, [x \mapsto 0], 0), \emptyset, \mathbf{this}.x = 30] \rangle \rightarrow$$

$$\langle [(c, a) \mapsto (C, [x \mapsto 30], 0) \cdot [\mathbf{this} \mapsto c], [(c, a) \mapsto (C, [x \mapsto 0], 0), \emptyset, 30] \rangle$$

$$\langle [(c, a) \mapsto (C, [x \mapsto 30], 0) \cdot [\mathbf{this} \mapsto c], [(c, a) \mapsto (C, [x \mapsto 0], 0), \emptyset, \mathbf{this}.x] \rangle \rightarrow$$

$$\langle [(c, a) \mapsto (C, [x \mapsto 30], 0) \cdot [\mathbf{this} \mapsto c], [(c, a) \mapsto (C, [x \mapsto 0], 0), \emptyset, 30] \rangle$$

Note how we initially have the state variable x with a value of 0. Using STATEASSIGN, we change the value to 30, and after that, we use STATEREAD to retrieve its value from the blockchain. In both operations, we make use of the variable **this** to access the contract's reference.

5.5.3 Constructor Rule

(NEW)

$$\frac{\begin{array}{l} (c, a) \notin \text{Dom}(\beta) \quad \text{mkstate}(\tilde{S}\tilde{C}, [\]) = \tilde{T}s \quad \text{c3linearization}(CT(C)) = \tilde{S}\tilde{C} \\ \text{mkcvars}(\tilde{S}\tilde{C}, [\], \tilde{v}s) = \tilde{v}s \quad \text{exec}(\tilde{S}\tilde{C}, \beta, \sigma, \tilde{v}s) = \langle \beta', e_s \rangle \quad \text{Top}(\sigma) \neq \emptyset \end{array}}{\langle \beta, \beta_i, \sigma, \text{new } C.\text{value}(n)(\tilde{v}) \rangle \rightarrow \langle \text{uptbal}(\beta', \text{Top}(\sigma), -n).[(c, a) \mapsto (C, \text{init}(\tilde{T}s), n)].[\tilde{x} \mapsto \tilde{v}], \beta_i, \sigma, e_s; c \rangle}$$

(NEW-R)

$$\frac{\text{uptbal}(\beta, \text{Top}(\sigma), -n) = \perp \quad \text{Top}(\sigma) \neq \emptyset}{\langle \beta, \beta_i, \sigma, \text{new } C.\text{value}(n)(\tilde{v}) \rangle \rightarrow \langle \beta, \beta_i, \sigma, \text{revert} \rangle}$$

NEW rule deploys a new contract on the blockchain. In these rules, we made a modification to Di Pirro's formalization. Instead of requiring a one-to-one relation between constructor arguments and state variables, there is a freedom in these aspect. So constructor body and its arguments are evaluated as a function and the contract's state variables are

initially set to a default value. Our specification is identical to what Solidity does. NEW-R is the case if there is no balance on the EOA or contract deploying.

We should also note that in the original formalization there was a NEW-2 rule, which increased the blockchain liquidity and was used to model a creation of external-owned accounts (EOAs). However we dropped this rule since we have in mind that contracts and EOA are distinct and accounts in blockchain is just a mapping between their addresses and balances.

In Section 5.5.8, we provide a detailed example illustrating how this rule applies in our environment.

5.5.4 Conditional and Sequential Statements

(IF-TRUE)

$$\frac{}{\langle \beta, \beta_i, \sigma, \text{if true then } e1 \text{ else } e2 \rangle \rightarrow \langle \beta, \beta_i, \sigma, e1 \rangle}$$

(IF-FALSE)

$$\frac{}{\langle \beta, \beta_i, \sigma, \text{if false then } e1 \text{ else } e2 \rangle \rightarrow \langle \beta, \beta_i, \sigma, e2 \rangle}$$

These rules have the standard meaning. Important to note that if we want to omit either branch the omitted branch should take the value unit internally.

Here is an example how this rule simplifies:

$$\langle \beta, \beta, \emptyset, \text{if "0x01" == "0x01" then 2 else 3} \rangle \rightarrow \langle \beta, \beta, \emptyset, \text{if true then 2 else 3} \rangle \rightarrow \langle \beta, \beta, \emptyset, 2 \rangle$$

(SEQ-C)

$$\frac{\sigma = \varepsilon}{\langle \beta, \beta_i, \sigma, v; e \rangle \rightarrow \langle \beta, \beta_i, \sigma, e \rangle}$$

(SEQ-R)

$$\frac{\sigma = \varepsilon}{\langle \beta, \beta_i, \sigma, \text{revert}; e \rangle \rightarrow \langle \beta_i, \beta_i, \sigma, \text{revert} \rangle}$$

(SEQ)

$$\frac{\text{Top}(\sigma) = a}{\langle \beta, \beta_i, \sigma, v; e \rangle \rightarrow \langle \beta, \beta_i, \sigma, e \rangle}$$

These rules have to take into account different scenarios. First, when an expression like $v; e$ is found at the top level ($\text{Top}(\sigma) = \emptyset$) then the last transaction was successful. In fact,

it evaluated to a value v , which can be discarded to proceed with the next transaction e . Furthermore, since evaluating to a value is equivalent to a commit, we apply the changes and proceed with β . This scenario is modeled by SEQ-C. Secondly, $v; e$ might be at an inner level ($\text{Top}(\sigma) = a$). In this case (SEQ) we cannot apply any changes, since the transaction is not over yet. Hence, we just proceed with the evaluation of e , discarding v . Lastly, a revert may have been thrown during the evaluation, so instead of $v; e$ the expression involving sequential composition is $\text{revert}; e$ (SEQ-R). In this case we do not further proceed in evaluating the other transactions and we replace β with β_i .

We make extensively use of these rules in our example in section 5.5.8.

5.5.5 Mappings

$$\begin{array}{c}
 \text{(MAPREAD)} \\
 \frac{\exists a. \beta((c, a)) = (C, \tilde{s}\tilde{v}, n)}{\langle \beta, \beta_i, \sigma, M[v1] \rangle \rightarrow \langle \beta, \beta_i, \sigma, M(v1) \rangle}
 \end{array}$$

$$\begin{array}{c}
 \text{(MAPREAD)} \\
 \frac{M[v1] = \perp}{\langle \beta, \beta_i, \sigma, M[v1] \rangle \rightarrow \langle \beta, \beta_i, \sigma, \text{defaultvalue}(T) \rangle}
 \end{array}$$

$$\begin{array}{c}
 \text{(MAPASSIGN)} \\
 \frac{M' = M \setminus \{(v1, M(v1))\} \cup \{(v1, v2)\}}{\langle \beta, \beta_i, \sigma, M[v1 \mapsto v2] \rangle \rightarrow \langle \beta, \beta_i, \sigma, M' \rangle}
 \end{array}$$

In our formalization, users have the capability to define mappings, just like in Solidity. The initial two rules ensure that our mappings read function operates as a total function: in the worst-case scenario, attempting to access a mapping position should return a default value corresponding to its return type. The ASSIGN rule, on the other hand, serves to replace the value associated with a specific key. Unlike Di Pirro's original formalization, we align with Solidity's behavior. Consequently, when a key lacks a value, we return the default value corresponding to the value type.

Next, we exemplify how these rules work in our language. In the first example, we demonstrate that attempting to access the value of a key in an empty map always yields a value in return. In this case, we assume a map where keys and values are of type `uint`, and we return 0 as it is the default value for this type. In the second example, we illustrate that if we insert a key with a value into an empty map, the returned map contains this key. Lastly, using the map we obtained, we attempt to access the key and obtain the desired return: the value associated with the key in the map.

$$\langle \beta, \beta, \emptyset, [][1] \rangle \rightarrow \langle \beta, \beta, \emptyset, 0 \rangle$$

$$\langle \beta, \beta, \emptyset, [][1 \mapsto 2] \rangle \rightarrow \langle \beta, \beta, \emptyset, [1 \mapsto 2] \rangle$$

$$\langle \beta, \beta, \emptyset, [1 \mapsto 2][1] \rangle \rightarrow \langle \beta, \beta, \emptyset, 2 \rangle$$

5.5.6 Functions

(BALANCE)

$$\exists c. \beta((c, a)) = (C, \tilde{s}v, n) \vee \beta(a) = n$$

$$\langle \beta, \beta_i, \sigma, \text{balance}(a) \rangle \rightarrow \langle \beta, \beta_i, \sigma, n \rangle$$

As in Solidity, our system includes an operation to retrieve a balance from an address. Above, we provide an operational semantic rule that models this behavior.

Next we present a practical example of this rule, assuming a contract game deployed in our blockchain:

$$\langle \beta_a \cdot [(c_{\text{game}}, a_{\text{game}}) \mapsto (\text{Game}, s_1, 200)], \beta_a \cdot [(c_{\text{game}}, a_{\text{game}}) \mapsto (\text{Game}, s_1, 200)], \text{balance}(a_{\text{game}}) \rangle \rightarrow$$

$$\langle \beta_a \cdot [(c_{\text{game}}, a_{\text{game}}) \mapsto (\text{Game}, s_1, 200)], \beta_a \cdot [(c_{\text{game}}, a_{\text{game}}) \mapsto (\text{Game}, s_1, 200)], 200 \rangle$$

(ADDRESS)

$$\exists a. \beta((c, a)) = (C, \tilde{s}v, n) \vee c = a$$

$$\langle \beta, \beta_i, \sigma, \text{address}(c) \rangle \rightarrow \langle \beta, \beta_i, \sigma, a \rangle$$

Similarly to the previous operation, this operation is also present in Solidity and it just casts a contract instance or EOA. We only can point to externally-owned accounts by their address, so in that case we cannot cast an instance of it, but an address, which does essentially nothing. Note that c can be either a contract instance or an address.

Now, let's consider a similar example to the one previously shown, but this time focusing on this operation:

$$\langle \beta_a \cdot [(c_{\text{game}}, a_{\text{game}}) \mapsto (\text{Game}, s_1, 200)], \beta_a \cdot [(c_{\text{game}}, a_{\text{game}}) \mapsto (\text{Game}, s_1, 200)], \emptyset, \text{address}(c_{\text{game}}) \rangle \rightarrow$$

$$\langle \beta_a \cdot [(c_{\text{game}}, a_{\text{game}}) \mapsto (\text{Game}, s_1, 200)], \beta_a \cdot [(c_{\text{game}}, a_{\text{game}}) \mapsto (\text{Game}, s_1, 200)], \emptyset, a_{\text{game}} \rangle$$

(TRANSFER)

$$\beta(a) = (C, sv, n)$$

$$\text{fbody}(C, fb, \{\}) = (\{\}, e) \quad \beta' = \text{uptbal}(\text{uptbal}(\beta, \text{Top}(\sigma), -n), a, n)$$

$$\langle \beta, \beta_i, \sigma, a.\text{transfer}(n) \rangle \rightarrow \langle \beta', \beta_i, \sigma * a, e\{\text{this} := c, \text{msg.sender} := \text{Top}(\sigma), \text{msg.value} := n\} \rangle$$

$$\begin{array}{c}
 \text{(TRANSFER-R)} \\
 \frac{uptbal(\beta, \text{Top}(\sigma), -n) = \perp}{\langle \beta, \beta_i, \sigma, v; e \rangle \rightarrow \langle \beta, \beta_i, \sigma, \text{revert} \rangle}
 \end{array}$$

These rules model the behavior of the transfer function in Solidity. This function is callable through an address pointer and is used to transfer Ether from the contract to the address specified in the function call. The value of the transfer is passed as an argument n . Notably, we have a second rule that causes this function to fail if the contract's balance does not contain the required amount for the transfer

In the example below, we illustrate how this rule works for a transfer between two contracts. Assuming we have two contracts deployed in our blockchain: Game and Dummy. Game has a balance of 200, and Dummy has a balance of 0. If we invoke a transfer function from Game to Dummy with a value of 200, we will have Game as our address at the top of the stack σ . The expression will be reduced to the body of the fallback function of the Dummy contract (e_{fb}) with all the variables set: `this`, `msg.sender`, and `msg.value`. We also transfer 200 from Game to Dummy, leaving Game with a balance of 0 and Dummy with a balance of 200.

$$\begin{array}{l}
 \langle \beta_a \cdot [(c_{\text{game}}, a_{\text{game}}) \mapsto (\text{Game}, s_1, 200)] \cdot [(c_{\text{dummy}}, a_{\text{dummy}}) \mapsto (\text{Dummy}, s_2, 0)], \\
 \beta_a \cdot [(c_{\text{game}}, a_{\text{game}}) \mapsto (\text{Game}, s_1, 200)] \cdot [(c_{\text{dummy}}, a_{\text{dummy}}) \mapsto (\text{Dummy}, s_2, 0)], [a_{\text{game}}], a_{\text{dummy}}.\text{transfer}(200) \rangle \rightarrow \\
 \langle \beta_a \cdot [(c_{\text{game}}, a_{\text{game}}) \mapsto (\text{Game}, s_1, 0)] \cdot [(c_{\text{dummy}}, a_{\text{dummy}}) \mapsto (\text{Dummy}, s_2, 200)], \\
 [\text{this} \mapsto c_{\text{dummy}}; \text{msg.sender} \mapsto a_{\text{game}}; \text{msg.value} \mapsto 200], \\
 \beta_a \cdot [(c_{\text{game}}, a_{\text{game}}) \mapsto (\text{Game}, s_1, 200)] \cdot [(c_{\text{dummy}}, a_{\text{dummy}}) \mapsto (\text{Dummy}, s_2, 0)], [a_{\text{game}}], e_{fb} \rangle
 \end{array}$$

$$\begin{array}{c}
 \text{(CONTRRETR)} \\
 \frac{\exists c. \beta(c, a) = C \quad \text{subtyping}(C, D, CT)}{\langle \beta, \beta_i, \sigma, D(a) \rangle \rightarrow \langle \beta, \beta_i, \sigma, c \rangle}
 \end{array}$$

$$\begin{array}{c}
 \text{(CONTRRETR-R)} \\
 \frac{\exists c. \beta(c, a) = C \quad \neg \text{subtyping}(C, D, CT)}{\langle \beta, \beta_i, \sigma, D(a) \rangle \rightarrow \langle \beta, \beta_i, \sigma, \text{revert} \rangle}
 \end{array}$$

Since there may not be possible to know what type of address we are passing before hand (i.e in compile time), we need to have a premise to ensure that the address we are casting

to a contract is pointing to a contract C that is a subtype of D . We also do this verification in the typechecker and it saves us this computation before hand.

This subtyping will be introduced in section 7.3.2 and so we do not develop more here.

Considering the next example:

$$\langle \beta_a \cdot [(c_{\text{game}}, a_{\text{game}}) \mapsto (\text{Game}, s_1, 200)], \beta_a \cdot [(c_{\text{game}}, a_{\text{game}}) \mapsto (\text{Game}, s_1, 200)], \emptyset, \text{Game}(a_{\text{game}}) \rangle \rightarrow$$

$$\langle \beta_a \cdot [(c_{\text{game}}, a_{\text{game}}) \mapsto (\text{Game}, s_1, 200)], \beta_a \cdot [(c_{\text{game}}, a_{\text{game}}) \mapsto (\text{Game}, s_1, 200)], \emptyset, c_{\text{game}} \rangle$$

As we can see, we use the cast operation applied to a_{game} , which is the address reference to our deployed contract, to obtain the value c_{game} - a reference to the contract associated with the address. Since our contract Game also inherits from the Context contract, we could also use Context to cast. However, using any other contract for casting, aside from these two, would have resulted in a revert.

(CALL)

$$\frac{\begin{array}{l} \exists a. \beta((c, a)) = (C, \tilde{s}\tilde{v}, n) \quad \text{fbody}(C, f, \tilde{v}) = (\tilde{x}, e) \\ \tilde{x} \notin \text{Dom}(\beta) \quad \beta' = \text{uptbal}(\text{uptbal}(\beta, \text{Top}(\sigma), -n), a, n) \cdot [\tilde{x} \mapsto \tilde{v}] \\ e_s = e\{\text{this} := c, \text{msg.sender} := \text{Top}(\sigma), \text{msg.value} := n\} \end{array}}{\langle \beta, \beta_i, \sigma, c.f.\text{value}(n)(\tilde{v}) \rangle \rightarrow \langle \beta, \beta_i, \sigma \cdot a, e_s \rangle}$$

(CALL-R)

$$\frac{\text{uptbal}(\beta, \text{Top}(\sigma), -n) = \perp}{\langle \beta, \beta_i, \sigma, C(a) \rangle \rightarrow \langle \beta, \beta_i, \sigma, \text{revert} \rangle}$$

These rules enable one contract to call a function within another contract. To execute this operation, we require a contract reference c , a function name f , a value n to pass as the value (which will represent the special value msg.value inside the called function) and a set of arguments (potential empty). This function call can fail if and only if, the contract where we are executing this, does not have enough balance to pass the value n .

(CALLTOPLEVEL)

$$\frac{\begin{array}{l} \exists a. \beta((c, a)) = (C, \tilde{s}\tilde{v}, n) \quad \text{fbody}(C, f, \tilde{v}) = (\tilde{x}, e) \\ \tilde{x} \notin \text{Dom}(\beta) \quad \beta' = \text{uptbal}(\text{uptbal}(\beta_w, a', -n), a, n) \cdot [\tilde{x} \mapsto \tilde{v}] \\ e_s = e\{\text{this} := c, \text{msg.sender} := a', \text{msg.value} := n\} \quad \text{Top}(\sigma) = \emptyset \end{array}}{\langle \beta, \beta_i, \sigma, c.f.\text{value}(n).\text{sender}(a')(\tilde{v}); e' \rangle \rightarrow \langle \beta', \beta_i, \sigma \cdot a, e_s; e' \rangle}$$

(CALLTOPLEVEL-R)

$$\frac{\text{uptbal}(\beta, a', -n) = \perp \quad \text{Top}(\sigma) = \emptyset}{\langle \beta, \beta_i, \sigma, c.f.\text{value}(n).\text{sender}(a')(\tilde{v}); e' \rangle \rightarrow \langle \beta', \beta_i, \sigma \cdot a, \text{revert}; e' \rangle}$$

While this rule does not inherently exist in Solidity, it represents our and Di Pirro's solution for initiating transactions. This rule is designed to be invoked by a specific externally-owned account, enabling them to interact with our deployed contracts.

In our example in Section 5.5.8, we use CALLTOPLEVEL and one can see how it works practically. It is worth noting that the only distinction between CALLTOPLEVEL and CALL lies in the ability of CALLTOPLEVEL to specify the value of `msg.sender` as one of the arguments in the function call. In contrast, CALL does not allow for the setting of `msg.sender` directly, and its value is determined by invoking the Top function.

(RETURN)

$$\frac{}{\langle \beta, \beta_i, \sigma \cdot a, \text{return } v \rangle \rightarrow \langle \beta', \beta_i, \sigma, v \rangle}$$

(RETURN-R)

$$\frac{}{\langle \beta, \beta_i, \sigma \cdot a, \text{return revert} \rangle \rightarrow \langle \beta_i, \beta_i, \sigma, \text{revert} \rangle}$$

We have straightforward rules for returning a value: we can either return a value or trigger a revert. The key distinction lies in the fact that when we initiate a revert, we replace β with β_i , which represents the blockchain's original state before the transaction that was reverted was executed.

5.5.7 Computation Rules

(REVERT)

$$\frac{}{\langle \beta, \beta_i, \sigma, \text{revert} \rangle \rightarrow \langle \beta, \beta_i, \sigma, \text{revert} \rangle}$$

(CONG)

$$\frac{\langle \beta, \beta_i, \sigma, e \rangle \rightarrow \langle \beta', \beta_i, \sigma', e' \rangle}{\langle \beta, \beta_i, \sigma, E[e] \rangle \rightarrow \langle \beta', \beta_i, \sigma', E[e'] \rangle}$$

These two rules above define our computation rules. $E[e]$ represents the expression obtained by replacing the hole in context E by the expression e . Below (Figure 5.6) we formally define all evaluation contexts.

$$\begin{aligned} E ::= & [] \mid \text{balance}(E) \mid \text{address}(E) \mid E.s \mid \\ & E.\text{transfer}(e) \mid a.\text{transfer}(E) \mid E;e \mid \\ & \text{new } C.\text{value}(E)(\tilde{e}) \mid \\ & C(e) \mid E.f.\text{value}(e)(\tilde{e}) \mid c.f.\text{value}(E)(\tilde{e}) \mid \\ & E.f.\text{value}(e).\text{sender}(e)(\tilde{e}) \mid c.f.\text{value}(E).\text{sender}(e)(\tilde{e}) \mid \\ & c.f.\text{value}(n).\text{sender}(E)(\tilde{e}) \mid \\ & T \ x = E;e \mid x = E \mid E.s = e \mid c.s = E \mid E[e] \mid \\ & M[E] \mid E[e \rightarrow e] \mid M[E \rightarrow e] \mid M[v \rightarrow E] \mid \\ & \text{if } E \text{ then } e \text{ else } e \mid \text{return } E \end{aligned}$$

Figure 5.6: Evaluation Contexts of Featherweight Solidity 2.0

5.5.8 Operational Semantics by Example

In this section we will show some examples of our operational semantics rules are applied. Let's assume that our initial state for all examples is:

$$\langle [0x01 \rightarrow 2000; 0x02 \rightarrow 2000], [0x01 \rightarrow 2000; 0x02 \rightarrow 2000], \emptyset, e \rangle$$

As one can see, our blockchain already have two EOAs, one with address 0x01 and another with 0x02 and both with 2000 balance.

To simplify and adjust to our needs we will redefine our Game contract

```

1  contract Game {
2      uint totalNfts;
3      mapping(uint => address) nfts;
4      uint price;
5      address owner;
6      constructor() {
7          this.owner = msgsender;
8      }
9      function fb() {}
10     function setNFTPrice(uint price) {
11         if(msgsender == this.owner) {
12             this.price = price;
13         }
14         else {
15             revert;
16         }
17     }
18     function buyNFT() returns (uint) {
19         if(msgvalue >= price){
20             this.nfts = (this.nfts[this.totalNfts] = msgsender);
21             this.totalNfts = this.totalNfts + 1;
22             return this.totalNfts;
23         }
24         else {
25             revert;
26         }
27     }
28     function getOwner(uint nft) returns (address) {
29         return this.nfts[nft];
30     }
31 }

```

Figure 5.7: A simpler NFT Game contract

Initially, the contract is deployed by the address "0x01". Subsequently, this address will proceed to set the price for each NFT at 200 by invoking setNFTPrice function. Following this, address "0x02" will execute the buyNFT function to acquire an NFT. Finally, it will invoke the getOwner function to confirm its ownership of the recently purchased token.

```

e := Game game = new Game();

game.setNFTPrice.value(0).sender("0x01")(200);

uint num = game.buyNFT.value(200).sender("0x02")();

game.getOwner.value(0).sender("0x02")(num)

```

Next we will exemplify the application of the previously mentioned rules and how they reduce our expressions. It is important to note that for the sake of simplicity, we will omit each variable after its initial appearance, assuming that they persist in β after their omission.

We start by applying together both the NEW and LET rule. Applying these rules we deploy the contract to the blockchain and store its reference on a variable game. Then applying the SEQ-C rule we commit our blockchain, i.e replace the blockchain replica by the actual blockchain. In the end we apply the CALL rule, simplifying our expression with its arguments values.

$$\langle \beta_a, \beta_a, \emptyset, \text{Game game} = \text{new Game.value}(0)(); e' \rangle \rightarrow$$

$$\langle \beta_a \cdot [(c_{\text{game}}, a_{\text{game}}) \mapsto (\text{Game}, s_1, 0)], \beta_a, \emptyset, \text{Game game} = c_{\text{game}}; e' \rangle \rightarrow \text{SEQ-C}$$

$$\langle \beta_a \cdot [(c_{\text{game}}, a_{\text{game}}) \mapsto (\text{Game}, s_1, 0)] \cdot [\text{game} \mapsto c_{\text{game}}], \beta_a \cdot [(c_{\text{game}}, (a_{\text{game}}) \mapsto (\text{Game}, s_1, 0)], \emptyset,$$

$$\text{game.setNFTPrice.value}(0).sender("0x01")(200); e' \rangle \rightarrow \text{CALLTOPLEVEL}$$

After we have our CALLTOPLEVEL rule applied we add to our variable environment two keys: msgsender and price and its corresponding values. Doing this helps us replacing in the expression wherever these values appear. Eventually we commit the result applying the SEQ-C rule again. After all these rules are applied we have our contract's state transitioned from s_1 to s_2 and we get the third expression to evaluate.

```

⟨βa · [(cgame, agame) ↦ (Game, s1, 0)] · [msgsender ↦ "0x01"; this ↦ cgame; msg.value ↦ 0; price ↦ 200],
βa · [(cgame, agame) ↦ (Game, s1, 0)], [agame],
if(msgsender == this.owner) then this.price = price else revert; e'⟩ →
⟨βa · [(cgame, agame) ↦ (Game, s1, 0)], βa · [(cgame, agame) ↦ (Game, s1, 0)], [agame],
if("0x01" == "0x01") then this.price = 200 else revert; e'⟩ →
⟨βa · [(cgame, agame) ↦ (Game, s1, 0)], βa · [(cgame, agame) ↦ (Game, s1, 0)], [agame]
if(true) then this.price = 200 else revert; e'⟩ →
⟨βa · [(cgame, agame) ↦ (Game, s2, 0)], βa · [(cgame, agame) ↦ (Game, s1, 0)], [agame],
this.price = 200; e'⟩ → SEQ-C
⟨βa · [(cgame, agame) ↦ (Game, s2, 0)], βa · [(cgame, agame) ↦ (Game, s2, 0)], ∅,
uint num = game.buyNFT.value(200).sender("0x02"); e'⟩ → CALLTOPLEVEL
    
```

In the next step, we will assess the expression in which address "0x02" purchases an NFT and stores the token ID in a variable num. We will exclude the evaluation of the function body for simplicity, as it closely resembles previous instances. Subsequently, by utilizing this variable as the argument for the getOwner function, we obtain the return value, which is the address "0x02".

$$\begin{aligned}
 & \langle \beta_a \cdot [(c_{\text{game}}, a_{\text{game}}) \mapsto (\text{Game}, s_2, 0)], \beta_a \cdot [(c_{\text{game}}, a_{\text{game}}) \mapsto (\text{Game}, s_2, 0)] \cdot [\text{msgsender} \mapsto "0x02"]; \\
 & \text{this} \mapsto c_{\text{game}}; \text{msg.value} \mapsto 200], [a_{\text{game}}], \\
 & \text{uint num} = \text{game.buyNFT.value}(200).\text{sender}("0x02")(); e' \rangle \rightarrow^* \\
 & \langle \beta_b \cdot [(c_{\text{game}}, a_{\text{game}}) \mapsto (\text{Game}, s_3, 0)], \beta_a \cdot [(c_{\text{game}}, a_{\text{game}}) \mapsto (\text{Game}, s_2, 0)], [a_{\text{game}}], \text{uint num} = 0; e' \rangle \rightarrow \\
 & \langle \beta_b \cdot [(c_{\text{game}}, a_{\text{game}}) \mapsto (\text{Game}, s_3, 200)], \beta_b \cdot [(c_{\text{game}}, a_{\text{game}}) \mapsto (\text{Game}, s_3, 200)], \emptyset, \\
 & \text{game.getOwner.value}(0).\text{sender}("0x02")(num) \rangle \rightarrow \\
 & \langle \beta_b \cdot [(c_{\text{game}}, a_{\text{game}}) \mapsto (\text{Game}, s_3, 200)] \cdot [\text{msgsender} \mapsto "0x02"; \text{this} \mapsto c_{\text{game}}; \text{msg.value} \mapsto 0; \text{nft} \mapsto 0], \\
 & \beta_b \cdot [(c_{\text{game}}, a_{\text{game}}) \mapsto (\text{Game}, s_3, 200)], [a_{\text{game}}], \text{return this.nfts}[\text{nft}] \rangle \rightarrow \\
 & \langle \beta_b \cdot [(c_{\text{game}}, a_{\text{game}}) \mapsto (\text{Game}, s_3, 200)], \beta_b \cdot [(c_{\text{game}}, a_{\text{game}}) \mapsto (\text{Game}, s_3, 200)], \emptyset, "0x02" \rangle \\
 & \beta_a = [0x01 \rightarrow 2000; 0x02 \rightarrow 2000] \\
 & \beta_b = [0x01 \rightarrow 2000; 0x02 \rightarrow 1800] \\
 & s_1 = [\text{nfts} \mapsto []; \text{totalNfts} \mapsto 0; \text{price} \mapsto 0; \text{owner} \mapsto "0x01"] \\
 & s_2 = [\text{nfts} \mapsto []; \text{totalNfts} \mapsto 0; \text{price} \mapsto 200; \text{owner} \mapsto "0x01"] \\
 & s_3 = [\text{nfts} \mapsto [0 \mapsto "0x02"]; \text{totalNfts} \mapsto 1; \text{price} \mapsto 200; \text{owner} \mapsto "0x01"]
 \end{aligned}$$

5.5.9 Implementation in OCaml

When choosing to implement operational semantics, two distinct approaches are available: small-step semantics and big-step semantics. The latter differs from the former by hiding even more execution details.

For our project, which aimed to develop a proof-of-concept interpreter, we opted for big step semantics due to several advantages aligned with this objective. First and foremost, it allowed us to create a simpler implementation of the interpreter, focusing on the overall result of program execution rather than the intricate details of each computational step. This simplicity facilitated a swift development process and enabled us to build a functional prototype efficiently.

Furthermore, big step semantics provided a higher level of abstraction, making it easier to capture the essential behavior of the language or system being modeled. This abstraction proved beneficial in conceptualizing and formalizing the core functionality of the interpreter, allowing us to abstract away from low-level operational intricacies.

Lastly, the use of big step semantics eased the transition from formalization to code implementation. The high-level nature of big step semantics made it more straightforward to translate our rules into executable code in a functional programming language, establishing a direct mapping between the theoretical model and the practical interpreter.

In summary, the choice of big step semantics was well-suited to our proof-of-concept interpreter project. It facilitated a simpler implementation, provided a higher level of abstraction, and streamlined the process of moving from formalization to executable code.

In Figure 5.9, a snippet of our implementation is presented. Note how we evaluate all necessary expressions recursively in a single step, rather than dividing them into multiple steps. Additionally, our function takes a contract table (ct) and a variable table (vars) as input, both implemented as hash tables. Our configuration consists of a quadruple: the first two elements are identical, with one representing our blockchain and the other a copy of the blockchain before a transaction. The third element is a stack of addresses (sigma), and the last one is an expression to evaluate.

Since blockchain definition is more robust than others, we provide its definition below:

```
1 type contracts = ((values * values), (string * (expr) StateVars.t * values))
2 type accounts = (values, values) Hashtbl.t (* addresses * balance *)
3 type blockchain = (contracts * accounts)
```

Figure 5.8: Blockchain definition in OCaml

As we can see, our blockchain is structured as a tuple, composed by two distinct types: contracts and accounts. The first type is designed to model contracts, and it closely resembles Di Pirro’s formalization of the blockchain. However, we have extended our implementation to include accounts, represented as a hash table that maps each account’s address to their respective balance.

```

1 let rec eval_expr
2   (ct: (string, contract_def) Hashtbl.t)
3   (vars: (string, expr) Hashtbl.t)
4   (conf: conf) : conf
5   =
6   let (blockchain, blockchain', sigma, e) = conf in
7   match e with
8   | Let (_, x, e1, e2) ->
9     begin if Hashtbl.mem vars x then
10      (
11       (blockchain, blockchain', sigma, Revert))
12      else
13       let (_, _, _, e1') = eval_expr ct vars (blockchain, blockchain',
14        sigma, e1) in
15       Hashtbl.add vars x e1'; eval_expr ct vars (blockchain, blockchain',
16        sigma, e2)
17      end
18   | Assign (x, e1) ->
19     let (_, _, _, e1') = eval_expr ct vars (blockchain, blockchain', sigma,
20      e1) in
21     Hashtbl.replace vars x e1';
22     Val VUnit
23   | If (e1, e2, e3) -> let (_, _, _, e1') = eval_expr ct vars
24      (blockchain, blockchain', sigma, e1) in
25      begin match e1' with
26      | Val (VBool b) -> begin match b with
27      | True -> eval_expr ct vars (blockchain, blockchain', sigma, e2)
28      | False -> eval_expr ct vars (blockchain, blockchain', sigma, e3)
29      end
30      | e1' -> assert false
31      end
32   .....

```

Figure 5.9: Snippet of Op Sem Ocaml implementation

As one can observe, we make extensively use of OCaml's pattern matching feature. The first two cases, LET and ASSIGN, deal with local variables. In the former, we check if the name x already exists in our table: if it exists we just return a *Revert*, otherwise it will be identical to ASSIGN, where we add the value evaluated to our hash table and call recursively the function with e_2 as an expression.

In conditional statements, we evaluate the condition: if we get True we return the evaluation of e_2 , if we get False we return the evaluation of e_3 .

PROPERTY-BASED TESTING

6.1 General Context

Property-based testing is a software testing technique that focuses on verifying the correctness of a program by specifying a set of properties or invariants that the program should satisfy. Rather than providing specific input values and expected outcomes, property-based testing tools, such as QuickCheck [0], automatically generate a wide range of test cases based on the defined properties.

While we did not test all of our system properties, the results for the ones we did test were highly encouraging, strongly indicating their correctness and robustness.

In this section, we will provide a brief overview of our accomplishments and outline what should be done in the futures to ensure the correctness of the compiler.

6.2 Arithmetic and Boolean Operations

One notable instance of property-based testing involved the validation of our arithmetic operations by employing the concept of an abelian group. By applying this property, we aimed to ensure that our arithmetic operations adhered to the fundamental principles of associativity, commutativity, and identity elements. In doing so, we sought to confirm that our system's arithmetic operations not only produced accurate results but also exhibited the expected algebraic properties that underpin their behavior.

In figure 6.1 one can see the definition on how our test will build the random expressions, in this case always arithmetic expressions.

Then in figure 6.5 we use these previously pre-built generators and test against three properties: commutative, associative and neutral element.

```

1 let leafgen_int = Gen.oneof[ Gen.map (fun i -> if i > 0 then Val(VUInt i)
   else Val(VUInt 0)) Gen.int]
2 let rec gen_arit_op_ast n = match n with
3   | 0 -> leafgen_int
4   | n -> Gen.oneof [
5     leafgen_int;
6     Gen.map2 (fun l r -> AritOp(Plus(l,r))) (gen_arit_op_ast (n/2))
   (gen_arit_op_ast (n/2));
7     Gen.map2 (fun l r -> AritOp(Div(l,r))) (gen_arit_op_ast (n/2))
   (gen_arit_op_ast (n/2));
8     Gen.map2 (fun l r -> AritOp(Times(l,r))) (gen_arit_op_ast (n/2))
   (gen_arit_op_ast (n/2));
9     Gen.map2 (fun l r -> AritOp(Minus(l,r))) (gen_arit_op_ast (n/2))
   (gen_arit_op_ast (n/2));
10    Gen.map2 (fun l r -> AritOp(Exp(l,r))) (gen_arit_op_ast (n/2))
   (gen_arit_op_ast (n/2));
11    Gen.map2 (fun l r -> AritOp(Mod(l,r))) (gen_arit_op_ast (n/2))
   (gen_arit_op_ast (n/2));
12  ]
13 let arb_tree_arit = make ~print:expr_to_string (gen_arit_op_ast 8)

```

Figure 6.1: Generating random arithmetic expressions

6.3 Conditional Expressions

Additionally, we conducted property-based testing on our system’s handling of conditional expressions, specifically exploring the commutative property of condition reversal. This test suite was designed to verify whether reversing the condition and changing the order of statements within an if expression would yield equal outcomes.

To implement this, we use the following generator:

```

1 let rec gen_if_expr n =
2   match n with
3   | 0 -> Gen.map3 (fun e1 e2 e3 -> If(e1, e2, e3)) (gen_bool_revert_safe_ast
   8) (gen_arit_op_ast 8) (gen_arit_op_ast 8)
4   | n -> Gen.map3 (fun e1 e2 e3 -> If(e1, e2, e3)) (gen_bool_revert_safe_ast
   8) (gen_if_expr (n/2)) (gen_if_expr (n/2))
5 let arb_if_expr = make ~print:expr_to_string (gen_if_expr 8)

```

Figure 6.3: Generating random conditional expressions

The function `gen_bool_revert_safe_ast` returns a random boolean expression that is safe from revert, as including revert would render our tests unable to verify any property. Additionally, we generate e_2 and e_3 as random arithmetic expressions using the previously presented generator.

As shown in Figure 6.5, when testing conditional expressions, we used the following properties:

```

1 let test_arit_plus_op = Test.make ~name:"test plus arithmetic operator"
2   (set_shrink tshrink arb_tree_arit)
3   (fun (e) ->
4     begin
5       let ct = Hashtbl.create 64 in
6       let vars = Hashtbl.create 64 in
7       let blockchain = (Hashtbl.create 64, Hashtbl.create 64) in
8       let sigma = Stack.create() in
9       let e' = AritOp(Plus(e,e)) in
10      (
11        eval_expr ct vars (blockchain, blockchain, sigma,
12          (AritOp(Plus(e',e))))
13      =
14        eval_expr ct vars (blockchain, blockchain, sigma,
15          (AritOp(Plus(e,e'))))
16      )
17      &&
18      (
19        eval_expr ct vars (blockchain, blockchain, sigma,
20          (AritOp(Plus((AritOp(Plus(e,e'))), e))))
21      =
22        eval_expr ct vars (blockchain, blockchain, sigma,
23          (AritOp(Plus((AritOp(Plus(e,e))), e'))))
24      )
25      &&
26      (
27        eval_expr ct vars (blockchain, blockchain, sigma,
28          (AritOp(Plus(Val(VUInt 0),e))))
29      =
30        eval_expr ct vars (blockchain, blockchain, sigma, e)
31      )
32    end
33  )

```

Figure 6.2: Plus operation property-based testing

$$\text{if false then } e_1 \text{ else } e_2 = \text{if true then } e_2 \text{ else } e_1$$

$$\text{if } b \text{ then } e_1 \text{ else } e_2 = \text{if } \neg b \text{ then } e_2 \text{ else } e_1$$

Essentially, reversing the condition and the order of the statements in a conditional expression results in an equivalent expression.

6.4 Future Work

As we have seen, leveraging the power of property-based testing is highly effective in ensuring the correctness of a given system. This approach enables comprehensive validation of the system's behavior across a wide range of scenarios, with a particular emphasis on critical properties essential for correctness and reliability.

The outcomes of our tests proved to be not only highly encouraging but also compellingly indicative of the correctness and robustness of select operational semantics rules.

```

1 let test_if = Test.make ~name:"test if operator"
2   (set_shrink tshrink arb_if_expr)
3   (fun (e) ->
4     begin
5       let ct = Hashtbl.create 64 in
6       let vars = Hashtbl.create 64 in
7       let blockchain = (Hashtbl.create 64, Hashtbl.create 64) in
8       let sigma = Stack.create() in
9       let (e1, e2, e3) = match e with
10        | If (e1, e2, e3) ->
11          let (_, _, _, e1') = eval_expr ct vars (blockchain, blockchain,
12            sigma, e1) in
13            if e1' <> Revert then (e1, e2, e3) else (Val(VBool False), e2,
14              e3)
15        | _ -> assert false
16      in
17        (
18          eval_expr ct vars (blockchain, blockchain, sigma, (If(Val(VBool
19            True), e2, e3)))
20          =
21          eval_expr ct vars (blockchain, blockchain, sigma, e2)
22        )
23        &&
24        (
25          eval_expr ct vars (blockchain, blockchain, sigma, (If(Val(VBool
26            False), e2, e3)))
27          =
28          eval_expr ct vars (blockchain, blockchain, sigma, e3)
29        )
30        &&
31        (
32          eval_expr ct vars (blockchain, blockchain, sigma, (If(e1, e2, e3)))
33          =
34          eval_expr ct vars (blockchain, blockchain, sigma,
35            (If(BoolOp(Neg(e1)), e3, e2)))
36        )
37        &&
38        (
39          let (_, _, _, e') = eval_expr ct vars (blockchain, blockchain,
40            sigma, (If(e1, e2, e3))) in
41          let (_, _, _, e2') = eval_expr ct vars (blockchain, blockchain,
42            sigma, e2) in
43          let (_, _, _, e3') = eval_expr ct vars (blockchain, blockchain,
44            sigma, e3) in
45          (e' = e2') || (e' = e3')
46        )
47      end
48    )
49  )

```

Figure 6.4: Conditional expressions property-based testing

it is important to note that while we successfully applied these tests to some rules, for an industrial tool we advise extending property-based testing to encompass all expressions within a system. By doing so, one can systematically evaluate the behavior and consistency of every aspect of the system, thus reinforcing its overall reliability and adherence to established properties.

```
qcheck random seed: 508888590
Testing `FS Operational Semantics Tests'.
This run has ID `UY4FU4R1'.

[OK] suite      0 test division by zero.
[OK] suite      1 test arithmetic operators.
[OK] suite      2 test boolean operators.
[OK] suite      3 test if operator.
[OK] suite      4 test let operator.
[OK] suite      5 test assign operator.
[OK] suite      6 test add contract to contract table.
[OK] suite      7 test new contract.
[OK] suite      8 test plus arithmetic operator.
[OK] suite      9 test minus arithmetic operator.
[OK] suite     10 test times arithmetic operator.
[OK] suite     11 test div arithmetic operator.
[OK] suite     12 test balance function.
[OK] suite     13 test address function.
[OK] suite     14 test state read.
[OK] suite     15 test state write.
[OK] suite     16 test typecheck arithmetic operations.
[OK] suite     17 test typecheck boolean operations.

Full test results in `~/featherweight-solidity/_build/default/test/_build/_tests/FS Operational Semantics Tests'.
Test Successful in 0.301s. 18 tests run.
```

Figure 6.5: Property-based testing results

In the figure above, we can see the output we get after running *dune runtest* in the project root.

FS 2.0: TYPE SYSTEM

7.1 General Context

In this section, we delve into the core of our study: the type system rules in Featherweight Solidity 2.0. A well-defined type system is the backbone of any programming language, governing how values and expressions are classified, checked, and enforced during program execution. These rules form the foundation upon which developers build robust, error-free software.

Our goal here is to elucidate the intricacies of the type system in FS 2.0, providing a comprehensive understanding of how types are assigned, inferred, and verified within the language. In particular, we explain in detail, how do we deal with multiple inheritance and subtyping in the type address. This exploration will lay the groundwork for subsequent discussions on type checking, type inference, and ultimately, the assurance of program correctness.

Next, we present the fundamental type system rules that shape the programming experience in FS 2.0.

7.2 Type Environment

Our environment differs slightly from Di Pirro's in that we incorporate an additional component to store the types of each state variable. In our model, we assume that a contract's state variables are only accessible internally, through the **this** keyword. In Solidity, state variables are also private by default, but when declared as **public**, the language automatically generates a getter function for external access. In our language, you need to explicitly create it.

$$(Type\ Environment)\ \Gamma ::= \emptyset \mid \Gamma, x : T \mid \Gamma, \mathbf{this}.s : T \mid \Gamma, a : \mathbf{address} \mid \Gamma, c : C$$

Figure 7.1: Type Environment of Featherweight Solidity 2.0

7.3 Rules

We will start by presenting the type system rules without inheritance in section 7.3.1. Since inheritance is only possible in two types: contracts and addresses, some rules will not change between both formalization and we chose to omit to avoid duplicating. Most of the rules, only need a less strict type matching: if we are expecting an expression with type T , but we instead get T' , if T' is a subtype of T , it is still valid. The rule to type check a cast operation is the one that will change the most.

7.3.1 Rules Without Inheritance

Our presentation begins with the foundational axioms that underpin our type system. These axioms serve as the cornerstone, shaping the system's core principles.

Next, we systematically categorize rules based on their similarity:

- Arithmetic and Boolean Operations - Emphasizing mathematical and logical operations.
- Variable Operations - Covering state and local variables.
- Conditional and Sequential Statements - If statements and sequence of expressions.
- Function-like Operations - Including castings, function calls, top-level calls, and balance calculations.

This classification framework provides a structured approach to explore our type system comprehensively.

7.3.1.1 Axioms

$$\begin{array}{c} \text{(REF)} \\ \frac{\Gamma, c : C, \Gamma' \vdash \langle \rangle}{\Gamma, c : C, \Gamma' \vdash c : C} \end{array}$$

$$\begin{array}{c} \text{(VAR)} \\ \frac{\Gamma, x : T, \Gamma' \vdash \langle \rangle}{\Gamma, x : T, \Gamma' \vdash x : T} \end{array}$$

$$\begin{array}{c} \text{(TRUE)} \\ \frac{\Gamma \vdash \langle \rangle}{\Gamma \vdash \text{true} : \text{bool}} \end{array}$$

$$\frac{\text{(FALSE)} \quad \Gamma \vdash \langle \rangle}{\Gamma \vdash \text{false} : \text{bool}}$$

$$\frac{\text{(UNIT)} \quad \Gamma \vdash \langle \rangle}{\Gamma \vdash u : \text{unit}}$$

$$\frac{\text{(REVERT)} \quad \Gamma \vdash \langle \rangle}{\Gamma \vdash \text{revert} : T}$$

$$\frac{\text{(UINT)} \quad n \in \mathbb{N}_0 \quad \Gamma \vdash \langle \rangle}{\Gamma \vdash n : \text{uint}}$$

$$\frac{\text{(ADDRESS)} \quad \Gamma, a : \text{address}, \Gamma' \vdash \langle \rangle}{\Gamma, a : \text{address}, \Gamma' \vdash a : \text{address}}$$

These rules are equivalent to Di Pirro, formalization, and so there is no much to say.

7.3.1.2 Arithmetic and Boolean

$$\frac{\text{(PLUS/MINUS/DIV/MULT/EXP/MOD)} \quad \Gamma \vdash e_1 : \text{uint} \quad \Gamma \vdash e_2 : \text{uint}}{\Gamma \vdash e_1 \diamond e_2 : \text{uint}}$$

$$\frac{\text{(EQUALS/INEQUALS/GREATER/LESSER)} \quad \Gamma \vdash e_1 : \text{uint} \quad \Gamma \vdash e_2 : \text{uint}}{\Gamma \vdash e_1 \diamond e_2 : \text{bool}}$$

$$\frac{\text{(EQUALS/INEQUALS)} \quad \Gamma \vdash e_1 : \text{address} \quad \Gamma \vdash e_2 : \text{address}}{\Gamma \vdash e_1 \diamond e_2 : \text{bool}}$$

$$\begin{array}{c}
(\text{CONJ/DISJ}) \\
\frac{\Gamma \vdash e_1 : \text{bool} \quad \Gamma \vdash e_2 : \text{bool}}{\Gamma \vdash e_1 \diamond e_2 : \text{bool}}
\end{array}$$

These rules establish criteria for well-typed expressions involving arithmetic and Boolean operators supported by our programming language. It is essential to note that our equality and inequality operations support comparisons between address and uint data types.

7.3.1.3 Variables

$$\begin{array}{c}
(\text{LET}) \\
\frac{\Gamma \vdash e_1 : T_1 \quad \Gamma, x : T_1 \vdash e_2 : T_2}{\Gamma \vdash T_1 \ x = e_1; e_2 : T_2}
\end{array}$$

$$\begin{array}{c}
(\text{ASSIGN}) \\
\frac{\Gamma \vdash x : T \quad \Gamma \vdash e : T}{\Gamma \vdash x = e : T}
\end{array}$$

$$\begin{array}{c}
(\text{STATEASSIGN}) \\
\frac{\Gamma \vdash e_1.s : T \quad \Gamma \vdash e_2 : T}{\Gamma \vdash e_1.s = e_2 : T}
\end{array}$$

$$\begin{array}{c}
(\text{STATEREAD}) \\
\frac{}{\Gamma \vdash e_1.s_i : T_i}
\end{array}$$

These rules handle either local variables (ASSIGN and LET) or state variables (STATEASSIGN and STATEREAD) and closely resemble the corresponding rules we previously introduced in Chapter 4. The only difference here is that, since we store our state variable types in our type environment, the STATEREAD rule is straightforward, and we can directly retrieve its value from our Γ .

7.3.1.4 Constructor

$$\begin{array}{c}
(\text{RETRIEVE CONSTRUCTOR PARAMETERS TYPES}) \\
\frac{SC = \text{contract } C \text{ is } \widetilde{D} \{(\widetilde{T} \ s) \ K \ \widetilde{F}\} \quad K = \text{constructor } ((\widetilde{T} \ x))(e)}{\text{ctypes}(SC) = \widetilde{T}}
\end{array}$$

$$\begin{array}{c}
\text{(NEW)} \\
\frac{\Gamma \vdash e' : \text{uint} \quad \Gamma \vdash \tilde{e} : \tilde{T} \quad \text{ctypes}(C) = \tilde{T}}{\Gamma \vdash \text{new } C.\text{value}(e')(\tilde{e}) : C}
\end{array}$$

To check if a statement with NEW is well-typed we need an auxiliary predicate, which appears first, where we retrieve for a contract its constructor parameters type (\tilde{T}). Then in the rule, we check that e' is an unsigned integer and check that argument in \tilde{e} respect the expected type in \tilde{T} .

7.3.1.5 Conditional and Sequential Statements

$$\begin{array}{c}
\text{(SEQ)} \\
\frac{\Gamma \vdash e_1 : T_1 \quad \Gamma \vdash e_2 : T_2}{\Gamma \vdash e_1; e_2 : T_2}
\end{array}$$

This rule ensures that e_1 is well-typed and has type T_1 and e_2 typechecks successfully with T_2 .

$$\begin{array}{c}
\text{(IF)} \\
\frac{\Gamma \vdash e_1 : \text{bool} \quad \Gamma \vdash e_2 : T \quad \Gamma \vdash e_3 : T}{\Gamma \vdash \text{if } e_1 \text{ then } e_2 \text{ else } e_3 : T}
\end{array}$$

This follows the traditional expect type check rule for conditional statements, e_1 has the condition must have a boolean type and e_2 and e_3 should have the same type.

7.3.1.6 Mappings

$$\begin{array}{c}
\text{(MAPPING)} \\
\frac{M = \{\widetilde{(k, v)}\} \quad \Gamma \vdash \tilde{k} : T_1 \quad \Gamma \vdash \tilde{v} : T_2}{\Gamma \vdash M : \text{mapping}(T_1 \Rightarrow T_2)}
\end{array}$$

$$\begin{array}{c}
\text{(MAPASSIGN)} \\
\frac{\Gamma \vdash e_1 : \text{mapping}(T_1 \rightarrow T_2) \quad \Gamma \vdash e_2 : T_1 \quad \Gamma \vdash e_3 : T_2}{\Gamma \vdash e_1[e_2 \rightarrow e_3] : \text{mapping}(T_1 \Rightarrow T_2)}
\end{array}$$

$$\begin{array}{c}
\text{(MAPREAD)} \\
\frac{\Gamma \vdash e_1 : \text{mapping}(T_1 \rightarrow T_2) \quad \Gamma \vdash e_2 : T_1}{\Gamma \vdash e_1[e_2] : T_2}
\end{array}$$

For mappings we have three rules. The first one, checks if a mapping is well-typed, by checking that each key k and each value v has type T_1 and T_2 respectively. When we are

assigning a value to a key in a certain expression e_1 , we check that e_1 is a mapping, e_2 has type T_1 and e_3 has type T_2 . When we are reading a value, we need to infer that e_1 can be typed to a mapping and check that e_2 has type T_1 , which is the type of keys of e_1 .

7.3.1.7 Functions

$$\begin{array}{c} \text{(BALANCE)} \\ \Gamma \vdash e : \text{address} \\ \hline \Gamma \vdash \text{balance}(e) : \text{uint} \end{array}$$

$$\begin{array}{c} \text{(ADDR)} \\ \Gamma \vdash e : C \\ \hline \Gamma \vdash \text{address}(e) : \text{address} \end{array}$$

$$\begin{array}{c} \text{(RETURN)} \\ \Gamma \vdash e : T \\ \hline \Gamma \vdash \text{return } e : T \end{array}$$

$$\begin{array}{c} \text{(CONTRRETR)} \\ \Gamma \vdash e : \text{address} \\ \hline \Gamma \vdash C(e) : C \end{array}$$

$$\begin{array}{c} \text{(TRANSFER)} \\ \Gamma \vdash e_1 : \text{address}\langle C \rangle \quad \text{ftype}(C, \text{fb}) = \{ \} \rightarrow \text{unit} \quad \Gamma \vdash e_2 : \text{uint} \\ \hline \Gamma \vdash e_1.\text{transfer}(e_2) : \text{unit} \end{array}$$

$$\begin{array}{c} \text{(CALLTOPLEVEL)} \\ \text{ftype}(C, f) = \tilde{T}_1 \rightarrow T_2 \quad \Gamma \vdash e_1 : C \quad \Gamma \vdash e_2 : \text{uint} \\ |\tilde{e}| = |\tilde{T}_1| \quad \Gamma \vdash e_3 : \text{address} \quad \Gamma \vdash \tilde{e} : \tilde{T}_e \quad \tilde{T}_e = \tilde{T}_1 \\ \hline \Gamma \vdash e_1.f.\text{value}(e_2).\text{sender}(e_3)(\tilde{e}) : T_2 \end{array}$$

$$\begin{array}{c} \text{(CALL)} \\ \Gamma \vdash e_1 : C \\ \Gamma \vdash e_2 : \text{uint} \quad \text{ftype}(C, f) = \tilde{T}_1 \rightarrow T_2 \quad |\tilde{e}| = |\tilde{T}_1| \quad \Gamma \vdash \tilde{e} : \tilde{T}_e \quad \tilde{T}_e = \tilde{T}_1 \\ \hline \Gamma \vdash e_1.f.\text{value}(e_2)(\tilde{e}) : T_2 \end{array}$$

$$\begin{array}{c}
 \text{(FUN)} \\
 \frac{\Gamma \vdash e_1 : \mathbf{C} \quad \text{ftype}(C, f) = \tilde{T}_1 \rightarrow T_2}{\Gamma \vdash c.f : \tilde{T}_1 \rightarrow T_2}
 \end{array}$$

$$\begin{array}{c}
 \text{(CALLVALUE)} \\
 \frac{\Gamma \vdash \tilde{T}_1 \rightarrow T_2 \quad \Gamma \vdash e_2 : \text{uint} \quad \Gamma \vdash \tilde{e}_1 : \tilde{T}_1 \quad |\tilde{e}| = \tilde{T}_1 \quad \Gamma \vdash \tilde{e} : \tilde{T}_e \quad \tilde{T}_e = \tilde{T}_1}{\Gamma \vdash e_1.\text{value}(e_2)(\tilde{e}) : T_2}
 \end{array}$$

Excluding inheritance, these rules are straightforward and identical to Di Pirro's formalization. We will delve deeper into CALLTOPLEVEL, CALL, CONTRRETR, FUN, and CALLVALUE as they are the ones that differ the most from either without inheritance rules or Di Pirro's rules.

7.3.2 Rules With Inheritance

In this section, we will delve into the concept of inheritance within our rule framework. It is crucial to emphasize the existence of a top type for contracts, denoted as CTOP. As such, it operates akin to the Object class in Java. However, users cannot declare this type themselves, as it serves as an internal representation encompassing all contracts.

$$\text{issc}(C_1, C_2, CT) = C_2 \in \text{c3linearization}(CT(C_1)) \vee C_2 = \text{CTOP}$$

$$\text{subtyping}(T_1, T_2, CT) = \begin{cases} \text{issc}(C_1, C_2, CT) & T_1 = C_1 \wedge T_2 = C_2 \\ \text{issc}(C, \text{CTOP}, CT) & T_1 = \text{address}\langle C \rangle \wedge T_2 = \text{address} \\ \text{issc}(\text{CTOP}, C, CT) & T_1 = \text{address} \wedge T_2 = \text{address}\langle C \rangle \\ \text{subtyping}(C_1, C_2, CT) & T_1 = \text{address}\langle C_1 \rangle \wedge T_2 = \text{address}\langle C_2 \rangle \\ T_1 = T_2 & \text{otherwise} \end{cases}$$

This algorithm serves as our method to determine if two types, denoted as T_1 and T_2 , exhibit a subtype relationship. It is essential to note that the order in which we conduct these comparisons holds significance because our goal is to establish whether T_1 is either identical or a subtype of T_2 . Our type hierarchy incorporates inheritance exclusively for contract and address types; therefore, for all other types, strict equality suffices.

To qualify as a subtype of contract C_t , one of two conditions must be met: either C_t is equal to the top-level contract, denoted as CTOP, or it is part of the linearization of the contract type under consideration. In the case of address types, the determination of subtype status hinges on comparing the enclosing annotation contract of the address type with the other type or, equivalently, assessing if the generic address is akin to the CTOP contract.

7.3.2.1 Axioms

$$\begin{array}{c}
\text{(ADDRESS)} \\
\frac{\Gamma, a : \text{address}\langle C \rangle, \Gamma' \vdash \langle \rangle}{\Gamma, a : \text{address}\langle C \rangle, \Gamma' \vdash a : \text{address}\langle C \rangle}
\end{array}$$

Just one axiom change to support inheritance type, since we introduce a new subtyping address and so every address should be annotated with it is contract type. For $C = \text{TOP}$ we have the generic type address, that already exists in Solidity.

7.3.2.2 Adapted Rules

$$\begin{array}{c}
\text{(LET)} \\
\frac{\Gamma \vdash e_1 : T'_1 \quad \text{subtyping}(T'_1, T_1, CT) \quad \Gamma, x : T_1 \vdash e_2 : T_2}{\Gamma \vdash T_1 \ x = e_1; e_2 : T_2}
\end{array}$$

$$\begin{array}{c}
\text{(ASSIGN)} \\
\frac{\Gamma \vdash x : T \quad \Gamma \vdash e : T_e \quad \text{subtyping}(T_e, T, CT)}{\Gamma \vdash x = e : T_e}
\end{array}$$

$$\begin{array}{c}
\text{(STATEASSIGN)} \\
\frac{\Gamma \vdash e_1.s : T \quad \Gamma \vdash e_2 : T' \quad \text{subtyping}(T', T, CT)}{\Gamma \vdash e_1.s = e_2 : T'}
\end{array}$$

$$\begin{array}{c}
\text{(IF)} \\
\frac{\Gamma \vdash e_1 : \text{bool} \quad \Gamma \vdash e_2 : T_1 \quad \Gamma \vdash e_3 : T_2 \quad \text{subtyping}(T_1, T, CT) \quad \text{subtyping}(T_2, T, CT)}{\Gamma \vdash \text{if } e_1 \text{ then } e_2 \text{ else } e_3 : T}
\end{array}$$

$$\begin{array}{c}
\text{(MAPPING)} \\
\frac{M = \{\widetilde{(k, v)}\} \quad \Gamma \vdash \widetilde{k} : T_k \quad \Gamma \vdash \widetilde{v} : T_v \quad \text{subtyping}(T_k, T_1, CT) \quad \text{subtyping}(T_v, T_2, CT)}{\Gamma \vdash M : \text{mapping}(T_1 \Rightarrow T_2)}
\end{array}$$

$$\begin{array}{c}
\text{(MAPASSIGN)} \\
\frac{\Gamma \vdash e_1 : \text{mapping}(T_1 \rightarrow T_2) \quad \Gamma \vdash e_2 : T'_1 \quad \Gamma \vdash e_3 : T'_2 \quad \text{subtyping}(T'_1, T_1, CT) \quad \text{subtyping}(T'_2, T_2, CT)}{\Gamma \vdash e_1[e_2 \rightarrow e_3] : \text{mapping}(T_1 \Rightarrow T_2)}
\end{array}$$

$$\begin{array}{c}
\text{(MAPREAD)} \\
\frac{\Gamma \vdash \mathbf{e1} : \text{mapping}(\mathbf{T}_1 \rightarrow \mathbf{T}_2) \quad \Gamma \vdash e2 : T'_1 \quad \text{subtyping}(T'_1, T_1, CT)}{\Gamma \vdash e1[e2] : T_2}
\end{array}$$

$$\begin{array}{c}
\text{(TRANSFER)} \\
\frac{\Gamma \vdash \mathbf{e1} : \text{address}\langle C \rangle \quad \text{ftype}(C, \text{fb}) = \{\} \rightarrow \text{unit} \quad \Gamma \vdash e2 : \text{uint}}{\Gamma \vdash e1.\text{transfer}(e2) : \text{unit}}
\end{array}$$

In all of these rules, we adjusted them, relaxing the condition where we get T' , but instead we expected a type T . In these cases, we just need to call our algorithm of subtyping and ensure that T' is either equal or a subtype of T : $\text{subtyping}(T', T, CT)$

7.3.2.3 Rules Changing the Most

$$\begin{array}{c}
\text{(CONTRRETR)} \\
\frac{\Gamma \vdash e : \text{address}\langle C \rangle \quad \text{subtyping}(C, D, CT)}{\Gamma \vdash D(e) : D}
\end{array}$$

With inheritance and subtyping, we are more restrictive in our casts. First our address is either annotated with a contract or it is a generic address. If it's a generic address it cannot be casted to a contract, since the result for $\text{subtyping}(CTOP, D, CT)$ will always be false. This is a good improvement from the current Solidity version, where we can typecast every address, even addresses referencing EOAs. However if we have an expression annotated with an address subtype, we can cast this address to every contract that is in the linearization of the said contract, as a result of the premise using subtyping function. We consider our address type safe, making our language type safe, contrary to Solidity.

$$\begin{array}{c}
\text{(CALLTOPLEVEL)} \\
\frac{\Gamma \vdash \mathbf{e1} : \mathbf{C} \quad \Gamma \vdash e2 : \text{uint} \quad \text{ftype}(C, f) = \tilde{T}_1 \rightarrow T_2 \quad |\tilde{e}| = |\tilde{T}_1| \quad \Gamma \vdash e3 : \text{address} \quad \Gamma \vdash \tilde{e} : \tilde{T}_e \quad \text{subtyping}(\tilde{T}_e, \tilde{T}_1, CT)}{\Gamma \vdash e1.f.\text{value}(e2).\text{sender}(e3)(\tilde{e}) : T_2}
\end{array}$$

As mentioned earlier, CALLTOPLEVEL is exclusively called by externally owned accounts (EOAs), ensuring that the value of $e3$ is always an address. Therefore, we can dispense with the premise $\text{subtyping}(C', \text{fsender}(C, f), CT)$ since we do not categorize external-owned accounts as contracts. Additionally, by default, the functions that can be invoked by this constructor should be those without any annotations, expecting only a generic address. The expression $\text{subtyping}(\tilde{T}_e, \tilde{T}_1, CT)$ is essentially a syntactic shortcut for applying our subtyping function to each element in \tilde{T}_e and \tilde{T}_1 individually:

$$\text{subtyping}(\tilde{T}_e, \tilde{T}_1, CT) \equiv \bigwedge_{i=1}^n \text{subtyping}(T_{e_i}, T_{1_i}, CT)$$

(CALL)

$$\frac{\Gamma \vdash e_1 : \mathbf{C} \quad \Gamma \vdash e_2 : \text{uint} \quad \text{ftype}(C, f) = \tilde{T}_1 \rightarrow T_2 \quad |\tilde{e}| = |\tilde{T}_1| \quad \Gamma \vdash \text{this} : C' \quad \text{subtyping}(C', \text{fsender}(C, f), CT) \quad \Gamma \vdash \tilde{e} : \tilde{T}_e \quad \text{subtyping}(\tilde{T}_e, \tilde{T}_1, CT)}{\Gamma \vdash e_1.f.\text{value}(e_2)(\tilde{e}) : T_2}$$

(FUN)

$$\frac{\Gamma \vdash e_1 : \mathbf{C} \quad \text{ftype}(C, f) = \tilde{T}_1 \rightarrow T_2 \quad \Gamma \vdash \text{this} : C' \quad \text{subtyping}(C', \text{fsender}(C, f), CT)}{\Gamma \vdash c.f : \tilde{T}_1 \rightarrow T_2}$$

(CALLVALUE)

$$\frac{\Gamma \vdash e_2 : \text{uint} \quad \tilde{T}'_1 <: T_1 \quad \Gamma \vdash \tilde{T}_1 \rightarrow T_2 \quad |\tilde{e}| = |\tilde{T}_1| \quad \Gamma \vdash \tilde{e} : \tilde{T}_e \quad \text{subtyping}(\tilde{T}_e, \tilde{T}_1, CT)}{\Gamma \vdash e_1.\text{value}(e_2)(\tilde{e}) : T_2}$$

These rules are designed to ensure well-typed function calls within contracts. They require that function calls are either unannotated or annotated with a contract that is either the contract from which we are invoking the function or a supertype of it.

Below, we present an OCaml snippet that demonstrates how we did implement the premises:

$$|\tilde{e}| = |\tilde{T}_1| \wedge \text{subtyping}(\tilde{T}_e, \tilde{T}_1, CT)$$

In this snippet, we leverage the power of OCaml's `iter2` function from the `List` library. This function allows us to iterate over two lists simultaneously. If they don't have the same length, it throws an exception. Inside this iterator, we check whether each element in \tilde{e} has the expected type according to the parameters, recursively calling our type-checking function.

```

1 let fun_check gamma cname fun_name le ct t =
2   let ftype = function_type cname fun_name ct in
3   let (t_es, rettype) = ftype in
4   if subtyping rettype t ct then
5     try
6       List.iter2 (fun t_e e' -> typecheck gamma e' t_e ct;) t_es le
7     with Invalid_argument _ ->
8       raise (Failure "Invalid arguments")
9   else
10    raise (TypeMismatch (rettype, t))
11 in

```

Figure 7.2: OCaml implementation type checking function calls

7.3.3 Type System by Example

In this section, we will use the contract example introduced previously, which was also utilized in the examples demonstrating the application of semantic operational rules (Figure 5.7). In the first example, we will present the rules applied when verifying the `setNFTPrice` function:

$$\text{IF} \frac{\text{---} \quad \text{---} \quad \text{---}}{T_1 \ T_2 \ T_3}}{\Gamma \vdash \text{if}(\text{msgsender} == \text{this.owner}) \text{ then } \text{this.price} = \text{price} \text{ else } \text{revert} : \text{unit}}$$

We begin by applying the IF rule to decompose our expression into simpler expressions. This process allows us to eventually validate it using an axiom or identify a point where we cannot proceed, resulting in a poorly typed expression. Applying this rule we obtain three sub trees which we show next:

$$\text{UINT} \frac{\frac{\text{---} \quad \text{---}}{\Gamma \vdash \text{this.price} : \text{uint}} \quad \frac{\text{---}}{\Gamma \vdash \text{price} : \text{uint}}}{\Gamma \vdash \text{this.price} = \text{price} : \text{unit}} \quad \text{UINT} \quad \text{STATEASSIGN}$$

$$\text{REVERT} \frac{\frac{\text{---}}{\Gamma \vdash \text{revert} : \text{unit}}}{\Gamma \vdash \text{revert} : T'}$$

$$\text{ADDRESS} \frac{\frac{\text{---} \quad \text{---}}{\Gamma \vdash \text{msgsender} : \text{address}} \quad \frac{\text{---}}{\Gamma \vdash \text{this.owner} : \text{address}}}{\Gamma \vdash \text{msgsender} == \text{this.owner} : \text{bool}} \quad \text{ADDRESS} \quad \text{EQUALS}$$

$\Gamma = [\text{this.price} \mapsto \text{uint}; \text{price} \mapsto \text{uint}; \text{this.owner} \mapsto \text{address}; \text{msgsender} \mapsto \text{address}; \text{this.totalNfts} \mapsto \text{uint}; \text{this.nfts} \mapsto \text{mapping}(\text{uint} \rightarrow \text{address})]$

Note that in all rules, we obtain an axiom that allows us to validate that our expressions are well-typed. The REVERT rule is particularly interesting, as it is well-typed regardless of the type T' we need to achieve. This is crucial since revert is often used to abort a transaction with an error and revert to the blockchain initial state. Consequently, it can be employed in conditional statements if the desired evaluation is not achieved, necessitating the use of REVERT with other supported types.

We also want to note that we omit the premises using the subtyping function. However, assume that after this derivation tree, there are two additional steps to verify:

$\text{subtyping}(\text{unit}, \text{unit}, CT)$

$\text{subtyping}(\text{unit}, \text{unit}, CT)$

Both of which obviously evaluate to true, since this expression is also well-typed if we did not have inheritance.

Now we see how we can apply our type system rules to successfully verify that `getOwner` function is correctly typed:

$$\begin{array}{c}
 \text{STATEREAD} \frac{\frac{\frac{}{\Gamma \vdash \text{this.nfts} : \text{mapping}(\text{uint} \rightarrow \text{address})}{} \quad \frac{\frac{}{\Gamma \vdash \text{nft} : \text{uint}}{} \quad \text{UINT}}{\text{subtyping}(\text{uint}, \text{uint}, CT)}}{\Gamma \vdash \text{this.nfts}[\text{nft}] : \text{address}}}{\Gamma \vdash \text{return this.nfts}[\text{nft}] : \text{address}} \\
 \text{RETURN} \frac{}{}
 \end{array}
 \text{MAPREAD}$$

$\Gamma = [\text{this.price} \mapsto \text{uint}; \text{nft} \mapsto \text{uint}; \text{this.owner} \mapsto \text{address}; \text{msgsender} \mapsto \text{address}; \text{this.totalNfts} \mapsto \text{uint}; \text{this.nfts} \mapsto \text{mapping}(\text{uint} \rightarrow \text{address})]$

7.3.4 Implementation in OCaml

Our approach to implementing the type system rules in OCaml leverage mutually recursive functions: `infer_type` and `typecheck` are mutually recursive (Figure 7.5). The `infer_type` function was necessary because some rules had premises that required inferring types (indicated in bold on the rules). However, to infer a type from an expression, we also needed to ensure that the expression was well-composed and correctly typed. Both functions took as parameters a gamma (*type environment*), an expression e , and a contract table.

```
1 type gamma_vars = (string, t_exp) Hashtbl.t
2 type gamma_state_vars = (string, t_exp) Hashtbl.t
3 type gamma_addresses = (values, t_exp) Hashtbl.t
4 type gamma_contracts = (values, t_exp) Hashtbl.t
5 type gamma = (gamma_vars * gamma_state_vars * gamma_addresses *
  gamma_contracts)
```

Figure 7.3: Gamma definition in OCaml

Since every expression e will most of the times reduce to an axiom, note how we have a function `typecheck_axioms` (Figure 7.4) where we enforce that the type of an axiom, t_e is either `TRevert` (line 7) or it evaluates to true calling the function: `subtyping(t_e , t , CT)` (line 8). This ensures inheritance is applied across most of the rules.

```
1
2 let typecheck_axioms (gamma: gamma) (e: expr) (t: t_exp) (ct:
  contract_table): unit =
3 let t_e = axioms gamma e in
4 begin match t_e with
5 | Ok(t_e) ->
6   begin match t_e with
7   | TRevert -> ()
8   | _ -> (if subtyping t_e t ct then () else (raise (TypeMismatch (t_e,
  t))))
9   (* | _ -> if t_e = t then () else raise (TypeMismatch (t_e, t)) *)
10 end
11 | Error(s) -> raise (Failure s)
12 end
```

Figure 7.4: Typecheck axioms in OCaml

In the `typecheck` function, we introduced an additional argument t , representing the expected type for an expression e . While `infer_type` returns either a type or an error, `typecheck` returns either a *Failure* or an *Exception*.

```

1 let rec infer_type (gamma: gamma) (e: expr) (ct: contract_table) : (t_exp,
  string) result =
2   match e with
3   | Val _ | Var _ | Revert | This None | MsgSender | MsgValue -> axioms gamma
  e
4   | Balance e1 -> typecheck gamma (Balance(e1)) UInt ct; Ok(UInt)
5   | Assign (s, e1) -> typecheck gamma (Assign(s, e1)) Unit ct; Ok(Unit)
6   | Address e1 -> typecheck gamma (Address e1) (Address None) ct; Ok(Address
  None)
7   | Return e1 -> let t_e1 = infer_type gamma e1 ct in t_e1
8   | If (e1, e2, e3) ->
9     begin
10      typecheck gamma e1 Bool ct; (* checks if condition is a boolean *)
11      let t_e2 = infer_type gamma e2 ct in (* infers e2 type *)
12      match t_e2 with (*if it can be inferred, checks e3 have the same type*)
13      | Ok(t_e2) -> typecheck gamma e3 t_e2 ct; Ok(t_e2)
14      | Error s -> Error s
15    end
16   ....
17 and typecheck (gamma: gamma) (e: expr) (t: t_exp) (ct: contract_table) : unit
  =
18   match e with
19   | Call (e1, s, e2, le) ->
20     begin
21       let t_e1 = infer_type gamma e1 ct in
22       match t_e1 with
23       | Ok(C name) ->
24         let (gv, _, _, _) = gamma in
25         let t_sender = match fsender name s ct with
26         | Ok t_sender -> t_sender
27         | Error s -> raise (Failure s)
28         in
29         Hashtbl.add gv "msg.sender" t_sender;
30         typecheck gamma e2 UInt ct;
31         fun_check gamma name s le ct t;
32       | Ok(CTop) -> raise (Failure "Can't reference a top class")
33       | Ok(t_e1) -> raise (TypeMismatch (t_e1, CTop))
34       | Error s -> raise (Failure s)
35     end
36   ....

```

Figure 7.5: Snippet of type check implementation

In line 4, we observe that we return the type `uint` if and only if the expression `balance(e)` is well-typed. The same principle applies to the `if` expression (lines 8 to 15). Initially, we check whether e_1 is a boolean expression. Subsequently, we infer the type of e_2 , and if e_3 has the same type as e_2 , we return the type of both as the result. Otherwise, we return an *Error*.

In the `typecheck` function, we can observe the implementation of type checking for the `CALL` rule. Firstly, in line 21, we infer the type of e_1 . Subsequently, we match this type with the possible values we can obtain (lines 22 to 35). If we obtain anything other than an expression typed with a contract, we raise a *Failure*. Otherwise, we proceed by storing the sender type annotation. If the sender type annotation is hidden, the sender type defaults

to just an address (lines 25 to 29). Following this, we ensure that e_2 has the type `uint` and finally call the previously presented function, `fun_check` (Figure 7.2), checking that each argument has a type that is equal or a subtype of the corresponding parameter type.

8.1 General Context

In this section we will present our main case study during our implementation. The contract which will be the one interacted by EOAs is in Figure 5.3. The following contracts are all used to build our Game contract:

```
1 contract Context {  
2  
3     constructor() {}  
4  
5     function fb() {}  
6     function msgSender() returns (address) {  
7         return msgsender;  
8     }  
9 }
```

Figure 8.1: Context Contract

```
1 import "./Context.sol";
2
3 contract Ownable is Context {
4     address owner;
5
6     constructor(address initialOwner) Context() {
7         this.owner = initialOwner;
8     }
9
10
11     function fb() {}
12     function checkOwner() {
13         if (this.owner != this.msgSender()) {
14             revert;
15         }
16     }
17     function onlyOwner() {
18         this.checkOwner();
19     }
20     function renounceOwnership() {
21         this.onlyOwner();
22         this.transferOwnership(address(0));
23     }
24     function transferOwnership(address newOwner) {
25         this.checkOwner();
26         if (newOwner == address(0)) {
27             revert;
28         } else {
29             this.transferOwnership(newOwner);
30         }
31     }
32 }
```

Figure 8.2: Ownable Contract

```
1 import "./Context.sol";
2
3 contract TinyERC721 is Context {
4
5     mapping(uint => address) owners;
6     mapping(address => uint) balances;
7
8     constructor() Context(){}
9
10    function fb() {}
11    function ownerOf(uint tokenId) returns (address) {
12        address ownr = this.owners[tokenId];
13        if (ownr == address(0)) {
14            return address(0);
15        }
16        else {
17            return ownr;
18        }
19    }
20    function balanceOf(address owner) returns (uint) {
21        if (owner == address(0)) {
22            revert;
23        }
24        else{
25            return this.balances[owner];
26        }
27    }
28    function exists(uint tokenId) returns (bool) {
29        return (this.ownerOf(tokenId) != address(0));
30    }
31    function transferFrom(address from, address to, uint tokenId) {
32        this.balances = (this.balances[from] = this.balances[from] - 1);
33        this.balances = ( this.balances[to] = this.balances[to] + 1);
34        this.owners = (this.owners[tokenId] = to);
35    }
36    function mint(address to, uint tokenId) {
37        if ((to == address(0)) || this.exists(tokenId)) {
38            revert;
39        }
40        else {
41            this.balances = ( this.balances[to] = this.balances[to] + 1);
42            this.owners = (this.owners[tokenId] = to);
43        }
44    }
45    function burn(address own, uint tokenId) {
46        this.owners = (this.owners[tokenId] = address(0));
47        uint currentBal = this.balances[own];
48        this.balances = (this.balances[own] = currentBal - 1);
49    }
50 }
```

Figure 8.3: ERC721 Contract

```
1 import "./Ownable.sol";
2 import "./ERC721.sol";
3
4 contract NFTStorage is Ownable, TinyERC721 {
5
6     uint lastUnsoldToken;
7     uint lastTokenId;
8     uint price;
9
10    constructor() Ownable(msgsender) TinyERC721() {}
11
12    function fb() {}
13    function setNFTPrice(uint val) {
14        this.onlyOwner();
15        this.price = val;
16    }
17    function createNFT(address to) {
18        this.onlyOwner();
19        this.mint(to, this.lastTokenId);
20        this.lastTokenId = this.lastTokenId + 1;
21    }
22    function buyNFT(address buyer) {
23        if(msgvalue >= this.price) {
24            this.transferFrom(address(0), buyer, this.lastUnsoldToken);
25        }
26        else {
27            revert;
28        }
29    }
30    function transferNFT(uint tokenId, address from, address to) {
31        if(from == this.ownerOf(tokenId)) {
32            this.transferFrom(from, to, tokenId);
33        }
34        else {
35            revert;
36        }
37    }
38    function destroyNFT(address owner, uint tokenId) {
39        if(owner == this.ownerOf(tokenId)) {
40            this.burn(owner, tokenId);
41        }
42        else {
43            revert;
44        }
45    }
46 }
```

Figure 8.4: NFT Storage Contract

If one prefers, there is also a GitHub repository, where all the contracts are conveniently organized. The repository can be found [here](#).

8.2 Contract Hierarchy and C3 Linearization Results

In the presented class hierarchy (figure 8.5), following the C3 Linearization algorithm we did implement, the linearization sequences for each class are as follows:

Context \rightarrow [Context]

Ownable \rightarrow [Ownable, Context]

TinyERC721 \rightarrow [TinyERC721, Context]

NFTStorage \rightarrow [NFTStorage, TinyERC721, Ownable, Context]

Game \rightarrow [Game, Context]

In the linearization of TinyERC721, the order of appearance for TinyERC721 and Ownable is not critical: both [NFTStorage, Ownable, TinyERC721, Context] and [NFTStorage, TinyERC721, Ownable, Context] are valid linearization sequences. However, in Solidity, the last-written superclass within a contract's definition takes precedence in the linearization order, which accounts for the specific output observed with our algorithm. It is worth noting that our algorithm does not consider the CTOP class in its computations. This decision is based on the rationale that CTOP will always occupy the last position in the hierarchy. It also does not have any methods or state variables, essentially serving as a top-level class for distinguishing contracts from externally owned accounts (EOAs) in addresses. This allows for the use of addresses of the form `address<CTOP>`, which exclusively accept contract instances.

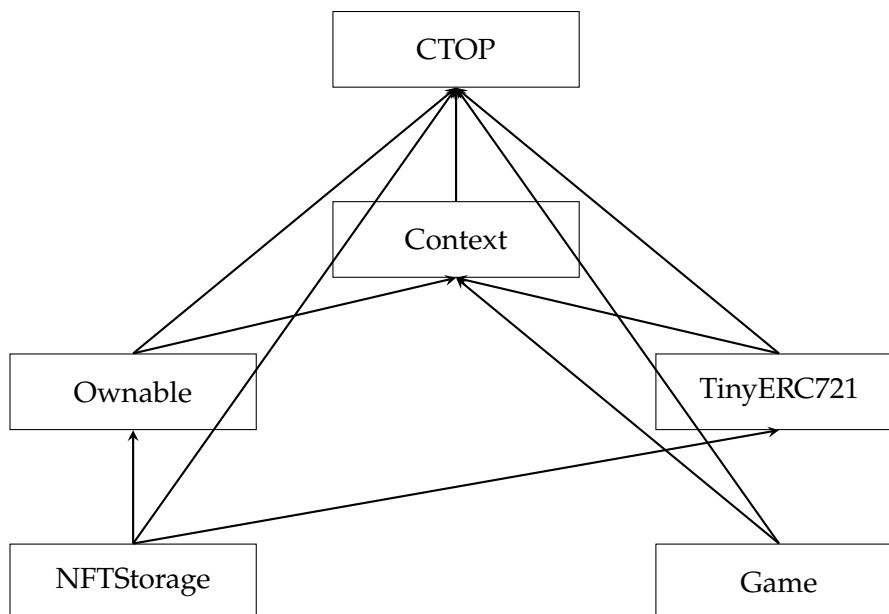


Figure 8.5: Case study class hierarchy

8.3 Execution Example

To illustrate the operational semantics, we show a possible interaction with these contracts. First, we deploy the `NFTStorage` contract into our mock blockchain, where we have two addresses, say `a1` (`0xe1abd8...`) and `a2` (`0x229005...`). Subsequently both will interact with the contract deployed and the whole blockchain through function calling (i.e., transactions).

First `a1` deploys the `Game` contract, creating an instance of it on our blockchain and we can referencings this instance either by their unique instance identifier or address. Then the same user creates a new store, invoking function `createStore`. After the transaction we have another contract on our blockchain, an instance of `NFTStorage` and it is stored in the state variable `stores`. We also return the address where this instance is located and store it because it is useful to pass on the next functions. It is important to say that this address is stored in the type environment as a subtype of the address type and therefore its refined type is `address<NFTStorage>`. Figure 8.6 shows the state of the blockchain after these first two transactions (for each contract shows the content of state variables and its balance). Next `a1` sets each NFT price to 322. In Figure 8.7 you can see that `a2` earned an NFT, by invoking the `buyNFT` function: balances and owners are two state variables of `NFTStorage` contract, both being maps. The former mapping each externally owned account to the amount of tokens this user is owner of and the latter is a mapping between each token and its owner. We repeat this transaction in Figure 8.8, making now `a2` holding two NFTs. Lastly `a2` transfers the first NFT bought to `a1`, using the `transferNFT` function, as you can see in Figure 8.9: key 0 of state variable `owners`, it changed the value and we can also note that `a1` now has a balance of 1 in state variable `balances`.

```
Contract 0, 0xa5a12d5c2d4b9af407325dd3b0d3491cf3f5304d, Contract Name: Game,
  State Variables:
stores ----> [
0x117c8569914f13732494006d1118888e8c4e9751 ---> Contract 1
]
Balance: 0

Contract 1, 0x117c8569914f13732494006d1118888e8c4e9751, Contract Name:
  NFTStorage, State Variables:
balances ----> [
]
lastTokenId ----> 0
lastUnsoldToken ----> 0
owner ----> 0xa5a12d5c2d4b9af407325dd3b0d3491cf3f5304d
owners ----> [
]
price ----> 0
Balance: 0

Account: 0x2290052e72090ecde135bde4c864732239f32c19 ,Balance: 100000
Account: 0xe1abd8e36dfeb46a79436f2d7a1bcfde8d70e097 ,Balance: 100000
```

Figure 8.6: Output of our tool after executing first two transactions

```

Contract 0, 0xa5a12d5c2d4b9af407325dd3b0d3491cf3f5304d, Contract Name: Game,
  State Variables:
stores ----> [
  0x117c8569914f13732494006d1118888e8c4e9751 ---> Contract 1
]
Balance: 0

Contract 1, 0x117c8569914f13732494006d1118888e8c4e9751, Contract Name:
  NFTStorage, State Variables:
balances ----> [
  0x2290052e72090ecde135bde4c864732239f32c19 ---> 1
]
lastTokenId ----> 1
lastUnsoldToken ----> 0
owner ----> 0xa5a12d5c2d4b9af407325dd3b0d3491cf3f5304d
owners ----> [
  0 ---> 0x2290052e72090ecde135bde4c864732239f32c19
]
price ----> 322
Balance: 0

Account: 0x2290052e72090ecde135bde4c864732239f32c19 ,Balance: 100000

Account: 0xe1abd8e36dfef46a79436f2d7a1bcfde8d70e097 ,Balance: 100000

```

Figure 8.7: Output after the fourth transaction

```

Contract 0, 0xa5a12d5c2d4b9af407325dd3b0d3491cf3f5304d, Contract Name: Game,
  State Variables:
stores ----> [
  0x117c8569914f13732494006d1118888e8c4e9751 ---> Contract 1
]
Balance: 0

Contract 1, 0x117c8569914f13732494006d1118888e8c4e9751, Contract Name:
  NFTStorage, State Variables:
balances ----> [
  0x2290052e72090ecde135bde4c864732239f32c19 ---> 2
]
lastTokenId ----> 2
lastUnsoldToken ----> 0
owner ----> 0xa5a12d5c2d4b9af407325dd3b0d3491cf3f5304d
owners ----> [
  0 ---> 0x2290052e72090ecde135bde4c864732239f32c19
  1 ---> 0x2290052e72090ecde135bde4c864732239f32c19
]
price ----> 322
Balance: 0

Account: 0x2290052e72090ecde135bde4c864732239f32c19 ,Balance: 100000

Account: 0xe1abd8e36dfef46a79436f2d7a1bcfde8d70e097 ,Balance: 100000

```

Figure 8.8: Output after the fifth transaction

```
Contract 0, 0xa5a12d5c2d4b9af407325dd3b0d3491cf3f5304d, Contract Name: Game,
  State Variables:
stores ----> [
0x117c8569914f13732494006d1118888e8c4e9751 ---> Contract 1
]
Balance: 0

Contract 1, 0x117c8569914f13732494006d1118888e8c4e9751, Contract Name:
  NFTStorage, State Variables:
balances ----> [
  0xe1abd8e36dfef46a79436f2d7a1bcfde8d70e097 ---> 1
  0x2290052e72090ecde135bde4c864732239f32c19 ---> 1
]
lastTokenId ----> 2
lastUnsoldToken ----> 0
owner ----> 0xa5a12d5c2d4b9af407325dd3b0d3491cf3f5304d
owners ----> [
  0 ---> 0xe1abd8e36dfef46a79436f2d7a1bcfde8d70e097
  1 ---> 0x2290052e72090ecde135bde4c864732239f32c19
]
price ----> 322
Balance: 0

Account: 0x2290052e72090ecde135bde4c864732239f32c19 ,Balance: 100000

Account: 0xe1abd8e36dfef46a79436f2d7a1bcfde8d70e097 ,Balance: 100000
```

Figure 8.9: Output after all transactions

CONCLUSIONS AND FUTURE WORK

In this final chapter, we summarize the key findings and contributions of our research in the domain of safety on smart contracts. We also outline potential enhancements and improvements for our PoC.

9.1 Conclusion

Smart contracts are still in an early stage of development, and their adoption is growing rapidly. Ensuring their correctness through formal verification is paramount due to their immutability and the financial stakes involved, specially because they are immutable and errors can possibly have catastrophic financial consequences. Formal verification emerges as a powerful tool to achieve this goal.

Correctness can be achieved through either making a language robust against vulnerabilities or developing tools to detect and mitigate potential issues. Both approaches are essential for a safer smart contract ecosystem.

We have also shown compelling evidence indicating that domain-specific languages are significantly more suitable for writing smart contracts.

Our work has focused on formalizing an extension to a Solidity's subset, making it type safe. We have taken an existing formalization and modified it to align with the latest version of Solidity, including features like multiple inheritance between contracts. Additionally, we have implemented our formalization in OCaml to demonstrate that our solution is feasible.

9.2 Future Work

One limitation of our formalization in the context of blockchains arises from the fact that smart contracts run on shared virtual machines, increasing the likelihood of contract name collisions. To address this issue, we propose an enhancement: instead of using only the contract name, we include the hash of the contract code as part of the address type, represented as $\text{address}\langle C, H \rangle$, where H represents the hash of the contract code.

Furthermore, it would be beneficial for blockchains to collaborate and adopt an universal language that is compatible across all platforms, akin to JavaScript's ubiquity in web development. Having a single language would concentrate the efforts of blockchain projects and developers on ensuring the safety of one language, rather than dealing with multiple languages.

When it comes to programming languages for writing smart contracts, it is also essential to offer tested and audited libraries and frameworks, similar to the established [OpenZeppelin](#), while promoting safe coding practices.

BIBLIOGRAPHY

- [1] V. Buterin. <https://twitter.com/VitalikButerin/status/1024265820849860609>. 2018 (cit. on p. 11).
- [2] V. Buterin. *Ethereum: A Next-Generation Smart Contract and Decentralized Application Platform*. https://ethereum.org/669c9e2e2027310b6b3cdce6e1c52962/Ethereum_Whitepaper_-_Buterin_2014.pdf. 2014 (cit. on p. 4).
- [3] V. Buterin. *Hard Fork Completed*. <https://blog.ethereum.org/2016/07/20/hard-fork-completed/>. 2016 (cit. on p. 22).
- [0] K. Claessen and J. Hughes. “QuickCheck: a lightweight tool for random testing of Haskell programs”. In: *Proceedings of the Fifth ACM SIGPLAN International Conference on Functional Programming (ICFP '00), Montreal, Canada, September 18-21, 2000*. Ed. by M. Odersky and P. Wadler. ACM, 2000, pp. 268–279. DOI: [10.1145/351240.351266](https://doi.org/10.1145/351240.351266). URL: <https://doi.org/10.1145/351240.351266> (cit. on p. 75).
- [4] *Consistency model*. https://en.wikipedia.org/wiki/Consistency_model. Accessed: 2022-05-10 (cit. on p. 5).
- [5] T. H. Cormen et al. *Introduction to Algorithms, 3rd Edition*. MIT Press, 2009. ISBN: 978-0-262-03384-8. URL: <http://mitpress.mit.edu/books/introduction-algorithms> (cit. on p. 6).
- [6] S. Crafa, M. D. Pirro, and E. Zucca. “Is Solidity Solid Enough?” In: *Financial Cryptography and Data Security - FC 2019 International Workshops, VOTING and WTSC, St. Kitts, St. Kitts and Nevis, February 18-22, 2019, Revised Selected Papers*. Ed. by A. Bracciali et al. Vol. 11599. Lecture Notes in Computer Science. Springer, 2019, pp. 138–153. DOI: [10.1007/978-3-030-43725-1_11](https://doi.org/10.1007/978-3-030-43725-1_11) (cit. on p. 2).
- [7] P. Daian. *Analysis of the DAO exploit*. <https://hackingdistributed.com/2016/06/18/analysis-of-the-dao-exploit/>. 2016 (cit. on p. 22).
- [8] C. Dwork and M. Naor. “Pricing via Processing or Combatting Junk Mail”. In: *Advances in Cryptology - CRYPTO '92, 12th Annual International Cryptology Conference, Santa Barbara, California, USA, August 16-20, 1992, Proceedings*. Ed. by E. F. Brickell.

- Vol. 740. Lecture Notes in Computer Science. Springer, 1992, pp. 139–147. DOI: [10.1007/3-540-48071-4_10](https://doi.org/10.1007/3-540-48071-4_10). URL: https://doi.org/10.1007/3-540-48071-4_10 (cit. on p. 7).
- [9] J. FRANKENFIELD. *51% Attack*. <https://www.investopedia.com/terms/1/51-attack.asp>. 2021 (cit. on p. 7).
- [10] L. Goodman. *Tezos — a self-amending crypto-ledger White paper*. <https://tezos.com/whitepaper.pdf>. 2014 (cit. on p. 18).
- [11] Á. Hajdu and D. Jovanović. “solc-verify: A Modular Verifier for Solidity Smart Contracts”. In: *Verified Software. Theories, Tools, and Experiments*. Ed. by S. Chakraborty and J. A. Navas. Vol. 12301. Lecture Notes in Computer Science. Springer, 2020, pp. 161–179. ISBN: 978-3-030-41600-3. DOI: [10.1007/978-3-030-41600-3_11](https://doi.org/10.1007/978-3-030-41600-3_11) (cit. on p. 30).
- [12] A. Igarashi, B. C. Pierce, and P. Wadler. “Featherweight Java: a minimal core calculus for Java and GJ”. In: *ACM Trans. Program. Lang. Syst.* 23.3 (2001), pp. 396–450. DOI: [10.1145/503502.503505](https://doi.org/10.1145/503502.503505). URL: <https://doi.org/10.1145/503502.503505> (cit. on p. 33).
- [0] *Implementing C3 Linearization*. <https://xivilization.net/~marek/blog/2014/12/08/implementing-c3-linearization/> (cit. on p. 54).
- [13] N. Insights. *Zilliqa: The Sharding Chain*. <https://novuminsights.medium.com/zilliqa-the-sharding-chain-d25bb81c3fdc>. 2022 (cit. on p. 20).
- [14] J. Jiao et al. “Executable Operational Semantics of Solidity”. In: *CoRR abs/1804.01295* (2018). arXiv: [1804.01295](https://arxiv.org/abs/1804.01295). URL: <http://arxiv.org/abs/1804.01295> (cit. on p. 2).
- [15] L. Lamport. “The Part-Time Parliament”. In: *ACM Trans. Comput. Syst.* 16.2 (1998), pp. 133–169. DOI: [10.1145/279227.279229](https://doi.org/10.1145/279227.279229). URL: <https://doi.org/10.1145/279227.279229> (cit. on p. 5).
- [16] L. Lamport, R. Shostak, and M. Pease. *The Byzantine Generals Problem*. <https://www.microsoft.com/en-us/research/uploads/prod/2016/12/The-Byzantine-Generals-Problem.pdf>. 1981 (cit. on p. 5).
- [17] N. A. Lynch. *Distributed Algorithms*. Morgan Kaufmann, 1996. ISBN: 1-55860-348-4 (cit. on p. 5).
- [18] R. C. Merkle. “A Digital Signature Based on a Conventional Encryption Function”. In: *Advances in Cryptology - CRYPTO '87, A Conference on the Theory and Applications of Cryptographic Techniques, Santa Barbara, California, USA, August 16-20, 1987, Proceedings*. Ed. by C. Pomerance. Vol. 293. Lecture Notes in Computer Science. Springer, 1987, pp. 369–378. DOI: [10.1007/3-540-48184-2_32](https://doi.org/10.1007/3-540-48184-2_32). URL: https://link.springer.com/chapter/10.1007/3-540-48184-2_32 (cit. on p. 6).
- [19] S. Nakamoto. *Bitcoin: A Peer-to-Peer Electronic Cash System*. <https://bitcoin.org/bitcoin.pdf>. 2009 (cit. on p. 4).

-
- [20] C. Parizo. *What are the 4 different types of blockchain technology?* <https://www.techtarget.com/searchcio/feature/What-are-the-4-different-types-of-blockchain-technology>. 2021 (cit. on p. 8).
- [21] M. D. Pirro. *How Solid is Solidity? An In-dept Study of Solidity's Type Safety*. Msc Thesis. Università degli Studi di Padova. 2018 (cit. on p. 2).
- [22] Z. D. Portal. *Sharding Mechanism*. <https://dev.zilliqa.com/docs/basics/basics-zil-sharding/>. 2022 (cit. on p. 20).
- [23] S. N. Sunny King. *PPCoin: Peer-to-Peer Crypto-Currency with Proof-of-Stake*. <https://decred.org/research/king2012.pdf>. 2012 (cit. on p. 7).
- [24] N. Szabo. *Smart Contracts: Building Blocks for Digital Markets*. https://www.fon.hum.uva.nl/rob/Courses/InformationInSpeech/CDROM/Literature/L0Twinterschool2006/szabo.best.vwh.net/smart_contracts_2.html. 1996 (cit. on p. 4).
- [25] A. S. Tanenbaum and M. van Steen. *Distributed systems - principles and paradigms, 3rd Edition*. distributed-systems.net, 2017. ISBN: 978-1543057386 (cit. on p. 4).
- [26] P. Technologies. *A Postmortem on the Parity Multi-Sig Library Self-Destruct*. <https://www.parity.io/blog/a-postmortem-on-the-parity-multi-sig-library-self-destruct/>. 2017 (cit. on p. 25).
- [27] R. M. Townsend. *Distributed Ledgers: Design and Regulation of Financial Infrastructure and Payment Systems*. <https://direct.mit.edu/books/book/4932/chapter/625100/Preface>. 2020 (cit. on p. 5).
- [28] *What Is the Byzantine Generals Problem?* <https://river.com/learn/what-is-the-byzantine-generals-problem/>. Accessed: 2022-05-10 (cit. on p. 5).
- [29] J. Zakrzewski. "Towards Verification of Ethereum Smart Contracts: A Formalization of Core of Solidity". In: *Verified Software. Theories, Tools, and Experiments*. Ed. by R. Piskac and P. Rümmer. Cham: Springer International Publishing, 2018, pp. 229–247. ISBN: 978-3-030-03592-1 (cit. on p. 2).



2023 Towards a Stronger Solidarity João Reis