

# Verification of Bitcoin Script in Agda using Weakest Preconditions for Access Control

Fahad F. Alhabardi ✉ 

Dept. of Computer Science, Dept. of Computer Science

Arnold Beckmann ✉ 

Swansea University, Dept. of Computer Science

Bogdan Lazar ✉

University of Bath

Anton Setzer ✉ 

Swansea University, Dept. of Computer Science

---

## Abstract

This paper contributes to the verification of programs written in Bitcoin’s smart contract language SCRIPT in the interactive theorem prover Agda. It focuses on the security property of access control for SCRIPT programs that govern the distribution of Bitcoins. It advocates that *weakest preconditions* in the context of Hoare triples are the appropriate notion for verifying access control. It aims at obtaining human-readable descriptions of weakest preconditions in order to close the validation gap between user requirements and formal specification of smart contracts.

As examples for the proposed approach, the paper focuses on two standard SCRIPT programs that govern the distribution of Bitcoins, *Pay to Public Key Hash (P2PKH)* and *Pay to Multisig (P2MS)*. The paper introduces an operational semantics of the SCRIPT commands used in P2PKH and P2MS, which is formalised in the Agda proof assistant and reasoned about using Hoare triples. Two methodologies for obtaining human-readable descriptions of weakest preconditions are discussed: (1) a step-by-step approach, which works backwards instruction by instruction through a script, sometimes grouping several instructions together; (2) symbolic execution of the code and translation into a nested case distinction, which allows to read off weakest preconditions as the disjunction of conjunctions of conditions along accepting paths. A syntax for equational reasoning with Hoare Triples is defined in order to formalise those approaches in Agda.

**2012 ACM Subject Classification** Theory of computation → Hoare logic; Theory of computation → Type theory; Theory of computation → Programming logic; Theory of computation → Interactive proof systems; Theory of computation → Operational semantics; Theory of computation → Denotational semantics; Security and privacy → Access control; Security and privacy → Logic and verification; Applied computing → Digital cash

**Keywords and phrases** Blockchain; Cryptocurrency; Bitcoin; Agda; Verification; Hoare logic; Bitcoin script; P2PKH; P2MS; Access control; Weakest precondition; Predicate transformer semantics; Provable correctness; Symbolic execution; Smart contracts

**Digital Object Identifier** 10.4230/LIPIcs.TYPES.2021.1

**Supplementary Material** *Software (Agda Source Code):*

<https://github.com/fahad1985lab/Smart--Contracts--Verification--With--Agda>

**Acknowledgements** We would like to thank the anonymous referees for valuable comments and suggestions.



© Fahad F. Alhabardi, Arnold Beckmann, Bogdan Lazar, Anton Setzer;  
licensed under Creative Commons License CC-BY 4.0

27th International Conference on Types for Proofs and Programs (TYPES 2021).

Editors: Henning Basold, Jesper Cockx, and Silvia Ghilezan; Article No. 1; pp. 1:1–1:27

Leibniz International Proceedings in Informatics



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

## 1 Introduction

Bitcoin, the first cryptocurrency, was introduced in 2008 by Satoshi Nakamoto [37] to provide a public payment mechanism, the blockchain, using pseudonymous keys in a peer-to-peer network with distributed control. Many other cryptocurrencies have been introduced since.

Bitcoin’s blockchain provides a scripting system for transactions called SCRIPT. Lists of instructions in SCRIPT are denoted *Bitcoin scripts* or simply *scripts*. The invention of Ethereum [13] strengthened Bitcoin by adding full (Turing complete) smart contracts to blockchains. In this context, smart contracts can be seen as programs that automatically execute when transactions are performed on a blockchain. Though not Turing complete as Ethereum [13], Bitcoin scripts can be viewed as a weak form of smart contracts, that provide important functionality, e.g. by governing the distribution of the Bitcoin cryptocurrency.

As smart contracts, including Bitcoin scripts, can control real world values and are immutable once deployed on the blockchain network, a method to demonstrate their security and correctness is needed [31, 49]. According to [4, 36, 19], there are two ways to verify their correctness: (1) by using mathematical methods like formal verification, which utilise theorem proving, model checking, and similar techniques, and (2) by employing testing. Theorem proving provides an extremely flexible verification method that can be applied to various types of systems including smart contracts. It can be done in interactive, automated, or hybrid mode.

In our approach, we use the interactive theorem prover Agda [2] for the verification of Bitcoin scripts. Agda is designed to be both an interactive theorem prover and a dependently typed programming language [39], therefore Agda allows us to define programs and reason about them in the same system. This reduces the danger of producing errors when translating programs from a programming language to a theorem prover, and allows to execute smart contracts in Agda directly. Another advantage of Agda is to have proofs that are checkable by hand. Other frameworks, such as Coq [42], use automatic proof search tools which usually do not provide proof certificates that could in principle be checked by hand. Human checkable proof certificates reduce the need to rely on a theorem prover being correct. The latter is desirable because of potential attacks that exploit errors in theorem provers, e.g. by creating a smart contract that contains a deliberate error together with a correctness proof that exploits the error in the theorem prover.<sup>1</sup> As a final point, there are some key distinctions between Agda and other theorem provers like Coq that suggest a different applicability of Agda. For example, Agda supports inductive-recursive types, while Coq does not [12]. Agda also has a more flexible pattern matching system than Coq, including support for copattern matching [12].

**Main contributions.** Our main contributions in this paper are:

- We argue that weakest preconditions are the appropriate notion to verify access control for Bitcoin scripts.
- We propose to aim for human-readable descriptions of weakest preconditions to support judging whether the security property of access control is satisfied.
- We describe two methods for achieving human-readable descriptions of weakest preconditions: a step-by-step approach, and a symbolic-execution-and-translation approach.

---

<sup>1</sup> See the forum discussion on [47] for a well documented list of incorrect protocols with false correctness proofs. To hide backdoors using deliberately false correctness proofs is certainly conceivable.

- We apply our proposed methodology to two standard Bitcoin scripts, providing fully formalised arguments in Agda.

In the following we explain our contributions in more detail. The paper introduces an operational semantics of the SCRIPT commands used in *Pay to Public Key Hash (P2PKH)* and *Pay to Multisig (P2MS)*, two standard scripts that govern the distribution of Bitcoins. We define the operational semantics as stack operations and reason about the correctness of such operations using Hoare triples utilising pre- and postconditions.

**Weakest precondition for access control.** Our verification focuses on the security property of *access control*. Access control is the restriction to access for a resource, which in our use case is access to cryptocurrencies like Bitcoin. We advocate that, in the context of Hoare triples, *weakest preconditions* are the appropriate notion to model access control: A (general) precondition expresses that when it is satisfied, access is granted, but there may be other ways to gain access without satisfying the precondition. The weakest precondition expresses that access is granted if and only if the condition is satisfied.

**Human-readable descriptions.** The weakest precondition can always be described in a direct way, for example as the set of states that after execution of the smart contract end in a state satisfying the given postcondition. However, such a description is meaningless to humans who want to convince themselves that the smart contract is secure, in the sense that they do not provide any further insights beyond the original smart contract.

It is known in software engineering, that failures of safety-critical systems are often due to incomplete requirements or specifications rather than coding errors.<sup>2</sup> The same applies to security related software.<sup>3</sup> It is not sufficient to have a proof of security of a protocol, if the statement does not express what is required. That the specification (here the formal statement of secure access control) guarantees that the requirements are fulfilled (namely that it is impossible for a hacker to access the resource, here the Bitcoin), needs to be checked by a human being, who needs to be able to read the specification and determine whether it really is what is expressed by the requirements. Thus, the challenge is to obtain simple, human-readable descriptions of the weakest precondition of a smart contract. This would allow to close the validation gap between user requirements and formal specification of smart contracts.

**Two methods for obtaining human-readable weakest preconditions.** We discuss two methods for obtaining readable weakest preconditions: The first, step-by-step approach, is obtained by working through the program backwards instruction by instruction. In some cases it is easier to group several instructions together and deal with them differently, as we will demonstrate with an example in Sect. 6.3. The second method, symbolic-execution-and-translation, evaluates the program in a symbolic way, and translates it into a nested case distinction. The case distinctions are made on variables (of type `nat` or `stack`) or on expressions formed from variables by applying basic functions to them such as hashing or checking for signature. From the resulting decision tree, the weakest precondition can be

---

<sup>2</sup> For instance, [32] writes: “Almost all accidents with serious consequences in which software was involved can be traced to requirements failures, and particularly to incomplete requirements.”

<sup>3</sup> The long list of protocols which were proven to be secure but had wrong proofs [20] demonstrates that a proof of correctness is not sufficient. We assume that most of the examples had correct proofs, but the statement shown was not sufficient to guarantee security.

read off as the disjunction of the conjunctions of the conditions that occur along branches that lead to a successful outcome.

For both methods, it is necessary to prove that the established weakest precondition is indeed the weakest precondition for the program under consideration. For the first method, this follows by stepwise operation. The second uses a proof that the original program is equivalent to the transformed program from which the weakest precondition has been established, or a direct proof which follows the case distinctions used in the symbolic evaluation.

**Application of our proposed methodology.** We demonstrate the feasibility of our approaches by carrying them out in Agda for concrete smart contracts, including P2PKH and P2MS.

Our approach also provides opportunities for further applications: The usage of the weakest precondition with explicit proofs can be seen as a method of building verified smart contracts that are *correct by construction*. Instead of constructing a program and then evaluating it, one can start with the intended weakest precondition and postcondition, add some intermediate conditions, and then develop the program between those conditions. Such an approach would extend the SPARK Ada framework [1] to use Hoare logic (without the weakest precondition) to check programs.

The remainder of this paper is organised as follows. In Sect. 2, we introduce related work on verification of smart contracts. Sect. 3 introduces Bitcoin SCRIPT and defines its operational semantics. In Sect. 4, we specify the security of Bitcoin SCRIPT using Hoare logic and weakest preconditions. We formalise these notions in Agda and introduce equational reasoning for Hoare triples to streamline our correctness proofs. Sect. 5 introduces our first, step-by-step method of developing human-readable weakest preconditions and proving correctness of P2PKH. In Sect. 6, we introduce our second method based on symbolic execution and apply it to various examples. In Sect. 7, we explain how to practically use Agda to determine and prove weakest preconditions using our library [46]. We conclude in Sect. 8.

**Notations and git repository.** The formulas can be presented as full Agda code, but often the formulas can also be presented in mathematical style. In order to switch between Agda code and mathematical code easy, we use the functional style for application (i.e. writing  $f a b c$  instead of  $f(a, b, c)$ ) and  $x : A$  instead of  $x \in A$ .  $s :: l$  denotes prepending an element onto a list. The original Agda definitions are also available [46]. Most display style Agda code presented in this paper has been automatically extracted from the Agda code, in some cases it was formatted by hand based on  $\text{\LaTeX}$  code generated by Agda to improve the presentation.

## 2 Related Work

In this section we describe research relevant to our approach. We start by discussing two papers introducing Hoare logic, predicate transformer semantics and weakest preconditions. We then review papers that address verification of smart contracts, and Bitcoin scripts. We present a number of approaches to use model-checking for the verification of smart contracts, and finish with work with employs Agda in the verification of smart contracts.

**Hoare Logic, Predicate Transformer Semantics and Weakest Preconditions.** Hoare [26] defines a formal system using logical rules for reasoning about the correctness of computer

programs. It uses so-called Hoare triples which combine two predicates, a pre- and a postcondition, with a program to express that if the precondition holds for a state and the program executes successfully, then the postcondition holds for the resulting state. Dijkstra [22] introduces predicate transformer semantics that assigns to each statement in an imperative programming paradigm a corresponding total function between two predicates on the state space of the statement. The predicate transformer defined by Dijkstra applied to a postcondition returns the weakest precondition.

**Verification of Smart Contracts.** A number of authors have addressed the verification of smart contracts in Ethereum and similar platforms. Hirai [25] used Isabelle/HOL theorem prover to validate Ethereum Virtual Machine (EVM) bytecode by developing a formal model for EVM using the Lem language [?]. They use this model to prove invariants and safety properties of Ethereum smart contracts. Amani et al. [5] extended Hirai's EVM formalisation in Isabelle/HOL by a sound program logic at bytecode level. To this end, they stored bytecode sequences into blocks of straight-line code, creating a program logic that could reason about these sequences. Ribeiro et al. [44] developed an imperative language for a relevant subset of Solidity in the context of Ethereum, using a big-step semantics system. Additionally, they formalised smart contracts in Isabelle/HOL, extending the existing work. Their formalisation of semantics is based on Hoare logic and the weakest precondition calculus. Their main contributions are proofs of soundness and relative completeness, as well as applications of their machinery to verify some smart contracts including modelling of smart contract vulnerabilities. Bernardo et al. [9] present Mi-Cho-Coq, a Coq framework which has been used to formalise Tezos smart contracts written in the stack-based language Michelson. The framework is composed of a Michelson interpreter implemented in Coq, and the weakest precondition calculus to verify Michelson smart contracts' functional correctness. O'Connor [40] introduces Simplicity, a low-level, typed functional language, which is Turing incomplete. The goal of Simplicity is to improve on existing blockchain-based languages, like Ethereum's EVM and Bitcoin SCRIPT, while avoiding some of their issues. Simplicity is based on formal semantics and is specified in the Coq proof assistant. Bhargavan et al. [10] provided formalisations of EVM bytecode in  $F^*$ , a functional programming language designed for program verification. They defined a smart contract verification architecture that can compile Solidity contracts, and decompile EVM bytecode into  $F^*$  using their shallow embedding, in order to express and analyse smart contracts.

**Verification of Bitcoin Scripts.** Klomp et al. [30] proposed a symbolic verification theory, and a tool to analyse and validate Bitcoin scripts, with a particular focus on characterising the conditions under which an output script, which controls the successful transfer of Bitcoins, will succeed. Bartoletti et al. [8] presented BitML, a high-level domain-specific language for designing smart contracts in Bitcoin. They provided a compiler to convert smart contracts into Bitcoin transactions, and proved the correctness of their compiler w.r.t. a symbolic model for BitML and a computational model, which has been defined as well in [7] for Bitcoin. Setzer [45] developed models of the Bitcoin blockchain in the interactive theorem prover Agda. This work focuses on the formalisation of basic primitives in Agda as a basis for future work on verifying the protocols of cryptocurrencies and developing verified smart contracts.

**Verification of Smart Contracts using Model Checking.** A number of papers discuss tools for analysing and verifying smart contracts that utilise model checking. Kalra et al. [28] developed a framework called ZEUS whose aim is to support automatic formal verification of

smart contracts using abstract interpretation and symbolic model checking. ZEUS starts from a high-level smart contract, and employs user assistance for capturing correctness and fairness requirements. The contract and policy specification are then transformed into an intermediate language with well defined execution semantics. ZEUS then performs static analysis on the intermediate level and uses external SMT solvers to evaluate any verification properties discovered. A main focus of the work is on reducing efficiently the state explosion problem inherent in any model checking approach. Park et al. [41] proposed a formal verification tool for EVM bytecode based on KEVM, a complete formal semantics of EVM bytecode developed in the K-framework. To address performance challenges, they define EVM-specific abstractions and lemmas, which they then utilise to verify a number of concrete smart contracts. Mavridou et al. [33] introduce the VeriSolid framework to support the verification of Ethereum smart contracts. VeriSolid is based on earlier work (FSolidM) which allows to graphically specify Ethereum smart contracts as transitions systems, and to generate Solidity code from those specification. It uses model checking to verify smart contract models. Luu et al. [31] provided operational semantics of a subset of Ethereum bytecode called EtherLite, which forms the bases of their symbolic execution tool Oyente for analysing Ethereum smart contracts. Based on their tool they discovered a number of weaknesses in deployed smart contracts, including the DAO bug [23]. Filiâtre et al. [24] introduced the Why3 system, which allows writing imperative programs in WhyML, an ML dialect used for programming and specification. The system can add pre-, post- and intermediate conditions to it but does not make use of weakest precondition. Why3 can generate verification conditions for Hoare triple, which are checked using variously automated and interactive theorem provers. Why3 is used in SPARK Ada to verify its verification conditions.

**Agda in the Verification of Blockchains.** Finally, Agda features in several papers discussing verification of blockchains. Chakravarty et al. [16] introduce Extended UTXO (EUTXO), which extends Bitcoin's UTXO model to enable more expressive forms of validation scripts. These scripts can express general state machines and reason about transaction chains: The authors introduce a new class of state machines based on Mealy machines which they call Constraint Emitting Machines (CEM). In addition to formalising CEMs using Agda proof assistant, they demonstrate its conversion to EUTXO, and give a weak bisimulation between both systems. In [14] Chakravarty et al. introduce a generalisation of the EUTXO ledger model using native tokens which they denote EUTXOma for EUTXO with multi-assets. They provide a formalisation of the multi-asset EUTXO model in Agda. Chakravarty et al. [15] introduce a version of EUTXOma aligned to Bitcoin's UTXO model, hence denoted UTXOma. They present a formal specification of the UTXOma ledger rules and formalise their model in Agda. Chapman et al. [17] formalise System  $F_{\omega\mu}$ , which is polymorphic  $\lambda$ -calculus with higher-kinded and arbitrary recursive types, in Agda. System  $F_{\omega\mu}$  corresponds to Plutus Core, which is the core of the smart contract language Plutus that features in the Cardano blockchain. Melkonian [34] introduces a formal Bitcoin transaction model to simulate transactions in the Bitcoin environment and to study their safety and correctness. The paper presents a formalisation of a process calculus for Bitcoin smart contracts, denoted BitML. The calculus can accept different types such as basic types, contracts, or small step semantics to outline a 'certified compiler' [35].

### 3 Operational Semantics for Bitcoin Script

We give a brief introduction of Bitcoin SCRIPT in Subsec. 3.1, before defining its operational semantics in Subsec. 3.2.

#### 3.1 Introduction to Bitcoin Script

The scripting language for Bitcoin is stack-based, inspired by the programming language Forth [43], with the stack being the only memory available. Elements on the stack are byte vectors, which we represent as natural numbers. Values on the stack are also interpreted as truth values, any value  $>0$  will be interpreted as true, and any other value as false. SCRIPT has its own set of commands called *opcodes*, which manipulate the stack. They are similar to machine instructions, although some instructions have a more complex behaviour. The instructions of SCRIPT are executed in sequence. In case of conditionals (which are not part of this paper) the execution of instructions might be ignored until the end of an if- or else-case has been reached, otherwise the script is executed from left to right. Execution of instructions might fail, in which case the execution of the script is aborted. A full list of instructions with their meaning can be found in [11], which is the defacto specification of SCRIPT.

The operational semantics of the opcodes can be found in the source code [46]. We introduce here a number of opcodes that are relevant to this paper. Execution of all opcodes fails, if there are not sufficiently many elements on the stack to perform the operation in question.

- OP\_DUP duplicates the top element of the stack.
- OP\_HASH takes the top item of the stack and replaces it with its hash.
- OP\_EQUAL pops the top two elements in the stack and checks whether they are equal or not, pushing the Boolean result on the stack.
- OP\_VERIFY invalidates the transaction if the top stack value is false. The top item on the stack will be removed.
- OP\_CHECKSIG pops two elements from the stack and checks whether they form a correct pair of a signature and a public key signing a serialised message obtained from the selected input and all outputs of the transaction, and pushes the Boolean result on the stack.
- OP\_CHECKLOCKTIMEVERIFY fails if the time on the stack is greater than the current time.
- OP\_MULTISIG is the multisig instruction, which will be discussed in detail in Sect. 6.2.
- There are a number of opcodes for pushing byte vectors of different lengths onto the stack. We write `<number>` for the opcode together with arguments pushing `number` onto the stack. In Agda we will have one instruction `opPush n` which pushes the number `n` on the stack.

Scripts can also contain control flow statements such as OP\_IF. The verification of scripts involving control statements is more involved and will be considered in a follow-up paper.

In Bitcoin we consider the interplay between a locking script `scriptPubKey` and an unlocking script `scriptSig`.<sup>4</sup> The locking script is provided by the sender of a transaction to lock the transaction, and the unlocking script is provided by the recipient to unlock it. The unlocking script pushes the data required to unlock the transaction on the stack, and

<sup>4</sup> We are using the terminology locking script and unlocking script from [6, Chapt 5].

the locking script then checks whether the stack contains the required data. Therefore, the unlocking script is executed first, followed by the locking script.<sup>5</sup>

The main example in this paper is the pay-to-public-key-hash (P2PKH) script consisting of the following locking and unlocking scripts:

```
scriptPubKey: OP_DUP OP_HASH160 <pubKeyHash> OP_EQUAL OP_VERIFY OP_CHECKSIG
scriptSig:    <sig> <pubKey>
```

The standard unlocking script `scriptSig` pushes a signature `sig` and a public key `pubKey` onto the stack. The locking script `scriptPubKey` checks whether `pubKey` provided by the unlocking script hashes to the provided `pubKeyHash`, and whether the signature is a signature for the message signed by the public key. Full details will be discussed in Sect. 5.

### 3.2 Operational Semantics

Opcodes like `OP_DUP` operate on the stack defined in Agda as a list of natural numbers `Stack`. Opcodes like `OP_CHECKSIG` check for signatures for the part of the transaction which is to be signed – what is to be signed is hard coded in Bitcoin. In order to abstract away from the precise format and the encoding, we define a message type `Msg` in Agda, which allows to represent messages such as those for the transaction to be signed, and is to be instantiated with the concrete message to be signed. Other opcodes like `OP_CHECKLOCKTIMEVERIFY` refer to the current time, for which we define a type `Time` in Agda. Therefore, the operational semantics of opcodes depends on `Time × Msg × Stack` which we define in Agda as the record type `StackState`.<sup>6</sup> Note that `Time` and `Msg` don’t change during the execution of a script.

The type of all opcodes is given as `InstructionBasic`.<sup>7</sup> Opcodes can also fail, for instance if there are not enough elements on the stack as required by the operation. Hence, the operational semantics of an instruction `op` : `InstructionBasic` is given as  $\llbracket op \rrbracket_s : \text{StackState} \rightarrow \text{Maybe StackState}$ .<sup>8</sup>

The message and time never change, so  $\llbracket p \rrbracket_s$  will, if executed successfully, only change the stack part of  $s$ . As an example, we can define the semantics of the instruction `opEqual`.

<sup>5</sup> In the original version of Bitcoin both scripts were concatenated and executed. However, because Bitcoin script has non-local instructions (e.g. the conditionals `OP_IF`, `OP_ELSE`, `OP_ENDIF`), when concatenating the two scripts any non-local opcode occurring in the locking script (for instance as part of data) could be interpreted when running as the counterpart of a non-local opcode in the locking script and therefore result in an unintended execution of the unlocking script. As a bug fix, in a later version of Bitcoin this was modified by having a break point in between the two, where only the stack is passed on. See Chapter 6, “Separate execution of unlocking and locking scripts” in [6, p. 136]. In this paper this problem doesn’t occur because we don’t consider non-local instructions.

<sup>6</sup> The idea of packaging all components of the state into one product type, which is then expanded into a more expanded state as more language constructs are added to the language, is inspired by Peter Mosses’ Modular SOS approach [21]. This approach was successful in creating a library of reusable components funcons for defining an executable operational semantics of language constructs, which require different sets of states. One outcome was a “component-based semantics for CAML LIGHT” [18].

<sup>7</sup> We are using in this paper a sublanguage `BitcoinScriptBasic` of Bitcoin, which doesn’t contain conditionals, because they require a more complex operational semantics and state (see the discussion in the conclusion). We make the distinction between the basic and full language in order to be compatible with the planned follow up papers based on code under development, which will extend the basic language. We sometimes use notations such as <sup>b</sup> to differentiate between functions referring to the basic and full language.

<sup>8</sup> For the reader not familiar with the `Maybe` type, a set theoretic notation can be given as `Maybe X` :=  $\{\text{nothing}\} \cup \{\text{just } x \mid x : X\}$ . Here, `nothing` denotes undefined, and `just x` denotes the defined element  $x$ . `Maybe` forms a monad, with `return` := `just` :  $A \rightarrow \text{Maybe } A$  and the bind operation  $(p \gg= q) : \text{Maybe } B$  for  $p : \text{Maybe } A$  and  $q : A \rightarrow \text{Maybe } B$  defined by  $(\text{nothing} \gg= q) = \text{nothing}$  and  $(\text{just } a \gg= q) = q a$ .



We first define a simpler function  $\llbracket \_ \rrbracket_s^s$ , which abstracts away the non-changing components `Time` and `Msg`:

$$\begin{aligned} \llbracket \_ \rrbracket_s^s &: \text{InstructionBasic} \rightarrow \text{Time} \rightarrow \text{Msg} \rightarrow \text{Stack} \rightarrow \text{Maybe Stack} \\ \llbracket \text{opEqual} \rrbracket_s^s \text{ time}_1 \text{ msg} &= \text{executeStackEquality} \end{aligned}$$

The function `executeStackEquality` : `Stack`  $\rightarrow$  `Maybe Stack` fails and returns `nothing` if the stack has height  $\leq 1$ , and otherwise compares the two top numbers on the stack, replacing them by 1 for true in case they are equal, and by 0 for false otherwise.

$\llbracket \_ \rrbracket_s^s$  is then lifted to the semantics of the instructions  $\llbracket \_ \rrbracket_s$  using a generic function `liftStackFun2StackState`:

$$\begin{aligned} \llbracket \_ \rrbracket_s &: \text{InstructionBasic} \rightarrow \text{StackState} \rightarrow \text{Maybe StackState} \\ \llbracket \text{op} \rrbracket_s &= \text{liftStackFun2StackState} \llbracket \text{op} \rrbracket_s^s \end{aligned}$$

As prerequisites for Sect 6.1, we define functions that define the operational semantics of further Bitcoin instructions used in this paper: `executeStackDup` : `Stack`  $\rightarrow$  `Maybe Stack` fails and returns `nothing` if the stack is empty; otherwise, a duplicate of the top element will be added onto the stack. The function `executeOpHash` : `Stack`  $\rightarrow$  `Maybe Stack` fails and returns `nothing` if the stack is empty; otherwise, the top element is replaced by its hash. `executeStackVerify` : `Stack`  $\rightarrow$  `Maybe Stack` fails and returns `nothing` if the stack is empty or the top element is 0; otherwise, it will remove the top element of the stack. `executeStackCheckSig` : `Stack`  $\rightarrow$  `Maybe Stack` fails and returns `nothing` if the height of the stack  $\leq 1$ . Otherwise it pops the two top elements from the stack, and considers them as a signature and public key. It decides whether the message given by the argument `msg` : `Msg` is correctly signed by these data, and pushes the Boolean result on the stack.

SCRIPT has instructions with more complex behaviour, an example is the instruction `OP_MULTISIG` which will be introduced in Sect. 6.2. Some instructions depend on cryptographic functions for hashing and checking signatures. We abstract away from their concrete definition and take them as parameters of the modules of the Agda code. This is not a problem in this paper, since the weakest preconditions only depend on the results returned by these functions, such as a check whether the part of the transaction to be signed is signed by a signature corresponding to a given public key.

General scripts are formalised in Agda as lists of instructions, `BitcoinScriptBasic`. Let  $p$  be a script. We define  $\llbracket p \rrbracket : \text{StackState} \rightarrow \text{Maybe StackState}$  by monadic composition, that is

- $\llbracket [] \rrbracket := \text{just}$ ,
  - for an instruction  $op$ , script  $q$  and  $s : \text{StackState}$  define  $\llbracket op :: q \rrbracket s := \llbracket op \rrbracket_s s \gg= \llbracket q \rrbracket$ .
- It follows that  $\forall s : \text{StackState}. \llbracket p ++ q \rrbracket s \equiv \llbracket p \rrbracket s \gg= \llbracket q \rrbracket$ .

We lift as well  $\llbracket p \rrbracket$  to  $s : \text{Maybe StackState}$  by defining  $\llbracket p \rrbracket^+ s := s \gg= \llbracket p \rrbracket$ .

Let

`StackStatePred` = `StackState`  $\rightarrow$  `Set`,

`StackPredicate` = `Time`  $\rightarrow$  `Msg`  $\rightarrow$  `Stack`  $\rightarrow$  `Set`, and

`stackPred2SPred` : `StackPredicate`  $\rightarrow$  `StackStatePred` be the obvious lifting.

## 4 Specifying Security of Bitcoin Scripts

In this section we argue that weakest precondition in the context of Hoare logic are the appropriate notion to express security properties in Subsect. 4.1. We provide a formalisation of weakest preconditions in Agda in Subsect. 4.2, and discuss how weakest preconditions can be generated automatically in Subsect. 4.3, leading to the claim that we need human-readable

descriptions of weakest preconditions. To support our verification, we develop a library for equational reasoning with Hoare triples in Subsect. 4.4.

#### 4.1 Weakest Precondition for Security

One widely used way to specify the correctness of imperative programs axiomatically is Hoare logic [26]. Hoare logic is based on pre- and postconditions. It works well for safety critical systems, where the set of inputs is controlled, and the aim is to guarantee a safe result. An example of a commercial system for writing safety critical systems using Hoare logic is SPARK 2014 [1].

However, when dealing with security aspects, in particular access control, Hoare logic in general is not sufficient. The issue is that for security it is necessary to guard against malicious entries to a program. We argue that weakest preconditions in the context of Hoare logic is an appropriate notion to specify security properties. A weakest precondition expresses that it is not only sufficient, but as well necessary for the postcondition to hold after executing the program.

To explain our point, we specify the intended correctness of the locking script `scriptPubKey` from Sect. 3. The intention, usually given by the user requirement, is that in order for a locking script to run successfully, we need to provide a public key *pbk* and a signature *sig* such that *pbk* hashes to the value `<pubKeyHash>` stored in the locking script, and that *sig* validates the signed message using *pbk*. The values *pbk* and *sig* need to be the top elements on the stack. If we also fix their order and allow the stack to have arbitrary values otherwise,<sup>9</sup> then we can express this condition as follows:

The two top elements of the stack are *pbk* and *sig*, *pbk* hashes to `<pubKeyHsh>`, and *sig* is a valid signature of the signed message w.r.t. *pbk*. (CondPBKH)

We can define the specification of the locking script `scriptPubKey` as the property that (CondPBKH) is the weakest precondition for the accepting postcondition. We will show in Sect. 5 that (CondPBKH) is indeed the weakest precondition of `scriptPubKey`, which verifies that `scriptPubKey` fulfils the specification.

Let us now consider a faulty locking script instead of `scriptPubKey`:

```
scriptPubKeyFaulty: OP_DUP OP_HASH160 <pubKeyHash> OP_EQUAL
```

To see that it does not fulfil the specification given above, consider the weakest precondition for `scriptPubKeyFaulty` for the accepting postcondition, which can be described by the following condition:

The top element of the stack is *pbk*, and *pbk* hashes to `<pubKeyHsh>`. (CondPBKHfaulty)

By inspection we see that (CondPBKHfaulty) is not equivalent to (CondPBKH), and therefore `scriptPubKeyFaulty` doesn't fulfil the specification. In fact we can identify states which satisfy (CondPBKHfaulty) but not (CondPBKH), e.g. a malicious attacker could just copy the public key of the sender onto the stack, which violates the user requirements of a locking script.

We observe that this example also demonstrates the inadequacy of general Hoare logic for the verification the security property of access control: Using standard Hoare logic, we can prove that (CondPBKH) is a precondition for the accepting postcondition for both `scriptPubKey` and `scriptPubKeyFaulty`.

---

<sup>9</sup> Bitcoin scripts do not put any requirements on the stack below the data required by the scripts.

As with all formal verification approaches, there remains a gap between the user's intention expressed as requirements, and what is expressed as a formal specification. This gap cannot be filled in a provably correct way, since requirements are a mental intention expressed in natural language. However, the gap can be narrowed by expressing the specification in a human-readable format so that the validation is as easy and clear as possible. Here, validation means showing that the specification guarantees the requirements, and is carried out by a human reader.

## 4.2 Formalising Weakest Preconditions in Agda

We now describe how weakest preconditions can be defined in Agda. Let a precondition  $\varphi$  and postcondition  $\psi$  be given, both of type `StackStatePred`. In order to accommodate `Maybe`, we define a postfix operator  $\_+$ , to lift  $\psi$  to  $(\psi^+)$ : `Maybe StackState`  $\rightarrow$  `Set`, defining  $(\psi^+) \text{ nothing} = \perp$  and  $(\psi^+) \circ \text{just} = \psi$ .

A Hoare triple, consisting of a precondition, a program, and a postcondition, expresses that if the precondition is satisfied before execution of the program, then the postcondition holds after executing it. We formalise Hoare triples as follows:

$$\langle \varphi \rangle p \langle \psi \rangle := \forall s \in \text{StackState}. \varphi(s) \rightarrow (\psi^+) (\llbracket p \rrbracket s)$$

Weakest preconditions express that the precondition not only is sufficient, but as well necessary for the postcondition to hold after executing the program:

$$\langle \varphi \rangle^{\text{iff}} p \langle \psi \rangle := \forall s \in \text{StackState}. \varphi(s) \leftrightarrow (\psi^+) (\llbracket p \rrbracket s)$$

Thus, for security the backwards direction of the equivalence in the previous formula is the important direction.

In Bitcoin we consider a locking script `scriptPubKey` and an unlocking script `scriptSig`, see Section 3.1. Let us fix an unlocking script `unlock` and a locking script `lock`. Let `init` be the initial state consisting of an empty stack, and let `acceptState` be the accepting condition expressing that the stack is non empty with top element being not false, i.e.  $>0$ . The combination of `unlock` and `lock` is accepted iff running `unlock` on `init` succeeds and running `lock` on the resulting stack results in a state that satisfies the accepting condition, i.e. iff  $(\text{acceptState}^+) (\llbracket lock \rrbracket^+ (\llbracket unlock \rrbracket \text{init}))$ . Note that Bitcoin does not run the concatenation of the two scripts, as it did in its first version, but runs first the unlocking scripts, and if it succeeds runs the locking script on the resulting stack. Let  $\varphi$  be the weakest precondition of `lock`, i.e.  $\langle \varphi \rangle^{\text{iff}} lock \langle \text{acceptState} \rangle$ . Then the acceptance condition is equivalent to  $(\varphi^+) (\llbracket unlock \rrbracket \text{init})$ . Thus, `unlock` succeeds iff running the unlocking script `unlock` on the initial state `init` produces a state fulfilling  $\varphi$ . Hence, by determining the weakest precondition for the locking script w.r.t. the accepting condition we have obtained a characterisation of the set of unlocking scripts which unlock the locking script. Note that we do not define inductively all successful unlocking scripts, since they could be arbitrary complex programs, but instead characterise them by the output they produce.

## 4.3 Automatically Generated Weakest Preconditions

We start by giving a direct method for defining the weakest precondition for any Bitcoin script by describing the set of states that lead to a given final state. We then apply this general method to a toy example to demonstrate that the description obtained in this way is usually not helpful for a human to judge whether the script has the right properties, thus making the case that the task must be to find (equivalent) human-readable descriptions.

Weakest preconditions can be defined by the simple definition

## 1:12 Verification of Bitcoin Script in Agda

```
weakestPreConds : BitcoinScriptBasic → StackStatePred → StackStatePred
weakestPreConds p φ s = (φ+) ([ p ] s)
```

Consider a simple toy program which removes the top element from the stack three times:

```
testprog = opDrop :: opDrop :: [ opDrop ]
```

Its weakest precondition can be computed as

```
weakestPreCondTestProg = weakestPreConds testprog acceptState
```

We obtain the following code (we slightly reformatted it to improve readability):

```
weakestPreCondTestProgNormalised s =
  (stackPred2SPred acceptStates +)
  (stackState2WithMaybe ⟨ currentTime s , msg s , executeStackDrop (stack s) ⟩
  >>= (λ s1 → stackState2WithMaybe ⟨ currentTime s1 , msg s1 , executeStackDrop (stack s1) ⟩
  >>= liftStackFun2StackState (λ time1 msg1 → executeStackDrop)))
```

This condition is difficult to understand. The reason is that each instruction may cause the program to abort in case the stack is empty. The condition expresses: if the stack is empty then the condition is false. Otherwise, if after dropping the top element the stack is empty the condition is false. Otherwise, if after dropping again the top element the stack is empty the condition is false. Otherwise the condition is true if after dropping again the top element the stack is non empty and the top element is not false. The readable condition would express that the height of the stack is  $\geq 4$  and the fourth element from the top is  $> 0$ . In this simple example simplifying the condition would be easy, but when using different instructions the situation becomes more complicated.

What we did using our methods to avoid this problem was to create the weakest precondition by starting from the end and improving it in each step, or by replacing the program by an easier program (which in case of this example would return nothing if the stack has height  $\leq 2$  and otherwise returns the result of dropping the first three elements off the stack). An interesting project for future work would be to automate the steps we carried out manually, and obtain readable weakest preconditions automatically.

### 4.4 Equational Reasoning with Hoare Triples

To support the verification of Bitcoin scripts with Hoare triples and weakest preconditions in Agda, we have developed a library in Agda for equational reasoning with Hoare triples. The library is inspired by what is described in Wadler et al. [48]. Let  $p, q$  be scripts and  $\phi, \phi', \psi, \psi' : \text{Predicate}$ . If we define  $\phi \langle \Rightarrow \rangle^P \psi := \forall s : \text{StackState}. \phi(s) \leftrightarrow \psi(s)$ , we can easily show

$$\begin{array}{l}
 \langle \phi \rangle^{\text{iff}} p \langle \psi \rangle \wedge \langle \psi \rangle^{\text{iff}} q \langle \rho \rangle \rightarrow \langle \phi \rangle^{\text{iff}} p ++ q \langle \rho \rangle \\
 \langle \phi \rangle^{\text{iff}} p \langle \psi \rangle \wedge \psi \langle \Rightarrow \rangle^P \psi' \rightarrow \langle \phi \rangle^{\text{iff}} p \langle \psi' \rangle \\
 \phi' \langle \Rightarrow \rangle^P \phi \wedge \langle \phi \rangle^{\text{iff}} p \langle \psi \rangle \rightarrow \langle \phi' \rangle^{\text{iff}} p \langle \psi \rangle
 \end{array}$$

We demonstrate our syntax by an example, assuming (using Agda postulate) programs `prog1`, `prog2`, `prog3`, and proofs

```
proof1 : < precondition >iff prog1 < intermediateCond1 >
proof2 : < intermediateCond1 >iff prog2 < intermediateCond2 >
proof3 : intermediateCond2 <=>P intermediateCond3
proof4 : < intermediateCond3 >iff prog3 < postcondition >
```

Then the proof for the Hoare triple for `prog1 ++ (prog2 ++ prog3)` is given in Agda as follows:<sup>10</sup>

```
theorem : < precondition >iff prog1 ++ (prog2 ++ prog3) < postcondition >
theorem = precondition      <><>< prog1 >\ proof1 >
      intermediateCond1 <><>< prog2 >\ proof2 >
      intermediateCond2 <=>< proof3 >
      intermediateCond3 <><>< prog3 >\ proof4 >e postcondition ■
```

## 5 Proof of Correctness of the P2PKH script using the Step-by-Step Approach

This section explains the usage of our approach by providing an example of how to prove the correctness of the P2PKH using step-by-step to obtain the weakest precondition. The P2PKH is the most used script in Bitcoin transactions. The locking script, which depends on a public key hash, is defined as follows:

```
scriptP2PKHb : (pbkh : ℕ) → BitcoinScriptBasic
scriptP2PKHb pbkh = opDup :: opHash :: (opPush pbkh) :: opEqual :: opVerify :: [ opCheckSig ]
```

In this section we develop a readable weakest precondition of the P2PKH script and prove its correctness by working backwards instruction by instruction.

Let `acceptState` be the accepting state where the stack is non-empty with top element  $>0$ . We define intermediate conditions `accept1`, `accept2`, etc, the weakest precondition `wPreCondP2PKH`, and proofs `correct-opCheckSig`, `correct-opVerify` etc of corresponding Hoare triples w.r.t. the instructions of the Bitcoin script, working backwards starting from the last instruction `opCheckSig`:

```
correct-opCheckSig : < accept1 >iff ([ opCheckSig ] < acceptState >
correct-opVerify : < accept2 >iff ([ opVerify ] < accept1 >
correct-opEqual : < accept3 >iff ([ opEqual ] < accept2 >
correct-opPush : (pbkh : ℕ) → < accept4 pbkh >iff ([ opPush pbkh ] < accept3 >
correct-opHash : (pbkh : ℕ) → < accept5 pbkh >iff ([ opHash ] < accept4 pbkh >
correct-opDup : (pbkh : ℕ) → < wPreCondP2PKH pbkh >iff ([ opDup ] < accept5 pbkh >
```

The intermediate conditions can be read off from the operations. We present them in mathematical notation below, using the following conventions and abbreviations:  $t : \mathbb{N}$  denotes time,  $m : \text{Msg}$ ,  $st, st' : \text{Stack}$ ,  $x : \mathbb{N}$ ; for brevity we omit types after  $\exists$  quantifiers. We use here and in the remaining paper <sup>s</sup> for operations where the `StackState` argument has been unfolded into its components.

<sup>10</sup>In the last step we use  $\rangle^e$  instead of  $\rangle$ . This avoids concatenating the program with `[]`. If we used  $\rangle$ , the theorem would prove the condition for program `prog1++(prog2++(prog3++[]))`, which is provably but not definitionally equal to the original program, requiring an additional proof step.

## 1:14 Verification of Bitcoin Script in Agda

$$\begin{aligned}
\text{acceptState}^s t m st &\Leftrightarrow \exists x, st'. && st \equiv x :: st' \wedge x > 0 \\
\text{accept}_1^s t m st &\Leftrightarrow \exists pbk, sig, st'. && st \equiv pbk :: sig :: st' \\
&&& \wedge \text{IsSigned } m \text{ sig } pbk \\
\text{accept}_2^s t m st &\Leftrightarrow \exists x, pbk, sig, st'. && st \equiv x :: pbk :: sig :: st' \\
&&& \wedge x > 0 \wedge \text{IsSigned } m \text{ sig } pbk \\
\text{accept}_3^s t m st &\Leftrightarrow \exists pbkh_2, pbkh_1, pbk, sig, st'. && st \equiv pbkh_2 :: pbkh_1 :: pbk :: sig :: st' \\
&&& \wedge pbkh_2 \equiv pbkh_1 \wedge \text{IsSigned } m \text{ sig } pbk \\
\text{accept}_4^s pbkh_1 t m st &\Leftrightarrow \exists pbkh_2, pbk, sig, st'. && st \equiv pbkh_2 :: pbk :: sig :: st' \\
&&& \wedge pbkh_2 \equiv pbkh_1 \wedge \text{IsSigned } m \text{ sig } pbk \\
\text{accept}_5^s pbkh_1 t m st &\Leftrightarrow \exists pbk_1, pbk, sig, st'. && st \equiv pbk_1 :: pbk :: sig :: st' \\
&&& \wedge \text{hashFun } pbk_1 \equiv pbkh_1 \wedge \text{IsSigned } m \text{ sig } pbk \\
\text{wPreCondP2PKH}^s pbkh_1 t m st &\Leftrightarrow \exists pbk, sig, st'. && st \equiv pbk :: sig :: st' \\
&&& \wedge \text{hashFun } pbk \equiv pbkh_1 \wedge \text{IsSigned } m \text{ sig } pbk
\end{aligned}$$

In Agda, these formulas are defined by case distinction on the stack. As examples, the code for the accept condition (`acceptState`) and the weakest precondition (`wPreCondP2PKHs`) is as follows:

```

acceptStates : StackPredicate
acceptStates time msg1 [] = ⊥
acceptStates time msg1 (x :: stack1) = NotFalse x

wPreCondP2PKHs : (pbkh : ℕ) → StackPredicate
wPreCondP2PKHs pbkh time m [] = ⊥
wPreCondP2PKHs pbkh time m (x :: []) = ⊥
wPreCondP2PKHs pbkh time m (pbk :: sig :: st) =
  (hashFun pbk ≡ pbkh) ∧ IsSigned m sig pbk

```

Using our syntax for equational reasoning, we can prove the weakest precondition for the P2PKH script as follows:

```

theoremP2PKH : (pbkh : ℕ) → < wPreCondP2PKH pbkh >iff scriptP2PKHb pbkh < acceptState >
theoremP2PKH pbkh = wPreCondP2PKH pbkh <><>< [ opDup ] >>< correct-opDup pbkh >
  accept5 pbkh <><>< [ opHash ] >>< correct-opHash pbkh >
  accept4 pbkh <><>< [ opPush pbkh ] >>< correct-opPush pbkh >
  accept3 <><>< [ opEqual ] >>< correct-opEqual >
  accept2 <><>< [ opVerify ] >>< correct-opVerify >
  accept1 <><>< [ opCheckSig ] >>< correct-opCheckSig >e
  acceptState ■p

```

The locking script will be accepted if, after executing the code starting with the stack returned by the unlocking script, the accept condition `acceptState` is fulfilled. The verification conditions and proofs were developed by working backwards starting from the last instruction and determining the weakest preconditions “`accepti`” w.r.t. the end piece of the script starting with that instruction and the accept condition as post-condition. The preconditions were obtained manually – one could automate this by determining for each instruction depending on the post-condition a corresponding pre-condition, where the challenge would be to simplify the resulting pre-conditions in order to avoid a blowup in size. We continued in this way until we reached the first instruction and obtained the weakest precondition for the locking script.

`theoremP2PKH` is using single instructions in order to prove the correctness of P2PKH. The proofs `correct-opCheckSig`, `correct-opVerify`, etc are done by following the case distinctions made in the corresponding verification conditions. The harder direction is to prove that they are actually *weakest* preconditions: Proving that the precondition implies the postcondition after running the program, is easier since we are used to mentally executing programs in forward direction. Proving the opposite direction requires showing that the only way, after running the program, to obtain the postcondition is to have the precondition fulfilled, which requires mentally reversing the execution of programs.

## 6 Proof of Correctness using Symbolic Execution

In this section, we will introduce a second method for obtaining readable representations of weakest preconditions of Bitcoin scripts. This method is based on symbolic execution [29] of the Bitcoin script, and investigating the sequence of case distinctions carried out during the execution. We will consider three examples: The first will be the P2PKH script which we analysed already. We use it to explain the method and provide a second approach to determine and verify the already obtained weakest precondition. The second example will consider the multisig script which is a direct application of the `OP_MULTISIG` instruction. The third example will see an application of a combination of both methods.

### 6.1 Example: P2PKH Script

When applying the symbolic evaluation method to the P2PKH script and analysing the sequence of case distinctions carried out, we will see that there will be exactly one path through the tree of case distinctions which results in an accepting condition. The conjunction of the cases that determine this path will form the weakest precondition. In examples with more than one accepting path we would take the disjunction of the conditions for each accepting path.<sup>11</sup> We will prove that the precondition is indeed the weakest by developing an equivalent program `p2pkhFunctionDecoded` and showing that it fulfils the weakest precondition.

We start by declaring (using Agda’s `postulate`) symbolic values `pbkh`, `msg1`, `stack1`, `x1`, etc for the parameters (postulates are typeset in blue). This allows us to evaluate expressions up to `executeStackVerify` symbolically by using the normalisation procedure of Agda and to determine the function `p2pkhFunctionDecoded`. (In Sect. 7 we will elaborate how to do this practically in Agda). Afterwards, we stop using those postulates (they were defined as `private`) and prove that the result of evaluating the P2PKH script for arbitrary parameters is equivalent to `p2pkhFunctionDecoded`.

When evaluating  $\llbracket \text{scriptP2PKH}^b \text{ pbkh} \rrbracket^s \text{ time}_1 \text{ msg}_1 \text{ stack}_1$  we obtain

```
executeStackDup stack1                >>=
λ stack2 → executeOpHash stack2      >>=
λ stack3 → executeStackEquality (pbkh :: stack3) >>=
```

<sup>11</sup>In our examples we got only a few accepting paths, since concrete scripts in use are designed to deal with a small number of different scenarios for unlocking them, so the majority of paths in the program are unsuccessful paths. It could happen however that with more advanced examples nested conditions result in an exponential blowup of the number of cases – if that occurs one would need to take an approach where the nested case distinctions are preserved at least partly and the resulting extracted formulas reflect those nested case distinctions rather than flattening them out. This would avoid the blowup in the size of the resulting weakest precondition.

## 1:16 Verification of Bitcoin Script in Agda

$$\begin{aligned} \lambda \text{ stack}_4 &\rightarrow \text{executeStackVerify } \text{stack}_4 && \gg= \\ \lambda \text{ stack}_5 &\rightarrow \text{executeStackCheckSig } \text{msg}_1 \text{ stack}_5 \end{aligned}$$

We can write it equivalently using the do notation<sup>12</sup>

```
do stack2 ← executeStackDup stack1
   stack3 ← executeOpHash stack2
   stack4 ← executeStackEquality (pbkh :: stack3)
   stack5 ← executeStackVerify stack4
   executeStackCheckSig msg1 stack5
```

At this point further reduction is blocked by the first line of the previous expression, because `executeStackDup stack1` makes a case distinction on `stack1`. Therefore, we introduce a symbolic case distinction on `stack1`:

- `[[ scriptP2PKHb pbkh ]]s time1 msg1 []` evaluates to `nothing`.
- `[[ scriptP2PKHb pbkh ]]s time1 msg1 (pbk :: stack1)` evaluates to what in do notation can be written as

```
do stack5 ← executeStackVerify (compareNaturals pbkh (hashFun pbk) :: pbk :: stack1)
   executeStackCheckSig msg1 stack5
```

Evaluation of the latter expression is blocked by the function `executeStackVerify` which makes a case distinction on the expression `compareNaturals pbkh (hashFun pbk)`. We define

```
abstrFun : (stack1 : Stack)(cmp : ℕ) → Maybe Stack
abstrFun stack1 cmp = do stack5 ← executeStackVerify (cmp :: pbk :: stack1)
   executeStackCheckSig msg1 stack5
```

hence `[[ scriptP2PKHb pbkh ]]s time1 msg1 (pbk :: stack1)` evaluates to `abstrFun stack1 (compareNaturals pbkh (hashFun pbk))`.

Next we carry out a symbolic case distinction on the argument `cmp` of `abstrFun`:

- `abstrFun stack1 0` evaluates to `nothing`.
- `abstrFun stack1 (suc x1)` evaluates to `executeStackCheckSig msg1 (pbk :: stack1)`.

In order to normalise further, `executeStackCheckSig` needs to make a case distinction on `stack1`, so we carry out a symbolic case distinction on that argument:

- `abstrFun [] (suc x1)` evaluates to `nothing`.
- `abstrFun (sig1 :: stack1) (suc x1)` evaluates to `just (boolToNat (isSigned msg1 sig1 pbk) :: stack1)`

We can now read off the weakest precondition. The only path which ends up in a `just` result is when the stack is non empty of the form `pbk :: stack1`, and `compareNaturals pbkh (hashFun pbk)` evaluates to `suc x1`, i.e. it must be `>0`. Furthermore, in this case `stack1` needs to be itself non empty. For `stack1 = sig1 :: stack2`, the result returned is `just (boolToNat (isSigned msg1 sig1 pbk) :: stack1)`, which fulfils the accept condition if `boolToNat (isSigned msg1 sig1 pbk) > 0`. The latter is the case if `isSigned msg1 sig1 pbk` is true.

---

<sup>12</sup>The do notation is a widely used Haskell notation adapted to Agda, which provides an alternative syntax for the same expression making it appear as an imperative program if one reads `←` as assignments. It demonstrates that we are consecutively executing the instructions, with the possibility of aborting in each step.



Furthermore, `compareNaturals`  $n\ m$  returns 1 if  $n, m$  are equal otherwise 0, so it is  $>0$  if  $n = m$ . Therefore the P2PKH locking script succeeds with an output stack fulfilling the acceptance condition, if and only if the input stack has height at least two, and if it is `pbk :: sig1 :: stack2`, then `pbkh` is equal to `hashFun pbk`, and `isSigned msg1 sig1 pbk` is true. That is the same as the weakest precondition that we determined using the first approach.

In order to prove correctness, we first determine a more Agda style formulation of the result of evaluation of the P2PKH script, which we derive from the previous symbolic evaluation:

```

p2pkhFunctionDecoded : (pbkh : ℕ)(msg1 : Msg)(stack1 : Stack) → Maybe Stack
p2pkhFunctionDecoded pbkh msg1 [] = nothing
p2pkhFunctionDecoded pbkh msg1 (pbk :: stack1) = p2pkhFunctionDecodedAux1 pbk msg1 stack1
                                                    (compareNaturals pbkh (hashFun pbk))

p2pkhFunctionDecodedAux1 : (pbk : ℕ)(msg1 : Msg)(stack1 : Stack)(cpRes : ℕ) → Maybe Stack
p2pkhFunctionDecodedAux1 pbk msg1 [] cpRes = nothing
p2pkhFunctionDecodedAux1 pbk msg1 (sig1 :: stack1) zero = nothing
p2pkhFunctionDecodedAux1 pbk msg1 (sig1 :: stack1) (suc cpRes) =
  just (boolToNat (isSigned msg1 sig1 pbk) :: stack1)

```

We prove that this function is equivalent to the result of evaluating the P2PKH script. The proof is a simple case distinction following the cases defining `p2pkhFunctionDecoded`:

```

p2pkhFunctionDecodedcor : (time1 : ℕ) (pbkh : ℕ)(msg1 : Msg)(stack1 : Stack)
  → [ scriptP2PKHb pbkh ]s time1 msg1 stack1 ≡ p2pkhFunctionDecoded pbkh msg1 stack1

```

We show that the extracted weakest precondition is a correct for the extracted program:<sup>13</sup>

```

lemmaPTKHcoraux : (pbkh : ℕ) → < weakestPreConditionP2PKHs pbkh >gs
  (λ time msg1 s → p2pkhFunctionDecoded pbkh msg1 s)
  < acceptStates >

```

Afterwards, this is transferred into a proof of the weakest precondition for the P2PKH script, using the equality proof from before:

```

theoPTPKHcor : (pbkh : ℕ) → < wPreCondP2PKH pbkh >iff scriptP2PKHb pbkh < acceptState >

```

Carrying out the symbolic execution was relatively easy, because Agda supports evaluation of terms very well. It only becomes relatively long in the Agda code [46] when documenting all the steps, which we did in order to explain how this is done in detail. What matters is the resulting program and a prove that it is equivalent, which was relatively short and easy. Maybe Agda's reflection mechanism [3], once it is more fully developed, could be of help to find the successful branches of the program more easily. To obtain a readable program rather than a machine generated program, and therefore readable verification conditions, would however require a lot of work, and probably require delegating some programming tasks from Agda (in which tactics need to be written) to its foreign language interface.

<sup>13</sup>  $\langle \_ \rangle^g \langle \_ \rangle$  is the generalisation of  $\langle \_ \rangle^{iff} \langle \_ \rangle$  where Bitcoin scripts are replaced by Agda functions `StackState`  $\rightarrow$  `Maybe StackState`;  $\langle \_ \rangle^{g^s} \langle \_ \rangle$  is the version, where the `StackState` is unfolded into its components.

## 6.2 Example: MultiSig Script (P2MS)

The `OP_MULTISIG` instruction is an instruction which has a more complex behaviour: It assumes that the top elements of the stack are as follows:

$$n :: pbk_n :: \dots :: pbk_2 :: pbk_1 :: m :: sig_m :: \dots :: sig_2 :: sig_1 :: dummy$$

`OP_MULTISIG` checks whether the  $m$  signatures are signatures corresponding to  $m$  of the  $n$  public keys for the `msg` to be signed, where the matching public keys are in the same order as the signatures. Observe that when pushed from a script, the public keys and signatures appear in reverse order on the stack, as  $pbk_1$  is pushed first onto the stack, etc. The *dummy* element occurs due to a mistake in the Bitcoin protocol, which has not been corrected as it would require a hard fork [6, p. 151-152].

The operational semantics is given by a function `executeMultiSig`, which fetches the data from the stack as described before. It fails if there are not enough elements on the stack and otherwise returns `just (boolToNat (cmpMultiSigs msg sigs pbks) :: restStack)`, where `sigs` and `pbks` are the signatures and public keys fetched from the stack in reverse order, and `restStack` is the remainder of the stack. The function `cmpMultiSigs` compares whether signatures correspond to public keys and is defined as follows:

```

cmpMultiSigs : (msg : Msg)(sigs pbks : List N) → Bool
cmpMultiSigs msg [] pubkeys           = true
cmpMultiSigs msg (sig :: sigs) []     = false
cmpMultiSigs msg (sig :: sigs) (pbk :: pbks) = cmpMultiSigsAux msg sigs pbks sig (isSigned msg sig pbk)

cmpMultiSigsAux : (msg : Msg)(sigs pbks : List N)(sig : N)(testRes : Bool) → Bool
cmpMultiSigsAux msg sigs pbks sig false = cmpMultiSigs msg (sig :: sigs) pbks
cmpMultiSigsAux msg sigs pbks sig true  = cmpMultiSigs msg sigs pbks

```

We define now a generic multisig function. First we define `opPushList`, which pushes a list of public keys on the stack:

```

opPushList : (pbkList : List N) → BitcoinScriptBasic
opPushList [] = []
opPushList (pbk1 :: pbkList) = opPush pbk1 :: opPushList pbkList

```

The  $m$  out of  $n$  multi-signature script P2MS ( $n = \text{length } pbkList$ ) is defined as follows:

```

multiSigScriptm-nb : (m : N)(pbkList : List N)(m < n : m < length pbkList)
  → BitcoinScriptBasic
multiSigScriptm-nb m pbkList m < n =
  opPush m :: (opPushList pbkList ++ (opPush (length pbkList) :: [ opMultiSig ]))

```

The locking script MultiSig script P2MS applies `OP_MULTISIG` to  $m$  signatures and  $n$  public keys. It pushes the number  $m$  of required signatures, then  $n$  public keys, and then the number  $n$  as the number of public keys, onto the stack, and executes `OP_MULTISIG`. If `OP_MULTISIG` finds that the  $m$  signatures are valid signature for the message to be signed for  $m$  out of the  $n$  public keys in the same order as they appear in the list of public keys, then the script will be unlocked. As unlocking script one can use `opPushList` applied to a list of  $m$  appropriate signatures. In order to verify the script we will consider the concrete example of the 2-out-of-4 P2MS, for which we obtain a very readable verification condition (the generic one becomes difficult to read).

We will use the second approach of determining a readable form of the weakest precondition and proving correctness by symbolic evaluation for the 2 out of 4 `multiSigScript2-4b`. The first approach is difficult to carry out since the instruction `opMultiSig` has a very complex precondition that is difficult to handle – it requires that the stack contains the number of public keys, then the public keys themselves, then the number of signatures and the signatures, and a dummy element, where the number of public keys and number of signatures can be arbitrary. It is much easier to handle the full `multiSigScript2-4b` script, since, after the data has been inputted, the number of required signatures is known, and the public keys are already provided by the script.

In order to demonstrate the first approach we will instead, in Subsect. 6.3, apply the step-by-step approach to a combined script, of which `multiSigScript2-4b` is one part. This way we obtain a readable form of the weakest precondition and can then prove its correctness. This will demonstrate that in some cases it is beneficial to interleave the two processes, and apply the second method to sequences of instructions while applying the first approach to the resulting sequences of instructions instead of single instructions.

We start the symbolic evaluation by computing the normal form of

```
[[ multiSigScript2-4b pbk1 pbk2 pbk3 pbk4 ]]s time1 msg1 stack1
```

and obtain

```
executeMultiSig3 msg1 (pbk1 :: pbk2 :: pbk3 :: [ pbk4 ]) 2 stack1 []
```

Here, `executeMultiSig3` is one of the auxiliary functions in the definition of `executeMultiSig`.

That expression makes a case distinction on `stack1` and returns:

- `nothing` when the stack has height at most 2 (obtained by evaluating it symbolically for stacks of height 0, 1, 2).
- Otherwise, the stack has height  $\geq 3$ , and, if it is of the form `sig2 :: sig1 :: dummy :: stack1`, it reduces to

```
just (boolToNat (cmpMultiSigsAux msg1 [ sig2 ] (pbk2 :: pbk3 :: [ pbk4 ]) sig1
(isSigned msg1 sig1 pbk1)) :: stack1)
```

The script has terminated, because we obtain `just` as a result of the evaluation. We now need to check whether the result fulfils the accept condition. For this the top element of the stack needs to be  $>0$ , which is the case if

```
cmpMultiSigsAux msg1 [ sig2 ] (pbk2 :: pbk3 :: [ pbk4 ]) sig1(isSigned msg1 sig1 pbk1)
```

returns `true`. Therefore, we perform symbolic case distinctions in the following way:

- In case `isSigned msg1 sig1 pbk1` evaluates to `true`, i.e. if we replace that expression by `true`, the reduction continues to `cmpMultiSigsAux msg1 [] (pbk3 :: [ pbk4 ]) sig2 (isSigned msg1 sig2 pbk2)`, which makes a case distinction on `isSigned msg1 sig2 pbk2`.
  - If that expression returns again `true`, we obtain `true`.
  - If it returns false, we obtain `cmpMultiSigsAux msg1 [] [ pbk4 ] sig2 (isSigned msg1 sig2 pbk3)` which makes a case distinction on `isSigned msg1 sig2 pbk3`
    - \* In case of `true`, we obtain `true`.
    - \* Otherwise the case distinctions continue, see the git repository [46] for full details.

In total we see that we obtain `true` iff one of the following cases holds:

- $(\text{isSigned } \text{msg}_1 \text{ sig}_1 \text{ pbk}_1) \wedge (\text{isSigned } \text{msg}_1 \text{ sig}_2 \text{ pbk}_2)$
- $(\text{isSigned } \text{msg}_1 \text{ sig}_1 \text{ pbk}_1) \wedge \neg (\text{isSigned } \text{msg}_1 \text{ sig}_2 \text{ pbk}_2) \wedge (\text{isSigned } \text{msg}_1 \text{ sig}_2 \text{ pbk}_3)$
- $(\text{isSigned } \text{msg}_1 \text{ sig}_1 \text{ pbk}_1) \wedge \neg (\text{isSigned } \text{msg}_1 \text{ sig}_2 \text{ pbk}_2) \wedge \neg (\text{isSigned } \text{msg}_1 \text{ sig}_2 \text{ pbk}_3) \wedge (\text{isSigned } \text{msg}_1 \text{ sig}_2 \text{ pbk}_4)$

- ... more cases.

These cases can be simplified to an equivalent disjunction of the following cases:

- $(\text{isSigned } \text{msg}_1 \text{ sig}_1 \text{ pbk}_1) \wedge (\text{isSigned } \text{msg}_1 \text{ sig}_2 \text{ pbk}_2)$
- $(\text{isSigned } \text{msg}_1 \text{ sig}_1 \text{ pbk}_1) \wedge (\text{isSigned } \text{msg}_1 \text{ sig}_2 \text{ pbk}_3)$
- $(\text{isSigned } \text{msg}_1 \text{ sig}_1 \text{ pbk}_1) \wedge (\text{isSigned } \text{msg}_1 \text{ sig}_2 \text{ pbk}_4)$
- ... more cases.

We obtain the following weakest precondition as a stack predicate:

```

weakestPreCondMultiSig-2-4s : (pbk1 pbk2 pbk3 pbk4 : ℕ) → StackPredicate
weakestPreCondMultiSig-2-4s pbk1 pbk2 pbk3 pbk4 time msg1 [] = ⊥
weakestPreCondMultiSig-2-4s pbk1 pbk2 pbk3 pbk4 time msg1 (x :: []) = ⊥
weakestPreCondMultiSig-2-4s pbk1 pbk2 pbk3 pbk4 time msg1 (x :: y :: []) = ⊥
weakestPreCondMultiSig-2-4s pbk1 pbk2 pbk3 pbk4 time msg1 (sig2 :: sig1 :: dummy :: stack1) =
  ( (IsSigned msg1 sig1 pbk1 ∧ IsSigned msg1 sig2 pbk2) ⊕
    (IsSigned msg1 sig1 pbk1 ∧ IsSigned msg1 sig2 pbk3) ⊕
    (IsSigned msg1 sig1 pbk1 ∧ IsSigned msg1 sig2 pbk4) ⊕
    (IsSigned msg1 sig1 pbk2 ∧ IsSigned msg1 sig2 pbk3) ⊕
    (IsSigned msg1 sig1 pbk2 ∧ IsSigned msg1 sig2 pbk4) ⊕
    (IsSigned msg1 sig1 pbk3 ∧ IsSigned msg1 sig2 pbk4))

```

It expresses that the stack must have height at least 3, and if it is of the form  $\text{sig}_2 :: \text{sig}_1 :: \text{dummy} :: \text{stack}_1$  then the signatures need to correspond to 2 out of the 4 public keys in the same order as the public keys. Using the same case distinctions as they occurred in the symbolic evaluation above we can now prove:

```

theoremCorrectnessMultiSig-2-4 : (pbk1 pbk2 pbk3 pbk4 : ℕ)
  → < stackPred2SPred (weakestPreCondMultiSig-2-4s pbk1 pbk2 pbk3 pbk4) >iff
    multiSigScript2-4b pbk1 pbk2 pbk3 pbk4
    < stackPred2SPred acceptStates >

```

From the theorem above, we have obtained a readable weakest precondition by symbolic execution, which will be used as a starting template for developing a generic verification. The next step would be to generalise the verification conditions and theorems to the generic case, however that would go beyond the scope of the current paper.

### 6.3 Example: Combining the two Methods

In this subsection we show how to verify a combined script which consists of a simple script checking a certain amount of time has passed and the multisig script from the previous subsection. To determine a readable form of the weakest precondition and proving correctness we will combine both of our techniques: The weakest precondition for the multisig script has been determined by symbolic evaluation in the previous subsection. The weakest precondition for the simple time checking script will be obtained directly, as it is very simple. When we consider the combined scripts we will use the first method of moving backwards step-by-step. However, instead of using single instructions in each step, we now use several instructions as a single step.

We define the checktime script as follows:

```

checkTimeScriptb : (time1 : Time) → BitcoinScriptBasic
checkTimeScriptb time1 = (opPush time1) :: opCHECKLOCKTIMEVERIFY :: [ opDrop ]

```

If we define

```
timeCheckPreCond : (time1 : Time) → StackPredicate
timeCheckPreCond time1 time2 msg stack1 = time1 ≤ time2
```

we can define its weakest precondition relative to a post condition  $\phi$  only affecting the stack as in the following theorem:

```
theoremCorrectnessTimeCheck : (ϕ : StackPredicate)(time1 : Time)
→ < stackPred2SPred (timeCheckPreCond time1 ∧ sp ϕ) >iff checkTimeScriptb time1
  < stackPred2SPred ϕ >
```

Now we can determine the weakest precondition for the combined script and prove its correctness as follows:

```
theoremCorrectnessCombinedMultiSigTimeCheck : (time1 : Time) (pbk1 pbk2 pbk3 pbk4 : ℕ)
→ < stackPred2SPred ( timeCheckPreCond time1 ∧ sp
  weakestPreCondMultiSig-2-4s pbk1 pbk2 pbk3 pbk4 ) >iff
  checkTimeScriptb time1 ++ multiSigScript2-4b pbk1 pbk2 pbk3 pbk4
  < acceptState >
theoremCorrectnessCombinedMultiSigTimeCheck time1 pbk1 pbk2 pbk3 pbk4 =
  stackPred2SPred (timeCheckPreCond time1 ∧ sp
    weakestPreCondMultiSig-2-4s pbk1 pbk2 pbk3 pbk4)
  <><>< checkTimeScriptb time1 >< theoremCorrectnessTimeCheck
    (weakestPreCondMultiSig-2-4s pbk1 pbk2 pbk3 pbk4) time1 >
  stackPred2SPred (weakestPreCondMultiSig-2-4s pbk1 pbk2 pbk3 pbk4)
  <><>< multiSigScript2-4b pbk1 pbk2 pbk3 pbk4
    >< theoremCorrectnessMultiSig-2-4 pbk1 pbk2 pbk3 pbk4 >e
  stackPred2SPred acceptStates ■ p
```

The weakest precondition states that the state time is  $\geq time_1$ , and that the weakest precondition for the multisig script is fulfilled ( $\wedge sp$  forms the conjunction of the two conditions). For proving it we used a combination of both methods, the second method was used to determine preconditions for the two parts of the scripts, and the first method, where we used whole scripts instead of basic instructions, was used to determine the combined weakest precondition.

## 7 Using Agda to Determine Readable Weakest Preconditions

Our library provides the operational semantics for (a subset of) Bitcoin SCRIPT, and a framework for specifying and reasoning about weakest preconditions. The Agda user has to specify the script to be verified, and then consider suitable pieces of the specified script and provide weakest preconditions. Agda will then create goals, which are unimplemented holes in the code. Agda will display the type of goals and list of assumptions available for solving them, and provide considerable additional support for resolving those goals. For instance, it allows to refine partial solutions provided by the user by applying it to sufficiently many new goals. Agda will as well automatically create case distinctions (such as whether an element of type `Maybe` is `just` or `nothing`). Agda can solve goals if the solution is unique and can be found in a direct way. Agda's automated theorem proving support for finding solutions which are not unique is not very strong due to the high complexity of the language.

Agda Reflection [3] is an ongoing project which already now provides a considerable library for inspecting code inside a goal and computing solutions as Agda code. The aim is to provide something similar to Coq's tactic language. In our code we frequently had to consider a nested case distinction for proving a goal, where most cases were solved because at one point one of the arguments became an element of the empty type. Automating this using Agda Reflection would make it much easier to use our library.

Finding a description of the weakest precondition has to be done manually at the moment. We plan to create a library which computes such descriptions for instructions or small pieces of instructions. Sometimes it is easier to provide weakest precondition for small pieces of code, for instance in case of the multisig instruction the weakest precondition for the instruction itself is very complex, whereas the weakest precondition for the P2MS script is much easier to display. Defining and simplifying the weakest preconditions in the intermediate steps has to be done manually at the moment. Proofs have to be done manually in Agda, but they are relatively easy because of Agda's support for developing proofs. It would be desirable to have a more automated support, where the user only needs to specify the verification conditions, but proofs are carried out automatically. In general our impression is that for writing programs and specifying verification conditions Agda is very suitable: one obtains code which is very readable and close to standard mathematical notations. Where Agda is lacking is in providing support for machine assisted proofs of the resulting conditions.

Regarding the question, which of the two approaches to use (working backwards step-by-step or using symbolic evaluation), we have only some heuristics at the moment. A good approach is that for pieces of code, where one has an intuition what the underlying program written in Agda could be, the symbolic evaluation is more suitable. For longer code, a good strategy is to cut the code into suitable pieces, for which one can find a symbolic program and weakest preconditions, and then work oneself backwards using the first approach starting from the acceptance condition. Note that symbolic execution can be done very fast: The user postulates variables for the arguments, applies the functions to be evaluated to those postulated arguments and then executes Agda's normalisation mechanism. Then the user needs to manually inspect the result to see which sub expression trigger the case distinction. It would be nice project to develop a procedure which automates that process of symbolic execution – this could be applicable to verification of other kinds of programs as well.

## **8** Conclusion

In this paper, we have implemented and tested two methods for developing human-readable weakest preconditions and proving their correctness. These methods can help smart contract developers to fill the validation gap between user requirements and formal specification. We have argued that weakest preconditions in Hoare logic is the correct notion for specifying the security property of access control. We have applied our approaches to P2PKH, P2MS, and a combination of P2MS with a time lock. The whole approach has been formalised in Agda [46].

In future work, we will treat non-local instructions such as `OP_IF`, `OP_ELSE`, and `OP_ENDIF`, and will formalise key instructions to extend our approach to the whole of Bitcoin `SCRIPT`. The difficulty is nesting of conditionals, and that Bitcoin scripts are not structured, and therefore some additional work needs to be done to find the matching of if-then-else instructions. In our approach we will use an expanded state space for dealing with those conditionals. Moreover, we plan to expand our library to support finding weakest preconditions for scripts having conditionals in a modular way. Furthermore, we plan the make the process of script

verification more user-friendly by using a text parser that can record the instructions used for verification.

In addition, we aim to generalise the verification of P2MS to arbitrary  $m$  out of  $n$  multiscripts, where the challenge is finding a suitable generic human-readable weakest precondition.

Another route for future research is to develop our approach into a framework for developing smart contracts that are correct by construction. One way to build such smart contracts is to use Hoare Type Theory [27, 38].

---

**References**

---

- 1 Adacore. SPARK 2014, Retrieved 9 November 2021. <https://www.adacore.com/about-spark>.
- 2 Agda Team. Agda documentation, Retrieved 21 April 2022. URL: <https://agda.readthedocs.io/en/latest/index.html>.
- 3 Agda Team. Agda Reflection, Retrieved 21 April 2022. URL: <https://agda.readthedocs.io/en/latest/language/reflection.html>.
- 4 Mouhamad Almakhour, Layth Sliman, Abed Ellatif Samhat, and Abdelhamid Mellouk. Verification of smart contracts: A survey. *Pervasive and Mobile Computing*, 67:1–19, 2020. doi:<http://dx.doi.org/10.1016/j.pmcj.2020.101227>.
- 5 Sidney Amani, Myriam Bégel, Maksym Bortin, and Mark Staples. Towards Verifying Ethereum Smart Contract Bytecode in Isabelle/HOL. In *Proceedings of the 7th ACM SIGPLAN International Conference on Certified Programs and Proofs, CPP 2018*, page 66–77, New York, NY, USA, 2018. Association for Computing Machinery. doi:<https://doi.org/10.1145/3167084>.
- 6 Andreas M Antonopoulos. *Mastering Bitcoin: Programming the open blockchain*. " Second ed. O'Reilly Media, Inc.", 2017.
- 7 Nicola Atzei, Massimo Bartoletti, Stefano Lande, and Roberto Zunino. A formal model of bitcoin transactions. In *Financial Cryptography and Data Security*, pages 541–560, Berlin, Heidelberg, 2018. Springer Berlin Heidelberg. doi:[https://doi.org/10.1007/978-3-662-58387-6\\_29](https://doi.org/10.1007/978-3-662-58387-6_29).
- 8 Massimo Bartoletti and Roberto Zunino. BitML: A Calculus for Bitcoin Smart Contracts. In *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security, CCS '18*, page 83–100, New York, NY, USA, 2018. Association for Computing Machinery. doi:<http://dx.doi.org/10.1145/3243734.3243795>.
- 9 Bruno Bernardo, Raphaël Cauderlier, Zhenlei Hu, Basile Pesin, and Julien Tesson. Mi-Chocoq, a Framework for Certifying Tezos Smart Contracts. In *Formal Methods. FM 2019 International Workshops*, pages 368–379, Cham, 2020. Springer International Publishing. doi:[https://doi.org/10.1007/978-3-030-54994-7\\_28](https://doi.org/10.1007/978-3-030-54994-7_28).
- 10 Karthikeyan Bhargavan, Antoine Delignat-Lavaud, Cédric Fournet, Anitha Gollamudi, Georges Gonthier, Nadim Kobeissi, Natalia Kulatova, Aseem Rastogi, Thomas Sibut-Pinote, Nikhil Swamy, and Santiago Zanella-Béguelin. Formal Verification of Smart Contracts: Short Paper. In *Proceedings of the 2016 ACM Workshop on Programming Languages and Analysis for Security, PLAS '16*, page 91–96, New York, NY, USA, 2016. Association for Computing Machinery. doi:<http://dx.doi.org/10.1145/2993600.2993611>.
- 11 Bitcoin Community. Welcome to the Bitcoin Wiki. Available from <https://en.bitcoin.it/wiki/Script>, 2010.
- 12 Ana Bove, Peter Dybjer, and Ulf Norell. A Brief Overview of Agda – A Functional Language with Dependent Types. In *Theorem Proving in Higher Order Logics*, pages 73–7, Berlin, Heidelberg, 2009. Springer Berlin Heidelberg. doi:[https://doi.org/10.1007/978-3-642-03359-9\\_6](https://doi.org/10.1007/978-3-642-03359-9_6).
- 13 Vitalik Buterin. Ethereum: A next-generation smart contract and decentralized application platform, 2014. URL: <https://ethereum.org/en/whitepaper>.
- 14 Manuel M. T. Chakravarty, James Chapman, Kenneth MacKenzie, Orestis Melkonian, Jann Müller, Michael Peyton Jones, Polina Vinogradova, and Philip Wadler. Native Custom Tokens in the Extended UTXO Model. In *Leveraging Applications of Formal Methods, Verification and Validation: Applications*, pages 89–111, Cham, 2020. Springer International Publishing. doi:[https://doi.org/10.1007/978-3-030-61467-6\\_7](https://doi.org/10.1007/978-3-030-61467-6_7).
- 15 Manuel M. T. Chakravarty, James Chapman, Kenneth MacKenzie, Orestis Melkonian, Jann Müller, Michael Peyton Jones, Polina Vinogradova, Philip Wadler, and Joachim Zahnentferner. UTXOma: UTXO with Multi-asset Support. In *Leveraging Applications of Formal Methods, Verification and Validation: Applications*, pages 112–130, Cham, 2020. Springer International Publishing. doi:[https://doi.org/10.1007/978-3-030-61467-6\\_8](https://doi.org/10.1007/978-3-030-61467-6_8).



- 16 Manuel M. T. Chakravarty, James Chapman, Kenneth MacKenzie, Orestis Melkonian, Michael Peyton Jones, and Philip Wadler. The Extended UTXO Model. In *Financial Cryptography and Data Security*, pages 525–539, Cham, 2020. Springer International Publishing.
- 17 James Chapman, Roman Kireev, Chad Nester, and Philip Wadler. System F in Agda, for Fun and Profit. In *Mathematics of Program Construction*, pages 255–297, Cham, 2019. Springer International Publishing. doi:[https://doi.org/10.1007/978-3-030-33636-3\\_10](https://doi.org/10.1007/978-3-030-33636-3_10).
- 18 Martin Churchill, Peter D. Mosses, Neil Sculthorpe, and Paolo Torrini. Reusable components of semantic specifications. In Shigeru Chiba, Éric Tanter, Erik Ernst, and Robert Hirschfeld, editors, *Transactions on Aspect-Oriented Software Development XII*, pages 132–179, Berlin, Heidelberg, 2015. Springer Berlin Heidelberg. doi:10.1007/978-3-662-46734-3\_4.
- 19 Edmund M. Clarke and Jeannette M. Wing. Formal methods: State of the art and future directions. *ACM Comput. Surv.*, 28(4):626–643, dec 1996. doi:<https://doi.org/10.1145/242223.242257>.
- 20 crypto.stackexchange. Is there any famous protocol that were proven secure but whose proof was wrong and lead to real world attacks?, Retrieved 22 April 2022. URL: <https://crypto.stackexchange.com/questions/98829/is-there-any-famous-protocol-that-were-proven-secure-but-whose-proof-was-wrong-a>.
- 21 Peter D and Mosses. Modular structural operational semantics. *Journal of Logic and Algebraic Programming*, 60-61(0):195 – 228, 2004. URL: <http://www.sciencedirect.com/science/article/pii/S156783260400027X>, doi:10.1016/j.jlap.2004.03.008.
- 22 Edsger W. Dijkstra. Guarded commands, nondeterminacy and formal derivation of programs. *Commun. ACM*, 18(8):453–457, aug 1975. doi:<https://doi.org/10.1145/360933.360975>.
- 23 Etherscan. TheDAO smart contract 2016, Retrieved 27 March 2022. Available from <http://etherscan.io/address/0xbb9bc244d798123fde783fcc1c72d3bb8c189413#code>.
- 24 Jean-Christophe Filliâtre and Andrei Paskevich. Why3 — Where Programs Meet Provers. In *Programming Languages and Systems*, pages 125–128, Berlin, Heidelberg, 2013. Springer Berlin Heidelberg. doi:[https://doi.org/10.1007/978-3-642-37036-6\\_8](https://doi.org/10.1007/978-3-642-37036-6_8).
- 25 Yoichi Hirai. Defining the Ethereum Virtual Machine for Interactive Theorem Provers. In *Financial Cryptography and Data Security*, pages 520–535, Cham, 2017. Springer International Publishing. doi:[https://doi.org/10.1007/978-3-319-70278-0\\_33](https://doi.org/10.1007/978-3-319-70278-0_33).
- 26 C. A. R. Hoare. An axiomatic basis for computer programming. *Communications of the ACM*, 12(10):576 – 585, October 1969. doi:10.1145/363235.363259.
- 27 IMDEA Software Institute. HTT: Hoare Type Theory, 10 March 2015. Available from <https://software.imdea.org/~aleks/htt/>.
- 28 Sukrit Kalra, Seep Goel, Mohan Dhawan, and Subodh Sharma. ZEUS: Analyzing Safety of Smart Contracts. In *25th Annual Network and Distributed System Security Symposium, NDSS 2018, San Diego, California, USA, February 18-21, 2018*, pages 1–15. The Internet Society, 2018. doi:<http://dx.doi.org/10.14722/ndss.2018.23082>.
- 29 James C. King. Symbolic execution and program testing. *Commun. ACM*, 19(7):385–394, July 1976. doi:10.1145/360248.360252.
- 30 Rick Klomp and Andrea Bracciali. On Symbolic Verification of Bitcoin’s script Language. In *Data Privacy Management, Cryptocurrencies and Blockchain Technology*, pages 38–56, Cham, 2018. Springer International Publishing. doi:[https://doi.org/10.1007/978-3-030-00305-0\\_3](https://doi.org/10.1007/978-3-030-00305-0_3).
- 31 Loi Luu, Duc-Hiep Chu, Hrishi Olickel, Prateek Saxena, and Aquinas Hobor. Making Smart Contracts Smarter. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security, CCS ’16*, page 254–269, New York, NY, USA, 2016. Association for Computing Machinery. doi:<http://dx.doi.org/10.1145/2976749.2978309>.
- 32 Luiz Eduardo G. Martins and Tony Gorschek. Requirements engineering for safety-critical systems: Overview and challenges. *IEEE Software*, 34(4):49–57, 2017. doi:10.1109/MS.2017.94.

- 33 Anastasia Mavridou, Aron Laszka, Emmanouela Stachtari, and Abhishek Dubey. VeriSolid: Correct-by-Design Smart Contracts for Ethereum. In *Financial Cryptography and Data Security*, pages 446–465, Cham, 2019. Springer International Publishing.
- 34 Orestis Melkonian. Formalizing BitML Calculus in Agda, 2019. Student Research Competition, Poster Session, ICFP’19. URL: <https://omelkonian.github.io/data/publications/formal-bitml.pdf>.
- 35 Orestis Melkonian. Formalizing Extended UTxO and BitML Calculus in Agda. Master’s thesis, Utrecht University, Department of Information and Computing Sciences, July 2019. URL: <https://studenttheses.uu.nl/handle/20.500.12932/32981>.
- 36 Yvonne Murray and David A. Anisi. Survey of Formal Verification Methods for Smart Contracts on Blockchain. In *2019 10th IFIP International Conference on New Technologies, Mobility and Security (NTMS)*, pages 1–6, 2019. doi:<http://dx.doi.org/10.1109/NTMS.2019.8763832>.
- 37 Satoshi Nakamoto. Bitcoin: A peer-to-peer electronic cash system. *Decentralized Business Review*, 2008. URL: <https://www.debr.io/article/21260.pdf>.
- 38 Nanevski, Aleksandar and Vafeiadis, Viktor and Berdine, Josh. Structuring the Verification of Heap-Manipulating Programs. *SIGPLAN Not.*, 45(1):261–274, jan 2010. doi:<https://doi.org/10.1145/1707801.1706331>.
- 39 Ulf Norell. Dependently typed programming in Agda. In *Advanced Functional Programming: 6th International School, AFP 2008, Heijen, The Netherlands, May 2008, Revised Lectures*, pages 230–266, Berlin, Heidelberg, 2009. Springer Berlin Heidelberg. doi:10.1007/978-3-642-04652-0\_5.
- 40 Russell O’Connor. Simplicity: A new language for blockchains. In *Proceedings of the 2017 Workshop on Programming Languages and Analysis for Security, PLAS ’17*, page 107–120, New York, NY, USA, 2017. Association for Computing Machinery. doi:<https://doi.org/10.1145/3139337.3139340>.
- 41 Daejun Park, Yi Zhang, Manasvi Saxena, Philip Daian, and Grigore Roşu. A Formal Verification Tool for Ethereum VM Bytecode. In *Proceedings of the 2018 26th ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering, ESEC/FSE 2018*, page 912–915, New York, NY, USA, 2018. Association for Computing Machinery. doi:<http://dx.doi.org/10.1145/3236024.3264591>.
- 42 Christine Paulin-Mohring. *Introduction to the Coq Proof-Assistant for Practical Software Verification*, pages 45–95. Springer Berlin Heidelberg, Berlin, Heidelberg, 2012. doi:[http://dx.doi.org/10.1007/978-3-642-35746-6\\_3](http://dx.doi.org/10.1007/978-3-642-35746-6_3).
- 43 Elizabeth D. Rather, Donald R. Colburn, and Charles H. Moore. *The Evolution of Forth*, page 625–670. Association for Computing Machinery, New York, NY, USA, 1996. doi:<https://doi.org/10.1145/234286.1057832>.
- 44 Maria Ribeiro, Pedro Adão, and Paulo Mateus. *Formal Verification of Ethereum Smart Contracts Using Isabelle/HOL*, pages 71–97. Springer International Publishing, Cham, 2020. doi:[https://doi.org/10.1007/978-3-030-62077-6\\_7](https://doi.org/10.1007/978-3-030-62077-6_7).
- 45 Anton Setzer. Modelling Bitcoin in Agda. *CoRR*, abs/1804.06398, 2018. URL: <http://arxiv.org/abs/1804.06398>, arXiv:1804.06398.
- 46 Anton Setzer, Fahad Alhabardi, and Bogdan Lazar. Verification Of Smart Contracts With Agda. Available from <https://github.com/fahad1985lab/Smart--Contracts--Verification--With--Agda>, 2021.
- 47 Stack Exchange Inc. provable security - Is there any famous protocol that were proven secure but whose proof was wrong and lead to real world attacks? , Retrieved 22 April 2022. Availabe from <https://crypto.stackexchange.com/questions/98829/is-there-any-famous-protocol-that-were-proven-secure-but-whose-proof-was-wrong-a>.
- 48 Philip Wadler, Wen Kokke, and Jeremy G. Siek. *Programming Language Foundations in Agda*. Online textbook, July 2020. URL: <https://plfa.github.io/Equality/>.
- 49 Maximilian Wohrer and Uwe Zdun. Smart Contracts: Security Patterns in the Ethereum Ecosystem and Solidity. In *2018 International Workshop on Blockchain Oriented Software*

*Engineering (IWBOSE)*, pages 2–8, 2018. doi:<http://dx.doi.org/10.1109/IWBOSE.2018.8327565>.