

---

# **scilla-doc Documentation**

*Release 0.11.0*

**Amrit Kumar**

**Aug 19, 2022**



<b>1</b>	<b>Development Status</b>	<b>3</b>
<b>2</b>	<b>Resources</b>	<b>5</b>
<b>3</b>	<b>Contents</b>	<b>7</b>
3.1	Scilla Design Principles . . . . .	7
3.2	Trying out Scilla . . . . .	8
3.2.1	Savant IDE . . . . .	8
3.2.2	Example Contracts . . . . .	8
3.3	Scilla by Example . . . . .	9
3.3.1	HelloWorld . . . . .	9
3.3.2	A Second Example: Crowdfunding . . . . .	14
3.3.3	A Third Example: A Simple Token Exchange . . . . .	21
3.4	Scilla in Depth . . . . .	42
3.4.1	Structure of a Scilla Contract . . . . .	42
3.4.2	Primitive Data Types & Operations . . . . .	52
3.4.3	Algebraic Datatypes . . . . .	59
3.4.4	More ADT examples . . . . .	68
3.4.5	Standard Libraries . . . . .	72
3.4.6	User-defined Libraries . . . . .	73
3.4.7	Scilla versions . . . . .	74
3.5	The Scilla Standard Library . . . . .	77
3.5.1	BoolUtils . . . . .	77
3.5.2	IntUtils . . . . .	77
3.5.3	ListUtils . . . . .	78
3.5.4	NatUtils . . . . .	82
3.5.5	PairUtils . . . . .	83
3.5.6	Conversions . . . . .	83
3.5.7	Polynetwork Support Library . . . . .	84
3.6	Scilla Tips and Tricks . . . . .	84
3.6.1	Performance . . . . .	84
3.6.2	Money Idioms . . . . .	86
3.6.3	Safety . . . . .	86
3.7	The Scilla checker . . . . .	88
3.7.1	Phases of the Scilla checker . . . . .	88
3.7.2	Cashflow Analysis . . . . .	91
3.8	Interpreter Interface . . . . .	94

3.8.1	Calling Interface . . . . .	94
3.8.2	Initializing the Immutable State . . . . .	95
3.8.3	Input Blockchain State . . . . .	97
3.8.4	Input Message . . . . .	98
3.8.5	Interpreter Output . . . . .	101
3.8.6	Input Mutable Contract State . . . . .	103
3.9	Contact . . . . .	104



*Scilla* (short for *Smart Contract Intermediate-Level Language*) is an intermediate-level smart contract language being developed for the [Zilliqa](#) blockchain. Scilla is designed as a principled language with smart contract safety in mind.

Scilla imposes a structure on smart contracts that will make applications less vulnerable to attacks by eliminating certain known vulnerabilities directly at the language-level. Furthermore, the principled structure of Scilla will make applications inherently more secure and amenable to formal verification.

The language is being developed hand-in-hand with formalization of its semantics and its embedding into the [Coq proof assistant](#) — a state-of-the-art tool for mechanized proofs about properties of programs. Coq is based on advanced dependently-typed theory and features a large set of mathematical libraries. It has been successfully applied previously to implement certified (i.e., fully mechanically verified) compilers, concurrent and distributed applications, including blockchains among others.

*Zilliqa* — the underlying blockchain platform on which Scilla contracts are run — has been designed to be scalable. It employs the idea of sharding to validate transactions in parallel. Zilliqa has an intrinsic token named *Zilling* (ZIL for short) that are required to run smart contracts on Zilliqa.



# CHAPTER 1

---

## Development Status

---

Scilla is under active research and development and hence parts of the specification described in this document are subject to change. Scilla currently comes with an interpreter binary that has been integrated into two Scilla-specific web-based IDEs. *Trying out Scilla* presents the features of the two IDEs.





There are several resources to learn about Scilla and Zilliqa. Some of these are given below:

### **Scilla**

- [Scilla Design Paper](#)
- [Scilla Slides](#)
- [Scilla Language Grammar](#)
- [Scilla Design Story Piece by Piece: Part 1 \(Why do we need a new language?\)](#)

### **Zilliqa**

- [The Zilliqa Design Story Piece by Piece: Part 1 \(Network Sharding\)](#)
- [The Zilliqa Design Story Piece by Piece: Part 2 \(Consensus Protocol\)](#)
- [The Zilliqa Design Story Piece by Piece: Part 3 \(Making Consensus Efficient\)](#)
- [Technical Whitepaper](#)
- [The Not-So-Short Zilliqa Technical FAQ](#)



### 3.1 Scilla Design Principles

*Smart contracts* provide a mechanism to express computations on a blockchain, i.e., a decentralized Byzantine-fault tolerant distributed ledger. With the advent of smart contracts, it has become possible to build what is referred to as *decentralized applications* or Dapps for short. These applications have their program and business logic coded in the form of a smart contract that can be run on a decentralized blockchain network.

Running applications on a decentralized network eliminates the need of a trusted centralized party or a server typical of other applications. These features of smart contracts have become so popular today that they now drive real-world economies through applications such as crowdfunding, games, decentralized exchanges, payment processors among many others.

However, experience over the last few years has shown that implemented operational semantics of smart contract languages admit rather subtle behaviour that diverge from the *intuitive understanding* of the language in the minds of contract developers. This divergence has led to some of the largest attacks on smart contracts, e.g., the attack on the DAO contract and Parity wallet among others. The problem becomes even more severe because smart contracts cannot directly be updated due to the immutable nature of blockchains. It is hence crucial to ensure that smart contracts that get deployed are safe to run.

Formal methods such as verification and model checking have proven to be effective in improving the safety of software systems in other disciplines and hence it is natural to explore their applicability in improving the readability and safety of smart contracts. Moreover, with formal methods, it becomes possible to produce rigorous guarantees about the behavior of a contract.

Applying formal verification tools with existing languages such as Solidity however is not an easy task because of the extreme expressivity typical of a Turing-complete language. Indeed, there is a trade-off between making a language simpler to understand and amenable to formal verification, and making it more expressive. For instance, Bitcoin's scripting language occupies the *simpler* end of the spectrum and does not handle stateful-objects. On the *expressive* side of the spectrum is a Turing-complete language such as Solidity.

*Scilla* is a new (intermediate-level) smart contract language that has been designed to achieve both *expressivity* and *tractability* at the same time, while enabling formal reasoning about contract behavior by adopting certain fundamental design principles as described below:

### Separation Between Computation and Communication

Contracts in Scilla are structured as communicating automata: every in-contract computation (e.g., changing its balance or computing a value of a function) is implemented as a standalone, atomic transition, i.e., without involving any other parties. Whenever such involvement is required (e.g., for transferring control to another party), a transition would end, with an explicit communication, by means of sending and receiving messages. The automata-based structure makes it possible to disentangle the contract-specific effects (i.e., transitions) from blockchain-wide interactions (i.e., sending/receiving funds and messages), thus providing a clean reasoning mechanism about contract composition and invariants.

### Separation Between Effectful and Pure Computations

Any in-contract computation happening within a transition has to terminate, and have a predictable effect on the state of the contract and the execution. In order to achieve this, Scilla draws inspiration from functional programming with effects in distinguishing between pure expressions (e.g., expressions with primitive data types and maps), impure local state manipulations (i.e., reading/writing into contract fields), and blockchain reflection (e.g., reading current block number). By carefully designing semantics of interaction between pure and impure language aspects, Scilla ensures a number of foundational properties about contract transitions, such as progress and type preservation, while also making them amenable to interactive and/or automatic verification with standalone tools.

### Separation Between Invocation and Continuation

Structuring contracts as communicating automata provides a computational model, which only allows *tail-calls*, i.e., every call to an external function (i.e., another contract) has to be done as the absolutely last instruction.

## 3.2 Trying out Scilla

Scilla is under active development. You can try out Scilla in the online IDE.

### 3.2.1 Savant IDE

[Neo Savant IDE](#) is a web-based development environment that allows you to interact with the simulated testnet environment, the live developer testnet, and the live mainnet. It is optimized for use in Chrome Web Browser. Neo Savant IDE allows you to import accounts from external wallets like Ledger or keystore files.

The IDE automatically request the faucets to disburse testnet \$ZIL to you when the wallet is successfully imported. On the simulated testnet environment, you will receive 10,000 \$ZIL. While on the developer testnet, you will receive 300 \$ZIL. There are no faucets for the live mainnet.

The Neo Savant IDE can act as a staging environment, before doing automated script testing with tools like [Isolated Server](#) and [Zilliqa-JS](#). To try out the Neo Savant IDE, users need to visit [Neo Savant IDE](#).

### 3.2.2 Example Contracts

Savant IDE comes with the following sample smart contracts written in Scilla:

- **HelloWorld** : It is a simple contract that allows a specified account denoted `owner` to set a welcome message. Setting the welcome message is done via `setHello (msg: String)`. The contract also provides an interface `getHello()` to allow any account to be returned with the welcome message when called.
- **BookStore** : A demonstration of a CRUD app. Only `owner` of the contract can add members. All members will have read/write access capability to create OR update books in the inventory with `book title`, `author`, and `bookID`.

- **CrowdFunding** : Crowdfunding implements a Kickstarter-style campaign where users can donate funds to the contract using `Donate()`. If the campaign is successful, i.e., enough money is raised within a given time period, the raised money can be sent to a predefined account `owner` via `GetFunds()`. Else, if the campaign fails, then contributors can take back their donations via the transition `ClaimBack()`.
- **Auction** : A simple open auction contract where bidders can make their bid using `Bid()`, and the highest and winning bid amount goes to a predefined account. Bidders who don't win can take back their bid using the transition `Withdraw()`. The organizer of the auction can claim the highest bid by invoking the transition `AuctionEnd()`.
- **FungibleToken** : ZRC-2 Fungible token standard contract for creating fungible digital assets such as stable-coins, utility tokens, and loyalty points.
- **NonFungibleToken** : ZRC-1 Non-fungible token standard contract for creating unique digital assets such as digital collectibles, music records, arts, and domains.
- **ZilGame** : A two-player game where the goal is to find the closest preimage of a given SHA256 digest (puzzle). More formally, given a digest  $d$ , and two values  $x$  and  $y$ ,  $x$  is said to be a closer preimage than  $y$  of  $d$  if  $\text{Distance}(\text{SHA-256}(x), d) < \text{Distance}(\text{SHA-256}(y), d)$ , for some *Distance* function. The game is played in two phases. In the first phase, players submit their hash, i.e.,  $\text{SHA-256}(x)$  and  $\text{SHA-256}(y)$  using the transition `Play(guess: ByStr32)`. Once the first player has submitted her hash, the second player has a bounded time to submit her hash. If the second player does not submit her hash within the stipulated time, then the first player may become the winner. In the second phase, players have to submit the corresponding values  $x$  or  $y$  using the transition `ClaimReward(solution: Int128)`. The player submitting the closest preimage is declared the winner and wins a reward. The contract also provides a transition `Withdraw()` to recover funds and send to a specified `owner` in case no player plays the game.
- **SchnorrTest** : A sample contract to test the generation of a Schnorr public/private key pairs, signing of a `msg` with the private keys, and verification of the signature.
- **ECDSATest** : A sample contract to test the generation of a ECDSA public/private keypairs, signing of a message with the private keys, and verification of the signature.

## 3.3 Scilla by Example

### 3.3.1 HelloWorld

We start off by writing a classical `HelloWorld.scilla` contract with the following specification:

- It should have an *immutable contract parameter* `owner` to be initialized by the creator of the contract. The parameter is immutable in the sense that once initialized during contract deployment, its value cannot be changed. `owner` will be of type `ByStr20` (a hexadecimal Byte String representing a 20 byte address).
- It should have a *mutable field* `welcome_msg` of type `String` initialized to `" "`. Mutability here refers to the possibility of modifying the value of a variable even after the contract has been deployed.
- The `owner` and **only her** should be able to modify `welcome_msg` through an interface `setHello`. The interface takes a `msg` (of type `String`) as input and allows the `owner` to set the value of `welcome_msg` to `msg`.
- It should have an interface `getHello` that welcomes any caller with `welcome_msg`. `getHello` will not take any input.

## Defining a Contract, its Immutable Parameters and Mutable Fields

A contract is declared using the `contract` keyword that starts the scope of the contract. The keyword is followed by the name of the contract which will be `HelloWorld` in our example. So, the following code fragment declares a `HelloWorld` contract.

```
contract HelloWorld
```

---

**Note:** In the current implementation, a Scilla contract can only contain a single contract declaration and hence any code that follows the `contract` keyword is part of the contract declaration. In other words, there is no explicit keyword to declare the end of the contract definition.

---

A contract declaration is followed by the declaration of its immutable parameters, the scope of which is defined by `()`. Each immutable parameter is declared in the following way: `vname : vtype`, where `vname` is the parameter name and `vtype` is the parameter type. Immutable parameters are separated by `,`. As per the specification, the contract will have only one immutable parameter `owner` of type `ByStr20` and hence the following code fragment.

```
(owner: ByStr20)
```

Mutable fields in a contract are declared through keyword `field`. Each mutable field is declared in the following way: `field vname : vtype = init_val`, where `vname` is the field name, `vtype` is its type and `init_val` is the value to which the field has to be initialized. The `HelloWorld` contract has one mutable field `welcome_msg` of type `String` initialized to `""`. This yields the following code fragment:

```
field welcome_msg : String = ""
```

At this stage, our `HelloWorld.scilla` contract will have the following form that includes the contract name, its immutable parameters and mutable fields:

```
contract HelloWorld
(owner: ByStr20)

field welcome_msg : String = ""
```

## Defining Interfaces aka Transitions

Interfaces like `setHello` are referred to as *transitions* in Scilla. Transitions are similar to *functions* or *methods* in other languages. There is an important difference, however, most languages allow their functions or methods to be “interrupted” by a thread running in parallel, but Scilla won’t let a transition to be interrupted ensuring there is no so-called reentrancy issues.

---

**Note:** The term *transition* comes from the underlying computation model in Scilla which follows a communicating automaton. A contract in Scilla is an automaton with some state. The state of an automaton can be changed using a transition that takes a previous state and an input and yields a new state. Check the [wikipedia entry](#) to read more about transition systems.

---

A transition is declared using the keyword `transition`. The end of a transition scope is declared using the keyword `end`. The `transition` keyword is followed by the transition name, which is `setHello` for our example. Then follows the input parameters within `()`. Each input parameter is separated by a `,` and is declared in the following format: `vname : vtype`. According to the specification, `setHello` takes only one parameter of name `msg` of type `String`. This yields the following code fragment:

```
transition setHello (msg : String)
```

What follows the transition signature is the body of the transition. Code for the first transition `setHello (msg : String)` to set `welcome_msg` is given below:

```

1 transition setHello (msg : String)
2   is_owner = builtin eq owner _sender;
3   match is_owner with
4   | False =>
5     e = {_eventname : "setHello"; code : not_owner_code};
6     event e
7   | True =>
8     welcome_msg := msg;
9     e = {_eventname : "setHello"; code : set_hello_code};
10    event e
11  end
12 end
```

At first, the caller of the transition is checked against the `owner` using the instruction `builtin eq owner _sender` in Line 2. In order to compare two addresses, we are using the function `eq` defined as a builtin operator. The operator returns a Boolean value `True` or `False`.

**Note:** Scilla internally defines some variables that have special semantics. These special variables are often prefixed by `_`. For instance, `_sender` in Scilla refers to the account address that called the current contract.

Depending on the output of the comparison, the transition takes a different path declared using *pattern matching*, the syntax of which is given in the fragment below.

```

match expr with
| pattern_1 => expr_1
| pattern_2 => expr_2
end
```

The above code checks whether `expr` evaluates to a value that matches `pattern_1` or `pattern_2`. If `expr` evaluates to a value matching `pattern_1`, then the next expression to be evaluated will be `expr_1`. Otherwise, if `expr` evaluates to a value matching `pattern_2`, then the next expression to be evaluated will be `expr_2`.

Hence, the following code block implements an `if-then-else` instruction:

```

match expr with
| True  => expr_1
| False => expr_2
end
```

### The Caller is Not the Owner

In case the caller is different from `owner`, the transition takes the `False` branch and the contract emits an event using the instruction `event`.

An event is a signal that gets stored on the blockchain for everyone to see. If a user uses a client application to invoke a transition on a contract, the client application can listen for events that the contract may emit, and alert the user.

More concretely, the output event in this case is:

```
e = {_eventname : "setHello"; code : not_owner_code};
```

An event is comprised of a number of `vname : value` pairs delimited by `;` inside a pair of curly braces `{}`. An event must contain the compulsory field `_eventname`, and may contain other fields such as the `code` field in the example above.

---

**Note:** In our example we have chosen to name the event after the transition that emits the event, but any name can be chosen. However, it is recommended that you name the events in a way that makes it easy to see which part of the code emitted the event.

---

## The Caller is the Owner

In case the caller is `owner`, the contract allows the caller to set the value of the mutable field `welcome_msg` to the input parameter `msg`. This is done through the following instruction:

```
welcome_msg := msg;
```

---

**Note:** Writing to a mutable field is done using the operator `:=`.

---

And as in the previous case, the contract then emits an event with the code `set_hello_code`.

## Libraries

A Scilla contract may come with some helper libraries that declare purely functional components of a contract, i.e., components with no state manipulation. A library is declared in the preamble of a contract using the keyword `library` followed by the name of the library. In our current example a library declaration would look as follows:

```
library HelloWorld
```

The library may include utility functions and program constants using the `let ident = expr` construct. In our example the library will only include the definition of error codes:

```
let not_owner_code = UInt32 1
let set_hello_code = UInt32 2
```

At this stage, our contract fragment will have the following form:

```
library HelloWorld

let not_owner_code = UInt32 1
let set_hello_code = UInt32 2

contract HelloWorld
(owner: ByStr20)

field welcome_msg : String = ""

transition setHello (msg : String)
  is_owner = builtin eq owner _sender;
```

(continues on next page)





```

(*****
(*           Associated library           *)
(*****
library HelloWorld

let not_owner_code = Uint32 1
let set_hello_code = Uint32 2

(*****
(*           The contract definition           *)
(*****

contract HelloWorld
(owner: ByStr20)

field welcome_msg : String = ""

transition setHello (msg : String)
  is_owner = builtin eq owner _sender;
  match is_owner with
  | False =>
    e = {_eventname : "setHello"; code : not_owner_code};
    event e
  | True =>
    welcome_msg := msg;
    e = {_eventname : "setHello"; code : set_hello_code};
    event e
  end
end

transition getHello ()
  r <- welcome_msg;
  e = {_eventname: "getHello"; msg: r};
  event e
end

```

### 3.3.2 A Second Example: Crowdfunding

In this section, we present a slightly more involved contract that runs a crowdfunding campaign. In a crowdfunding campaign, a project owner wishes to raise funds through donations from the community.

It is assumed that the owner (`owner`) wishes to run the campaign until a certain, predetermined block number is reached on the blockchain (`max_block`). The owner also wishes to raise a minimum amount of QA (`goal`) without which the project can not be started. The contract hence has three immutable parameters `owner`, `max_block` and `goal`.

The immutable parameters are provided when the contract is deployed. At that point we wish to add a sanity check that the `goal` is a strictly positive amount. If the contract is accidentally initialised with a `goal` of 0, then the contract should not be deployed.

The total amount that has been donated to the campaign so far is stored in a field `_balance`. Any contract in Scilla has an implicit `_balance` field of type `Uint128`, which is initialised to 0 when the contract is deployed, and which holds the amount of QA in the contract's account on the blockchain.

The campaign is deemed successful if the owner can raise the goal in the stipulated time. In case the campaign is

unsuccessful, the donations are returned to the project backers who contributed during the campaign. The backers are supposed to ask for refund explicitly.

The contract maintains two mutable fields:

- `backers`: a field map from a contributor's address (a `ByStr20` value) to the amount contributed, represented with a `Uint128` value. Since there are no backers initially, this map is initialized to an `Emp` (empty) map. The map enables the contract to register a donor, prevent multiple donations and to refund back the money if the campaign does not succeed.
- `funded`: a Boolean flag initialized to `False` that indicates whether the owner has already transferred the funds after the end of the campaign.

The contract contains three transitions: `Donate ()` that allows anyone to contribute to the crowdfunding campaign, `GetFunds ()` that allows **only the owner** to claim the donated amount and transfer it to `owner` and `ClaimBack ()` that allows contributors to claim back their donations in case the campaign is not successful.

### Sanity check for contract parameters

To ensure that the `goal` is a strictly positive amount, we use a *contract constraint*:

```
with
  let zero = Uint128 0 in
    builtin lt zero goal
=>
```

The Boolean expression between `with` and `=>` above is evaluated during contract deployment and the contract only gets deployed if the result of evaluation is `True`. This ensures that the contract cannot be deployed with a `goal` of 0 by mistake.

### Reading the Current Block Number

The deadline is given as a block number, so to check whether the deadline has passed, we must compare the deadline against the current block number.

The current block number is read as follows:

```
blk <- & BLOCKNUMBER;
```

Block numbers have a dedicated type `BNum` in Scilla, so as to not confuse them with regular unsigned integers.

---

**Note:** Reading data from the blockchain is done using the operator `<- &`. Blockchain data cannot be updated directly from the contract.

---

### Reading and Updating the Current Balance

The target for the campaign is specified by the owner in the immutable parameter `goal` when the contract is deployed. To check whether the target have been met, we must compare the total amount raised to the target.

The amount of QA raised is stored in the contract's account on the blockchain, and can be accessed through the implicitly declared `_balance` field as follows:

```
bal <- _balance;
```

Money is represented as values of type `Uint128`.

---

**Note:** The `_balance` field is read using the operator `<-` just like any other contract field. However, the `_balance` field can only be updated by accepting money from incoming messages (using the instruction `accept`), or by explicitly transferring money to other account (using the instruction `send` as explained below).

---

### Sending Messages

In Scilla, there are two ways that transitions can transmit data. One way is through events, as covered in the previous example. The other is through the sending of messages using the instruction `send`.

`send` is used to send messages to other accounts, either in order to invoke transitions on another smart contract, or to transfer money to user accounts. On the other hand, events are dispatched signals that smart contracts can use to transmit data to client applications.

To construct a message we use a similar syntax as when constructing events:

```
msg = {_tag : ""; _recipient : owner; _amount : bal; code : got_funds_code};
```

A message must contain the compulsory *message fields* `_tag`, `_recipient` and `_amount`. The `_recipient` message field is the blockchain address (of type `ByStr20`) that the message is to be sent to, and the `_amount` message field is the number of QA to be transferred to that account.

The value of the `_tag` message field is the name of the transition (of type `String`) that is to be invoked on the contract deployed at `_recipient` address. If `_recipient` is the address of a user account then the value of `_tag` is ignored, hence for simplicity we put `"` here.

---

**Note:** To make it possible to refund both contracts and user accounts (this is useful if a backer used a wallet contract to donate), use a standard transition name as per [ZRC-5](#), i.e. `AddFunds`.

---

In addition to the compulsory fields the message may contain other fields, such as `code` above. However, if the message recipient is a contract, the additional fields must have the same names and types as the parameters of the transition being invoked on the recipient contract.

Sending a message is done using the `send` instruction, which takes a list of messages as a parameter. Since we will only ever send one message at a time in the crowdfunding contract, we define a library function `one_msg` to construct a list consisting of one message:

```
let one_msg =  
  fun (msg : Message) =>  
    let nil_msg = Nil {Message} in  
    Cons {Message} msg nil_msg
```

To send out a message, we first construct the message, insert it into a list, and send it:

```
msg = {_tag : ""; _recipient : owner; _amount : bal; code : got_funds_code};  
msgs = one_msg msg;  
send msgs
```

### Procedures

The transitions of a Scilla contract often need to perform the same small sequence of instructions. In order to prevent code duplication a contract may define a number of *procedures*, which may be invoked from the contract's transitions.

Procedures also help divide the contract code into separate, self-contained pieces which are easier to read and reason about individually.

A procedure is declared using the keyword `procedure`. The end of a procedure is declared using the keyword `end`. The `procedure` keyword is followed by the transition name, then the input parameters within `()`, and then the statements of the procedure.

In our example the `Donate` transition will issue an event in three situations: An error event if the donation happens after the deadline, another error event if the backer has donated money previously, and a non-error event indicating a successful donation. Since much of the event issuing code is identical, we decide to define a procedure `DonationEvent` which is responsible for issuing the correct event:

```
procedure DonationEvent (failure : Bool, error_code : Int32)
  match failure with
  | False =>
    e = {_eventname : "DonationSuccess"; donor : _sender;
        amount : _amount; code : accepted_code};
    event e
  | True =>
    e = {_eventname : "DonationFailure"; donor : _sender;
        amount : _amount; code : error_code};
    event e
  end
end
```

The procedure takes two arguments: A `Bool` indicating whether the donation failed, and an error code indicating the type of failure if a failure did indeed occur.

The procedure performs a `match` on the `failure` argument. If the donation did not fail, the error code is ignored, and a `DonationSuccess` event is issued. Otherwise, if the donation failed, then a `DonationFailure` event is issued with the error code that was passed as the second argument to the procedure.

The following code shows how to invoke the `DonationEvent` procedure with the arguments `True` and `0`:

```
c = True;
err_code = Int32 0;
DonationEvent c err_code;
```

---

**Note:** The special parameters `_sender`, `_origin` and `_amount` are available to a procedure even though the procedure is invoked by a transition rather than by an incoming message. It is not necessary to pass these special parameters as arguments to the procedure.

---



---

**Note:** Procedures are similar to library functions in that they can be invoked from any transition (as long as the transition is defined after the procedure). However, procedures are different from library functions in that library functions cannot access the contract state, and procedures cannot return a value.

Procedures are similar to transitions in that they can access and change the contract state, as well as read the incoming messages and send outgoing messages. However, procedures cannot be invoked from the blockchain layer. Only transitions may be invoked from outside the contract, so procedures can be viewed as private transitions.

---

### Putting it All Together

The complete crowdfunding contract is given below.

```

(*****
(*                               *)
(Scilla version)
(*****

scilla_version 0

(*****
(*                               *)
(Associated library)
(*****
import BoolUtils

library Crowdfunding

let one_msg =
  fun (msg : Message) =>
    let nil_msg = Nil {Message} in
    Cons {Message} msg nil_msg

let blk_leq =
  fun (blk1 : BNum) =>
  fun (blk2 : BNum) =>
    let bc1 = builtin blt blk1 blk2 in
    let bc2 = builtin eq blk1 blk2 in
    orb bc1 bc2

let get_funds_allowed =
  fun (cur_block : BNum) =>
  fun (max_block : BNum) =>
  fun (balance : Uint128) =>
  fun (goal : Uint128) =>
    let in_time = blk_leq cur_block max_block in
    let deadline_passed = negb in_time in
    let target_not_reached = builtin lt balance goal in
    let target_reached = negb target_not_reached in
    andb deadline_passed target_reached

let claimback_allowed =
  fun (balance : Uint128) =>
  fun (goal : Uint128) =>
  fun (already_funded : Bool) =>
    let target_not_reached = builtin lt balance goal in
    let not_already_funded = negb already_funded in
    andb target_not_reached not_already_funded

let accepted_code = Int32 1
let missed_deadline_code = Int32 2
let already_backed_code = Int32 3
let not_owner_code = Int32 4
let too_early_code = Int32 5
let got_funds_code = Int32 6
let cannot_get_funds = Int32 7
let cannot_reclaim_code = Int32 8
let reclaimed_code = Int32 9

(*****
(*                               *)
(The contract definition)
(*****

```

(continues on next page)

(continued from previous page)

```

contract Crowdfunding

(* Parameters *)
(owner      : ByStr20,
max_block  : BNum,
goal       : Uint128)

(* Contract constraint *)
with
  let zero = Uint128 0 in
    builtin lt zero goal
=>

(* Mutable fields *)
field backers : Map ByStr20 Uint128 = Emp ByStr20 Uint128
field funded  : Bool = False

procedure DonationEvent (failure : Bool, error_code : Int32)
  match failure with
  | False =>
    e = {_eventname : "DonationSuccess"; donor : _sender;
        amount : _amount; code : accepted_code};
    event e
  | True =>
    e = {_eventname : "DonationFailure"; donor : _sender;
        amount : _amount; code : error_code};
    event e
  end
end

procedure PerformDonate ()
  c <- exists backers[_sender];
  match c with
  | False =>
    accept;
    backers[_sender] := _amount;
    DonationEvent c accepted_code
  | True =>
    DonationEvent c already_backed_code
  end
end

transition Donate ()
  blk <- & BLOCKNUMBER;
  in_time = blk_leq blk max_block;
  match in_time with
  | True =>
    PerformDonate
  | False =>
    t = True;
    DonationEvent t missed_deadline_code
  end
end

procedure GetFundsFailure (error_code : Int32)
  e = {_eventname : "GetFundsFailure"; caller : _sender;
      amount : _amount; code : error_code};

```

(continues on next page)

(continued from previous page)

```

    event e
end

procedure PerformGetFunds ()
    bal <- _balance;
    tt = True;
    funded := tt;
    msg = {_tag : ""; _recipient : owner; _amount : bal; code : got_funds_code};
    msgs = one_msg msg;
    send msgs
end

transition GetFunds ()
    is_owner = builtin eq owner _sender;
    match is_owner with
    | False =>
        GetFundsFailure not_owner_code
    | True =>
        blk <- & BLOCKNUMBER;
        bal <- _balance;
        allowed = get_funds_allowed blk max_block bal goal;
        match allowed with
        | False =>
            GetFundsFailure cannot_get_funds
        | True =>
            PerformGetFunds
        end
    end
end

procedure ClaimBackFailure (error_code : Int32)
    e = {_eventname : "ClaimBackFailure"; caller : _sender;
        amount : _amount; code : error_code};
    event e
end

procedure PerformClaimBack (amount : Uint128)
    delete backers[_sender];
    msg = {_tag : ""; _recipient : _sender; _amount : amount; code : reclaimed_code};
    msgs = one_msg msg;
    e = { _eventname : "ClaimBackSuccess"; caller : _sender; amount : amount; code : _
    ↪reclaimed_code};
    event e;
    send msgs
end

transition ClaimBack ()
    blk <- & BLOCKNUMBER;
    after_deadline = builtin blt max_block blk;
    match after_deadline with
    | False =>
        ClaimBackFailure too_early_code
    | True =>
        bal <- _balance;
        f <- funded;
        allowed = claimback_allowed bal goal f;
        match allowed with

```

(continues on next page)



(continued from previous page)

```

| False =>
  ClaimBackFailure cannot_reclaim_code
| True =>
  res <- backers[_sender];
  match res with
  | None =>
    (* Sender has not donated *)
    ClaimBackFailure cannot_reclaim_code
  | Some v =>
    PerformClaimBack v
  end
end
end
end
end

```

### 3.3.3 A Third Example: A Simple Token Exchange

As a third example we look at how contracts written in Scilla can interact by passing messages to each other, and by reading each other's states. As our example application we choose a simplified token exchange contracts in which users can place offers of swapping one type of fungible tokens for another type.

#### Fungible Tokens

Recall that a fungible token is one which is indistinguishable from another token of the same type. For example, a US \$1 bank note is indistinguishable from any other US \$1 bank note (for the purposes of using the bank note to pay for goods, services, or other tokens, at least).

The [Zilliqa Reference Contracts](#) library offers specifications and reference implementations of commonly used contract types, and the [ZRC2](#) standard specifies a standard for fungible tokens, which we will use for this example. We will not go into detail about how the token contract works, but only point out a few important aspects that will be needed in order to implement the token exchange.

#### Exchange Specification

We want our simple exchange to support the following functionality:

- The exchange has a number of listed tokens that can be freely swapped with each other. Each listed token is identified by its token code (e.g., “USD” for US dollars).
- The exchange should have an administrator at all times. The administrator is in charge of approving token contracts, and listing them on the exchange. The administrator may pass the administrator role on to someone else.
- Any user can place an order on the exchange. To place an order, the user specifies which token he wants to sell and how many of them he is offering, and which token he wants to buy and how many he wants in return. The contract keeps track of every active (unmatched) order.
- When a user attempts to place an order to sell some tokens, the exchange checks that the user actually has those tokens to sell. If he does, then the exchange claims those tokens and holds on to them until the order is matched.
- Any user can match an active order on the exchange. To match an order, the user specifies which order to match.
- When a user attempts to match an order, the exchange checks that the user actually has the tokens that the order placer wants to buy. If he does, then the exchange transfers the tokens that were claimed when the order was

placed to the order matcher, and transfers the tokens that the order placer wants to buy from the order matcher to the order placer. After the tokens have been transferred the exchange deletes the fulfilled order.

To keep the example brief our exchange will not support unlisting of tokens, cancellation of orders, orders with expiry time, prioritising orders so that the order matcher gets the best deal possible, partial matching of orders, securing the exchange against abuse, fees for trading on the exchange, etc.. We encourage the reader to implement additional features as a way to familiarise themselves even further with Scilla.

## The Administrator Role

The exchange must have an administrator at all times, including when it is first deployed. The administrator may change over time, so we define a mutable field `admin` to keep track of the current administrator, and initialise it to an `initial_admin`, which is given as an immutable parameter:

```
contract SimpleExchange
(
  initial_admin : ByStr20 with end
)

field admin : ByStr20 with end = initial_admin
```

The type of the `admin` field is `ByStr20 with end`, which is an *address type*. As in the earlier examples `ByStr20` is the type of byte strings of length 20, but we now add the additional requirement that when that byte string is interpreted as an address on the network, the address must be *in use*, and the contents at that address must satisfy whatever is between the `with` and `end` keywords.

In this case there is nothing between `with` and `end`, so we have no additional requirements. However, the address must be in use, either by a user or by another contract - otherwise Scilla will not accept it as having a legal address type. (We will go into more detail about address types when the exchange interacts with the listed token contracts.)

Multiple transitions will need to check that the `_sender` is the current `admin`, so let us define a procedure that checks that that is the case:

```
procedure CheckSenderIsAdmin()
  current_admin <- admin;
  is_admin = builtin eq _sender current_admin;
  match is_admin with
  | True => (* Nothing to do *)
  | False =>
    (* Construct an exception object and throw it *)
    e = { _exception : "SenderIsNotAdmin" };
    throw e
  end
end
```

If the `_sender` is the current administrator, then nothing happens, and whichever transition called this procedure can continue. If the `_sender` is someone else, however, the procedure throws an *exception* causing the current transaction to be aborted.

We want the administrator to be able to pass on the administrator role to someone else, so we define our first transition `SetAdmin` as follows:

```
transition SetAdmin(new_admin : ByStr20 with end)
  (* Only the former admin may appoint a new admin *)
  CheckSenderIsAdmin;
  admin := new_admin
end
```

The transition applies the `CheckSenderIsAdmin` procedure, and if no exception is thrown then the sender is indeed the current administrator, and is thus allowed to pass on the administrator role on to someone else. The new admin must once again be an address that is in use.

### Intermezzo: Transferring Tokens On Behalf Of The Token Owner

Before we continue adding features to our exchange we must first look at how token contracts transfer tokens between users.

The ZRC2 token standard defines a field `balances` which keeps track of how many tokens each user has:

```
field balances: Map ByStr20 UInt128
```

However, this is not particularly useful for our exchange, because the token contract won't allow the exchange to transfer tokens belonging to someone other than the exchange itself.

Instead, the ZRC2 standard defines a field `allowances`, which a user who owns tokens can use to allow another user partial access to the owner's tokens:

```
field allowances: Map ByStr20 (Map ByStr20 UInt128)
```

For instance, if Alice has given Bob an allowance of 100 tokens, then the `allowances` map in token contract will contain the value `allowances[<address of Alice>][<address of Bob>] = 100`. This allows Bob to spend 100 of Alice's tokens as if they were his own. (Alice can of course withdraw the allowance, as long as Bob hasn't yet spent the tokens).

Before a user places an order, the user should provide the exchange with an allowance of the token he wants to sell to cover the order. The user can then place the order, and the exchange can check that the allowance is sufficient. The exchange then transfers the tokens to its own account for holding until the order is matched.

Similarly, before a user matches an order, the user should provide the exchange with an allowance of the token that the order placer wants to buy. The user can then match the order, and the exchange can check that the allowance is sufficient. The exchange then transfers those tokens to the user who placed the order, and transfers to the matching user the tokens that it transferred to itself when the order was placed.

In order to check the current allowance that a user has given to the exchange, we will need to specify the `allowances` field in the token address type. We do this as follows:

```
ByStr20 with contract field allowances : Map ByStr20 (Map ByStr20 UInt128) end
```

As with the `admin` field we require that the address is in use. Additionally, the requirements between `with` and `end` must also be satisfied:

- The keyword `contract` specifies that the address must be in use by a contract, and not by a user.
- The keyword `field` specifies that the contract in question must contain a mutable field with the specified name and of the specified type.

### Listing a New Token

The exchange keeps track of its listed tokens, i.e., which tokens are allowed to be traded on the exchange. We do this by defining a map from the token code (a `String`) to the address of the token.

```
field listed_tokens :
  Map String (ByStr20 with contract
              field allowances : Map ByStr20 (Map ByStr20 UInt128)
              end)
```

(continues on next page)

(continued from previous page)

```

= Emp String (ByStr20 with contract
                field allowances : Map ByStr20 (Map ByStr20 Uint128)
                end)

```

Only the administrator is allowed to list new tokens, so we leverage the `CheckSenderIsAdmin` procedure again here.

Additionally, we only want to list tokens that have a different token code from the previously listed tokens. For this purpose we define a procedure `CheckIsTokenUnlisted` to check whether a token code is defined as a key in the `listed_tokens` map. :

```

library SimpleExchangeLib

let false = False

...

contract SimpleExchange (...)

...

procedure ThrowListingStatusException(
  token_code : String,
  expected_status : Bool,
  actual_status : Bool)
  e = { _exception : "UnexpectedListingStatus";
        token_code: token_code;
        expected   : expected_status;
        actual     : actual_status };
  throw e
end

procedure CheckIsTokenUnlisted(
  token_code : String
)
  (* Is the token code listed? *)
  token_code_is_listed <- exists listed_tokens[token_code];
  match token_code_is_listed with
  | True =>
    (* Incorrect listing status *)
    ThrowListingStatusException token_code false token_code_is_listed
  | False => (* Nothing to do *)
  end
end

```

This time we define a helper procedure `ThrowListingStatusException` which unconditionally throws an exception. This will be useful later when we later write the transition for placing orders, because we will need to check that the tokens involved in the order are listed.

We also define the constant `false` in the contract's library. This is due to the fact that Scilla requires all values to be named before they are used in computations. Defining constants in library code prevents us from cluttering the transition code with constant definitions:

```

(* Incorrect listing status *)
false = False; (* We don't want to do it like this *)
ThrowListingStatusException token_code false token_code_is_listed

```

With the helper procedures in place we are now ready to define the `ListToken` transition as follows:

```
transition ListToken(
  token_code : String,
  new_token : ByStr20 with contract field allowances : Map ByStr20 (Map ByStr20,
↳ Uint128) end
)
(* Only the admin may list new tokens. *)
CheckSenderIsAdmin;
(* Only new token codes are allowed. *)
CheckIsTokenUnlisted token_code;
(* Everything is ok. The token can be listed *)
listed_tokens[token_code] := new_token
end
```

### Placing an Order

To place an order a user must specify the token code and the amount of the token he wants to sell, and the token code and amount he wants to buy. We invoke the `ThrowListingStatusException` procedure if any of the token codes are unlisted:

```
transition PlaceOrder(
  token_code_sell : String,
  sell_amount : Uint128,
  token_code_buy: String,
  buy_amount : Uint128
)
(* Check that the tokens are listed *)
token_sell_opt <- listed_tokens[token_code_sell];
token_buy_opt <- listed_tokens[token_code_buy];
match token_sell_opt with
| Some token_sell =>
  match token_buy_opt with
  | Some token_buy =>
    ...
  | None =>
    (* Unlisted token *)
    ThrowListingStatusException token_code_buy true false
  end
| None =>
  (* Unlisted token *)
  ThrowListingStatusException token_code_sell true false
end
end
```

If both tokens are listed, we must first check that the user has supplied a sufficient allowance to the exchange. We will need a similar check when another user matches the order, so we define a helper procedure `CheckAllowance` to perform the check:

```
procedure CheckAllowance(
  token : ByStr20 with contract field allowances : Map ByStr20 (Map ByStr20 Uint128),
↳ end,
  expected : Uint128
)
...
end
```

To perform the check we will need to perform a *remote read* of the `allowances` field in the token contract. We are interested in the allowance given by the `_sender` to the exchange, whose address is given by a special immutable field `_this_address`, so we want to remote read the value of `allowances[_sender][_this_address]` in the token contract.

Remote reads in Scilla are performed using the operator `<- &`, and we use `.` notation to specify the contract that we want to remote read from. The entire statement for the remote read is therefore as follows:

```
actual_opt <-& token.allowances[_sender][_this_address];
```

Just as when we perform a local read of a map, the result of reading from a remote map is an optional value. If the result is `Some v` for some `v`, then the user has provided the exchange with an allowance of `v` tokens, and if the result is `None` the user has not supplied an allowance at all. We therefore need to pattern-match the result to get the actual allowance:

```
(* Find actual allowance. Use 0 if None is given *)
actual = match actual_opt with
  | Some x => x
  | None => zero
end;
```

Once again, we define the constant `zero = Uint128 0` in the contract library for convenience.

We can now compare the actual allowance to the allowance we are expecting, and throw an exception if the actual allowance is insufficient:

```
is_sufficient = uint128_le expected actual;
match is_sufficient with
| True => (* Nothing to do *)
| False =>
  ThrowInsufficientAllowanceException token expected actual
end
```

The function `uint128_le` is a utility function which performs a less-than-or-equal comparison on values of type `Uint128`. The function is defined in the `IntUtils` part of the standard library, so in order to use the function we must import `IntUtils` into the contract, which is done immediately after the `scilla_version` preamble, and before the contract library definitions:

```
scilla_version 0

import IntUtils

library SimpleExchangeLib
...
```

We also utilise a helper procedure `ThrowInsufficientAllowanceException` to throw an exception if the allowance is insufficient, so the `CheckAllowance` procedure ends up looking as follows:

```
procedure ThrowInsufficientAllowanceException (
  token : ByStr20,
  expected : Uint128,
  actual : Uint128)
  e = { _exception : "InsufficientAllowance";
        token: token;
        expected : expected;
        actual : actual };
  throw e
```

(continues on next page)

(continued from previous page)

```

end

procedure CheckAllowance(
  token : ByStr20 with contract field allowances : Map ByStr20 (Map ByStr20 Uint128)
  ↪end,
  expected : Uint128
)
  actual_opt <-& token.allowances[_sender][_this_address];
  (* Find actual allowance. Use 0 if None is given *)
  actual = match actual_opt with
    | Some x => x
    | None => zero
  end;
  is_sufficient = uint128_le expected actual;
  match is_sufficient with
  | True => (* Nothing to do *)
  | False =>
    ThrowInsufficientAllowanceException token expected actual
  end
end

transition PlaceOrder(
  token_code_sell : String,
  sell_amount : Uint128,
  token_code_buy: String,
  buy_amount : Uint128
)
  (* Check that the tokens are listed *)
  token_sell_opt <- listed_tokens[token_code_sell];
  token_buy_opt <- listed_tokens[token_code_buy];
  match token_sell_opt with
  | Some token_sell =>
    match token_buy_opt with
    | Some token_buy =>
      (* Check that the placer has allowed sufficient funds to be accessed *)
      CheckAllowance token_sell sell_amount;
      ...
    | None =>
      (* Unlisted token *)
      ThrowListingStatusException token_code_buy true false
    end
  | None =>
    (* Unlisted token *)
    ThrowListingStatusException token_code_sell true false
  end
end
end

```

If the user has given the exchange a sufficient allowance, the exchange can send a message to the token contract to perform the transfer of tokens from the allowance the exchange's own balance. The transition we need to invoke on the token contract is called `TransferFrom`, as opposed to `Transfer` which transfers funds from the sender's own token balance rather than from the sender's allowance of someone else's balance.

Since the message will look much like the messages that we need when an order is matched, we generate the message using helper functions in the contract library (we will also need a new constant `true`):

```
library SimpleExchangeLib
```

(continues on next page)

(continued from previous page)

```

let true = True

...

let one_msg : Message -> List Message =
  fun (msg : Message) =>
    let mty = Nil { Message } in
    Cons { Message } msg mty

let mk_transfer_msg : Bool -> ByStr20 -> ByStr20 -> ByStr20 -> Uint128 -> Message =
  fun (transfer_from : Bool) =>
  fun (token_address : ByStr20) =>
  fun (from : ByStr20) =>
  fun (to : ByStr20) =>
  fun (amount : Uint128) =>
    let tag = match transfer_from with
      | True => "TransferFrom"
      | False => "Transfer"
    end
  in
  { _recipient : token_address;
    _tag : tag;
    _amount : Uint128 0; (* No Zil are transferred, only custom tokens *)
    from : from;
    to : to;
    amount : amount }

let mk_place_order_msg : ByStr20 -> ByStr20 -> ByStr20 -> Uint128 -> List Message =
  fun (token_address : ByStr20) =>
  fun (from : ByStr20) =>
  fun (to : ByStr20) =>
  fun (amount : Uint128) =>
    (* Construct a TransferFrom message to transfer from seller's allowance to_
    ↪exchange *)
    let msg = mk_transfer_msg true token_address from to amount in
      (* Create a singleton list *)
      one_msg msg

contract SimpleExchange (...)

...

transition PlaceOrder(
  token_code_sell : String,
  sell_amount : Uint128,
  token_code_buy: String,
  buy_amount : Uint128
)
(* Check that the tokens are listed *)
token_sell_opt <- listed_tokens[token_code_sell];
token_buy_opt <- listed_tokens[token_code_buy];
match token_sell_opt with
| Some token_sell =>
  match token_buy_opt with
  | Some token_buy =>
    (* Check that the placer has allowed sufficient funds to be accessed *)
    CheckAllowance token_sell sell_amount;

```

(continues on next page)



(continued from previous page)

```

    (* Transfer the sell tokens to the exchange for holding. Construct a
↳TransferFrom message to the token contract. *)
    msg = mk_place_order_msg token_sell _sender _this_address sell_amount;
    (* Send message when the transition completes. *)
    send msg;
    ...
  | None =>
    (* Unlisted token *)
    ThrowListingStatusException token_code_buy true false
  end
  | None =>
    (* Unlisted token *)
    ThrowListingStatusException token_code_sell true false
  end
end
end

```

Finally, we need to store the new order, so that users may match the order in the future. For this we define a new type `Order`, which holds all the information needed when eventually the order is matched:

```

(* Order placer, sell token, sell amount, buy token, buy amount *)
type Order =
| Order of ByStr20
    (ByStr20 with contract field allowances : Map ByStr20 (Map ByStr20
↳Uint128) end)
    Uint128
    (ByStr20 with contract field allowances : Map ByStr20 (Map ByStr20
↳Uint128) end)
    Uint128

```

A value of type `Order` is given by the type constructor `Order`, a token address and an amount of tokens to sell, and a token address and an amount of tokens to buy.

We now need a field containing a map from order numbers (of type `Uint128`) to `Order`, which represents the currently active orders. Additionally, we will need a way to generate a unique order number, so we'll define a field which holds the next order number to use:

```

field active_orders : Map Uint128 Order = Emp Uint128 Order
field next_order_no : Uint128 = zero

```

To add a new order we need to generate a new order number, store the generated order number and the new order in the `active_orders` map, and finally increment the `next_order_no` field (using the library constant `one = Uint128 1`) so that it is ready for the next order to be placed. We will put that in a helper procedure `AddOrder`, and add a call to the procedure in the `PlaceOrder` transition:

```

procedure AddOrder(
  order : Order
)
  (* Get the next order number *)
  order_no <- next_order_no;
  (* Add the order *)
  active_orders[order_no] := order;
  (* Update the next_order_no field *)
  new_order_no = builtin add order_no one;
  next_order_no := new_order_no
end

```

(continues on next page)

(continued from previous page)

```

transition PlaceOrder(
  token_code_sell : String,
  sell_amount : Uint128,
  token_code_buy: String,
  buy_amount : Uint128
)
  (* Check that the tokens are listed *)
  token_sell_opt <- listed_tokens[token_code_sell];
  token_buy_opt <- listed_tokens[token_code_buy];
  match token_sell_opt with
  | Some token_sell =>
    match token_buy_opt with
    | Some token_buy =>
      (* Check that the placer has allowed sufficient funds to be accessed *)
      CheckAllowance token_sell sell_amount;
      (* Transfer the sell tokens to the exchange for holding. Construct a
      ↪TransferFrom message to the token contract. *)
      msg = mk_place_order_msg token_sell _sender _this_address sell_amount;
      (* Send message when the transition completes. *)
      send msg;
      (* Create order and add to list of active orders *)
      order = Order _sender token_sell sell_amount token_buy buy_amount;
      AddOrder order
    | None =>
      (* Unlisted token *)
      ThrowListingStatusException token_code_buy true false
    end
  | None =>
      (* Unlisted token *)
      ThrowListingStatusException token_code_sell true false
    end
  end
end

```

PlaceOrder is now complete, but there is still one thing missing. The ZRC2 token standard specifies that when a TransferFrom transition is executed, the token sends messages to the recipient and the `_sender` (known as the *initiator*) notifying them of the successful transfer. These notifications are known as *callbacks*. Since our exchange executes a TransferFrom transition on the sell token, and since the exchange is the recipient of those tokens, we will need to specify transitions that can handle both callbacks - if we don't, then the callbacks will not be recognised, causing the entire PlaceOrder transaction to fail.

Token contracts notify the recipients of token transfers because such notifications add an extra safeguard against the risk of transferring tokens to a contract that is unable to deal with token ownership. For instance, if someone were to transfer tokens to the HelloWorld contract in the first example in this section, then the tokens would be locked forever because the HelloWorld contract is incapable of doing anything with the tokens.

Our exchange is only capable of dealing with tokens for which there is an active order, but in principle there is nothing stopping a user from transferring funds to the exchange without placing an order, so we need to ensure that the exchange is only involved in token transfers that it itself has initiated. We therefore define a procedure CheckInitiator, which throws an exception if the exchange is involved in a token transfer that it itself did not initiate, and invoke that procedure from all callback transitions:

```

procedure CheckInitiator(
  initiator : ByStr20)
  initiator_is_this = builtin eq initiator _this_address;
  match initiator_is_this with

```

(continues on next page)

(continued from previous page)

```

| True => (* Do nothing *)
| False =>
  e = { _exception : "UnexpectedTransfer";
        token_address : _sender;
        initiator : initiator };
  throw e
  end
end

transition RecipientAcceptTransferFrom (
  initiator : ByStr20,
  sender : ByStr20,
  recipient : ByStr20,
  amount : Uint128)
  CheckInitiator initiator
end

transition TransferFromSuccessCallBack (
  initiator : ByStr20,
  sender : ByStr20,
  recipient : ByStr20,
  amount : Uint128)
  CheckInitiator initiator
end

```

## Matching an Order

For the `MatchOrder` transition we can leverage many of the helper functions and procedures defined in the previous section.

The user specifies an order he wishes to match. We then look up the order number in the `active_orders` map, and throw an exception if the order is not found:

```

transition MatchOrder(
  order_id : Uint128)
  order <- active_orders[order_id];
  match order with
  | Some (Order order_placer sell_token sell_amount buy_token buy_amount) =>
    ...
  | None =>
    e = { _exception : "UnknownOrder";
          order_id : order_id };
    throw e
  end
end

```

In order to match the order, the matcher has to provide sufficient allowance of the buy token. This is checked by the `CheckAllowance` procedure we defined earlier, so we simply reuse that procedure here:

```

transition MatchOrder(
  order_id : Uint128)
  order <- active_orders[order_id];
  match order with
  | Some (Order order_placer sell_token sell_amount buy_token buy_amount) =>
    (* Check that the placer has allowed sufficient funds to be accessed *)

```

(continues on next page)

(continued from previous page)

```

    CheckAllowance buy_token buy_amount;
    ...
| None =>
    e = { _exception : "UnknownOrder";
          order_id : order_id };
    throw e
    end
end

```

We now need to generate two transfer messages: One message is a `TransferFrom` message on the buy token, transferring the matcher's allowance to the user who placed the order, and the other message is a `Transfer` message on the sell token, transferring the tokens held by the exchange to the order matcher. Once again, we define helper functions to generate the messages:

```

library SimpleExchangeLib

...

let two_msgs : Message -> Message -> List Message =
  fun (msg1 : Message) =>
  fun (msg2 : Message) =>
    let first = one_msg msg1 in
    Cons { Message } msg2 first

let mk_make_order_msgs : ByStr20 -> Uint128 -> ByStr20 -> Uint128 ->
  ByStr20 -> ByStr20 -> ByStr20 -> List Message =
  fun (token_sell_address : ByStr20) =>
  fun (sell_amount : Uint128) =>
  fun (token_buy_address : ByStr20) =>
  fun (buy_amount : Uint128) =>
  fun (this_address : ByStr20) =>
  fun (order_placer : ByStr20) =>
  fun (order_maker : ByStr20) =>
    (* Construct a Transfer message to transfer from exchange to maker *)
    let sell_msg = mk_transfer_msg false token_sell_address this_address order_maker_
    ↪sell_amount in
    (* Construct a TransferFrom message to transfer from maker to placer *)
    let buy_msg = mk_transfer_msg true token_buy_address order_maker order_placer_
    ↪buy_amount in
    (* Create a singleton list *)
    two_msgs sell_msg buy_msg

...

contract SimpleExchange (...)

...

transition MatchOrder(
  order_id : Uint128)
  order <- active_orders[order_id];
  match order with
| Some (Order order_placer sell_token sell_amount buy_token buy_amount) =>
  (* Check that the placer has allowed sufficient funds to be accessed *)
  CheckAllowance buy_token buy_amount;
  (* Create the two transfer messages and send them *)

```

(continues on next page)

(continued from previous page)

```

    msgs = mk_make_order_msgs sell_token sell_amount buy_token buy_amount _this_
↪address order_placer _sender;
    send msgs;
    ...
| None =>
    e = { _exception : "UnkownOrder";
         order_id : order_id };
    throw e
end
end

```

Since the order has now been matched, it should no longer be listed as an active order, so we delete the entry in `active_orders`:

```

transition MatchOrder(
    order_id : Uint128)
    order <- active_orders[order_id];
match order with
| Some (Order order_placer sell_token sell_amount buy_token buy_amount) =>
    (* Check that the placer has allowed sufficient funds to be accessed *)
    CheckAllowance buy_token buy_amount;
    (* Create the two transfer messages and send them *)
    msgs = mk_make_order_msgs sell_token sell_amount buy_token buy_amount _this_
↪address order_placer _sender;
    send msgs;
    (* Order has now been matched, so remove it *)
    delete active_orders[order_id]
| None =>
    e = { _exception : "UnkownOrder";
         order_id : order_id };
    throw e
end
end

```

This concludes the `MatchOrder` transition, but we need to define one additional callback transition. When placing an order we executed a `TransferFrom` transition, but now we also execute a `Transfer` transition, which gives rise to a different callback:

```

transition TransferSuccessCallback (
    initiator : ByStr20,
    sender : ByStr20,
    recipient : ByStr20,
    amount : Uint128)
    (* The exchange only accepts transfers that it itself has initiated. *)
    CheckInitiator initiator
end

```

Note that we do not need to specify a transition handling the receipt of tokens from a `Transfer` transition, because the exchange never executes a `Transfer` with itself as the recipient. By not defining the callback transition at all, we also take care of the situation where a user performs a `Transfer` with the exchange as the recipient, because the recipient callback won't have a matching transition on the exchange, causing the entire transfer transaction to fail.

### Adding callbacks

Similar to how tokens execute callbacks whenever a transfer is performed, we will want to issue callbacks to the users and contracts that trade on our exchange. Callbacks should be issued in the following cases:

- When an order is placed: The address placing the order should receive a callback.
- When an order is matched: The address that placed the order and the address matching the order should each receive a callback.

We choose to only issue callbacks whenever the addresses involved are contracts. This is not strictly necessary, since a callback sent to a user address does not cause a failure, but we do this to illustrate the use of address type casts.

We add the following helper functions to the contract library. These simply construct the callback messages to the relevant addresses in the cases mentioned above:

```
let mk_place_order_callback_msg : ByStr20 -> ByStr20 -> Uint128 -> ByStr20 -> Uint128,
↳-> List Message =
  fun (order_placer : ByStr20) =>
  fun (token_sell : ByStr20) =>
  fun (sell_amount : Uint128) =>
  fun (token_buy : ByStr20) =>
  fun (buy_amount : Uint128) =>
    let msg = { _recipient : order_placer;
                _tag : "PlaceOrderSuccessful";
                _amount : Uint128 0;
                selling_token : token_sell;
                selling_amount : sell_amount;
                buying_token : token_buy;
                buying_amount : buy_amount }

    in
    one_msg msg

let mk_order_matched_callback_msg =
  fun (order_matcher : ByStr20) =>
  fun (token_sell : ByStr20) =>
  fun (sell_amount : Uint128) =>
  fun (token_buy : ByStr20) =>
  fun (buy_amount : Uint128) =>
    let msg = { _recipient : order_matcher;
                _tag : "MatchOrderSuccessful";
                _amount : Uint128 0;
                selling_token : token_sell;
                selling_amount : sell_amount;
                buying_token : token_buy;
                buying_amount : buy_amount }

    in
    one_msg msg

let mk_placed_order_matched_callback_msg =
  fun (order_placer : ByStr20) =>
  fun (token_sell : ByStr20) =>
  fun (sell_amount : Uint128) =>
  fun (token_buy : ByStr20) =>
  fun (buy_amount : Uint128) =>
    let msg = { _recipient : order_placer;
                _tag : "PlacedOrderMatched";
                _amount : Uint128 0;
                selling_token : token_sell;
                selling_amount : sell_amount;
                buying_token : token_buy;
                buying_amount : buy_amount }

    in
    one_msg msg
```

When an order is placed successfully we want to check if the placer is a contract. The way to check this is by using an *address type cast*, which is done as follows:

```
sender_as_contract_opt <-& _sender as ByStr20 with contract end;
```

If `_sender` satisfies the address type `ByStr20 with contract end`, then `sender_as_contract_opt` will be bound to the value `Some v`, where `v` is a value that is equal to `_sender`, but which has the type `ByStr20 with contract end`. If `_sender` does not satisfy the type, then `sender_as_contract_opt` will be bound to `None`.

Performing this type cast allows us to issue a callback in `PlaceOrder` when `_sender` is a contract:

```
match sender_as_contract_opt with
| Some sender_as_contract =>
  callback_msg = mk_place_order_callback_msg _sender token_sell sell_amount token_buy_
  ↪buy_amount;
  send callback_msg
| None => (* Do nothing *)
end
```

We add a similar check to `MatchOrder`, except that here we might need to issue callbacks both to the `_sender` and the order placer:

```
match sender_as_contract_opt with
| Some sender_as_contract =>
  callback_msg = mk_order_matched_callback_msg _sender sell_token sell_amount buy_
  ↪token buy_amount;
  send callback_msg
| None => (* Do nothing *)
end;
placer_as_contract_opt <-& order_placer as ByStr20 with contract end;
match placer_as_contract_opt with
| Some placer_as_contract =>
  callback_msg = mk_placed_order_matched_callback_msg order_placer sell_token sell_
  ↪amount buy_token buy_amount;
  send callback_msg
| None => (* Do nothing *)
end
```

## Putting it All Together

We now have everything in place to specify the entire contract:

```
scilla_version 0

import IntUtils

library SimpleExchangeLib

(* Order placer, sell token, sell amount, buy token, buy amount *)
type Order =
| Order of ByStr20
  (ByStr20 with contract field allowances : Map ByStr20 (Map ByStr20_
  ↪Uint128) end)
  Uint128
  (ByStr20 with contract field allowances : Map ByStr20 (Map ByStr20_
  ↪Uint128) end)
```

(continues on next page)

```

    Uint128

(* Helper values and functions *)
let true = True
let false = False

let zero = Uint128 0
let one = Uint128 1

let one_msg : Message -> List Message =
  fun (msg : Message) =>
    let mty = Nil { Message } in
    Cons { Message } msg mty

let two_msgs : Message -> Message -> List Message =
  fun (msg1 : Message) =>
  fun (msg2 : Message) =>
    let first = one_msg msg1 in
    Cons { Message } msg2 first

let mk_transfer_msg : Bool -> ByStr20 -> ByStr20 -> ByStr20 -> Uint128 -> Message =
  fun (transfer_from : Bool) =>
  fun (token_address : ByStr20) =>
  fun (from : ByStr20) =>
  fun (to : ByStr20) =>
  fun (amount : Uint128) =>
    let tag = match transfer_from with
      | True => "TransferFrom"
      | False => "Transfer"
    end
  in
  { _recipient : token_address;
    _tag : tag;
    _amount : Uint128 0; (* No Zil are transferred, only custom tokens *)
    from : from;
    to : to;
    amount : amount }

let mk_place_order_msg : ByStr20 -> ByStr20 -> ByStr20 -> Uint128 -> List Message =
  fun (token_address : ByStr20) =>
  fun (from : ByStr20) =>
  fun (to : ByStr20) =>
  fun (amount : Uint128) =>
    (* Construct a TransferFrom message to transfer from seller's allowance to_
    ↪exchange *)
    let msg = mk_transfer_msg true token_address from to amount in
    (* Create a singleton list *)
    one_msg msg

let mk_make_order_msgs : ByStr20 -> Uint128 -> ByStr20 -> Uint128 ->
  ByStr20 -> ByStr20 -> ByStr20 -> List Message =
  fun (token_sell_address : ByStr20) =>
  fun (sell_amount : Uint128) =>
  fun (token_buy_address : ByStr20) =>
  fun (buy_amount : Uint128) =>
  fun (this_address : ByStr20) =>
  fun (order_placer : ByStr20) =>

```

(continues on next page)



(continued from previous page)

```

fun (order_maker : ByStr20) =>
  (* Construct a Transfer message to transfer from exchange to maker *)
  let sell_msg = mk_transfer_msg false token_sell_address this_address order_maker_
↳sell_amount in
  (* Construct a TransferFrom message to transfer from maker to placer *)
  let buy_msg = mk_transfer_msg true token_buy_address order_maker order_placer buy_
↳amount in
  (* Create a singleton list *)
  two_msgs sell_msg buy_msg

(* Callback messages *)
let mk_place_order_callback_msg : ByStr20 -> ByStr20 -> Uint128 -> ByStr20 -> Uint128,
↳-> List Message =
  fun (order_placer : ByStr20) =>
  fun (token_sell : ByStr20) =>
  fun (sell_amount : Uint128) =>
  fun (token_buy : ByStr20) =>
  fun (buy_amount : Uint128) =>
    let msg = { _recipient : order_placer;
               _tag : "PlaceOrderSuccessful";
               _amount : Uint128 0;
               selling_token : token_sell;
               selling_amount : sell_amount;
               buying_token : token_buy;
               buying_amount : buy_amount }

    in
    one_msg msg

let mk_order_matched_callback_msg =
  fun (order_matcher : ByStr20) =>
  fun (token_sell : ByStr20) =>
  fun (sell_amount : Uint128) =>
  fun (token_buy : ByStr20) =>
  fun (buy_amount : Uint128) =>
    let msg = { _recipient : order_matcher;
               _tag : "MatchOrderSuccessful";
               _amount : Uint128 0;
               selling_token : token_sell;
               selling_amount : sell_amount;
               buying_token : token_buy;
               buying_amount : buy_amount }

    in
    one_msg msg

let mk_placed_order_matched_callback_msg =
  fun (order_placer : ByStr20) =>
  fun (token_sell : ByStr20) =>
  fun (sell_amount : Uint128) =>
  fun (token_buy : ByStr20) =>
  fun (buy_amount : Uint128) =>
    let msg = { _recipient : order_placer;
               _tag : "PlacedOrderMatched";
               _amount : Uint128 0;
               selling_token : token_sell;
               selling_amount : sell_amount;
               buying_token : token_buy;
               buying_amount : buy_amount }

```

(continues on next page)

(continued from previous page)

```

in
  one_msg msg

contract SimpleExchange
(
  initial_admin : ByStr20 with end
)

(* Active admin. *)
field admin : ByStr20 with end = initial_admin

(* Tokens listed on the exchange. *)
field listed_tokens :
  Map String (ByStr20 with contract
    field allowances : Map ByStr20 (Map ByStr20 Uint128)
    end)
  = Emp String (ByStr20 with contract
    field allowances : Map ByStr20 (Map ByStr20 Uint128)
    end)

(* Active orders, identified by the order number *)
field active_orders : Map Uint128 Order = Emp Uint128 Order

(* The order number to use when the next order is placed *)
field next_order_no : Uint128 = zero

procedure ThrowListingStatusException(
  token_code : String,
  expected_status : Bool,
  actual_status : Bool)
  e = { _exception : "UnexpectedListingStatus";
        token_code: token_code;
        expected : expected_status;
        actual : actual_status };
  throw e
end

procedure ThrowInsufficientAllowanceException(
  token : ByStr20,
  expected : Uint128,
  actual : Uint128)
  e = { _exception : "InsufficientAllowance";
        token: token;
        expected : expected;
        actual : actual };
  throw e
end

(* Check that _sender is the active admin. *)
(* If not, throw an error and abort the transaction *)
procedure CheckSenderIsAdmin()
  current_admin <- admin;
  is_admin = builtin eq _sender current_admin;
  match is_admin with
  | True => (* Nothing to do *)
  | False =>
    (* Construct an exception object and throw it *)

```

(continues on next page)

(continued from previous page)

```

    e = { _exception : "SenderIsNotAdmin" };
    throw e
  end
end

(* Change the active admin *)
transition SetAdmin(
  new_admin : ByStr20 with end
)
(* Only the former admin may appoint a new admin *)
CheckSenderIsAdmin;
admin := new_admin
end

(* Check that a given token code is not already listed. If it is, throw an error. *)
procedure CheckIsTokenUnlisted(
  token_code : String
)
(* Is the token code listed? *)
token_code_is_listed <- exists listed_tokens[token_code];
match token_code_is_listed with
| True =>
  (* Incorrect listing status *)
  ThrowListingStatusException token_code false token_code_is_listed
| False => (* Nothing to do *)
end
end

(* List a new token on the exchange. Only the admin may list new tokens. *)
(* If a token code is already in use, raise an error *)
transition ListToken(
  token_code : String,
  new_token : ByStr20 with contract field allowances : Map ByStr20 (Map ByStr20
↳ Uint128) end
)
(* Only the admin may list new tokens. *)
CheckSenderIsAdmin;
(* Only new token codes are allowed. *)
CheckIsTokenUnlisted token_code;
(* Everything is ok. The token can be listed *)
listed_tokens[token_code] := new_token
end

(* Check that the sender has allowed access to sufficient funds *)
procedure CheckAllowance(
  token : ByStr20 with contract field allowances : Map ByStr20 (Map ByStr20 Uint128)
↳ end,
  expected : Uint128
)
actual_opt <-& token.allowances[_sender][_this_address];
(* Find actual allowance. Use 0 if None is given *)
actual = match actual_opt with
  | Some x => x
  | None => zero
end;
is_sufficient = uint128_le expected actual;
match is_sufficient with

```

(continues on next page)

```

| True => (* Nothing to do *)
| False =>
    ThrowInsufficientAllowanceException token expected actual
end
end

procedure AddOrder(
  order : Order
)
(* Get the next order number *)
order_no <- next_order_no;
(* Add the order *)
active_orders[order_no] := order;
(* Update the next_order_no field *)
new_order_no = builtin add order_no one;
next_order_no := new_order_no
end

(* Place an order on the exchange *)
transition PlaceOrder(
  token_code_sell : String,
  sell_amount : Uint128,
  token_code_buy : String,
  buy_amount : Uint128
)
(* Check that the tokens are listed *)
token_sell_opt <- listed_tokens[token_code_sell];
token_buy_opt <- listed_tokens[token_code_buy];
match token_sell_opt with
| Some token_sell =>
  match token_buy_opt with
  | Some token_buy =>
    (* Check that the placer has allowed sufficient funds to be accessed *)
    CheckAllowance token_sell sell_amount;
    (* Transfer the sell tokens to the exchange for holding. Construct a_
    ↪TransferFrom message to the token contract. *)
    msg = mk_place_order_msg token_sell _sender _this_address sell_amount;
    (* Send message when the transition completes. *)
    send msg;
    (* Create order and add to list of active orders *)
    order = Order _sender token_sell sell_amount token_buy buy_amount;
    AddOrder order;
    (* Do a callback if the placer is a contract *)
    sender_as_contract_opt <-& _sender as ByStr20 with contract end;
    match sender_as_contract_opt with
    | Some sender_as_contract =>
      ↪callback_msg = mk_place_order_callback_msg _sender token_sell sell_amount_
      ↪token_buy buy_amount;
      send callback_msg
    | None => (* Do nothing *)
    end
  | None =>
    (* Unlisted token *)
    ThrowListingStatusException token_code_buy true false
    end
  | None =>
    (* Unlisted token *)

```

(continues on next page)

(continued from previous page)

```

    ThrowListingStatusException token_code_sell true false
end
end

transition MatchOrder(
  order_id : Uint128)
  order <- active_orders[order_id];
  match order with
  | Some (Order order_placer sell_token sell_amount buy_token buy_amount) =>
    (* Check that the placer has allowed sufficient funds to be accessed *)
    CheckAllowance buy_token buy_amount;
    (* Create the two transfer messages and send them *)
    msgs = mk_make_order_msgs sell_token sell_amount buy_token buy_amount _this_
    ↪address order_placer _sender;
    send msgs;
    (* Order has now been matched, so remove it *)
    delete active_orders[order_id];
    (* Do callbacks if the matcher or the placer were contracts *)
    sender_as_contract_opt <-& _sender as ByStr20 with contract end;
    match sender_as_contract_opt with
    | Some sender_as_contract =>
      callback_msg = mk_order_matched_callback_msg _sender sell_token sell_amount buy_
    ↪token buy_amount;
      send callback_msg
    | None => (* Do nothing *)
    end;
    placer_as_contract_opt <-& order_placer as ByStr20 with contract end;
    match placer_as_contract_opt with
    | Some placer_as_contract =>
      callback_msg = mk_placed_order_matched_callback_msg order_placer sell_token_
    ↪sell_amount buy_token buy_amount;
      send callback_msg
    | None => (* Do nothing *)
    end
  | None =>
    e = { _exception : "UnknownOrder";
          order_id : order_id };
    throw e
  end
end

procedure CheckInitiator(
  initiator : ByStr20)
  initiator_is_this = builtin eq initiator _this_address;
  match initiator_is_this with
  | True => (* Do nothing *)
  | False =>
    e = { _exception : "UnexpecedTransfer";
          token_address : _sender;
          initiator : initiator };
    throw e
  end
end

transition RecipientAcceptTransferFrom (
  initiator : ByStr20,
  sender : ByStr20,

```

(continues on next page)

(continued from previous page)

```

recipient : ByStr20,
amount : Uint128)
(* The exchange only accepts transfers that it itself has initiated. *)
CheckInitiator initiator
end

transition TransferFromSuccessCallBack (
  initiator : ByStr20,
  sender : ByStr20,
  recipient : ByStr20,
  amount : Uint128)
(* The exchange only accepts transfers that it itself has initiated. *)
CheckInitiator initiator
end

transition TransferSuccessCallBack (
  initiator : ByStr20,
  sender : ByStr20,
  recipient : ByStr20,
  amount : Uint128)
(* The exchange only accepts transfers that it itself has initiated. *)
CheckInitiator initiator
end

```

As mentioned in the introduction we have kept the exchange simplistic in order to keep the focus on Scilla features.

To further familiarise themselves with Scilla we encourage the reader to add additional features such as unlisting of tokens, cancellation of orders, orders with expiry time, prioritising orders so that the order matcher gets the best deal possible, partial matching of orders, securing the exchange against abuse, fees for trading on the exchange, etc..

## 3.4 Scilla in Depth

### 3.4.1 Structure of a Scilla Contract

The general structure of a Scilla contract is given in the code fragment below:

- The contract starts with the declaration of `scilla_version`, which indicates which major Scilla version the contract uses.
- Then follows the declaration of a `library` that contains purely mathematical functions, e.g., a function to compute the Boolean AND of two bits, or a function computing the factorial of a given natural number.
- Then follows the actual contract definition declared using the keyword `contract`.
- Within a contract, there are then four distinct parts:
  1. The first part declares the immutable parameters of the contract.
  2. The second part describes the contract's constraint, which must be valid when the contract is deployed.
  3. The third part declares the mutable fields.
  4. The fourth part contains all `transition` and `procedure` definitions.

```

(* Scilla contract structure *)

(*****

```

(continues on next page)

(continued from previous page)

```

(*                               Scilla version                               *)
(*****)

scilla_version 1

(*****)
(*                               Associated library                               *)
(*****)

library MyContractLib

(* Library code block follows *)

(*****)
(*                               Contract definition                               *)
(*****)

contract MyContract

(* Immutable contract parameters declaration *)

(vname_1 : vtype_1,
 vname_2 : vtype_2)

(* Contract constraint *)
with
  (* Constraint expression *)
=>

(* Mutable fields declaration *)

field vname_1 : vtype_1 = init_val_1
field vname_2 : vtype_2 = init_val_2

(* Transitions and procedures *)

(* Procedure signature *)
procedure firstProcedure (param_1 : type_1, param_2 : type_2)
  (* Procedure body *)

end

(* Transition signature *)
transition firstTransition (param_1 : type_1, param_2 : type_2)
  (* Transition body *)

end

(* Procedure signature *)
procedure secondProcedure (param_1 : type_1, param_2 : type_2)
  (* Procedure body *)

end

```

(continues on next page)

(continued from previous page)

```
transition secondTransition (param_1: type_1)
  (* Transition body *)

end
```

## Immutable Contract Parameters

*Immutable parameters* are the contract's initial parameters whose values are defined when the contract is deployed, and cannot be modified afterwards.

Immutable parameters are declared using the following syntax:

```
(vname_1 : vtype_1,
 vname_2 : vtype_2,
 ... )
```

Each declaration consists of a parameter name (an identifier) and followed by its type, separated by `:`. Multiple parameter declarations are separated by `,`. The initialization values for parameters are to be specified when the contract is deployed.

---

**Note:** In addition to the explicitly declared immutable parameters, a Scilla contract has the following implicitly declared immutable contract parameters

1. `_this_address` of type `ByStr20`, which is initialised to the address of the contract when the contract is deployed.
2. `_creation_block` of type `BNum`, which is initialized to the block number at which the contract is / was deployed.

These parameters can be freely read within the implementation without having to dereference it using `<-` and cannot be modified with `:=`.

---

## Contract Constraints

A *contract constraint* is a requirement placed on the contract's immutable parameters. A contract constraint provides a way of establishing a contract invariant as soon as the contract is deployed, thus preventing the contract being deployed with nonsensical parameters.

A contract constraint is declared using the following syntax:

```
with
  ...
=>
```

The constraint must be an expression of type `Bool`.

The constraint is checked when the contract is deployed. Contract deployment only succeeds if the constraint evaluates to `True`. If it evaluates to `False`, then the deployment fails.

Here is a simple example of using contract constraints to make sure a contract with a limited period of functioning is not deployed *after* that period:



```
contract Mortal(end_of_life : BNum)
with
  builtin blt _creation_block end_of_life
=>
```

The snippet above uses the implicit contract parameter `_creation_block` described in *Immutable Contract Parameters*.

**Note:** Declaring a contract constraint is optional. If no constraint is declared, then the constraint is assumed to simply be `True`.

## Mutable Fields

*Mutable fields* represent the mutable state (mutable variables) of the contract. They are declared after the immutable parameters, with each declaration prefixed with the keyword `field`.

```
field vname_1 : vtype_1 = expr_1
field vname_2 : vtype_2 = expr_2
...
```

Each expression here is an initialiser for the field in question. The definitions complete the initial state of the contract, at the time of creation. As the contract executes a transition, the values of these fields get modified.

**Note:** In addition to the explicitly declared mutable fields, a Scilla contract has an implicitly declared mutable field `_balance` of type `Uint128`, which is initialised to 0 when the contract is deployed. The `_balance` field keeps the amount of funds held by the contract, measured in QA (1 ZIL = 1,000,000,000,000 QA). This field can be freely read within the implementation, but can only be modified by explicitly transferring funds to other accounts (using `send`), or by accepting money from incoming messages (using `accept`).

**Note:** Both mutable fields and immutable parameters must be of a *storable* type:

- Messages, events and the special `Unit` type are not storable. All other primitive types like integers and strings are storable.
- Function types are not storable.
- Complex types involving uninstantiated type variables are not storable.
- Maps and ADT are storable if the types of their subvalues are storable. For maps this means that the key type and the value type must both be storable, and for ADTs this means that the type of every constructor argument must be storable.

## Units

The Zilliqa protocol supports three basic tokens units - ZIL, LI ( $10^{-6}$  ZIL) and QA ( $10^{-12}$  ZIL).

The base unit used in Scilla smart contracts is QA. Hence, when using money variables, it is important to attach the trailing zeroes that are needed to represent it in QAs.

```
(* fee is 1 QA *)  
let fee = Uint128 1  
  
(* fee is 1 LI *)  
let fee = Uint128 1000000  
  
(* fee is 1 ZIL *)  
let fee = Uint128 1000000000000
```

## Transitions

*Transitions* are a way to define how the state of the contract may change. The transitions of a contract define the public interface for the contract, since transitions may be invoked by sending a message to the contract.

Transitions are defined with the keyword `transition` followed by the parameters to be passed. The definition ends with the `end` keyword.

```
transition foo (vname_1 : vtype_1, vname_2 : vtype_2, ...)  
  ...  
end
```

where `vname` : `vtype` specifies the name and type of each parameter and multiple parameters are separated by `,`

---

**Note:** In addition to the parameters that are explicitly declared in the definition, each transition has the following implicit parameters:

- `_amount` : `Uint128` : Incoming amount, in QA (see section above on the units), sent by the sender. To transfer the money from the sender to the contract, the transition must explicitly accept the money using the `accept` instruction. The money transfer does not happen if the transition does not execute an `accept`.
- `_sender` : `ByStr20 with end` : The account address that triggered this transition. If the transition was called by a contract account instead of a user account, then `_sender` is the address of the contract that called this transition. In a chain call, this is the contract that sent the message invoking the current transition.
- `_origin` : `ByStr20 with end` : The account address that initiated the current transaction (which can possibly be a chain call). This is always a user address, since contracts can never initiate transactions.

The type `ByStr20 with end` is an *address type*. Address types are explained in detail in the [Addresses](#) section.

---

**Note:** Transition parameters must be of a *serialisable* type.

Serialisable:

- Byte strings, integers, strings, and block numbers are serialisable.
- Addresses are serialisable only as `ByStr20` values.
- ADT are serialisable if the types of their subvalues are serialisable. This means that the type of every constructor argument must be serialisable.

Not Serialisable:

- Messages, events and the special `Unit` type are not serialisable.
- Function types and map types are not serialisable.
- Complex types involving uninstantiated type variables are not serialisable.

## Procedures

*Procedures* are another way to define how the state of the contract may change, but in contrast to transitions, procedures are not part of the public interface of the contract, and may not be invoked by sending a message to the contract. The only way to invoke a procedure is to call it from a transition or from another procedure.

Procedures are defined with the keyword `procedure` followed by the parameters to be passed. The definition ends with the `end` keyword.

```
procedure foo (vname_1 : vtype_1, vname_2 : vtype_2, ...)
  ...
end
```

where `vname : vtype` specifies the name and type of each parameter and multiple parameters are separated by `,`

Once a procedure is defined it is available to be invoked from transitions and procedures in the rest of the contract file. It is not possible to invoke a procedure from transition or procedure defined earlier in the contract, nor is it possible for a procedure to call itself recursively.

Procedures are invoked using the name of the procedure followed by the actual arguments to the procedure:

```
v1 = ...;
v2 = ...;
foo v1 v2;
```

All arguments must be supplied when the procedure is invoked. A procedure does not return a result.

---

**Note:** The implicit transition parameters `_sender`, `_origin` and `_amount` are implicitly passed to all the procedures that a transition calls. There is therefore no need to declare those parameters explicitly when defining a procedure.

---



---

**Note:** Procedure parameters cannot be (or contain) maps. If a procedure needs to access a map, it is therefore necessary to either make the procedure directly access the contract field containing the map, or use a library function to perform the necessary computations on the map.

---

## Expressions

*Expressions* handle pure operations. Scilla contains the following types of expressions:

- `let x = f`: Give `f` the name `x` in the contract. The binding of `x` to `f` is **global** and extends to the end of the contract. The following code fragment defines a constant `one` whose value is 1 of type `Int32` throughout the contract.

```
let one = Int32 1
```

- `let x = f in expr`: Bind `f` to the name `x` within expression `expr`. The binding here is **local** to `expr` only. The following example binds the value of `one` to 1 of type `Int32` and `two` to 2 of type `Int32` in the expression `builtin add one two`, which adds 1 to 2 and hence evaluates to 3 of type `Int32`.

```
let sum =  
  let one = Int32 1 in  
  let two = Int32 2 in  
  builtin add one two
```

- { <entry>\_1 ; <entry>\_2 ... }: Message or event expression, where each entry has the following form:  $b : x$ . Here  $b$  is an identifier and  $x$  a variable, whose value is bound to the identifier in the message.
- $\text{fun } (x : T) \Rightarrow \text{expr}$ : A function that takes an input  $x$  of type  $T$  and returns the value to which expression  $\text{expr}$  evaluates.
- $f \ x$ : Apply the function  $f$  to the parameter  $x$ .
- $\text{tfun } 'T \Rightarrow \text{expr}$ : A type function that takes  $'T$  as a parametric type and returns the value to which expression  $\text{expr}$  evaluates. These are typically used to build library functions. See the implementation of *fst* for an example.

---

**Note:** Shadowing of type variables is not currently allowed. E.g.  $\text{tfun } 'T \Rightarrow \text{tfun } 'T \Rightarrow \text{expr}$  is not a valid expression.

---

- $@x \ T$ : Apply the type function  $x$  to the type  $T$ . This specialises the type function  $x$  by instantiating the first type variable of  $x$  to  $T$ . Type applications are typically used when a library function is about to be applied. See the example application of *fst* for an example.
- $\text{builtin } f \ x$ : Apply the built-in function  $f$  on  $x$ .
- $\text{match}$  expression: Matches a bound variable with patterns and evaluates the expression in that clause. The  $\text{match}$  expression is similar to the  $\text{match}$  expression in OCaml. The pattern to be matched can be an ADT constructor (see *ADTs*) with subpatterns, a variable, or a wildcard  $\_$ . An ADT constructor pattern matches values constructed with the same constructor if the subpatterns match the corresponding subvalues. A variable matches anything, and binds the variable to the value it matches in the expression of that clause. A wildcard matches anything, but the value is then ignored.

```
match x with  
| pattern_1 =>  
  expression_1 ...  
| pattern_2 =>  
  expression_2 ...  
| _ => (*Wildcard*)  
  expression ...  
end
```

---

**Note:** A pattern-match must be exhaustive, i.e., every legal (type-safe) value of  $x$  must be matched by a pattern. Additionally, every pattern must be reachable, i.e., for each pattern there must be a legal (type-safe) value of  $x$  that matches that pattern, and which does not match any pattern preceding it.

---

## Statements

Statements in Scilla are operations with effect, and hence not purely mathematical. Scilla contains the following types of statements:

- $x \leftarrow f$ : Fetch the value of the contract field  $f$ , and store it into the local variable  $x$ .

- `f := x` : Update the mutable contract field `f` with the value of `x`. `x` may be a local variable, or another contract field.
- `x <- & BLOCKNUMBER` : Fetch the value of the blockchain state variable `BLOCKNUMBER`, and store it into the local variable `x`.
- `x <- & c.f` : Remote fetch. Fetch the value of the contract field `f` at address `c`, and store it into the local variable `x`. Note that the type of `c` must be an address type containing the field `f`. See the section on [Addresses](#) for details on address types.
- `x <- & c as t` : Address type cast. Check whether `c` satisfies the type `t`. If it does, then store the value `v` into `x`, where `v` is the same value as `c`, but of type `t` rather than the type of `c`. If it does not, then store the value `None` into `x`. Note that `c` must be of type `ByStr20` or an address type, and that `t` must be an address type. See the section on [Addresses](#) for details on address types.
- `v = e` : Evaluate the expression `e`, and assign the value to the local variable `v`.
- `p x y z` : Invoke the procedure `p` with the arguments `x`, `y` and `z`. The number of arguments supplied must correspond to the number of arguments the procedure takes.
- `forall ls p` : Invoke procedure `p` for each element in the list `ls`. `p` should be defined to take exactly one argument whose type is equal to an element of the list `ls`.
- `match` : Pattern-matching at statement level:

```

match x with
| pattern_1 =>
  statement_11;
  statement_12;
  ...
| pattern_2 =>
  statement_21;
  statement_22;
  ...
| _ => (*Wildcard*)
  statement_n1;
  statement_n2;
  ...
end

```

- `accept` : Accept the QA of the message that invoked the transition. The amount is automatically added to the `_balance` field of the contract. If a message contains QA, but the invoked transition does not accept the money, the money is transferred back to the sender of the message. Not accepting the incoming amount (when it is non-zero) is not an error.
- `send` and `event` : Communication with the blockchain. See the next section for details.
- `In-place map operations` : Operations on contract fields of type `Map`. See the [Maps](#) section for details.

A sequence of statements must be separated by semicolons `;` :

```

transition T ()
  statement_1;
  statement_2;
  ...
  statement_n
end

```

Notice that the final statement does not have a trailing `;`, since `;` is used to separate statements rather than terminate them.

## Communication

A contract can communicate with other contract and user accounts through the `send` instruction:

- `send msgs` : Send a list of messages `msgs`.

The following code snippet defines a `msg` with four entries `_tag`, `_recipient`, `_amount` and `param`.

```
(*Assume contractAddress is the address of the contract being called and the_
↳contract contains the transition setHello*)
msg = { _tag : "setHello"; _recipient : contractAddress; _amount : Uint128 0;_
↳param : Uint32 0 };
```

A message passed to `send` must contain the mandatory fields `_tag`, `_recipient` and `_amount`.

The `_recipient` field (of type `ByStr20`) is the blockchain address that the message is to be sent to, and the `_amount` field (of type `Uint128`) is the number of QA to be transferred to that account.

The `_tag` field (of type `String`) is only used when the value of the `_recipient` field is the address of a contract. In this case, the value of the `_tag` field is the name of the transition that is to be invoked on the recipient contract. If the recipient is a user account, the `_tag` field is ignored.

---

**Note:** To make it possible to transfer funds from a contract to both contracts and user accounts, use a standard transition name as per [ZRC-5](#), i.e. `AddFunds`. Please make sure to check if a contract to which you intend to send funds is implemented in adherence with [ZRC-5](#) convention.

---

In addition to the compulsory fields the message may contain other fields (of any type), such as `param` above. However, if the message recipient is a contract, the additional fields must have the same names and types as the parameters of the transition being invoked on the recipient contract.

Here's an example that sends multiple messages.

```
msg1 = { _tag : "setFoo"; _recipient : contractAddress1; _amount : Uint128 0;
↳ foo : Uint32 101 };
msg2 = { _tag : "setBar"; _recipient : contractAddress2; _amount : Uint128 0;
↳ bar : Uint32 100 };
msgs =
  let nil = Nil {Message} in
  let m1 = Cons {Message} msg1 nil in
  Cons msg2 m1
;
send msgs
```

---

**Note:** A transition may execute a `send` at any point during execution (including during the execution of the procedures it invokes), but the messages will not be sent onwards until after the transition has finished. More details can be found in the [Chain Calls](#) section.

---

---

**Note:** Messages sent to a user-defined library are ignored. They do not cause the transaction to fail, but they have no effect.

---

A contract can also communicate to the outside world by emitting events. An event is a signal that gets stored on the blockchain for everyone to see. If a user uses a client application to invoke a transition on a contract, the client application can listen for events that the contract may emit, and alert the user.

- `event e`: Emit a message `e` as an event. The following code emits an event with name `e_name`.

```
e = { _eventname : "e_name"; <entry>_2 ; <entry>_3 };
event e
```

An emitted event must contain the compulsory field `_eventname` (of type `String`), and may contain other entries as well. The value of the `_eventname` entry must be a string literal. All events with the same name must have the same entry names and types.

---

**Note:** As with the sending of messages, a transition may emit events at any point during execution (including during the execution of the procedures it invokes), but the event will not be visible on the blockchain before the transition has finished. More details can be found in the *Chain Calls* section.

---

## Run-time Errors

A transition may encounter a run-time error during execution, such as out-of-gas errors, integer overflows, or deliberately thrown exceptions. A run-time error causes the transition to terminate abruptly, and the entire transaction to be aborted. However, gas is still charged up until the point of the error.

The syntax for throwing an exception is similar to that of events and messages.

```
e = { _exception : "InvalidInput"; <entry>_2; <entry>_3 };
throw e
```

Unlike that for `event` or `send`, The argument to `throw` is optional and can be omitted. An empty throw will result in an error that just conveys the location of where the `throw` happened without more information.

---

**Note:** If a run-time error occurs during the execution of a transition, then the entire transaction is aborted, and any state changes in both the current and other contracts are rolled back. (The state of other contracts may have changed due to a chain call).

In particular:

- All transferred funds are returned to their respective senders, even if an `accept` was executed before the error.
- The message queue is cleared, so that as yet unprocessed messages will no longer be sent onwards even if a `send` was executed before the error.
- The event list is cleared, so that no events are emitted even if an `event` was executed before the error.

Gas is still charged for the transaction up until the point the run-time error occurs.

---

**Note:** Scilla does not have exception handlers. Throwing an exception always aborts the entire transaction.

---

## Gas consumption in Scilla

Deploying contracts and executing transitions in them cost gas. The detailed cost mechanism is explained [here](#).

The [Nucleus Wallet](#) page can be used to estimate gas costs for some transactions .

## 3.4.2 Primitive Data Types & Operations

### Integer Types

Scilla defines signed and unsigned integer types of 32, 64, 128, and 256 bits. These integer types can be specified with the keywords `IntX` and `UintX` where `X` can be 32, 64, 128, or 256. For example, the type of an unsigned integer of 32 bits is `Uint32`.

The following code snippet declares a variable of type `Uint32`:

```
let x = Uint32 43
```

Scilla supports the following built-in operations on integers. Each operation takes two integers `IntX`/`UintX` (of the same type) as arguments. Exceptions are `pow` whose second argument is always `Uint32` and `isqrt` which takes in a single `UintX` argument.

- `builtin eq i1 i2`: Is `i1` equal to `i2`? Returns a `Bool`.
- `builtin add i1 i2`: Add integer values `i1` and `i2`. Returns an integer of the same type.
- `builtin sub i1 i2`: Subtract `i2` from `i1`. Returns an integer of the same type.
- `builtin mul i1 i2`: Integer product of `i1` and `i2`. Returns an integer of the same type.
- `builtin div i1 i2`: Integer division of `i1` by `i2`. Returns an integer of the same type.
- `builtin rem i1 i2`: The remainder of integer division of `i1` by `i2`. Returns an integer of the same type.
- `builtin lt i1 i2`: Is `i1` less than `i2`? Returns a `Bool`.
- `builtin pow i1 i2`: `i1` raised to the power of `i2`. Returns an integer of the same type as `i1`.
- `builtin isqrt i`: Computes the integer square root of `i`, i.e. the largest integer `j` such that `j * j <= i`. Returns an integer of the same type as `i`.
- `builtin to_nat i1`: Convert a value of type `Uint32` to the equivalent value of type `Nat`.
- `builtin to_(u)int32/64/128/256`: Convert a `UintX`/`IntX` or a `String` (that represents a decimal number) value to the result of `Option UintX` or `Option IntX` type. Returns `Some res` if the conversion succeeded and `None` otherwise. The conversion may fail when
  - there is not enough bits to represent the result;
  - when converting a negative integer (or a string representing a negative integer) into a value of an unsigned type;
  - the input string cannot be parsed as an integer.

Here is the list of concrete conversion builtins for better discoverability: `to_int32`, `to_int64`, `to_int128`, `to_int256`, `to_uint32`, `to_uint64`, `to_uint128`, `to_uint256`.

Addition, subtraction, multiplication, `pow`, division and remainder operations may raise integer overflow, underflow and `division_by_zero` errors. This aborts the execution of the current transition and unrolls all the state changes made so far.

---

**Note:** Variables related to blockchain money, such as the `_amount` entry of a message or the `_balance` field of a contract, are of type `Uint128`.

---



## Strings

String literals in Scilla are expressed using a sequence of characters enclosed in double quotes. Variables can be declared by specifying using keyword `String`.

The following code snippet declares a variable of type `String`:

```
let x = "Hello"
```

Scilla supports the following built-in operations on strings:

- `builtin eq s1 s2`: Is `s1` equal to `s2`? Returns a `Bool`. `s1` and `s2` must be of type `String`.
- `builtin concat s1 s2`: Concatenate string `s1` with string `s2`. Returns a `String`.
- `builtin substr s idx len`: Extract the substring of `s` of length `len` starting from position `idx`. `idx` and `len` must be of type `UInt32`. Character indices in strings start from 0. Returns a `String` or fails with a runtime error if the combination of the input parameters results in an invalid substring.
- `builtin to_string x`: Convert `x` to a string literal. Valid types of `x` are `IntX`, `UIntX`, `ByStrX` and `ByStr`. Returns a `String`. Byte strings are converted to textual hexadecimal representation.
- `builtin strlen s`: Calculate the length of `s` (of type `String`). Returns a `UInt32`.
- `builtin strrev s`: Returns the reverse of the string `s`.
- `builtin to_ascii h`: Reinterprets a byte string (`ByStr` or `ByStrX`) as a printable ASCII string and returns an equivalent `String` value. If the byte string contains any non-printable characters, a runtime error is raised.

## Byte strings

Byte strings in Scilla are represented using the types `ByStr` and `ByStrX`, where `X` is a number. `ByStr` refers to a byte string of arbitrary length, whereas for any `X`, `ByStrX` refers to a byte string of fixed length `X`. For instance, `ByStr20` is the type of byte strings of length 20, `ByStr32` is the type of byte strings of length 32, and so on.

Byte strings literals in Scilla are written using hexadecimal characters prefixed with `0x`. Note that it takes 2 hexadecimal characters to specify 1 byte, so a `ByStrX` literal requires  $2 * X$  hexadecimal characters. The following code snippet declares a variable of type `ByStr32`:

```
let x = 0x123456789012345678901234567890123456789012345678901234567890abff
```

Scilla supports the following built-in operations for computing on and converting between byte string types:

- `builtin to_bystr h`: Convert a value `h` of type `ByStrX` (for some known `X`) to one of arbitrary length of type `ByStr`.
- `builtin to_bystrX h`: (note that `X` is a numerical parameter here and not a part of the builtin name, see the examples below)
  - if the argument `h` is of type `ByStr`: Convert an arbitrary size byte string value `h` (of type `ByStr`) to a fixed sized byte string of type `ByStrX`, with length `X`. The result is of type `Option ByStrX` in this case: the builtin returns `Some res` if the length of the argument is equal to `X` and `None` otherwise. E.g. `builtin to_bystr42 bs` returns `Some bs'` if the length of `bs` is 42.
  - if the argument `h` is of type `UInt (32/64/128/256)`: Convert unsigned integers to their big endian byte representation, returning a `ByStr (4/8/16/32)` value (notice it's not an optional type in this case). For instance, `builtin to_bystr4 x` (this only typechecks if `x` has type `UInt32`) or `builtin to_bystr16 x` (this only typechecks if `x` is of type `UInt128`).

- builtin `to_uint (32/64/128/256) h` : Convert a fixed sized byte string value `h` to an equivalent value of type `Uint (32/64/128/256)`. `h` must be of type `ByStrX` for some known `X` less than or equal to `(4/8/16/32)`. A big-endian representation is assumed.
- builtin `concat h1 h2`: Concatenate byte strings `h1` and `h2`.
  - If `h1` has type `ByStrX` and `h2` has type `ByStrY`, then the result will have type `ByStr (X+Y)`.
  - If the arguments are of type `ByStr`, the result is also of type `ByStr`.
- builtin `strlen h`: The length of byte string (`ByStr`) `h`. Returns `Uint32`.
- builtin `eq a1 a2`: Is `a1` equal to `a2`? Returns a `Bool`.

## Addresses

Addresses on the Zilliqa network are strings of 20 bytes, and raw addresses are therefore represented by values of type `ByStr20`.

Additionally, Scilla supports structured address types, i.e., types that are equivalent to `ByStr20`, but which, when interpreted as an address on the network, provide additional information about the contents of that address. Address types are written using the form `ByStr20` with `<address contents>` end, where `<address contents>` refers to what the address contains.

The hierarchy of address types is as follows:

- `ByStr20`: A raw byte string of length 20. The type does not provide any guarantee as to what is located at the address. (Generally, `ByStr20` is not regarded as an address type, because it can refer to any byte string of length 20, whether it is meant to represent an address or not.)
- `ByStr20` with end: A `ByStr20` which, when interpreted as a network address, refers to an address that is in use. An address is in use if it either contains a contract or a library, or if the balance or the *nonce* of the address is greater than 0. (The balance of an address is the number of *Qa* held by the address account. The nonce of an address is the number of transactions that have been initiated from that address).
- `ByStr20` with contract end: A `ByStr20` which, when interpreted as a network address, refers to the address of a contract.
- `ByStr20` with contract field `f1 : t1, field f2 : t2, ... end`: A `ByStr20` which, when interpreted as a network address, refers to the address of a contract containing the mutable fields `f1` of type `t1`, `f2` of type `t2`, and so on. The contract in question may define more fields than the ones specified in the type, but the fields specified in the type must be defined in the contract.

---

**Note:** All addresses in use, and therefore by extension all contract addresses, implicitly define a mutable field `_balance : Uint128`. For user accounts the `_balance` field refers to the account balance.

---

---

**Note:** Address types specifying immutable parameters or transitions of a contract are not supported.

---

## Address subtyping

The hierarchy of address types defines a subtype relation:

- Any address type `ByStr20` with `... end` is as subtype of `ByStr20`. This means that any address type can be used in place of a `ByStr20`, e.g., for comparing equality using builtin `eq`, or as the `_recipient` value of a message.

- Any contract address type `ByStr20 with contract ... end` is a subtype of `ByStr20 with end`.
- Any contract address type specifying explicit fields `ByStr20 with contract field f1 : t11, field f2 : t12, ... end` is a subtype of a contract address type specifying a subset of those fields `ByStr20 with contract field f1 : t21, field f2 : t22, ... end`, provided that `t11` is a subtype of `t21`, `t12` is a subtype of `t22`, and so on for each field specified in both contract types.
- For ADTs with type parameters such as `List` or `Option`, an ADT `T t1 t2 ...` is a subtype of `S s1 s2 ...` if `T` is the same as `S`, and `t1` is a subtype of `s1`, `t2` is a subtype of `s2`, and so on.
- A map with key type `kt1` and value type `vt1` is a subtype of another map with key type `kt2` and value type `vt2` if `kt1` is a subtype of `kt2` and `vt1` is a subtype of `vt2`.

### Dynamic typecheck of addresses

In general, address types cannot be fully typechecked statically by the Scilla checker. This can happen, e.g., because a byte string is a transition parameter and thus not known statically, or because a byte string refers to an address that does not currently contain a contract, but which might contain a contract in the future.

For this reason immutable parameters (i.e., contract parameters supplied when the contract is deployed) and transition parameters of address types are typechecked dynamically, when the actual byte string is known.

For example, a contract might specify an immutable field `init_owner` as follows:

```
contract MyContract (init_owner : ByStr20 with end)
```

When the contract is deployed, the byte string supplied as `init_owner` is looked up as an address on the blockchain, and if the contents of that address matches the address type (in this case that the address is in use either by a user or by a contract), then deployment continues, and `init_owner` can be treated as a `ByStr20 with end` throughout the contract.

Similarly, a transition might specify a parameter `token_contract` as follows:

```
transition Transfer (
  token_contract : ByStr20 with contract
                    field balances : Map ByStr20 UInt128
                    end
)
```

When the transition is invoked, the byte string supplied as the `token_contract` parameter is looked up as an address on the blockchain, and if the contents of that address matches the address type (in this case that the address contains a contract with a field `balances` of a type that is assignable to `Map ByStr20 UInt128`), then the transition parameter is initialised successfully, and `token_contract` can be treated as a `ByStr20 with contract field balances : Map ByStr20 UInt128 end` throughout this transition invocation.

In either case, if the contents of the address does not match the specified type, then the dynamic typecheck is unsuccessful, causing deployment (for failed immutable parameters) or transition invocation (for transition parameters) to fail. A failed dynamic typecheck is considered a run-time error, causing the current transaction to abort. (For the purposes of dynamic typechecks of immutable fields the deployment of a contract is considered a transaction).

It is also possible to perform a dynamic typecheck as a statement, using the address type cast construct:

```
x <- & c as ByStr20 with ... end
```

If `c` satisfies the address type, then `x` is bound to `Some v`, where `v` is the same value as `c` but is treated as having the address type specified in `cast`. On the other hand, if `c` does not satisfy the address type, then `x` is bound to `None`.

The advantage of using an address type cast instead of a dynamic typecheck of transition parameters is that a type cast does not cause a run-time failure in case of a failed cast - a failed cast just results in a `None` value.

The disadvantage of address type casts is that they cost more gas. The type cast itself costs the same as a typecheck of a transition parameter, but one then needs to pattern-match on the result of the type cast, which costs additional gas.

**Note:** It is not possible to specify a `ByStr20` literal and have it interpreted as an address. In other words, the following code snippet will result in a static type error:

```
let x : ByStr20 with end = 0x1234567890123456789012345678901234567890
```

The only ways for a byte string to be validated against an address type is either to pass it as the value of an immutable field or as a transition parameter of the appropriate type, or to perform a cast statement to the appropriate type.

## Remote fetches

To perform a remote fetch `x <- & c.f`, the type of `c` must be some address type declaring the field `f`. For instance, if `c` has the type `ByStr20 with contract field paused : Bool end`, then the value of the field `paused` at address `c` can be fetched using the statement `x <- & c.paused`, whereas it is not possible to fetch the value of an undeclared field (e.g., `admin`) of `c`, even if the contract at address `c` does actually contain a field `admin`. To be able to fetch the value of the `admin` field, the type of `c` must contain the `admin` field as well, e.g., `ByStr20 with contract field paused : Bool, field admin : ByStr20 end`.

Remote fetches of map fields can be performed using in-place operations in the same way as for locally declared map fields, i.e., `x <- & c.m[key]`, `x <- & c.m[key1][key2]`, `x <- & exists m[key]`, etc. As with remote fetches of map fields, the remote map field must be declared in the type of `c`, e.g., `ByStr20 with contract field m : Map Uint128 (Map Uint32 Bool) end`.

Writing to a remote field is not allowed.

## Crypto Built-ins

A hash in Scilla is declared using the data type `ByStr32`. A `ByStr32` represents a hexadecimal byte string of 32 bytes (64 hexadecimal characters). A `ByStr32` literal is prefixed with `0x`.

Scilla supports the following built-in operations on hashes and other cryptographic primitives, including byte sequences.

- `builtin eq h1 h2`: Is `h1` equal to `h2`? Both inputs are of the same type `ByStrX` (or both are of type `ByStr`). Returns a `Bool`.
- `builtin sha256hash x`: Convert `x` of any non-closure type to its SHA256 hash. Returns a `ByStr32`.
- `builtin keccak256hash x`: Convert `x` of any non-closure type to its Keccak256 hash. Returns a `ByStr32`.
- `builtin ripemd160hash x`: Convert `x` of any non-closure type to its RIPEMD-160 hash. Returns a `ByStr20`.
- `builtin substr h idx len`: Extract the sub-byte-string of `h` of length `len` starting from position `idx`. `idx` and `len` must be of type `Uint32`. Character indices in byte strings start from 0. Returns a `ByStr` or fails with a runtime error.
- `builtin strrev h`: Reverse byte string (either `ByStr` or `ByStrX`). Returns a value of the same type as the argument.
- `builtin schnorr_verify pubk data sig`: Verify a signature `sig` of type `ByStr64` against a byte string `data` of type `ByStr` with the Schnorr public key `pubk` of type `ByStr33`.

- builtin `schnorr_get_address pubk`: Given a public key of type `ByStr33`, returns the `ByStr20` Zilliqa address that corresponds to that public key.
- builtin `ecdsa_verify pubk data sig`: Verify a signature `sig` of type `ByStr64` against a byte string `data` of type `ByStr` with the ECDSA public key `pubk` of type `ByStr33`.
- builtin `ecdsa_recover_pk data sig recid`: Recover data (of type `ByStr`), having signature `sig` (of type `ByStr64`) and a `Uint32` recovery integer `recid`, whose value is restricted to be 0, 1, 2 or 3, the uncompressed public key, returning a `ByStr65` value.
- builtin `bech32_to_bystr20 prefix addr`. The builtin takes a network specific prefix ("`zil`" / "`tzil`") of type `String` and an input bech32 string (of type `String`) and if the inputs are valid, converts it to a raw byte address (`ByStr20`). The return type is `Option ByStr20`. On success, `Some addr` is returned and on invalid inputs `None` is returned.
- builtin `bystr20_to_bech32 prefix addr`. The builtin takes a network specific prefix ("`zil`" / "`tzil`") of type `String` and an input `ByStr20` address, and if the inputs are valid, converts it to a bech32 address. The return type is `Option String`. On success, `Some addr` is returned and on invalid inputs `None` is returned.
- builtin `alt_bn128_G1_add p1 p2`. The builtin takes two points `p1, p2` on the `alt_bn128` curve and returns the sum of the points in the underlying group `G1`. The input points are each a `Pair {ByStr32 ByStr32}`. The return type is `Option (Pair (ByStr32) (ByStr32))`. On success, `Some point` is returned and on invalid inputs `None` is returned. Each scalar component `ByStr32` of a point is a big-endian encoded number. Also see <https://github.com/ethereum/EIPs/blob/master/EIPS/eip-196.md>
- builtin `alt_bn128_G1_mul p s`. The builtin takes a point `p` on the `alt_bn128` curve (as described previously), and a scalar `ByStr32` value `s` and returns the sum of the point `p` taken `s` times. The return type is `Option (Pair (ByStr32) (ByStr32))`. On success, `Some point` is returned and on invalid inputs `None` is returned.
- builtin `alt_bn128_G1_neg p`. The builtin takes a point `p` on the `alt_bn128` curve and returns the negation of the point. The result is `Option (Pair (ByStr32) (ByStr32))`. On success, `Some point` is returned and on invalid inputs `None` is returned.
- builtin `alt_bn128_pairing_product pairs`. This builtin takes in a list of pairs `pairs` of points. Each pair consists of a point in group `G1` (`Pair {ByStr32 ByStr32}`) as the first component and a point in group `G2` (`Pair {ByStr64 ByStr64}`) as the second component. Hence the argument has type `List {(Pair (Pair ByStr32 ByStr32) (Pair ByStr64 ByStr64)) }`. The function applies a pairing function on each point to check for equality and returns `True` or `False` depending on whether the pairing check succeeds or fails. Also see <https://github.com/ethereum/EIPs/blob/master/EIPS/eip-197.md>

## Maps

A value of type `Map kt vt` provides a key-value store where `kt` is the type of keys and `vt` is the type of values (in some other programming languages datatypes like Scilla's `Map` are called associative arrays, symbol tables, or dictionaries). The type of map keys `kt` may be any one of the following *primitive* types: `String`, `IntX`, `UintX`, `ByStrX`, `ByStr` or `BNum`. The type of values `vt` may be any type except a function type, this includes both builtin and user-defined algebraic datatypes.

Since compound types are not supported as map key types, the way to model, e.g. association of pairs of values to another value is by using *nested* maps. For instance, if one wants to associate with an account and a particular trusted user some money limit the trusted user is allowed to spend on behalf of the account, one can use the following nested map:

```
field operators: Map ByStr20 (Map ByStr20 Uint128)
  = Emp ByStr20 (Map ByStr20 Unit)
```

The first and the second key are of type `ByStr20` and represent accounts and the trusted users correspondingly. We represent the money limits with the `Uint128` type.

Scilla supports a number of operations on map, which can be categorized as

- *in-place* operations which modify *field* maps without making any copies, hence they belong to the imperative fragment of Scilla. These operations are efficient and recommended to use in almost all of the cases;
- *functional* map operations are intended to use in pure functions, e.g. when designing a Scilla library, because they never modify the original map they are called on. These operations may incur significant performance overhead as some of them create a new (modified) copy of the input map. Syntactically, the copying operations are all prefixed with `builtin` keyword (see below). Note that to call the functional builtins on a field map one first needs to make a *copy* of the field map using a command like so: `map_copy <- field_map`, which results in gas consumption proportional to the size of `field_map`.

### In-place map operations

- `m[k] := v`: *In-place* insert. Inserts a key `k` bound to a value `v` into a map `m`. If `m` already contains key `k`, the old value bound to `k` gets replaced by `v` in the map. `m` must refer to a mutable field in the current contract. Insertion into nested maps is supported with the syntax `m[k1][k2][...]` `:= v`. If the intermediate keys do not exist in the nested maps, they are freshly created along with the map values they are associated with.
- `x <- m[k]`: *In-place* local fetch. Fetches the value associated with the key `k` in the map `m`. `m` must refer to a mutable field in the current contract. If `k` has an associated value `v` in `m`, then the results of the fetch is `Some v` (see the `Option` type below), otherwise the result is `None`. After the fetch, the result gets bound to the local variable `x`. Fetching from nested maps is supported with the syntax `x <- m[k1][k2][...]`. If one or more of the intermediate keys do not exist in the corresponding map, the result of the fetch is `None`.
- `x <- & c.m[k]`: *In-place* remote fetch. Works in the same way as the local fetch operation, except that `m` must refer to a mutable field in the contract at address `c`.
- `x <- exists m[k]`: *In-place* local key existence check. `m` must refer to a mutable field in the current contract. If `k` has an associated value in the map `m` then the result (of type `Bool`) of the check is `True`, otherwise the result is `False`. After the check, the result gets bound to the local variable `x`. Existence checks through nested maps is supported with the syntax `x <- exists m[k1][k2][...]`. If one or more of the intermediate keys do not exist in the corresponding map, the result is `False`.
- `b <- & exists c.m[k]`: *In-place* remote key existence check. Works in the same way as the local key existence check, except that `m` must refer to a mutable field in the contract at address `c`.
- `delete m[k]`: *In-place* remove. Removes a key `k` and its associated value from the map `m`. The identifier `m` must refer to a mutable field in the current contract. Removal from nested maps is supported with the syntax `delete m[k1][k2][...]`. If one or more of the intermediate keys do not exist in the corresponding map, then the operation has no effect. Note that in the case of a nested removal `delete m[k1][...][kn-1][kn]`, only the key-value association of `kn` is removed. The key-value bindings of `k1` to `kn-1` will remain in the map.

### Functional map operations

- `builtin put m k v`: Insert a key `k` bound to a value `v` into a map `m`. Returns a new map which is a copy of the `m` but with `k` associated with `v`. If `m` already contains key `k`, the old value bound to `k` gets replaced by `v` in the result map. The value of `m` is unchanged. The `put` function is typically used in library functions. Note that `put` makes a copy of `m` before inserting the key-value pair.
- `builtin get m k`: Fetch the value associated with the key `k` in the map `m`. Returns an optional value (see the `Option` type below) – if `k` has an associated value `v` in `m`, then the result is `Some v`, otherwise the result is `None`. The `get` function is typically used in library functions.

- `builtin contains m k`: Is the key `k` associated with a value in the map `m`? Returns a `Bool`. The `contains` function is typically used in library functions.
- `builtin remove m k`: Remove a key `k` and its associated value from the map `m`. Returns a new map which is a copy of `m` but with `k` being unassociated with a value. The value of `m` is unchanged. If `m` does not contain key `k` the `remove` function simply returns a copy of `m` with no indication that `k` is missing. The `remove` function is typically used in library functions. Note that `remove` makes a copy of `m` before removing the key-value pair.
- `builtin to_list m`: Convert a map `m` to a `List (Pair kt vt)` where `kt` and `vt` are key and value types, respectively (see the `List` type below).
- `builtin size m`: Return the number of bindings in map `m`. The result type is `Uint32`. Calling this builtin consumes a small *constant* amount of gas. But calling it directly on a *field* map is not supported, meaning that getting the size of field maps while avoiding expensive copying needs some more scaffolding, which you can find out about in the *Field map size* section.

---

**Note:** Builtin functions `put` and `remove` return a new map, which is a possibly modified copy of the original map. This may affect performance!

---



---

**Note:** Empty maps can be constructed using the `Emp` keyword, specifying the key and value types as its arguments. This is the way to initialise `Map` fields to be empty. For example `field foomap : Map Uint128 String = Emp Uint128 String` declares a `Map` field with keys of type `Uint128` and values of type `String`, which is initialized to be the empty map.

---

## Block Numbers

Block numbers have a dedicated type `BNum` in Scilla. Variables of this type are specified with the keyword `BNum` followed by an integer value (for example `BNum 101`).

Scilla supports the following built-in operations on block numbers:

- `eq b1 b2`: Is `b1` equal to `b2`? Returns a `Bool`.
- `blt b1 b2`: Is `b1` less than `b2`? Returns a `Bool`.
- `badd b1 i1`: Add `i1` of type `UintX` to `b1` of type `BNum`. Returns a `BNum`.
- `bsub b1 b2`: Subtract `b2` from `b1`, both of type `BNum`. Returns an `Int256`.

### 3.4.3 Algebraic Datatypes

An *algebraic datatype* (ADT) is a composite type used commonly in functional programming. Each ADT is defined as a set of **constructors**. Each constructor takes a set of arguments of certain types.

Scilla is equipped with a number of built-in ADTs, which are described below. Additionally, Scilla allows users to define their own ADTs.

#### Boolean

Boolean values are specified using the type `Bool`. The `Bool` ADT has two constructors `True` and `False`, neither of which take any arguments. Thus the following code fragment constructs a value of type `Bool` by using the constructor `True`:

```
x = True
```

## Option

Optional values are specified using the type `Option t`, where `t` is some type. The `Option` ADT has two constructors:

- `Some` represents the presence of a value. The `Some` constructor takes one argument (the value, of type `t`).
- `None` represents the absence of a value. The `None` constructor takes no arguments.

The following code snippet constructs two optional values. The first value is an absent string value, constructed using `None`. The second value is the `Int32` value 10, which, because the value is present, is constructed using `Some`:

```
let none_value = None {String}

let some_value =
  let ten = Int32 10 in
  Some {Int32} ten
```

Optional values are useful for initialising fields where the value is not yet known:

```
field empty_bool : Option Bool = None {Bool}
```

Optional values are also useful for functions that might not have a result, such as the `get` function for maps:

```
getValue = builtin get m _sender;
match getValue with
| Some v =>
  (* _sender was associated with v in m *)
  v = v + v;
  ...
| None =>
  (* _sender was not associated with a value in m *)
  ...
end
```

## List

Lists of values are specified using the type `List t`, where `t` is some type. The `List` ADT has two constructors:

- `Nil` represents an empty list. The `Nil` constructor takes no arguments.
- `Cons` represents a non-empty list. The `Cons` constructor takes two arguments: The first element of the list (of type `t`), and another list (of type `List t`) representing the rest of the list.

All elements in a list must be of the same type `t`. In other words, two values of different types cannot be added to the same list.

The following example shows how to build a list of `Int32` values. First we create an empty list using the `Nil` constructor. We then add four other values one by one using the `Cons` constructor. Notice how the list is constructed backwards by adding the last element, then the second-to-last element, and so on, so that the final list is `[11; 10; 2; 1]`:



```

let one = Int32 1 in
let two = Int32 2 in
let ten = Int32 10 in
let eleven = Int32 11 in

let nil = Nil {Int32} in
let l1 = Cons {Int32} one nil in
let l2 = Cons {Int32} two l1 in
let l3 = Cons {Int32} ten l2 in
  Cons {Int32} eleven l3

```

Scilla provides three structural recursion primitives for lists, which can be used to traverse all the elements of any list:

- `list_foldl`: `('B -> 'A -> 'B) -> 'B -> (List 'A) -> 'B`: Recursively process the elements in a list from front to back, while keeping track of an *accumulator* (which can be thought of as a running total). `list_foldl` takes three arguments, which all depend on the two type variables `'A` and `'B`:
  - The function processing the elements. This function takes two arguments. The first argument is the current value of the accumulator (of type `'B`). The second argument is the next list element to be processed (of type `'A`). The result of the function is the next value of the accumulator (of type `'B`).
  - The initial value of the accumulator (of type `'B`).
  - The list of elements to be processed (of type `(List 'A)`).

The result of applying `list_foldl` is the value of the accumulator (of type `'B`) when all list elements have been processed.

- `list_foldr`: `('A -> 'B -> 'B) -> 'B -> (List 'A) -> 'B`: Similar to `list_foldl`, except the list elements are processed from back to front. Notice also that the processing function takes the list element and the accumulator in the opposite order from the order in `list_foldl`.
- `list_foldk`: `('B -> 'A -> ('B -> 'B) -> 'B) -> 'B -> (List 'A) -> 'B`: Recursively process the elements in a list according to a *folding function*, while keeping track of an *accumulator*. `list_foldk` is a more general version of the left and right folds, which, by the way, can be both implemented in terms of it. `list_foldk` takes three arguments, which all depend on the two type variables `'A` and `'B`:
  - The function describing the fold step. This function takes three arguments. The first argument is the current value of the accumulator (of type `'B`). The second argument is the next list element to be processed (of type `'A`). The third argument represents the postponed recursive call (of type `'B -> 'B`). The result of the function is the next value of the accumulator (of type `'B`). The computation *terminates* if the programmer does not invoke the postponed recursive call. This is a major difference between `list_foldk` and the left and right folds which process their input lists from the beginning to the end unconditionally.
  - The initial value of the accumulator `z` (of type `'B`).
  - The list of elements to be processed (of type `List 'A`).

---

**Note:** When an ADT takes type arguments (such as `List 'A`), and occurs inside a bigger type (such as the type of `list_foldl`), the ADT and its arguments must be grouped using parentheses `( )`. This is the case even when the ADT occurs as the only argument to another ADT. For instance, when constructing an empty list of optional values of type `Int32`, one must instantiate the list type using the syntax `Nil {(Option Int32)}`.

---

To further illustrate the `List` type in Scilla, we show a small example using `list_foldl` to count the number of elements in a list. For an example of `list_foldk` see [list\\_find](#).

```

1 let list_length : forall 'A. List 'A -> UInt32 =
2   tfun 'A =>

```

(continues on next page)

(continued from previous page)

```

3   fun (l : List 'A) =>
4   let foldl = @list_foldl 'A Uint32 in
5   let init = Uint32 0 in
6   let one = Uint32 1 in
7   let iter =
8     fun (z : Uint32) =>
9     fun (h : 'A) =>
10      builtin add one z
11   in
12   foldl iter init l

```

`list_length` defines a function that takes a type argument `'A`, and a normal (value) argument `l` of type `List 'A`. `'A` is a *type variable* which must be instantiated by the code that intends to use `list_length`. The type variable is specified in line 2.

In line 4 we instantiate the types for `list_foldl`. Since we are traversing a list of values of type `'A`, we pass `'A` as the first type argument to `list_foldl`, and since we are calculating the length of the list (a non-negative integer), we pass `Uint32` as the accumulator type.

In line 5 we define the initial value of the accumulator. Since an empty list has length 0, the initial value of the accumulator is 0 (of type `Uint32`, to match the accumulator type).

In lines 6-10 we specify the processing function `iter`, which takes the current accumulator value `z` and the current list element `h`. In this case the processing function ignores the list element, and increments the accumulator by 1. When all elements in the list have been processed, the accumulator will have been incremented as many times as there are elements in the list, and hence the final value of the accumulator will be equal to the length of the list.

In line 12 we apply the type-instantiated version of `list_foldl` from line 4 to the processing function, the initial accumulator, and the list of values.

Common utilities for the `List` type (including `list_length`) are provided in the `ListUtils` library as part of the standard library distribution for Scilla.

## Pair

Pairs of values are specified using the type `Pair t1 t2`, where `t1` and `t2` are types. The `Pair` ADT has one constructor:

- `Pair` represents a pair of values. The `Pair` constructor takes two arguments, namely the two values of the pair, of types `t1` and `t2`, respectively.

---

**Note:** `Pair` is both the name of a type and the name of a constructor of that type. An ADT and a constructor typically only share their names when the constructor is the only constructor of the ADT.

---

A `Pair` value may contain values of different types. In other words, `t1` and `t2` need not be the same type.

Below is an example where we declare a field `pp` of type `Pair String Uint32`, which we then initialise by constructing a pair consisting of a value of type `String` and a value of type `Uint32`:

```

field pp: Pair String Uint32 =
    let s1 = "Hello" in
    let num = Uint32 2 in
    Pair {String Uint32} s1 num

```

Notice the difference in how we specify the type of the field as `Pair 'A' 'B'`, and how we specify the types of values given to the constructor as `Pair { 'A' 'B' }`.

We now illustrate how pattern matching can be used to extract the first element from a `Pair`. The function `fst` shown below is defined in the `PairUtils` library of the Scilla standard library.

```
let fst =
  tfun 'A =>
  tfun 'B =>
  fun (p : Pair ('A) ('B)) =>
    match p with
    | Pair a b => a
  end
```

To apply `fst` to one must first instantiate the type variables `'A` and `'B`, which is done as follows:

```
p <- pp;
fst_specialised = @fst String Uint32;
p_fst = fst_specialised p
```

The value associated with the identifier `p_fst` will be the string `"Hello"`.

---

**Note:** Using `Pair` is generally discouraged. Instead, the programmer should define an ADT which is specialised to the particular type of pairs that is needed in the particular use case. See the section on *User-defined ADTs* below.

---

## Nat

Peano numbers are specified using the type `Nat`. The `Nat` ADT has two constructors:

- `Zero` represents the number 0. The `Zero` constructor takes no arguments.
- `Succ` represents the successor of another Peano number. The `Succ` constructor takes one argument (of type `Nat`) which represents the Peano number that is one less than the current number.

The following code shows how to build the Peano number corresponding to the integer 3:

```
let three =
  let zero = Zero in
  let one = Succ zero in
  let two = Succ one in
  Succ two
```

Scilla provides two structural recursion primitives for Peano numbers, which can be used to traverse all the Peano numbers from a given `Nat` down to `Zero`:

- `nat_fold`: `('A -> Nat -> 'A) -> 'A -> Nat -> 'A`: Recursively process the succession of numbers from a `Nat` down to `Zero`, while keeping track of an accumulator. `nat_fold` takes three arguments, two of which depend on the type variable `'A`:
  - The function processing the numbers. This function takes two arguments. The first argument is the current value of the accumulator (of type `'A`). The second argument is the next Peano number to be processed (of type `Nat`). Incidentally, the next number to be processed is the predecessor of the current number being processed. The result of the function is the next value of the accumulator (of type `'A`).
  - The initial value of the accumulator (of type `'A`).
  - The first Peano number to be processed (of type `Nat`).

The result of applying `nat_fold` is the value of the accumulator (of type `'A`) when all Peano numbers down to Zero have been processed.

- `nat_foldk`: `('A -> Nat -> ('A -> 'A) -> 'A) -> 'A -> Nat -> 'A`: Recursively process the Peano numbers down to zero according to a *folding function*, while keeping track of an *accumulator*. `nat_foldk` is a more general version of the left fold allowing for early termination. It takes three arguments, two depending on the type variable `'A`.
  - The function describing the fold step. This function takes three arguments. The first argument is the current value of the accumulator (of type `'A`). The second argument is the predecessor of the Peano number being processed (of type `Nat`). The third argument represents the postponed recursive call (of type `'A -> 'A`). The result of the function is the next value of the accumulator (of type `'A`). The computation *terminates* if the programmer does not invoke the postponed recursive call. Left folds inevitably process the whole list whereas `nat_foldk` can differ in this regard.
  - The initial value of the accumulator `z` (of type `'A`).
  - The Peano number to be processed (of type `Nat`).

To better understand `nat_foldk`, we explain how `nat_eq` works. `nat_eq` checks to see if two Peano numbers are equivalent. Below is the program, with line numbers and an explanation.

```

1 let nat_eq : Nat -> Nat -> Bool =
2 fun (n : Nat) => fun (m : Nat) =>
3   let foldk = @nat_foldk Nat in
4   let iter =
5     fun (n : Nat) => fun (ignore : Nat) => fun (recurse : Nat -> Nat) =>
6       match n with
7       | Succ n_pred => recurse n_pred
8       | Zero => m      (* m is not zero in this context *)
9     end in
10  let remaining = foldk iter n m in
11  match remaining with
12  | Zero => True
13  | _ => False
14  end

```

Line 2 specifies that we take two Peano numbers `m` and `n`. Line 3 instantiates the type of `nat_foldk`, we give it `Nat` because we will be passing a `Nat` value as the fold accumulator.

Lines 4 to 8 specify the fold description, this is the first argument that `nat_foldk` takes usually of type `'A -> Nat -> ('A -> 'A) -> 'A` but we have specified that `'A` is `Nat` in this case. Our function takes the accumulator `n` and `ignore : Nat` is the predecessor of the number being processed which we don't care about in this particular case.

Essentially, we start accumulating the end result from `n` and iterate at most `m` times (see line 10), decrementing both `n` and `m` at each recursive step (lines 4 - 9). The `m` variable gets decremented implicitly because this is how `nat_foldk` works under the hood. And we explicitly decrement `n` using pattern matching (lines 6, 7). To continue iteratively decrement both `m` and `n` we use `recurse` on line 7. If the two input numbers are equal, we will get the accumulator (`n`) equal to zero in the end. We call the final value of the accumulator `remaining` on line 10. At the end we will be checking to see if our accumulator ended up at `Zero` to say if the input numbers are equal. The last lines, return `True` when the result of the fold is `Zero` and `False` otherwise as described above.

In the case when accumulator `n` reaches zero (line 8) while `m` still has not been fully processed, we stop iteration (hence no `recurse` on that line) and return a non-zero natural number to indicate inequality. Any number (e.g. `Succ Zero`) would do, but to make the code concise we return the original input number `m` because we know `iter` gets called on `m` only if it's not zero.

In the symmetrical case when `m` reaches zero while the accumulator `n` is still strictly positive, we indicate inequality, because `remaining` gets this final value of `n`.

## User-defined ADTs

In addition to the built-in ADTs described above, Scilla supports user-defined ADTs.

ADT definitions may only occur in the library parts of a program, either in the library part of the contract, or in an imported library. An ADT definition is in scope in the entire library in which it is defined, except that an ADT definition may only refer to other ADT definitions defined earlier in the same library, or in imported libraries. In particular, an ADT definition may not refer to itself in an inductive/recursive manner.

Each ADT defines a set of constructors. Each constructor specifies a number of types which corresponds to the number and types of arguments that the constructor takes. A constructor may be specified as taking no arguments.

The ADTs of a contract must have distinct names, and the set of all constructors of all ADTs in a contract must also have distinct names. Both the ADT and constructor names must begin with a capital letter ('A' - 'Z'). However, a constructor and an ADT may have the same name, as is the case with the `Pair` type whose only constructor is also called `Pair`.

As an example of user-defined ADTs, consider the following type declarations from a contract implementing a chess-like game called Shogi or Japanese Chess (<https://en.wikipedia.org/wiki/Shogi>). When in turn, a player can choose to either move one of his pieces, place a previously captured piece back onto the board, or resign and award the victory to the opponent.

The pieces of the game can be defined using the following type `Piece`:

```
type Piece =
| King
| GoldGeneral
| SilverGeneral
| Knight
| Lance
| Pawn
| Rook
| Bishop
```

Each of the constructors represents a type of piece in the game. None of the constructors take any arguments.

The board is represented as a set of squares, where each square has two coordinates:

```
type Square =
| Square of Uint32 Uint32
```

The type `Square` is an example of a type where a constructor has the same name as the type. This usually happens when a type has only one constructor. The constructor `Square` takes two arguments, both of type `Uint32`, which are the coordinates (the row and the column) of the square on the board.

Similar to the definition of the type `Piece`, we can define the type of direction of movement using a constructor for each of the legal directions as follows:

```
type Direction =
| East
| SouthEast
| South
| SouthWest
| West
| NorthWest
| North
| NorthEast
```

We are now in a position to define the type of possible actions that a user may choose to perform when in turn:

```
type Action =  
| Move of Square Direction Uint32 Bool  
| Place of Piece Square  
| Resign
```

If a player chooses to move a piece, she should use the constructor `Move`, and provide four arguments:

- An argument of type `Square`, indicating the current position of the piece she wants to move.
- An argument of type `Direction`, indicating the direction of movement.
- An argument of type `Uint32`, indicating the distance the piece should move.
- An argument of type `Bool`, indicating whether the moved piece should be promoted after being moved.

If instead the player chooses to place a previously captured piece back onto the board, she should use the constructor `Place`, and provide two arguments:

- An argument of type `Piece`, indicating which piece to place on the board.
- An argument of type `Square`, indicating the position the piece should be placed in.

Finally, if the player chooses to resign and award the victory to her opponent, she should use the constructor `Resign`. Since `Resign` does not take any arguments, no arguments should be provided.

To check which action a player has chosen we use a match statement or a match expression:

```
transition PlayerAction (action : Action)  
...  
  match action with  
  | Resign =>  
    ...  
  | Place piece square =>  
    ...  
  | Move square direction distance promote =>  
    ...  
  end;  
...  
end
```

## Type identity for user-defined ADTs

---

**Note:** Due to a bug in the Scilla implementation the information in this section is only valid from Scilla version 0.10.0 and forwards. Contracts written in Scilla versions prior to 0.10.0 and which exploit this bug will have to be rewritten and redeployed, as they will no longer work from version 0.10.0 and onwards.

---

Each type declaration defines a unique type. In particular this means that even if two libraries both define identical types, the types are considered different.

As an example, consider the following two contracts:

```
library C1Lib  
  
type T =  
| C1 of Uint32  
| C2 of Bool
```

(continues on next page)

(continued from previous page)

```

contract Contract1()

field contract2_address : ByStr20 = 0x1234567890123456789012345678901234567890

transition Sending ()
  c2 <- contract2_address;
  x = Uint32 0;
  out = C1 x;
  msg = { _tag : "Receiving" ; _recipient : c2 ; _amount : Uint128 0 ;
         param : out };
  no_msg = Nil {Message};
  msgs = Cons {Message} msg no_msg;
  send msgs
end

(* ***** *)

(* Contract2 is deployed at address 0x1234567890123456789012345678901234567890 *)
library C2Lib

type T =
| C1 of Uint32
| C2 of Bool

contract Contract2()

transition Receiving (param : T)
  match param with
  | C1 v =>
  | C2 b =>
  end
end
end

```

Even though both contracts define identical types `T`, the two types are considered different in Scilla. In particular this means that the message sent from `Contract1` to `Contract2` will not trigger the `Receiving` transition, because the value sent as the `param` message field has the type `T` from `Contract1`, whereas the type `T` from `Contract2` is expected.

In order to pass a value of a user-defined ADT as a parameter to a transition, the type must be defined in a user-defined library, which both the sending and the receiving contract must import:

```

library MutualLib

type T =
| C1 of Uint32
| C2 of Bool

(* ***** *)

import MutualLib

library C1Lib

contract Contract1()

field contract2_address : ByStr20 = 0x1234567890123456789012345678901234567890

```

(continues on next page)

(continued from previous page)

```

transition Sending ()
  c2 <- contract2_address;
  x = Uint32 0;
  out = C1 x;
  msg = { _tag : "Receiving" ; _recipient : c2 ; _amount : Uint128 0 ;
         param : out };
  no_msg = Nil {Message};
  msgs = Cons {Message} msg no_msg;
  send msgs
end

(* ***** *)

(* Contract2 is deployed at address 0x1234567890123456789012345678901234567890 *)

scilla_version 0

import MutualLib

library C2Lib

contract Contract2()

transition Receiving (param : T)
  match param with
  | C1 v =>
  | C2 b =>
  end
end

```

The section *User-defined Libraries* has more information on how to define and use libraries.

### 3.4.4 More ADT examples

To further illustrate how ADTs can be used, we provide some more examples and describe them in detail. Versions of both the functions described below can be found in the `ListUtils` part of the *Scilla standard library*.

#### Computing the Head of a List

The function `list_head` returns the first element of a list.

Since a list may be empty, `list_head` may not always be able to compute a result, and thus should return a value of the `Option` type. If the list is non-empty, and the first element is `h`, then `list_head` should return `Some h`. Otherwise, if the list is empty, `list_head` should return `None`.

The following code snippet shows the implementation of `list_head`, and how to apply it:

```

1 let list_head =
2   tfun 'A =>
3     fun (l : List 'A) =>
4       match l with
5         | Cons h t =>
6           Some {'A} h
7         | Nil =>
8           None {'A}

```

(continues on next page)



(continued from previous page)

```

9     end
10
11 let int_head = @list_head Int32 in
12
13 let one = Int32 1 in
14 let two = Int32 2 in
15 let three = Int32 3 in
16 let nil = Nil {Int32} in
17
18 let l1 = Cons {Int32} three nil in
19 let l2 = Cons {Int32} two l1 in
20 let l3 = Cons {Int32} one l2 in
21 int_head l3

```

Line 2 specifies that 'A is a type parameter to the function, while line 3 specifies that l is a (value) parameter of type List 'A. In other words, lines 1-3 specify a function list\_head which can be instantiated for any type 'A, and which takes as an argument a value of type List 'A.

The pattern-match in lines 4-9 matches on the value of l. In line 5 we match on the list constructor Cons h t, where h is the first element of the list, and t is the rest of the list. If the list is not empty then the match is successful, and we return the first element as an optional value Some h. In line 7 we match on the list constructor Nil. If the list is empty then the match is successful, and we return the optional value None indicating that there was no head element of the list.

Line 11 instantiates the list\_head function for the type Int32, so that list\_head can be applied to values of type List Int32. Lines 13-20 build a list of type List Int32, and line 21 invokes the instantiated list\_head function on the list that was built.

### Computing a Left Fold

The function list\_foldl returns the result of a left fold given a function  $f : 'B \rightarrow 'A \rightarrow 'B$ , accumulator  $z : 'B$  and list  $xs : List 'A$ . This can be implemented as a recursion primitive or a list utility function.

A left fold is a recursive application of an accumulator  $z$  and next list element  $x : 'A$  with  $f$  repetitively until there are no more list elements. For example the left fold on  $[1, 2, 3]$  using subtraction starting with accumulator 0 would be  $((0-1)-2)-3 = -6$ . The left fold is explained in pseudocode below, note that the result is always the accumulator type.

```

1 list_foldl _ z [] = z
2 list_foldl f z (x:xs) = list_foldl f (f z x) xs

```

The same can be achieved with list\_foldk by partially applying a left fold description; this avoids illegal direct recursion. Our fold description left\_f : 'B -> 'A -> ('B -> 'B) -> 'B takes arguments accumulator, next list element and recursive call. The recursive call will be supplied by the list\_foldk function. An implementation is explained below.

```

1 let list_foldl : forall 'A. forall 'B. ('B -> 'A -> 'B) -> 'B -> List 'A -> 'B =
2   tfun 'A => tfun 'B =>
3     fun (f : 'B -> 'A -> 'B) =>
4       let left_f = fun (z: 'B) => fun (x: 'A) =>
5         fun (recurse : 'B -> 'B) => let res = f z x in
6           recurse res in
7       let folder = @list_foldk 'A 'B in
8       folder left_f

```

On line 1, we declare the name and type signature as according to the first paragraph. On the second line, we say that the function takes two types as arguments 'A and 'B. The third line says that we take some function `f` to process the list element and accumulator, as in paragraph two.

On line 4, we define the fold description using `f`. The fold description does not take a function but instead it should be implemented in terms of some function, as according to the type signature, `left_f : 'B -> 'A -> ('B -> 'B) -> 'B`. `left_f` takes arguments as described in paragraph two. We calculate the new accumulator `f z x` and call it `res`. Then we recursively call with the new accumulator.

On line 7, we instantiate an instance of `list_foldk` that has the right types for the job using a type application.

On line 8, we partially apply `folder` with the left fold description. . What is significant about `list_foldk` is that when calling the description, it provides a recursive call to itself, changing to the next element in the list and respective tail each time. This results in a function that just needs the user to provide the updated accumulator in the description.

## Computing a Right Fold

The function `list_foldr` returns the result of a right fold given some function `f : 'A -> 'B -> 'B`, accumulator `z : 'B` and list `xs : List 'A`. Like `list_foldl`, this can be a recursion primitive or a list utility function.

A right fold is similar to a left fold but is reversed in a way. The right fold applies a function `f` with an accumulator `z` starting from the end and then combines with the second last element, third last element, etc... until it reaches the beginning. For example a right fold on the list `[1, 2, 3]` with subtraction starting with accumulator `0` would be equal to  $1 - (2 - (3 - 0)) = 2$ . It is listed below in pseudocode, note that the result is always the accumulator type.

```
1 list_foldr _ z [] = z
2 list_foldr f z (x:xs) = f x (list_foldr f z xs)
```

Like before, the same can be achieved with `list_foldk` by partially applying a right fold description. The fold description takes arguments accumulator `z : 'B`, next list element `x : 'A` and recursive call `recurse : 'B -> 'B`. The recursive call will be supplied by the `list_foldk` function. An implementation is explained below.

```
1 let list_foldr : forall 'A. forall 'B. ('A -> 'B -> 'B) -> 'B -> List 'A -> 'B =
2   tfun 'A => tfun 'B =>
3   fun (f : 'A -> 'B -> 'B) =>
4   let right_f = fun (z : 'B) => fun (x : 'A) =>
5     fun (recurse : 'B -> 'B) => let res = recurse z in f x res in
6   let folder = @list_foldk 'A 'B in
7   folder right_f
```

This is very similar to before. On line 1 we declare the name and type signature, according to the first paragraph. On line 2, we take two type arguments 'A and 'B. The third line says that we take some function `f` to process the list element `x : 'A` and accumulator `z`. The argument order is necessarily different to that of a left fold.

Following that we write a fold description like before. `list_foldk` processes lists from left to right. But we need `list_foldr` to emulate the right-to-left traversal. By calling `recurse z` on line 5 as our first action, we postpone actual computation with the combining function `f` preserving the original accumulator until the very end. Once the recursive call reaches an empty list it returns the original accumulator. Then the function calls `f x res` (line 5) will evaluate outwards combining from the end to the beginning, see paragraph two.

The recursive call `recurse z` on line 5 may seem to be the same each time but what is changing is the list element we process.

On line 6, we instantiate `list_foldk` by applying the types 'A and 'B to make a type-specific function. The last line we partially apply `folder` with the right fold description. Like before what is special about `list_foldk` is that it calls this function with a recursive call to itself that each time slightly truncates the list; this provides the recursion.

## Checking for Existence in a List

The function `list_exists` takes a predicate function and a list, and returns a value indicating whether the predicate holds for at least one element in the list.

A predicate function is a function returning a Boolean value, and since we want to apply it to elements in the list, the argument type of the function should be the same as the element type of the list.

`list_exists` should return either `True` (if the predicate holds for at least one element) or `False` (if the predicate does not hold for any element in the list), so the return type of `list_exists` should be `Bool`.

The following code snippet shows the implementation of `list_exists`, and how to apply it:

```

1 let list_exists =
2   tfun 'A =>
3   fun (f : 'A -> Bool) =>
4   fun (l : List 'A) =>
5     let folder = @list_foldl 'A Bool in
6     let init = False in
7     let iter =
8       fun (z : Bool) =>
9         fun (h : 'A) =>
10          let res = f h in
11          match res with
12          | True =>
13            True
14          | False =>
15            z
16        end
17     in
18     folder iter init l
19
20 let int_exists = @list_exists Int128 in
21 let f =
22   fun (a : Int128) =>
23     let three = Int128 3 in
24     builtin lt a three
25
26 (* build list l3 similar to the previous example *)
27 ...
28
29 (* check if l3 has at least one element satisfying f *)
30 int_exists f l3

```

As in the previous example `'A` is a type variable to the function. The function takes two arguments:

- A predicate `f`, i.e., a function that returns a `Bool`. In this case, `f` will be applied to elements of the list, so the argument type of the predicate should be `'A`. Hence, `f` should have the type `'A -> Bool`.
- A list of elements `l` of type `List 'A`, so that the type of the elements in the list matches the argument type of `f`.

To traverse the elements of the input list `l` we use `list_foldl`. In line 5 we instantiate `list_foldl` for lists with elements of type `'A` and for the accumulator type `Bool`. In line 6 we set the initial accumulator value to `False` to indicate that no element satisfying the predicate has yet been seen.

The processing function `iter` defined in lines 7-16 tests the predicate on the current list element, and returns an updated accumulator. If an element has been found which satisfies the predicate, the accumulator is set to `True` and remains so for the rest of the traversal.

The final value of the accumulator is either `True`, indicating that `f` returned `True` for at least one element in the list, or `False`, indicating that `f` returned `False` for all elements in the list.

In line 20 we instantiate `list_exists` to work on lists of type `Int128`. In lines 21-24 we define the predicate, which returns `True` if its argument is less than 3, and returns `False` otherwise.

Omitted in line 27 is building the same list `l3` as in the previous example. In line 30 we apply the instantiated `list_exists` to the predicate and the list.

### Finding the first occurrence satisfying a predicate

The function `list_find` searches for the first occurrence in a list that satisfies some predicate `p : 'A -> Bool`. It takes the predicate and the list, returning `Some {'A} x :: Option 'A` if `x` is the first element such that `p x` and `None {'A} :: Option 'A` otherwise.

Below we have an implementation of `list_find` that illustrates how to use `list_foldk`.

```

1 let list_find : forall 'A. ('A -> Bool) -> List 'A -> Option 'A =
2 tfun 'A =>
3 fun (p : 'A -> Bool) =>
4   let foldk = @list_foldk 'A (Option 'A) in
5   let init = None {'A} in
6   (* continue fold on None, exit fold when Some compare st. p(compare) *)
7   let predicate_step =
8     fun (ignore : Option 'A) => fun (x : 'A) =>
9     fun (recurse : Option 'A -> Option 'A) =>
10      let p_x = p x in
11      match p_x with
12      | True => Some {'A} x
13      | False => recurse init
14    end in
15   foldk predicate_step init

```

Like before, we take a type variable `'A` on line 2 and take the predicate on the next line. We begin by using this type variable to instantiate `foldk`, by giving it our processing type and return type. The processing type being the list element type and the result type being `Option 'A`. The next line is our accumulator, we assume that at the start of the search there is no satisfier.

On line 7, we write a fold description for `foldk`. This embodies the order of the recursion and conditions for recursion. `predicate_step` has the type `Option 'A -> 'A -> (Option 'A -> Option 'A) -> Option 'A`. The first argument is the accumulator, the second `x` is the next element to process and the third `recurse` is the recursive call. We do not care what the accumulator `ignore` is since if it mattered we will have already terminated.

On lines 10 to 12 check for `p x` and if so return `Some {'A} x`. In the case that `p x` does not hold, try again from scratch with the next element and so on via recursion. `recurse init` is in pseudo-code equal to  $\lambda k. \text{foldk } \text{predicate\_step } \text{init } k \text{ } xs$  where `xs` is the tail of our list of to be processed elements.

With the final line we partially apply `foldk` so that it just takes a list argument and gives us our final answer. The first argument of `foldk` gives us the specific fold we want, for example if you wanted a left fold you would replace `predicate_step` with something else.

### 3.4.5 Standard Libraries

The Scilla standard library is *documented* separately.

### 3.4.6 User-defined Libraries

In addition to the standard library provided by Scilla, users are allowed to deploy library code on the blockchain. Library files are allowed to only contain pure Scilla code (which is the same restriction that in-contract library code has). Library files must use the `.scillib` file extension.

Below is an example of a user-defined library that defines a single function `add_if_equal` that adds to `Uint128` values if they are equal and returns 0 otherwise.

```
import IntUtils

library ExampleLib

let add_if_equal =
  fun (a : Uint128) => fun (b : Uint128) =>
    let eq = uint128_eq a b in
    match eq with
    | True => builtin add a b
    | False => Uint128 0
```

The structure of a library file is similar to the structure of the library part of a Scilla contract. A library file contains definitions of variables and pure library functions, but does not contain an actual contract definition with parameters, fields, transitions and so on.

Of particular importance is that a library cannot declare fields. Therefore, all libraries are stateless and can only contain pure code.

Similar to how contracts can import libraries, a library can import other libraries (including user-defined libraries) too. The scope of variables in an imported library is restricted to the immediate importer. So if `X` imports library `Y` which in turn imports library `Z`, then the names in `Z` are not in scope in `X`, but only in `Y`. Cyclic dependencies in imports are not allowed and flagged as errors during the checking phase.

#### Local Development with User-defined Libraries

To use variables and functions declared in an external (user-defined) library module, the command line argument to the Scilla executables must include a `-libdir` option, along with a list of directories as an argument. If the Scilla file imports a library `ALib`, then the Scilla executable will search for a library file called `ALib.scillib` in the directories provided. If more than one directory contains a file with the correct name, then the directories are given priority in the same order as they are provided to the Scilla executable. Alternatively, the environment variable `SCILLA_STDLIB_PATH` can be set to a list of library directories.

`scilla-checker` typechecks library modules in the same way as contract modules. Similarly, `scilla-runner` can deploy libraries. Note that `scilla-runner` takes a `blockchain.json` as argument (the way it does for *Contract Creation*) to be command line argument compatible with contract creation.

#### User-defined Libraries on the Blockchain

While the Zilliqa blockchain is designed to provide the standard Scilla libraries to an executing contract, it must be provided with extra information to support user-defined libraries.

The `init.json` of a library must include a `Bool` entry named `_library`, set to `True`. Additionally, A contract or a library that imports user-defined libraries must include in its `init.json` an entry named `_extlibs`, of Scilla type `List (Pair String ByStr20)`. Each entry in the list maps an imported library's name to its address in the blockchain.

Continuing the previous example, a contract or library that imports `ExampleLib` should have the following entry in its `init.json`:

```
[
  ...,
  {
    "vname" : "_library",
    "type" : "Bool",
    "value": { "constructor": "True", "argtypes": [], "arguments": [] }
  }
  {
    "vname" : "_extlibs",
    "type" : "List(Pair String ByStr20)",
    "value" : [
      {
        "constructor" : "Pair",
        "argtypes" : ["String", "ByStr20"],
        "arguments" : ["ExampleLib", "0x9865567890123456789012345678901234567890123456abcd"]
      },
      ...
    ]
  }
]
```

## Namespaces

Import statements can be used to define separate namespaces for imported names. To push the names from a library `Foo` into the namespace `Bar`, use the statement `import Foo as Bar`. Accessing a variable `v` in `Foo` must now be done using the qualified name `Bar.v`. This is useful when importing multiple libraries that define the same name.

The same variable name must not be defined more than once in the same namespace, so if multiple imported libraries define the same name, then at most one of the libraries may reside in the default (unqualified) namespace. All other conflicting libraries must be pushed to separate namespaces.

Extending our previous example, shown below is a contract that imports `ExampleLib` in namespace `Bar`, to use the function `add_if_equal`.

```
scilla_version 0

import ExampleLib as Bar

library MyContract

let adder = fun (a : Uint128) => fun (b : Uint128) =>
  Bar.add_if_equal a b

contract MyContract ()
...
```

## 3.4.7 Scilla versions

### Major and Minor versions

Scilla releases have a major version, minor version and a patch number, denoted as `X.Y.Z` where `X` is the major version, `Y` is the minor version, and `Z` the patch number.

- Patches are usually bug fixes that do not impact the behaviour of existing contracts. Patches are backward compatible.

- Minor versions typically include performance improvements and feature additions that do not affect the behaviour of existing contracts. Minor versions are backward compatible until the latest major version.
- Major versions are not backward compatible. It is expected that miners have access to implementations of each major version of Scilla for running contracts set to that major version.

Within a major version, miners are advised to use the latest minor revision.

The command `scilla-runner -version` will print major, minor and patch versions of the interpreter being invoked.

## Contract Syntax

Every Scilla contract must begin with a major version declaration. The syntax is shown below:

```
(*****  
(*          Scilla version          *)  
(*****  
  
scilla_version 0  
  
(*****  
(*          Associated library      *)  
(*****  
  
library MyContractLib  
  
...  
  
(*****  
(*          Contract definition     *)  
(*****  
  
contract MyContract  
  
...
```

When deploying a contract the output of the interpreter contains the field `scilla_version : X.Y.Z`, to be used by the blockchain code to keep track of the version of the contract. Similarly, `scilla-checker` also reports the version of the contract on a successful check.

## The `init.json` file

In addition to the version specified in the contract source code, it is also required that the contract's `init.json` specifies the same version when the contract is deployed and when the contract's transitions are invoked. This eases the process for the blockchain code to decide which interpreter to invoke.

A mismatch in the versions specified in `init.json` and the source code will lead to a gas-charged error by the interpreter.

An example `init.json`:

```
[  
  {  
    "vname" : "_creation_block",  
    "type"  : "BNum",  
    "value" : "1"  }  
]
```

(continues on next page)

(continued from previous page)

```
    },  
    {  
      "vname" : "_scilla_version",  
      "type" : "Uint32",  
      "value" : "1",  
    }  
  ]
```

## Chain Calls

When a user invokes a transition on a contract by sending a message with the contract's address as the recipient, then that transition may then send one or more messages onwards, possibly invoking other transitions on other contracts. The resulting collection of messages, fund transfers, transition invocations, and contract state changes are referred to as a *transaction*.

A transition that sends a message invoking another transition (typically on another contract) is referred to as a *chain call*.

During a transaction a LIFO queue (i.e., a stack) of unprocessed messages is maintained. Initially, the message queue contains only the single message sent by the original user, but additional messages may be added to the queue when a transition performs a chain call.

When a transition finishes, its outgoing messages are added to the message queue. The first message in the queue is then removed from the queue for processing. If there are no messages left to process, then the transaction finishes, and all state changes and fund transfers are committed to the blockchain.

When a transition sends multiple messages, the messages are added to the queue in the following order:

- If multiple `send` statements are executed, then the messages of the last `send` are added first. This means that the messages of the first `send` get processed first.
- If a `send` statement is given a list with multiple messages in it, then the messages of the tail of the list are added to the queue before the head of the list is added. This means that the first message in the list gets processed first.

Any run-time failure during the execution of a transaction causes the entire transaction to be aborted, with no further statements being executed, no further messages being processed, all state changes being rolled back, and all transferred funds returned to their respective senders. However, gas is still charged for the transaction up until the point of the failure.

The total number of messages that can be sent in a single transaction is currently set at 20. The number is subject to revision in the future.

Contracts of different Scilla versions may invoke transitions on each other. The semantics of message passing between contracts is guaranteed to be backward compatible between major versions.

## Accounting

For the transfer of native ZIL funds, Scilla follows an *acceptance semantics*. For a transfer to take place between contracts the funds must explicitly be accepted by the recipient by executing an `accept` statement - it is not sufficient that the sender executes a `send` statement.

When a contract executes a `send` statement the `_amount` values of the outgoing messages are deducted from the `_balance` field. However, the funds are not yet added to the recipient's `_balance` field, because they have not yet been explicitly accepted by the recipient.

When the recipient executes an `accept` statement the `_amount` of the incoming message is added to the contract's `_balance` field.



If the recipient does not execute and `accept` statement, the `_amount` is transferred back to the sender before the next message is processed.

While the funds in a message are “in transit”, i.e., between the point at which they is sent by the sender and the point at which they are accepted by the recipient (or, in the case of non-acceptance, the point at which the recipient transition finishes executing), the funds are stored away in the message itself, and therefore do not appear in neither the sender’s nor the recipient’s balances.

---

**Note:** A user account (i.e., an address that does not hold a contract) implicitly accepts all incoming funds, but does not do so until the message carrying the funds is processed.

---

If a contract executes a `send` statement on messages for which the total amount is more than the contract’s current balance, then a run-time error occurs, and the entire transaction is aborted. In other words, no account balance may drop below 0 at any point during a transaction.

## 3.5 The Scilla Standard Library

The Scilla standard library contains five libraries: `BoolUtils.scilla`, `IntUtils.scilla`, `ListUtils.scilla`, `NatUtils.scilla` and `PairUtils.scilla`. As the names suggests these contracts implement utility operations for the `Bool`, `IntX`, `List`, `Nat` and `Pair` types, respectively.

To use functions from the standard library in a contract, the relevant library file must be imported using the `import` declaration. The following code snippet shows how to import the functions from the `ListUtils` and `IntUtils` libraries:

```
import ListUtils IntUtils
```

The `import` declaration must occur immediately before the contract’s own library declaration, e.g.:

```
import ListUtils IntUtils

library WalletLib
... (* The declarations of the contract's own library values and functions *)

contract Wallet ( ... )
... (* The transitions and procedures of the contract *)
```

Below, we present the functions defined in each of the library.

### 3.5.1 BoolUtils

- `andb` : `Bool -> Bool -> Bool`: Computes the logical AND of two `Bool` values.
- `orb` : `Bool -> Bool -> Bool`: Computes the logical OR of two `Bool` values.
- `negb` : `Bool -> Bool`: Computes the logical negation of a `Bool` value.
- `bool_to_string` : `Bool -> String`: Transforms a `Bool` value into a `String` value. `True` is transformed into `"True"`, and `False` is transformed into `"False"`.

### 3.5.2 IntUtils

- `intX_eq` : `IntX -> IntX -> Bool`: Equality operator specialised for each `IntX` type.

```

let int_list_eq = @list_eq Int64 in

let one = Int64 1 in
let two = Int64 2 in
let ten = Int64 10 in
let eleven = Int64 11 in

let nil = Nil {Int64} in
let l1 = Cons {Int64} eleven nil in
let l2 = Cons {Int64} ten l1 in
let l3 = Cons {Int64} two l2 in
let l4 = Cons {Int64} one l3 in

let f = int64_eq in
(* See if [2,10,11] = [1,2,10,11] *)
int_list_eq f l3 l4

```

- `uintX_eq` : `UintX -> UintX -> Bool`: Equality operator specialised for each `UintX` type.
- `intX_lt` : `IntX -> IntX -> Bool`: Less-than operator specialised for each `IntX` type.
- `uintX_lt` : `UintX -> UintX -> Bool`: Less-than operator specialised for each `UintX` type.
- `intX_neq` : `IntX -> IntX -> Bool`: Not-equal operator specialised for each `IntX` type.
- `uintX_neq` : `UintX -> UintX -> Bool`: Not-equal operator specialised for each `UintX` type.
- `intX_le` : `IntX -> IntX -> Bool`: Less-than-or-equal operator specialised for each `IntX` type.
- `uintX_le` : `UintX -> UintX -> Bool`: Less-than-or-equal operator specialised for each `UintX` type.
- `intX_gt` : `IntX -> IntX -> Bool`: Greater-than operator specialised for each `IntX` type.
- `uintX_gt` : `UintX -> UintX -> Bool`: Greater-than operator specialised for each `UintX` type.
- `intX_ge` : `IntX -> IntX -> Bool`: Greater-than-or-equal operator specialised for each `IntX` type.
- `uintX_ge` : `UintX -> UintX -> Bool`: Greater-than-or-equal operator specialised for each `UintX` type.

### 3.5.3 ListUtils

- `list_map` : `('A -> 'B) -> List 'A -> : List 'B`.

Apply `f` : `'A -> 'B` to every element of `l` : `List 'A`, constructing a list (of type `List 'B`) of the results.

```

(* Library *)
let f =
  fun (a : Int32) =>
    builtin sha256hash a

(* Contract transition *)
(* Assume input is the list [ 1 ; 2 ; 3 ] *)
(* Apply f to all values in input *)

```

(continues on next page)

(continued from previous page)

```

hash_list_int32 = @list_map Int32 ByStr32;
hashed_list = hash_list_int32 f input;
(* hashed_list is now [ sha256hash 1 ; sha256hash 2 ; sha256hash 3 ] *)

```

- `list_filter` : ('A -> Bool) -> List 'A -> List 'A.

Filter out elements on the list based on the predicate `f` : 'A -> Bool. If an element satisfies `f`, it will be in the resultant list, otherwise it is removed. The order of the elements is preserved.

```

(*Library*)
let f =
  fun (a : Int32) =>
    let ten = Int32 10 in
    builtin lt a ten

(* Contract transition *)
(* Assume input is the list [ 1 ; 42 ; 2 ; 11 ; 12 ] *)
less_ten_int32 = @list_filter Int32;
less_ten_list = less_ten_int32 f l
(* less_ten_list is now [ 1 ; 2 ]*)

```

- `list_head` : (List 'A) -> (Option 'A).

Return the head element of a list `l` : List 'A as an optional value. If `l` is not empty with the first element `h`, the result is `Some h`. If `l` is empty, then the result is `None`.

- `list_tail` : (List 'A) -> (Option List 'A).

Return the tail of a list `l` : List 'A as an optional value. If `l` is a non-empty list of the form `Cons h t`, then the result is `Some t`. If `l` is empty, then the result is `None`.

- `list_foldl_while` : ('B -> 'A -> Option 'B) -> 'B -> List 'A -> 'B

Given a function `f` : 'B -> 'A -> Option 'B, accumulator `z` : 'B and list `ls` : List 'A execute a left fold when our given function returns `Some x` : Option 'B using `f z x` : 'B or list is empty but in the case of `None` : Option 'B terminate early, returning `z`.

```

(* assume zero = 0, one = 1, negb is in scope and ls = [10,12,9,7]
   given a max and list with elements a_0, a_1, ..., a_m
   find largest n s.t. sum of i from 0 to (n-1) a_i <= max *)
let prefix_step = fun (len_limit : Pair Uint32 Uint32) => fun (x : Uint32) =>
  match len_limit with
  | Pair len limit => let limit_lt_x = builtin lt limit x in
    let x_leq_limit = negb limit_lt_x in
    match x_leq_limit with

```

(continues on next page)

(continued from previous page)

```

| True => let len_succ = builtin add len one in let l_sub_x = builtin sub limit x_
↳in
    let res = Pair {Uint32 Uint32} len_succ l_sub_x in
        Some {(Pair Uint32 Uint32)} res
    | False => None {(Pair Uint32 Uint32)}
    end
end in
let fold_while = @list_foldl_while Uint32 (Pair Uint32 Uint32) in
let max = Uint32 31 in
let init = Pair {Uint32 Uint32} zero max in
let prefix_length = fold_while prefix_step init ls in
match prefix_length with
| Pair length _ => length
end

```

- `list_append` : (List 'A -> List 'A -> List 'A).

Append the first list to the front of the second list, keeping the order of the elements in both lists. Note that `list_append` has linear time complexity in the length of the first argument list.

- `list_reverse` : (List 'A -> List 'A).

Return the reverse of the input list. Note that `list_reverse` has linear time complexity in the length of the argument list.

- `list_flatten` : (List List 'A) -> List 'A.

Construct a list of all the elements in a list of lists. Each element (which has type `List 'A`) of the input list (which has type `List List 'A`) are all concatenated together, keeping the order of the input list. Note that `list_flatten` has linear time complexity in the total number of elements in all of the lists.

- `list_length` : List 'A -> Uint32

Count the number of elements in a list. Note that `list_length` has linear time complexity in the number of elements in the list.

- `list_eq` : ('A -> 'A -> Bool) -> List 'A -> List 'A -> Bool.

Compare two lists element by element, using a predicate function `f` : 'A -> 'A -> Bool. If `f` returns True for every pair of elements, then `list_eq` returns True. If `f` returns False for at least one pair of elements, or if the lists have different lengths, then `list_eq` returns False.

- `list_mem` : ('A -> 'A -> Bool) -> 'A -> List 'A -> Bool.

Checks whether an element `a : 'A` is an element in the list `l : List 'A`. `f : 'A -> 'A -> Bool` should be provided for equality comparison.

```
(* Library *)
let f =
  fun (a : Int32) =>
  fun (b : Int32) =>
    builtin eq a b

(* Contract transition *)
(* Assume input is the list [ 1 ; 2 ; 3 ; 4 ] *)
keynumber = Int32 5;
list_mem_int32 = @list_mem Int32;
check_result = list_mem_int32 f keynumber input;
(* check_result is now False *)
```

- `list_forall : ('A -> Bool) -> List 'A -> Bool`.

Check whether all elements of list `l : List 'A` satisfy the predicate `f : 'A -> Bool`. `list_forall` returns `True` if all elements satisfy `f`, and `False` if at least one element does not satisfy `f`.

- `list_exists : ('A -> Bool) -> List 'A -> Bool`.

Check whether at least one element of list `l : List 'A` satisfies the predicate `f : 'A -> Bool`. `list_exists` returns `True` if at least one element satisfies `f`, and `False` if none of the elements satisfy `f`.

- `list_sort : ('A -> 'A -> Bool) -> List 'A -> List 'A`.

Sort the input list `l : List 'A` using insertion sort. The comparison function `flt : 'A -> 'A -> Bool` provided must return `True` if its first argument is less than its second argument. `list_sort` has quadratic time complexity.

```
let int_sort = @list_sort Uint64 in

let flt =
  fun (a : Uint64) =>
  fun (b : Uint64) =>
    builtin lt a b

let zero = Uint64 0 in
let one = Uint64 1 in
let two = Uint64 2 in
let three = Uint64 3 in
let four = Uint64 4 in

(* l6 = [ 3 ; 2 ; 1 ; 2 ; 3 ; 4 ; 2 ] *)
let l6 =
  let nil = Nil {Uint64} in
```

(continues on next page)

(continued from previous page)

```

let 10 = Cons {Uint64} two nil in
let 11 = Cons {Uint64} four 10 in
let 12 = Cons {Uint64} three 11 in
let 13 = Cons {Uint64} two 12 in
let 14 = Cons {Uint64} one 13 in
let 15 = Cons {Uint64} two 14 in
Cons {Uint64} three 15

(* res1 = [ 1 ; 2 ; 2 ; 2 ; 3 ; 3 ; 4 ] *)
let res1 = int_sort flt 16

```

- `list_find` : ('A -> Bool) -> List 'A -> Option 'A.

Return the first element in a list `l` : List 'A satisfying the predicate `f` : 'A -> Bool. If at least one element in the list satisfies the predicate, and the first one of those elements is `x`, then the result is `Some x`. If no element satisfies the predicate, the result is `None`.

- `list_zip` : List 'A -> List 'B -> List (Pair 'A 'B).

Combine two lists element by element, resulting in a list of pairs. If the lists have different lengths, the trailing elements of the longest list are ignored.

- `list_zip_with` : ('A -> 'B -> 'C) -> List 'A -> List 'B -> List 'C ).

Combine two lists element by element using a combining function `f` : 'A -> 'B -> 'C. The result of `list_zip_with` is a list of the results of applying `f` to the elements of the two lists. If the lists have different lengths, the trailing elements of the longest list are ignored.

- `list_unzip` : List (Pair 'A 'B) -> Pair (List 'A) (List 'B).

Split a list of pairs into a pair of lists consisting of the elements of the pairs of the original list.

- `list_nth` : Uint32 -> List 'A -> Option 'A.

Return the element number `n` from a list. If the list has at least `n` elements, and the element number `n` is `x`, `list_nth` returns `Some x`. If the list has fewer than `n` elements, `list_nth` returns `None`.

### 3.5.4 NatUtils

- `nat_prev` : Nat -> Option Nat: Return the Peano number one less than the current one. If the current number is `Zero`, the result is `None`. If the current number is `Succ x`, then the result is `Some x`.

- `nat_fold_while` : `('T -> Nat -> Option 'T) -> 'T -> Nat -> 'T`: Takes arguments `f` : `'T -> Nat -> Option 'T`, `z` : `'T` and `m` : `Nat`. This is `nat_fold` with early termination. Continues recursing so long as `f` returns `Some y` with new accumulator `y`. Once `f` returns `None`, the recursion terminates.
- `is_some_zero` : `Nat -> Bool`: Zero check for Peano numbers.
- `nat_eq` : `Nat -> Nat -> Bool`: Equality check specialised for the `Nat` type.
- `nat_to_int` : `Nat -> Uint32`: Convert a Peano number to its equivalent `Uint32` integer.
- `uintX_to_nat` : `UintX -> Nat`: Convert a `UintX` integer to its equivalent Peano number. The integer must be small enough to fit into a `Uint32`. If it is not, then an overflow error will occur.
- `intX_to_nat` : `IntX -> Nat`: Convert an `IntX` integer to its equivalent Peano number. The integer must be non-negative, and must be small enough to fit into a `Uint32`. If it is not, then an underflow or overflow error will occur.

### 3.5.5 PairUtils

- `fst` : `Pair 'A 'B -> 'A`: Extract the first element of a `Pair`.

```
let fst_strings = @fst String String in
let nick_name = "toby" in
let dog = "dog" in
let tobias = Pair {String String} nick_name dog in
fst_strings tobias
```

- `snd` : `Pair 'A 'B -> 'B`: Extract the second element of a `Pair`.

### 3.5.6 Conversions

This library provides conversions b/w Scilla types, particularly between integers and byte strings.

To enable specifying the encoding of integer arguments to these functions, a type is defined for endianness.

```
type IntegerEncoding =
| LittleEndian
| BigEndian
```

The functions below, along with their primary result, also return `next_pos` : `Uint32` which indicates the position from which any further data extraction from the input `ByStr` value can proceed. This is useful when deserializing a byte stream. In other words, `next_pos` indicates where this function stopped reading bytes from the input byte string.

- `substr_safe` : `ByStr -> Uint32 -> Uint32 -> Option ByStr` While Scilla provides a builtin to extract substrings of byte strings (`ByStr`), it is not exception safe. When provided incorrect arguments, it throws exceptions. This library function is provided as an exception safe function to extract, from a string `s` : `ByStr`, a substring starting at position `pos` : `Uint32` and of length `len` : `Uint32`. It returns `Some ByStr` on success and `None` on failure.
- `extract_uint32` : `IntegerEncoding -> ByStr -> Uint32 -> Option (Pair Uint32 Uint32)` Extracts a `Uint32` value from `bs` : `ByStr`, starting at position `pos` : `Uint32`. On success, `Some extracted_uint32_value next_pos` is returned. `None` otherwise.
- `extract_uint64` : `IntegerEncoding -> ByStr -> Uint32 -> Option (Pair Uint64 Uint32)` Extracts a `Uint64` value from `bs` : `ByStr`, starting at position `pos` : `Uint32`. On success, `Some extracted_uint64_value next_pos` is returned. `None` otherwise.

- `extract_uint128` : `IntegerEncoding -> ByStr -> Uint32 -> Option (Pair Uint128 Uint32)` Extracts a `Uint128` value from `bs` : `ByStr`, starting at position `pos` : `Uint32`. On success, `Some extracted_uint128_value next_pos` is returned. None otherwise.
- `extract_uint256` : `IntegerEncoding -> ByStr -> Uint32 -> Option (Pair Uint256 Uint32)` Extracts a `Uint256` value from `bs` : `ByStr`, starting at position `pos` : `Uint32`. On success, `Some extracted_uint256_value next_pos` is returned. None otherwise.
- `extract_bystr1` : `ByStr -> Uint32 -> Option (Pair ByStr1 Uint32)` Extracts a `ByStr1` value from `bs` : `ByStr`, starting at position `pos` : `Uint32`. On success, `Some extracted_bystr1_value next_pos` is returned. None otherwise.
- `extract_bystr2` : `ByStr -> Uint32 -> Option (Pair ByStr2 Uint32)` Extracts a `ByStr2` value from `bs` : `ByStr`, starting at position `pos` : `Uint32`. On success, `Some extracted_bystr2_value next_pos` is returned. None otherwise.
- `extract_bystr20` : `ByStr -> Uint32 -> Option (Pair ByStr20 Uint32)` Extracts a `ByStr2` value from `bs` : `ByStr`, starting at position `pos` : `Uint32`. On success, `Some extracted_bystr20_value next_pos` is returned. None otherwise.
- `extract_bystr32` : `ByStr -> Uint32 -> Option (Pair ByStr32 Uint32)` Extracts a `ByStr2` value from `bs` : `ByStr`, starting at position `pos` : `Uint32`. On success, `Some extracted_bystr32_value next_pos` is returned. None otherwise.
- `append_uint32` : `IntegerEncoding -> ByStr -> Uint32 -> ByStr` Serialize a `Uint32` value (with the specified encoding) and append it to the provided `ByStr` and return the result `ByStr`.
- `append_uint64` : `IntegerEncoding -> ByStr -> Uint32 -> ByStr` Serialize a `Uint64` value (with the specified encoding) and append it to the provided `ByStr` and return the result `ByStr`.
- `append_uint128` : `IntegerEncoding -> ByStr -> Uint32 -> ByStr` Serialize a `Uint128` value (with the specified encoding) and append it to the provided `ByStr` and return the result `ByStr`.
- `append_uint256` : `IntegerEncoding -> ByStr -> Uint32 -> ByStr` Serialize a `Uint256` value (with the specified encoding) and append it to the provided `ByStr` and return the result `ByStr`.

### 3.5.7 Polynetwork Support Library

This library provides utility functions used in building the Zilliqa Polynetwork bridge. These functions are migrated from Polynetwork's [ethereum support](#), with the contract itself separately deployed.

## 3.6 Scilla Tips and Tricks

### 3.6.1 Performance

#### Field map size

If your contract needs to know the size of a field map, i.e. a field variable that is a map then the obvious implementation that reads a field map into a variable and applies the `size` builtin (as in the following code snippet) can be very inefficient as it requires making a copy of the map in question.



```

field accounts : Map ByStr20 Uint128 = Emp ByStr20 Uint128

transition Foo()
  ...
  accounts_copy <- accounts;
  accounts_size = builtin size accounts_copy;
  ...
end

```

In order to solve the issue one tracks the size information in a corresponding field variable as in the following code snippet. Notice that now instead of copying a map one just reads from `accounts_size` field variable.

```

let uint32_one = Uint32 1

field accounts : Map ByStr20 Uint128 = Emp ByStr20 Uint128
field accounts_size : Uint32 = 0

transition Foo()
  ...
  num_of_accounts <- accounts_size;
  ...
end

```

Now, to make sure that the map and its size stay in sync, one needs to update the size of the map when using the builtin in-place statements like `m[k] := v` or `delete m[k]` or better yet define and use systematically procedures that do exactly that.

Here is the definition of a procedure that updates a key/value pair in the `accounts` map and changes its size accordingly.

```

procedure insert_to_accounts (key : ByStr20, value : Uint128)
  already_exists <- exists accounts[key];
  match already_exists with
  | True =>
    (* do nothing as the size does not change *)
  | False =>
    size <- accounts_size;
    new_size = builtin add size uint32_one;
    accounts_size := new_size
  end;
  accounts[key] := value
end

```

And this is the definition of a procedure that removes a key/value pair from the `accounts` map and changes its size accordingly.

```

procedure delete_from_accounts (key : ByStr20)
  already_exists <- exists accounts[key];
  match already_exists with
  | False =>
    (* do nothing as the map and its size do not change *)
  | True =>
    size <- accounts_size;
    new_size = builtin sub size uint32_one;
    accounts_size := new_size
  end;
  delete accounts[key]
end

```

## 3.6.2 Money Idioms

### Partially accepting funds

Let's say you are working on a contract which lets people tip each other. Naturally, you'd like to avoid a situation when a person tips too much because of a typo. It would be nice to ask Scilla to accept incoming funds partially, but there is no `accept <cap>` builtin. You can either not accept at all or accept the funds fully. We can work around this restriction by fully accepting the incoming funds and then immediately refunding the tipper if the tip exceeds some cap.

It turns out we can encapsulate this kind of behavior as a reusable procedure.

```
procedure accept_with_cap (cap : Uint128)
  sent_more_than_necessary = builtin lt cap _amount;
  match sent_more_than_necessary with
  | True =>
    amount_to_refund = builtin sub _amount cap;
    accept;
    msg = { _tag : ""; _recipient: _sender; _amount: amount_to_refund };
    msgs = one_msg msg;
    send msgs
  | False =>
    accept
  end
end
```

Now, the `accept_with_cap` procedure can be used as follows.

```
<contract library and procedures here>

contract Tips (tip_cap : Uint128)

transition Tip (message_from_tipper : String)
  accept_with_cap tip_cap;
  e = { _eventname: "ThanksForTheTip" };
  event e
end
```

## 3.6.3 Safety

### Transfer Contract Ownership

If your contract has an owner (usually it means someone with admin powers like adding/removing user accounts or pausing/unpausing the contract) which can be changed at runtime then at first sight this can be formalized in the code as a field `owner` and a transition like `ChangeOwner`:

```
contract Foo (initial_owner : ByStr20)

field owner : ByStr20 = initial_owner

transition ChangeOwner (new_owner : ByStr20)
  (* continue executing the transition if _sender is the owner,
   * throw exception and abort otherwise *)
  isOwner;
  owner := new_owner
end
```

However, this might lead to a situation when the current owner gets locked out of the control over the contract. For instance, the current owner can potentially transfer contract ownership to a non-existing address due to a typo in the address parameter when calling the `ChangeOwner` transition. Here we provide a design pattern to circumvent this issue.

One way to ensure the new owner is active is to do ownership transfer in two phases:

- the current owner offers to transfer ownership to a new owner, note that at this point the current owner is still the contract owner;
- the future new owner accepts the pending ownership transfer and becomes the current owner;
- at any moment before the future new owner accepts the transfer the current owner can abort the ownership transfer.

Here is a possible implementation of the pattern outlined above (the code for the `isOwner` procedure is not shown):

```
contract OwnershipTransfer (initial_owner : ByStr20)

field owner : ByStr20 = initial_owner
field pending_owner : Option ByStr20 = None {ByStr20}

transition RequestOwnershipTransfer (new_owner : ByStr20)
  isOwner;
  po = Some {ByStr20} new_owner;
  pending_owner := po
end

transition ConfirmOwnershipTransfer ()
  optional_po <- pending_owner;
  match optional_po with
  | Some pend_owner =>
    caller_is_new_owner = builtin eq _sender pend_owner;
    match caller_is_new_owner with
    | True =>
      (* transfer ownership *)
      owner := pend_owner;
      none = None {ByStr20};
      pending_owner := none
    | False => (* the caller is not the new owner, do nothing *)
    end
  | None => (* ownership transfer is not in-progress, do nothing *)
  end
end
```

Ownership transfer for contracts with the above transition is supposed to happen as follows:

- the current owner calls `RequestOwnershipTransfer` transition with the new owner address as its only explicit parameter;
- the new owner calls `ConfirmOwnershipTransfer`.

Until the ownership transition is finalized, the current owner can abort it by calling `RequestOwnershipTransfer` with their own address. A (redundant) dedicated transition can be added to make it even more evident to the contract owners and users.

## 3.7 The Scilla checker

The Scilla checker (`scilla-checker`) works as a compiler frontend, parsing the contract and performing a number of static checks including typechecking.

### 3.7.1 Phases of the Scilla checker

The Scilla checker operates in distinct phases, each of which can perform checks (and potentially reject contracts that do not pass the checks) and add annotations to each piece of syntax:

- *Lexing and parsing* reads the contract code and builds an abstract syntax tree (AST). Each node in the tree is annotated with a location from the source file in the form of line and column numbers.
- *ADT checking* checks various constraints on user-defined ADTs.
- *Typechecking* checks that values in the contract are used in a way that is consistent with the type system. The typechecker also annotates each expression with its type.
- *Pattern-checking* checks that each pattern-match in the contract is exhaustive (so that execution will not fail due to no match being found), and that each pattern can be reached (so that the programmer does not inadvertently introduce a pattern branch that can never be reached).
- *Event-info* checks that messages and events in the contract contain all necessary fields. For events, if a contract emits two events with the same name (`_eventname`), then their structure (the fields and their types) must also be the same.
- *Cashflow analysis* analyzes the usage of variables, fields and ADTs, and attempts to determine which fields are used to represent (native) blockchain money. No checks are performed, but expressions, variables, fields and ADTs are annotated with tags indicating their usage.
- *Payment acceptance checking* checks contracts for payment acceptances. It raises a warning if a contract has no transitions which accept payment. Also, the check raises a warning if a transition has any code path with potentially multiple `accept` statements in it. This check does not raise an error since it is not possible via static analysis to know for sure in all cases whether multiple `accept` statements would be reached if present, since this can be dependent on conditions which are only known at run-time. The Scilla checker only performs this check if only `-cf` flag is specified on the command line, i.e. it is performed together with the cashflow analysis. For instance, fungible token contracts don't generally need `accept` statements, hence this check is not mandatory.
- *Sanity-checking* performs a number of minor checks, e.g., that all parameters to a transition or a procedure have distinct names.

### Annotations

Each phase in the Scilla checker can add an annotation to each node in the abstract syntax tree. The type of an annotation is specified through instantiations of the module signature `Rep`. `Rep` specifies the type `rep`, which is the type of the annotation:

```
module type Rep = sig
  type rep
  ...
end
```

In addition to the type of the annotation, the instantiation of `Rep` can declare helper functions that allow subsequent phases to access the annotations of previous phases. Some of these functions are declared in the `Rep` signature,

because they involve creating new abstract syntax nodes, which must be created with annotations from the parser onward:

```

module type Rep = sig
  ...

  val mk_id_address : string -> rep ident
  val mk_id_uint128 : string -> rep ident
  val mk_id_bnum    : string -> rep ident
  val mk_id_string  : string -> rep ident

  val rep_of_sexp : Sexp.t -> rep
  val sexp_of_rep : rep -> Sexp.t

  val parse_rep : string -> rep
  val get_rep_str: rep -> string
end

```

`mk_id_<type>` creates an identifier with an appropriate type annotation if annotation is one of the phases that has been executed. If the typechecker has not yet been executed, the functions simply create an (untyped) identifier with a dummy location.

`rep_of_sexp` and `sexp_of_rep` are used for pretty-printing. They are automatically generated if `rep` is defined with the `[@@deriving sexp]` directive.

`parse_rep` and `get_rep_str` are used for caching of typechecked libraries, so that they do not need to be checked again if they haven't changed. These will likely be removed in future versions of the Scilla checker.

As an example, consider the annotation module `TypecheckerERep`:

```

module TypecheckerERep (R : Rep) = struct
  type rep = PlainTypes.t inferred_type * R.rep
  [@@deriving sexp]

  let get_loc r = match r with | (_, rr) -> R.get_loc rr

  let mk_id s t =
    match s with
    | Ident (n, r) -> Ident (n, (PlainTypes.mk_qualified_type t, r))

  let mk_id_address s = mk_id (R.mk_id_address s) (bystrx_typ address_length)
  let mk_id_uint128 s = mk_id (R.mk_id_uint128 s) uint128_typ
  let mk_id_bnum    s = mk_id (R.mk_id_bnum s) bnum_typ
  let mk_id_string  s = mk_id (R.mk_id_string s) string_typ

  let mk_rep (r : R.rep) (t : PlainTypes.t inferred_type) = (t, r)

  let parse_rep s = (PlainTypes.mk_qualified_type uint128_typ, R.parse_rep s)
  let get_rep_str r = match r with | (_, rr) -> R.get_rep_str rr

  let get_type (r : rep) = fst r
end

```

The functor (parameterized structure) takes the annotation from the previous phase as the parameter `R`. In the Scilla checker this previous phase is the parser, but any phase could be added in-between the two phases by specifying the phase in the top-level runner.

The type `rep` specifies that the new annotation is a pair of a type and the previous annotation.

The function `get_loc` merely serves as a proxy for the `get_loc` function of the previous phase.

The function `mk_id` is a helper function for the `mk_id_<type>` functions, which create an identifier with the appropriate type annotation.

The `mk_rep` function is a helper function used by the typechecker.

Prettyprinting does not output the types of AST nodes, so the functions `parse_rep` and `get_rep_str` ignore the type annotations.

Finally, the function `get_type` provides access to type information for subsequent phases. This function is not mentioned in the `Rep` signature, since it is made available by the typechecker once type annotations have been added to the AST.

## Abstract syntax

The `ScillaSyntax` functor defines the AST node types. Each phase will instantiate the functor twice, once for the input syntax and once for the output syntax. These two syntax instantiations differ only in the type of annotations of each syntax node. If the phase produces no additional annotations, the two instantiations will be identical.

The parameters `SR` and `ER`, both of type `Rep`, define the annotations for statements and expressions, respectively.

```
module ScillaSyntax (SR : Rep) (ER : Rep) = struct

  type expr_annot = expr * ER.rep
  and expr = ...

  type stmt_annot = stmt * SR.rep
  and stmt = ...

end
```

## Initial annotation

The parser generates the initial annotation, which only contains information about where the syntax node is located in the source file. The function `get_loc` allows subsequent phases to access the location.

The `ParserRep` structure is used for annotations both of statements and expressions.

```
module ParserRep = struct
  type rep = loc
  [@@deriving sexp]

  let get_loc l = l
  ...
end
```

## Typical phase

Each phase that produces additional annotations will need to provide a new implementation of the `Rep` module type. The implementation should take the previous annotation type (as a structure implementing the `Rep` module type) as a parameter, so that the phase's annotations can be added to the annotations of the previous phases.

The typechecker adds a type to each expression node in the AST, but doesn't add anything to statement node annotations. Consequently, the typechecker only defines an annotation type for expressions.

In addition, the `Rep` implementation defines a function `get_type`, so that subsequent phases can access the type in the annotation.

```

module TypecheckerERep (R : Rep) = struct
  type rep = PlainTypes.t inferred_type * R.rep
  [@@deriving sexp]

  let get_loc r = match r with | (_, rr) -> R.get_loc rr

  ...
  let get_type (r : rep) = fst r
end

```

The Scilla typechecker takes the statement and expression annotations of the previous phase, and then instantiates `TypeCheckerERep` (creating the new annotation type), `ScillaSyntax` (creating the abstract syntax type for the previous phase, which serves as input to the typechecker), and `ScillaSyntax` again (creating the abstract syntax type that the typechecker outputs).

```

module ScillaTypechecker
  (SR : Rep)
  (ER : Rep) = struct

  (* No annotation added to statements *)
  module STR = SR
  (* TypecheckerERep is the new annotation for expressions *)
  module ETR = TypecheckerERep (ER)

  (* Instantiate ScillaSyntax with source and target annotations *)
  module UntypedSyntax = ScillaSyntax (SR) (ER)
  module TypedSyntax = ScillaSyntax (STR) (ETR)

  (* Expose target syntax and annotations for subsequent phases *)
  include TypedSyntax
  include ETR

  (* Instantiate helper functors *)
  module TU = TypeUtilities (SR) (ER)
  module TBuiltins = ScillaBuiltIns (SR) (ER)
  module TypeEnv = TU.MakeTEnv(PlainTypes) (ER)
  module CU = ScillaContractUtil (SR) (ER)
  ...
end

```

Crucially, the typechecker module exposes the annotations and the syntax type that it generates, so that they can be made available to the next phase.

The typechecker finally instantiates helper functors such as `TypeUtilities` and `ScillaBuiltIns`.

### 3.7.2 Cashflow Analysis

The cashflow analysis phase analyzes the usage of a contract's variables, fields, and ADT constructor, and attempts to determine which fields and ADTs are used to represent (native) blockchain money. Each contract field is annotated with a tag indicating the field's usage.

The resulting tags are an approximation based on the usage of the contract's fields, variables, and ADT constructors. The tags are not guaranteed to be accurate, but are intended as a tool to help the contract developer use her fields in the intended manner.

## Running the analysis

The cashflow analysis is activated by running `scilla-checker` with the option `-cf`. The analysis is not run by default, since it is only intended to be used during contract development.

A contract is never rejected due to the result of the cashflow analysis. It is up to the contract developer to determine whether the cashflow tags are consistent with the intended use of each contract field.

## The Analysis in Detail

The analysis works by continually analysing the transitions and procedures of the contract until no further information is gathered.

The starting point for the analysis is the incoming message that invokes the contract's transition, the outgoing messages and events that may be sent by the contract, the contract's account balance, and any field being read from the blockchain such as the current blocknumber.

Both incoming and outgoing messages contain a field `_amount` whose value is the amount of money being transferred between accounts by the message. Whenever the value of the `_amount` field of the incoming message is loaded into a local variable, that local variable is tagged as representing money. Similarly, a local variable used to initialise the `_amount` field of an outgoing message is also tagged as representing money.

Conversely, the message fields `_sender`, `_origin`, `_recipient`, and `_tag`, the event field `_eventname`, the exception field `_exception`, and the blockchain field `BLOCKNUMBER`, are known to not represent money, so any variable used to initialise those fields or to hold the value read from one of those fields is tagged as not representing money.

Once some variables have been tagged, their usage implies how other variables can be tagged. For instance, if two variables tagged as money are added to each other, the result is also deemed to represent money. Conversely, if two variables tagged as non-money are added, the result is deemed to represent non-money.

Tagging of contract fields happens when a local variable is used when loading or storing a contract field. In these cases, the field is deemed to have the same tag as the local variable.

Tagging of custom ADTs is done when they are used for constructing values, and when they are used in pattern-matching.

Once a transition or procedure has been analyzed, the local variables and their tags are saved and the analysis proceeds to the next transition or procedure while keeping the tags of the contract fields and ADTs. The analysis continues until all the transitions and procedures have been analysed without any existing tags having changed.

## Tags

The analysis uses the following set of tags:

- *No information*: No information has been gathered about the variable. This sometimes (but not always) indicates that the variable is not being used, indicating a potential bug.
- *Money*: The variable represents money.
- *Not money*: The variable represents something other than money.
- *Map  $t$*  (where  $t$  is a tag): The variable represents a map or a function whose co-domain is tagged with  $t$ . Hence, when performing a lookup in the map, or when applying a function on the values stored in the map, the result is tagged with  $t$ . Keys of maps are assumed to always be *Not money*. Using a variable as a function parameter does not give rise to a tag.



- *T t1 ... tn* (where *T* is an ADT, and *t1 ... tn* are tags): The variable represents a value of an ADT, such as *List* or *Option*. The tags *t1 ... tn* correspond to the tags of each type parameter of the ADT. (See the simple *example* further down.)
- *Inconsistent*: The variable has been used to represent both money and not money. Inconsistent usage indicates a bug.

Library and local functions are only partially supported, since no attempt is made to connect the tags of parameters to the tag of the result. Built-in functions are fully supported, however.

## A simple example

Consider the following code snippet:

```
match p with
| Nil =>
| Cons x xs =>
  msg = { _amount : x ; ... }
  ...
end
```

*x* is used to initialise the `_amount` field of a message, so *x* gets tagged with *Money*. Since *xs* is the tail of a list of which *x* is the first element, *xs* must be a list of elements with the same tag as *x*. *xs* therefore gets tagged with *List Money*, corresponding to the fact that the `List 'A` type has one type parameter.

Similarly, *p* is matched against the patterns `Nil` and `Cons x xs`. `Nil` is a list, but since the list is empty we don't know anything about the contents of the list, and so the `Nil` pattern corresponds to the tag *List (No information)*. `Cons x xs` is also a list, but this time we do know something about the contents, namely that the first element *x* is tagged with *Money*, and the tail of the list is tagged with *List Money*. Consequently, `Cons x xs` corresponds to *List Money*.

Unifying the two tags *List (No information)* and *List Money* gives the tag *List Money*, so *p* gets tagged with *List Money*.

## ADT constructor tagging

In addition to tagging fields and local variables, the cashflow analyser also tags constructors of custom ADTs.

To see how this works, consider the following custom ADT:

```
type Transaction =
| UserTransaction of ByStr20 Uint128
| ContractTransaction of ByStr20 String Uint128
```

A user transaction is a transaction where the recipient is a user account, so the `UserTransaction` constructor takes two arguments: An address of the recipient user account, and the amount to transfer.

A contract transaction is a transaction where the recipient is another contract, so the `ContractTransaction` takes three arguments: An address of the recipient contract, the name of the transition to invoke on the recipient contract, and the amount to transfer.

In terms of cashflow it is clear that the last argument of both constructors is used to represent an amount of money, whereas all other arguments are used to represent non-money. The cashflow analyser therefore attempts to tag the arguments of the two constructors with appropriate tags, using the principles described in the previous sections.

## A more elaborate example

As an example, consider a crowdfunding contract written in Scilla. Such a contract may declare the following immutable parameters and mutable fields:

```
contract Crowdfunding

(* Parameters *)
(owner      : ByStr20,
max_block  : BNum,
goal       : Uint128)

(* Mutable fields *)
field backers : Map ByStr20 Uint128 = ...
field funded  : Bool = ...
```

The `owner` parameter represents the address of the person deploying the contract. The `goal` parameter is the amount of money the owner is trying to raise, and the `max_block` parameter represents the deadline by which the goal is to be met.

The field `backers` is a map from the addresses of contributors to the amount of money contributed, and the field `funded` represents whether the goal has been reached.

Since the field `goal` represents an amount of money, `goal` should be tagged as *Money* by the analysis. Similarly, the `backers` field is a map with a co-domain representing *Money*, so `backers` should be tagged with *Map Money*.

Conversely, both `owner`, `max_block` and `funded` represent something other than money, so they should all be tagged with *Not money*.

The cashflow analysis will tag the parameters and fields according to how they are used in the contract's transitions and procedures, and if the resulting tags do not correspond to the expectation, then the contract likely contains a bug somewhere.

## 3.8 Interpreter Interface

The Scilla interpreter provides a calling interface that enables users to invoke transitions with specified inputs and obtain outputs. Execution of a transition with supplied inputs will result in a set of outputs, and a change in the smart contract mutable state.

### 3.8.1 Calling Interface

A transition defined in a contract can be called either by the issuance of a transaction, or by message calls from another contract. The same calling interface will be used to call the contract via external transactions and inter-contract message calls.

The inputs to the interpreter (`scilla-runner`) consists of four input JSON files as described below. Every invocation of the interpreter to execute a transition must be provided with these four JSON inputs:

```
./scilla-runner -init init.json -istate input_state.json -iblockchain input_
↪blockchain.json -imessage input_message.json -o output.json -i input.scilla
```

The interpreter executable can be run either to create a contract (denoted `CreateContract`) or to invoke a transition in a contract (`InvokeContract`). Depending on which of these two, some of the arguments will be absent. The table below outlays the arguments that should be present in each of these two cases. A `CreateContract` is distinguished from an `InvokeContract`, based on the presence of `input_message.json` and `input_state.json`. If

these arguments are absent, then the interpreter will evaluate it as a `CreateContract`. Else, it will treat it as an `InvokeContract`. Note that for `CreateContract`, the interpreter only performs basic checks such as matching the contract's immutable parameters with `init.json` and whether the contract definition is free of syntax errors.

Input	Description	Present	
		CreateContract	InvokeContract
<code>init.json</code>	Immutable contract parameters	Yes	Yes
<code>input_state.json</code>	Mutable contract state	No	Yes
<code>input_blockchain.json</code>	Blockchain state	Yes	Yes
<code>input_message.json</code>	Transition and parameters	No	Yes
<code>output.json</code>	Output	Yes	Yes
<code>input.scilla</code>	Input contract	Yes	Yes

In addition to the command line arguments provided above, the interpreter also expects a mandatory `-gaslimit X` argument (where `X` is a positive integer value). If the contract or library module imports other libraries (including the standard library), a `-libdir` option must be provided, with a list of directories (in the standard `PATH` format) as the argument, indicating directories to be searched for for finding the libraries.

### 3.8.2 Initializing the Immutable State

`init.json` defines the values of the immutable parameters of a contract. It does not change between invocations. The JSON is an array of objects, each of which contains the following fields:

Field	Description
<code>vname</code>	Name of the immutable contract parameter
<code>type</code>	Type of the immutable contract parameter
<code>value</code>	Value of the immutable contract parameter

`init.json` must specify `_scilla_version` with type `Uint32`, specifying a value that is the same as specified in the contract's source. Additionally, the blockchain will provide two implicit contract parameters `_this_address`, a `ByStr20` value denoting the address of the contract itself, and `_creation_block`, a `BNum` value denoting the block in which the contract is / was created. While working with the offline interpreter, you may need to provide these values in the `init.json` yourself.

#### Example 1

For the `HelloWorld.scilla` contract fragment given below, we have only one immutable parameter `owner`.

```
contract HelloWorld
(* Immutable parameters *)
(owner: ByStr20)
```

A sample `init.json` for this contract will look like the following:

```
[
  {
    "vname" : "_scilla_version",
    "type"  : "Uint32",
    "value" : "0"
  },
  {
    "vname" : "owner",
```

(continues on next page)

(continued from previous page)

```

    "type" : "ByStr20",
    "value" : "0x1234567890123456789012345678901234567890"
  },
  {
    "vname" : "_this_address",
    "type" : "ByStr20",
    "value" : "0xabfeccdc9012345678901234567890f777567890"
  },
  {
    "vname" : "_creation_block",
    "type" : "BNum",
    "value" : "1"
  }
]

```

## Example 2

For the `Crowdfunding.scilla` contract fragment given below, we have three immutable parameters `owner`, `max_block` and `goal`.

```

contract Crowdfunding
  (* Immutable parameters *)
  (owner      : ByStr20,
   max_block  : BNum,
   goal       : UInt128)

```

A sample `init.json` for this contract will look like the following:

```

[
  {
    "vname" : "_scilla_version",
    "type"  : "UInt32",
    "value" : "0"
  },
  {
    "vname" : "owner",
    "type"  : "ByStr20",
    "value" : "0x1234567890123456789012345678901234567890"
  },
  {
    "vname" : "max_block",
    "type"  : "BNum",
    "value" : "199"
  },
  {
    "vname" : "_this_address",
    "type"  : "ByStr20",
    "value" : "0xabfeccdc9012345678901234567890f777567890"
  },
  {
    "vname" : "goal",
    "type"  : "UInt128",
    "value" : "5000000000000000"
  },
  {

```

(continues on next page)

(continued from previous page)

```

    "vname" : "_creation_block",
    "type" : "BNum",
    "value" : "1"
  }
]

```

### Example 3: Using Address Types

Whenever a contract has an immutable parameter of an address type, the type `ByStr20` must be used in the to initialise the parameter.

For the `SimpleExchange` we have a single the immutable parameter, which has an address type:

```

contract SimpleExchange
(
  initial_admin : ByStr20 with end
)

```

The JSON entry for the `initial_admin` parameter must use the type `ByStr20` rather than the type `ByStr20 with end`, so an example `init.json` for this contract could like the following:

```

[
  {
    "vname" : "_scilla_version",
    "type" : "Uint32",
    "value" : "0"
  },
  {
    "vname" : "_this_address",
    "type" : "ByStr20",
    "value" : "0xabfecdc9012345678901234567890f777567890"
  },
  {
    "vname" : "_creation_block",
    "type" : "BNum",
    "value" : "1"
  },
  {
    "vname" : "initial_admin",
    "type" : "ByStr20",
    "value" : "0x1234567890123456789012345678901234567890"
  }
]

```

### 3.8.3 Input Blockchain State

`input_blockchain.json` feeds the current blockchain state to the interpreter. It is similar to `init.json`, except that it is a fixed size array of objects, where each object has `vname` fields only from a predetermined set (which correspond to actual blockchain state variables).

**Permitted JSON fields:** At the moment, the only blockchain value that is exposed to contracts is the current `BLOCKNUMBER`.

```
[
  {
    "vname" : "BLOCKNUMBER",
    "type"  : "BNum",
    "value" : "3265"
  }
]
```

### 3.8.4 Input Message

`input_message.json` contains the information required to invoke a transition. The json is an array containing the following four objects:

Field	Description
<code>_tag</code>	Transition to be invoked
<code>_amount</code>	Number of QA to be transferred
<code>_sender</code>	Address of the invoker (in a chain call, this is the immediate caller)
<code>_origin</code>	Address from which the transaction originated
<code>params</code>	An array of parameter objects

All the four fields are mandatory. `params` can be empty if the transition takes no parameters.

The `params` array is encoded similar to how `init.json` is encoded, with each parameter specifying the (`vname`, `type`, `value`) that has to be passed to the transition that is being invoked.

#### Example 1

For the following transition:

```
transition SayHello()
```

an example `input_message.json` is given below:

```
{
  "_tag"      : "SayHello",
  "_amount"   : "0",
  "_sender"   : "0x1234567890123456789012345678901234567890",
  "_origin"   : "0x1234567890123456789012345678901234567890",
  "params"    : []
}
```

#### Example 2

For the following transition:

```
transition TransferFrom (from : ByStr20, to : ByStr20, tokens : Uint128)
```

an example `input_message.json` is given below:

```

{
  "_tag"      : "TransferFrom",
  "_amount"  : "0",
  "_sender"  : "0x64345678901234567890123456789012345678cd",
  "_origin"  : "0x64345678901234567890123456789012345678cd",
  "params"   : [
    {
      "vname" : "from",
      "type"  : "ByStr20",
      "value" : "0x1234567890123456789012345678901234567890"
    },
    {
      "vname" : "to",
      "type"  : "ByStr20",
      "value" : "0x78345678901234567890123456789012345678cd"
    },
    {
      "vname" : "tokens",
      "type"  : "Uint128",
      "value" : "5000000000000000"
    }
  ]
}

```

### Example 3: Using user-defined types

---

**Note:** Due to a bug in the Scilla implementation the information in this section is only valid from Scilla version 0.10.0 and forwards. Contracts written in Scilla versions prior to 0.10.0 and which exploit this bug will have to be rewritten and redeployed, as they will no longer work from version 0.10.0 and onwards.

---

When passing a value of user-defined type through the interpreter interface, the json structure is identical to the one described in the previous example. However, in the interpreter interface all types must be fully qualified, which is defined as follows:

- For a user-defined type `T` defined in a module deployed at address `A`, the fully qualified name is `A.T`.
- For a user-defined constructor `C` defined in a module deployed at address `A`, the fully qualified name is `A.C`.

---

**Note:** For the purposes of offline development the address of a module is defined as the module's filename, without file extension. That is, if a contract defines a type `T` with a constructor `C` in a file `F.scilla`, then the fully qualified name of the type is `F.T`, and the fully qualified name of the constructor is `F.C`.

---

As an example consider a contract that implements a simple board game. The contract might define a type `Direction` and a transition `MoveAction` as follows:

```

type Direction =
| East
| South
| West
| North
...

```

(continues on next page)

(continued from previous page)

```
transition MoveAction (dir : Direction, spaces : Uint32)
  ...
```

Say that the contract has been deployed at address `0x1234567890123456789012345678906784567890`. To invoke the transition with parameters `East` and `2`, use the type name `0x1234567890123456789012345678906784567890.Direction` and the constructor name `0x1234567890123456789012345678906784567890.East` in the message json:

```
{
  "_tag": "MoveAction",
  "_amount": "0",
  "_sender" : "0x64345678901234567890123456789012345678cd",
  "_origin" : "0x64345678901234567890123456789012345678cd",
  "params": [
    {
      "vname" : "dir",
      "type" : "0x1234567890123456789012345678906784567890.Direction",
      "value" :
        {
          "constructor" : "0x1234567890123456789012345678906784567890.East",
          "argtypes" : [],
          "arguments" : []
        }
    },
    {
      "vname" : "spaces",
      "type" : "Uint32",
      "value" : "2"
    }
  ]
}
```

If a contract has immutable fields of user-defined types, then the fields must also be initialised using fully qualified names in the associated `init.json`.

#### Example 4: Using Address Types

When passing an address value the type `ByStr20` must be used. It is not possible to use address types (`ByStr20` with `... end`) in messages.

This means that for the following transition

```
transition ListToken(
  token_code : String,
  new_token : ByStr20 with contract field allowances : Map ByStr20 (Map ByStr20_
↪ Uint128) end
)
```

the `input_message.json` must use the type `ByStr20` for the `new_token` parameter, e.g., as follows:

```
{
  "_tag" : "ListToken",
  "_amount" : "0",
  "_sender" : "0x64345678901234567890123456789012345678cd",
  "_origin" : "0x64345678901234567890123456789012345678cd",
```

(continues on next page)



(continued from previous page)

```

"params" : [
  {
    "vname" : "token_code",
    "type" : "String",
    "value" : "XYZ"
  },
  {
    "vname" : "new_token",
    "type" : "ByStr20",
    "value" : "0x78345678901234567890123456789012345678cd"
  }
]
}

```

### 3.8.5 Interpreter Output

The interpreter will return a JSON object (`output.json`) with the following fields:

Field	Description
<code>scilla_major_version</code>	The major version of the Scilla language of this contract.
<code>gas_remaining</code>	The remaining gas after invoking or deploying a contract.
<code>_accepted</code>	Whether the incoming QA have been accepted (Either "true" or "false")
<code>message</code>	The message to be sent to another contract/non-contract account, if any.
<code>states</code>	An array of objects that form the new contract state
<code>events</code>	An array of events emitted by the transition and the procedures it invoked.

- `message` is a JSON object with a similar format to `input_message.json`, except that it has a `_recipient` field instead of the `_sender` field. The fields in `message` are given below:

Field	Description
<code>_tag</code>	Transition to be invoked
<code>_amount</code>	Number of QA to be transferred
<code>_recipient</code>	Address of the recipient
<code>params</code>	An array of parameter objects to be passed

The `params` array is encoded similar to how `init.json` is encoded, with each parameter specifying the (`vname`, `type`, `value`) that has to be passed to the transition that is being invoked.

- `states` is an array of objects that represents the mutable state of the contract. Each entry of the `states` array also specifies (`vname`, `type`, `value`).
- `events` is an array of objects that represents the events emitted by the transition. The fields in each object in the `events` array are given below:

Field	Description
<code>_eventname</code>	The name of the event
<code>params</code>	An array of additional event fields

The `params` array is encoded similar to how `init.json` is encoded, with each parameter specifying the (`vname`, `type`, `value`) of each event field.

## Example 1

An example of the output generated by `Crowdfunding.scilla` is given below. The example also shows the format for maps in contract states.

```
{
  "scilla_major_version": "0",
  "gas_remaining": "7365",
  "_accepted": "false",
  "message": {
    "_tag": "",
    "_amount": "1000000000000000",
    "_recipient": "0x12345678901234567890123456789012345678ab",
    "params": []
  },
  "states": [
    { "vname": "_balance", "type": "Uint128", "value": "3000000000000000" },
    {
      "vname": "backers",
      "type": "Map (ByStr20) (Uint128)",
      "value": [
        { "key": "0x12345678901234567890123456789012345678cd", "val": "2000000000000000"
↵" },
        { "key": "0x123456789012345678901234567890123456abcd", "val": "1000000000000000"
↵" }
      ]
    },
    {
      "vname": "funded",
      "type": "Bool",
      "value": { "constructor": "False", "argtypes": [], "arguments": [] }
    }
  ],
  "events": [
    {
      "_eventname": "ClaimBackSuccess",
      "params": [
        {
          "vname": "caller",
          "type": "ByStr20",
          "value": "0x12345678901234567890123456789012345678ab"
        },
        { "vname": "amount", "type": "Uint128", "value": "1000000000000000" },
        { "vname": "code", "type": "Int32", "value": "9" }
      ]
    }
  ]
}
```

## Example 2

For values of an ADT type, the `value` field contains three subfields:

- `constructor`: The name of the constructor used to construct the value.
- `argtypes`: An array of type instantiations. For the `List` and `Option` types, this array will contain one type, indicating the type of the list elements or the optional value, respectively. For the `Pair` type, the array will contain two types, indicating the types of the two values in the pair. For all other ADTs, the array will be empty.

- arguments: The arguments to the constructor.

The following example shows how values of the `List` and `Option` types are represented in the output json:

```
{
  "scilla_major_version": "0",
  "gas_remaining": "7733",
  "_accepted": "false",
  "message": null,
  "states": [
    { "vname": "_balance", "type": "Uint128", "value": "0" },
    {
      "vname": "gpair",
      "type": "Pair (List (Int64)) (Option (Bool))",
      "value": {
        "constructor": "Pair",
        "argtypes": [ "List (Int64)", "Option (Bool)" ],
        "arguments": [
          [],
          { "constructor": "None", "argtypes": [ "Bool" ], "arguments": [] }
        ]
      }
    },
    { "vname": "l1list", "type": "List (List (Int64))", "value": [] },
    { "vname": "plist", "type": "List (Option (Int32))", "value": [] },
    {
      "vname": "gnat",
      "type": "Nat",
      "value": { "constructor": "Zero", "argtypes": [], "arguments": [] }
    },
    {
      "vname": "gmap",
      "type": "Map (ByStr20) (Pair (Int32) (Int32))",
      "value": [
        {
          "key": "0x12345678901234567890123456789012345678ab",
          "val": {
            "constructor": "Pair",
            "argtypes": [ "Int32", "Int32" ],
            "arguments": [ "1", "2" ]
          }
        }
      ]
    }
  ],
  "events": []
}
```

### 3.8.6 Input Mutable Contract State

`input_state.json` contains the current value of mutable state variables. It has the same forms as the `states` field in `output.json`. An example of `input_state.json` for `Crowdfunding.scilla` is given below.

```
[
  {
    "vname": "backers",
    "type": "Map (ByStr20) (Uint128)",
```

(continues on next page)

(continued from previous page)

```
"value": [  
  {  
    "key": "0x12345678901234567890123456789012345678cd",  
    "val": "2000000000000000"  
  },  
  {  
    "key": "0x12345678901234567890123456789012345678ab",  
    "val": "1000000000000000"  
  }  
]  
},  
{  
  "vname": "funded",  
  "type": "Bool",  
  "value": {  
    "constructor": "False",  
    "argtypes": [],  
    "arguments": []  
  }  
},  
{  
  "vname": "_balance",  
  "type": "Uint128",  
  "value": "3000000000000000"  
}  
]
```

## 3.9 Contact

Questions? Talk to us on [Discord](#).