Solidity Documentation

Release 0.7.6

Ethereum

Nov 26, 2020

Basics

Solidity is an object-oriented, high-level language for implementing smart contracts. Smart contracts are programs which govern the behaviour of accounts within the Ethereum state.

Solidity was influenced by C++, Python and JavaScript and is designed to target the Ethereum Virtual Machine (EVM).

Solidity is statically typed, supports inheritance, libraries and complex user-defined types among other features.

With Solidity you can create contracts for uses such as voting, crowdfunding, blind auctions, and multi-signature wallets.

When deploying contracts, you should use the latest released version of Solidity. This is because breaking changes as well as new features and bug fixes are introduced regularly. We currently use a 0.x version number to indicate this fast pace of change.

> **Warning:** Solidity recently released the 0.7.x version that introduced a lot of breaking changes. Make sure you read *the full list*.

Ideas for improving Solidity or this documentation are always welcome, read our *contributors guide* for more details.

Getting Started

1. Understand the Smart Contract Basics

If you are new to the concept of smart contracts we recommend you to get started by digging into the "Introduction to Smart Contracts" section, which covers:

- *A simple example smart contract* written in Solidity.

- *Blockchain Basics*.

- *The Ethereum Virtual Machine*.

2. Get to Know Solidity

Once you are accustomed to the basics, we recommend you read the *"Solidity by Example"* and "Language Description" sections to understand the core concepts of the language.

3. Install the Solidity Compiler

There are various ways to install the Solidity compiler, simply choose your preferred option and follow the steps outlined on the *installation page*.

Hint: You can try out code examples directly in your browser with the Remix IDE. Remix is a web browser based IDE that allows you to write, deploy and administer Solidity smart contracts, without the need to install Solidity locally.

Warning: As humans write software, it can have bugs. You should follow established software development best-practices when writing your smart contracts. This includes code review, testing, audits, and correctness proofs. Smart contract users are sometimes more confident with code than their authors, and blockchains and smart contracts have their own unique issues to watch out for, so before working on production code, make sure you read the *Security Considerations* section.

4. Learn More

If you want to learn more about building decentralized applications on Ethereum, the Ethereum Developer Resources can help you with further general documentation around Ethereum, and a wide selection of tutorials, tools and development frameworks.

If you have any questions, you can try searching for answers or asking on the Ethereum StackExchange, or our Gitter channel.

Translations

Community volunteers help translate this documentation into several languages. They have varying degrees of completeness and up-to-dateness. The English version stands as a reference.

- French (in progress)
- Italian (in progress)
- Japanese
- Korean (in progress)
- Russian (rather outdated)
- Simplified Chinese (in progress)
- Spanish
- Turkish (partial)

Contents

3.1 Introduction to Smart Contracts

3.1.1 A Simple Smart Contract

Let us begin with a basic example that sets the value of a variable and exposes it for other contracts to access. It is fine if you do not understand everything right now, we will go into more detail later.

Storage Example

```
// SPDX-License-Identifier: GPL-3.0
pragma solidity >=0.4.16 <0.8.0;

contract SimpleStorage {
    uint storedData;

    function set(uint x) public {
        storedData = x;
    }

    function get() public view returns (uint) {
        return storedData;
    }
}
```

The first line tells you that the source code is licensed under the GPL version 3.0. Machine-readable license specifiers are important in a setting where publishing the source code is the default.

The next line specifies that the source code is written for Solidity version 0.4.16, or a newer version of the language up to, but not including version 0.8.0. This is to ensure that the contract is not compilable with a new (breaking) compiler

version, where it could behave differently. *Pragmas* are common instructions for compilers about how to treat the source code (e.g. pragma once).

A contract in the sense of Solidity is a collection of code (its *functions*) and data (its *state*) that resides at a specific address on the Ethereum blockchain. The line `uint storedData;` declares a state variable called `storedData` of type `uint` (*un*signed *int*eger of *256* bits). You can think of it as a single slot in a database that you can query and alter by calling functions of the code that manages the database. In this example, the contract defines the functions `set` and `get` that can be used to modify or retrieve the value of the variable.

To access a state variable, you do not need the prefix `this.` as is common in other languages.

This contract does not do much yet apart from (due to the infrastructure built by Ethereum) allowing anyone to store a single number that is accessible by anyone in the world without a (feasible) way to prevent you from publishing this number. Anyone could call `set` again with a different value and overwrite your number, but the number is still stored in the history of the blockchain. Later, you will see how you can impose access restrictions so that only you can alter the number.

> **Warning:** Be careful with using Unicode text, as similar looking (or even identical) characters can have different code points and as such are encoded as a different byte array.

> **Note:** All identifiers (contract names, function names and variable names) are restricted to the ASCII character set. It is possible to store UTF-8 encoded data in string variables.

Subcurrency Example

The following contract implements the simplest form of a cryptocurrency. The contract allows only its creator to create new coins (different issuance schemes are possible). Anyone can send coins to each other without a need for registering with a username and password, all you need is an Ethereum keypair.

```solidity
// SPDX-License-Identifier: GPL-3.0
pragma solidity >0.5.99 <0.8.0;

contract Coin {
    // The keyword "public" makes variables
    // accessible from other contracts
    address public minter;
    mapping (address => uint) public balances;

    // Events allow clients to react to specific
    // contract changes you declare
    event Sent(address from, address to, uint amount);

    // Constructor code is only run when the contract
    // is created
    constructor() {
        minter = msg.sender;
    }

    // Sends an amount of newly created coins to an address
    // Can only be called by the contract creator
    function mint(address receiver, uint amount) public {
        require(msg.sender == minter);
        require(amount < 1e60);
```

```
        balances[receiver] += amount;
    }

    // Sends an amount of existing coins
    // from any caller to an address
    function send(address receiver, uint amount) public {
        require(amount <= balances[msg.sender], "Insufficient balance.");
        balances[msg.sender] -= amount;
        balances[receiver] += amount;
        emit Sent(msg.sender, receiver, amount);
    }
}
```

This contract introduces some new concepts, let us go through them one by one.

The line `address public minter;` declares a state variable of type *address*. The `address` type is a 160-bit value that does not allow any arithmetic operations. It is suitable for storing addresses of contracts, or a hash of the public half of a keypair belonging to *external accounts*.

The keyword `public` automatically generates a function that allows you to access the current value of the state variable from outside of the contract. Without this keyword, other contracts have no way to access the variable. The code of the function generated by the compiler is equivalent to the following (ignore `external` and `view` for now):

```
function minter() external view returns (address) { return minter; }
```

You could add a function like the above yourself, but you would have a function and state variable with the same name. You do not need to do this, the compiler figures it out for you.

The next line, `mapping (address => uint) public balances;` also creates a public state variable, but it is a more complex datatype. The *mapping* type maps addresses to *unsigned integers*.

Mappings can be seen as hash tables which are virtually initialised such that every possible key exists from the start and is mapped to a value whose byte-representation is all zeros. However, it is neither possible to obtain a list of all keys of a mapping, nor a list of all values. Record what you added to the mapping, or use it in a context where this is not needed. Or even better, keep a list, or use a more suitable data type.

The *getter function* created by the `public` keyword is more complex in the case of a mapping. It looks like the following:

```
function balances(address _account) external view returns (uint) {
    return balances[_account];
}
```

You can use this function to query the balance of a single account.

The line `event Sent(address from, address to, uint amount);` declares an *"event"*, which is emitted in the last line of the function `send`. Ethereum clients such as web applications can listen for these events emitted on the blockchain without much cost. As soon as it is emitted, the listener receives the arguments `from`, `to` and `amount`, which makes it possible to track transactions.

To listen for this event, you could use the following JavaScript code, which uses web3.js to create the `Coin` contract object, and any user interface calls the automatically generated `balances` function from above:

```
Coin.Sent().watch({}, '', function(error, result) {
    if (!error) {
        console.log("Coin transfer: " + result.args.amount +
            " coins were sent from " + result.args.from +
            " to " + result.args.to + ".");
```

```
        console.log("Balances now:\n" +
            "Sender: " + Coin.balances.call(result.args.from) +
            "Receiver: " + Coin.balances.call(result.args.to));
    }
})
```

The *constructor* is a special function that is executed during the creation of the contract and cannot be called afterwards. In this case, it permanently stores the address of the person creating the contract. The `msg` variable (together with `tx` and `block`) is a *special global variable* that contains properties which allow access to the blockchain. `msg.sender` is always the address where the current (external) function call came from.

The functions that make up the contract, and that users and contracts can call are `mint` and `send`.

The `mint` function sends an amount of newly created coins to another address. The *require* function call defines conditions that reverts all changes if not met. In this example, `require(msg.sender == minter);` ensures that only the creator of the contract can call `mint`, and `require(amount < 1e60);` ensures a maximum amount of tokens. This ensures that there are no overflow errors in the future.

The `send` function can be used by anyone (who already has some of these coins) to send coins to anyone else. If the sender does not have enough coins to send, the `require` call fails and provides the sender with an appropriate error message string.

Note: If you use this contract to send coins to an address, you will not see anything when you look at that address on a blockchain explorer, because the record that you sent coins and the changed balances are only stored in the data storage of this particular coin contract. By using events, you can create a "blockchain explorer" that tracks transactions and balances of your new coin, but you have to inspect the coin contract address and not the addresses of the coin owners.

3.1.2 Blockchain Basics

Blockchains as a concept are not too hard to understand for programmers. The reason is that most of the complications (mining, hashing, elliptic-curve cryptography, peer-to-peer networks, etc.) are just there to provide a certain set of features and promises for the platform. Once you accept these features as given, you do not have to worry about the underlying technology - or do you have to know how Amazon's AWS works internally in order to use it?

Transactions

A blockchain is a globally shared, transactional database. This means that everyone can read entries in the database just by participating in the network. If you want to change something in the database, you have to create a so-called transaction which has to be accepted by all others. The word transaction implies that the change you want to make (assume you want to change two values at the same time) is either not done at all or completely applied. Furthermore, while your transaction is being applied to the database, no other transaction can alter it.

As an example, imagine a table that lists the balances of all accounts in an electronic currency. If a transfer from one account to another is requested, the transactional nature of the database ensures that if the amount is subtracted from one account, it is always added to the other account. If due to whatever reason, adding the amount to the target account is not possible, the source account is also not modified.

Furthermore, a transaction is always cryptographically signed by the sender (creator). This makes it straightforward to guard access to specific modifications of the database. In the example of the electronic currency, a simple check ensures that only the person holding the keys to the account can transfer money from it.

Blocks

One major obstacle to overcome is what (in Bitcoin terms) is called a "double-spend attack": What happens if two transactions exist in the network that both want to empty an account? Only one of the transactions can be valid, typically the one that is accepted first. The problem is that "first" is not an objective term in a peer-to-peer network.

The abstract answer to this is that you do not have to care. A globally accepted order of the transactions will be selected for you, solving the conflict. The transactions will be bundled into what is called a "block" and then they will be executed and distributed among all participating nodes. If two transactions contradict each other, the one that ends up being second will be rejected and not become part of the block.

These blocks form a linear sequence in time and that is where the word "blockchain" derives from. Blocks are added to the chain in rather regular intervals - for Ethereum this is roughly every 17 seconds.

As part of the "order selection mechanism" (which is called "mining") it may happen that blocks are reverted from time to time, but only at the "tip" of the chain. The more blocks are added on top of a particular block, the less likely this block will be reverted. So it might be that your transactions are reverted and even removed from the blockchain, but the longer you wait, the less likely it will be.

Note: Transactions are not guaranteed to be included in the next block or any specific future block, since it is not up to the submitter of a transaction, but up to the miners to determine in which block the transaction is included.

If you want to schedule future calls of your contract, you can use the alarm clock or a similar oracle service.

3.1.3 The Ethereum Virtual Machine

Overview

The Ethereum Virtual Machine or EVM is the runtime environment for smart contracts in Ethereum. It is not only sandboxed but actually completely isolated, which means that code running inside the EVM has no access to network, filesystem or other processes. Smart contracts even have limited access to other smart contracts.

Accounts

There are two kinds of accounts in Ethereum which share the same address space: **External accounts** that are controlled by public-private key pairs (i.e. humans) and **contract accounts** which are controlled by the code stored together with the account.

The address of an external account is determined from the public key while the address of a contract is determined at the time the contract is created (it is derived from the creator address and the number of transactions sent from that address, the so-called "nonce").

Regardless of whether or not the account stores code, the two types are treated equally by the EVM.

Every account has a persistent key-value store mapping 256-bit words to 256-bit words called **storage**.

Furthermore, every account has a **balance** in Ether (in "Wei" to be exact, `1 ether` is `10**18 wei`) which can be modified by sending transactions that include Ether.

Transactions

A transaction is a message that is sent from one account to another account (which might be the same or empty, see below). It can include binary data (which is called "payload") and Ether.

If the target account contains code, that code is executed and the payload is provided as input data.

If the target account is not set (the transaction does not have a recipient or the recipient is set to `null`), the transaction creates a **new contract**. As already mentioned, the address of that contract is not the zero address but an address derived from the sender and its number of transactions sent (the "nonce"). The payload of such a contract creation transaction is taken to be EVM bytecode and executed. The output data of this execution is permanently stored as the code of the contract. This means that in order to create a contract, you do not send the actual code of the contract, but in fact code that returns that code when executed.

Note: While a contract is being created, its code is still empty. Because of that, you should not call back into the contract under construction until its constructor has finished executing.

Gas

Upon creation, each transaction is charged with a certain amount of **gas**, whose purpose is to limit the amount of work that is needed to execute the transaction and to pay for this execution at the same time. While the EVM executes the transaction, the gas is gradually depleted according to specific rules.

The **gas price** is a value set by the creator of the transaction, who has to pay `gas_price * gas` up front from the sending account. If some gas is left after the execution, it is refunded to the creator in the same way.

If the gas is used up at any point (i.e. it would be negative), an out-of-gas exception is triggered, which reverts all modifications made to the state in the current call frame.

Storage, Memory and the Stack

The Ethereum Virtual Machine has three areas where it can store data- storage, memory and the stack, which are explained in the following paragraphs.

Each account has a data area called **storage**, which is persistent between function calls and transactions. Storage is a key-value store that maps 256-bit words to 256-bit words. It is not possible to enumerate storage from within a contract, it is comparatively costly to read, and even more to initialise and modify storage. Because of this cost, you should minimize what you store in persistent storage to what the contract needs to run. Store data like derived calculations, caching, and aggregates outside of the contract. A contract can neither read nor write to any storage apart from its own.

The second data area is called **memory**, of which a contract obtains a freshly cleared instance for each message call. Memory is linear and can be addressed at byte level, but reads are limited to a width of 256 bits, while writes can be either 8 bits or 256 bits wide. Memory is expanded by a word (256-bit), when accessing (either reading or writing) a previously untouched memory word (i.e. any offset within a word). At the time of expansion, the cost in gas must be paid. Memory is more costly the larger it grows (it scales quadratically).

The EVM is not a register machine but a stack machine, so all computations are performed on a data area called the **stack**. It has a maximum size of 1024 elements and contains words of 256 bits. Access to the stack is limited to the top end in the following way: It is possible to copy one of the topmost 16 elements to the top of the stack or swap the topmost element with one of the 16 elements below it. All other operations take the topmost two (or one, or more, depending on the operation) elements from the stack and push the result onto the stack. Of course it is possible to move stack elements to storage or memory in order to get deeper access to the stack, but it is not possible to just access arbitrary elements deeper in the stack without first removing the top of the stack.

Instruction Set

The instruction set of the EVM is kept minimal in order to avoid incorrect or inconsistent implementations which could cause consensus problems. All instructions operate on the basic data type, 256-bit words or on slices of memory (or other byte arrays). The usual arithmetic, bit, logical and comparison operations are present. Conditional and unconditional jumps are possible. Furthermore, contracts can access relevant properties of the current block like its number and timestamp.

For a complete list, please see the *list of opcodes* as part of the inline assembly documentation.

Message Calls

Contracts can call other contracts or send Ether to non-contract accounts by the means of message calls. Message calls are similar to transactions, in that they have a source, a target, data payload, Ether, gas and return data. In fact, every transaction consists of a top-level message call which in turn can create further message calls.

A contract can decide how much of its remaining **gas** should be sent with the inner message call and how much it wants to retain. If an out-of-gas exception happens in the inner call (or any other exception), this will be signaled by an error value put onto the stack. In this case, only the gas sent together with the call is used up. In Solidity, the calling contract causes a manual exception by default in such situations, so that exceptions "bubble up" the call stack.

As already said, the called contract (which can be the same as the caller) will receive a freshly cleared instance of memory and has access to the call payload - which will be provided in a separate area called the **calldata**. After it has finished execution, it can return data which will be stored at a location in the caller's memory preallocated by the caller. All such calls are fully synchronous.

Calls are **limited** to a depth of 1024, which means that for more complex operations, loops should be preferred over recursive calls. Furthermore, only 63/64th of the gas can be forwarded in a message call, which causes a depth limit of a little less than 1000 in practice.

Delegatecall / Callcode and Libraries

There exists a special variant of a message call, named **delegatecall** which is identical to a message call apart from the fact that the code at the target address is executed in the context of the calling contract and `msg.sender` and `msg.value` do not change their values.

This means that a contract can dynamically load code from a different address at runtime. Storage, current address and balance still refer to the calling contract, only the code is taken from the called address.

This makes it possible to implement the "library" feature in Solidity: Reusable library code that can be applied to a contract's storage, e.g. in order to implement a complex data structure.

Logs

It is possible to store data in a specially indexed data structure that maps all the way up to the block level. This feature called **logs** is used by Solidity in order to implement *events*. Contracts cannot access log data after it has been created, but they can be efficiently accessed from outside the blockchain. Since some part of the log data is stored in bloom filters, it is possible to search for this data in an efficient and cryptographically secure way, so network peers that do not download the whole blockchain (so-called "light clients") can still find these logs.

Create

Contracts can even create other contracts using a special opcode (i.e. they do not simply call the zero address as a transaction would). The only difference between these **create calls** and normal message calls is that the payload data

is executed and the result stored as code and the caller / creator receives the address of the new contract on the stack.

Deactivate and Self-destruct

The only way to remove code from the blockchain is when a contract at that address performs the `selfdestruct` operation. The remaining Ether stored at that address is sent to a designated target and then the storage and code is removed from the state. Removing the contract in theory sounds like a good idea, but it is potentially dangerous, as if someone sends Ether to removed contracts, the Ether is forever lost.

> **Warning:** Even if a contract is removed by `selfdestruct`, it is still part of the history of the blockchain and probably retained by most Ethereum nodes. So using `selfdestruct` is not the same as deleting data from a hard disk.

Note: Even if a contract's code does not contain a call to `selfdestruct`, it can still perform that operation using `delegatecall` or `callcode`.

If you want to deactivate your contracts, you should instead **disable** them by changing some internal state which causes all functions to revert. This makes it impossible to use the contract, as it returns Ether immediately.

3.2 Installing the Solidity Compiler

3.2.1 Versioning

Solidity versions follow semantic versioning and in addition to releases, **nightly development builds** are also made available. The nightly builds are not guaranteed to be working and despite best efforts they might contain undocumented and/or broken changes. We recommend using the latest release. Package installers below will use the latest release.

3.2.2 Remix

We recommend Remix for small contracts and for quickly learning Solidity.

Access Remix online, you do not need to install anything. If you want to use it without connection to the Internet, go to https://github.com/ethereum/remix-live/tree/gh-pages and download the `.zip` file as explained on that page. Remix is also a convenient option for testing nightly builds without installing multiple Solidity versions.

Further options on this page detail installing commandline Solidity compiler software on your computer. Choose a commandline compiler if you are working on a larger contract or if you require more compilation options.

3.2.3 npm / Node.js

Use `npm` for a convenient and portable way to install `solcjs`, a Solidity compiler. The *solcjs* program has fewer features than the ways to access the compiler described further down this page. The *Using the Commandline Compiler* documentation assumes you are using the full-featured compiler, `solc`. The usage of `solcjs` is documented inside its own repository.

Note: The solc-js project is derived from the C++ *solc* by using Emscripten which means that both use the same compiler source code. *solc-js* can be used in JavaScript projects directly (such as Remix). Please refer to the solc-js repository for instructions.

```
npm install -g solc
```

Note: The commandline executable is named `solcjs`.

The commandline options of `solcjs` are not compatible with `solc` and tools (such as `geth`) expecting the behaviour of `solc` will not work with `solcjs`.

3.2.4 Docker

Docker images of Solidity builds are available using the `solc` image from the `ethereum` organisation. Use the `stable` tag for the latest released version, and `nightly` for potentially unstable changes in the develop branch.

The Docker image runs the compiler executable, so you can pass all compiler arguments to it. For example, the command below pulls the stable version of the `solc` image (if you do not have it already), and runs it in a new container, passing the `--help` argument.

```
docker run ethereum/solc:stable --help
```

You can also specify release build versions in the tag, for example, for the 0.5.4 release.

```
docker run ethereum/solc:0.5.4 --help
```

To use the Docker image to compile Solidity files on the host machine mount a local folder for input and output, and specify the contract to compile. For example.

```
docker run -v /local/path:/sources ethereum/solc:stable -o /sources/output --abi --
→bin /sources/Contract.sol
```

You can also use the standard JSON interface (which is recommended when using the compiler with tooling). When using this interface it is not necessary to mount any directories.

```
docker run ethereum/solc:stable --standard-json < input.json > output.json
```

3.2.5 Linux Packages

Binary packages of Solidity are available at solidity/releases.

We also have PPAs for Ubuntu, you can get the latest stable version using the following commands:

```
sudo add-apt-repository ppa:ethereum/ethereum
sudo apt-get update
sudo apt-get install solc
```

The nightly version can be installed using these commands:

```
sudo add-apt-repository ppa:ethereum/ethereum
sudo add-apt-repository ppa:ethereum/ethereum-dev
sudo apt-get update
sudo apt-get install solc
```

We are also releasing a snap package, which is installable in all the supported Linux distros. To install the latest stable version of solc:

```
sudo snap install solc
```

If you want to help testing the latest development version of Solidity with the most recent changes, please use the following:

```
sudo snap install solc --edge
```

Note: The `solc` snap uses strict confinement. This is the most secure mode for snap packages but it comes with limitations, like accessing only the files in your `/home` and `/media` directories. For more information, go to Demystifying Snap Confinement.

Arch Linux also has packages, albeit limited to the latest development version:

```
pacman -S solidity
```

Gentoo Linux has an Ethereum overlay that contains a Solidity package. After the overlay is setup, `solc` can be installed in x86_64 architectures by:

```
emerge dev-lang/solidity
```

3.2.6 macOS Packages

We distribute the Solidity compiler through Homebrew as a build-from-source version. Pre-built bottles are currently not supported.

```
brew update
brew upgrade
brew tap ethereum/ethereum
brew install solidity
```

To install the most recent 0.4.x / 0.5.x version of Solidity you can also use `brew install solidity@4` and `brew install solidity@5`, respectively.

If you need a specific version of Solidity you can install a Homebrew formula directly from Github.

View solidity.rb commits on Github.

Copy the commit hash of the version you want and check it out on your machine.

```
git clone https://github.com/ethereum/homebrew-ethereum.git
cd homebrew-ethereum
git checkout <your-hash-goes-here>
```

Install it using `brew`:

```
brew unlink solidity
# eg. Install 0.4.8
brew install solidity.rb
```

3.2.7 Static Binaries

We maintain a repository containing static builds of past and current compiler versions for all supported platforms at solc-bin. This is also the location where you can find the nightly builds.

The repository is not only a quick and easy way for end users to get binaries ready to be used out-of-the-box but it is also meant to be friendly to third-party tools:

- The content is mirrored to https://binaries.soliditylang.org where it can be easily downloaded over HTTPS without any authentication, rate limiting or the need to use git.

- Content is served with correct *Content-Type* headers and lenient CORS configuration so that it can be directly loaded by tools running in the browser.

- Binaries do not require installation or unpacking (with the exception of older Windows builds bundled with necessary DLLs).

- We strive for a high level of backwards-compatibility. Files, once added, are not removed or moved without providing a symlink/redirect at the old location. They are also never modified in place and should always match the original checksum. The only exception would be broken or unusable files with a potential to cause more harm than good if left as is.

- Files are served over both HTTP and HTTPS. As long as you obtain the file list in a secure way (via git, HTTPS, IPFS or just have it cached locally) and verify hashes of the binaries after downloading them, you do not have to use HTTPS for the binaries themselves.

The same binaries are in most cases available on the Solidity release page on Github. The difference is that we do not generally update old releases on the Github release page. This means that we do not rename them if the naming convention changes and we do not add builds for platforms that were not supported at the time of release. This only happens in `solc-bin`.

The `solc-bin` repository contains several top-level directories, each representing a single platform. Each one contains a `list.json` file listing the available binaries. For example in `emscripten-wasm32/list.json` you will find the following information about version 0.7.4:

```json
{
  "path": "solc-emscripten-wasm32-v0.7.4+commit.3f05b770.js",
  "version": "0.7.4",
  "build": "commit.3f05b770",
  "longVersion": "0.7.4+commit.3f05b770",
  "keccak256": "0x300330ecd127756b824aa13e843cb1f43c473cb22eaf3750d5fb9c99279af8c3",
  "urls": [
    "bzzr://16c5f09109c793db99fe35f037c6092b061bd39260ee7a677c8a97f18c955ab1",
    "dweb:/ipfs/QmTLs5MuLEWXQkths41HiACoXDiH8zxyqBHGFDRSzVE5CS"
  ]
}
```

This means that:

- You can find the binary in the same directory under the name solc-emscripten-wasm32-v0.7.4+commit.3f05b770.js. Note that the file might be a symlink, and you will need to resolve it yourself if you are not using git to download it or your file system does not support symlinks.

- The binary is also mirrored at https://binaries.soliditylang.org/emscripten-wasm32/solc-emscripten-wasm32-v0.7.4+commit.3f05b770.js. In this case git is not necessary and symlinks are resolved transparently, either by serving a copy of the file or returning a HTTP redirect.

- The file is also available on IPFS at QmTLs5MuLEWXQkths41HiACoXDiH8zxyqBHGFDRSzVE5CS.

- The file might in future be available on Swarm at 16c5f09109c793db99fe35f037c6092b061bd39260ee7a677c8a97f18c955ab1.

- You can verify the integrity of the binary by comparing its keccak256 hash to `0x300330ecd127756b824aa13e843cb1f43c473cb22eaf3750d5fb9c99279af8c3`. The hash can be computed on the command line using `keccak256sum` utility provided by sha3sum or keccak256() function from ethereumjs-util in JavaScript.

> **Warning:** Due to the strong backwards compatibility requirement the repository contains some legacy elements
> but you should avoid using them when writing new tools:
>
> - Use `emscripten-wasm32/` (with a fallback to `emscripten-asmjs/`) instead of `bin/` if you want
>   the best performance. Until version 0.6.1 we only provided asm.js binaries. Starting with 0.6.2 we switched
>   to WebAssembly builds with much better performance. We have rebuilt the older versions for wasm but the
>   original asm.js files remain in `bin/`. The new ones had to be placed in a separate directory to avoid name
>   clashes.
>
> - Use `emscripten-asmjs/` and `emscripten-wasm32/` instead of `bin/` and `wasm/` directories if
>   you want to be sure whether you are downloading a wasm or an asm.js binary.
>
> - Use `list.json` instead of `list.js` and `list.txt`. The JSON list format contains all the information
>   from the old ones and more.
>
> - Use https://binaries.soliditylang.org instead of https://solc-bin.ethereum.org. To keep things simple we
>   moved almost everything related to the compiler under the new `soliditylang.org` domain and this
>   applies to `solc-bin` too. While the new domain is recommended, the old one is still fully supported and
>   guaranteed to point at the same location.

> **Warning:** The binaries are also available at https://ethereum.github.io/solc-bin/ but this page stopped being
> updated just after the release of version 0.7.2, will not receive any new releases or nightly builds for any platform
> and does not serve the new directory structure, including non-emscripten builds.
>
> If you are using it, please switch to https://binaries.soliditylang.org, which is a drop-in replacement. This allows us
> to make changes to the underlying hosting in a transparent way and minimize disruption. Unlike the `ethereum.`
> `github.io` domain, which we do not have any control over, `binaries.soliditylang.org` is guaranteed
> to work and maintain the same URL structure in the long-term.

3.2.8 Building from Source

Prerequisites - All Operating Systems

The following are dependencies for all builds of Solidity:

Software	Notes
CMake (version 3.9+)	Cross-platform build file generator.
Boost (version 1.65+)	C++ libraries.
Git	Command-line tool for retrieving source code.
z3 (version 4.6+, Optional)	For use with SMT checker.
cvc4 (Optional)	For use with SMT checker.

> **Note:** Solidity versions prior to 0.5.10 can fail to correctly link against Boost versions 1.70+. A possible workaround
> is to temporarily rename `<Boost install path>/lib/cmake/Boost-1.70.0` prior to running the cmake
> command to configure solidity.
>
> Starting from 0.5.10 linking against Boost 1.70+ should work without manual intervention.

Minimum compiler versions

The following C++ compilers and their minimum versions can build the Solidity codebase:

- GCC, version 5+
- Clang, version 3.4+
- MSVC, version 2019+

Prerequisites - macOS

For macOS builds, ensure that you have the latest version of Xcode installed. This contains the Clang C++ compiler, the Xcode IDE and other Apple development tools that are required for building C++ applications on OS X. If you are installing Xcode for the first time, or have just installed a new version then you will need to agree to the license before you can do command-line builds:

```
sudo xcodebuild -license accept
```

Our OS X build script uses the Homebrew package manager for installing external dependencies. Here's how to uninstall Homebrew, if you ever want to start again from scratch.

Prerequisites - Windows

You need to install the following dependencies for Windows builds of Solidity:

Software	Notes
Visual Studio 2019 Build Tools	C++ compiler
Visual Studio 2019 (Optional)	C++ compiler and dev environment.

If you already have one IDE and only need the compiler and libraries, you could install Visual Studio 2019 Build Tools.

Visual Studio 2019 provides both IDE and necessary compiler and libraries. So if you have not got an IDE and prefer to develop Solidity, Visual Studio 2019 may be a choice for you to get everything setup easily.

Here is the list of components that should be installed in Visual Studio 2019 Build Tools or Visual Studio 2019:

- Visual Studio C++ core features
- VC++ 2019 v141 toolset (x86,x64)
- Windows Universal CRT SDK
- Windows 8.1 SDK
- C++/CLI support

Dependencies Helper Script

We have a helper script which you can use to install all required external dependencies on macOS, Windows and on numerous Linux distros.

```
./scripts/install_deps.sh
```

Or, on Windows:

```
scripts\install_deps.ps1
```

Note that the latter command will install `boost` and `cmake` to the `deps` subdirectory, while the former command will attempt to install the dependencies globally.

Clone the Repository

To clone the source code, execute the following command:

```
git clone --recursive https://github.com/ethereum/solidity.git
cd solidity
```

If you want to help developing Solidity, you should fork Solidity and add your personal fork as a second remote:

```
git remote add personal git@github.com:[username]/solidity.git
```

Note: This method will result in a prerelease build leading to e.g. a flag being set in each bytecode produced by such a compiler. If you want to re-build a released Solidity compiler, then please use the source tarball on the github release page:

https://github.com/ethereum/solidity/releases/download/v0.X.Y/solidity_0.X.Y.tar.gz

(not the "Source code" provided by github).

Command-Line Build

Be sure to install External Dependencies (see above) before build.

Solidity project uses CMake to configure the build. You might want to install ccache to speed up repeated builds. CMake will pick it up automatically. Building Solidity is quite similar on Linux, macOS and other Unices:

```
mkdir build
cd build
cmake .. && make
```

or even easier on Linux and macOS, you can run:

```
#note: this will install binaries solc and soltest at usr/local/bin
./scripts/build.sh
```

Warning: BSD builds should work, but are untested by the Solidity team.

And for Windows:

```
mkdir build
cd build
cmake -G "Visual Studio 16 2019 Win64" ..
```

In case you want to use the version of boost installed by `./scripts/install_deps.ps1`, you will additionally need to pass `-DBoost_DIR="..\deps\boost\lib\cmake\Boost-*"` and `-DCMAKE_MSVC_RUNTIME_LIBRARY=MultiThreaded` as arguments to the call to `cmake`.

This should result in the creation of **solidity.sln** in that build directory. Double-clicking on that file should result in Visual Studio firing up. We suggest building **Release** configuration, but all others work.

Alternatively, you can build for Windows on the command-line, like so:

```
cmake --build . --config Release
```

3.2.9 CMake Options

If you are interested what CMake options are available run `cmake .. -LH`.

SMT Solvers

Solidity can be built against SMT solvers and will do so by default if they are found in the system. Each solver can be disabled by a *cmake* option.

Note: In some cases, this can also be a potential workaround for build failures.

Inside the build folder you can disable them, since they are enabled by default:

```
# disables only Z3 SMT Solver.
cmake .. -DUSE_Z3=OFF

# disables only CVC4 SMT Solver.
cmake .. -DUSE_CVC4=OFF

# disables both Z3 and CVC4
cmake .. -DUSE_CVC4=OFF -DUSE_Z3=OFF
```

3.2.10 The Version String in Detail

The Solidity version string contains four parts:

- the version number
- pre-release tag, usually set to `develop.YYYY.MM.DD` or `nightly.YYYY.MM.DD`
- commit in the format of `commit.GITHASH`
- platform, which has an arbitrary number of items, containing details about the platform and compiler

If there are local modifications, the commit will be postfixed with `.mod`.

These parts are combined as required by Semver, where the Solidity pre-release tag equals to the Semver pre-release and the Solidity commit and platform combined make up the Semver build metadata.

A release example: `0.4.8+commit.60cc1668.Emscripten.clang`.

A pre-release example: `0.4.9-nightly.2017.1.17+commit.6ecb4aa3.Emscripten.clang`

3.2.11 Important Information About Versioning

After a release is made, the patch version level is bumped, because we assume that only patch level changes follow. When changes are merged, the version should be bumped according to semver and the severity of the change. Finally, a release is always made with the version of the current nightly build, but without the `prerelease` specifier.

Example:

0. the 0.4.0 release is made

1. nightly build has a version of 0.4.1 from now on

2. non-breaking changes are introduced - no change in version

3. a breaking change is introduced - version is bumped to 0.5.0

4. the 0.5.0 release is made

This behaviour works well with the *version pragma*.

3.3 Solidity by Example

3.3.1 Voting

The following contract is quite complex, but showcases a lot of Solidity's features. It implements a voting contract. Of course, the main problems of electronic voting is how to assign voting rights to the correct persons and how to prevent manipulation. We will not solve all problems here, but at least we will show how delegated voting can be done so that vote counting is **automatic and completely transparent** at the same time.

The idea is to create one contract per ballot, providing a short name for each option. Then the creator of the contract who serves as chairperson will give the right to vote to each address individually.

The persons behind the addresses can then choose to either vote themselves or to delegate their vote to a person they trust.

At the end of the voting time, `winningProposal()` will return the proposal with the largest number of votes.

```solidity
// SPDX-License-Identifier: GPL-3.0
pragma solidity ^0.7.0;

/// @title Voting with delegation.
contract Ballot {
    // This declares a new complex type which will
    // be used for variables later.
    // It will represent a single voter.
    struct Voter {
        uint weight; // weight is accumulated by delegation
        bool voted;  // if true, that person already voted
        address delegate; // person delegated to
        uint vote;   // index of the voted proposal
    }

    // This is a type for a single proposal.
    struct Proposal {
        bytes32 name;   // short name (up to 32 bytes)
        uint voteCount; // number of accumulated votes
    }

    address public chairperson;

    // This declares a state variable that
    // stores a `Voter` struct for each possible address.
    mapping(address => Voter) public voters;

    // A dynamically-sized array of `Proposal` structs.
    Proposal[] public proposals;
```

```solidity
    /// Create a new ballot to choose one of `proposalNames`.
    constructor(bytes32[] memory proposalNames) {
        chairperson = msg.sender;
        voters[chairperson].weight = 1;

        // For each of the provided proposal names,
        // create a new proposal object and add it
        // to the end of the array.
        for (uint i = 0; i < proposalNames.length; i++) {
            // `Proposal({...})` creates a temporary
            // Proposal object and `proposals.push(...)`
            // appends it to the end of `proposals`.
            proposals.push(Proposal({
                name: proposalNames[i],
                voteCount: 0
            }));
        }
    }

    // Give `voter` the right to vote on this ballot.
    // May only be called by `chairperson`.
    function giveRightToVote(address voter) public {
        // If the first argument of `require` evaluates
        // to `false`, execution terminates and all
        // changes to the state and to Ether balances
        // are reverted.
        // This used to consume all gas in old EVM versions, but
        // not anymore.
        // It is often a good idea to use `require` to check if
        // functions are called correctly.
        // As a second argument, you can also provide an
        // explanation about what went wrong.
        require(
            msg.sender == chairperson,
            "Only chairperson can give right to vote."
        );
        require(
            !voters[voter].voted,
            "The voter already voted."
        );
        require(voters[voter].weight == 0);
        voters[voter].weight = 1;
    }

    /// Delegate your vote to the voter `to`.
    function delegate(address to) public {
        // assigns reference
        Voter storage sender = voters[msg.sender];
        require(!sender.voted, "You already voted.");

        require(to != msg.sender, "Self-delegation is disallowed.");

        // Forward the delegation as long as
        // `to` also delegated.
        // In general, such loops are very dangerous,
        // because if they run too long, they might
```

```solidity
        // need more gas than is available in a block.
        // In this case, the delegation will not be executed,
        // but in other situations, such loops might
        // cause a contract to get "stuck" completely.
        while (voters[to].delegate != address(0)) {
            to = voters[to].delegate;

            // We found a loop in the delegation, not allowed.
            require(to != msg.sender, "Found loop in delegation.");
        }

        // Since `sender` is a reference, this
        // modifies `voters[msg.sender].voted`
        sender.voted = true;
        sender.delegate = to;
        Voter storage delegate_ = voters[to];
        if (delegate_.voted) {
            // If the delegate already voted,
            // directly add to the number of votes
            proposals[delegate_.vote].voteCount += sender.weight;
        } else {
            // If the delegate did not vote yet,
            // add to her weight.
            delegate_.weight += sender.weight;
        }
    }

    /// Give your vote (including votes delegated to you)
    /// to proposal `proposals[proposal].name`.
    function vote(uint proposal) public {
        Voter storage sender = voters[msg.sender];
        require(sender.weight != 0, "Has no right to vote");
        require(!sender.voted, "Already voted.");
        sender.voted = true;
        sender.vote = proposal;

        // If `proposal` is out of the range of the array,
        // this will throw automatically and revert all
        // changes.
        proposals[proposal].voteCount += sender.weight;
    }

    /// @dev Computes the winning proposal taking all
    /// previous votes into account.
    function winningProposal() public view
            returns (uint winningProposal_)
    {
        uint winningVoteCount = 0;
        for (uint p = 0; p < proposals.length; p++) {
            if (proposals[p].voteCount > winningVoteCount) {
                winningVoteCount = proposals[p].voteCount;
                winningProposal_ = p;
            }
        }
    }

    // Calls winningProposal() function to get the index
```

```solidity
    // of the winner contained in the proposals array and then
    // returns the name of the winner
    function winnerName() public view
            returns (bytes32 winnerName_)
    {
        winnerName_ = proposals[winningProposal()].name;
    }
}
```

Possible Improvements

Currently, many transactions are needed to assign the rights to vote to all participants. Can you think of a better way?

3.3.2 Blind Auction

In this section, we will show how easy it is to create a completely blind auction contract on Ethereum. We will start with an open auction where everyone can see the bids that are made and then extend this contract into a blind auction where it is not possible to see the actual bid until the bidding period ends.

Simple Open Auction

The general idea of the following simple auction contract is that everyone can send their bids during a bidding period. The bids already include sending money / Ether in order to bind the bidders to their bid. If the highest bid is raised, the previously highest bidder gets their money back. After the end of the bidding period, the contract has to be called manually for the beneficiary to receive their money - contracts cannot activate themselves.

```solidity
// SPDX-License-Identifier: GPL-3.0
pragma solidity ^0.7.0;

contract SimpleAuction {
    // Parameters of the auction. Times are either
    // absolute unix timestamps (seconds since 1970-01-01)
    // or time periods in seconds.
    address payable public beneficiary;
    uint public auctionEndTime;

    // Current state of the auction.
    address public highestBidder;
    uint public highestBid;

    // Allowed withdrawals of previous bids
    mapping(address => uint) pendingReturns;

    // Set to true at the end, disallows any change.
    // By default initialized to `false`.
    bool ended;

    // Events that will be emitted on changes.
    event HighestBidIncreased(address bidder, uint amount);
    event AuctionEnded(address winner, uint amount);

    // The following is a so-called natspec comment,
```

```solidity
    // recognizable by the three slashes.
    // It will be shown when the user is asked to
    // confirm a transaction.

    /// Create a simple auction with `_biddingTime`
    /// seconds bidding time on behalf of the
    /// beneficiary address `_beneficiary`.
    constructor(
        uint _biddingTime,
        address payable _beneficiary
    ) {
        beneficiary = _beneficiary;
        auctionEndTime = block.timestamp + _biddingTime;
    }

    /// Bid on the auction with the value sent
    /// together with this transaction.
    /// The value will only be refunded if the
    /// auction is not won.
    function bid() public payable {
        // No arguments are necessary, all
        // information is already part of
        // the transaction. The keyword payable
        // is required for the function to
        // be able to receive Ether.

        // Revert the call if the bidding
        // period is over.
        require(
            block.timestamp <= auctionEndTime,
            "Auction already ended."
        );

        // If the bid is not higher, send the
        // money back (the failing require
        // will revert all changes in this
        // function execution including
        // it having received the money).
        require(
            msg.value > highestBid,
            "There already is a higher bid."
        );

        if (highestBid != 0) {
            // Sending back the money by simply using
            // highestBidder.send(highestBid) is a security risk
            // because it could execute an untrusted contract.
            // It is always safer to let the recipients
            // withdraw their money themselves.
            pendingReturns[highestBidder] += highestBid;
        }
        highestBidder = msg.sender;
        highestBid = msg.value;
        emit HighestBidIncreased(msg.sender, msg.value);
    }

    /// Withdraw a bid that was overbid.
```

```solidity
    function withdraw() public returns (bool) {
        uint amount = pendingReturns[msg.sender];
        if (amount > 0) {
            // It is important to set this to zero because the recipient
            // can call this function again as part of the receiving call
            // before `send` returns.
            pendingReturns[msg.sender] = 0;

            if (!msg.sender.send(amount)) {
                // No need to call throw here, just reset the amount owing
                pendingReturns[msg.sender] = amount;
                return false;
            }
        }
        return true;
    }

    /// End the auction and send the highest bid
    /// to the beneficiary.
    function auctionEnd() public {
        // It is a good guideline to structure functions that interact
        // with other contracts (i.e. they call functions or send Ether)
        // into three phases:
        // 1. checking conditions
        // 2. performing actions (potentially changing conditions)
        // 3. interacting with other contracts
        // If these phases are mixed up, the other contract could call
        // back into the current contract and modify the state or cause
        // effects (ether payout) to be performed multiple times.
        // If functions called internally include interaction with external
        // contracts, they also have to be considered interaction with
        // external contracts.

        // 1. Conditions
        require(block.timestamp >= auctionEndTime, "Auction not yet ended.");
        require(!ended, "auctionEnd has already been called.");

        // 2. Effects
        ended = true;
        emit AuctionEnded(highestBidder, highestBid);

        // 3. Interaction
        beneficiary.transfer(highestBid);
    }
}
```

Blind Auction

The previous open auction is extended to a blind auction in the following. The advantage of a blind auction is that there is no time pressure towards the end of the bidding period. Creating a blind auction on a transparent computing platform might sound like a contradiction, but cryptography comes to the rescue.

During the **bidding period**, a bidder does not actually send their bid, but only a hashed version of it. Since it is currently considered practically impossible to find two (sufficiently long) values whose hash values are equal, the bidder commits to the bid by that. After the end of the bidding period, the bidders have to reveal their bids: They send their values unencrypted and the contract checks that the hash value is the same as the one provided during the bidding

period.

Another challenge is how to make the auction **binding and blind** at the same time: The only way to prevent the bidder from just not sending the money after they won the auction is to make them send it together with the bid. Since value transfers cannot be blinded in Ethereum, anyone can see the value.

The following contract solves this problem by accepting any value that is larger than the highest bid. Since this can of course only be checked during the reveal phase, some bids might be **invalid**, and this is on purpose (it even provides an explicit flag to place invalid bids with high value transfers): Bidders can confuse competition by placing several high or low invalid bids.

```solidity
// SPDX-License-Identifier: GPL-3.0
pragma solidity ^0.7.0;

contract BlindAuction {
    struct Bid {
        bytes32 blindedBid;
        uint deposit;
    }

    address payable public beneficiary;
    uint public biddingEnd;
    uint public revealEnd;
    bool public ended;

    mapping(address => Bid[]) public bids;

    address public highestBidder;
    uint public highestBid;

    // Allowed withdrawals of previous bids
    mapping(address => uint) pendingReturns;

    event AuctionEnded(address winner, uint highestBid);

    /// Modifiers are a convenient way to validate inputs to
    /// functions. `onlyBefore` is applied to `bid` below:
    /// The new function body is the modifier's body where
    /// `_` is replaced by the old function body.
    modifier onlyBefore(uint _time) { require(block.timestamp < _time); _; }
    modifier onlyAfter(uint _time) { require(block.timestamp > _time); _; }

    constructor(
        uint _biddingTime,
        uint _revealTime,
        address payable _beneficiary
    ) {
        beneficiary = _beneficiary;
        biddingEnd = block.timestamp + _biddingTime;
        revealEnd = biddingEnd + _revealTime;
    }

    /// Place a blinded bid with `_blindedBid` =
    /// keccak256(abi.encodePacked(value, fake, secret)).
    /// The sent ether is only refunded if the bid is correctly
    /// revealed in the revealing phase. The bid is valid if the
    /// ether sent together with the bid is at least "value" and
    /// "fake" is not true. Setting "fake" to true and sending
```

```solidity
    /// not the exact amount are ways to hide the real bid but
    /// still make the required deposit. The same address can
    /// place multiple bids.
    function bid(bytes32 _blindedBid)
        public
        payable
        onlyBefore(biddingEnd)
    {
        bids[msg.sender].push(Bid({
            blindedBid: _blindedBid,
            deposit: msg.value
        }));
    }

    /// Reveal your blinded bids. You will get a refund for all
    /// correctly blinded invalid bids and for all bids except for
    /// the totally highest.
    function reveal(
        uint[] memory _values,
        bool[] memory _fake,
        bytes32[] memory _secret
    )
        public
        onlyAfter(biddingEnd)
        onlyBefore(revealEnd)
    {
        uint length = bids[msg.sender].length;
        require(_values.length == length);
        require(_fake.length == length);
        require(_secret.length == length);

        uint refund;
        for (uint i = 0; i < length; i++) {
            Bid storage bidToCheck = bids[msg.sender][i];
            (uint value, bool fake, bytes32 secret) =
                    (_values[i], _fake[i], _secret[i]);
            if (bidToCheck.blindedBid != keccak256(abi.encodePacked(value, fake,
→secret))) {
                // Bid was not actually revealed.
                // Do not refund deposit.
                continue;
            }
            refund += bidToCheck.deposit;
            if (!fake && bidToCheck.deposit >= value) {
                if (placeBid(msg.sender, value))
                    refund -= value;
            }
            // Make it impossible for the sender to re-claim
            // the same deposit.
            bidToCheck.blindedBid = bytes32(0);
        }
        msg.sender.transfer(refund);
    }

    /// Withdraw a bid that was overbid.
    function withdraw() public {
        uint amount = pendingReturns[msg.sender];
```

```solidity
        if (amount > 0) {
            // It is important to set this to zero because the recipient
            // can call this function again as part of the receiving call
            // before `transfer` returns (see the remark above about
            // conditions -> effects -> interaction).
            pendingReturns[msg.sender] = 0;

            msg.sender.transfer(amount);
        }
    }

    /// End the auction and send the highest bid
    /// to the beneficiary.
    function auctionEnd()
        public
        onlyAfter(revealEnd)
    {
        require(!ended);
        emit AuctionEnded(highestBidder, highestBid);
        ended = true;
        beneficiary.transfer(highestBid);
    }

    // This is an "internal" function which means that it
    // can only be called from the contract itself (or from
    // derived contracts).
    function placeBid(address bidder, uint value) internal
            returns (bool success)
    {
        if (value <= highestBid) {
            return false;
        }
        if (highestBidder != address(0)) {
            // Refund the previously highest bidder.
            pendingReturns[highestBidder] += highestBid;
        }
        highestBid = value;
        highestBidder = bidder;
        return true;
    }
}
```

3.3.3 Safe Remote Purchase

Purchasing goods remotely currently requires multiple parties that need to trust each other. The simplest configuration involves a seller and a buyer. The buyer would like to receive an item from the seller and the seller would like to get money (or an equivalent) in return. The problematic part is the shipment here: There is no way to determine for sure that the item arrived at the buyer.

There are multiple ways to solve this problem, but all fall short in one or the other way. In the following example, both parties have to put twice the value of the item into the contract as escrow. As soon as this happened, the money will stay locked inside the contract until the buyer confirms that they received the item. After that, the buyer is returned the value (half of their deposit) and the seller gets three times the value (their deposit plus the value). The idea behind this is that both parties have an incentive to resolve the situation or otherwise their money is locked forever.

This contract of course does not solve the problem, but gives an overview of how you can use state machine-like

constructs inside a contract.

```solidity
// SPDX-License-Identifier: GPL-3.0
pragma solidity ^0.7.0;

contract Purchase {
    uint public value;
    address payable public seller;
    address payable public buyer;

    enum State { Created, Locked, Release, Inactive }
    // The state variable has a default value of the first member, `State.created`
    State public state;

    modifier condition(bool _condition) {
        require(_condition);
        _;
    }

    modifier onlyBuyer() {
        require(
            msg.sender == buyer,
            "Only buyer can call this."
        );
        _;
    }

    modifier onlySeller() {
        require(
            msg.sender == seller,
            "Only seller can call this."
        );
        _;
    }

    modifier inState(State _state) {
        require(
            state == _state,
            "Invalid state."
        );
        _;
    }

    event Aborted();
    event PurchaseConfirmed();
    event ItemReceived();
    event SellerRefunded();

    // Ensure that `msg.value` is an even number.
    // Division will truncate if it is an odd number.
    // Check via multiplication that it wasn't an odd number.
    constructor() payable {
        seller = msg.sender;
        value = msg.value / 2;
        require((2 * value) == msg.value, "Value has to be even.");
    }

    /// Abort the purchase and reclaim the ether.
```

3.3. Solidity by Example

```solidity
    /// Can only be called by the seller before
    /// the contract is locked.
    function abort()
        public
        onlySeller
        inState(State.Created)
    {
        emit Aborted();
        state = State.Inactive;
        // We use transfer here directly. It is
        // reentrancy-safe, because it is the
        // last call in this function and we
        // already changed the state.
        seller.transfer(address(this).balance);
    }

    /// Confirm the purchase as buyer.
    /// Transaction has to include `2 * value` ether.
    /// The ether will be locked until confirmReceived
    /// is called.
    function confirmPurchase()
        public
        inState(State.Created)
        condition(msg.value == (2 * value))
        payable
    {
        emit PurchaseConfirmed();
        buyer = msg.sender;
        state = State.Locked;
    }

    /// Confirm that you (the buyer) received the item.
    /// This will release the locked ether.
    function confirmReceived()
        public
        onlyBuyer
        inState(State.Locked)
    {
        emit ItemReceived();
        // It is important to change the state first because
        // otherwise, the contracts called using `send` below
        // can call in again here.
        state = State.Release;

        buyer.transfer(value);
    }

    /// This function refunds the seller, i.e.
    /// pays back the locked funds of the seller.
    function refundSeller()
        public
        onlySeller
        inState(State.Release)
    {
        emit SellerRefunded();
        // It is important to change the state first because
        // otherwise, the contracts called using `send` below
```

```
        // can call in again here.
        state = State.Inactive;

        seller.transfer(3 * value);
    }
}
```

3.3.4 Micropayment Channel

In this section we will learn how to build an example implementation of a payment channel. It uses cryptographic signatures to make repeated transfers of Ether between the same parties secure, instantaneous, and without transaction fees. For the example, we need to understand how to sign and verify signatures, and setup the payment channel.

Creating and verifying signatures

Imagine Alice wants to send a quantity of Ether to Bob, i.e. Alice is the sender and the Bob is the recipient.

Alice only needs to send cryptographically signed messages off-chain (e.g. via email) to Bob and it is similar to writing checks.

Alice and Bob use signatures to authorise transactions, which is possible with smart contracts on Ethereum. Alice will build a simple smart contract that lets her transmit Ether, but instead of calling a function herself to initiate a payment, she will let Bob do that, and therefore pay the transaction fee.

The contract will work as follows:

1. Alice deploys the `ReceiverPays` contract, attaching enough Ether to cover the payments that will be made.

2. Alice authorises a payment by signing a message with their private key.

3. Alice sends the cryptographically signed message to Bob. The message does not need to be kept secret (explained later), and the mechanism for sending it does not matter.

4. Bob claims their payment by presenting the signed message to the smart contract, it verifies the authenticity of the message and then releases the funds.

Creating the signature

Alice does not need to interact with the Ethereum network to sign the transaction, the process is completely offline. In this tutorial, we will sign messages in the browser using web3.js and MetaMask, using the method described in EIP-762, as it provides a number of other security benefits.

```
/// Hashing first makes things easier
var hash = web3.utils.sha3("message to sign");
web3.eth.personal.sign(hash, web3.eth.defaultAccount, function () { console.log(
↪"Signed"); });
```

Note: The `web3.eth.personal.sign` prepends the length of the message to the signed data. Since we hash first, the message will always be exactly 32 bytes long, and thus this length prefix is always the same.

What to Sign

For a contract that fulfils payments, the signed message must include:

1. The recipient's address.

2. The amount to be transferred.

3. Protection against replay attacks.

A replay attack is when a signed message is reused to claim authorization for a second action. To avoid replay attacks we use the same technique as in Ethereum transactions themselves, a so-called nonce, which is the number of transactions sent by an account. The smart contract checks if a nonce is used multiple times.

Another type of replay attack can occur when the owner deploys a `ReceiverPays` smart contract, makes some payments, and then destroys the contract. Later, they decide to deploy the `RecipientPays` smart contract again, but the new contract does not know the nonces used in the previous deployment, so the attacker can use the old messages again.

Alice can protect against this attack by including the contract's address in the message, and only messages containing the contract's address itself will be accepted. You can find an example of this in the first two lines of the `claimPayment()` function of the full contract at the end of this section.

Packing arguments

Now that we have identified what information to include in the signed message, we are ready to put the message together, hash it, and sign it. For simplicity, we concatenate the data. The ethereumjs-abi library provides a function called `soliditySHA3` that mimics the behaviour of Solidity's `keccak256` function applied to arguments encoded using `abi.encodePacked`. Here is a JavaScript function that creates the proper signature for the `ReceiverPays` example:

```
// recipient is the address that should be paid.
// amount, in wei, specifies how much ether should be sent.
// nonce can be any unique number to prevent replay attacks
// contractAddress is used to prevent cross-contract replay attacks
function signPayment(recipient, amount, nonce, contractAddress, callback) {
    var hash = "0x" + abi.soliditySHA3(
        ["address", "uint256", "uint256", "address"],
        [recipient, amount, nonce, contractAddress]
    ).toString("hex");

    web3.eth.personal.sign(hash, web3.eth.defaultAccount, callback);
}
```

Recovering the Message Signer in Solidity

In general, ECDSA signatures consist of two parameters, `r` and `s`. Signatures in Ethereum include a third parameter called `v`, that you can use to verify which account's private key was used to sign the message, and the transaction's sender. Solidity provides a built-in function *ecrecover* that accepts a message along with the `r`, `s` and `v` parameters and returns the address that was used to sign the message.

Extracting the Signature Parameters

Signatures produced by web3.js are the concatenation of `r`, `s` and `v`, so the first step is to split these parameters apart. You can do this on the client-side, but doing it inside the smart contract means you only need to send one signature

parameter rather than three. Splitting apart a byte array into its constituent parts is a mess, so we use *inline assembly* to do the job in the `splitSignature` function (the third function in the full contract at the end of this section).

Computing the Message Hash

The smart contract needs to know exactly what parameters were signed, and so it must recreate the message from the parameters and use that for signature verification. The functions `prefixed` and `recoverSigner` do this in the `claimPayment` function.

The full contract

```solidity
// SPDX-License-Identifier: GPL-3.0
pragma solidity ^0.7.0;

contract ReceiverPays {
    address owner = msg.sender;

    mapping(uint256 => bool) usedNonces;

    constructor() payable {}

    function claimPayment(uint256 amount, uint256 nonce, bytes memory signature)
→public {
        require(!usedNonces[nonce]);
        usedNonces[nonce] = true;

        // this recreates the message that was signed on the client
        bytes32 message = prefixed(keccak256(abi.encodePacked(msg.sender, amount,
→nonce, this)));

        require(recoverSigner(message, signature) == owner);

        msg.sender.transfer(amount);
    }

    /// destroy the contract and reclaim the leftover funds.
    function shutdown() public {
        require(msg.sender == owner);
        selfdestruct(msg.sender);
    }

    /// signature methods.
    function splitSignature(bytes memory sig)
        internal
        pure
        returns (uint8 v, bytes32 r, bytes32 s)
    {
        require(sig.length == 65);

        assembly {
            // first 32 bytes, after the length prefix.
            r := mload(add(sig, 32))
            // second 32 bytes.
            s := mload(add(sig, 64))
            // final byte (first byte of the next 32 bytes).
```

(continues on next page)

```
            v := byte(0, mload(add(sig, 96)))
        }

        return (v, r, s);
    }

    function recoverSigner(bytes32 message, bytes memory sig)
        internal
        pure
        returns (address)
    {
        (uint8 v, bytes32 r, bytes32 s) = splitSignature(sig);

        return ecrecover(message, v, r, s);
    }

    /// builds a prefixed hash to mimic the behavior of eth_sign.
    function prefixed(bytes32 hash) internal pure returns (bytes32) {
        return keccak256(abi.encodePacked("\x19Ethereum Signed Message:\n32", hash));
    }
}
```

Writing a Simple Payment Channel

Alice now builds a simple but complete implementation of a payment channel. Payment channels use cryptographic signatures to make repeated transfers of Ether securely, instantaneously, and without transaction fees.

What is a Payment Channel?

Payment channels allow participants to make repeated transfers of Ether without using transactions. This means that you can avoid the delays and fees associated with transactions. We are going to explore a simple unidirectional payment channel between two parties (Alice and Bob). It involves three steps:

1. Alice funds a smart contract with Ether. This "opens" the payment channel.

2. Alice signs messages that specify how much of that Ether is owed to the recipient. This step is repeated for each payment.

3. Bob "closes" the payment channel, withdrawing their portion of the Ether and sending the remainder back to the sender.

Note: Only steps 1 and 3 require Ethereum transactions, step 2 means that the sender transmits a cryptographically signed message to the recipient via off chain methods (e.g. email). This means only two transactions are required to support any number of transfers.

Bob is guaranteed to receive their funds because the smart contract escrows the Ether and honours a valid signed message. The smart contract also enforces a timeout, so Alice is guaranteed to eventually recover their funds even if the recipient refuses to close the channel. It is up to the participants in a payment channel to decide how long to keep it open. For a short-lived transaction, such as paying an internet café for each minute of network access, the payment channel may be kept open for a limited duration. On the other hand, for a recurring payment, such as paying an employee an hourly wage, the payment channel may be kept open for several months or years.

Opening the Payment Channel

To open the payment channel, Alice deploys the smart contract, attaching the Ether to be escrowed and specifying the intended recipient and a maximum duration for the channel to exist. This is the function `SimplePaymentChannel` in the contract, at the end of this section.

Making Payments

Alice makes payments by sending signed messages to Bob. This step is performed entirely outside of the Ethereum network. Messages are cryptographically signed by the sender and then transmitted directly to the recipient.

Each message includes the following information:

- The smart contract's address, used to prevent cross-contract replay attacks.

- The total amount of Ether that is owed the recipient so far.

A payment channel is closed just once, at the end of a series of transfers. Because of this, only one of the messages sent is redeemed. This is why each message specifies a cumulative total amount of Ether owed, rather than the amount of the individual micropayment. The recipient will naturally choose to redeem the most recent message because that is the one with the highest total. The nonce per-message is not needed anymore, because the smart contract only honours a single message. The address of the smart contract is still used to prevent a message intended for one payment channel from being used for a different channel.

Here is the modified JavaScript code to cryptographically sign a message from the previous section:

```javascript
function constructPaymentMessage(contractAddress, amount) {
    return abi.soliditySHA3(
        ["address", "uint256"],
        [contractAddress, amount]
    );
}

function signMessage(message, callback) {
    web3.eth.personal.sign(
        "0x" + message.toString("hex"),
        web3.eth.defaultAccount,
        callback
    );
}

// contractAddress is used to prevent cross-contract replay attacks.
// amount, in wei, specifies how much Ether should be sent.

function signPayment(contractAddress, amount, callback) {
    var message = constructPaymentMessage(contractAddress, amount);
    signMessage(message, callback);
}
```

Closing the Payment Channel

When Bob is ready to receive their funds, it is time to close the payment channel by calling a `close` function on the smart contract. Closing the channel pays the recipient the Ether they are owed and destroys the contract, sending any remaining Ether back to Alice. To close the channel, Bob needs to provide a message signed by Alice.

The smart contract must verify that the message contains a valid signature from the sender. The process for doing this verification is the same as the process the recipient uses. The Solidity functions `isValidSignature` and `recoverSigner` work just like their JavaScript counterparts in the previous section, with the latter function borrowed from the `ReceiverPays` contract.

Only the payment channel recipient can call the `close` function, who naturally passes the most recent payment message because that message carries the highest total owed. If the sender were allowed to call this function, they could provide a message with a lower amount and cheat the recipient out of what they are owed.

The function verifies the signed message matches the given parameters. If everything checks out, the recipient is sent their portion of the Ether, and the sender is sent the rest via a `selfdestruct`. You can see the `close` function in the full contract.

Channel Expiration

Bob can close the payment channel at any time, but if they fail to do so, Alice needs a way to recover their escrowed funds. An *expiration* time was set at the time of contract deployment. Once that time is reached, Alice can call `claimTimeout` to recover their funds. You can see the `claimTimeout` function in the full contract.

After this function is called, Bob can no longer receive any Ether, so it is important that Bob closes the channel before the expiration is reached.

The full contract

```solidity
// SPDX-License-Identifier: GPL-3.0
pragma solidity ^0.7.0;

contract SimplePaymentChannel {
    address payable public sender;      // The account sending payments.
    address payable public recipient;   // The account receiving the payments.
    uint256 public expiration;  // Timeout in case the recipient never closes.

    constructor (address payable _recipient, uint256 duration)
        payable
    {
        sender = msg.sender;
        recipient = _recipient;
        expiration = block.timestamp + duration;
    }

    /// the recipient can close the channel at any time by presenting a
    /// signed amount from the sender. the recipient will be sent that amount,
    /// and the remainder will go back to the sender
    function close(uint256 amount, bytes memory signature) public {
        require(msg.sender == recipient);
        require(isValidSignature(amount, signature));

        recipient.transfer(amount);
        selfdestruct(sender);
    }

    /// the sender can extend the expiration at any time
    function extend(uint256 newExpiration) public {
        require(msg.sender == sender);
        require(newExpiration > expiration);
```

(continues on next page)

```solidity
        expiration = newExpiration;
    }

    /// if the timeout is reached without the recipient closing the channel,
    /// then the Ether is released back to the sender.
    function claimTimeout() public {
        require(block.timestamp >= expiration);
        selfdestruct(sender);
    }

    function isValidSignature(uint256 amount, bytes memory signature)
        internal
        view
        returns (bool)
    {
        bytes32 message = prefixed(keccak256(abi.encodePacked(this, amount)));

        // check that the signature is from the payment sender
        return recoverSigner(message, signature) == sender;
    }

    /// All functions below this are just taken from the chapter
    /// 'creating and verifying signatures' chapter.

    function splitSignature(bytes memory sig)
        internal
        pure
        returns (uint8 v, bytes32 r, bytes32 s)
    {
        require(sig.length == 65);

        assembly {
            // first 32 bytes, after the length prefix
            r := mload(add(sig, 32))
            // second 32 bytes
            s := mload(add(sig, 64))
            // final byte (first byte of the next 32 bytes)
            v := byte(0, mload(add(sig, 96)))
        }

        return (v, r, s);
    }

    function recoverSigner(bytes32 message, bytes memory sig)
        internal
        pure
        returns (address)
    {
        (uint8 v, bytes32 r, bytes32 s) = splitSignature(sig);

        return ecrecover(message, v, r, s);
    }

    /// builds a prefixed hash to mimic the behavior of eth_sign.
    function prefixed(bytes32 hash) internal pure returns (bytes32) {
        return keccak256(abi.encodePacked("\x19Ethereum Signed Message:\n32", hash));
```

```
    }
}
```

Note: The function `splitSignature` does not use all security checks. A real implementation should use a more rigorously tested library, such as openzepplin's version of this code.

Verifying Payments

Unlike in the previous section, messages in a payment channel aren't redeemed right away. The recipient keeps track of the latest message and redeems it when it's time to close the payment channel. This means it's critical that the recipient perform their own verification of each message. Otherwise there is no guarantee that the recipient will be able to get paid in the end.

The recipient should verify each message using the following process:

1. Verify that the contact address in the message matches the payment channel.

2. Verify that the new total is the expected amount.

3. Verify that the new total does not exceed the amount of Ether escrowed.

4. Verify that the signature is valid and comes from the payment channel sender.

We'll use the ethereumjs-util library to write this verification. The final step can be done a number of ways, and we use JavaScript. The following code borrows the `constructMessage` function from the signing **JavaScript code** above:

```javascript
// this mimics the prefixing behavior of the eth_sign JSON-RPC method.
function prefixed(hash) {
    return ethereumjs.ABI.soliditySHA3(
        ["string", "bytes32"],
        ["\x19Ethereum Signed Message:\n32", hash]
    );
}

function recoverSigner(message, signature) {
    var split = ethereumjs.Util.fromRpcSig(signature);
    var publicKey = ethereumjs.Util.ecrecover(message, split.v, split.r, split.s);
    var signer = ethereumjs.Util.pubToAddress(publicKey).toString("hex");
    return signer;
}

function isValidSignature(contractAddress, amount, signature, expectedSigner) {
    var message = prefixed(constructPaymentMessage(contractAddress, amount));
    var signer = recoverSigner(message, signature);
    return signer.toLowerCase() ==
        ethereumjs.Util.stripHexPrefix(expectedSigner).toLowerCase();
}
```

3.3.5 Modular Contracts

A modular approach to building your contracts helps you reduce the complexity and improve the readability which will help to identify bugs and vulnerabilities during development and code review. If you specify and control the

behaviour or each module in isolation, the interactions you have to consider are only those between the module specifications and not every other moving part of the contract. In the example below, the contract uses the `move` method of the `Balances` *library* to check that balances sent between addresses match what you expect. In this way, the `Balances` library provides an isolated component that properly tracks balances of accounts. It is easy to verify that the `Balances` library never produces negative balances or overflows and the sum of all balances is an invariant across the lifetime of the contract.

```solidity
// SPDX-License-Identifier: GPL-3.0
pragma solidity >=0.5.0 <0.8.0;

library Balances {
    function move(mapping(address => uint256) storage balances, address from, address
↪to, uint amount) internal {
        require(balances[from] >= amount);
        require(balances[to] + amount >= balances[to]);
        balances[from] -= amount;
        balances[to] += amount;
    }
}

contract Token {
    mapping(address => uint256) balances;
    using Balances for *;
    mapping(address => mapping (address => uint256)) allowed;

    event Transfer(address from, address to, uint amount);
    event Approval(address owner, address spender, uint amount);

    function transfer(address to, uint amount) public returns (bool success) {
        balances.move(msg.sender, to, amount);
        emit Transfer(msg.sender, to, amount);
        return true;

    }

    function transferFrom(address from, address to, uint amount) public returns (bool
↪success) {
        require(allowed[from][msg.sender] >= amount);
        allowed[from][msg.sender] -= amount;
        balances.move(from, to, amount);
        emit Transfer(from, to, amount);
        return true;
    }

    function approve(address spender, uint tokens) public returns (bool success) {
        require(allowed[msg.sender][spender] == 0, "");
        allowed[msg.sender][spender] = tokens;
        emit Approval(msg.sender, spender, tokens);
        return true;
    }

    function balanceOf(address tokenOwner) public view returns (uint balance) {
        return balances[tokenOwner];
    }
}
```

3.4 Layout of a Solidity Source File

Source files can contain an arbitrary number of *contract definitions*, *import* directives, *pragma directives* and *struct*, *enum*, *function* and *constant variable* definitions.

3.4.1 SPDX License Identifier

Trust in smart contract can be better established if their source code is available. Since making source code available always touches on legal problems with regards to copyright, the Solidity compiler encourages the use of machine-readable SPDX license identifiers. Every source file should start with a comment indicating its license:

```
// SPDX-License-Identifier:  MIT
```

The compiler does not validate that the license is part of the list allowed by SPDX, but it does include the supplied string in the *bytecode metadata*.

If you do not want to specify a license or if the source code is not open-source, please use the special value `UNLICENSED`.

Supplying this comment of course does not free you from other obligations related to licensing like having to mention a specific license header in each source file or the original copyright holder.

The comment is recognized by the compiler anywhere in the file at the file level, but it is recommended to put it at the top of the file.

More information about how to use SPDX license identifiers can be found at the SPDX website.

3.4.2 Pragmas

The `pragma` keyword is used to enable certain compiler features or checks. A pragma directive is always local to a source file, so you have to add the pragma to all your files if you want enable it in your whole project. If you *import* another file, the pragma from that file does *not* automatically apply to the importing file.

Version Pragma

Source files can (and should) be annotated with a version pragma to reject compilation with future compiler versions that might introduce incompatible changes. We try to keep these to an absolute minimum and introduce them in a way that changes in semantics also require changes in the syntax, but this is not always possible. Because of this, it is always a good idea to read through the changelog at least for releases that contain breaking changes. These releases always have versions of the form `0.x.0` or `x.0.0`.

The version pragma is used as follows: `pragma solidity ^0.5.2;`

A source file with the line above does not compile with a compiler earlier than version 0.5.2, and it also does not work on a compiler starting from version 0.6.0 (this second condition is added by using `^`). Because there will be no breaking changes until version `0.6.0`, you can be sure that your code compiles the way you intended. The exact version of the compiler is not fixed, so that bugfix releases are still possible.

It is possible to specify more complex rules for the compiler version, these follow the same syntax used by npm.

Note: Using the version pragma *does not* change the version of the compiler. It also *does not* enable or disable features of the compiler. It just instructs the compiler to check whether its version matches the one required by the pragma. If it does not match, the compiler issues an error.

ABI Coder Pragma

By using `pragma abicoder v1` or `pragma abicoder v2` you can select between the two implementations of the ABI encoder and decoder.

The new ABI coder (v2) is able to encode and decode arbitrarily nested arrays and structs. It might produce less optimal code and has not received as much testing as the old encoder, but is considered non-experimental as of Solidity 0.6.0. You still have to explicitly activate it using `pragma abicoder v2;`. Since it will be activated by default starting from Solidity 0.8.0, there is the option to select the old coder using `pragma abicoder v1;`.

The set of types supported by the new encoder is a strict superset of the ones supported by the old one. Contracts that use it can interact with ones that do not without limitations. The reverse is possible only as long as the non-`abicoder v2` contract does not try to make calls that would require decoding types only supported by the new encoder. The compiler can detect this and will issue an error. Simply enabling `abicoder v2` for your contract is enough to make the error go away.

Note: This pragma applies to all the code defined in the file where it is activated, regardless of where that code ends up eventually. This means that a contract whose source file is selected to compile with ABI coder v1 can still contain code that uses the new encoder by inheriting it from another contract. This is allowed if the new types are only used internally and not in external function signatures.

Note: Up to Solidity 0.7.4, it was possible to select the ABI coder v2 by using `pragma experimental ABIEncoderV2`, but it was not possible to explicitly select coder v1 because it was the default.

Experimental Pragma

The second pragma is the experimental pragma. It can be used to enable features of the compiler or language that are not yet enabled by default. The following experimental pragmas are currently supported:

ABIEncoderV2

Because the ABI coder v2 is not considered experimental anymore, it can be selected via `pragma abicoder v2` (please see above) since Solidity 0.7.4.

SMTChecker

This component has to be enabled when the Solidity compiler is built and therefore it is not available in all Solidity binaries. The *build instructions* explain how to activate this option. It is activated for the Ubuntu PPA releases in most versions, but not for the Docker images, Windows binaries or the statically-built Linux binaries. It can be activated for solc-js via the smtCallback if you have an SMT solver installed locally and run solc-js via node (not via the browser).

If you use `pragma experimental SMTChecker;`, then you get additional *safety warnings* which are obtained by querying an SMT solver. The component does not yet support all features of the Solidity language and likely outputs many warnings. In case it reports unsupported features, the analysis may not be fully sound.

3.4.3 Importing other Source Files

Syntax and Semantics

Solidity supports import statements to help modularise your code that are similar to those available in JavaScript (from ES6 on). However, Solidity does not support the concept of a default export.

At a global level, you can use import statements of the following form:

```
import "filename";
```

This statement imports all global symbols from "filename" (and symbols imported there) into the current global scope (different than in ES6 but backwards-compatible for Solidity). This form is not recommended for use, because it unpredictably pollutes the namespace. If you add new top-level items inside "filename", they automatically appear in all files that import like this from "filename". It is better to import specific symbols explicitly.

The following example creates a new global symbol symbolName whose members are all the global symbols from "filename":

```
import * as symbolName from "filename";
```

which results in all global symbols being available in the format symbolName.symbol.

A variant of this syntax that is not part of ES6, but possibly useful is:

```
import "filename" as symbolName;
```

which is equivalent to import * as symbolName from "filename";.

If there is a naming collision, you can rename symbols while importing. For example, the code below creates new global symbols alias and symbol2 which reference symbol1 and symbol2 from inside "filename", respectively.

```
import {symbol1 as alias, symbol2} from "filename";
```

Paths

In the above, filename is always treated as a path with / as directory separator, and . as the current and .. as the parent directory. When . or .. is followed by a character except /, it is not considered as the current or the parent directory. All path names are treated as absolute paths unless they start with the current . or the parent directory ...

To import a file filename from the same directory as the current file, use import "./filename" as symbolName;. If you use import "filename" as symbolName; instead, a different file could be referenced (in a global "include directory").

It depends on the compiler (see *Use in Actual Compilers*) how to actually resolve the paths. In general, the directory hierarchy does not need to strictly map onto your local filesystem, and the path can also map to resources such as ipfs, http or git.

Note: Always use relative imports like import "./filename.sol"; and avoid using .. in path specifiers. In the latter case, it is probably better to use global paths and set up remappings as explained below.

Use in Actual Compilers

When invoking the compiler, you can specify how to discover the first element of a path, and also path prefix remappings. For example you can setup a remapping so that everything imported from the virtual directory github.com/ethereum/dapp-bin/library would actually be read from your local directory /usr/local/dapp-bin/

library. If multiple remappings apply, the one with the longest key is tried first. An empty prefix is not allowed. The remappings can depend on a context, which allows you to configure packages to import e.g., different versions of a library of the same name.

solc:

For solc (the commandline compiler), you provide these path remappings as context:prefix=target arguments, where both the context: and the =target parts are optional (target defaults to prefix in this case). All remapping values that are regular files are compiled (including their dependencies).

This mechanism is backwards-compatible (as long as no filename contains = or :) and thus not a breaking change. All files in or below the context directory that import a file that starts with prefix are redirected by replacing prefix by target.

For example, if you clone github.com/ethereum/dapp-bin/ locally to /usr/local/dapp-bin, you can use the following in your source file:

```
import "github.com/ethereum/dapp-bin/library/iterable_mapping.sol" as it_mapping;
```

Then run the compiler:

```
solc github.com/ethereum/dapp-bin/=/usr/local/dapp-bin/ source.sol
```

As a more complex example, suppose you rely on a module that uses an old version of dapp-bin that you checked out to /usr/local/dapp-bin_old, then you can run:

```
solc module1:github.com/ethereum/dapp-bin/=/usr/local/dapp-bin/ \
     module2:github.com/ethereum/dapp-bin/=/usr/local/dapp-bin_old/ \
     source.sol
```

This means that all imports in module2 point to the old version but imports in module1 point to the new version.

Note: solc only allows you to include files from certain directories. They have to be in the directory (or subdirectory) of one of the explicitly specified source files or in the directory (or subdirectory) of a remapping target. If you want to allow direct absolute includes, add the remapping /=/.

If there are multiple remappings that lead to a valid file, the remapping with the longest common prefix is chosen.

Remix:

Remix provides an automatic remapping for GitHub and automatically retrieves the file over the network. You can import the iterable mapping as above, e.g.

```
import "github.com/ethereum/dapp-bin/library/iterable_mapping.sol" as it_mapping;
```

Remix may add other source code providers in the future.

3.4.4 Comments

Single-line comments (//) and multi-line comments (/*...*/) are possible.

```
// This is a single-line comment.

/*
This is a
multi-line comment.
*/
```

Note: A single-line comment is terminated by any unicode line terminator (LF, VF, FF, CR, NEL, LS or PS) in utf8 encoding. The terminator is still part of the source code after the comment, so if it is not an ascii symbol (these are NEL, LS and PS), it will lead to a parser error.

Additionally, there is another type of comment called a natspec comment, which is detailed in the *style guide*. They are written with a triple slash (///) or a double asterisk block(/** ... */) and they should be used directly above function declarations or statements. You can use Doxygen-style tags inside these comments to document functions, annotate conditions for formal verification, and provide a **confirmation text** which is shown to users when they attempt to invoke a function.

In the following example we document the title of the contract, the explanation for the two function parameters and two return variables.

```solidity
// SPDX-License-Identifier: GPL-3.0
pragma solidity >=0.4.21 <0.8.0;

/** @title Shape calculator. */
contract ShapeCalculator {
    /// @dev Calculates a rectangle's surface and perimeter.
    /// @param w Width of the rectangle.
    /// @param h Height of the rectangle.
    /// @return s The calculated surface.
    /// @return p The calculated perimeter.
    function rectangle(uint w, uint h) public pure returns (uint s, uint p) {
        s = w * h;
        p = 2 * (w + h);
    }
}
```

3.5 Structure of a Contract

Contracts in Solidity are similar to classes in object-oriented languages. Each contract can contain declarations of *State Variables*, *Functions*, *Function Modifiers*, *Events*, *Struct Types* and *Enum Types*. Furthermore, contracts can inherit from other contracts.

There are also special kinds of contracts called *libraries* and *interfaces*.

The section about *contracts* contains more details than this section, which serves to provide a quick overview.

3.5.1 State Variables

State variables are variables whose values are permanently stored in contract storage.

```solidity
// SPDX-License-Identifier: GPL-3.0
pragma solidity >=0.4.0 <0.8.0;

contract SimpleStorage {
    uint storedData; // State variable
    // ...
}
```

See the *Types* section for valid state variable types and *Visibility and Getters* for possible choices for visibility.

3.5.2 Functions

Functions are the executable units of code. Functions are usually defined inside a contract, but they can also be defined outside of contracts.

```solidity
// SPDX-License-Identifier: GPL-3.0
pragma solidity >0.7.0 <0.8.0;

contract SimpleAuction {
    function bid() public payable { // Function
        // ...
    }
}

// Helper function defined outside of a contract
function helper(uint x) pure returns (uint) {
    return x * 2;
}
```

Function Calls can happen internally or externally and have different levels of *visibility* towards other contracts. *Functions* accept *parameters and return variables* to pass parameters and values between them.

3.5.3 Function Modifiers

Function modifiers can be used to amend the semantics of functions in a declarative way (see *Function Modifiers* in the contracts section).

Overloading, that is, having the same modifier name with different parameters, is not possible.

Like functions, modifiers can be *overridden*.

```solidity
// SPDX-License-Identifier: GPL-3.0
pragma solidity >=0.4.22 <0.8.0;

contract Purchase {
    address public seller;

    modifier onlySeller() { // Modifier
        require(
            msg.sender == seller,
            "Only seller can call this."
        );
        _;
    }

    function abort() public view onlySeller { // Modifier usage
        // ...
    }
}
```

3.5.4 Events

Events are convenience interfaces with the EVM logging facilities.

```
// SPDX-License-Identifier: GPL-3.0
pragma solidity >=0.4.21 <0.8.0;

contract SimpleAuction {
    event HighestBidIncreased(address bidder, uint amount); // Event

    function bid() public payable {
        // ...
        emit HighestBidIncreased(msg.sender, msg.value); // Triggering event
    }
}
```

See *Events* in contracts section for information on how events are declared and can be used from within a dapp.

3.5.5 Struct Types

Structs are custom defined types that can group several variables (see *Structs* in types section).

```
// SPDX-License-Identifier: GPL-3.0
pragma solidity >=0.4.0 <0.8.0;

contract Ballot {
    struct Voter { // Struct
        uint weight;
        bool voted;
        address delegate;
        uint vote;
    }
}
```

3.5.6 Enum Types

Enums can be used to create custom types with a finite set of 'constant values' (see *Enums* in types section).

```
// SPDX-License-Identifier: GPL-3.0
pragma solidity >=0.4.0 <0.8.0;

contract Purchase {
    enum State { Created, Locked, Inactive } // Enum
}
```

3.6 Types

Solidity is a statically typed language, which means that the type of each variable (state and local) needs to be specified. Solidity provides several elementary types which can be combined to form complex types.

In addition, types can interact with each other in expressions containing operators. For a quick reference of the various operators, see *Order of Precedence of Operators*.

The concept of "undefined" or "null" values does not exist in Solidity, but newly declared variables always have a *default value* dependent on its type. To handle any unexpected values, you should use the *revert function* to revert the whole transaction, or return a tuple with a second `bool` value denoting success.

3.6.1 Value Types

The following types are also called value types because variables of these types will always be passed by value, i.e. they are always copied when they are used as function arguments or in assignments.

Booleans

`bool`: The possible values are constants `true` and `false`.

Operators:

- `!` (logical negation)

- `&&` (logical conjunction, "and")

- `||` (logical disjunction, "or")

- `==` (equality)

- `!=` (inequality)

The operators `||` and `&&` apply the common short-circuiting rules. This means that in the expression `f(x) || g(y)`, if `f(x)` evaluates to `true`, `g(y)` will not be evaluated even if it may have side-effects.

Integers

`int` / `uint`: Signed and unsigned integers of various sizes. Keywords `uint8` to `uint256` in steps of `8` (unsigned of 8 up to 256 bits) and `int8` to `int256`. `uint` and `int` are aliases for `uint256` and `int256`, respectively.

Operators:

- Comparisons: `<=`, `<`, `==`, `!=`, `>=`, `>` (evaluate to `bool`)

- Bit operators: `&`, `|`, `^` (bitwise exclusive or), `~` (bitwise negation)

- Shift operators: `<<` (left shift), `>>` (right shift)

- Arithmetic operators: `+`, `-`, unary `-`, `*`, `/`, `%` (modulo), `**` (exponentiation)

For an integer type X, you can use `type(X).min` and `type(X).max` to access the minimum and maximum value representable by the type.

> **Warning:** Integers in Solidity are restricted to a certain range. For example, with `uint32`, this is `0` up to `2**32 - 1`. If the result of some operation on those numbers does not fit inside this range, it is truncated. These truncations can have serious consequences that you should *be aware of and mitigate against*.

Comparisons

The value of a comparison is the one obtained by comparing the integer value.

Bit operations

Bit operations are performed on the two's complement representation of the number. This means that, for example `~int256(0) == int256(-1)`.

Shifts

The result of a shift operation has the type of the left operand, truncating the result to match the type. Right operand must be unsigned type. Trying to shift by signed type will produce a compilation error.

- For positive and negative x values, `x << y` is equivalent to `x * 2**y`.

- For positive x values, `x >> y` is equivalent to `x / 2**y`.

- For negative x values, `x >> y` is equivalent to `(x + 1) / 2**y - 1` (which is the same as dividing x by `2**y` while rounding down towards negative infinity).

> **Warning:** Before version `0.5.0` a right shift `x >> y` for negative x was equivalent to `x / 2**y`, i.e., right shifts used rounding up (towards zero) instead of rounding down (towards negative infinity).

Addition, Subtraction and Multiplication

Addition, subtraction and multiplication have the usual semantics. They wrap in two's complement representation, meaning that for example `uint256(0) - uint256(1) == 2**256 - 1`. You have to take these overflows into account when designing safe smart contracts.

The expression `-x` is equivalent to `(T(0) - x)` where `T` is the type of x. This means that `-x` will not be negative if the type of x is an unsigned integer type. Also, `-x` can be positive if x is negative. There is another caveat also resulting from two's complement representation:

```
int x = -2**255;
assert(-x == x);
```

This means that even if a number is negative, you cannot assume that its negation will be positive.

Division

Since the type of the result of an operation is always the type of one of the operands, division on integers always results in an integer. In Solidity, division rounds towards zero. This mean that `int256(-5) / int256(2) == int256(-2)`.

Note that in contrast, division on *literals* results in fractional values of arbitrary precision.

> **Note:** Division by zero causes a failing assert.

Modulo

The modulo operation `a % n` yields the remainder r after the division of the operand a by the operand n, where q `= int(a / n)` and `r = a - (n * q)`. This means that modulo results in the same sign as its left operand (or zero) and `a % n == -(-a % n)` holds for negative a:

- `int256(5) % int256(2) == int256(1)`

- `int256(5) % int256(-2) == int256(1)`

- `int256(-5) % int256(2) == int256(-1)`

- `int256(-5) % int256(-2) == int256(-1)`

Note: Modulo with zero causes a failing assert.

Exponentiation

Exponentiation is only available for unsigned types in the exponent. The resulting type of an exponentiation is always equal to the type of the base. Please take care that it is large enough to hold the result and prepare for potential wrapping behaviour.

Note: Note that `0**0` is defined by the EVM as `1`.

Fixed Point Numbers

Warning: Fixed point numbers are not fully supported by Solidity yet. They can be declared, but cannot be assigned to or from.

`fixed` / `ufixed`: Signed and unsigned fixed point number of various sizes. Keywords `ufixedMxN` and `fixedMxN`, where `M` represents the number of bits taken by the type and `N` represents how many decimal points are available. `M` must be divisible by 8 and goes from 8 to 256 bits. `N` must be between 0 and 80, inclusive. `ufixed` and `fixed` are aliases for `ufixed128x18` and `fixed128x18`, respectively.

Operators:

- Comparisons: `<=`, `<`, `==`, `!=`, `>=`, `>` (evaluate to `bool`)
- Arithmetic operators: `+`, `-`, unary `-`, `*`, `/`, `%` (modulo)

Note: The main difference between floating point (`float` and `double` in many languages, more precisely IEEE 754 numbers) and fixed point numbers is that the number of bits used for the integer and the fractional part (the part after the decimal dot) is flexible in the former, while it is strictly defined in the latter. Generally, in floating point almost the entire space is used to represent the number, while only a small number of bits define where the decimal point is.

Address

The address type comes in two flavours, which are largely identical:

- `address`: Holds a 20 byte value (size of an Ethereum address).
- `address payable`: Same as `address`, but with the additional members `transfer` and `send`.

The idea behind this distinction is that `address payable` is an address you can send Ether to, while a plain `address` cannot be sent Ether.

Type conversions:

Implicit conversions from `address payable` to `address` are allowed, whereas conversions from `address` to `address payable` must be explicit via `payable(<address>)`.

Address literals can be implicitly converted to `address payable`.

Explicit conversions to and from `address` are allowed for integers, integer literals, `bytes20` and contract types with the following caveat: The result of a conversion of the form `address(x)` has the type `address payable`, if `x` is of integer or fixed bytes type, a literal or a contract with a receive or payable fallback function. If `x` is a contract without a receive or payable fallback function, then `address(x)` will be of type `address`. In external function signatures `address` is used for both the `address` and the `address payable` type.

Only expressions of type `address` can be converted to type `address payable` via `payable(<address>)`.

Note: It might very well be that you do not need to care about the distinction between `address` and `address payable` and just use `address` everywhere. For example, if you are using the *withdrawal pattern*, you can (and should) store the address itself as `address`, because you invoke the `transfer` function on `msg.sender`, which is an `address payable`.

Operators:

- `<=`, `<`, `==`, `!=`, `>=` and `>`

Warning: If you convert a type that uses a larger byte size to an `address`, for example `bytes32`, then the `address` is truncated. To reduce conversion ambiguity version 0.4.24 and higher of the compiler force you make the truncation explicit in the conversion. Take for example the 32-byte value `0x111122223333444455556666777788889999AAAABBBBCCCCDDDDEEEEFFFFCCCC`.

You can use `address(uint160(bytes20(b)))`, which results in `0x111122223333444455556666777788889999aAaa`, or you can use `address(uint160(uint256(b)))`, which results in `0x777788889999AaAAbBbbCcccddDdeeeEfFFfCcCc`.

Note: The distinction between `address` and `address payable` was introduced with version 0.5.0. Also starting from that version, contracts do not derive from the address type, but can still be explicitly converted to `address` or to `address payable`, if they have a receive or payable fallback function.

Members of Addresses

For a quick reference of all members of address, see *Members of Address Types*.

- `balance` and `transfer`

It is possible to query the balance of an address using the property `balance` and to send Ether (in units of wei) to a payable address using the `transfer` function:

```solidity
address payable x = address(0x123);
address myAddress = address(this);
if (x.balance < 10 && myAddress.balance >= 10) x.transfer(10);
```

The `transfer` function fails if the balance of the current contract is not large enough or if the Ether transfer is rejected by the receiving account. The `transfer` function reverts on failure.

Note: If `x` is a contract address, its code (more specifically: its *Receive Ether Function*, if present, or otherwise its *Fallback Function*, if present) will be executed together with the `transfer` call (this is a feature of the EVM and cannot be prevented). If that execution runs out of gas or fails in any way, the Ether transfer will be reverted and the current contract will stop with an exception.

- send

Send is the low-level counterpart of `transfer`. If the execution fails, the current contract will not stop with an exception, but `send` will return `false`.

> **Warning:** There are some dangers in using `send`: The transfer fails if the call stack depth is at 1024 (this can always be forced by the caller) and it also fails if the recipient runs out of gas. So in order to make safe Ether transfers, always check the return value of `send`, use `transfer` or even better: use a pattern where the recipient withdraws the money.

- `call`, `delegatecall` and `staticcall`

In order to interface with contracts that do not adhere to the ABI, or to get more direct control over the encoding, the functions `call`, `delegatecall` and `staticcall` are provided. They all take a single `bytes memory` parameter and return the success condition (as a `bool`) and the returned data (`bytes memory`). The functions `abi.encode`, `abi.encodePacked`, `abi.encodeWithSelector` and `abi.encodeWithSignature` can be used to encode structured data.

Example:

```solidity
bytes memory payload = abi.encodeWithSignature("register(string)", "MyName");
(bool success, bytes memory returnData) = address(nameReg).call(payload);
require(success);
```

> **Warning:** All these functions are low-level functions and should be used with care. Specifically, any unknown contract might be malicious and if you call it, you hand over control to that contract which could in turn call back into your contract, so be prepared for changes to your state variables when the call returns. The regular way to interact with other contracts is to call a function on a contract object (`x.f()`).

Note: Previous versions of Solidity allowed these functions to receive arbitrary arguments and would also handle a first argument of type `bytes4` differently. These edge cases were removed in version 0.5.0.

It is possible to adjust the supplied gas with the `gas` modifier:

```solidity
address(nameReg).call{gas: 1000000}(abi.encodeWithSignature("register(string)",
↪"MyName"));
```

Similarly, the supplied Ether value can be controlled too:

```solidity
address(nameReg).call{value: 1 ether}(abi.encodeWithSignature("register(string)",
↪"MyName"));
```

Lastly, these modifiers can be combined. Their order does not matter:

```solidity
address(nameReg).call{gas: 1000000, value: 1 ether}(abi.encodeWithSignature(
↪"register(string)", "MyName"));
```

In a similar way, the function `delegatecall` can be used: the difference is that only the code of the given address is used, all other aspects (storage, balance, ...) are taken from the current contract. The purpose of `delegatecall` is to use library code which is stored in another contract. The user has to ensure that the layout of storage in both contracts is suitable for delegatecall to be used.

Note: Prior to homestead, only a limited variant called `callcode` was available that did not provide access to the original `msg.sender` and `msg.value` values. This function was removed in version 0.5.0.

Since byzantium `staticcall` can be used as well. This is basically the same as `call`, but will revert if the called function modifies the state in any way.

All three functions `call`, `delegatecall` and `staticcall` are very low-level functions and should only be used as a *last resort* as they break the type-safety of Solidity.

The `gas` option is available on all three methods, while the `value` option is not supported for `delegatecall`.

Note: It is best to avoid relying on hardcoded gas values in your smart contract code, regardless of whether state is read from or written to, as this can have many pitfalls. Also, access to gas might change in the future.

Note: All contracts can be converted to `address` type, so it is possible to query the balance of the current contract using `address(this).balance`.

Contract Types

Every *contract* defines its own type. You can implicitly convert contracts to contracts they inherit from. Contracts can be explicitly converted to and from the `address` type.

Explicit conversion to and from the `address payable` type is only possible if the contract type has a re-ceive or payable fallback function. The conversion is still performed using `address(x)`. If the contract type does not have a receive or payable fallback function, the conversion to `address payable` can be done using `payable(address(x))`. You can find more information in the section about the *address type*.

Note: Before version 0.5.0, contracts directly derived from the address type and there was no distinction between `address` and `address payable`.

If you declare a local variable of contract type (`MyContract c`), you can call functions on that contract. Take care to assign it from somewhere that is the same contract type.

You can also instantiate contracts (which means they are newly created). You can find more details in the *'Contracts via new'* section.

The data representation of a contract is identical to that of the `address` type and this type is also used in the *ABI*.

Contracts do not support any operators.

The members of contract types are the external functions of the contract including any state variables marked as `public`.

For a contract `C` you can use `type(C)` to access *type information* about the contract.

Fixed-size byte arrays

The value types `bytes1`, `bytes2`, `bytes3`, ..., `bytes32` hold a sequence of bytes from one to up to 32. `byte` is an alias for `bytes1`.

Operators:

- Comparisons: <=, <, ==, !=, >=, > (evaluate to `bool`)

- Bit operators: &, |, ^ (bitwise exclusive or), ~ (bitwise negation)

- Shift operators: << (left shift), >> (right shift)

- Index access: If x is of type `bytesI`, then `x[k]` for `0 <= k < I` returns the k th byte (read-only).

The shifting operator works with unsigned integer type as right operand (but returns the type of the left operand), which denotes the number of bits to shift by. Shifting by a signed type will produce a compilation error.

Members:

- `.length` yields the fixed length of the byte array (read-only).

Note: The type `byte[]` is an array of bytes, but due to padding rules, it wastes 31 bytes of space for each element (except in storage). It is better to use the `bytes` type instead.

Dynamically-sized byte array

bytes: Dynamically-sized byte array, see *Arrays*. Not a value-type!

string: Dynamically-sized UTF-8-encoded string, see *Arrays*. Not a value-type!

Address Literals

Hexadecimal literals that pass the address checksum test, for example `0xdCad3a6d3569DF655070DEd06cb7A1b2Ccd1D3AF` are of `address payable` type. Hexadecimal literals that are between 39 and 41 digits long and do not pass the checksum test produce an error. You can prepend (for integer types) or append (for bytesNN types) zeros to remove the error.

Note: The mixed-case address checksum format is defined in EIP-55.

Rational and Integer Literals

Integer literals are formed from a sequence of numbers in the range 0-9. They are interpreted as decimals. For example, `69` means sixty nine. Octal literals do not exist in Solidity and leading zeros are invalid.

Decimal fraction literals are formed by a `.` with at least one number on one side. Examples include `1.`, `.1` and `1.3`.

Scientific notation is also supported, where the base can have fractions and the exponent cannot. Examples include `2e10`, `-2e10`, `2e-10`, `2.5e1`.

Underscores can be used to separate the digits of a numeric literal to aid readability. For example, decimal `123_000`, hexadecimal `0x2eff_abde`, scientific decimal notation `1_2e345_678` are all valid. Underscores are only allowed between two digits and only one consecutive underscore is allowed. There is no additional semantic meaning added to a number literal containing underscores, the underscores are ignored.

Number literal expressions retain arbitrary precision until they are converted to a non-literal type (i.e. by using them together with a non-literal expression or by explicit conversion). This means that computations do not overflow and divisions do not truncate in number literal expressions.

For example, `(2**800 + 1) - 2**800` results in the constant `1` (of type `uint8`) although intermediate results would not even fit the machine word size. Furthermore, `.5 * 8` results in the integer `4` (although non-integers were used in between).

Any operator that can be applied to integers can also be applied to number literal expressions as long as the operands are integers. If any of the two is fractional, bit operations are disallowed and exponentiation is disallowed if the exponent is fractional (because that might result in a non-rational number).

Shifts and exponentiation with literal numbers as left (or base) operand and integer types as the right (exponent) operand are always performed in the `uint256` (for non-negative literals) or `int256` (for a negative literals) type, regardless of the type of the right (exponent) operand.

> **Warning:** Division on integer literals used to truncate in Solidity prior to version 0.4.0, but it now converts into a rational number, i.e. `5 / 2` is not equal to `2`, but to `2.5`.

Note: Solidity has a number literal type for each rational number. Integer literals and rational number literals belong to number literal types. Moreover, all number literal expressions (i.e. the expressions that contain only number literals and operators) belong to number literal types. So the number literal expressions `1 + 2` and `2 + 1` both belong to the same number literal type for the rational number three.

Note: Number literal expressions are converted into a non-literal type as soon as they are used with non-literal expressions. Disregarding types, the value of the expression assigned to `b` below evaluates to an integer. Because `a` is of type `uint128`, the expression `2.5 + a` has to have a proper type, though. Since there is no common type for the type of `2.5` and `uint128`, the Solidity compiler does not accept this code.

```
uint128 a = 1;
uint128 b = 2.5 + a + 0.5;
```

String Literals and Types

String literals are written with either double or single-quotes (`"foo"` or `'bar'`), and they can also be split into multiple consecutive parts (`"foo" "bar"` is equivalent to `"foobar"`) which can be helpful when dealing with long strings. They do not imply trailing zeroes as in C; `"foo"` represents three bytes, not four. As with integer literals, their type can vary, but they are implicitly convertible to `bytes1`, ..., `bytes32`, if they fit, to `bytes` and to `string`.

For example, with `bytes32 samevar = "stringliteral"` the string literal is interpreted in its raw byte form when assigned to a `bytes32` type.

String literals can only contain printable ASCII characters, which means the characters between and including 0x1F .. 0x7E.

Additionally, string literals also support the following escape characters:

- `\<newline>` (escapes an actual newline)
- `\\` (backslash)
- `\'` (single quote)
- `\"` (double quote)
- `\b` (backspace)
- `\f` (form feed)
- `\n` (newline)

- \r (carriage return)

- \t (tab)

- \v (vertical tab)

- \xNN (hex escape, see below)

- \uNNNN (unicode escape, see below)

\xNN takes a hex value and inserts the appropriate byte, while \uNNNN takes a Unicode codepoint and inserts an UTF-8 sequence.

The string in the following example has a length of ten bytes. It starts with a newline byte, followed by a double quote, a single quote a backslash character and then (without separator) the character sequence abcdef.

```
"\n\"\'\\abc\
def"
```

Any Unicode line terminator which is not a newline (i.e. LF, VF, FF, CR, NEL, LS, PS) is considered to terminate the string literal. Newline only terminates the string literal if it is not preceded by a \.

Unicode Literals

While regular string literals can only contain ASCII, Unicode literals – prefixed with the keyword unicode – can contain any valid UTF-8 sequence. They also support the very same escape sequences as regular string literals.

```
string memory a = unicode"Hello ";
```

Hexadecimal Literals

Hexadecimal literals are prefixed with the keyword hex and are enclosed in double or single-quotes (hex"001122FF", hex'0011_22_FF'). Their content must be hexadecimal digits which can optionally use a single underscore as separator between byte boundaries. The value of the literal will be the binary representation of the hexadecimal sequence.

Multiple hexadecimal literals separated by whitespace are concatenated into a single literal: hex"00112233" hex"44556677" is equivalent to hex"0011223344556677"

Hexadecimal literals behave like *string literals* and have the same convertibility restrictions.

Enums

Enums are one way to create a user-defined type in Solidity. They are explicitly convertible to and from all integer types but implicit conversion is not allowed. The explicit conversion from integer checks at runtime that the value lies inside the range of the enum and causes a failing assert otherwise. Enums require at least one member, and its default value when declared is the first member.

The data representation is the same as for enums in C: The options are represented by subsequent unsigned integer values starting from 0.

```
// SPDX-License-Identifier: GPL-3.0
pragma solidity >=0.4.16 <0.8.0;

contract test {
    enum ActionChoices { GoLeft, GoRight, GoStraight, SitStill }
```

(continues on next page)

```
    ActionChoices choice;
    ActionChoices constant defaultChoice = ActionChoices.GoStraight;

    function setGoStraight() public {
        choice = ActionChoices.GoStraight;
    }

    // Since enum types are not part of the ABI, the signature of "getChoice"
    // will automatically be changed to "getChoice() returns (uint8)"
    // for all matters external to Solidity. The integer type used is just
    // large enough to hold all enum values, i.e. if you have more than 256 values,
    // `uint16` will be used and so on.
    function getChoice() public view returns (ActionChoices) {
        return choice;
    }

    function getDefaultChoice() public pure returns (uint) {
        return uint(defaultChoice);
    }
}
```

Note: Enums can also be declared on the file level, outside of contract or library definitions.

Function Types

Function types are the types of functions. Variables of function type can be assigned from functions and function parameters of function type can be used to pass functions to and return functions from function calls. Function types come in two flavours - *internal* and *external* functions:

Internal functions can only be called inside the current contract (more specifically, inside the current code unit, which also includes internal library functions and inherited functions) because they cannot be executed outside of the context of the current contract. Calling an internal function is realized by jumping to its entry label, just like when calling a function of the current contract internally.

External functions consist of an address and a function signature and they can be passed via and returned from external function calls.

Function types are notated as follows:

```
function (<parameter types>) {internal|external} [pure|view|payable] [returns (
↪<return types>)]
```

In contrast to the parameter types, the return types cannot be empty - if the function type should not return anything, the whole `returns (<return types>)` part has to be omitted.

By default, function types are internal, so the `internal` keyword can be omitted. Note that this only applies to function types. Visibility has to be specified explicitly for functions defined in contracts, they do not have a default.

Conversions:

A function type `A` is implicitly convertible to a function type `B` if and only if their parameter types are identical, their return types are identical, their internal/external property is identical and the state mutability of `A` is not more restrictive than the state mutability of `B`. In particular:

- `pure` functions can be converted to `view` and `non-payable` functions

- `view` functions can be converted to `non-payable` functions

- `payable` functions can be converted to `non-payable` functions

No other conversions between function types are possible.

The rule about `payable` and `non-payable` might be a little confusing, but in essence, if a function is `payable`, this means that it also accepts a payment of zero Ether, so it also is `non-payable`. On the other hand, a `non-payable` function will reject Ether sent to it, so `non-payable` functions cannot be converted to `payable` functions.

If a function type variable is not initialised, calling it results in a failed assertion. The same happens if you call a function after using `delete` on it.

If external function types are used outside of the context of Solidity, they are treated as the `function` type, which encodes the address followed by the function identifier together in a single `bytes24` type.

Note that public functions of the current contract can be used both as an internal and as an external function. To use `f` as an internal function, just use `f`, if you want to use its external form, use `this.f`.

Members:

External (or public) functions have the following members:

- `.address` returns the address of the contract of the function.

- `.selector` returns the *ABI function selector*

Note: External (or public) functions used to have the additional members `.gas(uint)` and `.value(uint)`. These were deprecated in Solidity 0.6.2 and removed in Solidity 0.7.0. Instead use `{gas: ...}` and `{value: ...}` to specify the amount of gas or the amount of wei sent to a function, respectively. See *External Function Calls* for more information.

Example that shows how to use the members:

```solidity
// SPDX-License-Identifier: GPL-3.0
pragma solidity >=0.6.4 <0.8.0;

contract Example {
    function f() public payable returns (bytes4) {
        assert(this.f.address == address(this));
        return this.f.selector;
    }

    function g() public {
        this.f{gas: 10, value: 800}();
    }
}
```

Example that shows how to use internal function types:

```solidity
// SPDX-License-Identifier: GPL-3.0
pragma solidity >=0.4.16 <0.8.0;

library ArrayUtils {
    // internal functions can be used in internal library functions because
    // they will be part of the same code context
    function map(uint[] memory self, function (uint) pure returns (uint) f)
        internal
```

```solidity
        pure
        returns (uint[] memory r)
    {
        r = new uint[](self.length);
        for (uint i = 0; i < self.length; i++) {
            r[i] = f(self[i]);
        }
    }

    function reduce(
        uint[] memory self,
        function (uint, uint) pure returns (uint) f
    )
        internal
        pure
        returns (uint r)
    {
        r = self[0];
        for (uint i = 1; i < self.length; i++) {
            r = f(r, self[i]);
        }
    }

    function range(uint length) internal pure returns (uint[] memory r) {
        r = new uint[](length);
        for (uint i = 0; i < r.length; i++) {
            r[i] = i;
        }
    }
}


contract Pyramid {
    using ArrayUtils for *;

    function pyramid(uint l) public pure returns (uint) {
        return ArrayUtils.range(l).map(square).reduce(sum);
    }

    function square(uint x) internal pure returns (uint) {
        return x * x;
    }

    function sum(uint x, uint y) internal pure returns (uint) {
        return x + y;
    }
}
```

Another example that uses external function types:

```solidity
// SPDX-License-Identifier: GPL-3.0
pragma solidity >=0.4.22 <0.8.0;


contract Oracle {
    struct Request {
        bytes data;
```

```solidity
        function(uint) external callback;
    }

    Request[] private requests;
    event NewRequest(uint);

    function query(bytes memory data, function(uint) external callback) public {
        requests.push(Request(data, callback));
        emit NewRequest(requests.length - 1);
    }

    function reply(uint requestID, uint response) public {
        // Here goes the check that the reply comes from a trusted source
        requests[requestID].callback(response);
    }
}


contract OracleUser {
    Oracle constant private ORACLE_CONST = Oracle(0x1234567); // known contract
    uint private exchangeRate;

    function buySomething() public {
        ORACLE_CONST.query("USD", this.oracleResponse);
    }

    function oracleResponse(uint response) public {
        require(
            msg.sender == address(ORACLE_CONST),
            "Only oracle can call this."
        );
        exchangeRate = response;
    }
}
```

Note: Lambda or inline functions are planned but not yet supported.

3.6.2 Reference Types

Values of reference type can be modified through multiple different names. Contrast this with value types where you get an independent copy whenever a variable of value type is used. Because of that, reference types have to be handled more carefully than value types. Currently, reference types comprise structs, arrays and mappings. If you use a reference type, you always have to explicitly provide the data area where the type is stored: `memory` (whose lifetime is limited to an external function call), `storage` (the location where the state variables are stored, where the lifetime is limited to the lifetime of a contract) or `calldata` (special data location that contains the function arguments).

An assignment or type conversion that changes the data location will always incur an automatic copy operation, while assignments inside the same data location only copy in some cases for storage types.

Data location

Every reference type has an additional annotation, the "data location", about where it is stored. There are three data locations: `memory`, `storage` and `calldata`. Calldata is a non-modifiable, non-persistent area where function

arguments are stored, and behaves mostly like memory. It is required for parameters of external functions but can also be used for other variables.

Note: Prior to version 0.5.0 the data location could be omitted, and would default to different locations depending on the kind of variable, function type, etc., but all complex types must now give an explicit data location.

Note: If you can, try to use `calldata` as data location because it will avoid copies and also makes sure that the data cannot be modified. Arrays and structs with `calldata` data location can also be returned from functions, but it is not possible to allocate such types.

Data location and assignment behaviour

Data locations are not only relevant for persistency of data, but also for the semantics of assignments:

- Assignments between `storage` and `memory` (or from `calldata`) always create an independent copy.

- Assignments from `memory` to `memory` only create references. This means that changes to one memory variable are also visible in all other memory variables that refer to the same data.

- Assignments from `storage` to a **local** storage variable also only assign a reference.

- All other assignments to `storage` always copy. Examples for this case are assignments to state variables or to members of local variables of storage struct type, even if the local variable itself is just a reference.

```solidity
// SPDX-License-Identifier: GPL-3.0
pragma solidity >=0.5.0 <0.8.0;

contract C {
    // The data location of x is storage.
    // This is the only place where the
    // data location can be omitted.
    uint[] x;

    // The data location of memoryArray is memory.
    function f(uint[] memory memoryArray) public {
        x = memoryArray; // works, copies the whole array to storage
        uint[] storage y = x; // works, assigns a pointer, data location of y is␣
↪storage
        y[7]; // fine, returns the 8th element
        y.pop(); // fine, modifies x through y
        delete x; // fine, clears the array, also modifies y
        // The following does not work; it would need to create a new temporary /
        // unnamed array in storage, but storage is "statically" allocated:
        // y = memoryArray;
        // This does not work either, since it would "reset" the pointer, but there
        // is no sensible location it could point to.
        // delete y;
        g(x); // calls g, handing over a reference to x
        h(x); // calls h and creates an independent, temporary copy in memory
    }

    function g(uint[] storage) internal pure {}
    function h(uint[] memory) public pure {}
}
```

Arrays

Arrays can have a compile-time fixed size, or they can have a dynamic size.

The type of an array of fixed size `k` and element type `T` is written as `T[k]`, and an array of dynamic size as `T[]`.

For example, an array of 5 dynamic arrays of `uint` is written as `uint[][5]`. The notation is reversed compared to some other languages. In Solidity, `X[3]` is always an array containing three elements of type `X`, even if `X` is itself an array. This is not the case in other languages such as C.

Indices are zero-based, and access is in the opposite direction of the declaration.

For example, if you have a variable `uint[][5] memory x`, you access the second `uint` in the third dynamic array using `x[2][1]`, and to access the third dynamic array, use `x[2]`. Again, if you have an array `T[5] a` for a type `T` that can also be an array, then `a[2]` always has type `T`.

Array elements can be of any type, including mapping or struct. The general restrictions for types apply, in that mappings can only be stored in the `storage` data location and publicly-visible functions need parameters that are *ABI types*.

It is possible to mark state variable arrays `public` and have Solidity create a *getter*. The numeric index becomes a required parameter for the getter.

Accessing an array past its end causes a failing assertion. Methods `.push()` and `.push(value)` can be used to append a new element at the end of the array, where `.push()` appends a zero-initialized element and returns a reference to it.

bytes and strings as Arrays

Variables of type `bytes` and `string` are special arrays. A `bytes` is similar to `byte[]`, but it is packed tightly in calldata and memory. `string` is equal to `bytes` but does not allow length or index access.

Solidity does not have string manipulation functions, but there are third-party string libraries. You can also compare two strings by their keccak256-hash using `keccak256(abi.encodePacked(s1)) == keccak256(abi.encodePacked(s2))` and concatenate two strings using `abi.encodePacked(s1, s2)`.

You should use `bytes` over `byte[]` because it is cheaper, since `byte[]` adds 31 padding bytes between the elements. As a general rule, use `bytes` for arbitrary-length raw byte data and `string` for arbitrary-length string (UTF-8) data. If you can limit the length to a certain number of bytes, always use one of the value types `bytes1` to `bytes32` because they are much cheaper.

Note: If you want to access the byte-representation of a string `s`, use `bytes(s).length` / `bytes(s)[7] = 'x';`. Keep in mind that you are accessing the low-level bytes of the UTF-8 representation, and not the individual characters.

Allocating Memory Arrays

Memory arrays with dynamic length can be created using the `new` operator. As opposed to storage arrays, it is **not** possible to resize memory arrays (e.g. the `.push` member functions are not available). You either have to calculate the required size in advance or create a new memory array and copy every element.

```
// SPDX-License-Identifier: GPL-3.0
pragma solidity >=0.4.16 <0.8.0;

contract C {
```

(continues on next page)

```
    function f(uint len) public pure {
        uint[] memory a = new uint[](7);
        bytes memory b = new bytes(len);
        assert(a.length == 7);
        assert(b.length == len);
        a[6] = 8;
    }
}
```

Array Literals

An array literal is a comma-separated list of one or more expressions, enclosed in square brackets ([...]). For example [1, a, f(3)]. There must be a common type all elements can be implicitly converted to. This is the elementary type of the array.

Array literals are always statically-sized memory arrays.

In the example below, the type of [1, 2, 3] is uint8[3] memory. Because the type of each of these constants is uint8, if you want the result to be a uint[3] memory type, you need to convert the first element to uint.

```
// SPDX-License-Identifier: GPL-3.0
pragma solidity >=0.4.16 <0.8.0;

contract C {
    function f() public pure {
        g([uint(1), 2, 3]);
    }
    function g(uint[3] memory) public pure {
        // ...
    }
}
```

Fixed size memory arrays cannot be assigned to dynamically-sized memory arrays, i.e. the following is not possible:

```
// SPDX-License-Identifier: GPL-3.0
pragma solidity >=0.4.0 <0.8.0;

// This will not compile.
contract C {
    function f() public {
        // The next line creates a type error because uint[3] memory
        // cannot be converted to uint[] memory.
        uint[] memory x = [uint(1), 3, 4];
    }
}
```

It is planned to remove this restriction in the future, but it creates some complications because of how arrays are passed in the ABI.

If you want to initialize dynamically-sized arrays, you have to assign the individual elements:

```
// SPDX-License-Identifier: GPL-3.0
pragma solidity >=0.4.0 <0.8.0;

contract C {
```

```solidity
    function f() public pure {
        uint[] memory x = new uint[](3);
        x[0] = 1;
        x[1] = 3;
        x[2] = 4;
    }
}
```

Array Members

length: Arrays have a `length` member that contains their number of elements. The length of memory arrays is fixed (but dynamic, i.e. it can depend on runtime parameters) once they are created.

push(): Dynamic storage arrays and `bytes` (not `string`) have a member function called `push()` that you can use to append a zero-initialised element at the end of the array. It returns a reference to the element, so that it can be used like `x.push().t = 2` or `x.push() = b`.

push(x): Dynamic storage arrays and `bytes` (not `string`) have a member function called `push(x)` that you can use to append a given element at the end of the array. The function returns nothing.

pop: Dynamic storage arrays and `bytes` (not `string`) have a member function called `pop` that you can use to remove an element from the end of the array. This also implicitly calls *delete* on the removed element.

Note: Increasing the length of a storage array by calling `push()` has constant gas costs because storage is zero-initialised, while decreasing the length by calling `pop()` has a cost that depends on the "size" of the element being removed. If that element is an array, it can be very costly, because it includes explicitly clearing the removed elements similar to calling *delete* on them.

Note: To use arrays of arrays in external (instead of public) functions, you need to activate ABI coder v2.

Note: In EVM versions before Byzantium, it was not possible to access dynamic arrays return from function calls. If you call functions that return dynamic arrays, make sure to use an EVM that is set to Byzantium mode.

```solidity
// SPDX-License-Identifier: GPL-3.0
pragma solidity >=0.6.0 <0.8.0;

contract ArrayContract {
    uint[2**20] m_aLotOfIntegers;
    // Note that the following is not a pair of dynamic arrays but a
    // dynamic array of pairs (i.e. of fixed size arrays of length two).
    // Because of that, T[] is always a dynamic array of T, even if T
    // itself is an array.
    // Data location for all state variables is storage.
    bool[2][] m_pairsOfFlags;

    // newPairs is stored in memory - the only possibility
    // for public contract function arguments
    function setAllFlagPairs(bool[2][] memory newPairs) public {
        // assignment to a storage array performs a copy of ``newPairs`` and
        // replaces the complete array ``m_pairsOfFlags``.
```

```solidity
        m_pairsOfFlags = newPairs;
    }

    struct StructType {
        uint[] contents;
        uint moreInfo;
    }
    StructType s;

    function f(uint[] memory c) public {
        // stores a reference to ``s`` in ``g``
        StructType storage g = s;
        // also changes ``s.moreInfo``.
        g.moreInfo = 2;
        // assigns a copy because ``g.contents``
        // is not a local variable, but a member of
        // a local variable.
        g.contents = c;
    }

    function setFlagPair(uint index, bool flagA, bool flagB) public {
        // access to a non-existing index will throw an exception
        m_pairsOfFlags[index][0] = flagA;
        m_pairsOfFlags[index][1] = flagB;
    }

    function changeFlagArraySize(uint newSize) public {
        // using push and pop is the only way to change the
        // length of an array
        if (newSize < m_pairsOfFlags.length) {
            while (m_pairsOfFlags.length > newSize)
                m_pairsOfFlags.pop();
        } else if (newSize > m_pairsOfFlags.length) {
            while (m_pairsOfFlags.length < newSize)
                m_pairsOfFlags.push();
        }
    }

    function clear() public {
        // these clear the arrays completely
        delete m_pairsOfFlags;
        delete m_aLotOfIntegers;
        // identical effect here
        m_pairsOfFlags = new bool[2][](0);
    }

    bytes m_byteData;

    function byteArrays(bytes memory data) public {
        // byte arrays ("bytes") are different as they are stored without padding,
        // but can be treated identical to "uint8[]"
        m_byteData = data;
        for (uint i = 0; i < 7; i++)
            m_byteData.push();
        m_byteData[3] = 0x08;
        delete m_byteData[2];
    }
```

```solidity
    function addFlag(bool[2] memory flag) public returns (uint) {
        m_pairsOfFlags.push(flag);
        return m_pairsOfFlags.length;
    }

    function createMemoryArray(uint size) public pure returns (bytes memory) {
        // Dynamic memory arrays are created using `new`:
        uint[2][] memory arrayOfPairs = new uint[2][](size);

        // Inline arrays are always statically-sized and if you only
        // use literals, you have to provide at least one type.
        arrayOfPairs[0] = [uint(1), 2];

        // Create a dynamic byte array:
        bytes memory b = new bytes(200);
        for (uint i = 0; i < b.length; i++)
            b[i] = byte(uint8(i));
        return b;
    }
}
```

Array Slices

Array slices are a view on a contiguous portion of an array. They are written as `x[start:end]`, where `start` and `end` are expressions resulting in a uint256 type (or implicitly convertible to it). The first element of the slice is `x[start]` and the last element is `x[end - 1]`.

If `start` is greater than `end` or if `end` is greater than the length of the array, an exception is thrown.

Both `start` and `end` are optional: `start` defaults to `0` and `end` defaults to the length of the array.

Array slices do not have any members. They are implicitly convertible to arrays of their underlying type and support index access. Index access is not absolute in the underlying array, but relative to the start of the slice.

Array slices do not have a type name which means no variable can have an array slices as type, they only exist in intermediate expressions.

Note: As of now, array slices are only implemented for calldata arrays.

Array slices are useful to ABI-decode secondary data passed in function parameters:

```solidity
// SPDX-License-Identifier: GPL-3.0
pragma solidity ^0.7.0;

contract Proxy {
    /// @dev Address of the client contract managed by proxy i.e., this contract
    address client;

    constructor(address _client) {
        client = _client;
    }

    /// Forward call to "setOwner(address)" that is implemented by client
    /// after doing basic validation on the address argument.
```

```solidity
    function forward(bytes calldata _payload) external {
        // Since ABI decoding requires padded data, we cannot
        // use abi.decode(_payload[:4], (bytes4)).
        bytes4 sig =
            _payload[0] |
            (bytes4(_payload[1]) >> 8) |
            (bytes4(_payload[2]) >> 16) |
            (bytes4(_payload[3]) >> 24);
        if (sig == bytes4(keccak256("setOwner(address)"))) {
            address owner = abi.decode(_payload[4:], (address));
            require(owner != address(0), "Address of owner cannot be zero.");
        }
        (bool status,) = client.delegatecall(_payload);
        require(status, "Forwarded call failed.");
    }
}
```

Structs

Solidity provides a way to define new types in the form of structs, which is shown in the following example:

```solidity
// SPDX-License-Identifier: GPL-3.0
pragma solidity >=0.6.0 <0.8.0;

// Defines a new type with two fields.
// Declaring a struct outside of a contract allows
// it to be shared by multiple contracts.
// Here, this is not really needed.
struct Funder {
    address addr;
    uint amount;
}

contract CrowdFunding {
    // Structs can also be defined inside contracts, which makes them
    // visible only there and in derived contracts.
    struct Campaign {
        address payable beneficiary;
        uint fundingGoal;
        uint numFunders;
        uint amount;
        mapping (uint => Funder) funders;
    }

    uint numCampaigns;
    mapping (uint => Campaign) campaigns;

    function newCampaign(address payable beneficiary, uint goal) public returns (uint␣
↪campaignID) {
        campaignID = numCampaigns++; // campaignID is return variable
        // We cannot use "campaigns[campaignID] = Campaign(beneficiary, goal, 0, 0)"
        // because the RHS creates a memory-struct "Campaign" that contains a␣
↪mapping.
        Campaign storage c = campaigns[campaignID];
        c.beneficiary = beneficiary;
```

```
        c.fundingGoal = goal;
    }

    function contribute(uint campaignID) public payable {
        Campaign storage c = campaigns[campaignID];
        // Creates a new temporary memory struct, initialised with the given values
        // and copies it over to storage.
        // Note that you can also use Funder(msg.sender, msg.value) to initialise.
        c.funders[c.numFunders++] = Funder({addr: msg.sender, amount: msg.value});
        c.amount += msg.value;
    }

    function checkGoalReached(uint campaignID) public returns (bool reached) {
        Campaign storage c = campaigns[campaignID];
        if (c.amount < c.fundingGoal)
            return false;
        uint amount = c.amount;
        c.amount = 0;
        c.beneficiary.transfer(amount);
        return true;
    }
}
```

The contract does not provide the full functionality of a crowdfunding contract, but it contains the basic concepts necessary to understand structs. Struct types can be used inside mappings and arrays and they can themselves contain mappings and arrays.

It is not possible for a struct to contain a member of its own type, although the struct itself can be the value type of a mapping member or it can contain a dynamically-sized array of its type. This restriction is necessary, as the size of the struct has to be finite.

Note how in all the functions, a struct type is assigned to a local variable with data location `storage`. This does not copy the struct but only stores a reference so that assignments to members of the local variable actually write to the state.

Of course, you can also directly access the members of the struct without assigning it to a local variable, as in `campaigns[campaignID].amount = 0`.

3.6.3 Mapping Types

Mapping types use the syntax `mapping(_KeyType => _ValueType)` and variables of mapping type are declared using the syntax `mapping(_KeyType => _ValueType) _VariableName`. The `_KeyType` can be any built-in value type, `bytes`, `string`, or any contract or enum type. Other user-defined or complex types, such as mappings, structs or array types are not allowed. `_ValueType` can be any type, including mappings, arrays and structs.

You can think of mappings as hash tables, which are virtually initialised such that every possible key exists and is mapped to a value whose byte-representation is all zeros, a type's *default value*. The similarity ends there, the key data is not stored in a mapping, only its `keccak256` hash is used to look up the value.

Because of this, mappings do not have a length or a concept of a key or value being set, and therefore cannot be erased without extra information regarding the assigned keys (see *Clearing Mappings*).

Mappings can only have a data location of `storage` and thus are allowed for state variables, as storage reference types in functions, or as parameters for library functions. They cannot be used as parameters or return parameters of contract functions that are publicly visible. These restrictions are also true for arrays and structs that contain mappings.

You can mark state variables of mapping type as `public` and Solidity creates a *getter* for you. The `_KeyType` becomes a parameter for the getter. If `_ValueType` is a value type or a struct, the getter returns `_ValueType`. If `_ValueType` is an array or a mapping, the getter has one parameter for each `_KeyType`, recursively.

In the example below, the `MappingExample` contract defines a public `balances` mapping, with the key type an `address`, and a value type a `uint`, mapping an Ethereum address to an unsigned integer value. As `uint` is a value type, the getter returns a value that matches the type, which you can see in the `MappingUser` contract that returns the value at the specified address.

```solidity
// SPDX-License-Identifier: GPL-3.0
pragma solidity >=0.4.0 <0.8.0;

contract MappingExample {
    mapping(address => uint) public balances;

    function update(uint newBalance) public {
        balances[msg.sender] = newBalance;
    }
}

contract MappingUser {
    function f() public returns (uint) {
        MappingExample m = new MappingExample();
        m.update(100);
        return m.balances(address(this));
    }
}
```

The example below is a simplified version of an ERC20 token. `_allowances` is an example of a mapping type inside another mapping type. The example below uses `_allowances` to record the amount someone else is allowed to withdraw from your account.

```solidity
// SPDX-License-Identifier: GPL-3.0
pragma solidity >=0.4.22 <0.8.0;

contract MappingExample {

    mapping (address => uint256) private _balances;
    mapping (address => mapping (address => uint256)) private _allowances;

    event Transfer(address indexed from, address indexed to, uint256 value);
    event Approval(address indexed owner, address indexed spender, uint256 value);

    function allowance(address owner, address spender) public view returns (uint256) {
        return _allowances[owner][spender];
    }

    function transferFrom(address sender, address recipient, uint256 amount) public
→returns (bool) {
        _transfer(sender, recipient, amount);
        approve(sender, msg.sender, amount);
        return true;
    }

    function approve(address owner, address spender, uint256 amount) public returns
→(bool) {
        require(owner != address(0), "ERC20: approve from the zero address");
        require(spender != address(0), "ERC20: approve to the zero address");
```

(continues on next page)

```solidity
        _allowances[owner][spender] = amount;
        emit Approval(owner, spender, amount);
        return true;
    }

    function _transfer(address sender, address recipient, uint256 amount) internal {
        require(sender != address(0), "ERC20: transfer from the zero address");
        require(recipient != address(0), "ERC20: transfer to the zero address");

        _balances[sender] -= amount;
        _balances[recipient] += amount;
        emit Transfer(sender, recipient, amount);
    }
}
```

Iterable Mappings

You cannot iterate over mappings, i.e. you cannot enumerate their keys. It is possible, though, to implement a data structure on top of them and iterate over that. For example, the code below implements an `IterableMapping` library that the `User` contract then adds data too, and the `sum` function iterates over to sum all the values.

```solidity
// SPDX-License-Identifier: GPL-3.0
pragma solidity >=0.6.0 <0.8.0;

struct IndexValue { uint keyIndex; uint value; }
struct KeyFlag { uint key; bool deleted; }

struct itmap {
    mapping(uint => IndexValue) data;
    KeyFlag[] keys;
    uint size;
}

library IterableMapping {
    function insert(itmap storage self, uint key, uint value) internal returns (bool
→replaced) {
        uint keyIndex = self.data[key].keyIndex;
        self.data[key].value = value;
        if (keyIndex > 0)
            return true;
        else {
            keyIndex = self.keys.length;
            self.keys.push();
            self.data[key].keyIndex = keyIndex + 1;
            self.keys[keyIndex].key = key;
            self.size++;
            return false;
        }
    }

    function remove(itmap storage self, uint key) internal returns (bool success) {
        uint keyIndex = self.data[key].keyIndex;
        if (keyIndex == 0)
            return false;
```

```solidity
        delete self.data[key];
        self.keys[keyIndex - 1].deleted = true;
        self.size --;
    }

    function contains(itmap storage self, uint key) internal view returns (bool) {
        return self.data[key].keyIndex > 0;
    }

    function iterate_start(itmap storage self) internal view returns (uint keyIndex) {
        return iterate_next(self, uint(-1));
    }

    function iterate_valid(itmap storage self, uint keyIndex) internal view returns␣
→(bool) {
        return keyIndex < self.keys.length;
    }

    function iterate_next(itmap storage self, uint keyIndex) internal view returns␣
→(uint r_keyIndex) {
        keyIndex++;
        while (keyIndex < self.keys.length && self.keys[keyIndex].deleted)
            keyIndex++;
        return keyIndex;
    }

    function iterate_get(itmap storage self, uint keyIndex) internal view returns␣
→(uint key, uint value) {
        key = self.keys[keyIndex].key;
        value = self.data[key].value;
    }
}

// How to use it
contract User {
    // Just a struct holding our data.
    itmap data;
    // Apply library functions to the data type.
    using IterableMapping for itmap;

    // Insert something
    function insert(uint k, uint v) public returns (uint size) {
        // This calls IterableMapping.insert(data, k, v)
        data.insert(k, v);
        // We can still access members of the struct,
        // but we should take care not to mess with them.
        return data.size;
    }

    // Computes the sum of all stored data.
    function sum() public view returns (uint s) {
        for (
            uint i = data.iterate_start();
            data.iterate_valid(i);
            i = data.iterate_next(i)
        ) {
            (, uint value) = data.iterate_get(i);
```

```
            s += value;
        }
    }
}
```

3.6.4 Operators Involving LValues

If `a` is an LValue (i.e. a variable or something that can be assigned to), the following operators are available as shorthands:

`a += e` is equivalent to `a = a + e`. The operators `-=`, `*=`, `/=`, `%=`, `|=`, `&=` and `^=` are defined accordingly. `a++` and `a--` are equivalent to `a += 1` / `a -= 1` but the expression itself still has the previous value of `a`. In contrast, `--a` and `++a` have the same effect on `a` but return the value after the change.

delete

`delete a` assigns the initial value for the type to `a`. I.e. for integers it is equivalent to `a = 0`, but it can also be used on arrays, where it assigns a dynamic array of length zero or a static array of the same length with all elements set to their initial value. `delete a[x]` deletes the item at index `x` of the array and leaves all other elements and the length of the array untouched. This especially means that it leaves a gap in the array. If you plan to remove items, a *mapping* is probably a better choice.

For structs, it assigns a struct with all members reset. In other words, the value of `a` after `delete a` is the same as if `a` would be declared without assignment, with the following caveat:

`delete` has no effect on mappings (as the keys of mappings may be arbitrary and are generally unknown). So if you delete a struct, it will reset all members that are not mappings and also recurse into the members unless they are mappings. However, individual keys and what they map to can be deleted: If `a` is a mapping, then `delete a[x]` will delete the value stored at `x`.

It is important to note that `delete a` really behaves like an assignment to `a`, i.e. it stores a new object in `a`. This distinction is visible when `a` is reference variable: It will only reset `a` itself, not the value it referred to previously.

```solidity
// SPDX-License-Identifier: GPL-3.0
pragma solidity >=0.4.0 <0.8.0;

contract DeleteExample {
    uint data;
    uint[] dataArray;

    function f() public {
        uint x = data;
        delete x; // sets x to 0, does not affect data
        delete data; // sets data to 0, does not affect x
        uint[] storage y = dataArray;
        delete dataArray; // this sets dataArray.length to zero, but as uint[] is a␣
→complex object, also
        // y is affected which is an alias to the storage object
        // On the other hand: "delete y" is not valid, as assignments to local␣
→variables
        // referencing storage objects can only be made from existing storage␣
→objects.
        assert(y.length == 0);
    }
}
```

3.6.5 Conversions between Elementary Types

Implicit Conversions

An implicit type conversion is automatically applied by the compiler in some cases during assignments, when passing arguments to functions and when applying operators. In general, an implicit conversion between value-types is possible if it makes sense semantically and no information is lost.

For example, `uint8` is convertible to `uint16` and `int128` to `int256`, but `int8` is not convertible to `uint256`, because `uint256` cannot hold values such as `-1`.

If an operator is applied to different types, the compiler tries to implicitly convert one of the operands to the type of the other (the same is true for assignments). This means that operations are always performed in the type of one of the operands.

For more details about which implicit conversions are possible, please consult the sections about the types themselves.

In the example below, `y` and `z`, the operands of the addition, do not have the same type, but `uint8` can be implicitly converted to `uint16` and not vice-versa. Because of that, `y` is converted to the type of `z` before the addition is performed in the `uint16` type. The resulting type of the expression `y + z` is `uint16``. Because it is assigned to a variable of type ``uint32` another implicit conversion is performed after the addition.

```
uint8 y;
uint16 z;
uint32 x = y + z;
```

Explicit Conversions

If the compiler does not allow implicit conversion but you are confident a conversion will work, an explicit type conversion is sometimes possible. This may result in unexpected behaviour and allows you to bypass some security features of the compiler, so be sure to test that the result is what you want and expect!

Take the following example that converts a negative `int` to a `uint`:

```
int    y = -3;
uint x = uint(y);
```

At the end of this code snippet, `x` will have the value `0xfffff..fd` (64 hex characters), which is -3 in the two's complement representation of 256 bits.

If an integer is explicitly converted to a smaller type, higher-order bits are cut off:

```
uint32 a = 0x12345678;
uint16 b = uint16(a); // b will be 0x5678 now
```

If an integer is explicitly converted to a larger type, it is padded on the left (i.e., at the higher order end). The result of the conversion will compare equal to the original integer:

```
uint16 a = 0x1234;
uint32 b = uint32(a); // b will be 0x00001234 now
assert(a == b);
```

Fixed-size bytes types behave differently during conversions. They can be thought of as sequences of individual bytes and converting to a smaller type will cut off the sequence:

```
bytes2 a = 0x1234;
bytes1 b = bytes1(a); // b will be 0x12
```

If a fixed-size bytes type is explicitly converted to a larger type, it is padded on the right. Accessing the byte at a fixed index will result in the same value before and after the conversion (if the index is still in range):

```
bytes2 a = 0x1234;
bytes4 b = bytes4(a); // b will be 0x12340000
assert(a[0] == b[0]);
assert(a[1] == b[1]);
```

Since integers and fixed-size byte arrays behave differently when truncating or padding, explicit conversions between integers and fixed-size byte arrays are only allowed, if both have the same size. If you want to convert between integers and fixed-size byte arrays of different size, you have to use intermediate conversions that make the desired truncation and padding rules explicit:

```
bytes2 a = 0x1234;
uint32 b = uint16(a); // b will be 0x00001234
uint32 c = uint32(bytes4(a)); // c will be 0x12340000
uint8 d = uint8(uint16(a)); // d will be 0x34
uint8 e = uint8(bytes1(a)); // e will be 0x12
```

3.6.6 Conversions between Literals and Elementary Types

Integer Types

Decimal and hexadecimal number literals can be implicitly converted to any integer type that is large enough to represent it without truncation:

```
uint8 a = 12; // fine
uint32 b = 1234; // fine
uint16 c = 0x123456; // fails, since it would have to truncate to 0x3456
```

Fixed-Size Byte Arrays

Decimal number literals cannot be implicitly converted to fixed-size byte arrays. Hexadecimal number literals can be, but only if the number of hex digits exactly fits the size of the bytes type. As an exception both decimal and hexadecimal literals which have a value of zero can be converted to any fixed-size bytes type:

```
bytes2 a = 54321; // not allowed
bytes2 b = 0x12; // not allowed
bytes2 c = 0x123; // not allowed
bytes2 d = 0x1234; // fine
bytes2 e = 0x0012; // fine
bytes4 f = 0; // fine
bytes4 g = 0x0; // fine
```

String literals and hex string literals can be implicitly converted to fixed-size byte arrays, if their number of characters matches the size of the bytes type:

```
bytes2 a = hex"1234"; // fine
bytes2 b = "xy"; // fine
bytes2 c = hex"12"; // not allowed
bytes2 d = hex"123"; // not allowed
bytes2 e = "x"; // not allowed
bytes2 f = "xyz"; // not allowed
```

Addresses

As described in *Address Literals*, hex literals of the correct size that pass the checksum test are of `address` type. No other literals can be implicitly converted to the `address` type.

Explicit conversions from `bytes20` or any integer type to `address` result in `address payable`.

An `address a` can be converted to `address payable` via `payable(a)`.

3.7 Units and Globally Available Variables

3.7.1 Ether Units

A literal number can take a suffix of `wei`, `gwei` or `ether` to specify a subdenomination of Ether, where Ether numbers without a postfix are assumed to be Wei.

```
assert(1 wei == 1);
assert(1 gwei == 1e9);
assert(1 ether == 1e18);
```

The only effect of the subdenomination suffix is a multiplication by a power of ten.

Note: The denominations `finney` and `szabo` have been removed in version 0.7.0.

3.7.2 Time Units

Suffixes like `seconds`, `minutes`, `hours`, `days` and `weeks` after literal numbers can be used to specify units of time where seconds are the base unit and units are considered naively in the following way:

- `1 == 1 seconds`
- `1 minutes == 60 seconds`
- `1 hours == 60 minutes`
- `1 days == 24 hours`
- `1 weeks == 7 days`

Take care if you perform calendar calculations using these units, because not every year equals 365 days and not even every day has 24 hours because of leap seconds. Due to the fact that leap seconds cannot be predicted, an exact calendar library has to be updated by an external oracle.

Note: The suffix `years` has been removed in version 0.5.0 due to the reasons above.

These suffixes cannot be applied to variables. For example, if you want to interpret a function parameter in days, you can in the following way:

```
function f(uint start, uint daysAfter) public {
    if (block.timestamp >= start + daysAfter * 1 days) {
      // ...
    }
}
```

3.7.3 Special Variables and Functions

There are special variables and functions which always exist in the global namespace and are mainly used to provide information about the blockchain or are general-use utility functions.

Block and Transaction Properties

- `blockhash(uint blockNumber) returns (bytes32)`: hash of the given block - only works for 256 most recent, excluding current, blocks
- `block.coinbase (address payable)`: current block miner's address
- `block.difficulty (uint)`: current block difficulty
- `block.gaslimit (uint)`: current block gaslimit
- `block.number (uint)`: current block number
- `block.timestamp (uint)`: current block timestamp as seconds since unix epoch
- `gasleft() returns (uint256)`: remaining gas
- `msg.data (bytes calldata)`: complete calldata
- `msg.sender (address payable)`: sender of the message (current call)
- `msg.sig (bytes4)`: first four bytes of the calldata (i.e. function identifier)
- `msg.value (uint)`: number of wei sent with the message
- `tx.gasprice (uint)`: gas price of the transaction
- `tx.origin (address payable)`: sender of the transaction (full call chain)

Note: The values of all members of `msg`, including `msg.sender` and `msg.value` can change for every **external** function call. This includes calls to library functions.

Note: Do not rely on `block.timestamp` or `blockhash` as a source of randomness, unless you know what you are doing.

Both the timestamp and the block hash can be influenced by miners to some degree. Bad actors in the mining community can for example run a casino payout function on a chosen hash and just retry a different hash if they did not receive any money.

The current block timestamp must be strictly larger than the timestamp of the last block, but the only guarantee is that it will be somewhere between the timestamps of two consecutive blocks in the canonical chain.

Note: The block hashes are not available for all blocks for scalability reasons. You can only access the hashes of the most recent 256 blocks, all other values will be zero.

Note: The function `blockhash` was previously known as `block.blockhash`, which was deprecated in version 0.4.22 and removed in version 0.5.0.

Note: The function `gasleft` was previously known as `msg.gas`, which was deprecated in version 0.4.21 and removed in version 0.5.0.

Note: In version 0.7.0, the alias `now` (for `block.timestamp`) was removed.

ABI Encoding and Decoding Functions

- `abi.decode(bytes memory encodedData, (...)) returns (...)`: ABI-decodes the given data, while the types are given in parentheses as second argument. Example: `(uint a, uint[2] memory b, bytes memory c) = abi.decode(data, (uint, uint[2], bytes))`

- `abi.encode(...) returns (bytes memory)`: ABI-encodes the given arguments

- `abi.encodePacked(...) returns (bytes memory)`: Performs *packed encoding* of the given arguments. Note that packed encoding can be ambiguous!

- `abi.encodeWithSelector(bytes4 selector, ...) returns (bytes memory)`:   ABI-encodes the given arguments starting from the second and prepends the given four-byte selector

- `abi.encodeWithSignature(string memory signature, ...) returns (bytes memory)`: Equivalent to `abi.encodeWithSelector(bytes4(keccak256(bytes(signature)))), ...)` `

Note: These encoding functions can be used to craft data for external function calls without actually calling an external function. Furthermore, `keccak256(abi.encodePacked(a, b))` is a way to compute the hash of structured data (although be aware that it is possible to craft a "hash collision" using different function parameter types).

See the documentation about the *ABI* and the *tightly packed encoding* for details about the encoding.

Error Handling

See the dedicated section on *assert and require* for more details on error handling and when to use which function.

assert(bool condition) causes an invalid opcode and thus state change reversion if the condition is not met - to be used for internal errors.

require(bool condition) reverts if the condition is not met - to be used for errors in inputs or external components.

require(bool condition, string memory message) reverts if the condition is not met - to be used for errors in inputs or external components. Also provides an error message.

revert() abort execution and revert state changes

revert(string memory reason) abort execution and revert state changes, providing an explanatory string

Mathematical and Cryptographic Functions

addmod(uint x, uint y, uint k) returns (uint) compute `(x + y) % k` where the addition is performed with arbitrary precision and does not wrap around at `2**256`. Assert that `k != 0` starting from version 0.5.0.

---

mulmod(uint x, uint y, uint k) returns (uint) compute `(x * y) % k` where the multiplication is performed with arbitrary precision and does not wrap around at `2**256`. Assert that `k != 0` starting from version 0.5.0.

keccak256(bytes memory) returns (bytes32) compute the Keccak-256 hash of the input

---

Note: There used to be an alias for `keccak256` called `sha3`, which was removed in version 0.5.0.

---

sha256(bytes memory) returns (bytes32) compute the SHA-256 hash of the input

ripemd160(bytes memory) returns (bytes20) compute RIPEMD-160 hash of the input

ecrecover(bytes32 hash, uint8 v, bytes32 r, bytes32 s) returns (address) recover the address associated with the public key from elliptic curve signature or return zero on error. The function parameters correspond to ECDSA values of the signature:

- `r` = first 32 bytes of signature

- `s` = second 32 bytes of signature

- `v` = final 1 byte of signature

`ecrecover` returns an `address`, and not an `address payable`. See *address payable* for conversion, in case you need to transfer funds to the recovered address.

For further details, read example usage.

Warning: If you use `ecrecover`, be aware that a valid signature can be turned into a different valid signature without requiring knowledge of the corresponding private key. In the Homestead hard fork, this issue was fixed for _transaction_ signatures (see EIP-2), but the ecrecover function remained unchanged.

This is usually not a problem unless you require signatures to be unique or use them to identify items. OpenZeppelin have a ECDSA helper library that you can use as a wrapper for `ecrecover` without this issue.

Note: When running `sha256`, `ripemd160` or `ecrecover` on a *private blockchain*, you might encounter Out-of-Gas. This is because these functions are implemented as "precompiled contracts" and only really exist after they receive the first message (although their contract code is hardcoded). Messages to non-existing contracts are more expensive and thus the execution might run into an Out-of-Gas error. A workaround for this problem is to first send Wei (1 for example) to each of the contracts before you use them in your actual contracts. This is not an issue on the main or test net.

Members of Address Types

<address>.balance(uint256) balance of the *Address* in Wei

<address payable>.transfer(uint256 amount) send given amount of Wei to *Address*, reverts on failure, forwards 2300 gas stipend, not adjustable

<address payable>.send(uint256 amount) returns (bool) send given amount of Wei to *Address*, returns `false` on failure, forwards 2300 gas stipend, not adjustable

<address>.call(bytes memory) returns (bool, bytes memory) issue low-level `CALL` with the given payload, returns success condition and return data, forwards all available gas, adjustable

---

`<address>.delegatecall(bytes memory) returns (bool, bytes memory)` issue low-level
DELEGATECALL with the given payload, returns success condition and return data, forwards all available gas,
adjustable

`<address>.staticcall(bytes memory) returns (bool, bytes memory)` issue low-level
STATICCALL with the given payload, returns success condition and return data, forwards all available gas,
adjustable

For more information, see the section on *Address*.

> **Warning:** You should avoid using `.call()` whenever possible when executing another contract function as it
> bypasses type checking, function existence check, and argument packing.

> **Warning:** There are some dangers in using `send`: The transfer fails if the call stack depth is at 1024 (this can
> always be forced by the caller) and it also fails if the recipient runs out of gas. So in order to make safe Ether
> transfers, always check the return value of `send`, use `transfer` or even better: Use a pattern where the recipient
> withdraws the money.

> **Note:** Prior to version 0.5.0, Solidity allowed address members to be accessed by a contract instance, for example
> `this.balance`. This is now forbidden and an explicit conversion to address must be done: `address(this).
> balance`.

> **Note:** If state variables are accessed via a low-level delegatecall, the storage layout of the two contracts must align in
> order for the called contract to correctly access the storage variables of the calling contract by name. This is of course
> not the case if storage pointers are passed as function arguments as in the case for the high-level libraries.

> **Note:** Prior to version 0.5.0, `.call`, `.delegatecall` and `.staticcall` only returned the success condition
> and not the return data.

> **Note:** Prior to version 0.5.0, there was a member called `callcode` with similar but slightly different semantics than
> `delegatecall`.

Contract Related

`this` (current contract's type) the current contract, explicitly convertible to *Address*

`selfdestruct(address payable recipient)` Destroy the current contract, sending its funds to the given
Address and end execution. Note that `selfdestruct` has some peculiarities inherited from the EVM:

- the receiving contract's receive function is not executed.

- the contract is only really destroyed at the end of the transaction and `revert` s might "undo" the destruc-
tion.

Furthermore, all functions of the current contract are callable directly including the current function.

Note: Prior to version 0.5.0, there was a function called `suicide` with the same semantics as `selfdestruct`.

Type Information

The expression `type(X)` can be used to retrieve information about the type `X`. Currently, there is limited support for this feature (`X` can be either a contract or an integer type) but it might be expanded in the future.

The following properties are available for a contract type `C`:

`type(C).name` The name of the contract.

`type(C).creationCode` Memory byte array that contains the creation bytecode of the contract. This can be used in inline assembly to build custom creation routines, especially by using the `create2` opcode. This property can **not** be accessed in the contract itself or any derived contract. It causes the bytecode to be included in the bytecode of the call site and thus circular references like that are not possible.

`type(C).runtimeCode` Memory byte array that contains the runtime bytecode of the contract. This is the code that is usually deployed by the constructor of `C`. If `C` has a constructor that uses inline assembly, this might be different from the actually deployed bytecode. Also note that libraries modify their runtime bytecode at time of deployment to guard against regular calls. The same restrictions as with `.creationCode` also apply for this property.

In addition to the properties above, the following properties are available for an interface type `I`:

`type(I).interfaceId:` A `bytes4` value containing the [EIP-165](#) interface identifier of the given interface `I`. This identifier is defined as the `XOR` of all function selectors defined within the interface itself - excluding all inherited functions.

The following properties are available for an integer type `T`:

`type(T).min` The smallest value representable by type `T`.

`type(T).max` The largest value representable by type `T`.

3.8 Expressions and Control Structures

3.8.1 Control Structures

Most of the control structures known from curly-braces languages are available in Solidity:

There is: `if`, `else`, `while`, `do`, `for`, `break`, `continue`, `return`, with the usual semantics known from C or JavaScript.

Solidity also supports exception handling in the form of `try`/`catch`-statements, but only for *external function calls* and contract creation calls.

Parentheses can *not* be omitted for conditionals, but curly braces can be omitted around single-statement bodies.

Note that there is no type conversion from non-boolean to boolean types as there is in C and JavaScript, so `if (1) { ... }` is *not* valid Solidity.

3.8.2 Function Calls

Internal Function Calls

Functions of the current contract can be called directly ("internally"), also recursively, as seen in this nonsensical example:

```
// SPDX-License-Identifier: GPL-3.0
pragma solidity >=0.4.22 <0.8.0;

contract C {
    function g(uint a) public pure returns (uint ret) { return a + f(); }
    function f() internal pure returns (uint ret) { return g(7) + f(); }
}
```

These function calls are translated into simple jumps inside the EVM. This has the effect that the current memory is not cleared, i.e. passing memory references to internally-called functions is very efficient. Only functions of the same contract instance can be called internally.

You should still avoid excessive recursion, as every internal function call uses up at least one stack slot and there are only 1024 slots available.

External Function Calls

The expressions `this.g(8);` and `c.g(2);` (where `c` is a contract instance) are also valid function calls, but this time, the function will be called "externally", via a message call and not directly via jumps. Please note that function calls on `this` cannot be used in the constructor, as the actual contract has not been created yet.

Functions of other contracts have to be called externally. For an external call, all function arguments have to be copied to memory.

Note: A function call from one contract to another does not create its own transaction, it is a message call as part of the overall transaction.

When calling functions of other contracts, you can specify the amount of Wei or gas sent with the call with the special options `{value: 10, gas: 10000}`. Note that it is discouraged to specify gas values explicitly, since the gas costs of opcodes can change in the future. Any Wei you send to the contract is added to the total balance of that contract:

```
// SPDX-License-Identifier: GPL-3.0
pragma solidity >=0.6.2 <0.8.0;

contract InfoFeed {
    function info() public payable returns (uint ret) { return 42; }
}

contract Consumer {
    InfoFeed feed;
    function setFeed(InfoFeed addr) public { feed = addr; }
    function callFeed() public { feed.info{value: 10, gas: 800}(); }
}
```

You need to use the modifier `payable` with the `info` function because otherwise, the `value` option would not be available.

> **Warning:** Be careful that `feed.info{value: 10, gas: 800}` only locally sets the `value` and amount of `gas` sent with the function call, and the parentheses at the end perform the actual call. So in this case, the function is not called and the `value` and `gas` settings are lost.

Due to the fact that the EVM considers a call to a non-existing contract to always succeed, Solidity uses the `extcodesize` opcode to check that the contract that is about to be called actually exists (it contains code) and causes an exception if it does not.

Function calls also cause exceptions if the called contract itself throws an exception or goes out of gas.

> **Warning:** Any interaction with another contract imposes a potential danger, especially if the source code of the contract is not known in advance. The current contract hands over control to the called contract and that may potentially do just about anything. Even if the called contract inherits from a known parent contract, the inheriting contract is only required to have a correct interface. The implementation of the contract, however, can be completely arbitrary and thus, pose a danger. In addition, be prepared in case it calls into other contracts of your system or even back into the calling contract before the first call returns. This means that the called contract can change state variables of the calling contract via its functions. Write your functions in a way that, for example, calls to external functions happen after any changes to state variables in your contract so your contract is not vulnerable to a reentrancy exploit.

> **Note:** Before Solidity 0.6.2, the recommended way to specify the value and gas was to use `f.value(x).gas(g)()`. This was deprecated in Solidity 0.6.2 and is no longer possible since Solidity 0.7.0.

Named Calls and Anonymous Function Parameters

Function call arguments can be given by name, in any order, if they are enclosed in `{ }` as can be seen in the following example. The argument list has to coincide by name with the list of parameters from the function declaration, but can be in arbitrary order.

```solidity
// SPDX-License-Identifier: GPL-3.0
pragma solidity >=0.4.0 <0.8.0;

contract C {
    mapping(uint => uint) data;

    function f() public {
        set({value: 2, key: 3});
    }

    function set(uint key, uint value) public {
        data[key] = value;
    }

}
```

Omitted Function Parameter Names

The names of unused parameters (especially return parameters) can be omitted. Those parameters will still be present on the stack, but they are inaccessible.

```
// SPDX-License-Identifier: GPL-3.0
pragma solidity >=0.4.22 <0.8.0;

contract C {
    // omitted name for parameter
    function func(uint k, uint) public pure returns(uint) {
        return k;
    }
}
```

3.8.3 Creating Contracts via `new`

A contract can create other contracts using the `new` keyword. The full code of the contract being created has to be known when the creating contract is compiled so recursive creation-dependencies are not possible.

```
// SPDX-License-Identifier: GPL-3.0
pragma solidity ^0.7.0;

contract D {
    uint public x;
    constructor(uint a) payable {
        x = a;
    }
}

contract C {
    D d = new D(4); // will be executed as part of C's constructor

    function createD(uint arg) public {
        D newD = new D(arg);
        newD.x();
    }

    function createAndEndowD(uint arg, uint amount) public payable {
        // Send ether along with the creation
        D newD = new D{value: amount}(arg);
        newD.x();
    }
}
```

As seen in the example, it is possible to send Ether while creating an instance of `D` using the `value` option, but it is not possible to limit the amount of gas. If the creation fails (due to out-of-stack, not enough balance or other problems), an exception is thrown.

Salted contract creations / create2

When creating a contract, the address of the contract is computed from the address of the creating contract and a counter that is increased with each contract creation.

If you specify the option `salt` (a bytes32 value), then contract creation will use a different mechanism to come up with the address of the new contract:

It will compute the address from the address of the creating contract, the given salt value, the (creation) bytecode of the created contract and the constructor arguments.

In particular, the counter ("nonce") is not used. This allows for more flexibility in creating contracts: You are able to derive the address of the new contract before it is created. Furthermore, you can rely on this address also in case the creating contracts creates other contracts in the meantime.

The main use-case here is contracts that act as judges for off-chain interactions, which only need to be created if there is a dispute.

```solidity
// SPDX-License-Identifier: GPL-3.0
pragma solidity ^0.7.0;

contract D {
    uint public x;
    constructor(uint a) {
        x = a;
    }
}

contract C {
    function createDSalted(bytes32 salt, uint arg) public {
        // This complicated expression just tells you how the address
        // can be pre-computed. It is just there for illustration.
        // You actually only need ``new D{salt: salt}(arg)``.
        address predictedAddress = address(uint(keccak256(abi.encodePacked(
            byte(0xff),
            address(this),
            salt,
            keccak256(abi.encodePacked(
                type(D).creationCode,
                arg
            ))
        )))));

        D d = new D{salt: salt}(arg);
        require(address(d) == predictedAddress);
    }
}
```

> **Warning:** There are some peculiarities in relation to salted creation. A contract can be re-created at the same address after having been destroyed. Yet, it is possible for that newly created contract to have a different deployed bytecode even though the creation bytecode has been the same (which is a requirement because otherwise the address would change). This is due to the fact that the compiler can query external state that might have changed between the two creations and incorporate that into the deployed bytecode before it is stored.

3.8.4 Order of Evaluation of Expressions

The evaluation order of expressions is not specified (more formally, the order in which the children of one node in the expression tree are evaluated is not specified, but they are of course evaluated before the node itself). It is only guaranteed that statements are executed in order and short-circuiting for boolean expressions is done.

3.8.5 Assignment

Destructuring Assignments and Returning Multiple Values

Solidity internally allows tuple types, i.e. a list of objects of potentially different types whose number is a constant at compile-time. Those tuples can be used to return multiple values at the same time. These can then either be assigned to newly declared variables or to pre-existing variables (or LValues in general).

Tuples are not proper types in Solidity, they can only be used to form syntactic groupings of expressions.

```solidity
// SPDX-License-Identifier: GPL-3.0
pragma solidity >=0.5.0 <0.8.0;

contract C {
    uint index;

    function f() public pure returns (uint, bool, uint) {
        return (7, true, 2);
    }

    function g() public {
        // Variables declared with type and assigned from the returned tuple,
        // not all elements have to be specified (but the number must match).
        (uint x, , uint y) = f();
        // Common trick to swap values -- does not work for non-value storage types.
        (x, y) = (y, x);
        // Components can be left out (also for variable declarations).
        (index, , ) = f(); // Sets the index to 7
    }
}
```

It is not possible to mix variable declarations and non-declaration assignments, i.e. the following is not valid: `(x, uint y) = (1, 2);`

Note: Prior to version 0.5.0 it was possible to assign to tuples of smaller size, either filling up on the left or on the right side (which ever was empty). This is now disallowed, so both sides have to have the same number of components.

Warning: Be careful when assigning to multiple variables at the same time when reference types are involved, because it could lead to unexpected copying behaviour.

Complications for Arrays and Structs

The semantics of assignments are more complicated for non-value types like arrays and structs, including `bytes` and `string`, see *Data location and assignment behaviour* for details.

In the example below the call to `g(x)` has no effect on `x` because it creates an independent copy of the storage value in memory. However, `h(x)` successfully modifies `x` because only a reference and not a copy is passed.

```solidity
// SPDX-License-Identifier: GPL-3.0
pragma solidity >=0.4.22 <0.8.0;

contract C {
    uint[20] x;

    function f() public {
```

```
        g(x);
        h(x);
    }

    function g(uint[20] memory y) internal pure {
        y[2] = 3;
    }

    function h(uint[20] storage y) internal {
        y[3] = 4;
    }
}
```

3.8.6 Scoping and Declarations

A variable which is declared will have an initial default value whose byte-representation is all zeros. The "default values" of variables are the typical "zero-state" of whatever the type is. For example, the default value for a `bool` is `false`. The default value for the `uint` or `int` types is `0`. For statically-sized arrays and `bytes1` to `bytes32`, each individual element will be initialized to the default value corresponding to its type. For dynamically-sized arrays, `bytes` and `string`, the default value is an empty array or string. For the `enum` type, the default value is its first member.

Scoping in Solidity follows the widespread scoping rules of C99 (and many other languages): Variables are visible from the point right after their declaration until the end of the smallest `{ }`-block that contains the declaration. As an exception to this rule, variables declared in the initialization part of a for-loop are only visible until the end of the for-loop.

Variables that are parameter-like (function parameters, modifier parameters, catch parameters, . . .) are visible inside the code block that follows - the body of the function/modifier for a function and modifier parameter and the catch block for a catch parameter.

Variables and other items declared outside of a code block, for example functions, contracts, user-defined types, etc., are visible even before they were declared. This means you can use state variables before they are declared and call functions recursively.

As a consequence, the following examples will compile without warnings, since the two variables have the same name but disjoint scopes.

```
// SPDX-License-Identifier: GPL-3.0
pragma solidity >=0.5.0 <0.8.0;
contract C {
    function minimalScoping() pure public {
        {
            uint same;
            same = 1;
        }

        {
            uint same;
            same = 3;
        }
    }
}
```

As a special example of the C99 scoping rules, note that in the following, the first assignment to `x` will actually assign the outer and not the inner variable. In any case, you will get a warning about the outer variable being shadowed.

```
// SPDX-License-Identifier: GPL-3.0
pragma solidity >=0.5.0 <0.8.0;
// This will report a warning
contract C {
    function f() pure public returns (uint) {
        uint x = 1;
        {
            x = 2; // this will assign to the outer variable
            uint x;
        }
        return x; // x has value 2
    }
}
```

> **Warning:** Before version 0.5.0 Solidity followed the same scoping rules as JavaScript, that is, a variable declared anywhere within a function would be in scope for the entire function, regardless where it was declared. The following example shows a code snippet that used to compile but leads to an error starting from version 0.5.0.

```
// SPDX-License-Identifier: GPL-3.0
pragma solidity >=0.5.0 <0.8.0;
// This will not compile
contract C {
    function f() pure public returns (uint) {
        x = 2;
        uint x;
        return x;
    }
}
```

3.8.7 Error handling: Assert, Require, Revert and Exceptions

Solidity uses state-reverting exceptions to handle errors. Such an exception undoes all changes made to the state in the current call (and all its sub-calls) and flags an error to the caller.

When exceptions happen in a sub-call, they "bubble up" (i.e., exceptions are rethrown) automatically. Exceptions to this rule are `send` and the low-level functions `call`, `delegatecall` and `staticcall`: they return `false` as their first return value in case of an exception instead of "bubbling up".

> **Warning:** The low-level functions `call`, `delegatecall` and `staticcall` return `true` as their first return value if the account called is non-existent, as part of the design of the EVM. Account existence must be checked prior to calling if needed.

Exceptions can be caught with the `try`/`catch` statement.

assert and require

The convenience functions `assert` and `require` can be used to check for conditions and throw an exception if the condition is not met.

The `assert` function should only be used to test for internal errors, and to check invariants. Properly functioning code should never reach a failing `assert` statement; if this happens there is a bug in your contract which you should

fix. Language analysis tools can evaluate your contract to identify the conditions and function calls which will reach a failing `assert`.

An `assert`-style exception is generated in the following situations:

1. If you access an array or an array slice at a too large or negative index (i.e. `x[i]` where `i >= x.length` or `i < 0`).

2. If you access a fixed-length `bytesN` at a too large or negative index.

3. If you divide or modulo by zero (e.g. `5 / 0` or `23 % 0`).

4. If you convert a value too big or negative into an enum type.

5. If you call a zero-initialized variable of internal function type.

6. If you call `assert` with an argument that evaluates to false.

The `require` function should be used to ensure valid conditions that cannot be detected until execution time. This includes conditions on inputs or return values from calls to external contracts.

A `require`-style exception is generated in the following situations:

1. Calling `require` with an argument that evaluates to `false`.

2. If you call a function via a message call but it does not finish properly (i.e., it runs out of gas, has no matching function, or throws an exception itself), except when a low level operation `call`, `send`, `delegatecall`, `callcode` or `staticcall` is used. The low level operations never throw exceptions but indicate failures by returning `false`.

3. If you create a contract using the `new` keyword but the contract creation *does not finish properly*.

4. If you perform an external function call targeting a contract that contains no code.

5. If your contract receives Ether via a public function without `payable` modifier (including the constructor and the fallback function).

6. If your contract receives Ether via a public getter function.

7. If a `.transfer()` fails.

You can optionally provide a message string for `require`, but not for `assert`.

The following example shows how you can use `require` to check conditions on inputs and `assert` for internal error checking.

```solidity
// SPDX-License-Identifier: GPL-3.0
pragma solidity >=0.5.0 <0.8.0;

contract Sharer {
    function sendHalf(address payable addr) public payable returns (uint balance) {
        require(msg.value % 2 == 0, "Even value required.");
        uint balanceBeforeTransfer = address(this).balance;
        addr.transfer(msg.value / 2);
        // Since transfer throws an exception on failure and
        // cannot call back here, there should be no way for us to
        // still have half of the money.
        assert(address(this).balance == balanceBeforeTransfer - msg.value / 2);
        return address(this).balance;
    }
}
```

Internally, Solidity performs a revert operation (instruction `0xfd`) for a `require`-style exception and executes an invalid operation (instruction `0xfe`) to throw an `assert`-style exception. In both cases, this causes the EVM to revert

all changes made to the state. The reason for reverting is that there is no safe way to continue execution, because an expected effect did not occur. Because we want to keep the atomicity of transactions, the safest action is to revert all changes and make the whole transaction (or at least call) without effect.

In both cases, the caller can react on such failures using `try`/`catch` (in the failing `assert`-style exception only if enough gas is left), but the changes in the caller will always be reverted.

Note: `assert`-style exceptions consume all gas available to the call, while `require`-style exceptions do not consume any gas starting from the Metropolis release.

---

revert

The `revert` function is another way to trigger exceptions from within other code blocks to flag an error and revert the current call. The function takes an optional string message containing details about the error that is passed back to the caller.

The following example shows how to use an error string together with `revert` and the equivalent `require`:

```solidity
// SPDX-License-Identifier: GPL-3.0
pragma solidity >=0.5.0 <0.8.0;

contract VendingMachine {
    function buy(uint amount) public payable {
        if (amount > msg.value / 2 ether)
            revert("Not enough Ether provided.");
        // Alternative way to do it:
        require(
            amount <= msg.value / 2 ether,
            "Not enough Ether provided."
        );
        // Perform the purchase.
    }
}
```

If you provide the reason string directly, then the two syntax options are equivalent, it is the developer's preference which one to use.

Note: The `require` function is evaluated just as any other function. This means that all arguments are evaluated before the function itself is executed. In particular, in `require(condition, f())` the function `f` is executed even if `condition` is true.

The provided string is *abi-encoded* as if it were a call to a function `Error(string)`. In the above example, `revert("Not enough Ether provided.");` returns the following hexadecimal as error return data:

```
0x08c379a0                                                          // Function␣
↪selector for Error(string)
0x0000000000000000000000000000000000000000000000000000000000000020 // Data offset
0x000000000000000000000000000000000000000000000000000000000000001a // String length
0x4e6f7420656e6f75676820457468657220070726f76696465642e000000000000 // String data
```

The provided message can be retrieved by the caller using `try`/`catch` as shown below.

Note: There used to be a keyword called `throw` with the same semantics as `revert()` which was deprecated in

---

version 0.4.13 and removed in version 0.5.0.

`try`/`catch`

A failure in an external call can be caught using a try/catch statement, as follows:

```solidity
// SPDX-License-Identifier: GPL-3.0
pragma solidity >=0.6.0 <0.8.0;

interface DataFeed { function getData(address token) external returns (uint value); }

contract FeedConsumer {
    DataFeed feed;
    uint errorCount;
    function rate(address token) public returns (uint value, bool success) {
        // Permanently disable the mechanism if there are
        // more than 10 errors.
        require(errorCount < 10);
        try feed.getData(token) returns (uint v) {
            return (v, true);
        } catch Error(string memory /*reason*/) {
            // This is executed in case
            // revert was called inside getData
            // and a reason string was provided.
            errorCount++;
            return (0, false);
        } catch (bytes memory /*lowLevelData*/) {
            // This is executed in case revert() was used
            // or there was a failing assertion, division
            // by zero, etc. inside getData.
            errorCount++;
            return (0, false);
        }
    }
}
```

The `try` keyword has to be followed by an expression representing an external function call or a contract creation (`new ContractName()`). Errors inside the expression are not caught (for example if it is a complex expression that also involves internal function calls), only a revert happening inside the external call itself. The `returns` part (which is optional) that follows declares return variables matching the types returned by the external call. In case there was no error, these variables are assigned and the contract's execution continues inside the first success block. If the end of the success block is reached, execution continues after the `catch` blocks.

Currently, Solidity supports different kinds of catch blocks depending on the type of error. If the error was caused by `revert("reasonString")` or `require(false, "reasonString")` (or an internal error that causes such an exception), then the catch clause of the type `catch Error(string memory reason)` will be executed.

It is planned to support other types of error data in the future. The string `Error` is currently parsed as is and is not treated as an identifier.

The clause `catch (bytes memory lowLevelData)` is executed if the error signature does not match any other clause, there was an error during decoding of the error message, if there was a failing assertion in the external call (for example due to a division by zero or a failing `assert()`) or if no error data was provided with the exception. The declared variable provides access to the low-level error data in that case.

If you are not interested in the error data, you can just use `catch { ... }` (even as the only catch clause).

In order to catch all error cases, you have to have at least the clause `catch { ...}` or the clause `catch (bytes memory lowLevelData) { ... }`.

The variables declared in the `returns` and the `catch` clause are only in scope in the block that follows.

Note: If an error happens during the decoding of the return data inside a try/catch-statement, this causes an exception in the currently executing contract and because of that, it is not caught in the catch clause. If there is an error during decoding of `catch Error(string memory reason)` and there is a low-level catch clause, this error is caught there.

Note: If execution reaches a catch-block, then the state-changing effects of the external call have been reverted. If execution reaches the success block, the effects were not reverted. If the effects have been reverted, then execution either continues in a catch block or the execution of the try/catch statement itself reverts (for example due to decoding failures as noted above or due to not providing a low-level catch clause).

Note: The reason behind a failed call can be manifold. Do not assume that the error message is coming directly from the called contract: The error might have happened deeper down in the call chain and the called contract just forwarded it. Also, it could be due to an out-of-gas situation and not a deliberate error condition: The caller always retains 63/64th of the gas in a call and thus even if the called contract goes out of gas, the caller still has some gas left.

3.9 Contracts

Contracts in Solidity are similar to classes in object-oriented languages. They contain persistent data in state variables, and functions that can modify these variables. Calling a function on a different contract (instance) will perform an EVM function call and thus switch the context such that state variables in the calling contract are inaccessible. A contract and its functions need to be called for anything to happen. There is no "cron" concept in Ethereum to call a function at a particular event automatically.

3.9.1 Creating Contracts

Contracts can be created "from outside" via Ethereum transactions or from within Solidity contracts.

IDEs, such as Remix, make the creation process seamless using UI elements.

One way to create contracts programmatically on Ethereum is via the JavaScript API web3.js. It has a function called web3.eth.Contract to facilitate contract creation.

When a contract is created, its *constructor* (a function declared with the `constructor` keyword) is executed once.

A constructor is optional. Only one constructor is allowed, which means overloading is not supported.

After the constructor has executed, the final code of the contract is stored on the blockchain. This code includes all public and external functions and all functions that are reachable from there through function calls. The deployed code does not include the constructor code or internal functions only called from the constructor.

Internally, constructor arguments are passed *ABI encoded* after the code of the contract itself, but you do not have to care about this if you use `web3.js`.

If a contract wants to create another contract, the source code (and the binary) of the created contract has to be known to the creator. This means that cyclic creation dependencies are impossible.

```solidity
// SPDX-License-Identifier: GPL-3.0
pragma solidity >=0.4.22 <0.8.0;


contract OwnedToken {
    // `TokenCreator` is a contract type that is defined below.
    // It is fine to reference it as long as it is not used
    // to create a new contract.
    TokenCreator creator;
    address owner;
    bytes32 name;

    // This is the constructor which registers the
    // creator and the assigned name.
    constructor(bytes32 _name) {
        // State variables are accessed via their name
        // and not via e.g. `this.owner`. Functions can
        // be accessed directly or through `this.f`,
        // but the latter provides an external view
        // to the function. Especially in the constructor,
        // you should not access functions externally,
        // because the function does not exist yet.
        // See the next section for details.
        owner = msg.sender;

        // We perform an explicit type conversion from `address`
        // to `TokenCreator` and assume that the type of
        // the calling contract is `TokenCreator`, there is
        // no real way to verify that.
        // This does not create a new contract.
        creator = TokenCreator(msg.sender);
        name = _name;
    }

    function changeName(bytes32 newName) public {
        // Only the creator can alter the name.
        // We compare the contract based on its
        // address which can be retrieved by
        // explicit conversion to address.
        if (msg.sender == address(creator))
            name = newName;
    }

    function transfer(address newOwner) public {
        // Only the current owner can transfer the token.
        if (msg.sender != owner) return;

        // We ask the creator contract if the transfer
        // should proceed by using a function of the
        // `TokenCreator` contract defined below. If
        // the call fails (e.g. due to out-of-gas),
        // the execution also fails here.
        if (creator.isTokenTransferOK(owner, newOwner))
            owner = newOwner;
    }
}
```

```solidity
contract TokenCreator {
    function createToken(bytes32 name)
        public
        returns (OwnedToken tokenAddress)
    {
        // Create a new `Token` contract and return its address.
        // From the JavaScript side, the return type
        // of this function is `address`, as this is
        // the closest type available in the ABI.
        return new OwnedToken(name);
    }

    function changeName(OwnedToken tokenAddress, bytes32 name) public {
        // Again, the external type of `tokenAddress` is
        // simply `address`.
        tokenAddress.changeName(name);
    }

    // Perform checks to determine if transferring a token to the
    // `OwnedToken` contract should proceed
    function isTokenTransferOK(address currentOwner, address newOwner)
        public
        pure
        returns (bool ok)
    {
        // Check an arbitrary condition to see if transfer should proceed
        return keccak256(abi.encodePacked(currentOwner, newOwner))[0] == 0x7f;
    }
}
```

3.9.2 Visibility and Getters

Solidity knows two kinds of function calls: internal ones that do not create an actual EVM call (also called a "message call") and external ones that do. Because of that, there are four types of visibility for functions and state variables.

Functions have to be specified as being external, public, internal or private. For state variables, external is not possible.

external External functions are part of the contract interface, which means they can be called from other contracts and via transactions. An external function f cannot be called internally (i.e. f() does not work, but this.f() works). External functions are sometimes more efficient when they receive large arrays of data, because the data is not copied from calldata to memory.

public Public functions are part of the contract interface and can be either called internally or via messages. For public state variables, an automatic getter function (see below) is generated.

internal Those functions and state variables can only be accessed internally (i.e. from within the current contract or contracts deriving from it), without using this.

private Private functions and state variables are only visible for the contract they are defined in and not in derived contracts.

Note: Everything that is inside a contract is visible to all observers external to the blockchain. Making something private only prevents other contracts from reading or modifying the information, but it will still be visible to the

whole world outside of the blockchain.

The visibility specifier is given after the type for state variables and between parameter list and return parameter list for functions.

```solidity
// SPDX-License-Identifier: GPL-3.0
pragma solidity >=0.4.16 <0.8.0;

contract C {
    function f(uint a) private pure returns (uint b) { return a + 1; }
    function setData(uint a) internal { data = a; }
    uint public data;
}
```

In the following example, `D`, can call `c.getData()` to retrieve the value of `data` in state storage, but is not able to call `f`. Contract `E` is derived from `C` and, thus, can call `compute`.

```solidity
// SPDX-License-Identifier: GPL-3.0
pragma solidity >=0.4.16 <0.8.0;

contract C {
    uint private data;

    function f(uint a) private pure returns(uint b) { return a + 1; }
    function setData(uint a) public { data = a; }
    function getData() public view returns(uint) { return data; }
    function compute(uint a, uint b) internal pure returns (uint) { return a + b; }
}

// This will not compile
contract D {
    function readData() public {
        C c = new C();
        uint local = c.f(7); // error: member `f` is not visible
        c.setData(3);
        local = c.getData();
        local = c.compute(3, 5); // error: member `compute` is not visible
    }
}

contract E is C {
    function g() public {
        C c = new C();
        uint val = compute(3, 5); // access to internal member (from derived to
↪parent contract)
    }
}
```

Getter Functions

The compiler automatically creates getter functions for all **public** state variables. For the contract given below, the compiler will generate a function called `data` that does not take any arguments and returns a `uint`, the value of the state variable `data`. State variables can be initialized when they are declared.

```solidity
// SPDX-License-Identifier: GPL-3.0
pragma solidity >=0.4.16 <0.8.0;

contract C {
    uint public data = 42;
}

contract Caller {
    C c = new C();
    function f() public view returns (uint) {
        return c.data();
    }
}
```

The getter functions have external visibility. If the symbol is accessed internally (i.e. without `this.`), it evaluates to a state variable. If it is accessed externally (i.e. with `this.`), it evaluates to a function.

```solidity
// SPDX-License-Identifier: GPL-3.0
pragma solidity >=0.4.0 <0.8.0;

contract C {
    uint public data;
    function x() public returns (uint) {
        data = 3; // internal access
        return this.data(); // external access
    }
}
```

If you have a `public` state variable of array type, then you can only retrieve single elements of the array via the generated getter function. This mechanism exists to avoid high gas costs when returning an entire array. You can use arguments to specify which individual element to return, for example `data(0)`. If you want to return an entire array in one call, then you need to write a function, for example:

```solidity
// SPDX-License-Identifier: GPL-3.0
pragma solidity >=0.4.16 <0.8.0;

contract arrayExample {
    // public state variable
    uint[] public myArray;

    // Getter function generated by the compiler
    /*
    function myArray(uint i) public view returns (uint) {
        return myArray[i];
    }
    */

    // function that returns entire array
    function getArray() public view returns (uint[] memory) {
        return myArray;
    }
}
```

Now you can use `getArray()` to retrieve the entire array, instead of `myArray(i)`, which returns a single element per call.

The next example is more complex:

```solidity
// SPDX-License-Identifier: GPL-3.0
pragma solidity >=0.4.0 <0.8.0;

contract Complex {
    struct Data {
        uint a;
        bytes3 b;
        mapping (uint => uint) map;
    }
    mapping (uint => mapping(bool => Data[])) public data;
}
```

It generates a function of the following form. The mapping in the struct is omitted because there is no good way to provide the key for the mapping:

```solidity
function data(uint arg1, bool arg2, uint arg3) public returns (uint a, bytes3 b) {
    a = data[arg1][arg2][arg3].a;
    b = data[arg1][arg2][arg3].b;
}
```

3.9.3 Function Modifiers

Modifiers can be used to change the behaviour of functions in a declarative way. For example, you can use a modifier to automatically check a condition prior to executing the function.

Modifiers are inheritable properties of contracts and may be overridden by derived contracts, but only if they are marked `virtual`. For details, please see *Modifier Overriding*.

```solidity
// SPDX-License-Identifier: GPL-3.0
pragma solidity >0.7.0 <0.8.0;

contract owned {
    constructor() { owner = msg.sender; }
    address payable owner;

    // This contract only defines a modifier but does not use
    // it: it will be used in derived contracts.
    // The function body is inserted where the special symbol
    // `_;` in the definition of a modifier appears.
    // This means that if the owner calls this function, the
    // function is executed and otherwise, an exception is
    // thrown.
    modifier onlyOwner {
        require(
            msg.sender == owner,
            "Only owner can call this function."
        );
        _;
    }
}

contract destructible is owned {
    // This contract inherits the `onlyOwner` modifier from
    // `owned` and applies it to the `destroy` function, which
    // causes that calls to `destroy` only have an effect if
    // they are made by the stored owner.
```

```solidity
    function destroy() public onlyOwner {
        selfdestruct(owner);
    }
}

contract priced {
    // Modifiers can receive arguments:
    modifier costs(uint price) {
        if (msg.value >= price) {
            _;
        }
    }
}

contract Register is priced, destructible {
    mapping (address => bool) registeredAddresses;
    uint price;

    constructor(uint initialPrice) { price = initialPrice; }

    // It is important to also provide the
    // `payable` keyword here, otherwise the function will
    // automatically reject all Ether sent to it.
    function register() public payable costs(price) {
        registeredAddresses[msg.sender] = true;
    }

    function changePrice(uint _price) public onlyOwner {
        price = _price;
    }
}

contract Mutex {
    bool locked;
    modifier noReentrancy() {
        require(
            !locked,
            "Reentrant call."
        );
        locked = true;
        _;
        locked = false;
    }

    /// This function is protected by a mutex, which means that
    /// reentrant calls from within `msg.sender.call` cannot call `f` again.
    /// The `return 7` statement assigns 7 to the return value but still
    /// executes the statement `locked = false` in the modifier.
    function f() public noReentrancy returns (uint) {
        (bool success,) = msg.sender.call("");
        require(success);
        return 7;
    }
}
```

Multiple modifiers are applied to a function by specifying them in a whitespace-separated list and are evaluated in the
order presented.

> **Warning:** In an earlier version of Solidity, `return` statements in functions having modifiers behaved differently.

Explicit returns from a modifier or function body only leave the current modifier or function body. Return variables are assigned and control flow continues after the "_" in the preceding modifier.

Arbitrary expressions are allowed for modifier arguments and in this context, all symbols visible from the function are visible in the modifier. Symbols introduced in the modifier are not visible in the function (as they might change by overriding).

3.9.4 Constant and Immutable State Variables

State variables can be declared as `constant` or `immutable`. In both cases, the variables cannot be modified after the contract has been constructed. For `constant` variables, the value has to be fixed at compile-time, while for `immutable`, it can still be assigned at construction time.

It is also possible to define `constant` variables at the file level.

The compiler does not reserve a storage slot for these variables, and every occurrence is replaced by the respective value.

Compared to regular state variables, the gas costs of constant and immutable variables are much lower. For a constant variable, the expression assigned to it is copied to all the places where it is accessed and also re-evaluated each time. This allows for local optimizations. Immutable variables are evaluated once at construction time and their value is copied to all the places in the code where they are accessed. For these values, 32 bytes are reserved, even if they would fit in fewer bytes. Due to this, constant values can sometimes be cheaper than immutable values.

Not all types for constants and immutables are implemented at this time. The only supported types are *strings* (only for constants) and *value types*.

```solidity
// SPDX-License-Identifier: GPL-3.0
pragma solidity >0.7.2;

uint constant X = 32**22 + 8;

contract C {
    string constant TEXT = "abc";
    bytes32 constant MY_HASH = keccak256("abc");
    uint immutable decimals;
    uint immutable maxBalance;
    address immutable owner = msg.sender;

    constructor(uint _decimals, address _reference) {
        decimals = _decimals;
        // Assignments to immutables can even access the environment.
        maxBalance = _reference.balance;
    }

    function isBalanceTooHigh(address _other) public view returns (bool) {
        return _other.balance > maxBalance;
    }
}
```

Constant

For `constant` variables, the value has to be a constant at compile time and it has to be assigned where the variable is declared. Any expression that accesses storage, blockchain data (e.g. `block.timestamp`, `address(this).balance` or `block.number`) or execution data (`msg.value` or `gasleft()`) or makes calls to external contracts is disallowed. Expressions that might have a side-effect on memory allocation are allowed, but those that might have a side-effect on other memory objects are not. The built-in functions `keccak256`, `sha256`, `ripemd160`, `ecrecover`, `addmod` and `mulmod` are allowed (even though, with the exception of `keccak256`, they do call external contracts).

The reason behind allowing side-effects on the memory allocator is that it should be possible to construct complex objects like e.g. lookup-tables. This feature is not yet fully usable.

Immutable

Variables declared as `immutable` are a bit less restricted than those declared as `constant`: Immutable variables can be assigned an arbitrary value in the constructor of the contract or at the point of their declaration. They cannot be read during construction time and can only be assigned once.

The contract creation code generated by the compiler will modify the contract's runtime code before it is returned by replacing all references to immutables by the values assigned to the them. This is important if you are comparing the runtime code generated by the compiler with the one actually stored in the blockchain.

3.9.5 Functions

Functions can be defined inside and outside of contracts.

Functions outside of a contract, also called "free functions", always have implicit `internal` *visibility*. Their code is included in all contracts that call them, similar to internal library functions.

```solidity
// SPDX-License-Identifier: GPL-3.0
pragma solidity >0.7.0 <0.8.0;

function sum(uint[] memory _arr) pure returns (uint s) {
    for (uint i = 0; i < _arr.length; i++)
        s += _arr[i];
}

contract ArrayExample {
    bool found;
    function f(uint[] memory _arr) public {
        // This calls the free function internally.
        // The compiler will add its code to the contract.
        uint s = sum(_arr);
        require(s >= 10);
        found = true;
    }
}
```

Note: Functions defined outside a contract are still always executed in the context of a contract. They still have access to the variable `this`, can call other contracts, send them Ether and destroy the contract that called them, among other things. The main difference to functions defined inside a contract is that free functions do not have direct access to storage variables and functions not in their scope.

Function Parameters and Return Variables

Functions take typed parameters as input and may, unlike in many other languages, also return an arbitrary number of values as output.

Function Parameters

Function parameters are declared the same way as variables, and the name of unused parameters can be omitted.

For example, if you want your contract to accept one kind of external call with two integers, you would use something like the following:

```
// SPDX-License-Identifier: GPL-3.0
pragma solidity >=0.4.16 <0.8.0;

contract Simple {
    uint sum;
    function taker(uint _a, uint _b) public {
        sum = _a + _b;
    }
}
```

Function parameters can be used as any other local variable and they can also be assigned to.

Note: An *external function* cannot accept a multi-dimensional array as an input parameter. This functionality is possible if you enable the ABI coder v2 by adding `pragma abicoder v2;` to your source file.

An *internal function* can accept a multi-dimensional array without enabling the feature.

Return Variables

Function return variables are declared with the same syntax after the `returns` keyword.

For example, suppose you want to return two results: the sum and the product of two integers passed as function parameters, then you use something like:

```
// SPDX-License-Identifier: GPL-3.0
pragma solidity >=0.4.16 <0.8.0;

contract Simple {
    function arithmetic(uint _a, uint _b)
        public
        pure
        returns (uint o_sum, uint o_product)
    {
        o_sum = _a + _b;
        o_product = _a * _b;
    }
}
```

The names of return variables can be omitted. Return variables can be used as any other local variable and they are initialized with their *default value* and have that value until they are (re-)assigned.

You can either explicitly assign to return variables and then leave the function as above, or you can provide return values (either a single or *multiple ones*) directly with the `return` statement:

```
// SPDX-License-Identifier: GPL-3.0
pragma solidity >=0.4.16 <0.8.0;

contract Simple {
    function arithmetic(uint _a, uint _b)
        public
        pure
        returns (uint o_sum, uint o_product)
    {
        return (_a + _b, _a * _b);
    }
}
```

If you use an early `return` to leave a function that has return variables, you must provide return values together with the return statement.

Note: You cannot return some types from non-internal functions, notably multi-dimensional dynamic arrays and structs. If you enable the ABI coder v2 by adding `pragma abicoder v2;` to your source file then more types are available, but `mapping` types are still limited to inside a single contract and you cannot transfer them.

Returning Multiple Values

When a function has multiple return types, the statement `return (v0, v1, ..., vn)` can be used to return multiple values. The number of components must be the same as the number of return variables and their types have to match, potentially after an *implicit conversion*.

View Functions

Functions can be declared `view` in which case they promise not to modify the state.

Note: If the compiler's EVM target is Byzantium or newer (default) the opcode `STATICCALL` is used when `view` functions are called, which enforces the state to stay unmodified as part of the EVM execution. For library `view` functions `DELEGATECALL` is used, because there is no combined `DELEGATECALL` and `STATICCALL`. This means library `view` functions do not have run-time checks that prevent state modifications. This should not impact security negatively because library code is usually known at compile-time and the static checker performs compile-time checks.

The following statements are considered modifying the state:

1. Writing to state variables.

2. *Emitting events*.

3. *Creating other contracts*.

4. Using `selfdestruct`.

5. Sending Ether via calls.

6. Calling any function not marked `view` or `pure`.

7. Using low-level calls.

8. Using inline assembly that contains certain opcodes.

```
// SPDX-License-Identifier: GPL-3.0
pragma solidity >=0.5.0 <0.8.0;

contract C {
    function f(uint a, uint b) public view returns (uint) {
        return a * (b + 42) + block.timestamp;
    }
}
```

Note: `constant` on functions used to be an alias to `view`, but this was dropped in version 0.5.0.

Note: Getter methods are automatically marked `view`.

Note: Prior to version 0.5.0, the compiler did not use the `STATICCALL` opcode for `view` functions. This enabled state modifications in `view` functions through the use of invalid explicit type conversions. By using `STATICCALL` for `view` functions, modifications to the state are prevented on the level of the EVM.

Pure Functions

Functions can be declared `pure` in which case they promise not to read from or modify the state.

Note: If the compiler's EVM target is Byzantium or newer (default) the opcode `STATICCALL` is used, which does not guarantee that the state is not read, but at least that it is not modified.

In addition to the list of state modifying statements explained above, the following are considered reading from the state:

1. Reading from state variables.

2. Accessing `address(this).balance` or `<address>.balance`.

3. Accessing any of the members of `block`, `tx`, `msg` (with the exception of `msg.sig` and `msg.data`).

4. Calling any function not marked `pure`.

5. Using inline assembly that contains certain opcodes.

```
// SPDX-License-Identifier: GPL-3.0
pragma solidity >=0.5.0 <0.8.0;

contract C {
    function f(uint a, uint b) public pure returns (uint) {
        return a * (b + 42);
    }
}
```

Pure functions are able to use the `revert()` and `require()` functions to revert potential state changes when an *error occurs*.

Reverting a state change is not considered a "state modification", as only changes to the state made previously in code that did not have the `view` or `pure` restriction are reverted and that code has the option to catch the `revert` and not

pass it on.

This behaviour is also in line with the `STATICCALL` opcode.

> **Warning:** It is not possible to prevent functions from reading the state at the level of the EVM, it is only possible to prevent them from writing to the state (i.e. only `view` can be enforced at the EVM level, `pure` can not).

> **Note:** Prior to version 0.5.0, the compiler did not use the `STATICCALL` opcode for `pure` functions. This enabled state modifications in `pure` functions through the use of invalid explicit type conversions. By using `STATICCALL` for `pure` functions, modifications to the state are prevented on the level of the EVM.

> **Note:** Prior to version 0.4.17 the compiler did not enforce that `pure` is not reading the state. It is a compile-time type check, which can be circumvented doing invalid explicit conversions between contract types, because the compiler can verify that the type of the contract does not do state-changing operations, but it cannot check that the contract that will be called at runtime is actually of that type.

Receive Ether Function

A contract can have at most one `receive` function, declared using `receive() external payable { ... }` (without the `function` keyword). This function cannot have arguments, cannot return anything and must have `external` visibility and `payable` state mutability. It can be virtual, can override and can have modifiers.

The receive function is executed on a call to the contract with empty calldata. This is the function that is executed on plain Ether transfers (e.g. via `.send()` or `.transfer()`). If no such function exists, but a payable *fallback function* exists, the fallback function will be called on a plain Ether transfer. If neither a receive Ether nor a payable fallback function is present, the contract cannot receive Ether through regular transactions and throws an exception.

In the worst case, the `receive` function can only rely on 2300 gas being available (for example when `send` or `transfer` is used), leaving little room to perform other operations except basic logging. The following operations will consume more gas than the 2300 gas stipend:

- Writing to storage
- Creating a contract
- Calling an external function which consumes a large amount of gas
- Sending Ether

> **Warning:** Contracts that receive Ether directly (without a function call, i.e. using `send` or `transfer`) but do not define a receive Ether function or a payable fallback function throw an exception, sending back the Ether (this was different before Solidity v0.4.0). So if you want your contract to receive Ether, you have to implement a receive Ether function (using payable fallback functions for receiving Ether is not recommended, since it would not fail on interface confusions).

> **Warning:** A contract without a receive Ether function can receive Ether as a recipient of a *coinbase transaction* (aka *miner block reward*) or as a destination of a `selfdestruct`.

A contract cannot react to such Ether transfers and thus also cannot reject them. This is a design choice of the EVM and Solidity cannot work around it.

It also means that `address(this).balance` can be higher than the sum of some manual accounting implemented in a contract (i.e. having a counter updated in the receive Ether function).

Below you can see an example of a Sink contract that uses function `receive`.

```solidity
// SPDX-License-Identifier: GPL-3.0
pragma solidity >=0.6.0 <0.8.0;

// This contract keeps all Ether sent to it with no way
// to get it back.
contract Sink {
    event Received(address, uint);
    receive() external payable {
        emit Received(msg.sender, msg.value);
    }
}
```

Fallback Function

A contract can have at most one `fallback` function, declared using either `fallback () external [payable]` or `fallback (bytes calldata _input) external [payable] returns (bytes memory _output)` (both without the `function` keyword). This function must have `external` visibility. A fallback function can be virtual, can override and can have modifiers.

The fallback function is executed on a call to the contract if none of the other functions match the given function signature, or if no data was supplied at all and there is no *receive Ether function*. The fallback function always receives data, but in order to also receive Ether it must be marked `payable`.

If the version with parameters is used, `_input` will contain the full data sent to the contract (equal to `msg.data`) and can return data in `_output`. The returned data will not be ABI-encoded. Instead it will be returned without modifications (not even padding).

In the worst case, if a payable fallback function is also used in place of a receive function, it can only rely on 2300 gas being available (see *receive Ether function* for a brief description of the implications of this).

Like any function, the fallback function can execute complex operations as long as there is enough gas passed on to it.

> **Warning:** A `payable` fallback function is also executed for plain Ether transfers, if no *receive Ether function* is present. It is recommended to always define a receive Ether function as well, if you define a payable fallback function to distinguish Ether transfers from interface confusions.

> **Note:** If you want to decode the input data, you can check the first four bytes for the function selector and then you can use `abi.decode` together with the array slice syntax to decode ABI-encoded data: `(c, d) = abi.decode(_input[4:], (uint256, uint256));` Note that this should only be used as a last resort and proper functions should be used instead.

```solidity
// SPDX-License-Identifier: GPL-3.0
pragma solidity >=0.6.2 <0.8.0;
```

```solidity
contract Test {
    // This function is called for all messages sent to
    // this contract (there is no other function).
    // Sending Ether to this contract will cause an exception,
    // because the fallback function does not have the `payable`
    // modifier.
    fallback() external { x = 1; }
    uint x;
}

contract TestPayable {
    // This function is called for all messages sent to
    // this contract, except plain Ether transfers
    // (there is no other function except the receive function).
    // Any call with non-empty calldata to this contract will execute
    // the fallback function (even if Ether is sent along with the call).
    fallback() external payable { x = 1; y = msg.value; }

    // This function is called for plain Ether transfers, i.e.
    // for every call with empty calldata.
    receive() external payable { x = 2; y = msg.value; }
    uint x;
    uint y;
}

contract Caller {
    function callTest(Test test) public returns (bool) {
        (bool success,) = address(test).call(abi.encodeWithSignature(
→"nonExistingFunction()"));
        require(success);
        // results in test.x becoming == 1.

        // address(test) will not allow to call ``send`` directly, since ``test`` has␣
→no payable
        // fallback function.
        // It has to be converted to the ``address payable`` type to even allow␣
→calling ``send`` on it.
        address payable testPayable = payable(address(test));

        // If someone sends Ether to that contract,
        // the transfer will fail, i.e. this returns false here.
        return testPayable.send(2 ether);
    }

    function callTestPayable(TestPayable test) public returns (bool) {
        (bool success,) = address(test).call(abi.encodeWithSignature(
→"nonExistingFunction()"));
        require(success);
        // results in test.x becoming == 1 and test.y becoming 0.
        (success,) = address(test).call{value: 1}(abi.encodeWithSignature(
→"nonExistingFunction()"));
        require(success);
        // results in test.x becoming == 1 and test.y becoming 1.

        // If someone sends Ether to that contract, the receive function in␣
→TestPayable will be called.
```

```
        require(address(test).send(2 ether));
        // results in test.x becoming == 2 and test.y becoming 2 ether.

        return true;
    }
}
```

Function Overloading

A contract can have multiple functions of the same name but with different parameter types. This process is called "overloading" and also applies to inherited functions. The following example shows overloading of the function f in the scope of contract A.

```
// SPDX-License-Identifier: GPL-3.0
pragma solidity >=0.4.16 <0.8.0;

contract A {
    function f(uint _in) public pure returns (uint out) {
        out = _in;
    }

    function f(uint _in, bool _really) public pure returns (uint out) {
        if (_really)
            out = _in;
    }
}
```

Overloaded functions are also present in the external interface. It is an error if two externally visible functions differ by their Solidity types but not by their external types.

```
// SPDX-License-Identifier: GPL-3.0
pragma solidity >=0.4.16 <0.8.0;

// This will not compile
contract A {
    function f(B _in) public pure returns (B out) {
        out = _in;
    }

    function f(address _in) public pure returns (address out) {
        out = _in;
    }
}

contract B {
}
```

Both f function overloads above end up accepting the address type for the ABI although they are considered different inside Solidity.

Overload resolution and Argument matching

Overloaded functions are selected by matching the function declarations in the current scope to the arguments supplied in the function call. Functions are selected as overload candidates if all arguments can be implicitly converted to the

---

expected types. If there is not exactly one candidate, resolution fails.

Note: Return parameters are not taken into account for overload resolution.

```solidity
// SPDX-License-Identifier: GPL-3.0
pragma solidity >=0.4.16 <0.8.0;

contract A {
    function f(uint8 _in) public pure returns (uint8 out) {
        out = _in;
    }

    function f(uint256 _in) public pure returns (uint256 out) {
        out = _in;
    }
}
```

Calling `f(50)` would create a type error since `50` can be implicitly converted both to `uint8` and `uint256` types. On another hand `f(256)` would resolve to `f(uint256)` overload as `256` cannot be implicitly converted to `uint8`.

3.9.6 Events

Solidity events give an abstraction on top of the EVM's logging functionality. Applications can subscribe and listen to these events through the RPC interface of an Ethereum client.

Events are inheritable members of contracts. When you call them, they cause the arguments to be stored in the transaction's log - a special data structure in the blockchain. These logs are associated with the address of the contract, are incorporated into the blockchain, and stay there as long as a block is accessible (forever as of now, but this might change with Serenity). The Log and its event data is not accessible from within contracts (not even from the contract that created them).

It is possible to request a Merkle proof for logs, so if an external entity supplies a contract with such a proof, it can check that the log actually exists inside the blockchain. You have to supply block headers because the contract can only see the last 256 block hashes.

You can add the attribute `indexed` to up to three parameters which adds them to a special data structure known as *"topics"* instead of the data part of the log. If you use arrays (including `string` and `bytes`) as indexed arguments, its Keccak-256 hash is stored as a topic instead, this is because a topic can only hold a single word (32 bytes).

All parameters without the `indexed` attribute are *ABI-encoded* into the data part of the log.

Topics allow you to search for events, for example when filtering a sequence of blocks for certain events. You can also filter events by the address of the contract that emitted the event.

For example, the code below uses the web3.js `subscribe("logs")` method to filter logs that match a topic with a certain address value:

```javascript
var options = {
    fromBlock: 0,
    address: web3.eth.defaultAccount,
    topics: ["0x0000000000000000000000000000000000000000000000000000000000000000",
    null, null]
};
web3.eth.subscribe('logs', options, function (error, result) {
    if (!error)
        console.log(result);
```

(continues on next page)

```
})
    .on("data", function (log) {
        console.log(log);
    })
    .on("changed", function (log) {
});
```

The hash of the signature of the event is one of the topics, except if you declared the event with the `anonymous` specifier. This means that it is not possible to filter for specific anonymous events by name, you can only filter by the contract address. The advantage of anonymous events is that they are cheaper to deploy and call.

```solidity
// SPDX-License-Identifier: GPL-3.0
pragma solidity >=0.4.21 <0.8.0;

contract ClientReceipt {
    event Deposit(
        address indexed _from,
        bytes32 indexed _id,
        uint _value
    );

    function deposit(bytes32 _id) public payable {
        // Events are emitted using `emit`, followed by
        // the name of the event and the arguments
        // (if any) in parentheses. Any such invocation
        // (even deeply nested) can be detected from
        // the JavaScript API by filtering for `Deposit`.
        emit Deposit(msg.sender, _id, msg.value);
    }
}
```

The use in the JavaScript API is as follows:

```javascript
var abi = /* abi as generated by the compiler */;
var ClientReceipt = web3.eth.contract(abi);
var clientReceipt = ClientReceipt.at("0x1234...ab67" /* address */);

var event = clientReceipt.Deposit();

// watch for changes
event.watch(function(error, result){
    // result contains non-indexed arguments and topics
    // given to the `Deposit` call.
    if (!error)
        console.log(result);
});


// Or pass a callback to start watching immediately
var event = clientReceipt.Deposit(function(error, result) {
    if (!error)
        console.log(result);
});
```

The output of the above looks like the following (trimmed):

```
{
    "returnValues": {
        "_from": "0x1111...FFFFCCCC",
        "_id": "0x50...sd5adb20",
        "_value": "0x420042"
    },
    "raw": {
        "data": "0x7f...91385",
        "topics": ["0xfd4...b4ead7", "0x7f...1a91385"]
    }
}
```

Low-Level Interface to Logs

It is also possible to access the low-level interface to the logging mechanism via the functions `log0`, `log1`, `log2`, `log3` and `log4`. Each function `logi` takes `i + 1` parameter of type `bytes32`, where the first argument will be used for the data part of the log and the others as topics. The event call above can be performed in the same way as

```solidity
// SPDX-License-Identifier: GPL-3.0
pragma solidity >=0.4.10 <0.8.0;

contract C {
    function f() public payable {
        uint256 _id = 0x420042;
        log3(
            bytes32(msg.value),
↪bytes32(0x50cb9fe53daa9737b786ab3646f04d0150dc50ef4e75f59509d83667ad5adb20),
            bytes32(uint256(msg.sender)),
            bytes32(_id)
        );
    }
}
```

where the long hexadecimal number is equal to `keccak256("Deposit(address,bytes32,uint256)")`, the signature of the event.

Additional Resources for Understanding Events

- Javascript documentation
- Example usage of events
- How to access them in js

3.9.7 Inheritance

Solidity supports multiple inheritance including polymorphism.

Polymorphism means that a function call (internal and external) always executes the function of the same name (and parameter types) in the most derived contract in the inheritance hierarchy. This has to be explicitly enabled on each function in the hierarchy using the `virtual` and `override` keywords. See *Function Overriding* for more details.

It is possible to call functions further up in the inheritance hierarchy internally by explicitly specifying the contract using `ContractName.functionName()` or using `super.functionName()` if you want to call the function one level higher up in the flattened inheritance hierarchy (see below).

When a contract inherits from other contracts, only a single contract is created on the blockchain, and the code from all the base contracts is compiled into the created contract. This means that all internal calls to functions of base contracts also just use internal function calls (`super.f(..)` will use JUMP and not a message call).

State variable shadowing is considered as an error. A derived contract can only declare a state variable `x`, if there is no visible state variable with the same name in any of its bases.

The general inheritance system is very similar to Python's, especially concerning multiple inheritance, but there are also some *differences*.

Details are given in the following example.

```solidity
// SPDX-License-Identifier: GPL-3.0
pragma solidity ^0.7.0;


contract Owned {
    constructor() { owner = msg.sender; }
    address payable owner;
}


// Use `is` to derive from another contract. Derived
// contracts can access all non-private members including
// internal functions and state variables. These cannot be
// accessed externally via `this`, though.
contract Destructible is Owned {
    // The keyword `virtual` means that the function can change
    // its behaviour in derived classes ("overriding").
    function destroy() virtual public {
        if (msg.sender == owner) selfdestruct(owner);
    }
}


// These abstract contracts are only provided to make the
// interface known to the compiler. Note the function
// without body. If a contract does not implement all
// functions it can only be used as an interface.
abstract contract Config {
    function lookup(uint id) public virtual returns (address adr);
}


abstract contract NameReg {
    function register(bytes32 name) public virtual;
    function unregister() public virtual;
}


// Multiple inheritance is possible. Note that `owned` is
// also a base class of `Destructible`, yet there is only a single
// instance of `owned` (as for virtual inheritance in C++).
contract Named is Owned, Destructible {
    constructor(bytes32 name) {
```

(continues on next page)

```
        Config config = Config(0xD5f9D8D94886E70b06E474c3fB14Fd43E2f23970);
        NameReg(config.lookup(1)).register(name);
    }


    // Functions can be overridden by another function with the same name and
    // the same number/types of inputs.  If the overriding function has different
    // types of output parameters, that causes an error.
    // Both local and message-based function calls take these overrides
    // into account.
    // If you want the function to override, you need to use the
    // `override` keyword. You need to specify the `virtual` keyword again
    // if you want this function to be overridden again.
    function destroy() public virtual override {
        if (msg.sender == owner) {
            Config config = Config(0xD5f9D8D94886E70b06E474c3fB14Fd43E2f23970);
            NameReg(config.lookup(1)).unregister();
            // It is still possible to call a specific
            // overridden function.
            Destructible.destroy();
        }
    }
}


// If a constructor takes an argument, it needs to be
// provided in the header or modifier-invocation-style at
// the constructor of the derived contract (see below).
contract PriceFeed is Owned, Destructible, Named("GoldFeed") {
    function updateInfo(uint newInfo) public {
        if (msg.sender == owner) info = newInfo;
    }

    // Here, we only specify `override` and not `virtual`.
    // This means that contracts deriving from `PriceFeed`
    // cannot change the behaviour of `destroy` anymore.
    function destroy() public override(Destructible, Named) { Named.destroy(); }
    function get() public view returns(uint r) { return info; }

    uint info;
}
```

Note that above, we call `Destructible.destroy()` to "forward" the destruction request. The way this is done
is problematic, as seen in the following example:

```
// SPDX-License-Identifier: GPL-3.0
pragma solidity ^0.7.0;

contract owned {
    constructor() { owner = msg.sender; }
    address payable owner;
}

contract Destructible is owned {
    function destroy() public virtual {
        if (msg.sender == owner) selfdestruct(owner);
    }
```

```
}

contract Base1 is Destructible {
    function destroy() public virtual override { /* do cleanup 1 */ Destructible.
↪destroy(); }
}

contract Base2 is Destructible {
    function destroy() public virtual override { /* do cleanup 2 */ Destructible.
↪destroy(); }
}

contract Final is Base1, Base2 {
    function destroy() public override(Base1, Base2) { Base2.destroy(); }
}
```

A call to `Final.destroy()` will call `Base2.destroy` because we specify it explicitly in the final override, but this function will bypass `Base1.destroy`. The way around this is to use `super`:

```
// SPDX-License-Identifier: GPL-3.0
pragma solidity ^0.7.0;

contract owned {
    constructor() { owner = msg.sender; }
    address payable owner;
}

contract Destructible is owned {
    function destroy() virtual public {
        if (msg.sender == owner) selfdestruct(owner);
    }
}

contract Base1 is Destructible {
    function destroy() public virtual override { /* do cleanup 1 */ super.destroy(); }
}


contract Base2 is Destructible {
    function destroy() public virtual override { /* do cleanup 2 */ super.destroy(); }
}

contract Final is Base1, Base2 {
    function destroy() public override(Base1, Base2) { super.destroy(); }
}
```

If `Base2` calls a function of `super`, it does not simply call this function on one of its base contracts. Rather, it calls this function on the next base contract in the final inheritance graph, so it will call `Base1.destroy()` (note that the final inheritance sequence is – starting with the most derived contract: Final, Base2, Base1, Destructible, owned). The actual function that is called when using super is not known in the context of the class where it is used, although its type is known. This is similar for ordinary virtual method lookup.

Function Overriding

Base functions can be overridden by inheriting contracts to change their behavior if they are marked as `virtual`. The overriding function must then use the `override` keyword in the function header. The overriding function may

only change the visibility of the overridden function from `external` to `public`. The mutability may be changed to a more strict one following the order: `nonpayable` can be overridden by `view` and `pure`. `view` can be overridden by `pure`. `payable` is an exception and cannot be changed to any other mutability.

The following example demonstrates changing mutability and visibility:

```solidity
// SPDX-License-Identifier: GPL-3.0
pragma solidity ^0.7.0;

contract Base
{
    function foo() virtual external view {}
}

contract Middle is Base {}

contract Inherited is Middle
{
    function foo() override public pure {}
}
```

For multiple inheritance, the most derived base contracts that define the same function must be specified explicitly after the `override` keyword. In other words, you have to specify all base contracts that define the same function and have not yet been overridden by another base contract (on some path through the inheritance graph). Additionally, if a contract inherits the same function from multiple (unrelated) bases, it has to explicitly override it:

```solidity
// SPDX-License-Identifier: GPL-3.0
pragma solidity >=0.6.0 <0.8.0;

contract Base1
{
    function foo() virtual public {}
}

contract Base2
{
    function foo() virtual public {}
}

contract Inherited is Base1, Base2
{
    // Derives from multiple bases defining foo(), so we must explicitly
    // override it
    function foo() public override(Base1, Base2) {}
}
```

An explicit override specifier is not required if the function is defined in a common base contract or if there is a unique function in a common base contract that already overrides all other functions.

```solidity
// SPDX-License-Identifier: GPL-3.0
pragma solidity >=0.6.0 <0.8.0;

contract A { function f() public pure{} }
contract B is A {}
contract C is A {}
// No explicit override required
contract D is B, C {}
```

More formally, it is not required to override a function (directly or indirectly) inherited from multiple bases if there

is a base contract that is part of all override paths for the signature, and (1) that base implements the function and no paths from the current contract to the base mentions a function with that signature or (2) that base does not implement the function and there is at most one mention of the function in all paths from the current contract to that base.

In this sense, an override path for a signature is a path through the inheritance graph that starts at the contract under consideration and ends at a contract mentioning a function with that signature that does not override.

If you do not mark a function that overrides as `virtual`, derived contracts can no longer change the behaviour of that function.

Note: Functions with the `private` visibility cannot be `virtual`.

Note: Functions without implementation have to be marked `virtual` outside of interfaces. In interfaces, all functions are automatically considered `virtual`.

Public state variables can override external functions if the parameter and return types of the function matches the getter function of the variable:

```
// SPDX-License-Identifier: GPL-3.0
pragma solidity >=0.6.0 <0.8.0;

contract A
{
    function f() external view virtual returns(uint) { return 5; }
}

contract B is A
{
    uint public override f;
}
```

Note: While public state variables can override external functions, they themselves cannot be overridden.

Modifier Overriding

Function modifiers can override each other. This works in the same way as *function overriding* (except that there is no overloading for modifiers). The `virtual` keyword must be used on the overridden modifier and the `override` keyword must be used in the overriding modifier:

```
// SPDX-License-Identifier: GPL-3.0
pragma solidity >=0.6.0 <0.8.0;

contract Base
{
    modifier foo() virtual {_;}
}

contract Inherited is Base
{
    modifier foo() override {_;}
}
```

In case of multiple inheritance, all direct base contracts must be specified explicitly:

```
// SPDX-License-Identifier: GPL-3.0
pragma solidity >=0.6.0 <0.8.0;

contract Base1
{
    modifier foo() virtual {_;}
}

contract Base2
{
    modifier foo() virtual {_;}
}

contract Inherited is Base1, Base2
{
    modifier foo() override(Base1, Base2) {_;}
}
```

Constructors

A constructor is an optional function declared with the `constructor` keyword which is executed upon contract creation, and where you can run contract initialisation code.

Before the constructor code is executed, state variables are initialised to their specified value if you initialise them inline, or zero if you do not.

After the constructor has run, the final code of the contract is deployed to the blockchain. The deployment of the code costs additional gas linear to the length of the code. This code includes all functions that are part of the public interface and all functions that are reachable from there through function calls. It does not include the constructor code or internal functions that are only called from the constructor.

If there is no constructor, the contract will assume the default constructor, which is equivalent to `constructor()` `{}`. For example:

```
// SPDX-License-Identifier: GPL-3.0
pragma solidity ^0.7.0;

abstract contract A {
    uint public a;

    constructor(uint _a) {
        a = _a;
    }
}

contract B is A(1) {
    constructor() {}
}
```

You can use internal parameters in a constructor (for example storage pointers). In this case, the contract has to be marked *abstract*, because these parameters cannot be assigned valid values from outside but only through the constructors of derived contracts.

> **Warning:** Prior to version 0.4.22, constructors were defined as functions with the same name as the contract. This syntax was deprecated and is not allowed anymore in version 0.5.0.

> **Warning:** Prior to version 0.7.0, you had to specify the visibility of constructors as either `internal` or `public`.

Arguments for Base Constructors

The constructors of all the base contracts will be called following the linearization rules explained below. If the base constructors have arguments, derived contracts need to specify all of them. This can be done in two ways:

```solidity
// SPDX-License-Identifier: GPL-3.0
pragma solidity ^0.7.0;

contract Base {
    uint x;
    constructor(uint _x) { x = _x; }
}

// Either directly specify in the inheritance list...
contract Derived1 is Base(7) {
    constructor() {}
}

// or through a "modifier" of the derived constructor.
contract Derived2 is Base {
    constructor(uint _y) Base(_y * _y) {}
}
```

One way is directly in the inheritance list (`is Base(7)`). The other is in the way a modifier is invoked as part of the derived constructor (`Base(_y * _y)`). The first way to do it is more convenient if the constructor argument is a constant and defines the behaviour of the contract or describes it. The second way has to be used if the constructor arguments of the base depend on those of the derived contract. Arguments have to be given either in the inheritance list or in modifier-style in the derived constructor. Specifying arguments in both places is an error.

If a derived contract does not specify the arguments to all of its base contracts' constructors, it will be abstract.

Multiple Inheritance and Linearization

Languages that allow multiple inheritance have to deal with several problems. One is the Diamond Problem. Solidity is similar to Python in that it uses "C3 Linearization" to force a specific order in the directed acyclic graph (DAG) of base classes. This results in the desirable property of monotonicity but disallows some inheritance graphs. Especially, the order in which the base classes are given in the `is` directive is important: You have to list the direct base contracts in the order from "most base-like" to "most derived". Note that this order is the reverse of the one used in Python.

Another simplifying way to explain this is that when a function is called that is defined multiple times in different contracts, the given bases are searched from right to left (left to right in Python) in a depth-first manner, stopping at the first match. If a base contract has already been searched, it is skipped.

In the following code, Solidity will give the error "Linearization of inheritance graph impossible".

```
// SPDX-License-Identifier: GPL-3.0
pragma solidity >=0.4.0 <0.8.0;

contract X {}
contract A is X {}
// This will not compile
contract C is A, X {}
```

The reason for this is that C requests X to override A (by specifying A, X in this order), but A itself requests to override X, which is a contradiction that cannot be resolved.

Due to the fact that you have to explicitly override a function that is inherited from multiple bases without a unique override, C3 linearization is not too important in practice.

One area where inheritance linearization is especially important and perhaps not as clear is when there are multiple constructors in the inheritance hierarchy. The constructors will always be executed in the linearized order, regardless of the order in which their arguments are provided in the inheriting contract's constructor. For example:

```
// SPDX-License-Identifier: GPL-3.0
pragma solidity ^0.7.0;

contract Base1 {
    constructor() {}
}

contract Base2 {
    constructor() {}
}

// Constructors are executed in the following order:
//  1 - Base1
//  2 - Base2
//  3 - Derived1
contract Derived1 is Base1, Base2 {
    constructor() Base1() Base2() {}
}

// Constructors are executed in the following order:
//  1 - Base2
//  2 - Base1
//  3 - Derived2
contract Derived2 is Base2, Base1 {
    constructor() Base2() Base1() {}
}

// Constructors are still executed in the following order:
//  1 - Base2
//  2 - Base1
//  3 - Derived3
contract Derived3 is Base2, Base1 {
    constructor() Base1() Base2() {}
}
```

Inheriting Different Kinds of Members of the Same Name

It is an error when any of the following pairs in a contract have the same name due to inheritance:

- a function and a modifier

- a function and an event

- an event and a modifier

As an exception, a state variable getter can override an external function.

3.9.8 Abstract Contracts

Contracts need to be marked as abstract when at least one of their functions is not implemented. Contracts may be marked as abstract even though all functions are implemented.

This can be done by using the `abstract` keyword as shown in the following example. Note that this contract needs to be defined as abstract, because the function `utterance()` was defined, but no implementation was provided (no implementation body `{ }` was given).:

```solidity
// SPDX-License-Identifier: GPL-3.0
pragma solidity >=0.6.0 <0.8.0;

abstract contract Feline {
    function utterance() public virtual returns (bytes32);
}
```

Such abstract contracts can not be instantiated directly. This is also true, if an abstract contract itself does implement all defined functions. The usage of an abstract contract as a base class is shown in the following example:

```solidity
// SPDX-License-Identifier: GPL-3.0
pragma solidity >=0.6.0 <0.8.0;

abstract contract Feline {
    function utterance() public pure virtual returns (bytes32);
}

contract Cat is Feline {
    function utterance() public pure override returns (bytes32) { return "miaow"; }
}
```

If a contract inherits from an abstract contract and does not implement all non-implemented functions by overriding, it needs to be marked as abstract as well.

Note that a function without implementation is different from a *Function Type* even though their syntax looks very similar.

Example of function without implementation (a function declaration):

```solidity
function foo(address) external returns (address);
```

Example of a declaration of a variable whose type is a function type:

```solidity
function(address) external returns (address) foo;
```

Abstract contracts decouple the definition of a contract from its implementation providing better extensibility and self-documentation and facilitating patterns like the Template method and removing code duplication. Abstract contracts are useful in the same way that defining methods in an interface is useful. It is a way for the designer of the abstract contract to say "any child of mine must implement this method".

Note: Abstract contracts cannot override an implemented virtual function with an unimplemented one.

3.9.9 Interfaces

Interfaces are similar to abstract contracts, but they cannot have any functions implemented. There are further restrictions:

- They cannot inherit from other contracts, but they can inherit from other interfaces.

- All declared functions must be external.

- They cannot declare a constructor.

- They cannot declare state variables.

Some of these restrictions might be lifted in the future.

Interfaces are basically limited to what the Contract ABI can represent, and the conversion between the ABI and an interface should be possible without any information loss.

Interfaces are denoted by their own keyword:

```solidity
// SPDX-License-Identifier: GPL-3.0
pragma solidity >=0.6.2 <0.8.0;

interface Token {
    enum TokenType { Fungible, NonFungible }
    struct Coin { string obverse; string reverse; }
    function transfer(address recipient, uint amount) external;
}
```

Contracts can inherit interfaces as they would inherit other contracts.

All functions declared in interfaces are implicitly `virtual`, which means that they can be overridden. This does not automatically mean that an overriding function can be overridden again - this is only possible if the overriding function is marked `virtual`.

Interfaces can inherit from other interfaces. This has the same rules as normal inheritance.

```solidity
// SPDX-License-Identifier: GPL-3.0
pragma solidity >=0.6.2 <0.8.0;

interface ParentA {
    function test() external returns (uint256);
}

interface ParentB {
    function test() external returns (uint256);
}

interface SubInterface is ParentA, ParentB {
    // Must redefine test in order to assert that the parent
    // meanings are compatible.
    function test() external override(ParentA, ParentB) returns (uint256);
}
```

Types defined inside interfaces and other contract-like structures can be accessed from other contracts: `Token.TokenType` or `Token.Coin`.

3.9.10 Libraries

Libraries are similar to contracts, but their purpose is that they are deployed only once at a specific address and their code is reused using the `DELEGATECALL` (`CALLCODE` until Homestead) feature of the EVM. This means that if library functions are called, their code is executed in the context of the calling contract, i.e. `this` points to the calling contract, and especially the storage from the calling contract can be accessed. As a library is an isolated piece of source code, it can only access state variables of the calling contract if they are explicitly supplied (it would have no way to name them, otherwise). Library functions can only be called directly (i.e. without the use of `DELEGATECALL`) if they do not modify the state (i.e. if they are `view` or `pure` functions), because libraries are assumed to be stateless. In particular, it is not possible to destroy a library.

Note: Until version 0.4.20, it was possible to destroy libraries by circumventing Solidity's type system. Starting from that version, libraries contain a *mechanism* that disallows state-modifying functions to be called directly (i.e. without `DELEGATECALL`).

Libraries can be seen as implicit base contracts of the contracts that use them. They will not be explicitly visible in the inheritance hierarchy, but calls to library functions look just like calls to functions of explicit base contracts (using qualified access like `L.f()`). Of course, calls to internal functions use the internal calling convention, which means that all internal types can be passed and types *stored in memory* will be passed by reference and not copied. To realize this in the EVM, code of internal library functions and all functions called from therein will at compile time be included in the calling contract, and a regular `JUMP` call will be used instead of a `DELEGATECALL`.

The following example illustrates how to use libraries (but using a manual method, be sure to check out *using for* for a more advanced example to implement a set).

```solidity
// SPDX-License-Identifier: GPL-3.0
pragma solidity >=0.6.0 <0.8.0;


// We define a new struct datatype that will be used to
// hold its data in the calling contract.
struct Data {
    mapping(uint => bool) flags;
}

library Set {
    // Note that the first parameter is of type "storage
    // reference" and thus only its storage address and not
    // its contents is passed as part of the call.  This is a
    // special feature of library functions.  It is idiomatic
    // to call the first parameter `self`, if the function can
    // be seen as a method of that object.
    function insert(Data storage self, uint value)
        public
        returns (bool)
    {
        if (self.flags[value])
            return false; // already there
        self.flags[value] = true;
        return true;
    }

    function remove(Data storage self, uint value)
        public
        returns (bool)
```

```solidity
    {
        if (!self.flags[value])
            return false; // not there
        self.flags[value] = false;
        return true;
    }

    function contains(Data storage self, uint value)
        public
        view
        returns (bool)
    {
        return self.flags[value];
    }
}


contract C {
    Data knownValues;

    function register(uint value) public {
        // The library functions can be called without a
        // specific instance of the library, since the
        // "instance" will be the current contract.
        require(Set.insert(knownValues, value));
    }
    // In this contract, we can also directly access knownValues.flags, if we want.
}
```

Of course, you do not have to follow this way to use libraries: they can also be used without defining struct data types. Functions also work without any storage reference parameters, and they can have multiple storage reference parameters and in any position.

The calls to Set.contains, Set.insert and Set.remove are all compiled as calls (DELEGATECALL) to an external contract/library. If you use libraries, be aware that an actual external function call is performed. msg.sender, msg.value and this will retain their values in this call, though (prior to Homestead, because of the use of CALLCODE, msg.sender and msg.value changed, though).

The following example shows how to use *types stored in memory* and internal functions in libraries in order to implement custom types without the overhead of external function calls:

```solidity
// SPDX-License-Identifier: GPL-3.0
pragma solidity >=0.6.0 <0.8.0;

struct bigint {
    uint[] limbs;
}

library BigInt {
    function fromUint(uint x) internal pure returns (bigint memory r) {
        r.limbs = new uint[](1);
        r.limbs[0] = x;
    }

    function add(bigint memory _a, bigint memory _b) internal pure returns (bigint
→memory r) {
        r.limbs = new uint[](max(_a.limbs.length, _b.limbs.length));
```

```solidity
        uint carry = 0;
        for (uint i = 0; i < r.limbs.length; ++i) {
            uint a = limb(_a, i);
            uint b = limb(_b, i);
            r.limbs[i] = a + b + carry;
            if (a + b < a || (a + b == uint(-1) && carry > 0))
                carry = 1;
            else
                carry = 0;
        }
        if (carry > 0) {
            // too bad, we have to add a limb
            uint[] memory newLimbs = new uint[](r.limbs.length + 1);
            uint i;
            for (i = 0; i < r.limbs.length; ++i)
                newLimbs[i] = r.limbs[i];
            newLimbs[i] = carry;
            r.limbs = newLimbs;
        }
    }

    function limb(bigint memory _a, uint _limb) internal pure returns (uint) {
        return _limb < _a.limbs.length ? _a.limbs[_limb] : 0;
    }

    function max(uint a, uint b) private pure returns (uint) {
        return a > b ? a : b;
    }
}

contract C {
    using BigInt for bigint;

    function f() public pure {
        bigint memory x = BigInt.fromUint(7);
        bigint memory y = BigInt.fromUint(uint(-1));
        bigint memory z = x.add(y);
        assert(z.limb(1) > 0);
    }
}
```

It is possible to obtain the address of a library by converting the library type to the `address` type, i.e. using `address(LibraryName)`.

As the compiler does not know the address where the library will be deployed, the compiled hex code will contain placeholders of the form `__$30bbc0abd4d6364515865950d3e0d10953$__`. The placeholder is a 34 character prefix of the hex encoding of the keccak256 hash of the fully qualified library name, which would be for example `libraries/bigint.sol:BigInt` if the library was stored in a file called `bigint.sol` in a `libraries/` directory. Such bytecode is incomplete and should not be deployed. Placeholders need to be replaced with actual addresses. You can do that by either passing them to the compiler when the library is being compiled or by using the linker to update an already compiled binary. See *Library linking* for information on how to use the commandline compiler for linking.

In comparison to contracts, libraries are restricted in the following ways:

- they cannot have state variables

- they cannot inherit nor be inherited

- they cannot receive Ether

- they cannot be destroyed

(These might be lifted at a later point.)

Function Signatures and Selectors in Libraries

While external calls to public or external library functions are possible, the calling convention for such calls is considered to be internal to Solidity and not the same as specified for the regular *contract ABI*. External library functions support more argument types than external contract functions, for example recursive structs and storage pointers. For that reason, the function signatures used to compute the 4-byte selector are computed following an internal naming schema and arguments of types not supported in the contract ABI use an internal encoding.

The following identifiers are used for the types in the signatures:

- Value types, non-storage `string` and non-storage `bytes` use the same identifiers as in the contract ABI.

- Non-storage array types follow the same convention as in the contract ABI, i.e. `<type>[]` for dynamic arrays and `<type>[M]` for fixed-size arrays of `M` elements.

- Non-storage structs are referred to by their fully qualified name, i.e. `C.S` for `contract C { struct S { ... } }`.

- Storage pointer mappings use `mapping(<keyType> => <valueType>) storage` where `<keyType>` and `<valueType>` are the identifiers for the key and value types of the mapping, respectively.

- Other storage pointer types use the type identifier of their corresponding non-storage type, but append a single space followed by `storage` to it.

The argument encoding is the same as for the regular contract ABI, except for storage pointers, which are encoded as a `uint256` value referring to the storage slot to which they point.

Similarly to the contract ABI, the selector consists of the first four bytes of the Keccak256-hash of the signature. Its value can be obtained from Solidity using the `.selector` member as follows:

```solidity
// SPDX-License-Identifier: GPL-3.0
pragma solidity >=0.5.14 <0.8.0;

library L {
    function f(uint256) external {}
}

contract C {
    function g() public pure returns (bytes4) {
        return L.f.selector;
    }
}
```

Call Protection For Libraries

As mentioned in the introduction, if a library's code is executed using a `CALL` instead of a `DELEGATECALL` or `CALLCODE`, it will revert unless a `view` or `pure` function is called.

The EVM does not provide a direct way for a contract to detect whether it was called using `CALL` or not, but a contract can use the `ADDRESS` opcode to find out "where" it is currently running. The generated code compares this address to the address used at construction time to determine the mode of calling.

More specifically, the runtime code of a library always starts with a push instruction, which is a zero of 20 bytes at compilation time. When the deploy code runs, this constant is replaced in memory by the current address and this modified code is stored in the contract. At runtime, this causes the deploy time address to be the first constant to be pushed onto the stack and the dispatcher code compares the current address against this constant for any non-view and non-pure function.

This means that the actual code stored on chain for a library is different from the code reported by the compiler as `deployedBytecode`.

3.9.11 Using For

The directive `using A for B;` can be used to attach library functions (from the library A) to any type (B) in the context of a contract. These functions will receive the object they are called on as their first parameter (like the `self` variable in Python).

The effect of `using A for *;` is that the functions from the library A are attached to *any* type.

In both situations, *all* functions in the library are attached, even those where the type of the first parameter does not match the type of the object. The type is checked at the point the function is called and function overload resolution is performed.

The `using A for B;` directive is active only within the current contract, including within all of its functions, and has no effect outside of the contract in which it is used. The directive may only be used inside a contract, not inside any of its functions.

Let us rewrite the set example from the *Libraries* in this way:

```solidity
// SPDX-License-Identifier: GPL-3.0
pragma solidity >=0.6.0 <0.8.0;


// This is the same code as before, just without comments
struct Data { mapping(uint => bool) flags; }

library Set {
    function insert(Data storage self, uint value)
        public
        returns (bool)
    {
        if (self.flags[value])
            return false; // already there
        self.flags[value] = true;
        return true;
    }

    function remove(Data storage self, uint value)
        public
        returns (bool)
    {
        if (!self.flags[value])
            return false; // not there
        self.flags[value] = false;
        return true;
    }

    function contains(Data storage self, uint value)
        public
        view
```

```solidity
        returns (bool)
    {
        return self.flags[value];
    }
}


contract C {
    using Set for Data; // this is the crucial change
    Data knownValues;

    function register(uint value) public {
        // Here, all variables of type Data have
        // corresponding member functions.
        // The following function call is identical to
        // `Set.insert(knownValues, value)`
        require(knownValues.insert(value));
    }
}
```

It is also possible to extend elementary types in that way:

```solidity
// SPDX-License-Identifier: GPL-3.0
pragma solidity >=0.4.16 <0.8.0;

library Search {
    function indexOf(uint[] storage self, uint value)
        public
        view
        returns (uint)
    {
        for (uint i = 0; i < self.length; i++)
            if (self[i] == value) return i;
        return uint(-1);
    }
}

contract C {
    using Search for uint[];
    uint[] data;

    function append(uint value) public {
        data.push(value);
    }

    function replace(uint _old, uint _new) public {
        // This performs the library function call
        uint index = data.indexOf(_old);
        if (index == uint(-1))
            data.push(_new);
        else
            data[index] = _new;
    }
}
```

Note that all external library calls are actual EVM function calls. This means that if you pass memory or value types, a copy will be performed, even of the `self` variable. The only situation where no copy will be performed is when

storage reference variables are used or when internal library functions are called.

3.10 Inline Assembly

You can interleave Solidity statements with inline assembly in a language close to the one of the Ethereum virtual machine. This gives you more fine-grained control, which is especially useful when you are enhancing the language by writing libraries.

The language used for inline assembly in Solidity is called *Yul* and it is documented in its own section. This section will only cover how the inline assembly code can interface with the surrounding Solidity code.

> **Warning:** Inline assembly is a way to access the Ethereum Virtual Machine at a low level. This bypasses several important safety features and checks of Solidity. You should only use it for tasks that need it, and only if you are confident with using it.

An inline assembly block is marked by `assembly { ... }`, where the code inside the curly braces is code in the *Yul* language.

The inline assembly code can access local Solidity variables as explained below.

Different inline assembly blocks share no namespace, i.e. it is not possible to call a Yul function or access a Yul variable defined in a different inline assembly block.

3.10.1 Example

The following example provides library code to access the code of another contract and load it into a `bytes` variable. This is not possible with "plain Solidity" and the idea is that reusable assembly libraries can enhance the Solidity language without a compiler change.

```solidity
// SPDX-License-Identifier: GPL-3.0
pragma solidity >=0.4.16 <0.8.0;

library GetCode {
    function at(address _addr) public view returns (bytes memory o_code) {
        assembly {
            // retrieve the size of the code, this needs assembly
            let size := extcodesize(_addr)
            // allocate output byte array - this could also be done without assembly
            // by using o_code = new bytes(size)
            o_code := mload(0x40)
            // new "memory end" including padding
            mstore(0x40, add(o_code, and(add(add(size, 0x20), 0x1f), not(0x1f))))
            // store length in memory
            mstore(o_code, size)
            // actually retrieve the code, this needs assembly
            extcodecopy(_addr, add(o_code, 0x20), 0, size)
        }
    }
}
```

Inline assembly is also beneficial in cases where the optimizer fails to produce efficient code, for example:

```
// SPDX-License-Identifier: GPL-3.0
pragma solidity >=0.4.16 <0.8.0;


library VectorSum {
    // This function is less efficient because the optimizer currently fails to
    // remove the bounds checks in array access.
    function sumSolidity(uint[] memory _data) public pure returns (uint sum) {
        for (uint i = 0; i < _data.length; ++i)
            sum += _data[i];
    }

    // We know that we only access the array in bounds, so we can avoid the check.
    // 0x20 needs to be added to an array because the first slot contains the
    // array length.
    function sumAsm(uint[] memory _data) public pure returns (uint sum) {
        for (uint i = 0; i < _data.length; ++i) {
            assembly {
                sum := add(sum, mload(add(add(_data, 0x20), mul(i, 0x20))))
            }
        }
    }

    // Same as above, but accomplish the entire code within inline assembly.
    function sumPureAsm(uint[] memory _data) public pure returns (uint sum) {
        assembly {
            // Load the length (first 32 bytes)
            let len := mload(_data)

            // Skip over the length field.
            //
            // Keep temporary variable so it can be incremented in place.
            //
            // NOTE: incrementing _data would result in an unusable
            //       _data variable after this assembly block
            let data := add(_data, 0x20)

            // Iterate until the bound is not met.
            for
                { let end := add(data, mul(len, 0x20)) }
                lt(data, end)
                { data := add(data, 0x20) }
            {
                sum := add(sum, mload(data))
            }
        }
    }
}
```

3.10.2 Access to External Variables, Functions and Libraries

You can access Solidity variables and other identifiers by using their name.

Local variables of value type are directly usable in inline assembly.

Local variables that refer to memory or calldata evaluate to the address of the variable in memory, resp. calldata, not the value itself.

For local storage variables or state variables, a single Yul identifier is not sufficient, since they do not necessarily occupy a single full storage slot. Therefore, their "address" is composed of a slot and a byte-offset inside that slot. To retrieve the slot pointed to by the variable `x`, you use `x.slot`, and to retrieve the byte-offset you use `x.offset`. Using `x` itself will result in an error.

For dynamic calldata arrays, you can access their calldata offset (in bytes) and length (number of elements) using `x.offset` and `x.length`. Both expressions can also be assigned to.

Local Solidity variables are available for assignments, for example:

```solidity
// SPDX-License-Identifier: GPL-3.0
pragma solidity ^0.7.0;

contract C {
    uint b;
    function f(uint x) public view returns (uint r) {
        assembly {
            // We ignore the storage slot offset, we know it is zero
            // in this special case.
            r := mul(x, sload(b.slot))
        }
    }
}
```

> **Warning:** If you access variables of a type that spans less than 256 bits (for example `uint64`, `address`, `bytes16` or `byte`), you cannot make any assumptions about bits not part of the encoding of the type. Especially, do not assume them to be zero. To be safe, always clear the data properly before you use it in a context where this is important: `uint32 x = f(); assembly { x := and(x, 0xffffffff) /* now use x */ }` To clean signed types, you can use the `signextend` opcode: `assembly { signextend(<num_bytes_of_x_minus_one>, x) }`

Since Solidity 0.6.0 the name of a inline assembly variable may not shadow any declaration visible in the scope of the inline assembly block (including variable, contract and function declarations).

Since Solidity 0.7.0, variables and functions declared inside the inline assembly block may not contain `.`, but using `.` is valid to access Solidity variables from outside the inline assembly block.

Assignments are possible to assembly-local variables and to function-local variables. Take care that when you assign to variables that point to memory or storage, you will only change the pointer and not the data.

You can assign to the `.slot` part of a local storage variable pointer. For these (structs, arrays or mappings), the `.offset` part is always zero. It is not possible to assign to the `.slot` or `.offset` part of a state variable, though.

3.10.3 Things to Avoid

Inline assembly might have a quite high-level look, but it actually is extremely low-level. Function calls, loops, ifs and switches are converted by simple rewriting rules and after that, the only thing the assembler does for you is re-arranging functional-style opcodes, counting stack height for variable access and removing stack slots for assembly-local variables when the end of their block is reached.

3.10.4 Conventions in Solidity

In contrast to EVM assembly, Solidity has types which are narrower than 256 bits, e.g. `uint24`. For efficiency, most arithmetic operations ignore the fact that types can be shorter than 256 bits, and the higher-order bits are cleaned when

necessary, i.e., shortly before they are written to memory or before comparisons are performed. This means that if you access such a variable from within inline assembly, you might have to manually clean the higher-order bits first.

Solidity manages memory in the following way. There is a "free memory pointer" at position `0x40` in memory. If you want to allocate memory, use the memory starting from where this pointer points at and update it. There is no guarantee that the memory has not been used before and thus you cannot assume that its contents are zero bytes. There is no built-in mechanism to release or free allocated memory. Here is an assembly snippet you can use for allocating memory that follows the process outlined above:

```
function allocate(length) -> pos {
  pos := mload(0x40)
  mstore(0x40, add(pos, length))
}
```

The first 64 bytes of memory can be used as "scratch space" for short-term allocation. The 32 bytes after the free memory pointer (i.e., starting at `0x60`) are meant to be zero permanently and is used as the initial value for empty dynamic memory arrays. This means that the allocatable memory starts at `0x80`, which is the initial value of the free memory pointer.

Elements in memory arrays in Solidity always occupy multiples of 32 bytes (this is even true for `byte[]`, but not for `bytes` and `string`). Multi-dimensional memory arrays are pointers to memory arrays. The length of a dynamic array is stored at the first slot of the array and followed by the array elements.

> **Warning:** Statically-sized memory arrays do not have a length field, but it might be added later to allow better convertibility between statically- and dynamically-sized arrays, so do not rely on this.

3.11 Cheatsheet

3.11.1 Order of Precedence of Operators

The following is the order of precedence for operators, listed in order of evaluation.

Precedence	Description	Operator
1	Postfix increment and decrement	`++, --`
	New expression	`new <typename>`
	Array subscripting	`<array>[<index>]`
	Member access	`<object>.<member>`
	Function-like call	`<func>(<args...>)`
	Parentheses	`(<statement>)`
2	Prefix increment and decrement	`++, --`
	Unary minus	`-`
	Unary operations	`delete`
	Logical NOT	`!`
	Bitwise NOT	`~`
3	Exponentiation	`**`
4	Multiplication, division and modulo	`*, /, %`
5	Addition and subtraction	`+, -`
6	Bitwise shift operators	`<<, >>`
7	Bitwise AND	`&`
8	Bitwise XOR	`^`
9	Bitwise OR	`|`
10	Inequality operators	`<, >, <=, >=`
11	Equality operators	`==, !=`
12	Logical AND	`&&`
13	Logical OR	`||`
14	Ternary operator	`<conditional> ? <if-true> : <if-false>`
	Assignment operators	`=, |=, ^=, &=, <<=, >>=, +=, -=, *=, /=, %=`
15	Comma operator	`,`

3.11.2 Global Variables

- `abi.decode(bytes memory encodedData, (...)) returns (...)`: *ABI*-decodes the provided data. The types are given in parentheses as second argument. Example: `(uint a, uint[2] memory b, bytes memory c) = abi.decode(data, (uint, uint[2], bytes))`

- `abi.encode(...) returns (bytes memory)`: *ABI*-encodes the given arguments

- `abi.encodePacked(...) returns (bytes memory)`: Performs *packed encoding* of the given arguments. Note that this encoding can be ambiguous!

- `abi.encodeWithSelector(bytes4 selector, ...) returns (bytes memory)`: *ABI*-encodes the given arguments starting from the second and prepends the given four-byte selector

- `abi.encodeWithSignature(string memory signature, ...) returns (bytes memory)`: Equivalent to `abi.encodeWithSelector(bytes4(keccak256(bytes(signature)), ...)`

- `block.coinbase (address payable)`: current block miner's address

- `block.difficulty (uint)`: current block difficulty

- `block.gaslimit (uint)`: current block gaslimit

- `block.number (uint)`: current block number

- `block.timestamp (uint)`: current block timestamp

- `gasleft() returns (uint256)`: remaining gas

- `msg.data` (`bytes`): complete calldata

- `msg.sender` (`address payable`): sender of the message (current call)

- `msg.value` (`uint`): number of wei sent with the message

- `tx.gasprice` (`uint`): gas price of the transaction

- `tx.origin` (`address payable`): sender of the transaction (full call chain)

- `assert(bool condition)`: abort execution and revert state changes if condition is `false` (use for internal error)

- `require(bool condition)`: abort execution and revert state changes if condition is `false` (use for malformed input or error in external component)

- `require(bool condition, string memory message)`: abort execution and revert state changes if condition is `false` (use for malformed input or error in external component). Also provide error message.

- `revert()`: abort execution and revert state changes

- `revert(string memory message)`: abort execution and revert state changes providing an explanatory string

- `blockhash(uint blockNumber) returns (bytes32)`: hash of the given block - only works for 256 most recent blocks

- `keccak256(bytes memory) returns (bytes32)`: compute the Keccak-256 hash of the input

- `sha256(bytes memory) returns (bytes32)`: compute the SHA-256 hash of the input

- `ripemd160(bytes memory) returns (bytes20)`: compute the RIPEMD-160 hash of the input

- `ecrecover(bytes32 hash, uint8 v, bytes32 r, bytes32 s) returns (address)`: recover address associated with the public key from elliptic curve signature, return zero on error

- `addmod(uint x, uint y, uint k) returns (uint)`: compute $(x + y)$ % k where the addition is performed with arbitrary precision and does not wrap around at $2**256$. Assert that `k != 0` starting from version 0.5.0.

- `mulmod(uint x, uint y, uint k) returns (uint)`: compute $(x * y)$ % k where the multiplication is performed with arbitrary precision and does not wrap around at $2**256$. Assert that `k != 0` starting from version 0.5.0.

- `this` (current contract's type): the current contract, explicitly convertible to `address` or `address payable`

- `super`: the contract one level higher in the inheritance hierarchy

- `selfdestruct(address payable recipient)`: destroy the current contract, sending its funds to the given address

- `<address>.balance` (`uint256`): balance of the *Address* in Wei

- `<address payable>.send(uint256 amount) returns (bool)`: send given amount of Wei to *Address*, returns `false` on failure

- `<address payable>.transfer(uint256 amount)`: send given amount of Wei to *Address*, throws on failure

- `type(C).name` (`string`): the name of the contract

- `type(C).creationCode` (`bytes memory`): creation bytecode of the given contract, see *Type Information*.

- `type(C).runtimeCode` (`bytes memory`): runtime bytecode of the given contract, see *Type Information*.

- `type(I).interfaceId(bytes4)`: value containing the EIP-165 interface identifier of the given interface, see *Type Information*.
- `type(T).min(T)`: the minimum value representable by the integer type `T`, see *Type Information*.
- `type(T).max(T)`: the maximum value representable by the integer type `T`, see *Type Information*.

Note: Do not rely on `block.timestamp` or `blockhash` as a source of randomness, unless you know what you are doing.

Both the timestamp and the block hash can be influenced by miners to some degree. Bad actors in the mining community can for example run a casino payout function on a chosen hash and just retry a different hash if they did not receive any money.

The current block timestamp must be strictly larger than the timestamp of the last block, but the only guarantee is that it will be somewhere between the timestamps of two consecutive blocks in the canonical chain.

Note: The block hashes are not available for all blocks for scalability reasons. You can only access the hashes of the most recent 256 blocks, all other values will be zero.

Note: In version 0.5.0, the following aliases were removed: `suicide` as alias for `selfdestruct`, `msg.gas` as alias for `gasleft`, `block.blockhash` as alias for `blockhash` and `sha3` as alias for `keccak256`.

Note: In version 0.7.0, the alias `now` (for `block.timestamp`) was removed.

3.11.3 Function Visibility Specifiers

```
function myFunction() <visibility specifier> returns (bool) {
    return true;
}
```

- `public`: visible externally and internally (creates a *getter function* for storage/state variables)
- `private`: only visible in the current contract
- `external`: only visible externally (only for functions) - i.e. can only be message-called (via `this.func`)
- `internal`: only visible internally

3.11.4 Modifiers

- `pure` for functions: Disallows modification or access of state.
- `view` for functions: Disallows modification of state.
- `payable` for functions: Allows them to receive Ether together with a call.
- `constant` for state variables: Disallows assignment (except initialisation), does not occupy storage slot.
- `immutable` for state variables: Allows exactly one assignment at construction time and is constant afterwards. Is stored in code.

- `anonymous` for events: Does not store event signature as topic.

- `indexed` for event parameters: Stores the parameter as topic.

- `virtual` for functions and modifiers: Allows the function's or modifier's behaviour to be changed in derived contracts.

- `override`: States that this function, modifier or public state variable changes the behaviour of a function or modifier in a base contract.

3.11.5 Reserved Keywords

These keywords are reserved in Solidity. They might become part of the syntax in the future:

`after`, `alias`, `apply`, `auto`, `case`, `copyof`, `default`, `define`, `final`, `immutable`, `implements`, `in`, `inline`, `let`, `macro`, `match`, `mutable`, `null`, `of`, `partial`, `promise`, `reference`, `relocatable`, `sealed`, `sizeof`, `static`, `supports`, `switch`, `typedef`, `typeof`, `unchecked`.

3.12 Layout of State Variables in Storage

Statically-sized variables (everything except mapping and dynamically-sized array types) are laid out contiguously in storage starting from position `0`. Multiple, contiguous items that need less than 32 bytes are packed into a single storage slot if possible, according to the following rules:

- The first item in a storage slot is stored lower-order aligned.

- Elementary types use only as many bytes as are necessary to store them.

- If an elementary type does not fit the remaining part of a storage slot, it is moved to the next storage slot.

- Structs and array data always start a new slot and occupy whole slots (but items inside a struct or array are packed tightly according to these rules).

For contracts that use inheritance, the ordering of state variables is determined by the C3-linearized order of contracts starting with the most base-ward contract. If allowed by the above rules, state variables from different contracts do share the same storage slot.

The elements of structs and arrays are stored after each other, just as if they were given explicitly.

> **Warning:** When using elements that are smaller than 32 bytes, your contract's gas usage may be higher. This is because the EVM operates on 32 bytes at a time. Therefore, if the element is smaller than that, the EVM must use more operations in order to reduce the size of the element from 32 bytes to the desired size.
>
> It is only beneficial to use reduced-size arguments if you are dealing with storage values because the compiler will pack multiple elements into one storage slot, and thus, combine multiple reads or writes into a single operation. When dealing with function arguments or memory values, there is no inherent benefit because the compiler does not pack these values.
>
> Finally, in order to allow the EVM to optimize for this, ensure that you try to order your storage variables and `struct` members such that they can be packed tightly. For example, declaring your storage variables in the order of `uint128`, `uint128`, `uint256` instead of `uint128`, `uint256`, `uint128`, as the former will only take up two slots of storage whereas the latter will take up three.

> **Note:** The layout of state variables in storage is considered to be part of the external interface of Solidity due to the fact that storage pointers can be passed to libraries. This means that any change to the rules outlined in this section is

considered a breaking change of the language and due to its critical nature should be considered very carefully before being executed.

3.12.1 Mappings and Dynamic Arrays

Due to their unpredictable size, mapping and dynamically-sized array types use a Keccak-256 hash computation to find the starting position of the value or the array data. These starting positions are always full stack slots.

The mapping or the dynamic array itself occupies a slot in storage at some position `p` according to the above rule (or by recursively applying this rule for mappings of mappings or arrays of arrays). For dynamic arrays, this slot stores the number of elements in the array (byte arrays and strings are an exception, see *below*). For mappings, the slot is unused (but it is needed so that two equal mappings after each other will use a different hash distribution). Array data is located at `keccak256(p)` and the value corresponding to a mapping key `k` is located at `keccak256(k . p)` where `.` is concatenation. If the value is again a non-elementary type, the positions are found by adding an offset of `keccak256(k . p)`.

So for the following contract snippet the position of `data[4][9].b` is at `keccak256(uint256(9) . keccak256(uint256(4) . uint256(1))) + 1`:

```
// SPDX-License-Identifier: GPL-3.0
pragma solidity >=0.4.0 <0.8.0;


contract C {
    struct S { uint a; uint b; }
    uint x;
    mapping(uint => mapping(uint => S)) data;
}
```

bytes and string

`bytes` and `string` are encoded identically. For short byte arrays, they store their data in the same slot where the length is also stored. In particular: if the data is at most `31` bytes long, it is stored in the higher-order bytes (left aligned) and the lowest-order byte stores `length * 2`. For byte arrays that store data which is `32` or more bytes long, the main slot stores `length * 2 + 1` and the data is stored as usual in `keccak256(slot)`. This means that you can distinguish a short array from a long array by checking if the lowest bit is set: short (not set) and long (set).

Note: Handling invalidly encoded slots is currently not supported but may be added in the future.

3.12.2 JSON Output

The storage layout of a contract can be requested via the *standard JSON interface*. The output is a JSON object containing two keys, `storage` and `types`. The `storage` object is an array where each element has the following form:

```
{
    "astId": 2,
    "contract": "fileA:A",
    "label": "x",
```

```
    "offset": 0,
    "slot": "0",
    "type": "t_uint256"
}
```

The example above is the storage layout of `contract A { uint x; }` from source unit `fileA` and

- `astId` is the id of the AST node of the state variable's declaration

- `contract` is the name of the contract including its path as prefix

- `label` is the name of the state variable

- `offset` is the offset in bytes within the storage slot according to the encoding

- `slot` is the storage slot where the state variable resides or starts. This number may be very large and therefore its JSON value is represented as a string.

- `type` is an identifier used as key to the variable's type information (described in the following)

The given `type`, in this case `t_uint256` represents an element in `types`, which has the form:

```
{
    "encoding": "inplace",
    "label": "uint256",
    "numberOfBytes": "32",
}
```

where

- `encoding` how the data is encoded in storage, where the possible values are:

    - `inplace`: data is laid out contiguously in storage (see *above*).

    - `mapping`: Keccak-256 hash-based method (see *above*).

    - `dynamic_array`: Keccak-256 hash-based method (see *above*).

    - `bytes`: single slot or Keccak-256 hash-based depending on the data size (see *above*).

- `label` is the canonical type name.

- `numberOfBytes` is the number of used bytes (as a decimal string). Note that if `numberOfBytes > 32` this means that more than one slot is used.

Some types have extra information besides the four above. Mappings contain its `key` and `value` types (again referencing an entry in this mapping of types), arrays have its `base` type, and structs list their `members` in the same format as the top-level `storage` (see *above*).

Note: The JSON output format of a contract's storage layout is still considered experimental and is subject to change in non-breaking releases of Solidity.

The following example shows a contract and its storage layout, containing value and reference types, types that are encoded packed, and nested types.

```
// SPDX-License-Identifier: GPL-3.0
pragma solidity >=0.4.0 <0.8.0;
contract A {
    struct S {
        uint128 a;
```

```
        uint128 b;
        uint[2] staticArray;
        uint[] dynArray;
    }

    uint x;
    uint y;
    S s;
    address addr;
    mapping (uint => mapping (address => bool)) map;
    uint[] array;
    string s1;
    bytes b1;
}
```

```
"storageLayout": {
  "storage": [
    {
      "astId": 14,
      "contract": "fileA:A",
      "label": "x",
      "offset": 0,
      "slot": "0",
      "type": "t_uint256"
    },
    {
      "astId": 16,
      "contract": "fileA:A",
      "label": "y",
      "offset": 0,
      "slot": "1",
      "type": "t_uint256"
    },
    {
      "astId": 18,
      "contract": "fileA:A",
      "label": "s",
      "offset": 0,
      "slot": "2",
      "type": "t_struct(S)12_storage"
    },
    {
      "astId": 20,
      "contract": "fileA:A",
      "label": "addr",
      "offset": 0,
      "slot": "6",
      "type": "t_address"
    },
    {
      "astId": 26,
      "contract": "fileA:A",
      "label": "map",
      "offset": 0,
      "slot": "7",
      "type": "t_mapping(t_uint256,t_mapping(t_address,t_bool))"
```

```
    },
    {
      "astId": 29,
      "contract": "fileA:A",
      "label": "array",
      "offset": 0,
      "slot": "8",
      "type": "t_array(t_uint256)dyn_storage"
    },
    {
      "astId": 31,
      "contract": "fileA:A",
      "label": "s1",
      "offset": 0,
      "slot": "9",
      "type": "t_string_storage"
    },
    {
      "astId": 33,
      "contract": "fileA:A",
      "label": "b1",
      "offset": 0,
      "slot": "10",
      "type": "t_bytes_storage"
    }
  ],
  "types": {
    "t_address": {
      "encoding": "inplace",
      "label": "address",
      "numberOfBytes": "20"
    },
    "t_array(t_uint256)2_storage": {
      "base": "t_uint256",
      "encoding": "inplace",
      "label": "uint256[2]",
      "numberOfBytes": "64"
    },
    "t_array(t_uint256)dyn_storage": {
      "base": "t_uint256",
      "encoding": "dynamic_array",
      "label": "uint256[]",
      "numberOfBytes": "32"
    },
    "t_bool": {
      "encoding": "inplace",
      "label": "bool",
      "numberOfBytes": "1"
    },
    "t_bytes_storage": {
      "encoding": "bytes",
      "label": "bytes",
      "numberOfBytes": "32"
    },
    "t_mapping(t_address,t_bool)": {
      "encoding": "mapping",
      "key": "t_address",
```

```
          "label": "mapping(address => bool)",
          "numberOfBytes": "32",
          "value": "t_bool"
        },
        "t_mapping(t_uint256,t_mapping(t_address,t_bool))": {
          "encoding": "mapping",
          "key": "t_uint256",
          "label": "mapping(uint256 => mapping(address => bool))",
          "numberOfBytes": "32",
          "value": "t_mapping(t_address,t_bool)"
        },
        "t_string_storage": {
          "encoding": "bytes",
          "label": "string",
          "numberOfBytes": "32"
        },
        "t_struct(S)12_storage": {
          "encoding": "inplace",
          "label": "struct A.S",
          "members": [
            {
              "astId": 2,
              "contract": "fileA:A",
              "label": "a",
              "offset": 0,
              "slot": "0",
              "type": "t_uint128"
            },
            {
              "astId": 4,
              "contract": "fileA:A",
              "label": "b",
              "offset": 16,
              "slot": "0",
              "type": "t_uint128"
            },
            {
              "astId": 8,
              "contract": "fileA:A",
              "label": "staticArray",
              "offset": 0,
              "slot": "1",
              "type": "t_array(t_uint256)2_storage"
            },
            {
              "astId": 11,
              "contract": "fileA:A",
              "label": "dynArray",
              "offset": 0,
              "slot": "3",
              "type": "t_array(t_uint256)dyn_storage"
            }
          ],
          "numberOfBytes": "128"
        },
        "t_uint128": {
          "encoding": "inplace",
```

3.12. Layout of State Variables in Storage 139

```
        "label": "uint128",
        "numberOfBytes": "16"
    },
    "t_uint256": {
        "encoding": "inplace",
        "label": "uint256",
        "numberOfBytes": "32"
    }
  }
}
```

3.13 Layout in Memory

Solidity reserves four 32-byte slots, with specific byte ranges (inclusive of endpoints) being used as follows:

- `0x00 - 0x3f` (64 bytes): scratch space for hashing methods
- `0x40 - 0x5f` (32 bytes): currently allocated memory size (aka. free memory pointer)
- `0x60 - 0x7f` (32 bytes): zero slot

Scratch space can be used between statements (i.e. within inline assembly). The zero slot is used as initial value for dynamic memory arrays and should never be written to (the free memory pointer points to `0x80` initially).

Solidity always places new objects at the free memory pointer and memory is never freed (this might change in the future).

Elements in memory arrays in Solidity always occupy multiples of 32 bytes (this is even true for `byte[]`, but not for `bytes` and `string`). Multi-dimensional memory arrays are pointers to memory arrays. The length of a dynamic array is stored at the first slot of the array and followed by the array elements.

> **Warning:** There are some operations in Solidity that need a temporary memory area larger than 64 bytes and therefore will not fit into the scratch space. They will be placed where the free memory points to, but given their short lifetime, the pointer is not updated. The memory may or may not be zeroed out. Because of this, one should not expect the free memory to point to zeroed out memory.
>
> While it may seem like a good idea to use `msize` to arrive at a definitely zeroed out memory area, using such a pointer non-temporarily without updating the free memory pointer can have unexpected results.

3.13.1 Differences to Layout in Storage

As described above the layout in memory is different from the layout in *storage*. Below there are some examples.

Example for Difference in Arrays

The following array occupies 32 bytes (1 slot) in storage, but 128 bytes (4 items with 32 bytes each) in memory.

```
uint8[4] a;
```

Example for Difference in Struct Layout

The following struct occupies 96 bytes (3 slots of 32 bytes) in storage, but 128 bytes (4 items with 32 bytes each) in memory.

```solidity
struct S {
    uint a;
    uint b;
    uint8 c;
    uint8 d;
}
```

3.14 Layout of Call Data

The input data for a function call is assumed to be in the format defined by the *ABI specification*. Among others, the ABI specification requires arguments to be padded to multiples of 32 bytes. The internal function calls use a different convention.

Arguments for the constructor of a contract are directly appended at the end of the contract's code, also in ABI encoding. The constructor will access them through a hard-coded offset, and not by using the `codesize` opcode, since this of course changes when appending data to the code.

3.15 Cleaning Up Variables

When a value is shorter than 256 bit, in some cases the remaining bits must be cleaned. The Solidity compiler is designed to clean such remaining bits before any operations that might be adversely affected by the potential garbage in the remaining bits. For example, before writing a value to memory, the remaining bits need to be cleared because the memory contents can be used for computing hashes or sent as the data of a message call. Similarly, before storing a value in the storage, the remaining bits need to be cleaned because otherwise the garbled value can be observed.

On the other hand, we do not clean the bits if the immediately following operation is not affected. For instance, since any non-zero value is considered `true` by `JUMPI` instruction, we do not clean the boolean values before they are used as the condition for `JUMPI`.

In addition to the design principle above, the Solidity compiler cleans input data when it is loaded onto the stack.

Different types have different rules for cleaning up invalid values:

Type	Valid Values	Invalid Values Mean
enum of n members	0 until n - 1	exception
bool	0 or 1	1
signed integers	sign-extended word	currently silently wraps; in the future exceptions will be thrown
unsigned integers	higher bits zeroed	currently silently wraps; in the future exceptions will be thrown

3.16 Source Mappings

As part of the AST output, the compiler provides the range of the source code that is represented by the respective node in the AST. This can be used for various purposes ranging from static analysis tools that report errors based on the AST and debugging tools that highlight local variables and their uses.

Furthermore, the compiler can also generate a mapping from the bytecode to the range in the source code that generated the instruction. This is again important for static analysis tools that operate on bytecode level and for displaying the current position in the source code inside a debugger or for breakpoint handling. This mapping also contains other information, like the jump type and the modifier depth (see below).

Both kinds of source mappings use integer identifiers to refer to source files. The identifier of a source file is stored in `output['sources'][sourceName]['id']` where `output` is the output of the standard-json compiler interface parsed as JSON. For some utility routines, the compiler generates "internal" source files that are not part of the original input but are referenced from the source mappings. These source files together with their identifiers can be obtained via `output['contracts'][sourceName][contractName]['evm']['bytecode']['generatedSources']`.

Note: In the case of instructions that are not associated with any particular source file, the source mapping assigns an integer identifier of `-1`. This may happen for bytecode sections stemming from compiler-generated inline assembly statements.

The source mappings inside the AST use the following notation:

`s:l:f`

Where `s` is the byte-offset to the start of the range in the source file, `l` is the length of the source range in bytes and `f` is the source index mentioned above.

The encoding in the source mapping for the bytecode is more complicated: It is a list of `s:l:f:j:m` separated by `;`. Each of these elements corresponds to an instruction, i.e. you cannot use the byte offset but have to use the instruction offset (push instructions are longer than a single byte). The fields `s`, `l` and `f` are as above. `j` can be either `i`, `o` or `-` signifying whether a jump instruction goes into a function, returns from a function or is a regular jump as part of e.g. a loop. The last field, `m`, is an integer that denotes the "modifier depth". This depth is increased whenever the placeholder statement (`_`) is entered in a modifier and decreased when it is left again. This allows debuggers to track tricky cases like the same modifier being used twice or multiple placeholder statements being used in a single modifier.

In order to compress these source mappings especially for bytecode, the following rules are used:

- If a field is empty, the value of the preceding element is used.
- If a `:` is missing, all following fields are considered empty.

This means the following source mappings represent the same information:

`1:2:1;1:9:1;2:1:2;2:1:2;2:1:2`

`1:2:1;:9;2:1:2;;`

3.17 The Optimiser

This section discusses the optimiser that was first added to Solidity, which operates on opcode streams. For information on the new Yul-based optimiser, please see the readme on github.

The Solidity optimiser operates on assembly. It splits the sequence of instructions into basic blocks at `JUMPs` and `JUMPDESTs`. Inside these blocks, the optimiser analyses the instructions and records every modification to the stack, memory, or storage as an expression which consists of an instruction and a list of arguments which are pointers to other expressions. The optimiser uses a component called "CommonSubexpressionEliminator" that amongst other tasks, finds expressions that are always equal (on every input) and combines them into an expression class. The optimiser first tries to find each new expression in a list of already known expressions. If this does not work, it simplifies the expression according to rules like `constant + constant = sum_of_constants` or `X * 1 = X`. Since this is a recursive process, we can also apply the latter rule if the second factor is a more complex expression where we know that it always evaluates to one. Modifications to storage and memory locations have to erase knowledge about

storage and memory locations which are not known to be different. If we first write to location x and then to location y and both are input variables, the second could overwrite the first, so we do not know what is stored at x after we wrote to y. If simplification of the expression x - y evaluates to a non-zero constant, we know that we can keep our knowledge about what is stored at x.

After this process, we know which expressions have to be on the stack at the end, and have a list of modifications to memory and storage. This information is stored together with the basic blocks and is used to link them. Furthermore, knowledge about the stack, storage and memory configuration is forwarded to the next block(s). If we know the targets of all JUMP and JUMPI instructions, we can build a complete control flow graph of the program. If there is only one target we do not know (this can happen as in principle, jump targets can be computed from inputs), we have to erase all knowledge about the input state of a block as it can be the target of the unknown JUMP. If the optimiser finds a JUMPI whose condition evaluates to a constant, it transforms it to an unconditional jump.

As the last step, the code in each block is re-generated. The optimiser creates a dependency graph from the expressions on the stack at the end of the block, and it drops every operation that is not part of this graph. It generates code that applies the modifications to memory and storage in the order they were made in the original code (dropping modifications which were found not to be needed). Finally, it generates all values that are required to be on the stack in the correct place.

These steps are applied to each basic block and the newly generated code is used as replacement if it is smaller. If a basic block is split at a JUMPI and during the analysis, the condition evaluates to a constant, the JUMPI is replaced depending on the value of the constant. Thus code like

```solidity
uint x = 7;
data[7] = 9;
if (data[x] != x + 2)
  return 2;
else
  return 1;
```

still simplifies to code which you can compile even though the instructions contained a jump in the beginning of the process:

```solidity
data[7] = 9;
return 1;
```

3.18 Contract Metadata

The Solidity compiler automatically generates a JSON file, the contract metadata, that contains information about the compiled contract. You can use this file to query the compiler version, the sources used, the ABI and NatSpec documentation to more safely interact with the contract and verify its source code.

The compiler appends by default the IPFS hash of the metadata file to the end of the bytecode (for details, see below) of each contract, so that you can retrieve the file in an authenticated way without having to resort to a centralized data provider. The other available options are the Swarm hash and not appending the metadata hash to the bytecode. These can be configured via the *Standard JSON Interface*.

You have to publish the metadata file to IPFS, Swarm, or another service so that others can access it. You create the file by using the solc --metadata command that generates a file called ContractName_meta.json. It contains IPFS and Swarm references to the source code, so you have to upload all source files and the metadata file.

The metadata file has the following format. The example below is presented in a human-readable way. Properly formatted metadata should use quotes correctly, reduce whitespace to a minimum and sort the keys of all objects to arrive at a unique formatting. Comments are not permitted and used here only for explanatory purposes.

```
{
  // Required: The version of the metadata format
  version: "1",
  // Required: Source code language, basically selects a "sub-version"
  // of the specification
  language: "Solidity",
  // Required: Details about the compiler, contents are specific
  // to the language.
  compiler: {
    // Required for Solidity: Version of the compiler
    version: "0.4.6+commit.2dabbdf0.Emscripten.clang",
    // Optional: Hash of the compiler binary which produced this output
    keccak256: "0x123..."
  },
  // Required: Compilation source files/source units, keys are file names
  sources:
  {
    "myFile.sol": {
      // Required: keccak256 hash of the source file
      "keccak256": "0x123...",
      // Required (unless "content" is used, see below): Sorted URL(s)
      // to the source file, protocol is more or less arbitrary, but a
      // Swarm URL is recommended
      "urls": [ "bzzr://56ab..." ],
      // Optional: SPDX license identifier as given in the source file
      "license": "MIT"
    },
    "destructible": {
      // Required: keccak256 hash of the source file
      "keccak256": "0x234...",
      // Required (unless "url" is used): literal contents of the source file
      "content": "contract destructible is owned { function destroy() { if (msg.
→sender == owner) selfdestruct(owner); } }"
    }
  },
  // Required: Compiler settings
  settings:
  {
    // Required for Solidity: Sorted list of remappings
    remappings: [ ":g=/dir" ],
    // Optional: Optimizer settings. The fields "enabled" and "runs" are deprecated
    // and are only given for backwards-compatibility.
    optimizer: {
      enabled: true,
      runs: 500,
      details: {
        // peephole defaults to "true"
        peephole: true,
        // jumpdestRemover defaults to "true"
        jumpdestRemover: true,
        orderLiterals: false,
        deduplicate: false,
        cse: false,
        constantOptimizer: false,
        yul: true,
        // Optional: Only present if "yul" is "true"
        yulDetails: {
```

(continues on next page)

```
          stackAllocation: false,
          optimizerSteps: "dhfoDgvulfnTUtnIf..."
        }
      }
    },
    metadata: {
      // Reflects the setting used in the input json, defaults to false
      useLiteralContent: true,
      // Reflects the setting used in the input json, defaults to "ipfs"
      bytecodeHash: "ipfs"
    }
    // Required for Solidity: File and name of the contract or library this
    // metadata is created for.
    compilationTarget: {
      "myFile.sol": "MyContract"
    },
    // Required for Solidity: Addresses for libraries used
    libraries: {
      "MyLib": "0x123123..."
    }
  },
  // Required: Generated information about the contract.
  output:
  {
    // Required: ABI definition of the contract
    abi: [ ... ],
    // Required: NatSpec user documentation of the contract
    userdoc: [ ... ],
    // Required: NatSpec developer documentation of the contract
    devdoc: [ ... ],
  }
}
```

> **Warning:** Since the bytecode of the resulting contract contains the metadata hash by default, any change to the metadata might result in a change of the bytecode. This includes changes to a filename or path, and since the metadata includes a hash of all the sources used, a single whitespace change results in different metadata, and different bytecode.

> **Note:** The ABI definition above has no fixed order. It can change with compiler versions. Starting from Solidity version 0.5.12, though, the array maintains a certain order.

3.18.1 Encoding of the Metadata Hash in the Bytecode

Because we might support other ways to retrieve the metadata file in the future, the mapping `{"ipfs": <IPFS hash>, "solc": <compiler version>}` is stored CBOR-encoded. Since the mapping might contain more keys (see below) and the beginning of that encoding is not easy to find, its length is added in a two-byte big-endian encoding. The current version of the Solidity compiler usually adds the following to the end of the deployed bytecode:

```
0xa2
0x64 'i' 'p' 'f' 's' 0x58 0x22 <34 bytes IPFS hash>
```

```
0x64 's' 'o' 'l' 'c' 0x43 <3 byte version encoding>
0x00 0x33
```

So in order to retrieve the data, the end of the deployed bytecode can be checked to match that pattern and use the IPFS hash to retrieve the file.

Whereas release builds of solc use a 3 byte encoding of the version as shown above (one byte each for major, minor and patch version number), prerelease builds will instead use a complete version string including commit hash and build date.

Note: The CBOR mapping can also contain other keys, so it is better to fully decode the data instead of relying on it starting with `0xa264`. For example, if any experimental features that affect code generation are used, the mapping will also contain `"experimental": true`.

Note: The compiler currently uses the IPFS hash of the metadata by default, but it may also use the bzzr1 hash or some other hash in the future, so do not rely on this sequence to start with `0xa2 0x64 'i' 'p' 'f' 's'`. We might also add additional data to this CBOR structure, so the best option is to use a proper CBOR parser.

3.18.2 Usage for Automatic Interface Generation and NatSpec

The metadata is used in the following way: A component that wants to interact with a contract (e.g. Mist or any wallet) retrieves the code of the contract, from that the IPFS/Swarm hash of a file which is then retrieved. That file is JSON-decoded into a structure like above.

The component can then use the ABI to automatically generate a rudimentary user interface for the contract.

Furthermore, the wallet can use the NatSpec user documentation to display a confirmation message to the user whenever they interact with the contract, together with requesting authorization for the transaction signature.

For additional information, read *Ethereum Natural Language Specification (NatSpec) format*.

3.18.3 Usage for Source Code Verification

In order to verify the compilation, sources can be retrieved from IPFS/Swarm via the link in the metadata file. The compiler of the correct version (which is checked to be part of the "official" compilers) is invoked on that input with the specified settings. The resulting bytecode is compared to the data of the creation transaction or `CREATE` opcode data. This automatically verifies the metadata since its hash is part of the bytecode. Excess data corresponds to the constructor input data, which should be decoded according to the interface and presented to the user.

In the repository sourcify (npm package) you can see example code that shows how to use this feature.

3.19 Contract ABI Specification

3.19.1 Basic Design

The Contract Application Binary Interface (ABI) is the standard way to interact with contracts in the Ethereum ecosystem, both from outside the blockchain and for contract-to-contract interaction. Data is encoded according to its type, as described in this specification. The encoding is not self describing and thus requires a schema in order to decode.

We assume the interface functions of a contract are strongly typed, known at compilation time and static. We assume that all contracts will have the interface definitions of any contracts they call available at compile-time.

This specification does not address contracts whose interface is dynamic or otherwise known only at run-time.

3.19.2 Function Selector

The first four bytes of the call data for a function call specifies the function to be called. It is the first (left, high-order in big-endian) four bytes of the Keccak-256 hash of the signature of the function. The signature is defined as the canonical expression of the basic prototype without data location specifier, i.e. the function name with the parenthesised list of parameter types. Parameter types are split by a single comma - no spaces are used.

Note: The return type of a function is not part of this signature. In *Solidity's function overloading* return types are not considered. The reason is to keep function call resolution context-independent. The *JSON description of the ABI* however contains both inputs and outputs.

3.19.3 Argument Encoding

Starting from the fifth byte, the encoded arguments follow. This encoding is also used in other places, e.g. the return values and also event arguments are encoded in the same way, without the four bytes specifying the function.

3.19.4 Types

The following elementary types exist:

- `uint<M>`: unsigned integer type of `M` bits, `0 < M <= 256`, `M % 8 == 0`. e.g. `uint32`, `uint8`, `uint256`.

- `int<M>`: two's complement signed integer type of `M` bits, `0 < M <= 256`, `M % 8 == 0`.

- `address`: equivalent to `uint160`, except for the assumed interpretation and language typing. For computing the function selector, `address` is used.

- `uint`, `int`: synonyms for `uint256`, `int256` respectively. For computing the function selector, `uint256` and `int256` have to be used.

- `bool`: equivalent to `uint8` restricted to the values 0 and 1. For computing the function selector, `bool` is used.

- `fixed<M>x<N>`: signed fixed-point decimal number of `M` bits, `8 <= M <= 256`, `M % 8 == 0`, and `0 < N <= 80`, which denotes the value `v` as `v / (10 ** N)`.

- `ufixed<M>x<N>`: unsigned variant of `fixed<M>x<N>`.

- `fixed`, `ufixed`: synonyms for `fixed128x18`, `ufixed128x18` respectively. For computing the function selector, `fixed128x18` and `ufixed128x18` have to be used.

- `bytes<M>`: binary type of `M` bytes, `0 < M <= 32`.

- `function`: an address (20 bytes) followed by a function selector (4 bytes). Encoded identical to `bytes24`.

The following (fixed-size) array type exists:

- `<type>[M]`: a fixed-length array of `M` elements, `M >= 0`, of the given type.

> **Note:** While this ABI specification can express fixed-length arrays with zero elements, they're not supported by the compiler.

The following non-fixed-size types exist:

- `bytes`: dynamic sized byte sequence.

- `string`: dynamic sized unicode string assumed to be UTF-8 encoded.

- `<type>[]`: a variable-length array of elements of the given type.

Types can be combined to a tuple by enclosing them inside parentheses, separated by commas:

- `(T1,T2,...,Tn)`: tuple consisting of the types `T1, ..., Tn, n >= 0`

It is possible to form tuples of tuples, arrays of tuples and so on. It is also possible to form zero-tuples (where `n == 0`).

Mapping Solidity to ABI types

Solidity supports all the types presented above with the same names with the exception of tuples. On the other hand, some Solidity types are not supported by the ABI. The following table shows on the left column Solidity types that are not part of the ABI, and on the right column the ABI types that represent them.

Solidity	ABI
address payable	`address`
contract	`address`
enum	smallest `uint` type that is large enough to hold all values For example, an `enum` of 256 values or less is mapped to `uint8` and an `enum` of more than 256 values is mapped to `uint16`.
struct	`tuple`

3.19.5 Design Criteria for the Encoding

The encoding is designed to have the following properties, which are especially useful if some arguments are nested arrays:

1. The number of reads necessary to access a value is at most the depth of the value inside the argument array structure, i.e. four reads are needed to retrieve `a_i[k][l][r]`. In a previous version of the ABI, the number of reads scaled linearly with the total number of dynamic parameters in the worst case.

2. The data of a variable or array element is not interleaved with other data and it is relocatable, i.e. it only uses relative "addresses".

3.19.6 Formal Specification of the Encoding

We distinguish static and dynamic types. Static types are encoded in-place and dynamic types are encoded at a separately allocated location after the current block.

Definition: The following types are called "dynamic":

- `bytes`

- `string`

- `T[]` for any `T`

- `T[k]` for any dynamic `T` and any `k >= 0`

- `(T1,...,Tk)` if `Ti` is dynamic for some `1 <= i <= k`

All other types are called "static".

Definition: `len(a)` is the number of bytes in a binary string `a`. The type of `len(a)` is assumed to be `uint256`.

We define `enc`, the actual encoding, as a mapping of values of the ABI types to binary strings such that `len(enc(X))` depends on the value of `X` if and only if the type of `X` is dynamic.

Definition: For any ABI value `X`, we recursively define `enc(X)`, depending on the type of `X` being

- `(T1,...,Tk)` for `k >= 0` and any types `T1, ..., Tk`

  `enc(X) = head(X(1)) ... head(X(k)) tail(X(1)) ... tail(X(k))`

  where `X = (X(1), ..., X(k))` and `head` and `tail` are defined for `Ti` as follows:

  if `Ti` is static:

      `head(X(i)) = enc(X(i))` and `tail(X(i)) = ""` (the empty string)

  otherwise, i.e. if `Ti` is dynamic:

      `head(X(i)) = enc(len( head(X(1)) ... head(X(k)) tail(X(1)) ...`
      `tail(X(i-1)) ))tail(X(i)) = enc(X(i))`

  Note that in the dynamic case, `head(X(i))` is well-defined since the lengths of the head parts only depend on the types and not the values. The value of `head(X(i))` is the offset of the beginning of `tail(X(i))` relative to the start of `enc(X)`.

- `T[k]` for any `T` and `k`:

  `enc(X) = enc((X[0], ..., X[k-1]))`

  i.e. it is encoded as if it were a tuple with `k` elements of the same type.

- `T[]` where `X` has `k` elements (`k` is assumed to be of type `uint256`):

  `enc(X) = enc(k) enc([X[0], ..., X[k-1]])`

  i.e. it is encoded as if it were an array of static size `k`, prefixed with the number of elements.

- `bytes`, of length `k` (which is assumed to be of type `uint256`):

  `enc(X) = enc(k) pad_right(X)`, i.e. the number of bytes is encoded as a `uint256` followed by the actual value of `X` as a byte sequence, followed by the minimum number of zero-bytes such that `len(enc(X))` is a multiple of 32.

- `string`:

  `enc(X) = enc(enc_utf8(X))`, i.e. `X` is utf-8 encoded and this value is interpreted as of `bytes` type and encoded further. Note that the length used in this subsequent encoding is the number of bytes of the utf-8 encoded string, not its number of characters.

- `uint<M>`: `enc(X)` is the big-endian encoding of `X`, padded on the higher-order (left) side with zero-bytes such that the length is 32 bytes.

- `address`: as in the `uint160` case

- `int<M>`: `enc(X)` is the big-endian two's complement encoding of `X`, padded on the higher-order (left) side with `0xff` bytes for negative `X` and with zero-bytes for non-negative `X` such that the length is 32 bytes.

- `bool`: as in the `uint8` case, where `1` is used for `true` and `0` for `false`

- `fixed<M>x<N>`: `enc(X)` is `enc(X * 10**N)` where `X * 10**N` is interpreted as a `int256`.

- `fixed`: as in the `fixed128x18` case

- `ufixed<M>x<N>`: `enc(X)` is `enc(X * 10**N)` where `X * 10**N` is interpreted as a `uint256`.

- `ufixed`: as in the `ufixed128x18` case

- `bytes<M>`: `enc(X)` is the sequence of bytes in `X` padded with trailing zero-bytes to a length of 32 bytes.

Note that for any X, `len(enc(X))` is a multiple of 32.

3.19.7 Function Selector and Argument Encoding

All in all, a call to the function `f` with parameters `a_1, ..., a_n` is encoded as

```
function_selector(f) enc((a_1, ..., a_n))
```

and the return values `v_1, ..., v_k` of `f` are encoded as

```
enc((v_1, ..., v_k))
```

i.e. the values are combined into a tuple and encoded.

3.19.8 Examples

Given the contract:

```solidity
// SPDX-License-Identifier: GPL-3.0
pragma solidity >=0.4.16 <0.8.0;

contract Foo {
    function bar(bytes3[2] memory) public pure {}
    function baz(uint32 x, bool y) public pure returns (bool r) { r = x > 32 || y; }
    function sam(bytes memory, bool, uint[] memory) public pure {}
}
```

Thus for our `Foo` example if we wanted to call `baz` with the parameters `69` and `true`, we would pass 68 bytes total, which can be broken down into:

- `0xcdcd77c0`: the Method ID. This is derived as the first 4 bytes of the Keccak hash of the ASCII form of the signature `baz(uint32,bool)`.

- `0x0000000000000000000000000000000000000000000000000000000000000045`: the first parameter, a uint32 value `69` padded to 32 bytes

- `0x0000000000000000000000000000000000000000000000000000000000000001`: the second parameter - boolean `true`, padded to 32 bytes

In total:

```
0xcdcd77c000000000000000000000000000000000000000000000000000000000000000450000000000000000000000000
```

It returns a single `bool`. If, for example, it were to return `false`, its output would be the single byte array `0x0000000000000000000000000000000000000000000000000000000000000000`, a single bool.

If we wanted to call `bar` with the argument `["abc", "def"]`, we would pass 68 bytes total, broken down into:

- `0xfce353f6`: the Method ID. This is derived from the signature `bar(bytes3[2])`.

- `0x6162630000000000000000000000000000000000000000000000000000000000`: the first part of the first parameter, a `bytes3` value `"abc"` (left-aligned).

- `0x6465660000000000000000000000000000000000000000000000000000000000`: the second part of the first parameter, a `bytes3` value `"def"` (left-aligned).

In total:

```
0xfce353f661626300000000000000000000000000000000000000000000000000000000006465660000000000000000000000000
```

If we wanted to call `sam` with the arguments `"dave"`, `true` and `[1,2,3]`, we would pass 292 bytes total, broken down into:

- `0xa5643bf2`: the Method ID. This is derived from the signature `sam(bytes,bool,uint256[])`. Note that `uint` is replaced with its canonical representation `uint256`.

- `0x0000000000000000000000000000000000000000000000000000000000000060`: the location of the data part of the first parameter (dynamic type), measured in bytes from the start of the arguments block. In this case, `0x60`.

- `0x0000000000000000000000000000000000000000000000000000000000000001`: the second parameter: boolean true.

- `0x00000000000000000000000000000000000000000000000000000000000000a0`: the location of the data part of the third parameter (dynamic type), measured in bytes. In this case, `0xa0`.

- `0x0000000000000000000000000000000000000000000000000000000000000004`: the data part of the first argument, it starts with the length of the byte array in elements, in this case, 4.

- `0x6461766500000000000000000000000000000000000000000000000000000000`: the contents of the first argument: the UTF-8 (equal to ASCII in this case) encoding of `"dave"`, padded on the right to 32 bytes.

- `0x0000000000000000000000000000000000000000000000000000000000000003`: the data part of the third argument, it starts with the length of the array in elements, in this case, 3.

- `0x0000000000000000000000000000000000000000000000000000000000000001`: the first entry of the third parameter.

- `0x0000000000000000000000000000000000000000000000000000000000000002`: the second entry of the third parameter.

- `0x0000000000000000000000000000000000000000000000000000000000000003`: the third entry of the third parameter.

In total:

```
0xa5643bf200000000000000000000000000000000000000000000000000000000000000600000000000000000000000000
```

3.19.9 Use of Dynamic Types

A call to a function with the signature `f(uint,uint32[],bytes10,bytes)` with values (`0x123, [0x456, 0x789], "1234567890", "Hello, world!"`) is encoded in the following way:

We take the first four bytes of `sha3("f(uint256,uint32[],bytes10,bytes)")`, i.e. `0x8be65246`. Then we encode the head parts of all four arguments. For the static types `uint256` and `bytes10`, these are directly the values we want to pass, whereas for the dynamic types `uint32[]` and `bytes`, we use the offset in bytes to the start of their data area, measured from the start of the value encoding (i.e. not counting the first four bytes containing the hash of the function signature). These are:

- `0x0000000000000000000000000000000000000000000000000000000000000123` (`0x123` padded to 32 bytes)

- `0x0000000000000000000000000000000000000000000000000000000000000080` (offset to start of data part of second parameter, 4*32 bytes, exactly the size of the head part)

- `0x3132333435363738393000000000000000000000000000000000000000000000` (`"1234567890"` padded to 32 bytes on the right)

- `0x00000000000000000000000000000000000000000000000000000000000000e0` (offset to start of data part of fourth parameter = offset to start of data part of first dynamic parameter + size of data part of first dynamic parameter = 4*32 + 3*32 (see below))

After this, the data part of the first dynamic argument, `[0x456, 0x789]` follows:

- `0x0000000000000000000000000000000000000000000000000000000000000002` (number of elements of the array, 2)

- `0x0000000000000000000000000000000000000000000000000000000000000456` (first element)

- `0x0000000000000000000000000000000000000000000000000000000000000789` (second element)

Finally, we encode the data part of the second dynamic argument, `"Hello, world!"`:

- `0x000000000000000000000000000000000000000000000000000000000000000d` (number of elements (bytes in this case): 13)

- `0x48656c6c6f2c20776f726c6421000000000000000000000000000000000000000` (`"Hello, world!"` padded to 32 bytes on the right)

All together, the encoding is (newline after function selector and each 32-bytes for clarity):

```
0x8be65246
  0000000000000000000000000000000000000000000000000000000000000123
  0000000000000000000000000000000000000000000000000000000000000080
  3132333435363738393000000000000000000000000000000000000000000000
  00000000000000000000000000000000000000000000000000000000000000e0
  0000000000000000000000000000000000000000000000000000000000000002
  0000000000000000000000000000000000000000000000000000000000000456
  0000000000000000000000000000000000000000000000000000000000000789
  000000000000000000000000000000000000000000000000000000000000000d
  48656c6c6f2c20776f726c6421000000000000000000000000000000000000000
```

Let us apply the same principle to encode the data for a function with a signature `g(uint[][],string[])` with values (`[[1, 2], [3]], ["one", "two", "three"]`) but start from the most atomic parts of the encoding:

First we encode the length and data of the first embedded dynamic array `[1, 2]` of the first root array `[[1, 2], [3]]`:

- `0x0000000000000000000000000000000000000000000000000000000000000002` (number of elements in the first array, 2; the elements themselves are `1` and `2`)

- `0x0000000000000000000000000000000000000000000000000000000000000001` (first element)

- `0x0000000000000000000000000000000000000000000000000000000000000002` (second element)

Then we encode the length and data of the second embedded dynamic array `[3]` of the first root array `[[1, 2], [3]]`:

- `0x0000000000000000000000000000000000000000000000000000000000000001` (number of elements in the second array, 1; the element is `3`)

- `0x0000000000000000000000000000000000000000000000000000000000000003` (first element)

Then we need to find the offsets `a` and `b` for their respective dynamic arrays `[1, 2]` and `[3]`. To calculate the offsets we can take a look at the encoded data of the first root array `[[1, 2], [3]]` enumerating each line in the encoding:

```
0 - a                                                              - offset of [1,␣
↪2]
1 - b                                                              - offset of [3]
2 - 0000000000000000000000000000000000000000000000000000000000000002 - count for [1,␣
↪2]
3 - 0000000000000000000000000000000000000000000000000000000000000001 - encoding of 1
4 - 0000000000000000000000000000000000000000000000000000000000000002 - encoding of 2
5 - 0000000000000000000000000000000000000000000000000000000000000001 - count for [3]
6 - 0000000000000000000000000000000000000000000000000000000000000003 - encoding of 3
```

Offset `a` points to the start of the content of the array `[1, 2]` which is line 2 (64 bytes); thus `a = 0x0000000000000000000000000000000000000000000000000000000000000040`.

Offset `b` points to the start of the content of the array `[3]` which is line 5 (160 bytes); thus `b = 0x00000000000000000000000000000000000000000000000000000000000000a0`.

Then we encode the embedded strings of the second root array:

- `0x0000000000000000000000000000000000000000000000000000000000000003` (number of characters in word `"one"`)

- `0x6f6e650000000000000000000000000000000000000000000000000000000000` (utf8 representation of word `"one"`)

- `0x0000000000000000000000000000000000000000000000000000000000000003` (number of characters in word `"two"`)

- `0x74776f0000000000000000000000000000000000000000000000000000000000` (utf8 representation of word `"two"`)

- `0x0000000000000000000000000000000000000000000000000000000000000005` (number of characters in word `"three"`)

- `0x7468726565000000000000000000000000000000000000000000000000000000` (utf8 representation of word `"three"`)

In parallel to the first root array, since strings are dynamic elements we need to find their offsets `c`, `d` and `e`:

```
0 - c                                                              - offset for "one
↪"
1 - d                                                              - offset for "two
↪"
2 - e                                                              - offset for
↪"three"
3 - 0000000000000000000000000000000000000000000000000000000000000003 - count for "one"
4 - 6f6e650000000000000000000000000000000000000000000000000000000000 - encoding of
↪"one"
5 - 0000000000000000000000000000000000000000000000000000000000000003 - count for "two"
6 - 74776f0000000000000000000000000000000000000000000000000000000000 - encoding of
↪"two"
```

```
 7 - 0000000000000000000000000000000000000000000000000000000000000005 - count for
↪"three"
 8 - 7468726565000000000000000000000000000000000000000000000000000000 - encoding of
↪"three"
```

Offset `c` points to the start of the content of the string `"one"` which is line 3 (96 bytes); thus `c =`
`0x0000000000000000000000000000000000000000000000000000000000000060`.

Offset `d` points to the start of the content of the string `"two"` which is line 5 (160 bytes); thus `d =`
`0x00000000000000000000000000000000000000000000000000000000000000a0`.

Offset `e` points to the start of the content of the string `"three"` which is line 7 (224 bytes); thus `e =`
`0x00000000000000000000000000000000000000000000000000000000000000e0`.

Note that the encodings of the embedded elements of the root arrays are not dependent on each other and have the same encodings for a function with a signature `g(string[],uint[][])`.

Then we encode the length of the first root array:

- `0x0000000000000000000000000000000000000000000000000000000000000002` (number of elements in the first root array, 2; the elements themselves are `[1, 2]` and `[3]`)

Then we encode the length of the second root array:

- `0x0000000000000000000000000000000000000000000000000000000000000003` (number of strings in the second root array, 3; the strings themselves are `"one"`, `"two"` and `"three"`)

Finally we find the offsets `f` and `g` for their respective root dynamic arrays `[[1, 2], [3]]` and `["one",`
`"two", "three"]`, and assemble parts in the correct order:

```
0x2289b18c                                                           - function␣
↪signature
 0 - f                                                               - offset of [[1,
↪ 2], [3]]
 1 - g                                                               - offset of [
↪"one", "two", "three"]
 2 - 0000000000000000000000000000000000000000000000000000000000000002 - count for [[1,
↪ 2], [3]]
 3 - 0000000000000000000000000000000000000000000000000000000000000040 - offset of [1,␣
↪2]
 4 - 00000000000000000000000000000000000000000000000000000000000000a0 - offset of [3]
 5 - 0000000000000000000000000000000000000000000000000000000000000002 - count for [1,␣
↪2]
 6 - 0000000000000000000000000000000000000000000000000000000000000001 - encoding of 1
 7 - 0000000000000000000000000000000000000000000000000000000000000002 - encoding of 2
 8 - 0000000000000000000000000000000000000000000000000000000000000001 - count for [3]
 9 - 0000000000000000000000000000000000000000000000000000000000000003 - encoding of 3
10 - 0000000000000000000000000000000000000000000000000000000000000003 - count for [
↪"one", "two", "three"]
11 - 0000000000000000000000000000000000000000000000000000000000000060 - offset for
↪"one"
12 - 00000000000000000000000000000000000000000000000000000000000000a0 - offset for
↪"two"
13 - 00000000000000000000000000000000000000000000000000000000000000e0 - offset for
↪"three"
14 - 0000000000000000000000000000000000000000000000000000000000000003 - count for "one
↪"
15 - 6f6e65000000000000000000000000000000000000000000000000000000000000 - encoding of
↪"one"
```

```
16 - 0000000000000000000000000000000000000000000000000000000000000003 - count for "two
↪"
17 - 74776f0000000000000000000000000000000000000000000000000000000000 - encoding of
↪"two"
18 - 0000000000000000000000000000000000000000000000000000000000000005 - count for
↪"three"
19 - 7468726565000000000000000000000000000000000000000000000000000000 - encoding of
↪"three"
```

Offset `f` points to the start of the content of the array `[[1, 2], [3]]` which is line 2 (64 bytes); thus `f = 0x0000000000000000000000000000000000000000000000000000000000000040`.

Offset `g` points to the start of the content of the array `["one", "two", "three"]` which is line 10 (320 bytes); thus `g = 0x0000000000000000000000000000000000000000000000000000000000000140`.

3.19.10 Events

Events are an abstraction of the Ethereum logging/event-watching protocol. Log entries provide the contract's address, a series of up to four topics and some arbitrary length binary data. Events leverage the existing function ABI in order to interpret this (together with an interface spec) as a properly typed structure.

Given an event name and series of event parameters, we split them into two sub-series: those which are indexed and those which are not. Those which are indexed, which may number up to 3, are used alongside the Keccak hash of the event signature to form the topics of the log entry. Those which are not indexed form the byte array of the event.

In effect, a log entry using this ABI is described as:

- `address`: the address of the contract (intrinsically provided by Ethereum);
- `topics[0]`: `keccak(EVENT_NAME+"("+EVENT_ARGS.map(canonical_type_of).join(", ")+")")` (`canonical_type_of` is a function that simply returns the canonical type of a given argument, e.g. for `uint indexed foo`, it would return `uint256`). If the event is declared as `anonymous` the `topics[0]` is not generated;
- `topics[n]`: `abi_encode(EVENT_INDEXED_ARGS[n - 1])` (`EVENT_INDEXED_ARGS` is the series of `EVENT_ARGS` that are indexed);
- `data`: ABI encoding of `EVENT_NON_INDEXED_ARGS` (`EVENT_NON_INDEXED_ARGS` is the series of `EVENT_ARGS` that are not indexed, `abi_encode` is the ABI encoding function used for returning a series of typed values from a function, as described above).

For all types of length at most 32 bytes, the `EVENT_INDEXED_ARGS` array contains the value directly, padded or sign-extended (for signed integers) to 32 bytes, just as for regular ABI encoding. However, for all "complex" types or types of dynamic length, including all arrays, `string`, `bytes` and structs, `EVENT_INDEXED_ARGS` will contain the *Keccak hash* of a special in-place encoded value (see *Encoding of Indexed Event Parameters*), rather than the encoded value directly. This allows applications to efficiently query for values of dynamic-length types (by setting the hash of the encoded value as the topic), but leaves applications unable to decode indexed values they have not queried for. For dynamic-length types, application developers face a trade-off between fast search for predetermined values (if the argument is indexed) and legibility of arbitrary values (which requires that the arguments not be indexed). Developers may overcome this tradeoff and achieve both efficient search and arbitrary legibility by defining events with two arguments — one indexed, one not — intended to hold the same value.

3.19.11 JSON

The JSON format for a contract's interface is given by an array of function and/or event descriptions. A function description is a JSON object with the fields:

- type: `"function"`, `"constructor"`, `"receive"` (the *"receive Ether" function*) or `"fallback"` (the *"default" function*);

- name: the name of the function;

- inputs: an array of objects, each of which contains:

    - name: the name of the parameter.

    - type: the canonical type of the parameter (more below).

    - components: used for tuple types (more below).

- outputs: an array of objects similar to inputs.

- stateMutability: a string with one of the following values: pure (*specified to not read blockchain state*), view (*specified to not modify the blockchain state*), nonpayable (function does not accept Ether - the default) and payable (function accepts Ether).

Constructor and fallback function never have name or outputs. Fallback function doesn't have inputs either.

Note: Sending non-zero Ether to non-payable function will revert the transaction.

Note: The state mutability nonpayable is reflected in Solidity by not specifying a state mutability modifier at all.

An event description is a JSON object with fairly similar fields:

- type: always `"event"`

- name: the name of the event.

- inputs: an array of objects, each of which contains:

    - name: the name of the parameter.

    - type: the canonical type of the parameter (more below).

    - components: used for tuple types (more below).

    - indexed: `true` if the field is part of the log's topics, `false` if it one of the log's data segment.

- anonymous: `true` if the event was declared as anonymous.

For example,

```solidity
// SPDX-License-Identifier: GPL-3.0
pragma solidity ^0.7.0;


contract Test {
    constructor() { b = hex"12345678901234567890123456789012"; }
    event Event(uint indexed a, bytes32 b);
    event Event2(uint indexed a, bytes32 b);
    function foo(uint a) public { emit Event(a, b); }
    bytes32 b;
}
```

would result in the JSON:

```
[{
"type":"event",
"inputs": [{"name":"a","type":"uint256","indexed":true},{"name":"b","type":"bytes32",
↪"indexed":false}],
"name":"Event"
}, {
"type":"event",
"inputs": [{"name":"a","type":"uint256","indexed":true},{"name":"b","type":"bytes32",
↪"indexed":false}],
"name":"Event2"
}, {
"type":"function",
"inputs": [{"name":"a","type":"uint256"}],
"name":"foo",
"outputs": []
}]
```

Handling tuple types

Despite that names are intentionally not part of the ABI encoding they do make a lot of sense to be included in the JSON to enable displaying it to the end user. The structure is nested in the following way:

An object with members `name`, `type` and potentially `components` describes a typed variable. The canonical type is determined until a tuple type is reached and the string description up to that point is stored in `type` prefix with the word `tuple`, i.e. it will be `tuple` followed by a sequence of `[]` and `[k]` with integers k. The components of the tuple are then stored in the member `components`, which is of array type and has the same structure as the top-level object except that `indexed` is not allowed there.

As an example, the code

```solidity
// SPDX-License-Identifier: GPL-3.0
pragma solidity >0.7.4;
pragma abicoder v2;

contract Test {
    struct S { uint a; uint[] b; T[] c; }
    struct T { uint x; uint y; }
    function f(S memory, T memory, uint) public pure {}
    function g() public pure returns (S memory, T memory, uint) {}
}
```

would result in the JSON:

```json
[
  {
    "name": "f",
    "type": "function",
    "inputs": [
      {
        "name": "s",
        "type": "tuple",
        "components": [
          {
            "name": "a",
            "type": "uint256"
          },
```

(continues on next page)

```
        {
          "name": "b",
          "type": "uint256[]"
        },
        {
          "name": "c",
          "type": "tuple[]",
          "components": [
            {
              "name": "x",
              "type": "uint256"
            },
            {
              "name": "y",
              "type": "uint256"
            }
          ]
        }
      ]
    },
    {
      "name": "t",
      "type": "tuple",
      "components": [
        {
          "name": "x",
          "type": "uint256"
        },
        {
          "name": "y",
          "type": "uint256"
        }
      ]
    },
    {
      "name": "a",
      "type": "uint256"
    }
  ],
  "outputs": []
  }
]
```

3.19.12 Strict Encoding Mode

Strict encoding mode is the mode that leads to exactly the same encoding as defined in the formal specification above. This means offsets have to be as small as possible while still not creating overlaps in the data areas and thus no gaps are allowed.

Usually, ABI decoders are written in a straightforward way just following offset pointers, but some decoders might enforce strict mode. The Solidity ABI decoder currently does not enforce strict mode, but the encoder always creates data in strict mode.

3.19.13 Non-standard Packed Mode

Through `abi.encodePacked()`, Solidity supports a non-standard packed mode where:

- types shorter than 32 bytes are neither zero padded nor sign extended and

- dynamic types are encoded in-place and without the length.

- array elements are padded, but still encoded in-place

Furthermore, structs as well as nested arrays are not supported.

As an example, the encoding of `int16(-1)`, `bytes1(0x42)`, `uint16(0x03)`, `string("Hello, world!")` results in:

```
0xffff42000348656c6c6f2c20776f726c6421
  ^^^^                                  int16(-1)
    ^^                                  bytes1(0x42)
     ^^^^                               uint16(0x03)
         ^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^ string("Hello, world!") without a length field
```

More specifically:

- During the encoding, everything is encoded in-place. This means that there is no distinction between head and tail, as in the ABI encoding, and the length of an array is not encoded.

- The direct arguments of `abi.encodePacked` are encoded without padding, as long as they are not arrays (or `string` or `bytes`).

- The encoding of an array is the concatenation of the encoding of its elements **with** padding.

- Dynamically-sized types like `string`, `bytes` or `uint[]` are encoded without their length field.

- The encoding of `string` or `bytes` does not apply padding at the end unless it is part of an array or struct (then it is padded to a multiple of 32 bytes).

In general, the encoding is ambiguous as soon as there are two dynamically-sized elements, because of the missing length field.

If padding is needed, explicit type conversions can be used: `abi.encodePacked(uint16(0x12)) == hex"0012"`.

Since packed encoding is not used when calling functions, there is no special support for prepending a function selector. Since the encoding is ambiguous, there is no decoding function.

> **Warning:** If you use `keccak256(abi.encodePacked(a, b))` and both `a` and `b` are dynamic types, it is easy to craft collisions in the hash value by moving parts of `a` into `b` and vice-versa. More specifically, `abi.encodePacked("a", "bc") == abi.encodePacked("ab", "c")`. If you use `abi.encodePacked` for signatures, authentication or data integrity, make sure to always use the same types and check that at most one of them is dynamic. Unless there is a compelling reason, `abi.encode` should be preferred.

3.19.14 Encoding of Indexed Event Parameters

Indexed event parameters that are not value types, i.e. arrays and structs are not stored directly but instead a keccak256-hash of an encoding is stored. This encoding is defined as follows:

- the encoding of a `bytes` and `string` value is just the string contents without any padding or length prefix.

- the encoding of a struct is the concatenation of the encoding of its members, always padded to a multiple of 32 bytes (even `bytes` and `string`).

- the encoding of an array (both dynamically- and statically-sized) is the concatenation of the encoding of its elements, always padded to a multiple of 32 bytes (even `bytes` and `string`) and without any length prefix

In the above, as usual, a negative number is padded by sign extension and not zero padded. `bytesNN` types are padded on the right while `uintNN` / `intNN` are padded on the left.

> **Warning:** The encoding of a struct is ambiguous if it contains more than one dynamically-sized array. Because of that, always re-check the event data and do not rely on the search result based on the indexed parameters alone.

3.20 Solidity v0.5.0 Breaking Changes

This section highlights the main breaking changes introduced in Solidity version 0.5.0, along with the reasoning behind the changes and how to update affected code. For the full list check the release changelog.

> **Note:** Contracts compiled with Solidity v0.5.0 can still interface with contracts and even libraries compiled with older versions without recompiling or redeploying them. Changing the interfaces to include data locations and visibility and mutability specifiers suffices. See the *Interoperability With Older Contracts* section below.

3.20.1 Semantic Only Changes

This section lists the changes that are semantic-only, thus potentially hiding new and different behavior in existing code.

- Signed right shift now uses proper arithmetic shift, i.e. rounding towards negative infinity, instead of rounding towards zero. Signed and unsigned shift will have dedicated opcodes in Constantinople, and are emulated by Solidity for the moment.

- The `continue` statement in a `do...while` loop now jumps to the condition, which is the common behavior in such cases. It used to jump to the loop body. Thus, if the condition is false, the loop terminates.

- The functions `.call()`, `.delegatecall()` and `.staticcall()` do not pad anymore when given a single `bytes` parameter.

- Pure and view functions are now called using the opcode `STATICCALL` instead of `CALL` if the EVM version is Byzantium or later. This disallows state changes on the EVM level.

- The ABI encoder now properly pads byte arrays and strings from calldata (`msg.data` and external function parameters) when used in external function calls and in `abi.encode`. For unpadded encoding, use `abi.encodePacked`.

- The ABI decoder reverts in the beginning of functions and in `abi.decode()` if passed calldata is too short or points out of bounds. Note that dirty higher order bits are still simply ignored.

- Forward all available gas with external function calls starting from Tangerine Whistle.

3.20.2 Semantic and Syntactic Changes

This section highlights changes that affect syntax and semantics.

- The functions `.call()`, `.delegatecall()`, `staticcall()`, `keccak256()`, `sha256()` and `ripemd160()` now accept only a single `bytes` argument. Moreover, the argument is not padded. This was changed to make more explicit and clear how the arguments are concatenated. Change every `.call()` (and family) to a `.call("")` and every `.call(signature, a, b, c)` to use `.call(abi.encodeWithSignature(signature, a, b, c))` (the last one only works for value types). Change every `keccak256(a, b, c)` to `keccak256(abi.encodePacked(a, b, c))`. Even though it is not a breaking change, it is suggested that developers change `x.call(bytes4(keccak256("f(uint256)")), a, b)` to `x.call(abi.encodeWithSignature("f(uint256)", a, b))`.

- Functions `.call()`, `.delegatecall()` and `.staticcall()` now return `(bool, bytes memory)` to provide access to the return data. Change `bool success = otherContract.call("f")` to `(bool success, bytes memory data) = otherContract.call("f")`.

- Solidity now implements C99-style scoping rules for function local variables, that is, variables can only be used after they have been declared and only in the same or nested scopes. Variables declared in the initialization block of a `for` loop are valid at any point inside the loop.

3.20.3 Explicitness Requirements

This section lists changes where the code now needs to be more explicit. For most of the topics the compiler will provide suggestions.

- Explicit function visibility is now mandatory. Add `public` to every function and constructor, and `external` to every fallback or interface function that does not specify its visibility already.

- Explicit data location for all variables of struct, array or mapping types is now mandatory. This is also applied to function parameters and return variables. For example, change `uint[] x = m_x` to `uint[] storage x = m_x`, and `function f(uint[][] x)` to `function f(uint[][] memory x)` where `memory` is the data location and might be replaced by `storage` or `calldata` accordingly. Note that `external` functions require parameters with a data location of `calldata`.

- Contract types do not include `address` members anymore in order to separate the namespaces. Therefore, it is now necessary to explicitly convert values of contract type to addresses before using an `address` member. Example: if `c` is a contract, change `c.transfer(...)` to `address(c).transfer(...)`, and `c.balance` to `address(c).balance`.

- Explicit conversions between unrelated contract types are now disallowed. You can only convert from a contract type to one of its base or ancestor types. If you are sure that a contract is compatible with the contract type you want to convert to, although it does not inherit from it, you can work around this by converting to `address` first. Example: if `A` and `B` are contract types, `B` does not inherit from `A` and `b` is a contract of type `B`, you can still convert `b` to type `A` using `A(address(b))`. Note that you still need to watch out for matching payable fallback functions, as explained below.

- The `address` type was split into `address` and `address payable`, where only `address payable` provides the `transfer` function. An `address payable` can be directly converted to an `address`, but the other way around is not allowed. Converting `address` to `address payable` is possible via conversion through `uint160`. If `c` is a contract, `address(c)` results in `address payable` only if `c` has a payable fallback function. If you use the *withdraw pattern*, you most likely do not have to change your code because `transfer` is only used on `msg.sender` instead of stored addresses and `msg.sender` is an `address payable`.

- Conversions between `bytesX` and `uintY` of different size are now disallowed due to `bytesX` padding on the right and `uintY` padding on the left which may cause unexpected conversion results. The size must now be adjusted within the type before the conversion. For example, you can convert a `bytes4` (4 bytes) to a `uint64` (8 bytes) by first converting the `bytes4` variable to `bytes8` and then to `uint64`. You get the opposite padding when converting through `uint32`.

- Using `msg.value` in non-payable functions (or introducing it via a modifier) is disallowed as a security feature. Turn the function into `payable` or create a new internal function for the program logic that uses `msg.value`.

- For clarity reasons, the command line interface now requires – if the standard input is used as source.

3.20.4 Deprecated Elements

This section lists changes that deprecate prior features or syntax. Note that many of these changes were already enabled in the experimental mode `v0.5.0`.

Command Line and JSON Interfaces

- The command line option `--formal` (used to generate Why3 output for further formal verification) was deprecated and is now removed. A new formal verification module, the SMTChecker, is enabled via `pragma experimental SMTChecker;`.

- The command line option `--julia` was renamed to `--yul` due to the renaming of the intermediate language `Julia` to `Yul`.

- The `--clone-bin` and `--combined-json clone-bin` command line options were removed.

- Remappings with empty prefix are disallowed.

- The JSON AST fields `constant` and `payable` were removed. The information is now present in the `stateMutability` field.

- The JSON AST field `isConstructor` of the `FunctionDefinition` node was replaced by a field called `kind` which can have the value `"constructor"`, `"fallback"` or `"function"`.

- In unlinked binary hex files, library address placeholders are now the first 36 hex characters of the keccak256 hash of the fully qualified library name, surrounded by `$...$`. Previously, just the fully qualified library name was used. This reduces the chances of collisions, especially when long paths are used. Binary files now also contain a list of mappings from these placeholders to the fully qualified names.

Constructors

- Constructors must now be defined using the `constructor` keyword.

- Calling base constructors without parentheses is now disallowed.

- Specifying base constructor arguments multiple times in the same inheritance hierarchy is now disallowed.

- Calling a constructor with arguments but with wrong argument count is now disallowed. If you only want to specify an inheritance relation without giving arguments, do not provide parentheses at all.

Functions

- Function `callcode` is now disallowed (in favor of `delegatecall`). It is still possible to use it via inline assembly.

- `suicide` is now disallowed (in favor of `selfdestruct`).

- `sha3` is now disallowed (in favor of `keccak256`).

- `throw` is now disallowed (in favor of `revert`, `require` and `assert`).

Conversions

- Explicit and implicit conversions from decimal literals to `bytesXX` types is now disallowed.

- Explicit and implicit conversions from hex literals to `bytesXX` types of different size is now disallowed.

Literals and Suffixes

- The unit denomination `years` is now disallowed due to complications and confusions about leap years.

- Trailing dots that are not followed by a number are now disallowed.

- Combining hex numbers with unit denominations (e.g. `0x1e wei`) is now disallowed.

- The prefix `0X` for hex numbers is disallowed, only `0x` is possible.

Variables

- Declaring empty structs is now disallowed for clarity.

- The `var` keyword is now disallowed to favor explicitness.

- Assignments between tuples with different number of components is now disallowed.

- Values for constants that are not compile-time constants are disallowed.

- Multi-variable declarations with mismatching number of values are now disallowed.

- Uninitialized storage variables are now disallowed.

- Empty tuple components are now disallowed.

- Detecting cyclic dependencies in variables and structs is limited in recursion to 256.

- Fixed-size arrays with a length of zero are now disallowed.

Syntax

- Using `constant` as function state mutability modifier is now disallowed.

- Boolean expressions cannot use arithmetic operations.

- The unary + operator is now disallowed.

- Literals cannot anymore be used with `abi.encodePacked` without prior conversion to an explicit type.

- Empty return statements for functions with one or more return values are now disallowed.

- The "loose assembly" syntax is now disallowed entirely, that is, jump labels, jumps and non-functional instructions cannot be used anymore. Use the new `while`, `switch` and `if` constructs instead.

- Functions without implementation cannot use modifiers anymore.

- Function types with named return values are now disallowed.

- Single statement variable declarations inside if/while/for bodies that are not blocks are now disallowed.

- New keywords: `calldata` and `constructor`.

- New reserved keywords: `alias`, `apply`, `auto`, `copyof`, `define`, `immutable`, `implements`, `macro`, `mutable`, `override`, `partial`, `promise`, `reference`, `sealed`, `sizeof`, `supports`, `typedef` and `unchecked`.

3.20.5 Interoperability With Older Contracts

It is still possible to interface with contracts written for Solidity versions prior to v0.5.0 (or the other way around) by defining interfaces for them. Consider you have the following pre-0.5.0 contract already deployed:

```
// SPDX-License-Identifier: GPL-3.0
pragma solidity ^0.4.25;
// This will report a warning until version 0.4.25 of the compiler
// This will not compile after 0.5.0
contract OldContract {
    function someOldFunction(uint8 a) {
        //...
    }
    function anotherOldFunction() constant returns (bool) {
        //...
    }
    // ...
}
```

This will no longer compile with Solidity v0.5.0. However, you can define a compatible interface for it:

```
// SPDX-License-Identifier: GPL-3.0
pragma solidity >=0.5.0 <0.8.0;
interface OldContract {
    function someOldFunction(uint8 a) external;
    function anotherOldFunction() external returns (bool);
}
```

Note that we did not declare `anotherOldFunction` to be `view`, despite it being declared `constant` in the original contract. This is due to the fact that starting with Solidity v0.5.0 `staticcall` is used to call `view` functions. Prior to v0.5.0 the `constant` keyword was not enforced, so calling a function declared `constant` with `staticcall` may still revert, since the `constant` function may still attempt to modify storage. Consequently, when defining an interface for older contracts, you should only use `view` in place of `constant` in case you are absolutely sure that the function will work with `staticcall`.

Given the interface defined above, you can now easily use the already deployed pre-0.5.0 contract:

```
// SPDX-License-Identifier: GPL-3.0
pragma solidity >=0.5.0 <0.8.0;

interface OldContract {
    function someOldFunction(uint8 a) external;
    function anotherOldFunction() external returns (bool);
}

contract NewContract {
    function doSomething(OldContract a) public returns (bool) {
        a.someOldFunction(0x42);
        return a.anotherOldFunction();
    }
}
```

Similarly, pre-0.5.0 libraries can be used by defining the functions of the library without implementation and supplying the address of the pre-0.5.0 library during linking (see *Using the Commandline Compiler* for how to use the commandline compiler for linking):

```solidity
// This will not compile after 0.6.0
// SPDX-License-Identifier: GPL-3.0
pragma solidity >=0.5.0 <0.5.99;

library OldLibrary {
    function someFunction(uint8 a) public returns(bool);
}

contract NewContract {
    function f(uint8 a) public returns (bool) {
        return OldLibrary.someFunction(a);
    }
}
```

3.20.6 Example

The following example shows a contract and its updated version for Solidity v0.5.0 with some of the changes listed in this section.

Old version:

```solidity
// SPDX-License-Identifier: GPL-3.0
pragma solidity ^0.4.25;
// This will not compile after 0.5.0

contract OtherContract {
    uint x;
    function f(uint y) external {
        x = y;
    }
    function() payable external {}
}

contract Old {
    OtherContract other;
    uint myNumber;

    // Function mutability not provided, not an error.
    function someInteger() internal returns (uint) { return 2; }

    // Function visibility not provided, not an error.
    // Function mutability not provided, not an error.
    function f(uint x) returns (bytes) {
        // Var is fine in this version.
        var z = someInteger();
        x += z;
        // Throw is fine in this version.
        if (x > 100)
            throw;
        bytes memory b = new bytes(x);
        y = -3 >> 1;
        // y == -1 (wrong, should be -2)
        do {
            x += 1;
            if (x > 10) continue;
            // 'Continue' causes an infinite loop.
```

(continues on next page)

```solidity
        } while (x < 11);
        // Call returns only a Bool.
        bool success = address(other).call("f");
        if (!success)
            revert();
        else {
            // Local variables could be declared after their use.
            int y;
        }
        return b;
    }

    // No need for an explicit data location for 'arr'
    function g(uint[] arr, bytes8 x, OtherContract otherContract) public {
        otherContract.transfer(1 ether);

        // Since uint32 (4 bytes) is smaller than bytes8 (8 bytes),
        // the first 4 bytes of x will be lost. This might lead to
        // unexpected behavior since bytesX are right padded.
        uint32 y = uint32(x);
        myNumber += y + msg.value;
    }
}
```

New version:

```solidity
// SPDX-License-Identifier: GPL-3.0
pragma solidity >=0.5.0 <0.5.99;
// This will not compile after 0.6.0

contract OtherContract {
    uint x;
    function f(uint y) external {
        x = y;
    }
    function() payable external {}
}

contract New {
    OtherContract other;
    uint myNumber;

    // Function mutability must be specified.
    function someInteger() internal pure returns (uint) { return 2; }

    // Function visibility must be specified.
    // Function mutability must be specified.
    function f(uint x) public returns (bytes memory) {
        // The type must now be explicitly given.
        uint z = someInteger();
        x += z;
        // Throw is now disallowed.
        require(x <= 100);
        int y = -3 >> 1;
        require(y == -2);
        do {
            x += 1;
```

```solidity
            if (x > 10) continue;
                // 'Continue' jumps to the condition below.
        } while (x < 11);

        // Call returns (bool, bytes).
        // Data location must be specified.
        (bool success, bytes memory data) = address(other).call("f");
        if (!success)
            revert();
        return data;
    }

    using address_make_payable for address;
    // Data location for 'arr' must be specified
    function g(uint[] memory /* arr */, bytes8 x, OtherContract otherContract,␣
→address unknownContract) public payable {
        // 'otherContract.transfer' is not provided.
        // Since the code of 'OtherContract' is known and has the fallback
        // function, address(otherContract) has type 'address payable'.
        address(otherContract).transfer(1 ether);

        // 'unknownContract.transfer' is not provided.
        // 'address(unknownContract).transfer' is not provided
        // since 'address(unknownContract)' is not 'address payable'.
        // If the function takes an 'address' which you want to send
        // funds to, you can convert it to 'address payable' via 'uint160'.
        // Note: This is not recommended and the explicit type
        // 'address payable' should be used whenever possible.
        // To increase clarity, we suggest the use of a library for
        // the conversion (provided after the contract in this example).
        address payable addr = unknownContract.make_payable();
        require(addr.send(1 ether));

        // Since uint32 (4 bytes) is smaller than bytes8 (8 bytes),
        // the conversion is not allowed.
        // We need to convert to a common size first:
        bytes4 x4 = bytes4(x); // Padding happens on the right
        uint32 y = uint32(x4); // Conversion is consistent
        // 'msg.value' cannot be used in a 'non-payable' function.
        // We need to make the function payable
        myNumber += y + msg.value;
    }
}

// We can define a library for explicitly converting ``address``
// to ``address payable`` as a workaround.
library address_make_payable {
    function make_payable(address x) internal pure returns (address payable) {
        return address(uint160(x));
    }
}
```

```solidity
            if (x > 10) continue;
                // 'Continue' jumps to the condition below.
        } while (x < 11);

        // Call returns (bool, bytes).
        // Data location must be specified.
        (bool success, bytes memory data) = address(other).call("f");
        if (!success)
            revert();
        return data;
    }

    using address_make_payable for address;
    // Data location for 'arr' must be specified
    function g(uint[] memory /* arr */, bytes8 x, OtherContract otherContract,␣
→address unknownContract) public payable {
        // 'otherContract.transfer' is not provided.
        // Since the code of 'OtherContract' is known and has the fallback
        // function, address(otherContract) has type 'address payable'.
        address(otherContract).transfer(1 ether);

        // 'unknownContract.transfer' is not provided.
        // 'address(unknownContract).transfer' is not provided
        // since 'address(unknownContract)' is not 'address payable'.
        // If the function takes an 'address' which you want to send
        // funds to, you can convert it to 'address payable' via 'uint160'.
        // Note: This is not recommended and the explicit type
        // 'address payable' should be used whenever possible.
        // To increase clarity, we suggest the use of a library for
        // the conversion (provided after the contract in this example).
        address payable addr = unknownContract.make_payable();
        require(addr.send(1 ether));

        // Since uint32 (4 bytes) is smaller than bytes8 (8 bytes),
        // the conversion is not allowed.
        // We need to convert to a common size first:
        bytes4 x4 = bytes4(x); // Padding happens on the right
        uint32 y = uint32(x4); // Conversion is consistent
        // 'msg.value' cannot be used in a 'non-payable' function.
        // We need to make the function payable
        myNumber += y + msg.value;
    }
}

// We can define a library for explicitly converting ``address``
// to ``address payable`` as a workaround.
library address_make_payable {
    function make_payable(address x) internal pure returns (address payable) {
        return address(uint160(x));
    }
}
```

3.21 Solidity v0.6.0 Breaking Changes

This section highlights the main breaking changes introduced in Solidity version 0.6.0, along with the reasoning behind the changes and how to update affected code. For the full list check the release changelog.

3.21.1 Changes the Compiler Might not Warn About

This section lists changes where the behaviour of your code might change without the compiler telling you about it.

- The resulting type of an exponentiation is the type of the base. It used to be the smallest type that can hold both the type of the base and the type of the exponent, as with symmetric operations. Additionally, signed types are allowed for the base of the exponentiation.

3.21.2 Explicitness Requirements

This section lists changes where the code now needs to be more explicit, but the semantics do not change. For most of the topics the compiler will provide suggestions.

- Functions can now only be overridden when they are either marked with the `virtual` keyword or defined in an interface. Functions without implementation outside an interface have to be marked `virtual`. When overriding a function or modifier, the new keyword `override` must be used. When overriding a function or modifier defined in multiple parallel bases, all bases must be listed in parentheses after the keyword like so: `override(Base1, Base2)`.

- Member-access to `length` of arrays is now always read-only, even for storage arrays. It is no longer possible to resize storage arrays by assigning a new value to their length. Use `push()`, `push(value)` or `pop()` instead, or assign a full array, which will of course overwrite the existing content. The reason behind this is to prevent storage collisions of gigantic storage arrays.

- The new keyword `abstract` can be used to mark contracts as abstract. It has to be used if a contract does not implement all its functions. Abstract contracts cannot be created using the `new` operator, and it is not possible to generate bytecode for them during compilation.

- Libraries have to implement all their functions, not only the internal ones.

- The names of variables declared in inline assembly may no longer end in `_slot` or `_offset`.

- Variable declarations in inline assembly may no longer shadow any declaration outside the inline assembly block. If the name contains a dot, its prefix up to the dot may not conflict with any declaration outside the inline assembly block.

- State variable shadowing is now disallowed. A derived contract can only declare a state variable x, if there is no visible state variable with the same name in any of its bases.

3.21.3 Semantic and Syntactic Changes

This section lists changes where you have to modify your code and it does something else afterwards.

- Conversions from external function types to `address` are now disallowed. Instead external function types have a member called `address`, similar to the existing `selector` member.

- The function `push(value)` for dynamic storage arrays does not return the new length anymore (it returns nothing).

- The unnamed function commonly referred to as "fallback function" was split up into a new fallback function that is defined using the `fallback` keyword and a receive ether function defined using the `receive` keyword.

– If present, the receive ether function is called whenever the call data is empty (whether or not ether is received). This function is implicitly `payable`.

– The new fallback function is called when no other function matches (if the receive ether function does not exist then this includes calls with empty call data). You can make this function `payable` or not. If it is not `payable` then transactions not matching any other function which send value will revert. You should only need to implement the new fallback function if you are following an upgrade or proxy pattern.

3.21.4 New Features

This section lists things that were not possible prior to Solidity 0.6.0 or were more difficult to achieve.

- The *try/catch statement* allows you to react on failed external calls.

- `struct` and `enum` types can be declared at file level.

- Array slices can be used for calldata arrays, for example `abi.decode(msg.data[4:], (uint, uint))` is a low-level way to decode the function call payload.

- Natspec supports multiple return parameters in developer documentation, enforcing the same naming check as `@param`.

- Yul and Inline Assembly have a new statement called `leave` that exits the current function.

- Conversions from `address` to `address payable` are now possible via `payable(x)`, where `x` must be of type `address`.

3.21.5 Interface Changes

This section lists changes that are unrelated to the language itself, but that have an effect on the interfaces of the compiler. These may change the way how you use the compiler on the command line, how you use its programmable interface, or how you analyze the output produced by it.

New Error Reporter

A new error reporter was introduced, which aims at producing more accessible error messages on the command line. It is enabled by default, but passing `--old-reporter` falls back to the the deprecated old error reporter.

Metadata Hash Options

The compiler now appends the IPFS hash of the metadata file to the end of the bytecode by default (for details, see documentation on *contract metadata*). Before 0.6.0, the compiler appended the Swarm hash by default, and in order to still support this behaviour, the new command line option `--metadata-hash` was introduced. It allows you to select the hash to be produced and appended, by passing either `ipfs` or `swarm` as value to the `--metadata-hash` command line option. Passing the value `none` completely removes the hash.

These changes can also be used via the *Standard JSON Interface* and effect the metadata JSON generated by the compiler.

The recommended way to read the metadata is to read the last two bytes to determine the length of the CBOR encoding and perform a proper decoding on that data block as explained in the *metadata section*.

Yul Optimizer

Together with the legacy bytecode optimizer, the *Yul* optimizer is now enabled by default when you call the compiler with `--optimize`. It can be disabled by calling the compiler with `--no-optimize-yul`. This mostly affects code that uses ABI coder v2.

C API Changes

The client code that uses the C API of `libsolc` is now in control of the memory used by the compiler. To make this change consistent, `solidity_free` was renamed to `solidity_reset`, the functions `solidity_alloc` and `solidity_free` were added and `solidity_compile` now returns a string that must be explicitly freed via `solidity_free()`.

3.21.6 How to update your code

This section gives detailed instructions on how to update prior code for every breaking change.

- Change `address(f)` to `f.address` for `f` being of external function type.

- Replace `function () external [payable] { ... }` by either `receive() external payable { ... }`, `fallback() external [payable] { ... }` or both. Prefer using a `receive` function only, whenever possible.

- Change `uint length = array.push(value)` to `array.push(value);`. The new length can be accessed via `array.length`.

- Change `array.length++` to `array.push()` to increase, and use `pop()` to decrease the length of a storage array.

- For every named return parameter in a function's `@dev` documentation define a `@return` entry which contains the parameter's name as the first word. E.g. if you have function `f()` defined like `function f() public returns (uint value)` and a `@dev` annotating it, document its return parameters like so: `@return value The return value..` You can mix named and un-named return parameters documentation so long as the notices are in the order they appear in the tuple return type.

- Choose unique identifiers for variable declarations in inline assembly that do not conflict with declarations outside the inline assembly block.

- Add `virtual` to every non-interface function you intend to override. Add `virtual` to all functions without implementation outside interfaces. For single inheritance, add `override` to every overriding function. For multiple inheritance, add `override(A, B, ..)`, where you list all contracts that define the overridden function in the parentheses. When multiple bases define the same function, the inheriting contract must override all conflicting functions.

3.22 Solidity v0.7.0 Breaking Changes

This section highlights the main breaking changes introduced in Solidity version 0.7.0, along with the reasoning behind the changes and how to update affected code. For the full list check the release changelog.

3.22.1 Silent Changes of the Semantics

- Exponentiation and shifts of literals by non-literals (e.g. `1 << x` or `2 ** x`) will always use either the type `uint256` (for non-negative literals) or `int256` (for negative literals) to perform the operation. Previously, the

operation was performed in the type of the shift amount / the exponent which can be misleading.

3.22.2 Changes to the Syntax

- In external function and contract creation calls, Ether and gas is now specified using a new syntax: `x. f{gas: 10000, value: 2 ether}(arg1, arg2)`. The old syntax – `x.f.gas(10000). value(2 ether)(arg1, arg2)` – will cause an error.

- The global variable `now` is deprecated, `block.timestamp` should be used instead. The single identifier `now` is too generic for a global variable and could give the impression that it changes during transaction processing, whereas `block.timestamp` correctly reflects the fact that it is just a property of the block.

- NatSpec comments on variables are only allowed for public state variables and not for local or internal variables.

- The token `gwei` is a keyword now (used to specify, e.g. `2 gwei` as a number) and cannot be used as an identifier.

- String literals now can only contain printable ASCII characters and this also includes a variety of escape sequences, such as hexadecimal (`\xff`) and unicode escapes (`\u20ac`).

- Unicode string literals are supported now to accommodate valid UTF-8 sequences. They are identified with the `unicode` prefix: `unicode"Hello "`.

- State Mutability: The state mutability of functions can now be restricted during inheritance: Functions with default state mutability can be overridden by `pure` and `view` functions while `view` functions can be overridden by `pure` functions. At the same time, public state variables are considered `view` and even `pure` if they are constants.

Inline Assembly

- Disallow `.` in user-defined function and variable names in inline assembly. It is still valid if you use Solidity in Yul-only mode.

- Slot and offset of storage pointer variable `x` are accessed via `x.slot` and `x.offset` instead of `x_slot` and `x_offset`.

3.22.3 Removal of Unused or Unsafe Features

Mappings outside Storage

- If a struct or array contains a mapping, it can only be used in storage. Previously, mapping members were silently skipped in memory, which is confusing and error-prone.

- Assignments to structs or arrays in storage does not work if they contain mappings. Previously, mappings were silently skipped during the copy operation, which is misleading and error-prone.

Functions and Events

- Visibility (`public` / `external`) is not needed for constructors anymore: To prevent a contract from being created, it can be marked `abstract`. This makes the visibility concept for constructors obsolete.

- Type Checker: Disallow `virtual` for library functions: Since libraries cannot be inherited from, library functions should not be virtual.

- Multiple events with the same name and parameter types in the same inheritance hierarchy are disallowed.

- `using A for B` only affects the contract it is mentioned in. Previously, the effect was inherited. Now, you have to repeat the `using` statement in all derived contracts that make use of the feature.

Expressions

- Shifts by signed types are disallowed. Previously, shifts by negative amounts were allowed, but reverted at runtime.

- The `finney` and `szabo` denominations are removed. They are rarely used and do not make the actual amount readily visible. Instead, explicit values like `1e20` or the very common `gwei` can be used.

Declarations

- The keyword `var` cannot be used anymore. Previously, this keyword would parse but result in a type error and a suggestion about which type to use. Now, it results in a parser error.

3.22.4 Interface Changes

- JSON AST: Mark hex string literals with `kind:  "hexString"`.
- JSON AST: Members with value `null` are removed from JSON output.
- NatSpec: Constructors and functions have consistent userdoc output.

3.22.5 How to update your code

This section gives detailed instructions on how to update prior code for every breaking change.

- Change `x.f.value(...)()` to `x.f{value:  ...}()`. Similarly `(new C).value(...)()` to `new C{value:  ...}()` and `x.f.gas(...).value(...)()` to `x.f{gas:  ...,  value:  ... }()`.
- Change `now` to `block.timestamp`.
- Change types of right operand in shift operators to unsigned types. For example change `x >> (256 - y)` to `x >> uint(256 - y)`.
- Repeat the `using A for B` statements in all derived contracts if needed.
- Remove the `public` keyword from every constructor.
- Remove the `internal` keyword from every constructor and add `abstract` to the contract (if not already present).
- Change `_slot` and `_offset` suffixes in inline assembly to `.slot` and `.offset`, respectively.

3.23 NatSpec Format

Solidity contracts can use a special form of comments to provide rich documentation for functions, return variables and more. This special form is named the Ethereum Natural Language Specification Format (NatSpec).

This documentation is segmented into developer-focused messages and end-user-facing messages. These messages may be shown to the end user (the human) at the time that they will interact with the contract (i.e. sign a transaction).

It is recommended that Solidity contracts are fully annotated using NatSpec for all public interfaces (everything in the ABI).

NatSpec includes the formatting for comments that the smart contract author will use, and which are understood by the Solidity compiler. Also detailed below is output of the Solidity compiler, which extracts these comments into a machine-readable format.

3.23.1 Documentation Example

Documentation is inserted above each `class`, `interface` and `function` using the doxygen notation format.

Note: a `public` state variable is equivalent to a `function` for the purposes of NatSpec.

- For Solidity you may choose `///` for single or multi-line comments, or `/**` and ending with `*/`.

- For Vyper, use `"""` indented to the inner contents with bare comments. See Vyper documentation.

The following example shows a contract and a function using all available tags.

Note: The Solidity compiler only interprets tags if they are external or public. You are welcome to use similar comments for your internal and private functions, but those will not be parsed.

```solidity
// SPDX-License-Identifier: GPL-3.0
pragma solidity >0.6.10 <0.8.0;

/// @title A simulator for trees
/// @author Larry A. Gardner
/// @notice You can use this contract for only the most basic simulation
/// @dev All function calls are currently implemented without side effects
contract Tree {
    /// @notice Calculate tree age in years, rounded up, for live trees
    /// @dev The Alexandr N. Tetearing algorithm could increase precision
    /// @param rings The number of rings from dendrochronological sample
    /// @return age in years, rounded up for partial years
    function age(uint256 rings) external virtual pure returns (uint256) {
        return rings + 1;
    }

    /// @notice Returns the amount of leaves the tree has.
    /// @dev Returns only a fixed number.
    function leaves() external virtual pure returns(uint256) {
        return 2;
    }
}

contract Plant {
    function leaves() external virtual pure returns(uint256) {
        return 3;
    }
}

contract KumquatTree is Tree, Plant {
    function age(uint256 rings) external override pure returns (uint256) {
        return rings + 2;
    }
```

```
    /// Return the amount of leaves that this specific kind of tree has
    /// @inheritdoc Tree
    function leaves() external override(Tree, Plant) pure returns(uint256) {
        return 3;
    }
}
```

3.23.2 Tags

All tags are optional. The following table explains the purpose of each NatSpec tag and where it may be used. As a special case, if no tags are used then the Solidity compiler will interpret a `///` or `/**` comment in the same way as if it were tagged with `@notice`.

Tag		Context
`@title`	A title that should describe the contract/interface	contract, interface
`@author`	The name of the author	contract, interface
`@notice`	Explain to an end user what this does	contract, interface, function, public state variable, event
`@dev`	Explain to a developer any extra details	contract, interface, function, state variable, event
`@param`	Documents a parameter just like in doxygen (must be followed by parameter name)	function, event
`@return`	Documents the return variables of a contract's function	function, public state variable
`@inheritdoc`	Copies all missing tags from the base function (must be followed by the contract name)	function, public state variable

If your function returns multiple values, like `(int quotient, int remainder)` then use multiple `@return` statements in the same format as the `@param` statements.

Dynamic expressions

The Solidity compiler will pass through NatSpec documentation from your Solidity source code to the JSON output as described in this guide. The consumer of this JSON output, for example the end-user client software, may present this to the end-user directly or it may apply some pre-processing.

For example, some client software will render:

```
/// @notice This function will multiply `a` by 7
```

to the end-user as:

```
This function will multiply 10 by 7
```

if a function is being called and the input `a` is assigned a value of 10.

Specifying these dynamic expressions is outside the scope of the Solidity documentation and you may read more at the radspec project.

Inheritance Notes

Functions without NatSpec will automatically inherit the documentation of their base function. Exceptions to this are:

- When the parameter names are different.

- When there is more than one base function.

- When there is an explicit `@inheritdoc` tag which specifies which contract should be used to inherit.

3.23.3 Documentation Output

When parsed by the compiler, documentation such as the one from the above example will produce two different JSON files. One is meant to be consumed by the end user as a notice when a function is executed and the other to be used by the developer.

If the above contract is saved as `ex1.sol` then you can generate the documentation using:

```
solc --userdoc --devdoc ex1.sol
```

And the output is below.

User Documentation

The above documentation will produce the following user documentation JSON file as output:

```
{
  "methods" :
  {
    "age(uint256)" :
    {
      "notice" : "Calculate tree age in years, rounded up, for live trees"
    }
  },
  "notice" : "You can use this contract for only the most basic simulation"
}
```

Note that the key by which to find the methods is the function's canonical signature as defined in the *Contract ABI* and not simply the function's name.

Developer Documentation

Apart from the user documentation file, a developer documentation JSON file should also be produced and should look like this:

```
{
  "author" : "Larry A. Gardner",
  "details" : "All function calls are currently implemented without side effects",
  "methods" :
  {
    "age(uint256)" :
    {
      "details" : "The Alexandr N. Tetearing algorithm could increase precision",
      "params" :
      {
        "rings" : "The number of rings from dendrochronological sample"
      },
      "return" : "age in years, rounded up for partial years"
    }
```

(continues on next page)

```
  },
  "title" : "A simulator for trees"
}
```

3.24 Security Considerations

While it is usually quite easy to build software that works as expected, it is much harder to check that nobody can use it in a way that was **not** anticipated.

In Solidity, this is even more important because you can use smart contracts to handle tokens or, possibly, even more valuable things. Furthermore, every execution of a smart contract happens in public and, in addition to that, the source code is often available.

Of course you always have to consider how much is at stake: You can compare a smart contract with a web service that is open to the public (and thus, also to malicious actors) and perhaps even open source. If you only store your grocery list on that web service, you might not have to take too much care, but if you manage your bank account using that web service, you should be more careful.

This section will list some pitfalls and general security recommendations but can, of course, never be complete. Also, keep in mind that even if your smart contract code is bug-free, the compiler or the platform itself might have a bug. A list of some publicly known security-relevant bugs of the compiler can be found in the *list of known bugs*, which is also machine-readable. Note that there is a bug bounty program that covers the code generator of the Solidity compiler.

As always, with open source documentation, please help us extend this section (especially, some examples would not hurt)!

NOTE: In addition to the list below, you can find more security recommendations and best practices in Guy Lando's knowledge list and the Consensys GitHub repo.

3.24.1 Pitfalls

Private Information and Randomness

Everything you use in a smart contract is publicly visible, even local variables and state variables marked `private`.

Using random numbers in smart contracts is quite tricky if you do not want miners to be able to cheat.

Re-Entrancy

Any interaction from a contract (A) with another contract (B) and any transfer of Ether hands over control to that contract (B). This makes it possible for B to call back into A before this interaction is completed. To give an example, the following code contains a bug (it is just a snippet and not a complete contract):

```solidity
// SPDX-License-Identifier: GPL-3.0
pragma solidity >=0.4.0 <0.8.0;

// THIS CONTRACT CONTAINS A BUG - DO NOT USE
contract Fund {
    /// @dev Mapping of ether shares of the contract.
    mapping(address => uint) shares;
    /// Withdraw your share.
    function withdraw() public {
```

```
            if (msg.sender.send(shares[msg.sender]))
                shares[msg.sender] = 0;
        }
}
```

The problem is not too serious here because of the limited gas as part of `send`, but it still exposes a weakness: Ether transfer can always include code execution, so the recipient could be a contract that calls back into `withdraw`. This would let it get multiple refunds and basically retrieve all the Ether in the contract. In particular, the following contract will allow an attacker to refund multiple times as it uses `call` which forwards all remaining gas by default:

```solidity
// SPDX-License-Identifier: GPL-3.0
pragma solidity >=0.6.2 <0.8.0;

// THIS CONTRACT CONTAINS A BUG - DO NOT USE
contract Fund {
    /// @dev Mapping of ether shares of the contract.
    mapping(address => uint) shares;
    /// Withdraw your share.
    function withdraw() public {
        (bool success,) = msg.sender.call{value: shares[msg.sender]}("");
        if (success)
            shares[msg.sender] = 0;
    }
}
```

To avoid re-entrancy, you can use the Checks-Effects-Interactions pattern as outlined further below:

```solidity
// SPDX-License-Identifier: GPL-3.0
pragma solidity >=0.4.11 <0.8.0;

contract Fund {
    /// @dev Mapping of ether shares of the contract.
    mapping(address => uint) shares;
    /// Withdraw your share.
    function withdraw() public {
        uint share = shares[msg.sender];
        shares[msg.sender] = 0;
        msg.sender.transfer(share);
    }
}
```

Note that re-entrancy is not only an effect of Ether transfer but of any function call on another contract. Furthermore, you also have to take multi-contract situations into account. A called contract could modify the state of another contract you depend on.

Gas Limit and Loops

Loops that do not have a fixed number of iterations, for example, loops that depend on storage values, have to be used carefully: Due to the block gas limit, transactions can only consume a certain amount of gas. Either explicitly or just due to normal operation, the number of iterations in a loop can grow beyond the block gas limit which can cause the complete contract to be stalled at a certain point. This may not apply to `view` functions that are only executed to read data from the blockchain. Still, such functions may be called by other contracts as part of on-chain operations and stall those. Please be explicit about such cases in the documentation of your contracts.

Sending and Receiving Ether

- Neither contracts nor "external accounts" are currently able to prevent that someone sends them Ether. Contracts can react on and reject a regular transfer, but there are ways to move Ether without creating a message call. One way is to simply "mine to" the contract address and the second way is using `selfdestruct(x)`.

- If a contract receives Ether (without a function being called), either the *receive Ether* or the *fallback* function is executed. If it does not have a receive nor a fallback function, the Ether will be rejected (by throwing an exception). During the execution of one of these functions, the contract can only rely on the "gas stipend" it is passed (2300 gas) being available to it at that time. This stipend is not enough to modify storage (do not take this for granted though, the stipend might change with future hard forks). To be sure that your contract can receive Ether in that way, check the gas requirements of the receive and fallback functions (for example in the "details" section in Remix).

- There is a way to forward more gas to the receiving contract using `addr.call{value: x}("")`. This is essentially the same as `addr.transfer(x)`, only that it forwards all remaining gas and opens up the ability for the recipient to perform more expensive actions (and it returns a failure code instead of automatically propagating the error). This might include calling back into the sending contract or other state changes you might not have thought of. So it allows for great flexibility for honest users but also for malicious actors.

- Use the most precise units to represent the wei amount as possible, as you lose any that is rounded due to a lack of precision.

- If you want to send Ether using `address.transfer`, there are certain details to be aware of:

  1. If the recipient is a contract, it causes its receive or fallback function to be executed which can, in turn, call back the sending contract.

  2. Sending Ether can fail due to the call depth going above 1024. Since the caller is in total control of the call depth, they can force the transfer to fail; take this possibility into account or use `send` and make sure to always check its return value. Better yet, write your contract using a pattern where the recipient can withdraw Ether instead.

  3. Sending Ether can also fail because the execution of the recipient contract requires more than the allotted amount of gas (explicitly by using *require*, *assert*, *revert* or because the operation is too expensive) - it "runs out of gas" (OOG). If you use `transfer` or `send` with a return value check, this might provide a means for the recipient to block progress in the sending contract. Again, the best practice here is to use a *"withdraw" pattern instead of a "send" pattern*.

Callstack Depth

External function calls can fail any time because they exceed the maximum call stack of 1024. In such situations, Solidity throws an exception. Malicious actors might be able to force the call stack to a high value before they interact with your contract.

Note that `.send()` does **not** throw an exception if the call stack is depleted but rather returns `false` in that case. The low-level functions `.call()`, `.delegatecall()` and `.staticcall()` behave in the same way.

tx.origin

Never use tx.origin for authorization. Let's say you have a wallet contract like this:

```solidity
// SPDX-License-Identifier: GPL-3.0
pragma solidity ^0.7.0;

// THIS CONTRACT CONTAINS A BUG - DO NOT USE
```

```solidity
contract TxUserWallet {
    address owner;

    constructor() {
        owner = msg.sender;
    }

    function transferTo(address payable dest, uint amount) public {
        require(tx.origin == owner);
        dest.transfer(amount);
    }
}
```

Now someone tricks you into sending Ether to the address of this attack wallet:

```solidity
// SPDX-License-Identifier: GPL-3.0
pragma solidity ^0.7.0;

interface TxUserWallet {
    function transferTo(address payable dest, uint amount) external;
}

contract TxAttackWallet {
    address payable owner;

    constructor() {
        owner = msg.sender;
    }

    receive() external payable {
        TxUserWallet(msg.sender).transferTo(owner, msg.sender.balance);
    }
}
```

If your wallet had checked `msg.sender` for authorization, it would get the address of the attack wallet, instead of the owner address. But by checking `tx.origin`, it gets the original address that kicked off the transaction, which is still the owner address. The attack wallet instantly drains all your funds.

Two's Complement / Underflows / Overflows

As in many programming languages, Solidity's integer types are not actually integers. They resemble integers when the values are small, but behave differently if the numbers are larger. For example, the following is true: `uint8(255) + uint8(1) == 0`. This situation is called an *overflow*. It occurs when an operation is performed that requires a fixed size variable to store a number (or piece of data) that is outside the range of the variable's data type. An *underflow* is the converse situation: `uint8(0) - uint8(1) == 255`.

In general, read about the limits of two's complement representation, which even has some more special edge cases for signed numbers.

Try to use `require` to limit the size of inputs to a reasonable range and use the *SMT checker* to find potential overflows, or use a library like SafeMath if you want all overflows to cause a revert.

Code such as `require((balanceOf[_to] + _value) >= balanceOf[_to])` can also help you check if values are what you expect.

Clearing Mappings

The Solidity type `mapping` (see *Mapping Types*) is a storage-only key-value data structure that does not keep track of the keys that were assigned a non-zero value. Because of that, cleaning a mapping without extra information about the written keys is not possible. If a `mapping` is used as the base type of a dynamic storage array, deleting or popping the array will have no effect over the `mapping` elements. The same happens, for example, if a `mapping` is used as the type of a member field of a `struct` that is the base type of a dynamic storage array. The `mapping` is also ignored in assignments of structs or arrays containing a `mapping`.

```solidity
// SPDX-License-Identifier: GPL-3.0
pragma solidity >=0.6.0 <0.8.0;

contract Map {
    mapping (uint => uint)[] array;

    function allocate(uint _newMaps) public {
        for (uint i = 0; i < _newMaps; i++)
            array.push();
    }

    function writeMap(uint _map, uint _key, uint _value) public {
        array[_map][_key] = _value;
    }

    function readMap(uint _map, uint _key) public view returns (uint) {
        return array[_map][_key];
    }

    function eraseMaps() public {
        delete array;
    }
}
```

Consider the example above and the following sequence of calls: `allocate(10)`, `writeMap(4, 128, 256)`. At this point, calling `readMap(4, 128)` returns 256. If we call `eraseMaps`, the length of state variable `array` is zeroed, but since its `mapping` elements cannot be zeroed, their information stays alive in the contract's storage. After deleting `array`, calling `allocate(5)` allows us to access `array[4]` again, and calling `readMap(4, 128)` returns 256 even without another call to `writeMap`.

If your `mapping` information must be deleted, consider using a library similar to iterable mapping, allowing you to traverse the keys and delete their values in the appropriate `mapping`.

Minor Details

- Types that do not occupy the full 32 bytes might contain "dirty higher order bits". This is especially important if you access `msg.data` - it poses a malleability risk: You can craft transactions that call a function `f(uint8 x)` with a raw byte argument of `0xff000001` and with `0x00000001`. Both are fed to the contract and both will look like the number `1` as far as `x` is concerned, but `msg.data` will be different, so if you use `keccak256(msg.data)` for anything, you will get different results.

3.24.2 Recommendations

Take Warnings Seriously

If the compiler warns you about something, you should better change it. Even if you do not think that this particular warning has security implications, there might be another issue buried beneath it. Any compiler warning we issue can be silenced by slight changes to the code.

Always use the latest version of the compiler to be notified about all recently introduced warnings.

Restrict the Amount of Ether

Restrict the amount of Ether (or other tokens) that can be stored in a smart contract. If your source code, the compiler or the platform has a bug, these funds may be lost. If you want to limit your loss, limit the amount of Ether.

Keep it Small and Modular

Keep your contracts small and easily understandable. Single out unrelated functionality in other contracts or into libraries. General recommendations about source code quality of course apply: Limit the amount of local variables, the length of functions and so on. Document your functions so that others can see what your intention was and whether it is different than what the code does.

Use the Checks-Effects-Interactions Pattern

Most functions will first perform some checks (who called the function, are the arguments in range, did they send enough Ether, does the person have tokens, etc.). These checks should be done first.

As the second step, if all checks passed, effects to the state variables of the current contract should be made. Interaction with other contracts should be the very last step in any function.

Early contracts delayed some effects and waited for external function calls to return in a non-error state. This is often a serious mistake because of the re-entrancy problem explained above.

Note that, also, calls to known contracts might in turn cause calls to unknown contracts, so it is probably better to just always apply this pattern.

Include a Fail-Safe Mode

While making your system fully decentralised will remove any intermediary, it might be a good idea, especially for new code, to include some kind of fail-safe mechanism:

You can add a function in your smart contract that performs some self-checks like "Has any Ether leaked?", "Is the sum of the tokens equal to the balance of the contract?" or similar things. Keep in mind that you cannot use too much gas for that, so help through off-chain computations might be needed there.

If the self-check fails, the contract automatically switches into some kind of "failsafe" mode, which, for example, disables most of the features, hands over control to a fixed and trusted third party or just converts the contract into a simple "give me back my money" contract.

Ask for Peer Review

The more people examine a piece of code, the more issues are found. Asking people to review your code also helps as a cross-check to find out whether your code is easy to understand - a very important criterion for good smart contracts.

3.24.3 Formal Verification

Using formal verification, it is possible to perform an automated mathematical proof that your source code fulfills a certain formal specification. The specification is still formal (just as the source code), but usually much simpler.

Note that formal verification itself can only help you understand the difference between what you did (the specification) and how you did it (the actual implementation). You still need to check whether the specification is what you wanted and that you did not miss any unintended effects of it.

Solidity implements a formal verification approach based on SMT solving. The SMTChecker module automatically tries to prove that the code satisfies the specification given by `require/assert` statements. That is, it considers `require` statements as assumptions and tries to prove that the conditions inside `assert` statements are always true. If an assertion failure is found, a counterexample is given to the user, showing how the assertion can be violated.

The SMTChecker also checks automatically for arithmetic underflow/overflow, trivial conditions and unreachable code. It is currently an experimental feature, therefore in order to use it you need to enable it via *a pragma directive*.

The SMTChecker traverses the Solidity AST creating and collecting program constraints. When it encounters a verification target, an SMT solver is invoked to determine the outcome. If a check fails, the SMTChecker provides specific input values that lead to the failure.

While the SMTChecker encodes Solidity code into SMT constraints, it contains two reasoning engines that use that encoding in different ways.

SMT Encoding

The SMT encoding tries to be as precise as possible, mapping Solidity types and expressions to their closest SMT-LIB representation, as shown in the table below.

Solidity type	SMT sort	Theories (quantifier-free)
Boolean	Bool	Bool
intN, uintN, address, bytesN, enum	Integer	LIA, NIA
array, mapping, bytes, string	Tuple (Array elements, Integer length)	Datatypes, Arrays, LIA
struct	Tuple	Datatypes
other types	Integer	LIA

Types that are not yet supported are abstracted by a single 256-bit unsigned integer, where their unsupported operations are ignored.

For more details on how the SMT encoding works internally, see the paper SMT-based Verification of Solidity Smart Contracts.

Model Checking Engines

The SMTChecker module implements two different reasoning engines that use the SMT encoding above, a Bounded Model Checker (BMC) and a system of Constrained Horn Clauses (CHC). Both engines are currently under development, and have different characteristics.

Bounded Model Checker (BMC)

The BMC engine analyzes functions in isolation, that is, it does not take the overall behavior of the contract throughout many transactions into account when analyzing each function. Loops are also ignored in this engine at the moment. Internal function calls are inlined as long as they are not recursive, direct or indirectly. External function calls are inlined if possible, and knowledge that is potentially affected by reentrancy is erased.

The characteristics above make BMC easily prone to reporting false positives, but it is also lightweight and should be able to quickly find small local bugs.

Constrained Horn Clauses (CHC)

The Solidity contract's Control Flow Graph (CFG) is modelled as a system of Horn clauses, where the lifecycle of the contract is represented by a loop that can visit every public/external function non-deterministically. This way, the behavior of the entire contract over an unbounded number of transactions is taken into account when analyzing any function. Loops are fully supported by this engine. Internal function calls are supported, but external function calls are currently unsupported.

The CHC engine is much more powerful than BMC in terms of what it can prove, and might require more computing resources.

Abstraction and False Positives

The SMTChecker implements abstractions in an incomplete and sound way: If a bug is reported, it might be a false positive introduced by abstractions (due to erasing knowledge or using a non-precise type). If it determines that a verification target is safe, it is indeed safe, that is, there are no false negatives (unless there is a bug in the SMTChecker).

In the BMC engine, function calls to the same contract (or base contracts) are inlined when possible, that is, when their implementation is available. Calls to functions in other contracts are not inlined even if their code is available, since we cannot guarantee that the actual deployed code is the same.

The CHC engine creates nonlinear Horn clauses that use summaries of the called functions to support internal function calls. The same approach can and will be used for external function calls, but the latter requires more work regarding the entire state of the blockchain and is still unimplemented.

Complex pure functions are abstracted by an uninterpreted function (UF) over the arguments.

Functions	SMT behavior
`assert`	Verification target
`require`	Assumption
internal	BMC: Inline function call CHC: Function summaries
external	BMC: Inline function call or erase knowledge about state variables and local storage references. CHC: Function summaries and erase state knowledge.
`gasleft`, `blockhash`, `keccak256`, `ecrecover` `ripemd160`, `addmod`, `mulmod`	Abstracted with UF
pure functions without implementation (external or complex)	Abstracted with UF
external functions without implementation	BMC: Unsupported CHC: Nondeterministic summary
others	Currently unsupported

Using abstraction means loss of precise knowledge, but in many cases it does not mean loss of proving power.

```
// SPDX-License-Identifier: GPL-3.0
pragma solidity >=0.5.0;
pragma experimental SMTChecker;
// This may report a warning if no SMT solver available.
```

```
contract Recover
{
    function f(
        bytes32 hash,
        uint8 _v1, uint8 _v2,
        bytes32 _r1, bytes32 _r2,
        bytes32 _s1, bytes32 _s2
    ) public pure returns (address) {
        address a1 = ecrecover(hash, _v1, _r1, _s1);
        require(_v1 == _v2);
        require(_r1 == _r2);
        require(_s1 == _s2);
        address a2 = ecrecover(hash, _v2, _r2, _s2);
        assert(a1 == a2);
        return a1;
    }
}
```

In the example above, the SMTChecker is not expressive enough to actually compute `ecrecover`, but by modelling the function calls as uninterpreted functions we know that the return value is the same when called on equivalent parameters. This is enough to prove that the assertion above is always true.

Abstracting a function call with an UF can be done for functions known to be deterministic, and can be easily done for pure functions. It is however difficult to do this with general external functions, since they might depend on state variables.

External function calls also imply that any current knowledge that the SMTChecker might have regarding mutable state variables needs to be erased to guarantee no false negatives, since the called external function might direct or indirectly call a function in the analyzed contract that changes state variables.

Reference Types and Aliasing

Solidity implements aliasing for reference types with the same *data location*. That means one variable may be modified through a reference to the same data area. The SMTChecker does not keep track of which references refer to the same data. This implies that whenever a local reference or state variable of reference type is assigned, all knowledge regarding variables of the same type and data location is erased. If the type is nested, the knowledge removal also includes all the prefix base types.

```
// SPDX-License-Identifier: GPL-3.0
pragma solidity >=0.5.0;
pragma experimental ABIEncoderV2;
pragma experimental SMTChecker;
// This will report a warning

contract Aliasing
{
    uint[] array1;
    uint[][] array2;
    function f(
        uint[] memory a,
        uint[] memory b,
        uint[][] memory c,
        uint[] storage d
    ) internal {
        array1[0] = 42;
```

```
        a[0] = 2;
        c[0][0] = 2;
        b[0] = 1;
        // Erasing knowledge about memory references should not
        // erase knowledge about state variables.
        assert(array1[0] == 42);
        // However, an assignment to a storage reference will erase
        // storage knowledge accordingly.
        d[0] = 2;
        // Fails as false positive because of the assignment above.
        assert(array1[0] == 42);
        // Fails because `a == b` is possible.
        assert(a[0] == 2);
        // Fails because `c[i] == b` is possible.
        assert(c[0][0] == 2);
        assert(d[0] == 2);
        assert(b[0] == 1);
    }
    function g(
        uint[] memory a,
        uint[] memory b,
        uint[][] memory c,
        uint x
    ) public {
        f(a, b, c, array2[x]);
    }
}
```

After the assignment to `b[0]`, we need to clear knowledge about `a` since it has the same type (`uint[]`) and data location (memory). We also need to clear knowledge about `c`, since its base type is also a `uint[]` located in memory. This implies that some `c[i]` could refer to the same data as `b` or `a`.

Notice that we do not clear knowledge about `array` and `d` because they are located in storage, even though they also have type `uint[]`. However, if `d` was assigned, we would need to clear knowledge about `array` and vice-versa.

Real World Assumptions

Some scenarios can be expressed in Solidity and the EVM, but are expected to never occur in practice. One of such cases is the length of a dynamic storage array overflowing during a push: If the `push` operation is applied to an array of length 2^256 - 1, its length silently overflows. However, this is unlikely to happen in practice, since the operations required to grow the array to that point would take billions of years to execute.

3.25 Resources

3.25.1 General

- Ethereum
- Changelog
- Source Code
- Ethereum Stackexchange
- Language Users Chat

- Compiler Developers Chat

3.25.2 Solidity Integrations

- Generic:

    - **EthFiddle** Solidity IDE in the Browser. Write and share your Solidity code. Uses server-side components.

    - **Remix** Browser-based IDE with integrated compiler and Solidity runtime environment without server-side components.

    - **Solhint** Solidity linter that provides security, style guide and best practice rules for smart contract validation.

    - **Solidity IDE** Browser-based IDE with integrated compiler, Ganache and local file system support.

    - **Ethlint** Linter to identify and fix style and security issues in Solidity.

    - **Superblocks Lab** Browser-based IDE. Built-in browser-based VM and Metamask integration (one click deployment to Testnet/Mainnet).

- Atom:

    - **Etheratom** Plugin for the Atom editor that features syntax highlighting, compilation and a runtime environment (Backend node & VM compatible).

    - **Atom Solidity Linter** Plugin for the Atom editor that provides Solidity linting.

    - **Atom Solium Linter** Configurable Solidity linter for Atom using Solium (now Ethlint) as a base.

- Eclipse:

    - **YAKINDU Solidity Tools** Eclipse based IDE. Features context sensitive code completion and help, code navigation, syntax coloring, built in compiler, quick fixes and templates.

- Emacs:

    - **Emacs Solidity** Plugin for the Emacs editor providing syntax highlighting and compilation error reporting.

- IntelliJ:

    - **IntelliJ IDEA plugin** Solidity plugin for IntelliJ IDEA (and all other JetBrains IDEs)

- Sublime:

    - **Package for SublimeText - Solidity language syntax** Solidity syntax highlighting for SublimeText editor.

- Vim:

    - **Vim Solidity** Plugin for the Vim editor providing syntax highlighting.

    - **Vim Syntastic** Plugin for the Vim editor providing compile checking.

- Visual Studio Code:

    - **Visual Studio Code extension** Solidity plugin for Microsoft Visual Studio Code that includes syntax highlighting and the Solidity compiler.

3.25.3 Solidity Tools

- **ABI to Solidity interface converter** A script for generating contract interfaces from the ABI of a smart contract.

- **Dapp** Build tool, package manager, and deployment assistant for Solidity.

- **Doxity** Documentation Generator for Solidity.

- **evmdis** EVM Disassembler that performs static analysis on the bytecode to provide a higher level of abstraction than raw EVM operations.

- **EVM Lab** Rich tool package to interact with the EVM. Includes a VM, Etherchain API, and a trace-viewer with gas cost display.

- **leafleth** A documentation generator for Solidity smart-contracts.

- **PIET** A tool to develop, audit and use Solidity smart contracts through a simple graphical interface.

- **solc-select** A script to quickly switch between Solidity compiler versions.

- **Solidity prettier plugin** A Prettier Plugin for Solidity.

- **Solidity REPL** Try Solidity instantly with a command-line Solidity console.

- **solgraph** Visualize Solidity control flow and highlight potential security vulnerabilities.

- **Securify** Fully automated online static analyzer for smart contracts, providing a security report based on vulnerability patterns.

- **Sūrya** Utility tool for smart contract systems, offering a number of visual outputs and information about the contracts' structure. Also supports querying the function call graph.

- **Universal Mutator** A tool for mutation generation, with configurable rules and support for Solidity and Vyper.

3.25.4 Third-Party Solidity Parsers and Grammars

- **Solidity Parser for JavaScript** A Solidity parser for JS built on top of a robust ANTLR4 grammar.

3.26 Using the compiler

3.26.1 Using the Commandline Compiler

Note: This section does not apply to *solcjs*, not even if it is used in commandline mode.

Basic usage

One of the build targets of the Solidity repository is `solc`, the solidity commandline compiler. Using `solc --help` provides you with an explanation of all options. The compiler can produce various outputs, ranging from simple binaries and assembly over an abstract syntax tree (parse tree) to estimations of gas usage. If you only want to compile a single file, you run it as `solc --bin sourceFile.sol` and it will print the binary. If you want to get some of the more advanced output variants of `solc`, it is probably better to tell it to output everything to separate files using `solc -o outputDirectory --bin --ast-json --asm sourceFile.sol`.

Optimizer options

Before you deploy your contract, activate the optimizer when compiling using `solc --optimize --bin sourceFile.sol`. By default, the optimizer will optimize the contract assuming it is called 200 times across its lifetime (more specifically, it assumes each opcode is executed around 200 times). If you want the initial contract deployment to be cheaper and the later function executions to be more expensive, set it to `--optimize-runs=1`. If you expect many transactions and do not care for higher deployment cost and output size, set `--optimize-runs` to a high number. This parameter has effects on the following (this might change in the future):

- the size of the binary search in the function dispatch routine
- the way constants like large numbers or strings are stored

Path remapping

The commandline compiler will automatically read imported files from the filesystem, but it is also possible to provide path redirects using `prefix=path` in the following way:

```
solc github.com/ethereum/dapp-bin/=/usr/local/lib/dapp-bin/ file.sol
```

This essentially instructs the compiler to search for anything starting with `github.com/ethereum/dapp-bin/` under `/usr/local/lib/dapp-bin`. `solc` will not read files from the filesystem that lie outside of the remapping targets and outside of the directories where explicitly specified source files reside, so things like `import "/etc/passwd";` only work if you add `/=/` as a remapping.

An empty remapping prefix is not allowed.

If there are multiple matches due to remappings, the one with the longest common prefix is selected.

When accessing the filesystem to search for imports, all paths are treated as if they were fully qualified paths. This behaviour can be customized by adding the command line option `--base-path` with a path to be prepended before each filesystem access for imports is performed. Furthermore, the part added via `--base-path` will not appear in the contract metadata.

For security reasons the compiler has restrictions what directories it can access. Paths (and their subdirectories) of source files specified on the commandline and paths defined by remappings are allowed for import statements, but everything else is rejected. Additional paths (and their subdirectories) can be allowed via the `--allow-paths /sample/path,/another/sample/path` switch.

Everything inside the path specified via `--base-path` is always allowed.

Library linking

If your contracts use *libraries*, you will notice that the bytecode contains substrings of the form `__$53aea86b7d70b31448b230b20ae141a537$__`. These are placeholders for the actual library addresses. The placeholder is a 34 character prefix of the hex encoding of the keccak256 hash of the fully qualified library name. The bytecode file will also contain lines of the form `// <placeholder> -> <fq library name>` at the end to help identify which libraries the placeholders represent. Note that the fully qualified library name is the path of its source file and the library name separated by `:`. You can use `solc` as a linker meaning that it will insert the library addresses for you at those points:

Either add `--libraries "file.sol:Math:0x1234567890123456789012345678901234567890 file.sol:Heap:0xabCD567890123456789012345678901234567890"` to your command to provide an address for each library (use commas or spaces as separators) or store the string in a file (one library per line) and run `solc` using `--libraries fileName`.

If `solc` is called with the option `--standard-json`, it will expect a JSON input (as explained below) on the standard input, and return a JSON output on the standard output. This is the recommended interface for more complex and especially automated uses. The process will always terminate in a "success" state and report any errors via the JSON output. The option `--base-path` is also processed in standard-json mode.

If `solc` is called with the option `--link`, all input files are interpreted to be unlinked binaries (hex-encoded) in the `__$53aea86b7d70b31448b230b20ae141a537$__`-format given above and are linked in-place (if the input is read from stdin, it is written to stdout). All options except `--libraries` are ignored (including `-o`) in this case.

> **Warning:** Manually linking libraries on the generated bytecode is discouraged because it does not update contract metadata. Since metadata contains a list of libraries specified at the time of compilation and bytecode contains a metadata hash, you will get different binaries, depending on when linking is performed.
>
> You should ask the compiler to link the libraries at the time a contract is compiled by either using the `--libraries` option of `solc` or the `libraries` key if you use the standard-JSON interface to the compiler.

> **Note:** The library placeholder used to be the fully qualified name of the library itself instead of the hash of it. This format is still supported by `solc --link` but the compiler will no longer output it. This change was made to reduce the likelihood of a collision between libraries, since only the first 36 characters of the fully qualified library name could be used.

3.26.2 Setting the EVM version to target

When you compile your contract code you can specify the Ethereum virtual machine version to compile for to avoid particular features or behaviours.

> **Warning:** Compiling for the wrong EVM version can result in wrong, strange and failing behaviour. Please ensure, especially if running a private chain, that you use matching EVM versions.

On the command line, you can select the EVM version as follows:

```
solc --evm-version <VERSION> contract.sol
```

In the *standard JSON interface*, use the `"evmVersion"` key in the `"settings"` field:

```
{
  "sources": { ... },
  "settings": {
    "optimizer": { ... },
    "evmVersion": "<VERSION>"
  }
}
```

Target options

Below is a list of target EVM versions and the compiler-relevant changes introduced at each version. Backward compatibility is not guaranteed between each version.

- **homestead**

- – (oldest version)

- **tangerineWhistle**

    - – Gas cost for access to other accounts increased, relevant for gas estimation and the optimizer.

    - – All gas sent by default for external calls, previously a certain amount had to be retained.

- **spuriousDragon**

    - – Gas cost for the `exp` opcode increased, relevant for gas estimation and the optimizer.

- **byzantium**

    - – Opcodes `returndatacopy`, `returndatasize` and `staticcall` are available in assembly.

    - – The `staticcall` opcode is used when calling non-library view or pure functions, which prevents the functions from modifying state at the EVM level, i.e., even applies when you use invalid type conversions.

    - – It is possible to access dynamic data returned from function calls.

    - – `revert` opcode introduced, which means that `revert()` will not waste gas.

- **constantinople**

    - – Opcodes `create2`,` ``extcodehash, `shl`, `shr` and `sar` are available in assembly.

    - – Shifting operators use shifting opcodes and thus need less gas.

- **petersburg**

    - – The compiler behaves the same way as with constantinople.

- **istanbul** (**default**)

    - – Opcodes `chainid` and `selfbalance` are available in assembly.

- `berlin` (**experimental**)

3.26.3 Compiler Input and Output JSON Description

The recommended way to interface with the Solidity compiler especially for more complex and automated setups is the so-called JSON-input-output interface. The same interface is provided by all distributions of the compiler.

The fields are generally subject to change, some are optional (as noted), but we try to only make backwards compatible changes.

The compiler API expects a JSON formatted input and outputs the compilation result in a JSON formatted output. The standard error output is not used and the process will always terminate in a "success" state, even if there were errors. Errors are always reported as part of the JSON output.

The following subsections describe the format through an example. Comments are of course not permitted and used here only for explanatory purposes.

Input Description

```
{
  // Required: Source code language. Currently supported are "Solidity" and "Yul".
  "language": "Solidity",
  // Required
  "sources":
```

```
{
  // The keys here are the "global" names of the source files,
  // imports can use other files via remappings (see below).
  "myFile.sol":
  {
    // Optional: keccak256 hash of the source file
    // It is used to verify the retrieved content if imported via URLs.
    "keccak256": "0x123...",
    // Required (unless "content" is used, see below): URL(s) to the source file.
    // URL(s) should be imported in this order and the result checked against the
    // keccak256 hash (if available). If the hash doesn't match or none of the
    // URL(s) result in success, an error should be raised.
    // Using the commandline interface only filesystem paths are supported.
    // With the JavaScript interface the URL will be passed to the user-supplied
    // read callback, so any URL supported by the callback can be used.
    "urls":
    [
      "bzzr://56ab...",
      "ipfs://Qma...",
      "/tmp/path/to/file.sol"
      // If files are used, their directories should be added to the command line
→via
      // `--allow-paths <path>`.
    ]
  },
  "destructible":
  {
    // Optional: keccak256 hash of the source file
    "keccak256": "0x234...",
    // Required (unless "urls" is used): literal contents of the source file
    "content": "contract destructible is owned { function shutdown() { if (msg.
→sender == owner) selfdestruct(owner); } }"
  }
},
// Optional
"settings":
{
  // Optional: Stop compilation after the given stage. Currently only "parsing" is
→valid here
  "stopAfter": "parsing",
  // Optional: Sorted list of remappings
  "remappings": [ ":g=/dir" ],
  // Optional: Optimizer settings
  "optimizer": {
    // disabled by default
    "enabled": true,
    // Optimize for how many times you intend to run the code.
    // Lower values will optimize more for initial deployment cost, higher
    // values will optimize more for high-frequency usage.
    "runs": 200,
    // Switch optimizer components on or off in detail.
    // The "enabled" switch above provides two defaults which can be
    // tweaked here. If "details" is given, "enabled" can be omitted.
    "details": {
      // The peephole optimizer is always on if no details are given,
      // use details to switch it off.
      "peephole": true,
```

(continued from previous page)

```
        // The unused jumpdest remover is always on if no details are given,
        // use details to switch it off.
        "jumpdestRemover": true,
        // Sometimes re-orders literals in commutative operations.
        "orderLiterals": false,
        // Removes duplicate code blocks
        "deduplicate": false,
        // Common subexpression elimination, this is the most complicated step but
        // can also provide the largest gain.
        "cse": false,
        // Optimize representation of literal numbers and strings in code.
        "constantOptimizer": false,
        // The new Yul optimizer. Mostly operates on the code of ABI coder v2
        // and inline assembly.
        // It is activated together with the global optimizer setting
        // and can be deactivated here.
        // Before Solidity 0.6.0 it had to be activated through this switch.
        "yul": false,
        // Tuning options for the Yul optimizer.
        "yulDetails": {
          // Improve allocation of stack slots for variables, can free up stack slots␣
→early.
          // Activated by default if the Yul optimizer is activated.
          "stackAllocation": true,
          // Select optimization steps to be applied.
          // Optional, the optimizer will use the default sequence if omitted.
          "optimizerSteps": "dhfoDgvulfnTUtnIf..."
        }
      }
    },
    // Version of the EVM to compile for.
    // Affects type checking and code generation. Can be homestead,
    // tangerineWhistle, spuriousDragon, byzantium, constantinople, petersburg,␣
→istanbul or berlin
    "evmVersion": "byzantium",
    // Optional: Change compilation pipeline to go through the Yul intermediate␣
→representation.
    // This is a highly EXPERIMENTAL feature, not to be used for production. This is␣
→false by default.
    "viaIR": true,
    // Optional: Debugging settings
    "debug": {
      // How to treat revert (and require) reason strings. Settings are
      // "default", "strip", "debug" and "verboseDebug".
      // "default" does not inject compiler-generated revert strings and keeps user-
→supplied ones.
      // "strip" removes all revert strings (if possible, i.e. if literals are used)␣
→keeping side-effects
      // "debug" injects strings for compiler-generated internal reverts, implemented␣
→for ABI encoders V1 and V2 for now.
      // "verboseDebug" even appends further information to user-supplied revert␣
→strings (not yet implemented)
      "revertStrings": "default"
    }
    // Metadata settings (optional)
    "metadata": {
      // Use only literal content and not URLs (false by default)
```

(continues on next page)

```
      "useLiteralContent": true,
      // Use the given hash method for the metadata hash that is appended to the␣
→bytecode.
      // The metadata hash can be removed from the bytecode via option "none".
      // The other options are "ipfs" and "bzzr1".
      // If the option is omitted, "ipfs" is used by default.
      "bytecodeHash": "ipfs"
   },
   // Addresses of the libraries. If not all libraries are given here,
   // it can result in unlinked objects whose output data is different.
   "libraries": {
      // The top level key is the the name of the source file where the library is␣
→used.
      // If remappings are used, this source file should match the global path
      // after remappings were applied.
      // If this key is an empty string, that refers to a global level.
      "myFile.sol": {
        "MyLib": "0x123123..."
      }
   }
   // The following can be used to select desired outputs based
   // on file and contract names.
   // If this field is omitted, then the compiler loads and does type checking,
   // but will not generate any outputs apart from errors.
   // The first level key is the file name and the second level key is the contract␣
→name.
   // An empty contract name is used for outputs that are not tied to a contract
   // but to the whole source file like the AST.
   // A star as contract name refers to all contracts in the file.
   // Similarly, a star as a file name matches all files.
   // To select all outputs the compiler can possibly generate, use
   // "outputSelection: { "*": { "*": [ "*" ], "": [ "*" ] } }"
   // but note that this might slow down the compilation process needlessly.
   //
   // The available output types are as follows:
   //
   // File level (needs empty string as contract name):
   //   ast - AST of all source files
   //   legacyAST - legacy AST of all source files
   //
   // Contract level (needs the contract name or "*"):
   //   abi - ABI
   //   devdoc - Developer documentation (natspec)
   //   userdoc - User documentation (natspec)
   //   metadata - Metadata
   //   ir - Yul intermediate representation of the code before optimization
   //   irOptimized - Intermediate representation after optimization
   //   storageLayout - Slots, offsets and types of the contract's state variables.
   //   evm.assembly - New assembly format
   //   evm.legacyAssembly - Old-style assembly format in JSON
   //   evm.bytecode.object - Bytecode object
   //   evm.bytecode.opcodes - Opcodes list
   //   evm.bytecode.sourceMap - Source mapping (useful for debugging)
   //   evm.bytecode.linkReferences - Link references (if unlinked object)
   //   evm.bytecode.generatedSources - Sources generated by the compiler
   //   evm.deployedBytecode* - Deployed bytecode (has all the options that evm.
→bytecode has)
```

```
    //   evm.deployedBytecode.immutableReferences - Map from AST ids to bytecode␣
↪ranges that reference immutables
    //   evm.methodIdentifiers - The list of function hashes
    //   evm.gasEstimates - Function gas estimates
    //   ewasm.wast - Ewasm in WebAssembly S-expressions format
    //   ewasm.wasm - Ewasm in WebAssembly binary format
    //
    // Note that using a using `evm`, `evm.bytecode`, `ewasm`, etc. will select every
    // target part of that output. Additionally, `*` can be used as a wildcard to␣
↪request everything.
    //
    "outputSelection": {
      "*": {
        "*": [
          "metadata", "evm.bytecode" // Enable the metadata and bytecode outputs of␣
↪every single contract.
          , "evm.bytecode.sourceMap" // Enable the source map output of every single␣
↪contract.
        ],
        "": [
          "ast" // Enable the AST output of every single file.
        ]
      },
      // Enable the abi and opcodes output of MyContract defined in file def.
      "def": {
        "MyContract": [ "abi", "evm.bytecode.opcodes" ]
      }
    }
  },
  "modelCheckerSettings":
  {
    // Choose which model checker engine to use: all (default), bmc, chc, none.
    "engine": "chc",
    // Timeout for each SMT query in milliseconds.
    // If this option is not given, the SMTChecker will use a deterministic
    // resource limit by default.
    // A given timeout of 0 means no resource/time restrictions for any query.
    "timeout": 20000
  }
}
```

Output Description

```
{
  // Optional: not present if no errors/warnings were encountered
  "errors": [
    {
      // Optional: Location within the source file.
      "sourceLocation": {
        "file": "sourceFile.sol",
        "start": 0,
        "end": 100
      ],
      // Optional: Further locations (e.g. places of conflicting declarations)
      "secondarySourceLocations": [
```

```
        {
          "file": "sourceFile.sol",
          "start": 64,
          "end": 92,
          "message": "Other declaration is here:"
        }
      ],
      // Mandatory: Error type, such as "TypeError", "InternalCompilerError",
→"Exception", etc.
      // See below for complete list of types.
      "type": "TypeError",
      // Mandatory: Component where the error originated, such as "general", "ewasm",
→etc.
      "component": "general",
      // Mandatory ("error" or "warning")
      "severity": "error",
      // Optional: unique code for the cause of the error
      "errorCode": "3141",
      // Mandatory
      "message": "Invalid keyword",
      // Optional: the message formatted with source location
      "formattedMessage": "sourceFile.sol:100: Invalid keyword"
    }
  ],
  // This contains the file-level outputs.
  // It can be limited/filtered by the outputSelection settings.
  "sources": {
    "sourceFile.sol": {
      // Identifier of the source (used in source maps)
      "id": 1,
      // The AST object
      "ast": {},
      // The legacy AST object
      "legacyAST": {}
    }
  },
  // This contains the contract-level outputs.
  // It can be limited/filtered by the outputSelection settings.
  "contracts": {
    "sourceFile.sol": {
      // If the language used has no contract names, this field should equal to an
→empty string.
      "ContractName": {
        // The Ethereum Contract ABI. If empty, it is represented as an empty array.
        // See https://docs.soliditylang.org/en/develop/abi-spec.html
        "abi": [],
        // See the Metadata Output documentation (serialised JSON string)
        "metadata": "{...}",
        // User documentation (natspec)
        "userdoc": {},
        // Developer documentation (natspec)
        "devdoc": {},
        // Intermediate representation (string)
        "ir": "",
        // See the Storage Layout documentation.
        "storageLayout": {"storage": [...], "types": {...} },
        // EVM-related outputs
```

```
      "evm": {
        // Assembly (string)
        "assembly": "",
        // Old-style assembly (object)
        "legacyAssembly": {},
        // Bytecode and related details.
        "bytecode": {
          // The bytecode as a hex string.
          "object": "00fe",
          // Opcodes list (string)
          "opcodes": "",
          // The source mapping as a string. See the source mapping definition.
          "sourceMap": "",
          // Array of sources generated by the compiler. Currently only
          // contains a single Yul file.
          "generatedSources": [{
            // Yul AST
            "ast": { ... }
            // Source file in its text form (may contain comments)
            "contents":"{ function abi_decode(start, end) -> data { data :=␣
→calldataload(start) } }",
            // Source file ID, used for source references, same "namespace" as the␣
→Solidity source files
            "id": 2,
            "language": "Yul",
            "name": "#utility.yul"
          }]
          // If given, this is an unlinked object.
          "linkReferences": {
            "libraryFile.sol": {
              // Byte offsets into the bytecode.
              // Linking replaces the 20 bytes located there.
              "Library1": [
                { "start": 0, "length": 20 },
                { "start": 200, "length": 20 }
              ]
            }
          }
        },
        "deployedBytecode": {
          ..., // The same layout as above.
          "immutableReferences": {
            // There are two references to the immutable with AST ID 3, both 32␣
→bytes long. One is
            // at bytecode offset 42, the other at bytecode offset 80.
            "3": [{ "start": 42, "length": 32 }, { "start": 80, "length": 32 }]
          }
        },
        // The list of function hashes
        "methodIdentifiers": {
          "delegate(address)": "5c19a95c"
        },
        // Function gas estimates
        "gasEstimates": {
          "creation": {
            "codeDepositCost": "420000",
            "executionCost": "infinite",
```

```
          "totalCost": "infinite"
        },
        "external": {
          "delegate(address)": "25000"
        },
        "internal": {
          "heavyLifting()": "infinite"
        }
      }
    },
    // Ewasm related outputs
    "ewasm": {
      // S-expressions format
      "wast": "",
      // Binary format (hex string)
      "wasm": ""
    }
  }
 }
}
```

Error types

1. `JSONError`: JSON input doesn't conform to the required format, e.g. input is not a JSON object, the language is not supported, etc.

2. `IOError`: IO and import processing errors, such as unresolvable URL or hash mismatch in supplied sources.

3. `ParserError`: Source code doesn't conform to the language rules.

4. `DocstringParsingError`: The NatSpec tags in the comment block cannot be parsed.

5. `SyntaxError`: Syntactical error, such as `continue` is used outside of a `for` loop.

6. `DeclarationError`: Invalid, unresolvable or clashing identifier names. e.g. `Identifier not found`

7. `TypeError`: Error within the type system, such as invalid type conversions, invalid assignments, etc.

8. `UnimplementedFeatureError`: Feature is not supported by the compiler, but is expected to be supported in future versions.

9. `InternalCompilerError`: Internal bug triggered in the compiler - this should be reported as an issue.

10. `Exception`: Unknown failure during compilation - this should be reported as an issue.

11. `CompilerError`: Invalid use of the compiler stack - this should be reported as an issue.

12. `FatalError`: Fatal error not processed correctly - this should be reported as an issue.

13. `Warning`: A warning, which didn't stop the compilation, but should be addressed if possible.

3.26.4 Compiler tools

solidity-upgrade

`solidity-upgrade` can help you to semi-automatically upgrade your contracts to breaking language changes. While it does not and cannot implement all required changes for every breaking release, it still supports the ones, that

would need plenty of repetitive manual adjustments otherwise.

Note: `solidity-upgrade` carries out a large part of the work, but your contracts will most likely need further manual adjustments. We recommend using a version control system for your files. This helps reviewing and eventually rolling back the changes made.

> **Warning:** `solidity-upgrade` is not considered to be complete or free from bugs, so please use with care.

How it works

You can pass (a) Solidity source file(s) to `solidity-upgrade [files]`. If these make use of `import` statement which refer to files outside the current source file's directory, you need to specify directories that are allowed to read and import files from, by passing `--allow-paths [directory]`. You can ignore missing files by passing `--ignore-missing`.

`solidity-upgrade` is based on `libsolidity` and can parse, compile and analyse your source files, and might find applicable source upgrades in them.

Source upgrades are considered to be small textual changes to your source code. They are applied to an in-memory representation of the source files given. The corresponding source file is updated by default, but you can pass `--dry-run` to simulate to whole upgrade process without writing to any file.

The upgrade process itself has two phases. In the first phase source files are parsed, and since it is not possible to upgrade source code on that level, errors are collected and can be logged by passing `--verbose`. No source upgrades available at this point.

In the second phase, all sources are compiled and all activated upgrade analysis modules are run alongside compilation. By default, all available modules are activated. Please read the documentation on *available modules* for further details.

This can result in compilation errors that may be fixed by source upgrades. If no errors occur, no source upgrades are being reported and you're done. If errors occur and some upgrade module reported a source upgrade, the first reported one gets applied and compilation is triggered again for all given source files. The previous step is repeated as long as source upgrades are reported. If errors still occur, you can log them by passing `--verbose`. If no errors occur, your contracts are up to date and can be compiled with the latest version of the compiler.

Available upgrade modules

Module	Version	Description
`constructor`	0.5.0	Constructors must now be defined using the `constructor` keyword.
`visibility`	0.5.0	Explicit function visibility is now mandatory, defaults to `public`.
`abstract`	0.6.0	The keyword `abstract` has to be used if a contract does not implement all its functions.
`virtual`	0.6.0	Functions without implementation outside an interface have to be marked `virtual`.
`override`	0.6.0	When overriding a function or modifier, the new keyword `override` must be used.
`dotsyntax`	0.7.0	The following syntax is deprecated: `f.gas(...)()`, `f.value(...)()` and `(new C).value(...)()`. Replace these calls by `f{gas: ..., value: ...}()` and `(new C){value: ...}()`.
`now`	0.7.0	The `now` keyword is deprecated. Use `block.timestamp` instead.
`constructor-visibility`	0.7.0	Removes visibility of constructors.

Please read *0.5.0 release notes*, *0.6.0 release notes* and *0.7.0 release notes* for further details.

Synopsis

```
Usage: solidity-upgrade [options] contract.sol

Allowed options:
    --help                Show help message and exit.
    --version             Show version and exit.
    --allow-paths path(s)
                          Allow a given path for imports. A list of paths can be
                          supplied by separating them with a comma.
    --ignore-missing      Ignore missing files.
    --modules module(s)   Only activate a specific upgrade module. A list of
                          modules can be supplied by separating them with a comma.
    --dry-run             Apply changes in-memory only and don't write to input
                          file.
    --verbose             Print logs, errors and changes. Shortens output of
                          upgrade patches.
    --unsafe              Accept *unsafe* changes.
```

Bug Reports / Feature requests

If you found a bug or if you have a feature request, please file an issue on Github.

Example

Assume that you have the following contract in `Source.sol`:

```solidity
pragma solidity >=0.6.0 <0.6.4;
// This will not compile after 0.7.0
// SPDX-License-Identifier: GPL-3.0
contract C {
    // FIXME: remove constructor visibility and make the contract abstract
    constructor() internal {}
}

contract D {
    uint time;

    function f() public payable {
        // FIXME: change now to block.timestamp
        time = now;
    }
}

contract E {
    D d;

    // FIXME: remove constructor visibility
    constructor() public {}

    function g() public {
```

```
        // FIXME: change .value(5) =>  {value: 5}
        d.f.value(5)();
    }
}
```

Required changes

The above contract will not compile starting from 0.7.0. To bring the contract up to date with the current Solidity version, the following upgrade modules have to be executed: `constructor-visibility`, `now` and `dotsyntax`. Please read the documentation on *available modules* for further details.

Running the upgrade

It is recommended to explicitly specify the upgrade modules by using `--modules` argument.

```
$ solidity-upgrade --modules constructor-visibility,now,dotsyntax Source.sol
```

The command above applies all changes as shown below. Please review them carefully (the pragmas will have to be updated manually.)

```solidity
pragma solidity ^0.7.0;
// SPDX-License-Identifier: GPL-3.0
abstract contract C {
    // FIXME: remove constructor visibility and make the contract abstract
    constructor() {}
}

contract D {
    uint time;

    function f() public payable {
        // FIXME: change now to block.timestamp
        time = block.timestamp;
    }
}

contract E {
    D d;

    // FIXME: remove constructor visibility
    constructor() {}

    function g() public {
        // FIXME: change .value(5) =>  {value: 5}
        d.f{value: 5}();
    }
}
```

3.27 Yul

Yul (previously also called JULIA or IULIA) is an intermediate language that can be compiled to bytecode for different backends.

Support for EVM 1.0, EVM 1.5 and Ewasm is planned, and it is designed to be a usable common denominator of all three platforms. It can already be used in stand-alone mode and for "inline assembly" inside Solidity and there is an experimental implementation of the Solidity compiler that uses Yul as an intermediate language. Yul is a good target for high-level optimisation stages that can benefit all target platforms equally.

3.27.1 Motivation and High-level Description

The design of Yul tries to achieve several goals:

1. Programs written in Yul should be readable, even if the code is generated by a compiler from Solidity or another high-level language.

2. Control flow should be easy to understand to help in manual inspection, formal verification and optimization.

3. The translation from Yul to bytecode should be as straightforward as possible.

4. Yul should be suitable for whole-program optimization.

In order to achieve the first and second goal, Yul provides high-level constructs like `for` loops, `if` and `switch` statements and function calls. These should be sufficient for adequately representing the control flow for assembly programs. Therefore, no explicit statements for `SWAP`, `DUP`, `JUMPDEST`, `JUMP` and `JUMPI` are provided, because the first two obfuscate the data flow and the last two obfuscate control flow. Furthermore, functional statements of the form `mul(add(x, y), 7)` are preferred over pure opcode statements like `7 y x add mul` because in the first form, it is much easier to see which operand is used for which opcode.

Even though it was designed for stack machines, Yul does not expose the complexity of the stack itself. The programmer or auditor should not have to worry about the stack.

The third goal is achieved by compiling the higher level constructs to bytecode in a very regular way. The only non-local operation performed by the assembler is name lookup of user-defined identifiers (functions, variables, . . .) and cleanup of local variables from the stack.

To avoid confusions between concepts like values and references, Yul is statically typed. At the same time, there is a default type (usually the integer word of the target machine) that can always be omitted to help readability.

To keep the language simple and flexible, Yul does not have any built-in operations, functions or types in its pure form. These are added together with their semantics when specifying a dialect of Yul, which allows specializing Yul to the requirements of different target platforms and feature sets.

Currently, there is only one specified dialect of Yul. This dialect uses the EVM opcodes as builtin functions (see below) and defines only the type `u256`, which is the native 256-bit type of the EVM. Because of that, we will not provide types in the examples below.

3.27.2 Simple Example

The following example program is written in the EVM dialect and computes exponentiation. It can be compiled using `solc --strict-assembly`. The builtin functions `mul` and `div` compute product and division, respectively.

```
{
    function power(base, exponent) -> result
    {
        switch exponent
```

```
        case 0 { result := 1 }
        case 1 { result := base }
        default
        {
            result := power(mul(base, base), div(exponent, 2))
            switch mod(exponent, 2)
                case 1 { result := mul(base, result) }
        }
    }
}
```

It is also possible to implement the same function using a for-loop instead of with recursion. Here, `lt(a, b)` computes whether `a` is less than `b`. less-than comparison.

```
{
    function power(base, exponent) -> result
    {
        result := 1
        for { let i := 0 } lt(i, exponent) { i := add(i, 1) }
        {
            result := mul(result, base)
        }
    }
}
```

At the *end of the section*, a complete implementation of the ERC-20 standard can be found.

3.27.3 Stand-Alone Usage

You can use Yul in its stand-alone form in the EVM dialect using the Solidity compiler. This will use the *Yul object notation* so that it is possible to refer to code as data to deploy contracts. This Yul mode is available for the commandline compiler (use `--strict-assembly`) and for the *standard-json interface*:

```
{
    "language": "Yul",
    "sources": { "input.yul": { "content": "{ sstore(0, 1) }" } },
    "settings": {
        "outputSelection": { "*": { "*": ["*"], "": [ "*" ] } },
        "optimizer": { "enabled": true, "details": { "yul": true } }
    }
}
```

> **Warning:** Yul is in active development and bytecode generation is only fully implemented for the EVM dialect of Yul with EVM 1.0 as target.

3.27.4 Informal Description of Yul

In the following, we will talk about each individual aspect of the Yul language. In examples, we will use the default EVM dialect.

Syntax

Yul parses comments, literals and identifiers in the same way as Solidity, so you can e.g. use `//` and `/* */` to denote comments. There is one exception: Identifiers in Yul can contain dots: `.`.

Yul can specify "objects" that consist of code, data and sub-objects. Please see *Yul Objects* below for details on that. In this section, we are only concerned with the code part of such an object. This code part always consists of a curly-braces delimited block. Most tools support specifying just a code block where an object is expected.

Inside a code block, the following elements can be used (see the later sections for more details):

- literals, i.e. `0x123`, `42` or `"abc"` (strings up to 32 characters)
- calls to builtin functions, e.g. `add(1, mload(0))`
- variable declarations, e.g. `let x := 7`, `let x := add(y, 3)` or `let x` (initial value of 0 is assigned)
- identifiers (variables), e.g. `add(3, x)`
- assignments, e.g. `x := add(y, 3)`
- blocks where local variables are scoped inside, e.g. `{ let x := 3 { let y := add(x, 1) } }`
- if statements, e.g. `if lt(a, b) { sstore(0, 1) }`
- switch statements, e.g. `switch mload(0) case 0 { revert() } default { mstore(0, 1) }`
- for loops, e.g. `for { let i := 0} lt(i, 10) { i := add(i, 1) } { mstore(i, 7) }`
- function definitions, e.g. `function f(a, b) -> c { c := add(a, b) }`\`

Multiple syntactical elements can follow each other simply separated by whitespace, i.e. there is no terminating `;` or newline required.

Literals

You can use integer constants in decimal or hexadecimal notation. When compiling for the EVM, this will be translated into an appropriate `PUSHi` instruction. In the following example, `3` and `2` are added resulting in 5 and then the bitwise `and` with the string "abc" is computed. The final value is assigned to a local variable called `x`.

Strings are stored left-aligned and cannot be longer than 32 bytes. The limit does not apply to string literals passed to builtin functions that require literal arguments (e.g. `setimmutable` or `loadimmutable`). Those strings never end up in the generated bytecode.

```
let x := and("abc", add(3, 2))
```

Unless it is the default type, the type of a literal has to be specified after a colon:

```
let x := and("abc":uint32, add(3:uint256, 2:uint256))
```

Function Calls

Both built-in and user-defined functions (see below) can be called in the same way as shown in the previous example. If the function returns a single value, it can be directly used inside an expression again. If it returns multiple values, they have to be assigned to local variables.

```
mstore(0x80, add(mload(0x80), 3))
// Here, the user-defined function `f` returns
// two values. The definition of the function
// is missing from the example.
let x, y := f(1, mload(0))
```

For built-in functions of the EVM, functional expressions can be directly translated to a stream of opcodes: You just read the expression from right to left to obtain the opcodes. In the case of the first line in the example, this is `PUSH1 3 PUSH1 0x80 MLOAD ADD PUSH1 0x80 MSTORE`.

For calls to user-defined functions, the arguments are also put on the stack from right to left and this is the order in which argument lists are evaluated. The return values, though, are expected on the stack from left to right, i.e. in this example, `y` is on top of the stack and `x` is below it.

Variable Declarations

You can use the `let` keyword to declare variables. A variable is only visible inside the `{...}`-block it was defined in. When compiling to the EVM, a new stack slot is created that is reserved for the variable and automatically removed again when the end of the block is reached. You can provide an initial value for the variable. If you do not provide a value, the variable will be initialized to zero.

Since variables are stored on the stack, they do not directly influence memory or storage, but they can be used as pointers to memory or storage locations in the built-in functions `mstore`, `mload`, `sstore` and `sload`. Future dialects might introduce specific types for such pointers.

When a variable is referenced, its current value is copied. For the EVM, this translates to a `DUP` instruction.

```
{
    let zero := 0
    let v := calldataload(zero)
    {
        let y := add(sload(v), 1)
        v := y
    } // y is "deallocated" here
    sstore(v, zero)
} // v and zero are "deallocated" here
```

If the declared variable should have a type different from the default type, you denote that following a colon. You can also declare multiple variables in one statement when you assign from a function call that returns multiple values.

```
{
    let zero:uint32 := 0:uint32
    let v:uint256, t:uint32 := f()
    let x, y := g()
}
```

Depending on the optimiser settings, the compiler can free the stack slots already after the variable has been used for the last time, even though it is still in scope.

Assignments

Variables can be assigned to after their definition using the `:=` operator. It is possible to assign multiple variables at the same time. For this, the number and types of the values have to match. If you want to assign the values returned from a function that has multiple return parameters, you have to provide multiple variables. The same variable may not occur multiple times on the left-hand side of an assignment, e.g. `x, x := f()` is invalid.

```
let v := 0
// re-assign v
v := 2
let t := add(v, 2)
function f() -> a, b { }
// assign multiple values
v, t := f()
```

If

The if statement can be used for conditionally executing code. No "else" block can be defined. Consider using "switch" instead (see below) if you need multiple alternatives.

```
if eq(value, 0) { revert(0, 0) }
```

The curly braces for the body are required.

Switch

You can use a switch statement as an extended version of the if statement. It takes the value of an expression and compares it to several literal constants. The branch corresponding to the matching constant is taken. Contrary to other programming languages, for safety reasons, control flow does not continue from one case to the next. There can be a fallback or default case called `default` which is taken if none of the literal constants matches.

```
{
    let x := 0
    switch calldataload(4)
    case 0 {
        x := calldataload(0x24)
    }
    default {
        x := calldataload(0x44)
    }
    sstore(0, div(x, 2))
}
```

The list of cases is not enclosed by curly braces, but the body of a case does require them.

Loops

Yul supports for-loops which consist of a header containing an initializing part, a condition, a post-iteration part and a body. The condition has to be an expression, while the other three are blocks. If the initializing part declares any variables at the top level, the scope of these variables extends to all other parts of the loop.

The `break` and `continue` statements can be used in the body to exit the loop or skip to the post-part, respectively.

The following example computes the sum of an area in memory.

```
{
    let x := 0
    for { let i := 0 } lt(i, 0x100) { i := add(i, 0x20) } {
        x := add(x, mload(i))
    }
}
```

For loops can also be used as a replacement for while loops: Simply leave the initialization and post-iteration parts empty.

```
{
    let x := 0
    let i := 0
    for { } lt(i, 0x100) { } {      // while(i < 0x100)
        x := add(x, mload(i))
        i := add(i, 0x20)
    }
}
```

Function Declarations

Yul allows the definition of functions. These should not be confused with functions in Solidity since they are never part of an external interface of a contract and are part of a namespace separate from the one for Solidity functions.

For the EVM, Yul functions take their arguments (and a return PC) from the stack and also put the results onto the stack. User-defined functions and built-in functions are called in exactly the same way.

Functions can be defined anywhere and are visible in the block they are declared in. Inside a function, you cannot access local variables defined outside of that function.

Functions declare parameters and return variables, similar to Solidity. To return a value, you assign it to the return variable(s).

If you call a function that returns multiple values, you have to assign them to multiple variables using `a, b := f(x)` or `let a, b := f(x)`.

The `leave` statement can be used to exit the current function. It works like the `return` statement in other languages just that it does not take a value to return, it just exits the functions and the function will return whatever values are currently assigned to the return variable(s).

Note that the EVM dialect has a built-in function called `return` that quits the full execution context (internal message call) and not just the current yul function.

The following example implements the power function by square-and-multiply.

```
{
    function power(base, exponent) -> result {
        switch exponent
        case 0 { result := 1 }
        case 1 { result := base }
        default {
            result := power(mul(base, base), div(exponent, 2))
            switch mod(exponent, 2)
                case 1 { result := mul(base, result) }
        }
    }
}
```

3.27.5 Specification of Yul

This chapter describes Yul code formally. Yul code is usually placed inside Yul objects, which are explained in their own chapter.

```
Block = '{' Statement* '}'
Statement =
    Block |
    FunctionDefinition |
    VariableDeclaration |
    Assignment |
    If |
    Expression |
    Switch |
    ForLoop |
    BreakContinue |
    Leave
FunctionDefinition =
    'function' Identifier '(' TypedIdentifierList? ')'
    ( '->' TypedIdentifierList )? Block
VariableDeclaration =
    'let' TypedIdentifierList ( ':=' Expression )?
Assignment =
    IdentifierList ':=' Expression
Expression =
    FunctionCall | Identifier | Literal
If =
    'if' Expression Block
Switch =
    'switch' Expression ( Case+ Default? | Default )
Case =
    'case' Literal Block
Default =
    'default' Block
ForLoop =
    'for' Block Expression Block Block
BreakContinue =
    'break' | 'continue'
Leave = 'leave'
FunctionCall =
    Identifier '(' ( Expression ( ',' Expression )* )? ')'
Identifier = [a-zA-Z_$] [a-zA-Z_$0-9.]*
IdentifierList = Identifier ( ',' Identifier)*
TypeName = Identifier
TypedIdentifierList = Identifier ( ':' TypeName )? ( ',' Identifier ( ':' TypeName )?␣
↪)*
Literal =
    (NumberLiteral | StringLiteral | TrueLiteral | FalseLiteral) ( ':' TypeName )?
NumberLiteral = HexNumber | DecimalNumber
StringLiteral = '"' ([^"\r\n\\] | '\\' .)* '"'
TrueLiteral = 'true'
FalseLiteral = 'false'
HexNumber = '0x' [0-9a-fA-F]+
DecimalNumber = [0-9]+
```

Restrictions on the Grammar

Apart from those directly imposed by the grammar, the following restrictions apply:

Switches must have at least one case (including the default case). All case values need to have the same type and distinct values. If all possible values of the expression type are covered, a default case is not allowed (i.e. a switch with a `bool` expression that has both a true and a false case do not allow a default case).

Every expression evaluates to zero or more values. Identifiers and Literals evaluate to exactly one value and function calls evaluate to a number of values equal to the number of return variables of the function called.

In variable declarations and assignments, the right-hand-side expression (if present) has to evaluate to a number of values equal to the number of variables on the left-hand-side. This is the only situation where an expression evaluating to more than one value is allowed. The same variable name cannot occur more than once in the left-hand-side of an assignment or variable declaration.

Expressions that are also statements (i.e. at the block level) have to evaluate to zero values.

In all other situations, expressions have to evaluate to exactly one value.

The `continue` and `break` statements can only be used inside loop bodies and have to be in the same function as the loop (or both have to be at the top level). The `continue` and `break` statements cannot be used in other parts of a loop, not even when it is scoped inside a second loop's body.

The condition part of the for-loop has to evaluate to exactly one value.

The `leave` statement can only be used inside a function.

Functions cannot be defined anywhere inside for loop init blocks.

Literals cannot be larger than their type. The largest type defined is 256-bit wide.

During assignments and function calls, the types of the respective values have to match. There is no implicit type conversion. Type conversion in general can only be achieved if the dialect provides an appropriate built-in function that takes a value of one type and returns a value of a different type.

Scoping Rules

Scopes in Yul are tied to Blocks (exceptions are functions and the for loop as explained below) and all declarations (`FunctionDefinition`, `VariableDeclaration`) introduce new identifiers into these scopes.

Identifiers are visible in the block they are defined in (including all sub-nodes and sub-blocks).

As an exception, the scope of the "init" part of the or-loop (the first block) extends across all other parts of the for loop. This means that variables declared in the init part (but not inside a block inside the init part) are visible in all other parts of the for-loop.

Identifiers declared in the other parts of the for loop respect the regular syntactical scoping rules.

This means a for-loop of the form `for { I... } C { P... } { B... }` is equivalent to `{ I... for {} C { P... } { B... } }`.

The parameters and return parameters of functions are visible in the function body and their names have to be distinct.

Variables can only be referenced after their declaration. In particular, variables cannot be referenced in the right hand side of their own variable declaration. Functions can be referenced already before their declaration (if they are visible).

Shadowing is disallowed, i.e. you cannot declare an identifier at a point where another identifier with the same name is also visible, even if it is not accessible.

Inside functions, it is not possible to access a variable that was declared outside of that function.

Formal Specification

We formally specify Yul by providing an evaluation function E overloaded on the various nodes of the AST. As builtin functions can have side effects, E takes two state objects and the AST node and returns two new state objects and a variable number of other values. The two state objects are the global state object (which in the context of the EVM is the memory, storage and state of the blockchain) and the local state object (the state of local variables, i.e. a segment of the stack in the EVM).

If the AST node is a statement, E returns the two state objects and a "mode", which is used for the `break`, `continue` and `leave` statements. If the AST node is an expression, E returns the two state objects and as many values as the expression evaluates to.

The exact nature of the global state is unspecified for this high level description. The local state `L` is a mapping of identifiers `i` to values `v`, denoted as `L[i] = v`.

For an identifier `v`, let `$v` be the name of the identifier.

We will use a destructuring notation for the AST nodes.

```
E(G, L, <{St1, ..., Stn}>: Block) =
    let G1, L1, mode = E(G, L, St1, ..., Stn)
    let L2 be a restriction of L1 to the identifiers of L
    G1, L2, mode
E(G, L, St1, ..., Stn: Statement) =
    if n is zero:
        G, L, regular
    else:
        let G1, L1, mode = E(G, L, St1)
        if mode is regular then
            E(G1, L1, St2, ..., Stn)
        otherwise
            G1, L1, mode
E(G, L, FunctionDefinition) =
    G, L, regular
E(G, L, <let var_1, ..., var_n := rhs>: VariableDeclaration) =
    E(G, L, <var_1, ..., var_n := rhs>: Assignment)
E(G, L, <let var_1, ..., var_n>: VariableDeclaration) =
    let L1 be a copy of L where L1[$var_i] = 0 for i = 1, ..., n
    G, L1, regular
E(G, L, <var_1, ..., var_n := rhs>: Assignment) =
    let G1, L1, v1, ..., vn = E(G, L, rhs)
    let L2 be a copy of L1 where L2[$var_i] = vi for i = 1, ..., n
    G, L2, regular
E(G, L, <for { i1, ..., in } condition post body>: ForLoop) =
    if n >= 1:
        let G1, L, mode = E(G, L, i1, ..., in)
        // mode has to be regular or leave due to the syntactic restrictions
        if mode is leave then
            G1, L1 restricted to variables of L, leave
        otherwise
            let G2, L2, mode = E(G1, L1, for {} condition post body)
            G2, L2 restricted to variables of L, mode
    else:
        let G1, L1, v = E(G, L, condition)
        if v is false:
            G1, L1, regular
        else:
            let G2, L2, mode = E(G1, L, body)
            if mode is break:
                G2, L2, regular
            otherwise if mode is leave:
                G2, L2, leave
            else:
                G3, L3, mode = E(G2, L2, post)
                if mode is leave:
                    G2, L3, leave
                otherwise
```

```
                        E(G3, L3, for {} condition post body)
E(G, L, break: BreakContinue) =
    G, L, break
E(G, L, continue: BreakContinue) =
    G, L, continue
E(G, L, leave: Leave) =
    G, L, leave
E(G, L, <if condition body>: If) =
    let G0, L0, v = E(G, L, condition)
    if v is true:
        E(G0, L0, body)
    else:
        G0, L0, regular
E(G, L, <switch condition case l1:t1 st1 ... case ln:tn stn>: Switch) =
    E(G, L, switch condition case l1:t1 st1 ... case ln:tn stn default {})
E(G, L, <switch condition case l1:t1 st1 ... case ln:tn stn default st'>: Switch) =
    let G0, L0, v = E(G, L, condition)
    // i = 1 .. n
    // Evaluate literals, context doesn't matter
    let _, _, v1 = E(G0, L0, l1)
    ...
    let _, _, vn = E(G0, L0, ln)
    if there exists smallest i such that vi = v:
        E(G0, L0, sti)
    else:
        E(G0, L0, st')

E(G, L, <name>: Identifier) =
    G, L, L[$name]
E(G, L, <fname(arg1, ..., argn)>: FunctionCall) =
    G1, L1, vn = E(G, L, argn)
    ...
    G(n-1), L(n-1), v2 = E(G(n-2), L(n-2), arg2)
    Gn, Ln, v1 = E(G(n-1), L(n-1), arg1)
    Let <function fname (param1, ..., paramn) -> ret1, ..., retm block>
    be the function of name $fname visible at the point of the call.
    Let L' be a new local state such that
    L'[$parami] = vi and L'[$reti] = 0 for all i.
    Let G'', L'', mode = E(Gn, L', block)
    G'', Ln, L''[$ret1], ..., L''[$retm]
E(G, L, l: StringLiteral) = G, L, utf8EncodeLeftAligned(l),
    where utf8EncodeLeftAligned performs a utf8 encoding of l
    and aligns it left into 32 bytes
E(G, L, n: HexNumber) = G, L, hex(n)
    where hex is the hexadecimal decoding function
E(G, L, n: DecimalNumber) = G, L, dec(n),
    where dec is the decimal decoding function
```

EVM Dialect

The default dialect of Yul currently is the EVM dialect for the currently selected version of the EVM. with a version of the EVM. The only type available in this dialect is u256, the 256-bit native type of the Ethereum Virtual Machine. Since it is the default type of this dialect, it can be omitted.

The following table lists all builtin functions (depending on the EVM version) and provides a short description of the semantics of the function / opcode. This document does not want to be a full description of the Ethereum virtual

machine. Please refer to a different document if you are interested in the precise semantics.

Opcodes marked with − do not return a result and all others return exactly one value. Opcodes marked with F, H, B, C or I are present since Frontier, Homestead, Byzantium, Constantinople or Istanbul, respectively.

In the following, mem[a...b) signifies the bytes of memory starting at position a up to but not including position b and storage[p] signifies the storage contents at slot p.

Since Yul manages local variables and control-flow, opcodes that interfere with these features are not available. This includes the dup and swap instructions as well as jump instructions, labels and the push instructions.

Instruction			Explanation
stop()	-	F	stop execution, identical to return(0, 0)
add(x, y)		F	x + y
sub(x, y)		F	x - y
mul(x, y)		F	x * y
div(x, y)		F	x / y or 0 if y == 0
sdiv(x, y)		F	x / y, for signed numbers in two's complement, 0 if y == 0
mod(x, y)		F	x % y, 0 if y == 0
smod(x, y)		F	x % y, for signed numbers in two's complement, 0 if y == 0
exp(x, y)		F	x to the power of y
not(x)		F	bitwise "not" of x (every bit of x is negated)
lt(x, y)		F	1 if x < y, 0 otherwise
gt(x, y)		F	1 if x > y, 0 otherwise
slt(x, y)		F	1 if x < y, 0 otherwise, for signed numbers in two's complement
sgt(x, y)		F	1 if x > y, 0 otherwise, for signed numbers in two's complement
eq(x, y)		F	1 if x == y, 0 otherwise
iszero(x)		F	1 if x == 0, 0 otherwise
and(x, y)		F	bitwise "and" of x and y
or(x, y)		F	bitwise "or" of x and y
xor(x, y)		F	bitwise "xor" of x and y
byte(n, x)		F	nth byte of x, where the most significant byte is the 0th byte
shl(x, y)		C	logical shift left y by x bits
shr(x, y)		C	logical shift right y by x bits
sar(x, y)		C	signed arithmetic shift right y by x bits
addmod(x, y, m)		F	(x + y) % m with arbitrary precision arithmetic, 0 if m == 0
mulmod(x, y, m)		F	(x * y) % m with arbitrary precision arithmetic, 0 if m == 0
signextend(i, x)		F	sign extend from (i*8+7)th bit counting from least significant
keccak256(p, n)		F	keccak(mem[p...(p+n)))
pc()		F	current position in code
pop(x)	-	F	discard value x
mload(p)		F	mem[p...(p+32))
mstore(p, v)	-	F	mem[p...(p+32)) := v
mstore8(p, v)	-	F	mem[p] := v & 0xff (only modifies a single byte)
sload(p)		F	storage[p]
sstore(p, v)	-	F	storage[p] := v
msize()		F	size of memory, i.e. largest accessed memory index
gas()		F	gas still available to execution
address()		F	address of the current contract / execution context
balance(a)		F	wei balance at address a
selfbalance()		I	equivalent to balance(address()), but cheaper
caller()		F	call sender (excluding delegatecall)

Instruction			Explanation
callvalue()		F	wei sent together with the current call
calldataload(p)		F	call data starting from position p (32 bytes)
calldatasize()		F	size of call data in bytes
calldatacopy(t, f, s)	-	F	copy s bytes from calldata at position f to mem at position t
codesize()		F	size of the code of the current contract / execution context
codecopy(t, f, s)	-	F	copy s bytes from code at position f to mem at position t
extcodesize(a)		F	size of the code at address a
extcodecopy(a, t, f, s)	-	F	like codecopy(t, f, s) but take code at address a
returndatasize()		B	size of the last returndata
returndatacopy(t, f, s)	-	B	copy s bytes from returndata at position f to mem at position t
extcodehash(a)		C	code hash of address a
create(v, p, n)		F	create new contract with code mem[p. . . (p+n)) and send v wei and return the new
create2(v, p, n, s)		C	create new contract with code mem[p. . . (p+n)) at address keccak256(0xff . this . s
call(g, a, v, in, insize, out, outsize)		F	call contract at address a with input mem[in. . . (in+insize)) providing g gas and v
callcode(g, a, v, in, insize, out, outsize)		F	identical to `call` but only use the code from a and stay in the context of the curre
delegatecall(g, a, in, insize, out, outsize)		H	identical to `callcode` but also keep `caller` and `callvalue` *See more*
staticcall(g, a, in, insize, out, outsize)		B	identical to `call(g, a, 0, in, insize, out, outsize)` but do not
return(p, s)	-	F	end execution, return data mem[p. . . (p+s))
revert(p, s)	-	B	end execution, revert state changes, return data mem[p. . . (p+s))
selfdestruct(a)	-	F	end execution, destroy current contract and send funds to a
invalid()	-	F	end execution with invalid instruction
log0(p, s)	-	F	log without topics and data mem[p. . . (p+s))
log1(p, s, t1)	-	F	log with topic t1 and data mem[p. . . (p+s))
log2(p, s, t1, t2)	-	F	log with topics t1, t2 and data mem[p. . . (p+s))
log3(p, s, t1, t2, t3)	-	F	log with topics t1, t2, t3 and data mem[p. . . (p+s))
log4(p, s, t1, t2, t3, t4)	-	F	log with topics t1, t2, t3, t4 and data mem[p. . . (p+s))
chainid()		I	ID of the executing chain (EIP 1344)
origin()		F	transaction sender
gasprice()		F	gas price of the transaction
blockhash(b)		F	hash of block nr b - only for last 256 blocks excluding current
coinbase()		F	current mining beneficiary
timestamp()		F	timestamp of the current block in seconds since the epoch
number()		F	current block number
difficulty()		F	difficulty of the current block
gaslimit()		F	block gas limit of the current block

Note: The `call*` instructions use the `out` and `outsize` parameters to define an area in memory where the return data is placed. This area is written to depending on how many bytes the called contract returns. If it returns more data, only the first `outsize` bytes are written. You can access the rest of the data using the `returndatacopy` opcode. If it returns less data, then the remaining bytes are not touched at all. You need to use the `returndatasize` opcode to check which part of this memory area contains the return data. The remaining bytes will retain their values as of before the call. If the call fails (it returns `0`), nothing is written to that area, but you can still retrieve the failure data using `returndatacopy`.

In some internal dialects, there are additional functions:

datasize, dataoffset, datacopy

The functions `datasize(x)`, `dataoffset(x)` and `datacopy(t, f, l)` are used to access other parts of a Yul object.

`datasize` and `dataoffset` can only take string literals (the names of other objects) as arguments and return the size and offset in the data area, respectively. For the EVM, the `datacopy` function is equivalent to `codecopy`.

setimmutable, loadimmutable

The functions `setimmutable("name", value)` and `loadimmutable("name")` are used for the immutable mechanism in Solidity and do not nicely map to pure Yul. The function `setimmutable` assumes that the runtime code of a contract is currently copied to memory at offset zero. The call to `setimmutable("name", value)` will store `value` at all points in memory that contain a call to `loadimmutable("name")`.

linkersymbol

The function `linkersymbol("fq_library_name")` is a placeholder for an address literal to be substituted by the linker. Its first and only argument must be a string literal and represents the fully qualified library name used with the `--libraries` option.

For example this code

```
let a := linkersymbol("file.sol:Math")
```

is equivalent to

```
let a := 0x1234567890123456789012345678901234567890
```

when the linker is invoked with `--libraries "file.sol:Math:0x1234567890123456789012345678901234567890` option.

See *Using the Commandline Compiler* for details about the Solidity linker.

memoryguard

This function is available in the EVM dialect with objects. The caller of `let ptr := memoryguard(size)` (where `size` has to be a literal number) promises that they only use memory in either the range `[0, size)` or the unbounded range starting at `ptr`.

Since the presence of a `memoryguard` call indicates that all memory access adheres to this restriction, it allows the optimizer to perform additional optimization steps, for example the stack limit evader, which attempts to move stack variables that would otherwise be unreachable to memory.

The Yul optimizer promises to only use the memory range `[size, ptr)` for its purposes. If the optimizer does not need to reserve any memory, it holds that `ptr == size`.

`memoryguard` can be called multiple times, but needs to have the same literal as argument within one Yul subobject. If at least one `memoryguard` call is found in a subobject, the additional optimiser steps will be run on it.

3.27.6 Specification of Yul Object

Yul objects are used to group named code and data sections. The functions `datasize`, `dataoffset` and `datacopy` can be used to access these sections from within code. Hex strings can be used to specify data in hex encoding, regular strings in native encoding. For code, `datacopy` will access its assembled binary representation.

```
Object = 'object' StringLiteral '{' Code ( Object | Data )* '}'
Code = 'code' Block
Data = 'data' StringLiteral ( HexLiteral | StringLiteral )
HexLiteral = 'hex' ('"' ([0-9a-fA-F]{2})* '"' | '\'' ([0-9a-fA-F]{2})* '\'')
StringLiteral = '"' ([^"\r\n\\] | '\\' .)* '"'
```

Above, `Block` refers to `Block` in the Yul code grammar explained in the previous chapter.

An example Yul Object is shown below:

```
// A contract consists of a single object with sub-objects representing
// the code to be deployed or other contracts it can create.
// The single "code" node is the executable code of the object.
// Every (other) named object or data section is serialized and
// made accessible to the special built-in functions datacopy / dataoffset / datasize
// The current object, sub-objects and data items inside the current object
// are in scope.
object "Contract1" {
    // This is the constructor code of the contract.
    code {
        function allocate(size) -> ptr {
            ptr := mload(0x40)
            if iszero(ptr) { ptr := 0x60 }
            mstore(0x40, add(ptr, size))
        }

        // first create "Contract2"
        let size := datasize("Contract2")
        let offset := allocate(size)
        // This will turn into codecopy for EVM
        datacopy(offset, dataoffset("Contract2"), size)
        // constructor parameter is a single number 0x1234
        mstore(add(offset, size), 0x1234)
        pop(create(offset, add(size, 32), 0))

        // now return the runtime object (the currently
        // executing code is the constructor code)
        size := datasize("runtime")
        offset := allocate(size)
        // This will turn into a memory->memory copy for Ewasm and
        // a codecopy for EVM
        datacopy(offset, dataoffset("runtime"), size)
        return(offset, size)
    }

    data "Table2" hex"4123"

    object "runtime" {
        code {
            function allocate(size) -> ptr {
                ptr := mload(0x40)
                if iszero(ptr) { ptr := 0x60 }
```

```
                mstore(0x40, add(ptr, size))
            }

            // runtime code

            mstore(0, "Hello, World!")
            return(0, 0x20)
        }
    }

    // Embedded object. Use case is that the outside is a factory contract,
    // and Contract2 is the code to be created by the factory
    object "Contract2" {
        code {
            // code here ...
        }

        object "runtime" {
            code {
                // code here ...
            }
        }

        data "Table1" hex"4123"
    }
}
```

3.27.7 Yul Optimizer

The Yul optimizer operates on Yul code and uses the same language for input, output and intermediate states. This allows for easy debugging and verification of the optimizer.

Please see the documentation in the source code for more details about its internals.

If you want to use Solidity in stand-alone Yul mode, you activate the optimizer using `--optimize`:

```
solc --strict-assembly --optimize
```

In Solidity mode, the Yul optimizer is activated together with the regular optimizer.

Optimization step sequence

By default the Yul optimizer applies its predefined sequence of optimization steps to the generated assembly. You can override this sequence and supply your own using the `--yul-optimizations` option:

```
solc --optimize --ir-optimized --yul-optimizations 'dhfoD[xarrscLMcCTU]uljmul'
```

The order of steps is significant and affects the quality of the output. Moreover, applying a step may uncover new optimization opportunities for others that were already applied so repeating steps is often beneficial. By enclosing part of the sequence in square brackets (`[]`) you tell the optimizer to repeatedly apply that part until it no longer improves the size of the resulting assembly. You can use brackets multiple times in a single sequence but they cannot be nested.

The following optimization steps are available:

Abbreviation	Full name
f	BlockFlattener
l	CircularReferencesPruner
c	CommonSubexpressionEliminator
C	ConditionalSimplifier
U	ConditionalUnsimplifier
n	ControlFlowSimplifier
D	DeadCodeEliminator
v	EquivalentFunctionCombiner
e	ExpressionInliner
j	ExpressionJoiner
s	ExpressionSimplifier
x	ExpressionSplitter
I	ForLoopConditionIntoBody
O	ForLoopConditionOutOfBody
o	ForLoopInitRewriter
i	FullInliner
g	FunctionGrouper
h	FunctionHoister
T	LiteralRematerialiser
L	LoadResolver
M	LoopInvariantCodeMotion
r	RedundantAssignEliminator
R	ReasoningBasedSimplifier - highly experimental
m	Rematerialiser
V	SSAReverser
a	SSATransform
t	StructuralSimplifier
u	UnusedPruner
d	VarDeclInitializer

Some steps depend on properties ensured by `BlockFlattener`, `FunctionGrouper`, `ForLoopInitRewriter`. For this reason the Yul optimizer always applies them before applying any steps supplied by the user.

The ReasoningBasedSimplifier is an optimizer step that is currently not enabled in the default set of steps. It uses an SMT solver to simplify arithmetic expressions and boolean conditions. It has not received thorough testing or validation yet and can produce non-reproducible results, so please use with care!

3.27.8 Complete ERC20 Example

```
object "Token" {
    code {
        // Store the creator in slot zero.
        sstore(0, caller())

        // Deploy the contract
        datacopy(0, dataoffset("runtime"), datasize("runtime"))
        return(0, datasize("runtime"))
    }
    object "runtime" {
```

(continues on next page)

```
code {
    // Protection against sending Ether
    require(iszero(callvalue()))

    // Dispatcher
    switch selector()
    case 0x70a08231 /* "balanceOf(address)" */ {
        returnUint(balanceOf(decodeAsAddress(0)))
    }
    case 0x18160ddd /* "totalSupply()" */ {
        returnUint(totalSupply())
    }
    case 0xa9059cbb /* "transfer(address,uint256)" */ {
        transfer(decodeAsAddress(0), decodeAsUint(1))
        returnTrue()
    }
    case 0x23b872dd /* "transferFrom(address,address,uint256)" */ {
        transferFrom(decodeAsAddress(0), decodeAsAddress(1), decodeAsUint(2))
        returnTrue()
    }
    case 0x095ea7b3 /* "approve(address,uint256)" */ {
        approve(decodeAsAddress(0), decodeAsUint(1))
        returnTrue()
    }
    case 0xdd62ed3e /* "allowance(address,address)" */ {
        returnUint(allowance(decodeAsAddress(0), decodeAsAddress(1)))
    }
    case 0x40c10f19 /* "mint(address,uint256)" */ {
        mint(decodeAsAddress(0), decodeAsUint(1))
        returnTrue()
    }
    default {
        revert(0, 0)
    }

    function mint(account, amount) {
        require(calledByOwner())

        mintTokens(amount)
        addToBalance(account, amount)
        emitTransfer(0, account, amount)
    }
    function transfer(to, amount) {
        executeTransfer(caller(), to, amount)
    }
    function approve(spender, amount) {
        revertIfZeroAddress(spender)
        setAllowance(caller(), spender, amount)
        emitApproval(caller(), spender, amount)
    }
    function transferFrom(from, to, amount) {
        decreaseAllowanceBy(from, caller(), amount)
        executeTransfer(from, to, amount)
    }

    function executeTransfer(from, to, amount) {
        revertIfZeroAddress(to)
```

```
            deductFromBalance(from, amount)
            addToBalance(to, amount)
            emitTransfer(from, to, amount)
        }


        /* ---------- calldata decoding functions ---------- */
        function selector() -> s {
            s := div(calldataload(0),␣
→0x100000000000000000000000000000000000000000000000000000000)
        }

        function decodeAsAddress(offset) -> v {
            v := decodeAsUint(offset)
            if iszero(iszero(and(v,␣
→not(0xffffffffffffffffffffffffffffffffffffffff)))) {
                revert(0, 0)
            }
        }
        function decodeAsUint(offset) -> v {
            let pos := add(4, mul(offset, 0x20))
            if lt(calldatasize(), add(pos, 0x20)) {
                revert(0, 0)
            }
            v := calldataload(pos)
        }
        /* ---------- calldata encoding functions ---------- */
        function returnUint(v) {
            mstore(0, v)
            return(0, 0x20)
        }
        function returnTrue() {
            returnUint(1)
        }


        /* -------- events ---------- */
        function emitTransfer(from, to, amount) {
            let signatureHash :=␣
→0xddf252ad1be2c89b69c2b068fc378daa952ba7f163c4a11628f55a4df523b3ef
            emitEvent(signatureHash, from, to, amount)
        }
        function emitApproval(from, spender, amount) {
            let signatureHash :=␣
→0x8c5be1e5ebec7d5bd14f71427d1e84f3dd0314c0f7b2291e5b200ac8c7c3b925
            emitEvent(signatureHash, from, spender, amount)
        }
        function emitEvent(signatureHash, indexed1, indexed2, nonIndexed) {
            mstore(0, nonIndexed)
            log3(0, 0x20, signatureHash, indexed1, indexed2)
        }


        /* -------- storage layout ---------- */
        function ownerPos() -> p { p := 0 }
        function totalSupplyPos() -> p { p := 1 }
        function accountToStorageOffset(account) -> offset {
            offset := add(0x1000, account)
        }
```

```
        function allowanceStorageOffset(account, spender) -> offset {
            offset := accountToStorageOffset(account)
            mstore(0, offset)
            mstore(0x20, spender)
            offset := keccak256(0, 0x40)
        }

        /* -------- storage access ---------- */
        function owner() -> o {
            o := sload(ownerPos())
        }
        function totalSupply() -> supply {
            supply := sload(totalSupplyPos())
        }
        function mintTokens(amount) {
            sstore(totalSupplyPos(), safeAdd(totalSupply(), amount))
        }
        function balanceOf(account) -> bal {
            bal := sload(accountToStorageOffset(account))
        }
        function addToBalance(account, amount) {
            let offset := accountToStorageOffset(account)
            sstore(offset, safeAdd(sload(offset), amount))
        }
        function deductFromBalance(account, amount) {
            let offset := accountToStorageOffset(account)
            let bal := sload(offset)
            require(lte(amount, bal))
            sstore(offset, sub(bal, amount))
        }
        function allowance(account, spender) -> amount {
            amount := sload(allowanceStorageOffset(account, spender))
        }
        function setAllowance(account, spender, amount) {
            sstore(allowanceStorageOffset(account, spender), amount)
        }
        function decreaseAllowanceBy(account, spender, amount) {
            let offset := allowanceStorageOffset(account, spender)
            let currentAllowance := sload(offset)
            require(lte(amount, currentAllowance))
            sstore(offset, sub(currentAllowance, amount))
        }

        /* ---------- utility functions ---------- */
        function lte(a, b) -> r {
            r := iszero(gt(a, b))
        }
        function gte(a, b) -> r {
            r := iszero(lt(a, b))
        }
        function safeAdd(a, b) -> r {
            r := add(a, b)
            if or(lt(r, a), lt(r, b)) { revert(0, 0) }
        }
        function calledByOwner() -> cbo {
            cbo := eq(owner(), caller())
        }
```

```solidity
            function revertIfZeroAddress(addr) {
                require(addr)
            }
            function require(condition) {
                if iszero(condition) { revert(0, 0) }
            }
        }
    }
}
```

3.28 Style Guide

3.28.1 Introduction

This guide is intended to provide coding conventions for writing solidity code. This guide should be thought of as an evolving document that will change over time as useful conventions are found and old conventions are rendered obsolete.

Many projects will implement their own style guides. In the event of conflicts, project specific style guides take precedence.

The structure and many of the recommendations within this style guide were taken from python's pep8 style guide.

The goal of this guide is *not* to be the right way or the best way to write solidity code. The goal of this guide is *consistency*. A quote from python's pep8 captures this concept well.

Note: A style guide is about consistency. Consistency with this style guide is important. Consistency within a project is more important. Consistency within one module or function is most important.

But most importantly: **know when to be inconsistent** – sometimes the style guide just doesn't apply. When in doubt, use your best judgement. Look at other examples and decide what looks best. And don't hesitate to ask!

3.28.2 Code Layout

Indentation

Use 4 spaces per indentation level.

Tabs or Spaces

Spaces are the preferred indentation method.

Mixing tabs and spaces should be avoided.

Blank Lines

Surround top level declarations in solidity source with two blank lines.

Yes:

```
// SPDX-License-Identifier: GPL-3.0
pragma solidity >=0.4.0 <0.8.0;


contract A {
    // ...
}


contract B {
    // ...
}


contract C {
    // ...
}
```

No:

```
// SPDX-License-Identifier: GPL-3.0
pragma solidity >=0.4.0 <0.8.0;


contract A {
    // ...
}
contract B {
    // ...
}

contract C {
    // ...
}
```

Within a contract surround function declarations with a single blank line.

Blank lines may be omitted between groups of related one-liners (such as stub functions for an abstract contract)

Yes:

```
// SPDX-License-Identifier: GPL-3.0
pragma solidity >=0.6.0 <0.8.0;


abstract contract A {
    function spam() public virtual pure;
    function ham() public virtual pure;
}


contract B is A {
    function spam() public pure override {
        // ...
    }

    function ham() public pure override {
        // ...
    }
}
```

No:

```
// SPDX-License-Identifier: GPL-3.0
pragma solidity >=0.6.0 <0.8.0;


abstract contract A {
    function spam() virtual pure public;
    function ham() public virtual pure;
}



contract B is A {
    function spam() public pure override {
        // ...
    }
    function ham() public pure override {
        // ...
    }
}
```

Maximum Line Length

Keeping lines under the PEP 8 recommendation to a maximum of 79 (or 99) characters helps readers easily parse the code.

Wrapped lines should conform to the following guidelines.

1. The first argument should not be attached to the opening parenthesis.

2. One, and only one, indent should be used.

3. Each argument should fall on its own line.

4. The terminating element,); , should be placed on the final line by itself.

Function Calls

Yes:

```
thisFunctionCallIsReallyLong(
    longArgument1,
    longArgument2,
    longArgument3
);
```

No:

```
thisFunctionCallIsReallyLong(longArgument1,
                             longArgument2,
                             longArgument3
);

thisFunctionCallIsReallyLong(longArgument1,
    longArgument2,
    longArgument3
);

thisFunctionCallIsReallyLong(
    longArgument1, longArgument2,
    longArgument3
```

```
);

thisFunctionCallIsReallyLong(
longArgument1,
longArgument2,
longArgument3
);

thisFunctionCallIsReallyLong(
    longArgument1,
    longArgument2,
    longArgument3);
```

Assignment Statements

Yes:

```
thisIsALongNestedMapping[being][set][to_some_value] = someFunction(
    argument1,
    argument2,
    argument3,
    argument4
);
```

No:

```
thisIsALongNestedMapping[being][set][to_some_value] = someFunction(argument1,
                                                                    argument2,
                                                                    argument3,
                                                                    argument4);
```

Event Definitions and Event Emitters

Yes:

```
event LongAndLotsOfArgs(
    address sender,
    address recipient,
    uint256 publicKey,
    uint256 amount,
    bytes32[] options
);

LongAndLotsOfArgs(
    sender,
    recipient,
    publicKey,
    amount,
    options
);
```

No:

```
event LongAndLotsOfArgs(address sender,
                        address recipient,
                        uint256 publicKey,
                        uint256 amount,
```

```
                                    bytes32[] options);

LongAndLotsOfArgs(sender,
                  recipient,
                  publicKey,
                  amount,
                  options);
```

Source File Encoding

UTF-8 or ASCII encoding is preferred.

Imports

Import statements should always be placed at the top of the file.

Yes:

```solidity
// SPDX-License-Identifier: GPL-3.0
pragma solidity >=0.4.0 <0.8.0;

import "./Owned.sol";

contract A {
    // ...
}

contract B is Owned {
    // ...
}
```

No:

```solidity
// SPDX-License-Identifier: GPL-3.0
pragma solidity >=0.4.0 <0.8.0;

contract A {
    // ...
}


import "./Owned.sol";


contract B is Owned {
    // ...
}
```

Order of Functions

Ordering helps readers identify which functions they can call and to find the constructor and fallback definitions easier.

Functions should be grouped according to their visibility and ordered:

- constructor

- receive function (if exists)

- fallback function (if exists)

- external

- public

- internal

- private

Within a grouping, place the `view` and `pure` functions last.

Yes:

```solidity
// SPDX-License-Identifier: GPL-3.0
pragma solidity ^0.7.0;

contract A {
    constructor() {
        // ...
    }

    receive() external payable {
        // ...
    }

    fallback() external {
        // ...
    }

    // External functions
    // ...

    // External functions that are view
    // ...

    // External functions that are pure
    // ...

    // Public functions
    // ...

    // Internal functions
    // ...

    // Private functions
    // ...
}
```

No:

```solidity
// SPDX-License-Identifier: GPL-3.0
pragma solidity ^0.7.0;

contract A {

    // External functions
```

(continues on next page)

```solidity
    // ...

    fallback() external {
        // ...
    }
    receive() external payable {
        // ...
    }

    // Private functions
    // ...

    // Public functions
    // ...

    constructor() {
        // ...
    }

    // Internal functions
    // ...
}
```

Whitespace in Expressions

Avoid extraneous whitespace in the following situations:

Immediately inside parenthesis, brackets or braces, with the exception of single line function declarations.

Yes:

```solidity
spam(ham[1], Coin({name: "ham"}));
```

No:

```solidity
spam( ham[ 1 ], Coin( { name: "ham" } ) );
```

Exception:

```solidity
function singleLine() public { spam(); }
```

Immediately before a comma, semicolon:

Yes:

```solidity
function spam(uint i, Coin coin) public;
```

No:

```solidity
function spam(uint i , Coin coin) public ;
```

More than one space around an assignment or other operator to align with another:

Yes:

```
x = 1;
y = 2;
long_variable = 3;
```

No:

```
x             = 1;
y             = 2;
long_variable = 3;
```

Don't include a whitespace in the receive and fallback functions:

Yes:

```solidity
receive() external payable {
    ...
}

fallback() external {
    ...
}
```

No:

```solidity
receive () external payable {
    ...
}

fallback () external {
    ...
}
```

Control Structures

The braces denoting the body of a contract, library, functions and structs should:

- open on the same line as the declaration
- close on their own line at the same indentation level as the beginning of the declaration.
- The opening brace should be preceded by a single space.

Yes:

```solidity
// SPDX-License-Identifier: GPL-3.0
pragma solidity >=0.4.0 <0.8.0;

contract Coin {
    struct Bank {
        address owner;
        uint balance;
    }
}
```

No:

```
// SPDX-License-Identifier: GPL-3.0
pragma solidity >=0.4.0 <0.8.0;


contract Coin
{
    struct Bank {
        address owner;
        uint balance;
    }
}
```

The same recommendations apply to the control structures `if`, `else`, `while`, and `for`.

Additionally there should be a single space between the control structures `if`, `while`, and `for` and the parenthetic block representing the conditional, as well as a single space between the conditional parenthetic block and the opening brace.

Yes:

```
if (...) {
    ...
}

for (...) {
    ...
}
```

No:

```
if (...)
{
    ...
}

while(...){
}

for (...) {
    ...;}
```

For control structures whose body contains a single statement, omitting the braces is ok *if* the statement is contained on a single line.

Yes:

```
if (x < 10)
    x += 1;
```

No:

```
if (x < 10)
    someArray.push(Coin({
        name: 'spam',
        value: 42
    }));
```

For `if` blocks which have an `else` or `else if` clause, the `else` should be placed on the same line as the `if`'s closing brace. This is an exception compared to the rules of other block-like structures.

Yes:

```
if (x < 3) {
    x += 1;
} else if (x > 7) {
    x -= 1;
} else {
    x = 5;
}


if (x < 3)
    x += 1;
else
    x -= 1;
```

No:

```
if (x < 3) {
    x += 1;
}
else {
    x -= 1;
}
```

Function Declaration

For short function declarations, it is recommended for the opening brace of the function body to be kept on the same line as the function declaration.

The closing brace should be at the same indentation level as the function declaration.

The opening brace should be preceded by a single space.

Yes:

```
function increment(uint x) public pure returns (uint) {
    return x + 1;
}

function increment(uint x) public pure onlyOwner returns (uint) {
    return x + 1;
}
```

No:

```
function increment(uint x) public pure returns (uint)
{
    return x + 1;
}

function increment(uint x) public pure returns (uint){
    return x + 1;
}

function increment(uint x) public pure returns (uint) {
    return x + 1;
```

(continues on next page)

```solidity
    }

function increment(uint x) public pure returns (uint) {
    return x + 1;}
```

The modifier order for a function should be:

1. Visibility

2. Mutability

3. Virtual

4. Override

5. Custom modifiers

Yes:

```solidity
function balance(uint from) public view override returns (uint)  {
    return balanceOf[from];
}

function shutdown() public onlyOwner {
    selfdestruct(owner);
}
```

No:

```solidity
function balance(uint from) public override view returns (uint)  {
    return balanceOf[from];
}

function shutdown() onlyOwner public {
    selfdestruct(owner);
}
```

For long function declarations, it is recommended to drop each argument onto it's own line at the same indentation level as the function body. The closing parenthesis and opening bracket should be placed on their own line as well at the same indentation level as the function declaration.

Yes:

```solidity
function thisFunctionHasLotsOfArguments(
    address a,
    address b,
    address c,
    address d,
    address e,
    address f
)
    public
{
    doSomething();
}
```

No:

```
function thisFunctionHasLotsOfArguments(address a, address b, address c,
    address d, address e, address f) public {
    doSomething();
}

function thisFunctionHasLotsOfArguments(address a,
                                        address b,
                                        address c,
                                        address d,
                                        address e,
                                        address f) public {
    doSomething();
}

function thisFunctionHasLotsOfArguments(
    address a,
    address b,
    address c,
    address d,
    address e,
    address f) public {
    doSomething();
}
```

If a long function declaration has modifiers, then each modifier should be dropped to its own line.

Yes:

```
function thisFunctionNameIsReallyLong(address x, address y, address z)
    public
    onlyOwner
    priced
    returns (address)
{
    doSomething();
}

function thisFunctionNameIsReallyLong(
    address x,
    address y,
    address z,
)
    public
    onlyOwner
    priced
    returns (address)
{
    doSomething();
}
```

No:

```
function thisFunctionNameIsReallyLong(address x, address y, address z)
                                      public
                                      onlyOwner
                                      priced
                                      returns (address) {
    doSomething();
```

```solidity
}

function thisFunctionNameIsReallyLong(address x, address y, address z)
    public onlyOwner priced returns (address)
{
    doSomething();
}

function thisFunctionNameIsReallyLong(address x, address y, address z)
    public
    onlyOwner
    priced
    returns (address) {
    doSomething();
}
```

Multiline output parameters and return statements should follow the same style recommended for wrapping long lines found in the *Maximum Line Length* section.

Yes:

```solidity
function thisFunctionNameIsReallyLong(
    address a,
    address b,
    address c
)
    public
    returns (
        address someAddressName,
        uint256 LongArgument,
        uint256 Argument
    )
{
    doSomething()

    return (
        veryLongReturnArg1,
        veryLongReturnArg2,
        veryLongReturnArg3
    );
}
```

No:

```solidity
function thisFunctionNameIsReallyLong(
    address a,
    address b,
    address c
)
    public
    returns (address someAddressName,
            uint256 LongArgument,
            uint256 Argument)
{
    doSomething()

    return (veryLongReturnArg1,
```

```
            veryLongReturnArg1,
            veryLongReturnArg1);
}
```

For constructor functions on inherited contracts whose bases require arguments, it is recommended to drop the base constructors onto new lines in the same manner as modifiers if the function declaration is long or hard to read.

Yes:

```solidity
// SPDX-License-Identifier: GPL-3.0
pragma solidity ^0.7.0;

// Base contracts just to make this compile
contract B {
    constructor(uint) {
    }
}
contract C {
    constructor(uint, uint) {
    }
}
contract D {
    constructor(uint) {
    }
}

contract A is B, C, D {
    uint x;

    constructor(uint param1, uint param2, uint param3, uint param4, uint param5)
        B(param1)
        C(param2, param3)
        D(param4)
    {
        // do something with param5
        x = param5;
    }
}
```

No:

```solidity
// SPDX-License-Identifier: GPL-3.0
pragma solidity ^0.7.0;


// Base contracts just to make this compile
contract B {
    constructor(uint) {
    }
}


contract C {
    constructor(uint, uint) {
    }
}
```

```
contract D {
    constructor(uint) {
    }
}


contract A is B, C, D {
    uint x;

    constructor(uint param1, uint param2, uint param3, uint param4, uint param5)
    B(param1)
    C(param2, param3)
    D(param4) {
        x = param5;
    }
}


contract X is B, C, D {
    uint x;

    constructor(uint param1, uint param2, uint param3, uint param4, uint param5)
        B(param1)
        C(param2, param3)
        D(param4) {
            x = param5;
        }
}
```

When declaring short functions with a single statement, it is permissible to do it on a single line.

Permissible:

```
function shortFunction() public { doSomething(); }
```

These guidelines for function declarations are intended to improve readability. Authors should use their best judgement as this guide does not try to cover all possible permutations for function declarations.

Mappings

In variable declarations, do not separate the keyword `mapping` from its type by a space. Do not separate any nested `mapping` keyword from its type by whitespace.

Yes:

```
mapping(uint => uint) map;
mapping(address => bool) registeredAddresses;
mapping(uint => mapping(bool => Data[])) public data;
mapping(uint => mapping(uint => s)) data;
```

No:

```
mapping (uint => uint) map;
mapping( address => bool ) registeredAddresses;
```

```
mapping (uint => mapping (bool => Data[])) public data;
mapping(uint => mapping (uint => s)) data;
```

Variable Declarations

Declarations of array variables should not have a space between the type and the brackets.

Yes:

```
uint[] x;
```

No:

```
uint [] x;
```

Other Recommendations

- Strings should be quoted with double-quotes instead of single-quotes.

Yes:

```
str = "foo";
str = "Hamlet says, 'To be or not to be...'";
```

No:

```
str = 'bar';
str = '"Be yourself; everyone else is already taken." -Oscar Wilde';
```

- Surround operators with a single space on either side.

Yes:

```
x = 3;
x = 100 / 10;
x += 3 + 4;
x |= y && z;
```

No:

```
x=3;
x = 100/10;
x += 3+4;
x |= y&&z;
```

- Operators with a higher priority than others can exclude surrounding whitespace in order to denote precedence. This is meant to allow for improved readability for complex statement. You should always use the same amount of whitespace on either side of an operator:

Yes:

```
x = 2**3 + 5;
x = 2*y + 3*z;
x = (a+b) * (a-b);
```

No:

```
x = 2** 3 + 5;
x = y+z;
x +=1;
```

3.28.3 Order of Layout

Layout contract elements in the following order:

1. Pragma statements

2. Import statements

3. Interfaces

4. Libraries

5. Contracts

Inside each contract, library or interface, use the following order:

1. Type declarations

2. State variables

3. Events

4. Functions

Note: It might be clearer to declare types close to their use in events or state variables.

3.28.4 Naming Conventions

Naming conventions are powerful when adopted and used broadly. The use of different conventions can convey significant *meta* information that would otherwise not be immediately available.

The naming recommendations given here are intended to improve the readability, and thus they are not rules, but rather guidelines to try and help convey the most information through the names of things.

Lastly, consistency within a codebase should always supersede any conventions outlined in this document.

Naming Styles

To avoid confusion, the following names will be used to refer to different naming styles.

- `b` (single lowercase letter)

- `B` (single uppercase letter)

- `lowercase`

- `lower_case_with_underscores`

- `UPPERCASE`

- `UPPER_CASE_WITH_UNDERSCORES`

- `CapitalizedWords` (or CapWords)

- `mixedCase` (differs from CapitalizedWords by initial lowercase character!)

- `Capitalized_Words_With_Underscores`

Note: When using initialisms in CapWords, capitalize all the letters of the initialisms. Thus HTTPServerError is better than HttpServerError. When using initialisms in mixedCase, capitalize all the letters of the initialisms, except keep the first one lower case if it is the beginning of the name. Thus xmlHTTPRequest is better than XMLHTTPRequest.

Names to Avoid

- `l` - Lowercase letter el

- `O` - Uppercase letter oh

- `I` - Uppercase letter eye

Never use any of these for single letter variable names. They are often indistinguishable from the numerals one and zero.

Contract and Library Names

- Contracts and libraries should be named using the CapWords style. Examples: `SimpleToken`, `SmartBank`, `CertificateHashRepository`, `Player`, `Congress`, `Owned`.

- Contract and library names should also match their filenames.

- If a contract file includes multiple contracts and/or libraries, then the filename should match the *core contract*. This is not recommended however if it can be avoided.

As shown in the example below, if the contract name is `Congress` and the library name is `Owned`, then their associated filenames should be `Congress.sol` and `Owned.sol`.

Yes:

```solidity
// SPDX-License-Identifier: GPL-3.0
pragma solidity ^0.7.0;


// Owned.sol
contract Owned {
    address public owner;

    constructor() {
        owner = msg.sender;
    }

    modifier onlyOwner {
        require(msg.sender == owner);
        _;
    }

    function transferOwnership(address newOwner) public onlyOwner {
        owner = newOwner;
    }
}
```

and in `Congress.sol`:

```solidity
// SPDX-License-Identifier: GPL-3.0
pragma solidity >=0.4.0 <0.8.0;

import "./Owned.sol";


contract Congress is Owned, TokenRecipient {
    //...
}
```

No:

```solidity
// SPDX-License-Identifier: GPL-3.0
pragma solidity ^0.7.0;


// owned.sol
contract owned {
    address public owner;

    constructor() {
        owner = msg.sender;
    }

    modifier onlyOwner {
        require(msg.sender == owner);
        _;
    }

    function transferOwnership(address newOwner) public onlyOwner {
        owner = newOwner;
    }
}
```

and in `Congress.sol`:

```solidity
import "./owned.sol";


contract Congress is owned, tokenRecipient {
    //...
}
```

Struct Names

Structs should be named using the CapWords style. Examples: `MyCoin`, `Position`, `PositionXY`.

Event Names

Events should be named using the CapWords style. Examples: `Deposit`, `Transfer`, `Approval`, `BeforeTransfer`, `AfterTransfer`.

Function Names

Functions should use mixedCase. Examples: `getBalance`, `transfer`, `verifyOwner`, `addMember`, `changeOwner`.

Function Argument Names

Function arguments should use mixedCase. Examples: `initialSupply`, `account`, `recipientAddress`, `senderAddress`, `newOwner`.

When writing library functions that operate on a custom struct, the struct should be the first argument and should always be named `self`.

Local and State Variable Names

Use mixedCase. Examples: `totalSupply`, `remainingSupply`, `balancesOf`, `creatorAddress`, `isPreSale`, `tokenExchangeRate`.

Constants

Constants should be named with all capital letters with underscores separating words. Examples: `MAX_BLOCKS`, `TOKEN_NAME`, `TOKEN_TICKER`, `CONTRACT_VERSION`.

Modifier Names

Use mixedCase. Examples: `onlyBy`, `onlyAfter`, `onlyDuringThePreSale`.

Enums

Enums, in the style of simple type declarations, should be named using the CapWords style. Examples: `TokenGroup`, `Frame`, `HashStyle`, `CharacterLocation`.

Avoiding Naming Collisions

- `single_trailing_underscore_`

This convention is suggested when the desired name collides with that of a built-in or otherwise reserved name.

3.28.5 NatSpec

Solidity contracts can have a form of comments that are the basis of the Ethereum Natural Language Specification Format.

Add comments above functions or contracts following doxygen notation of one or multiple lines starting with `///` or a multiline comment starting with `/**` and ending with `*/`.

For example, the contract from *a simple smart contract* with the comments added looks like the one below:

```solidity
// SPDX-License-Identifier: GPL-3.0
pragma solidity >=0.4.16 <0.8.0;


/// @author The Solidity Team
/// @title A simple storage example
contract SimpleStorage {
    uint storedData;

    /// Store `x`.
    /// @param x the new value to store
    /// @dev stores the number in the state variable `storedData`
    function set(uint x) public {
        storedData = x;
    }

    /// Return the stored value.
    /// @dev retrieves the value of the state variable `storedData`
    /// @return the stored value
    function get() public view returns (uint) {
        return storedData;
    }
}
```

It is recommended that Solidity contracts are fully annotated using *NatSpec* for all public interfaces (everything in the ABI).

Please see the section about *NatSpec* for a detailed explanation.

3.29 Common Patterns

3.29.1 Withdrawal from Contracts

The recommended method of sending funds after an effect is using the withdrawal pattern. Although the most intuitive method of sending Ether, as a result of an effect, is a direct `transfer` call, this is not recommended as it introduces a potential security risk. You may read more about this on the *Security Considerations* page.

The following is an example of the withdrawal pattern in practice in a contract where the goal is to send the most money to the contract in order to become the "richest", inspired by King of the Ether.

In the following contract, if you are no longer the richest, you receive the funds of the person who is now the richest.

```solidity
// SPDX-License-Identifier: GPL-3.0
pragma solidity ^0.7.0;

contract WithdrawalContract {
    address public richest;
    uint public mostSent;

    mapping (address => uint) pendingWithdrawals;

    constructor() payable {
        richest = msg.sender;
        mostSent = msg.value;
    }
```

<div style="text-align: right">(continues on next page)</div>

```
    function becomeRichest() public payable {
        require(msg.value > mostSent, "Not enough money sent.");
        pendingWithdrawals[richest] += msg.value;
        richest = msg.sender;
        mostSent = msg.value;
    }

    function withdraw() public {
        uint amount = pendingWithdrawals[msg.sender];
        // Remember to zero the pending refund before
        // sending to prevent re-entrancy attacks
        pendingWithdrawals[msg.sender] = 0;
        msg.sender.transfer(amount);
    }
}
```

This is as opposed to the more intuitive sending pattern:

```
// SPDX-License-Identifier: GPL-3.0
pragma solidity ^0.7.0;

contract SendContract {
    address payable public richest;
    uint public mostSent;

    constructor() payable {
        richest = msg.sender;
        mostSent = msg.value;
    }

    function becomeRichest() public payable {
        require(msg.value > mostSent, "Not enough money sent.");
        // This line can cause problems (explained below).
        richest.transfer(msg.value);
        richest = msg.sender;
        mostSent = msg.value;
    }
}
```

Notice that, in this example, an attacker could trap the contract into an unusable state by causing `richest` to be the address of a contract that has a receive or fallback function which fails (e.g. by using `revert()` or by just consuming more than the 2300 gas stipend transferred to them). That way, whenever `transfer` is called to deliver funds to the "poisoned" contract, it will fail and thus also `becomeRichest` will fail, with the contract being stuck forever.

In contrast, if you use the "withdraw" pattern from the first example, the attacker can only cause his or her own withdraw to fail and not the rest of the contract's workings.

3.29.2 Restricting Access

Restricting access is a common pattern for contracts. Note that you can never restrict any human or computer from reading the content of your transactions or your contract's state. You can make it a bit harder by using encryption, but if your contract is supposed to read the data, so will everyone else.

You can restrict read access to your contract's state by **other contracts**. That is actually the default unless you declare your state variables `public`.

Furthermore, you can restrict who can make modifications to your contract's state or call your contract's functions and this is what this section is about.

The use of **function modifiers** makes these restrictions highly readable.

```solidity
// SPDX-License-Identifier: GPL-3.0
pragma solidity >=0.4.22 <0.8.0;

contract AccessRestriction {
    // These will be assigned at the construction
    // phase, where `msg.sender` is the account
    // creating this contract.
    address public owner = msg.sender;
    uint public creationTime = block.timestamp;

    // Modifiers can be used to change
    // the body of a function.
    // If this modifier is used, it will
    // prepend a check that only passes
    // if the function is called from
    // a certain address.
    modifier onlyBy(address _account)
    {
        require(
            msg.sender == _account,
            "Sender not authorized."
        );
        // Do not forget the "_;"! It will
        // be replaced by the actual function
        // body when the modifier is used.
        _;
    }

    /// Make `_newOwner` the new owner of this
    /// contract.
    function changeOwner(address _newOwner)
        public
        onlyBy(owner)
    {
        owner = _newOwner;
    }

    modifier onlyAfter(uint _time) {
        require(
            block.timestamp >= _time,
            "Function called too early."
        );
        _;
    }

    /// Erase ownership information.
    /// May only be called 6 weeks after
    /// the contract has been created.
    function disown()
        public
        onlyBy(owner)
        onlyAfter(creationTime + 6 weeks)
    {
```

(continues on next page)

```solidity
        delete owner;
    }

    // This modifier requires a certain
    // fee being associated with a function call.
    // If the caller sent too much, he or she is
    // refunded, but only after the function body.
    // This was dangerous before Solidity version 0.4.0,
    // where it was possible to skip the part after `_;`.
    modifier costs(uint _amount) {
        require(
            msg.value >= _amount,
            "Not enough Ether provided."
        );
        _;
        if (msg.value > _amount)
            msg.sender.transfer(msg.value - _amount);
    }

    function forceOwnerChange(address _newOwner)
        public
        payable
        costs(200 ether)
    {
        owner = _newOwner;
        // just some example condition
        if (uint(owner) & 0 == 1)
            // This did not refund for Solidity
            // before version 0.4.0.
            return;
        // refund overpaid fees
    }
}
```

A more specialised way in which access to function calls can be restricted will be discussed in the next example.

3.29.3 State Machine

Contracts often act as a state machine, which means that they have certain **stages** in which they behave differently or in which different functions can be called. A function call often ends a stage and transitions the contract into the next stage (especially if the contract models **interaction**). It is also common that some stages are automatically reached at a certain point in **time**.

An example for this is a blind auction contract which starts in the stage "accepting blinded bids", then transitions to "revealing bids" which is ended by "determine auction outcome".

Function modifiers can be used in this situation to model the states and guard against incorrect usage of the contract.

Example

In the following example, the modifier atStage ensures that the function can only be called at a certain stage.

Automatic timed transitions are handled by the modifier timeTransitions, which should be used for all functions.

Note: Modifier Order Matters. If atStage is combined with timedTransitions, make sure that you mention it after the latter, so that the new stage is taken into account.

Finally, the modifier `transitionNext` can be used to automatically go to the next stage when the function finishes.

Note: Modifier May be Skipped. This only applies to Solidity before version 0.4.0: Since modifiers are applied by simply replacing code and not by using a function call, the code in the transitionNext modifier can be skipped if the function itself uses return. If you want to do that, make sure to call nextStage manually from those functions. Starting with version 0.4.0, modifier code will run even if the function explicitly returns.

```solidity
// SPDX-License-Identifier: GPL-3.0
pragma solidity >=0.4.22 <0.8.0;

contract StateMachine {
    enum Stages {
        AcceptingBlindedBids,
        RevealBids,
        AnotherStage,
        AreWeDoneYet,
        Finished
    }

    // This is the current stage.
    Stages public stage = Stages.AcceptingBlindedBids;

    uint public creationTime = block.timestamp;

    modifier atStage(Stages _stage) {
        require(
            stage == _stage,
            "Function cannot be called at this time."
        );
        _;
    }

    function nextStage() internal {
        stage = Stages(uint(stage) + 1);
    }

    // Perform timed transitions. Be sure to mention
    // this modifier first, otherwise the guards
    // will not take the new stage into account.
    modifier timedTransitions() {
        if (stage == Stages.AcceptingBlindedBids &&
                    block.timestamp >= creationTime + 10 days)
            nextStage();
        if (stage == Stages.RevealBids &&
                block.timestamp >= creationTime + 12 days)
            nextStage();
        // The other stages transition by transaction
        _;
    }

    // Order of the modifiers matters here!
```

(continues on next page)

```solidity
    function bid()
        public
        payable
        timedTransitions
        atStage(Stages.AcceptingBlindedBids)
    {
        // We will not implement that here
    }

    function reveal()
        public
        timedTransitions
        atStage(Stages.RevealBids)
    {
    }

    // This modifier goes to the next stage
    // after the function is done.
    modifier transitionNext()
    {
        _;
        nextStage();
    }

    function g()
        public
        timedTransitions
        atStage(Stages.AnotherStage)
        transitionNext
    {
    }

    function h()
        public
        timedTransitions
        atStage(Stages.AreWeDoneYet)
        transitionNext
    {
    }

    function i()
        public
        timedTransitions
        atStage(Stages.Finished)
    {
    }
}
```

3.30 List of Known Bugs

Below, you can find a JSON-formatted list of some of the known security-relevant bugs in the Solidity compiler. The file itself is hosted in the Github repository. The list stretches back as far as version 0.3.0, bugs known to be present only in versions preceding that are not listed.

There is another file called bugs_by_version.json, which can be used to check which bugs affect a specific version of

the compiler.

Contract source verification tools and also other tools interacting with contracts should consult this list according to the following criteria:

- It is mildly suspicious if a contract was compiled with a nightly compiler version instead of a released version. This list does not keep track of unreleased or nightly versions.

- It is also mildly suspicious if a contract was compiled with a version that was not the most recent at the time the contract was created. For contracts created from other contracts, you have to follow the creation chain back to a transaction and use the date of that transaction as creation date.

- It is highly suspicious if a contract was compiled with a compiler that contains a known bug and the contract was created at a time where a newer compiler version containing a fix was already released.

The JSON file of known bugs below is an array of objects, one for each bug, with the following keys:

name Unique name given to the bug

summary Short description of the bug

description Detailed description of the bug

link URL of a website with more detailed information, optional

introduced The first published compiler version that contained the bug, optional

fixed The first published compiler version that did not contain the bug anymore

publish The date at which the bug became known publicly, optional

severity Severity of the bug: very low, low, medium, high. Takes into account discoverability in contract tests, likelihood of occurrence and potential damage by exploits.

conditions Conditions that have to be met to trigger the bug. The following keys can be used: `optimizer`, Boolean value which means that the optimizer has to be switched on to enable the bug. `evmVersion`, a string that indicates which EVM version compiler settings trigger the bug. The string can contain comparison operators. For example, `">=constantinople"` means that the bug is present when the EVM version is set to `constantinople` or later. If no conditions are given, assume that the bug is present.

check This field contains different checks that report whether the smart contract contains the bug or not. The first type of check are Javascript regular expressions that are to be matched against the source code ("source-regex") if the bug is present. If there is no match, then the bug is very likely not present. If there is a match, the bug might be present. For improved accuracy, the checks should be applied to the source code after stripping comments. The second type of check are patterns to be checked on the compact AST of the Solidity program ("ast-compact-json-path"). The specified search query is a JsonPath expression. If at least one path of the Solidity AST matches the query, the bug is likely present.

```
[
    {
        "name": "EmptyByteArrayCopy",
        "summary": "Copying an empty byte array (or string) from memory or calldata
→to storage can result in data corruption if the target array's length is increased
→subsequently without storing new data.",
        "description": "The routine that copies byte arrays from memory or calldata
→to storage stores unrelated data from after the source array in the storage slot if
→the source array is empty. If the storage array's length is subsequently increased
→either by using ``.push()`` or by assigning to its ``.length`` attribute (only
→before 0.6.0), the newly created byte array elements will not be zero-initialized,
→but contain the unrelated data. You are not affected if you do not assign to ``.
→length`` and do not use ``.push()`` on byte arrays, or only use ``.push(<arg>)`` or
→manually initialize the new elements.",
```

(continues on next page)

```
             "fixed": "0.7.4",
             "severity": "medium"
         },
         {
             "name": "DynamicArrayCleanup",
             "summary": "When assigning a dynamically-sized array with types of size at␣
→most 16 bytes in storage causing the assigned array to shrink, some parts of␣
→deleted slots were not zeroed out.",
             "description": "Consider a dynamically-sized array in storage whose base-type␣
→is small enough such that multiple values can be packed into a single slot, such as␣
→`uint128[]`. Let us define its length to be `l`. When this array gets assigned from␣
→another array with a smaller length, say `m`, the slots between elements `m` and␣
→`l` have to be cleaned by zeroing them out. However, this cleaning was not␣
→performed properly. Specifically, after the slot corresponding to `m`, only the␣
→first packed value was cleaned up. If this array gets resized to a length larger␣
→than `m`, the indices corresponding to the unclean parts of the slot contained the␣
→original value, instead of 0. The resizing here is performed by assigning to the␣
→array `length`, by a `push()` or via inline assembly. You are not affected if you␣
→are only using `.push(<arg>)` or if you assign a value (even zero) to the new␣
→elements after increasing the length of the array.",
             "fixed": "0.7.3",
             "severity": "medium"
         },
         {
             "name": "FreeFunctionRedefinition",
             "summary": "The compiler does not flag an error when two or more free␣
→functions with the same name and parameter types are defined in a source unit or␣
→when an imported free function alias shadows another free function with a different␣
→name but identical parameter types.",
             "description": "In contrast to functions defined inside contracts, free␣
→functions with identical names and parameter types did not create an error. Both␣
→definition of free functions with identical name and parameter types and an␣
→imported free function with an alias that shadows another function with a different␣
→name but identical parameter types were permitted due to which a call to either the␣
→multiply defined free function or the imported free function alias within a␣
→contract led to the execution of that free function which was defined first within␣
→the source unit. Subsequently defined identical free function definitions were␣
→silently ignored and their code generation was skipped.",
             "introduced": "0.7.1",
             "fixed": "0.7.2",
             "severity": "low"
         },
         {
             "name": "UsingForCalldata",
             "summary": "Function calls to internal library functions with calldata␣
→parameters called via ``using for`` can result in invalid data being read.",
             "description": "Function calls to internal library functions using the␣
→``using for`` mechanism copied all calldata parameters to memory first and passed␣
→them on like that, regardless of whether it was an internal or an external call.␣
→Due to that, the called function would receive a memory pointer that is interpreted␣
→as a calldata pointer. Since dynamically sized arrays are passed using two stack␣
→slots for calldata, but only one for memory, this can lead to stack corruption. An␣
→affected library call will consider the JUMPDEST to which it is supposed to return␣
→as part of its arguments and will instead jump out to whatever was on the stack␣
→before the call.",
             "introduced": "0.6.9",
             "fixed": "0.6.10",
```

3.30. List of Known Bugs

```
        "severity": "very low"
    },
    {
        "name": "MissingEscapingInFormatting",
        "summary": "String literals containing double backslash characters passed␣
↪directly to external or encoding function calls can lead to a different string␣
↪being used when ABIEncoderV2 is enabled.",
        "description": "When ABIEncoderV2 is enabled, string literals passed directly␣
↪to encoding functions or external function calls are stored as strings in the␣
↪intemediate code. Characters outside the printable range are handled correctly, but␣
↪backslashes are not escaped in this procedure. This leads to double backslashes␣
↪being reduced to single backslashes and consequently re-interpreted as escapes␣
↪potentially resulting in a different string being encoded.",
        "introduced": "0.5.14",
        "fixed": "0.6.8",
        "severity": "very low",
        "conditions": {
            "ABIEncoderV2": true
        }
    },
    {
        "name": "ArraySliceDynamicallyEncodedBaseType",
        "summary": "Accessing array slices of arrays with dynamically encoded base␣
↪types (e.g. multi-dimensional arrays) can result in invalid data being read.",
        "description": "For arrays with dynamically sized base types, index range␣
↪accesses that use a start expression that is non-zero will result in invalid array␣
↪slices. Any index access to such array slices will result in data being read from␣
↪incorrect calldata offsets. Array slices are only supported for dynamic calldata␣
↪types and all problematic type require ABIEncoderV2 to be enabled.",
        "introduced": "0.6.0",
        "fixed": "0.6.8",
        "severity": "very low",
        "conditions": {
            "ABIEncoderV2": true
        }
    },
    {
        "name": "ImplicitConstructorCallvalueCheck",
        "summary": "The creation code of a contract that does not define a␣
↪constructor but has a base that does define a constructor did not revert for calls␣
↪with non-zero value.",
        "description": "Starting from Solidity 0.4.5 the creation code of contracts␣
↪without explicit payable constructor is supposed to contain a callvalue check that␣
↪results in contract creation reverting, if non-zero value is passed. However, this␣
↪check was missing in case no explicit constructor was defined in a contract at all,␣
↪but the contract has a base that does define a constructor. In these cases it is␣
↪possible to send value in a contract creation transaction or using inline assembly␣
↪without revert, even though the creation code is supposed to be non-payable.",
        "introduced": "0.4.5",
        "fixed": "0.6.8",
        "severity": "very low"
    },
    {
        "name": "TupleAssignmentMultiStackSlotComponents",
        "summary": "Tuple assignments with components that occupy several stack slots,␣
↪ i.e. nested tuples, pointers to external functions or references to dynamically␣
↪sized calldata arrays, can result in invalid values.",
```

```
        "description": "Tuple assignments did not correctly account for tuple␣
→components that occupy multiple stack slots in case the number of stack slots␣
→differs between left-hand-side and right-hand-side. This can either happen in the␣
→presence of nested tuples or if the right-hand-side contains external function␣
→pointers or references to dynamic calldata arrays, while the left-hand-side␣
→contains an omission.",
        "introduced": "0.1.6",
        "fixed": "0.6.6",
        "severity": "very low"
    },
    {
        "name": "MemoryArrayCreationOverflow",
        "summary": "The creation of very large memory arrays can result in␣
→overlapping memory regions and thus memory corruption.",
        "description": "No runtime overflow checks were performed for the length of␣
→memory arrays during creation. In cases for which the memory size of an array in␣
→bytes, i.e. the array length times 32, is larger than 2^256-1, the memory␣
→allocation will overflow, potentially resulting in overlapping memory areas. The␣
→length of the array is still stored correctly, so copying or iterating over such an␣
→array will result in out-of-gas.",
        "link": "https://solidity.ethereum.org/2020/04/06/memory-creation-overflow-
→bug/",
        "introduced": "0.2.0",
        "fixed": "0.6.5",
        "severity": "low"
    },
    {
        "name": "YulOptimizerRedundantAssignmentBreakContinue",
        "summary": "The Yul optimizer can remove essential assignments to variables␣
→declared inside for loops when Yul's continue or break statement is used. You are␣
→unlikely to be affected if you do not use inline assembly with for loops and␣
→continue and break statements.",
        "description": "The Yul optimizer has a stage that removes assignments to␣
→variables that are overwritten again or are not used in all following control-flow␣
→branches. This logic incorrectly removes such assignments to variables declared␣
→inside a for loop if they can be removed in a control-flow branch that ends with␣
→``break`` or ``continue`` even though they cannot be removed in other control-flow␣
→branches. Variables declared outside of the respective for loop are not affected.",
        "introduced": "0.6.0",
        "fixed": "0.6.1",
        "severity": "medium",
        "conditions": {
            "yulOptimizer": true
        }
    },
    {
        "name": "privateCanBeOverridden",
        "summary": "Private methods can be overridden by inheriting contracts.",
        "description": "While private methods of base contracts are not visible and␣
→cannot be called directly from the derived contract, it is still possible to␣
→declare a function of the same name and type and thus change the behaviour of the␣
→base contract's function.",
        "introduced": "0.3.0",
        "fixed": "0.5.17",
        "severity": "low"
    },
    {
```

```
            "name": "YulOptimizerRedundantAssignmentBreakContinue0.5",
            "summary": "The Yul optimizer can remove essential assignments to variables␣
→declared inside for loops when Yul's continue or break statement is used. You are␣
→unlikely to be affected if you do not use inline assembly with for loops and␣
→continue and break statements.",
            "description": "The Yul optimizer has a stage that removes assignments to␣
→variables that are overwritten again or are not used in all following control-flow␣
→branches. This logic incorrectly removes such assignments to variables declared␣
→inside a for loop if they can be removed in a control-flow branch that ends with␣
→``break`` or ``continue`` even though they cannot be removed in other control-flow␣
→branches. Variables declared outside of the respective for loop are not affected.",
            "introduced": "0.5.8",
            "fixed": "0.5.16",
            "severity": "low",
            "conditions": {
                "yulOptimizer": true
            }
    },
    {
            "name": "ABIEncoderV2LoopYulOptimizer",
            "summary": "If both the experimental ABIEncoderV2 and the experimental Yul␣
→optimizer are activated, one component of the Yul optimizer may reuse data in␣
→memory that has been changed in the meantime.",
            "description": "The Yul optimizer incorrectly replaces ``mload`` and␣
→``sload`` calls with values that have been previously written to the load location␣
→(and potentially changed in the meantime) if all of the following conditions are␣
→met: (1) there is a matching ``mstore`` or ``sstore`` call before; (2) the contents␣
→of memory or storage is only changed in a function that is called (directly or␣
→indirectly) in between the first store and the load call; (3) called function␣
→contains a for loop where the same memory location is changed in the condition or␣
→the post or body block. When used in Solidity mode, this can only happen if the␣
→experimental ABIEncoderV2 is activated and the experimental Yul optimizer has been␣
→activated manually in addition to the regular optimizer in the compiler settings.",
            "introduced": "0.5.14",
            "fixed": "0.5.15",
            "severity": "low",
            "conditions": {
                "ABIEncoderV2": true,
                "optimizer": true,
                "yulOptimizer": true
            }
    },
    {
            "name":
→"ABIEncoderV2CalldataStructsWithStaticallySizedAndDynamicallyEncodedMembers",
            "summary": "Reading from calldata structs that contain dynamically encoded,␣
→but statically-sized members can result in incorrect values.",
            "description": "When a calldata struct contains a dynamically encoded, but␣
→statically-sized member, the offsets for all subsequent struct members are␣
→calculated incorrectly. All reads from such members will result in invalid values.␣
→Only calldata structs are affected, i.e. this occurs in external functions with␣
→such structs as argument. Using affected structs in storage or memory or as␣
→arguments to public functions on the other hand works correctly.",
            "introduced": "0.5.6",
            "fixed": "0.5.11",
            "severity": "low",
            "conditions": {
```

```
            "ABIEncoderV2": true
        }
    },
    {
        "name": "SignedArrayStorageCopy",
        "summary": "Assigning an array of signed integers to a storage array of
→different type can lead to data corruption in that array.",
        "description": "In two's complement, negative integers have their higher
→order bits set. In order to fit into a shared storage slot, these have to be set to
→zero. When a conversion is done at the same time, the bits to set to zero were
→incorrectly determined from the source and not the target type. This means that
→such copy operations can lead to incorrect values being stored.",
        "link": "https://blog.ethereum.org/2019/06/25/solidity-storage-array-bugs/",
        "introduced": "0.4.7",
        "fixed": "0.5.10",
        "severity": "low/medium"
    },
    {
        "name": "ABIEncoderV2StorageArrayWithMultiSlotElement",
        "summary": "Storage arrays containing structs or other statically-sized
→arrays are not read properly when directly encoded in external function calls or in
→abi.encode*.",
        "description": "When storage arrays whose elements occupy more than a single
→storage slot are directly encoded in external function calls or using abi.encode*,
→their elements are read in an overlapping manner, i.e. the element pointer is not
→properly advanced between reads. This is not a problem when the storage data is
→first copied to a memory variable or if the storage array only contains value types
→or dynamically-sized arrays.",
        "link": "https://blog.ethereum.org/2019/06/25/solidity-storage-array-bugs/",
        "introduced": "0.4.16",
        "fixed": "0.5.10",
        "severity": "low",
        "conditions": {
            "ABIEncoderV2": true
        }
    },
    {
        "name": "DynamicConstructorArgumentsClippedABIV2",
        "summary": "A contract's constructor that takes structs or arrays that
→contain dynamically-sized arrays reverts or decodes to invalid data.",
        "description": "During construction of a contract, constructor parameters are
→copied from the code section to memory for decoding. The amount of bytes to copy
→was calculated incorrectly in case all parameters are statically-sized but contain
→dynamically-sized arrays as struct members or inner arrays. Such types are only
→available if ABIEncoderV2 is activated.",
        "introduced": "0.4.16",
        "fixed": "0.5.9",
        "severity": "very low",
        "conditions": {
            "ABIEncoderV2": true
        }
    },
    {
        "name": "UninitializedFunctionPointerInConstructor",
        "summary": "Calling uninitialized internal function pointers created in the
→constructor does not always revert and can cause unexpected behaviour.",
        "description": "Uninitialized internal function pointers point to a special
→piece of code that causes a revert when called. Jump target positions are used else
→during construction and after deployment, but the code for setting this special
→jump target only considered the situation after deployment.",
```

```
        "introduced": "0.5.0",
        "fixed": "0.5.8",
        "severity": "very low"
    },
    {

        "name": "UninitializedFunctionPointerInConstructor_0.4.x",
        "summary": "Calling uninitialized internal function pointers created in the␣
↪constructor does not always revert and can cause unexpected behaviour.",
        "description": "Uninitialized internal function pointers point to a special␣
↪piece of code that causes a revert when called. Jump target positions are different␣
↪during construction and after deployment, but the code for setting this special␣
↪jump target only considered the situation after deployment.",
        "introduced": "0.4.5",
        "fixed": "0.4.26",
        "severity": "very low"
    },
    {
        "name": "IncorrectEventSignatureInLibraries",
        "summary": "Contract types used in events in libraries cause an incorrect␣
↪event signature hash",
        "description": "Instead of using the type `address` in the hashed signature,␣
↪the actual contract name was used, leading to a wrong hash in the logs.",
        "introduced": "0.5.0",
        "fixed": "0.5.8",
        "severity": "very low"
    },
    {
        "name": "IncorrectEventSignatureInLibraries_0.4.x",
        "summary": "Contract types used in events in libraries cause an incorrect␣
↪event signature hash",
        "description": "Instead of using the type `address` in the hashed signature,␣
↪the actual contract name was used, leading to a wrong hash in the logs.",
        "introduced": "0.3.0",
        "fixed": "0.4.26",
        "severity": "very low"
    },
    {
        "name": "ABIEncoderV2PackedStorage",
        "summary": "Storage structs and arrays with types shorter than 32 bytes can␣
↪cause data corruption if encoded directly from storage using the experimental␣
↪ABIEncoderV2.",
        "description": "Elements of structs and arrays that are shorter than 32 bytes␣
↪are not properly decoded from storage when encoded directly (i.e. not via a memory␣
↪type) using ABIEncoderV2. This can cause corruption in the values themselves but␣
↪can also overwrite other parts of the encoded data.",
        "link": "https://blog.ethereum.org/2019/03/26/solidity-optimizer-and-
↪abiencoderv2-bug/",
        "introduced": "0.5.0",
        "fixed": "0.5.7",
        "severity": "low",
        "conditions": {
            "ABIEncoderV2": true
        }
    },
    {
        "name": "ABIEncoderV2PackedStorage_0.4.x",
        "summary": "Storage structs and arrays with types shorter than 32 bytes can␣
↪cause data corruption if encoded directly from storage using the experimental␣
↪ABIEncoderV2.",
```

```
        "description": "Elements of structs and arrays that are shorter than 32 bytes␣
→are not properly decoded from storage when encoded directly (i.e. not via a memory␣
→type) using ABIEncoderV2. This can cause corruption in the values themselves but␣
→can also overwrite other parts of the encoded data.",
        "link": "https://blog.ethereum.org/2019/03/26/solidity-optimizer-and-
→abiencoderv2-bug/",
        "introduced": "0.4.19",
        "fixed": "0.4.26",
        "severity": "low",
        "conditions": {
            "ABIEncoderV2": true
        }
    },
    {
        "name": "IncorrectByteInstructionOptimization",
        "summary": "The optimizer incorrectly handles byte opcodes whose second␣
→argument is 31 or a constant expression that evaluates to 31. This can result in␣
→unexpected values.",
        "description": "The optimizer incorrectly handles byte opcodes that use the␣
→constant 31 as second argument. This can happen when performing index access on␣
→bytesNN types with a compile-time constant value (not index) of 31 or when using␣
→the byte opcode in inline assembly.",
        "link": "https://blog.ethereum.org/2019/03/26/solidity-optimizer-and-
→abiencoderv2-bug/",
        "introduced": "0.5.5",
        "fixed": "0.5.7",
        "severity": "very low",
        "conditions": {
            "optimizer": true
        }
    },
    {
        "name": "DoubleShiftSizeOverflow",
        "summary": "Double bitwise shifts by large constants whose sum overflows 256␣
→bits can result in unexpected values.",
        "description": "Nested logical shift operations whose total shift size is␣
→2**256 or more are incorrectly optimized. This only applies to shifts by numbers of␣
→bits that are compile-time constant expressions.",
        "link": "https://blog.ethereum.org/2019/03/26/solidity-optimizer-and-
→abiencoderv2-bug/",
        "introduced": "0.5.5",
        "fixed": "0.5.6",
        "severity": "low",
        "conditions": {
            "optimizer": true,
            "evmVersion": ">=constantinople"
        }
    },
    {
        "name": "ExpExponentCleanup",
        "summary": "Using the ** operator with an exponent of type shorter than 256␣
→bits can result in unexpected values.",
        "description": "Higher order bits in the exponent are not properly cleaned␣
→before the EXP opcode is applied if the type of the exponent expression is smaller␣
→than 256 bits and not smaller than the type of the base. In that case, the result␣
→might be larger than expected if the exponent is assumed to lie within the value␣
→range of the type. Literal numbers as exponents are unaffected as are exponents or␣
→bases of type uint256.",
```

```
        "link": "https://blog.ethereum.org/2018/09/13/solidity-bugfix-release/",
        "fixed": "0.4.25",
        "severity": "medium/high",
        "check": {"regex-source": "[^/]\\*\\* *[^/0-9 ]"}
    },
    {
        "name": "EventStructWrongData",
        "summary": "Using structs in events logged wrong data.",
        "description": "If a struct is used in an event, the address of the struct is␣
→logged instead of the actual data.",
        "link": "https://blog.ethereum.org/2018/09/13/solidity-bugfix-release/",
        "introduced": "0.4.17",
        "fixed": "0.4.25",
        "severity": "very low",
        "check": {"ast-compact-json-path": "$..[?(@.nodeType === 'EventDefinition')]..
→[?(@.nodeType === 'UserDefinedTypeName' && @.typeDescriptions.typeString.startsWith(
→'struct'))]"}
    },
    {
        "name": "NestedArrayFunctionCallDecoder",
        "summary": "Calling functions that return multi-dimensional fixed-size arrays␣
→can result in memory corruption.",
        "description": "If Solidity code calls a function that returns a multi-
→dimensional fixed-size array, array elements are incorrectly interpreted as memory␣
→pointers and thus can cause memory corruption if the return values are accessed.␣
→Calling functions with multi-dimensional fixed-size arrays is unaffected as is␣
→returning fixed-size arrays from function calls. The regular expression only checks␣
→if such functions are present, not if they are called, which is required for the␣
→contract to be affected.",
        "link": "https://blog.ethereum.org/2018/09/13/solidity-bugfix-release/",
        "introduced": "0.1.4",
        "fixed": "0.4.22",
        "severity": "medium",
        "check": {"regex-source": "returns[^;{]*\\[\\s*[^\\] \\t\\r\\n\\v\\f][^
→\\]]*\\]\\s*\\[\\s*[^\\] \\t\\r\\n\\v\\f][^\\]]*\\][^{;]*[;{]"}
    },
    {
        "name": "OneOfTwoConstructorsSkipped",
        "summary": "If a contract has both a new-style constructor (using the␣
→constructor keyword) and an old-style constructor (a function with the same name as␣
→the contract) at the same time, one of them will be ignored.",
        "description": "If a contract has both a new-style constructor (using the␣
→constructor keyword) and an old-style constructor (a function with the same name as␣
→the contract) at the same time, one of them will be ignored. There will be a␣
→compiler warning about the old-style constructor, so contracts only using new-style␣
→constructors are fine.",
        "introduced": "0.4.22",
        "fixed": "0.4.23",
        "severity": "very low"
    },
    {
        "name": "ZeroFunctionSelector",
        "summary": "It is possible to craft the name of a function such that it is␣
→executed instead of the fallback function in very specific circumstances.",
        "description": "If a function has a selector consisting only of zeros, is␣
→payable and part of a contract that does not have a fallback function and at most␣
→five external functions in total, this function is called instead of the fallback␣
→function if Ether is sent to the contract without data.",
```

```
            "fixed": "0.4.18",
            "severity": "very low"
    },
    {
            "name": "DelegateCallReturnValue",
            "summary": "The low-level .delegatecall() does not return the execution␣
→outcome, but converts the value returned by the functioned called to a boolean␣
→instead.",
            "description": "The return value of the low-level .delegatecall() function is␣
→taken from a position in memory, where the call data or the return data resides.␣
→This value is interpreted as a boolean and put onto the stack. This means if the␣
→called function returns at least 32 zero bytes, .delegatecall() returns false even␣
→if the call was successful.",
            "introduced": "0.3.0",
            "fixed": "0.4.15",
            "severity": "low"
    },
    {
            "name": "ECRecoverMalformedInput",
            "summary": "The ecrecover() builtin can return garbage for malformed input.",
            "description": "The ecrecover precompile does not properly signal failure for␣
→malformed input (especially in the 'v' argument) and thus the Solidity function can␣
→return data that was previously present in the return area in memory.",
            "fixed": "0.4.14",
            "severity": "medium"
    },
    {
            "name": "SkipEmptyStringLiteral",
            "summary": "If \"\" is used in a function call, the following function␣
→arguments will not be correctly passed to the function.",
            "description": "If the empty string literal \"\" is used as an argument in a␣
→function call, it is skipped by the encoder. This has the effect that the encoding␣
→of all arguments following this is shifted left by 32 bytes and thus the function␣
→call data is corrupted.",
            "fixed": "0.4.12",
            "severity": "low"
    },
    {
            "name": "ConstantOptimizerSubtraction",
            "summary": "In some situations, the optimizer replaces certain numbers in the␣
→code with routines that compute different numbers.",
            "description": "The optimizer tries to represent any number in the bytecode␣
→by routines that compute them with less gas. For some special numbers, an incorrect␣
→routine is generated. This could allow an attacker to e.g. trick victims about a␣
→specific amount of ether, or function calls to call different functions (or none at␣
→all).",
            "link": "https://blog.ethereum.org/2017/05/03/solidity-optimizer-bug/",
            "fixed": "0.4.11",
            "severity": "low",
            "conditions": {
                "optimizer": true
            }
    },
    {
            "name": "IdentityPrecompileReturnIgnored",
            "summary": "Failure of the identity precompile was ignored.",
            "description": "Calls to the identity contract, which is used for copying␣
→memory, ignored its return value. On the public chain, calls to the identity␣
→precompile can be made in a way that they never fail, but this might be different␣
→on private chains.",
```

```
        "severity": "low",
        "fixed": "0.4.7"
    },
    {
        "name": "OptimizerStateKnowledgeNotResetForJumpdest",
        "summary": "The optimizer did not properly reset its internal state at jump␣
→destinations, which could lead to data corruption.",
        "description": "The optimizer performs symbolic execution at certain stages.␣
→At jump destinations, multiple code paths join and thus it has to compute a common␣
→state from the incoming edges. Computing this common state was simplified to just␣
→use the empty state, but this implementation was not done properly. This bug can␣
→cause data corruption.",
        "severity": "medium",
        "introduced": "0.4.5",
        "fixed": "0.4.6",
        "conditions": {
            "optimizer": true
        }
    },
    {
        "name": "HighOrderByteCleanStorage",
        "summary": "For short types, the high order bytes were not cleaned properly␣
→and could overwrite existing data.",
        "description": "Types shorter than 32 bytes are packed together into the same␣
→32 byte storage slot, but storage writes always write 32 bytes. For some types, the␣
→higher order bytes were not cleaned properly, which made it sometimes possible to␣
→overwrite a variable in storage when writing to another one.",
        "link": "https://blog.ethereum.org/2016/11/01/security-alert-solidity-
→variables-can-overwritten-storage/",
        "severity": "high",
        "introduced": "0.1.6",
        "fixed": "0.4.4"
    },
    {
        "name": "OptimizerStaleKnowledgeAboutSHA3",
        "summary": "The optimizer did not properly reset its knowledge about SHA3␣
→operations resulting in some hashes (also used for storage variable positions) not␣
→being calculated correctly.",
        "description": "The optimizer performs symbolic execution in order to save re-
→evaluating expressions whose value is already known. This knowledge was not␣
→properly reset across control flow paths and thus the optimizer sometimes thought␣
→that the result of a SHA3 operation is already present on the stack. This could␣
→result in data corruption by accessing the wrong storage slot.",
        "severity": "medium",
        "fixed": "0.4.3",
        "conditions": {
            "optimizer": true
        }
    },
    {
        "name": "LibrariesNotCallableFromPayableFunctions",
        "summary": "Library functions threw an exception when called from a call that␣
→received Ether.",
        "description": "Library functions are protected against sending them Ether␣
→through a call. Since the DELEGATECALL opcode forwards the information about how␣
→much Ether was sent with a call, the library function incorrectly assumed that␣
→Ether was sent to the library and threw an exception.",
```

```
            "severity": "low",
            "introduced": "0.4.0",
            "fixed": "0.4.2"
        },
        {
            "name": "SendFailsForZeroEther",
            "summary": "The send function did not provide enough gas to the recipient if
↪no Ether was sent with it.",
            "description": "The recipient of an Ether transfer automatically receives a
↪certain amount of gas from the EVM to handle the transfer. In the case of a zero-
↪transfer, this gas is not provided which causes the recipient to throw an exception.
↪",
            "severity": "low",
            "fixed": "0.4.0"
        },
        {
            "name": "DynamicAllocationInfiniteLoop",
            "summary": "Dynamic allocation of an empty memory array caused an infinite
↪loop and thus an exception.",
            "description": "Memory arrays can be created provided a length. If this
↪length is zero, code was generated that did not terminate and thus consumed all gas.
↪",
            "severity": "low",
            "fixed": "0.3.6"
        },
        {
            "name": "OptimizerClearStateOnCodePathJoin",
            "summary": "The optimizer did not properly reset its internal state at jump
↪destinations, which could lead to data corruption.",
            "description": "The optimizer performs symbolic execution at certain stages.
↪At jump destinations, multiple code paths join and thus it has to compute a common
↪state from the incoming edges. Computing this common state was not done correctly.
↪This bug can cause data corruption, but it is probably quite hard to use for
↪targeted attacks.",
            "severity": "low",
            "fixed": "0.3.6",
            "conditions": {
                "optimizer": true
            }
        },
        {
            "name": "CleanBytesHigherOrderBits",
            "summary": "The higher order bits of short bytesNN types were not cleaned
↪before comparison.",
            "description": "Two variables of type bytesNN were considered different if
↪their higher order bits, which are not part of the actual value, were different. An
↪attacker might use this to reach seemingly unreachable code paths by providing
↪incorrectly formatted input data.",
            "severity": "medium/high",
            "fixed": "0.3.3"
        },
        {
            "name": "ArrayAccessCleanHigherOrderBits",
            "summary": "Access to array elements for arrays of types with less than 32
↪bytes did not correctly clean the higher order bits, causing corruption in other
↪array elements.",
            "description": "Multiple elements of an array of values that are shorter than
↪17 bytes are packed into the same storage slot. Writing to a single element of such
↪an array did not properly clean the higher order bytes and thus could lead to data
↪corruption.",
```

```
        "severity": "medium/high",
        "fixed": "0.3.1"
    },
    {
        "name": "AncientCompiler",
        "summary": "This compiler version is ancient and might contain several␣
↪undocumented or undiscovered bugs.",
        "description": "The list of bugs is only kept for compiler versions starting␣
↪from 0.3.0, so older versions might contain undocumented bugs.",
        "severity": "high",
        "fixed": "0.3.0"
    }
]
```

3.31 Contributing

Help is always welcome and there are plenty of options how you can contribute to Solidity.

In particular, we appreciate support in the following areas:

- Reporting issues.
- Fixing and responding to Solidity's GitHub issues, especially those tagged as good first issue which are meant as introductory issues for external contributors.
- Improving the documentation.
- Translating the documentation into more languages.
- Responding to questions from other users on StackExchange and the Solidity Gitter Chat.
- Getting involved in the language design process by joining language design calls, proposing language changes or new features and providing feedback.

To get started, you can try *Building from Source* in order to familiarize yourself with the components of Solidity and the build process. Also, it may be useful to become well-versed at writing smart-contracts in Solidity.

Please note that this project is released with a Contributor Code of Conduct. By participating in this project - in the issues, pull requests, or Gitter channels - you agree to abide by its terms.

3.31.1 Team Calls

If you have issues or pull requests to discuss, or are interested in hearing what the team and contributors are working on, you can join our public team calls:

- Mondays at 3pm CET/CEST.
- Wednesdays at 2pm CET/CEST.

Both calls take place on Google Meet.

3.31.2 How to Report Issues

To report an issue, please use the GitHub issues tracker. When reporting issues, please mention the following details:

- Which version of Solidity you are using.

- What was the source code (if applicable).

- Which platform are you running on.

- How to reproduce the issue.

- What was the result of the issue.

- What the expected behaviour is.

Reducing the source code that caused the issue to a bare minimum is always very helpful and sometimes even clarifies a misunderstanding.

3.31.3 Workflow for Pull Requests

In order to contribute, please fork off of the `develop` branch and make your changes there. Your commit messages should detail *why* you made your change in addition to *what* you did (unless it is a tiny change).

If you need to pull in any changes from `develop` after making your fork (for example, to resolve potential merge conflicts), please avoid using `git merge` and instead, `git rebase` your branch. This will help us review your change more easily.

Additionally, if you are writing a new feature, please ensure you add appropriate test cases under `test/` (see below).

However, if you are making a larger change, please consult with the Solidity Development Gitter channel (different from the one mentioned above, this one is focused on compiler and language development instead of language usage) first.

New features and bugfixes should be added to the `Changelog.md` file: please follow the style of previous entries, when applicable.

Finally, please make sure you respect the coding style for this project. Also, even though we do CI testing, please test your code and ensure that it builds locally before submitting a pull request.

Thank you for your help!

3.31.4 Running the Compiler Tests

Prerequisites

Some tests require the evmone library, others require libz3. The test script tries to discover the location of the `evmone` library, which can be located in the current directory, installed on the system level, or the `deps` folder in the project top level. The required file is called `libevmone.so` on Linux systems, `evmone.dll` on Windows systems and `libevmone.dylib` on macOS.

On macOS some of the testing scripts expect GNU coreutils to be installed. This can be easiest accomplished using Homebrew: `brew install coreutils`.

Running the Tests

Solidity includes different types of tests, most of them bundled into the Boost C++ Test Framework application `soltest`. Running `build/test/soltest` or its wrapper `scripts/soltest.sh` is sufficient for most changes.

The `./scripts/tests.sh` script executes most Solidity tests automatically, including those bundled into the Boost C++ Test Framework application `soltest` (or its wrapper `scripts/soltest.sh`), as well as command line tests and compilation tests.

The test system automatically tries to discover the location of the `evmone` library starting from the current directory. The required file is called `libevmone.so` on Linux systems, `evmone.dll` on Windows systems and `libevmone.dylib` on macOS. If it is not found, tests that use it are skipped. These tests are `libsolididty/ semanticTests`, `libsolidity/GasCosts`, `libsolidity/SolidityEndToEndTest`, part of the soltest suite. To run all tests, download the library from GitHub and place it in the project root path or inside the `deps` folder.

If the `libz3` library is not installed on your system, you should disable the SMT tests by exporting `SMT_FLAGS=--no-smt` before running `./scripts/tests.sh` or running `./scripts/ soltest.sh --no-smt`. These tests are `libsolidity/smtCheckerTests` and `libsolidity/ smtCheckerTestsJSON`.

Note: To get a list of all unit tests run by Soltest, run `./build/test/soltest --list_content=HRF`.

For quicker results you can run a subset of, or specific tests.

To run a subset of tests, you can use filters: `./scripts/soltest.sh -t TestSuite/TestName`, where `TestName` can be a wildcard `*`.

Or, for example, to run all the tests for the yul disambiguator: `./scripts/soltest.sh -t "yulOptimizerTests/disambiguator/*" --no-smt`.

`./build/test/soltest --help` has extensive help on all of the options available.

See especially:

- show_progress (-p) to show test completion,

- run_test (-t) to run specific tests cases, and

- report-level (-r) give a more detailed report.

Note: Those working in a Windows environment wanting to run the above basic sets without libz3. Using Git Bash, you use: `./build/test/Release/soltest.exe -- --no-smt`. If you are running this in plain Command Prompt, use `.\build\test\Release\soltest.exe -- --no-smt`.

If you want to debug using GDB, make sure you build differently than the "usual". For example, you could run the following command in your `build` folder:

```
cmake -DCMAKE_BUILD_TYPE=Debug ..
make
```

This creates symbols so that when you debug a test using the `--debug` flag, you have access to functions and variables in which you can break or print with.

The CI runs additional tests (including `solc-js` and testing third party Solidity frameworks) that require compiling the Emscripten target.

Writing and Running Syntax Tests

Syntax tests check that the compiler generates the correct error messages for invalid code and properly accepts valid code. They are stored in individual files inside the `tests/libsolidity/syntaxTests` folder. These files must contain annotations, stating the expected result(s) of the respective test. The test suite compiles and checks them against the given expectations.

For example: `./test/libsolidity/syntaxTests/double_stateVariable_declaration.sol`

```
contract test {
    uint256 variable;
    uint128 variable;
}
// ----
// DeclarationError: (36-52): Identifier already declared.
```

A syntax test must contain at least the contract under test itself, followed by the separator `// ----`. The comments that follow the separator are used to describe the expected compiler errors or warnings. The number range denotes the location in the source where the error occurred. If you want the contract to compile without any errors or warning you can leave out the separator and the comments that follow it.

In the above example, the state variable `variable` was declared twice, which is not allowed. This results in a `DeclarationError` stating that the identifier was already declared.

The `isoltest` tool is used for these tests and you can find it under `./build/test/tools/`. It is an interactive tool which allows editing of failing contracts using your preferred text editor. Let's try to break this test by removing the second declaration of `variable`:

```
contract test {
    uint256 variable;
}
// ----
// DeclarationError: (36-52): Identifier already declared.
```

Running `./build/test/isoltest` again results in a test failure:

```
syntaxTests/double_stateVariable_declaration.sol: FAIL
    Contract:
        contract test {
            uint256 variable;
        }

    Expected result:
        DeclarationError: (36-52): Identifier already declared.
    Obtained result:
        Success
```

`isoltest` prints the expected result next to the obtained result, and also provides a way to edit, update or skip the current contract file, or quit the application.

It offers several options for failing tests:

- `edit`: `isoltest` tries to open the contract in an editor so you can adjust it. It either uses the editor given on the command line (as `isoltest --editor /path/to/editor`), in the environment variable `EDITOR` or just `/usr/bin/editor` (in that order).

- `update`: Updates the expectations for contract under test. This updates the annotations by removing unmet expectations and adding missing expectations. The test is then run again.

- `skip`: Skips the execution of this particular test.

- `quit`: Quits `isoltest`.

All of these options apply to the current contract, expect `quit` which stops the entire testing process.

Automatically updating the test above changes it to

```
contract test {
    uint256 variable;
}
// ----
```

and re-run the test. It now passes again:

```
Re-running test case...
syntaxTests/double_stateVariable_declaration.sol: OK
```

Note: Choose a name for the contract file that explains what it tests, e.g. `double_variable_declaration.sol`. Do not put more than one contract into a single file, unless you are testing inheritance or cross-contract calls. Each file should test one aspect of your new feature.

3.31.5 Running the Fuzzer via AFL

Fuzzing is a technique that runs programs on more or less random inputs to find exceptional execution states (segmentation faults, exceptions, etc). Modern fuzzers are clever and run a directed search inside the input. We have a specialized binary called `solfuzzer` which takes source code as input and fails whenever it encounters an internal compiler error, segmentation fault or similar, but does not fail if e.g., the code contains an error. This way, fuzzing tools can find internal problems in the compiler.

We mainly use AFL for fuzzing. You need to download and install the AFL packages from your repositories (afl, afl-clang) or build them manually. Next, build Solidity (or just the `solfuzzer` binary) with AFL as your compiler:

```
cd build
# if needed
make clean
cmake .. -DCMAKE_C_COMPILER=path/to/afl-gcc -DCMAKE_CXX_COMPILER=path/to/afl-g++
make solfuzzer
```

At this stage you should be able to see a message similar to the following:

```
Scanning dependencies of target solfuzzer
[ 98%] Building CXX object test/tools/CMakeFiles/solfuzzer.dir/fuzzer.cpp.o
afl-cc 2.52b by <lcamtuf@google.com>
afl-as 2.52b by <lcamtuf@google.com>
[+] Instrumented 1949 locations (64-bit, non-hardened mode, ratio 100%).
[100%] Linking CXX executable solfuzzer
```

If the instrumentation messages did not appear, try switching the cmake flags pointing to AFL's clang binaries:

```
# if previously failed
make clean
cmake .. -DCMAKE_C_COMPILER=path/to/afl-clang -DCMAKE_CXX_COMPILER=path/to/afl-clang++
make solfuzzer
```

Otherwise, upon execution the fuzzer halts with an error saying binary is not instrumented:

```
afl-fuzz 2.52b by <lcamtuf@google.com>
... (truncated messages)
[*] Validating target binary...
```

(continues on next page)

```
[-] Looks like the target binary is not instrumented! The fuzzer depends on
    compile-time instrumentation to isolate interesting test cases while
    mutating the input data. For more information, and for tips on how to
    instrument binaries, please see /usr/share/doc/afl-doc/docs/README.

    When source code is not available, you may be able to leverage QEMU
    mode support. Consult the README for tips on how to enable this.
    (It is also possible to use afl-fuzz as a traditional, "dumb" fuzzer.
    For that, you can use the -n option - but expect much worse results.)

[-] PROGRAM ABORT : No instrumentation detected
        Location : check_binary(), afl-fuzz.c:6920
```

Next, you need some example source files. This makes it much easier for the fuzzer to find errors. You can either copy some files from the syntax tests or extract test files from the documentation or the other tests:

```
mkdir /tmp/test_cases
cd /tmp/test_cases
# extract from tests:
path/to/solidity/scripts/isolate_tests.py path/to/solidity/test/libsolidity/
↪SolidityEndToEndTest.cpp
# extract from documentation:
path/to/solidity/scripts/isolate_tests.py path/to/solidity/docs docs
```

The AFL documentation states that the corpus (the initial input files) should not be too large. The files themselves should not be larger than 1 kB and there should be at most one input file per functionality, so better start with a small number of. There is also a tool called `afl-cmin` that can trim input files that result in similar behaviour of the binary.

Now run the fuzzer (the `-m` extends the size of memory to 60 MB):

```
afl-fuzz -m 60 -i /tmp/test_cases -o /tmp/fuzzer_reports -- /path/to/solfuzzer
```

The fuzzer creates source files that lead to failures in `/tmp/fuzzer_reports`. Often it finds many similar source files that produce the same error. You can use the tool `scripts/uniqueErrors.sh` to filter out the unique errors.

3.31.6 Whiskers

Whiskers is a string templating system similar to Mustache. It is used by the compiler in various places to aid readability, and thus maintainability and verifiability, of the code.

The syntax comes with a substantial difference to Mustache. The template markers `{{` and `}}` are replaced by `<` and `>` in order to aid parsing and avoid conflicts with *Yul* (The symbols `<` and `>` are invalid in inline assembly, while `{` and `}` are used to delimit blocks). Another limitation is that lists are only resolved one depth and they do not recurse. This may change in the future.

A rough specification is the following:

Any occurrence of `<name>` is replaced by the string-value of the supplied variable `name` without any escaping and without iterated replacements. An area can be delimited by `<#name>...</name>`. It is replaced by as many concatenations of its contents as there were sets of variables supplied to the template system, each time replacing any `<inner>` items by their respective value. Top-level variables can also be used inside such areas.

There are also conditionals of the form `<?name>...<!name>...</name>`, where template replacements continue recursively either in the first or the second segment depending on the value of the boolean parameter `name`. If `<?+name>...<!+name>...</+name>` is used, then the check is whether the string parameter `name` is non-empty.

3.31.7 Documentation Style Guide

In the following section you find style recommendations specifically focusing on documentation contributions to Solidity.

English Language

Use English, with British English spelling preferred, unless using project or brand names. Try to reduce the usage of local slang and references, making your language as clear to all readers as possible. Below are some references to help:

- Simplified technical English
- International English
- British English spelling

Note: While the official Solidity documentation is written in English, there are community contributed *Translations* in other languages available.

Title Case for Headings

Use title case for headings. This means capitalise all principal words in titles, but not articles, conjunctions, and prepositions unless they start the title.

For example, the following are all correct:

- Title Case for Headings.
- For Headings Use Title Case.
- Local and State Variable Names.
- Order of Layout.

Expand Contractions

Use expanded contractions for words, for example:

- "Do not" instead of "Don't".
- "Can not" instead of "Can't".

Active and Passive Voice

Active voice is typically recommended for tutorial style documentation as it helps the reader understand who or what is performing a task. However, as the Solidity documentation is a mixture of tutorials and reference content, passive voice is sometimes more applicable.

As a summary:

- Use passive voice for technical reference, for example language definition and internals of the Ethereum VM.
- Use active voice when describing recommendations on how to apply an aspect of Solidity.

For example, the below is in passive voice as it specifies an aspect of Solidity:

Functions can be declared `pure` in which case they promise not to read from or modify the state.

For example, the below is in active voice as it discusses an application of Solidity:

When invoking the compiler, you can specify how to discover the first element of a path, and also path prefix remappings.

Common Terms

- "Function parameters" and "return variables", not input and output parameters.

Code Examples

A CI process tests all code block formatted code examples that begin with `pragma solidity`, `contract`, `library` or `interface` using the `./test/cmdlineTests.sh` script when you create a PR. If you are adding new code examples, ensure they work and pass tests before creating the PR.

Ensure that all code examples begin with a `pragma` version that spans the largest where the contract code is valid. For example `pragma solidity >=0.4.0 <0.8.0;`.

Running Documentation Tests

Make sure your contributions pass our documentation tests by running `./scripts/docs.sh` that installs dependencies needed for documentation and checks for any problems such as broken links or syntax issues.

3.31.8 Solidity Language Design

If you want to get involved in the language design process and share your ideas, please join the solidity-users forum, where existing properties of the language and proposals for new features can be discussed.

We regularly host language design discussion calls, in which selected topics, issues or feature implementations are debated in detail. The invitation to those calls is shared via the aforementioned forum. We are also sharing feedback surveys and other language design relevant content in this forum.

For ad-hoc cases and questions you can reach out to us via the Solidity-dev Gitter channel, a dedicated chatroom for conversations around the Solidity compiler and language development.

You can follow the implementation status of new features in the Solidity Github project. Issues in the design backlog need further specification and will either be discussed in a language design call or in a regular team call. You can see the upcoming changes for the next breaking release by changing from the default branch (*develop*) to the breaking branch.

We are happy to hear your thoughts on how we can improve the language design process to be even more collaborative and transparent.

3.32 Solidity Brand Guide

This brand guide features information on Solidity's brand policy and logo usage guidelines.

3.32.1 The Solidity Brand

The Solidity programming language is an open-source, community project governed by a core team. The core team is sponsored by the Ethereum Foundation.

This document aims to provide information about how to best use the Solidity brand name and logo.

We encourage you to read this document carefully before using the brand name or the logo. Your cooperation is highly appreciated!

3.32.2 Solidity Brand Name

"Solidity" should be used to refer to the Solidity programming language solely.

Please do not use "Solidity":

- To refer to any other programming language.

- In a way that is misleading or may imply association of unrelated modules, tools, documentation, or other resources with the Solidity programming language.

- In ways that confuse the community as to whether the Solidity programming language is open-source and free to use.

3.32.3 Solidity Logo License

The Solidity logo is distributed and licensed under a Creative Commons Attribution 4.0 International License.

This is the most permissive Creative Commons license and allows reuse and modifications for any purpose.

You are free to:

- **Share** — Copy and redistribute the material in any medium or format.
- **Adapt** — Remix, transform, and build upon the material for any purpose, even commercially.

Under the following terms:

- **Attribution** — You must give appropriate credit, provide a link to the license, and indicate if changes were made. You may do so in any reasonable manner, but not in any way that suggests the the Solidity core team endorses you or your use.

When using the Solidity logo, please respect the Solidity logo guidelines.

3.32.4 Solidity Logo Guidelines

(Right click on the logo to download it.)

Please do not:

- Change the ratio of the logo (do not stretch it or cut it).

- Change the colors of the logo, unless it is absolutely necessary.

3.32.5 Credits

This document was, in parts, derived from the Python Software Foundation Trademark Usage Policy and the Rust Media Guide.

Index

A

abi, 77, 78, 146
abstract contract, **119**
access
    restricting, 241
account, **11**
addmod, 78, 131
address, 11, 51, 55
anonymous, 133
application binary interface, 146
array, 61, **62**, 101
    allocating, **63**
    length, **65**
    literals, **64**
    pop, **65**
    push, **65**
    slice, **67**
array of strings, 101
asm, **127**, **201**
assembly, **127**, **201**
assert, 78, **88**, 131
assignment, 73, **85**
    destructuring, **85**
auction
    blind, 25
    open, 25

B

balance, 11, 51, 79, 131
ballot, 22
base
    constructor, **117**
base class, **110**
blind auction, 25
block, **10**, 77, 131
    number, 77, 131
    timestamp, 77, 131
bool, **49**
break, 81

B (cont.)

Bugs, 245
byte array, 54
bytes, 57, **63**
bytes32, 54

C

C3 linearization, **117**
call, 51, 79
callcode, 13, 79, 120
cast, **73**
coding style, 220
coin, 10
coinbase, 77, 131
commandline compiler, **187**
comment, **45**
common subexpression elimination, 142
compile target, 189
compiler
    commandline, 187
constant, **99**, 133
constant propagation, 142
constructor, 92, **116**
    arguments, 92
continue, 81
contract, 46, **92**
    abstract, **119**
    base, **110**
    creation, **92**
    interface, **120**
    modular, 40
contract creation, 13
contract type, **54**
contract verification, 143
contracts
    creating, 84
creationCode, 81
cryptography, 78, 131

D

data, 77, 131

Y