
Solidity Documentation

Release 0.4.19

Ethereum

Dec 06, 2017

Contents

1	Translations	3
2	Useful links	5
3	Available Solidity Integrations	7
4	Solidity Tools	9
5	Third-Party Solidity Parsers and Grammars	11
6	Language Documentation	13
7	Contents	15
7.1	Introduction to Smart Contracts	15
7.2	Installing Solidity	21
7.3	Solidity by Example	25
7.4	Solidity in Depth	35
7.5	Security Considerations	110
7.6	Using the compiler	115
7.7	Contract Metadata	120
7.8	Application Binary Interface Specification	122
7.9	Joyfully Universal Language for (Inline) Assembly	131
7.10	Style Guide	139
7.11	Common Patterns	151
7.12	List of Known Bugs	156
7.13	Contributing	160
7.14	Frequently Asked Questions	162

Solidity is a contract-oriented, high-level language for implementing smart contracts. It 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.

As you will see, it is possible to create contracts for voting, crowdfunding, blind auctions, multi-signature wallets and more.

Note: The best way to try out Solidity right now is using [Remix](#) (it can take a while to load, please be patient).

CHAPTER 1

Translations

This documentation is translated into several languages by community volunteers, but the English version stands as a reference.

- [Spanish](#)
- [Russian](#) (rather outdated)

CHAPTER 2

Useful links

- [Ethereum](#)
- [Changelog](#)
- [Story Backlog](#)
- [Source Code](#)
- [Ethereum Stackexchange](#)
- [Gitter Chat](#)

Available Solidity Integrations

- **Remix** Browser-based IDE with integrated compiler and Solidity runtime environment without server-side components.
- **IntelliJ IDEA plugin** Solidity plugin for IntelliJ IDEA (and all other JetBrains IDEs)
- **Visual Studio Extension** Solidity plugin for Microsoft Visual Studio that includes the Solidity compiler.
- **Package for SublimeText — Solidity language syntax** Solidity syntax highlighting for SublimeText editor.
- **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 as a base.
- **Solium** A commandline linter for Solidity which strictly follows the rules prescribed by the [Solidity Style Guide](#).
- **Solhint** Solidity linter that provides security, style guide and best practice rules for smart contract validation.
- **Visual Studio Code extension** Solidity plugin for Microsoft Visual Studio Code that includes syntax highlighting and the Solidity compiler.
- **Emacs Solidity** Plugin for the Emacs editor providing syntax highlighting and compilation error reporting.
- **Vim Solidity** Plugin for the Vim editor providing syntax highlighting.
- **Vim Syntastic** Plugin for the Vim editor providing compile checking.

Discontinued:

- **Mix IDE** Qt based IDE for designing, debugging and testing solidity smart contracts.
- **Ethereum Studio** Specialized web IDE that also provides shell access to a complete Ethereum environment.

CHAPTER 4

Solidity Tools

- **Dapp** Build tool, package manager, and deployment assistant for Solidity.
- **Solidity REPL** Try Solidity instantly with a command-line Solidity console.
- **solgraph** Visualize Solidity control flow and highlight potential security vulnerabilities.
- **evmdis** EVM Disassembler that performs static analysis on the bytecode to provide a higher level of abstraction than raw EVM operations.
- **Doxity** Documentation Generator for Solidity.

Third-Party Solidity Parsers and Grammars

- **solidity-parser** Solidity parser for JavaScript
- **Solidity Grammar for ANTLR 4** Solidity grammar for the ANTLR 4 parser generator

Language Documentation

On the next pages, we will first see a *simple smart contract* written in Solidity followed by the basics about *blockchains* and the *Ethereum Virtual Machine*.

The next section will explain several *features* of Solidity by giving useful *example contracts*. Remember that you can always try out the contracts [in your browser!](#)

The last and most extensive section will cover all aspects of Solidity in depth.

If you still have questions, you can try searching or asking on the [Ethereum Stackexchange](#) site, or come to our [gitter channel](#). Ideas for improving Solidity or this documentation are always welcome!

[Keyword Index](#), [Search Page](#)

7.1 Introduction to Smart Contracts

7.1.1 A Simple Smart Contract

Let us begin with the most basic example. It is fine if you do not understand everything right now, we will go into more detail later.

Storage

```
pragma solidity ^0.4.0;

contract SimpleStorage {
    uint storedData;

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

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

The first line simply tells that the source code is written for Solidity version 0.4.0 or anything newer that does not break functionality (up to, but not including, version 0.5.0). This is to ensure that the contract does not suddenly behave differently with a new compiler version. The keyword `pragma` is called that way because, in general, pragmas are instructions for the compiler 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` (unsigned integer of 256 bits). You can think of it as a single slot in a database that can be queried and altered by calling functions of the code that manages the database. In the case of Ethereum, this is always the owning contract. And in this case, the functions `set` and `get` 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 (due to the infrastructure built by Ethereum) apart from 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. Of course, anyone could just call `set` again with a different value and overwrite your number, but the number will still be stored in the history of the blockchain. Later, we will see how you can impose access restrictions so that only you can alter the number.

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.

Warning: Be careful with using Unicode text as similarly looking (or even identical) characters can have different code points and as such will be encoded as a different byte array.

Subcurrency Example

The following contract will implement the simplest form of a cryptocurrency. It is possible to generate coins out of thin air, but only the person that created the contract will be able to do that (it is trivial to implement a different issuance scheme). Furthermore, anyone can send coins to each other without any need for registering with username and password - all you need is an Ethereum keypair.

```
pragma solidity ^0.4.0;

contract Coin {
    // The keyword "public" makes those variables
    // readable from outside.
    address public minter;
    mapping (address => uint) public balances;

    // Events allow light clients to react on
    // changes efficiently.
    event Sent(address from, address to, uint amount);

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

    function mint(address receiver, uint amount) {
        if (msg.sender != minter) return;
        balances[receiver] += amount;
    }

    function send(address receiver, uint amount) {
        if (balances[msg.sender] < amount) return;
        balances[msg.sender] -= amount;
    }
}
```

```

        balances[receiver] += amount;
        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` that is publicly accessible. The `address` type is a 160-bit value that does not allow any arithmetic operations. It is suitable for storing addresses of contracts or keypairs belonging to external persons. The keyword `public` automatically generates a function that allows you to access the current value of the state variable. Without this keyword, other contracts have no way to access the variable. The function will look something like this:

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

Of course, adding a function exactly like that will not work because we would have a function and a state variable with the same name, but hopefully, you get the idea - the compiler figures that 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 type maps addresses to unsigned integers. Mappings can be seen as [hash tables](#) which are virtually initialized such that every possible key exists and is mapped to a value whose byte-representation is all zeros. This analogy does not go too far, though, as it is neither possible to obtain a list of all keys of a mapping, nor a list of all values. So either keep in mind (or better, keep a list or use a more advanced data type) what you added to the mapping or use it in a context where this is not needed, like this one. The *getter function* created by the `public` keyword is a bit more complex in this case. It roughly looks like the following:

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

As you see, you can use this function to easily query the balance of a single account.

The line `event Sent(address from, address to, uint amount);` declares a so-called “event” which is fired in the last line of the function `send`. User interfaces (as well as server applications of course) can listen for those events being fired on the blockchain without much cost. As soon as it is fired, the listener will also receive the arguments `from`, `to` and `amount`, which makes it easy to track transactions. In order to listen for this event, you would use

```

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));
    }
})

```

Note how the automatically generated function `balances` is called from the user interface.

The special function `Coin` is the constructor which is run during creation of the contract and cannot be called afterwards. It permanently stores the address of the person creating the contract: `msg` (together with `tx` and `block`) is a magic global variable that contains some properties which allow access to the blockchain. `msg.sender` is always the address where the current (external) function call came from.

Finally, the functions that will actually end up with the contract and can be called by users and contracts alike are `mint` and `send`. If `mint` is called by anyone except the account that created the contract, nothing will happen. On

the other hand, `send` can be used by anyone (who already has some of these coins) to send coins to anyone else. Note that 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 fact that you sent coins and the changed balances are only stored in the data storage of this particular coin contract. By the use of events it is relatively easy to create a “blockchain explorer” that tracks transactions and balances of your new coin.

7.1.2 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. 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 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, a so-called conflict?

The abstract answer to this is that you do not have to care. An order of the transactions will be selected for you, 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 that are added on top, the less likely it is. 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.

7.1.3 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) 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 the special zero-account, see below). It can include binary data (its 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 the zero-account (the account with the address 0), 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 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.

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. 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 in the same way.

If the gas is used up at any point (i.e. it is 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

Each account has a persistent memory area which is called **storage**. 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 and it is comparatively costly to read and even more so, to modify storage. A contract can neither read nor write to any storage apart from its own.

The second memory 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 (ie. 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 an 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, 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 implementations which could cause consensus problems. All instructions operate on the basic data type, 256-bit words. 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.

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 signalled 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.

Calls are **limited** to a depth of 1024, which means that for more complex operations, loops should be preferred over recursive calls.

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 (“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). 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.

Self-destruct

The only possibility that code is removed 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.

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

Note: The pruning of old contracts may or may not be implemented by Ethereum clients. Additionally, archive nodes could choose to keep the contract storage and code indefinitely.

Note: Currently **external accounts** cannot be removed from the state.

7.2 Installing Solidity

7.2.1 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.

7.2.2 Remix

If you just want to try Solidity for small contracts, you can try **Remix** which does not need any installation. If you want to use it without connection to the Internet, you can go to <https://github.com/ethereum/browser-solidity/tree/gh-pages> and download the .ZIP file as explained on that page.

7.2.3 npm / Node.js

This is probably the most portable and most convenient way to install Solidity locally.

A platform-independent JavaScript library is provided by compiling the C++ source into JavaScript using Emscripten. It can be used in projects directly (such as Remix). Please refer to the [solc-js](#) repository for instructions.

It also contains a commandline tool called *solcjs*, which can be installed via npm:

```
npm install -g solc
```

Note: The comandline options of *solcjs* are not compatible with *solc* and tools (such as *geth*) expecting the behaviour of *solc* will not work with *solcjs*.

7.2.4 Docker

We provide up to date docker builds for the compiler. The `stable` repository contains released versions while the `nightly` repository contains potentially unstable changes in the `develop` branch.

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

Currently, the docker image only contains the compiler executable, so you have to do some additional work to link in the source and output directories.

7.2.5 Binary Packages

Binary packages of Solidity available at [solidity/releases](#).

We also have PPAs for Ubuntu. For the latest stable version.

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

If you want to use the cutting edge developer version:

```
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
```

Or if you want to help testing the unstable solc with the most recent changes from the development branch:

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

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

```
pacman -S solidity
```

Homebrew is missing pre-built bottles at the time of writing, following a Jenkins to TravisCI migration, but Homebrew should still work just fine as a means to build-from-source. We will re-add the pre-built bottles soon.

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

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

View [solidity.rb commits on Github](#).

Follow the history links until you have a raw file link of a specific commit of `solidity.rb`.

Install it using brew:

```
brew unlink solidity
# Install 0.4.8
brew install https://raw.githubusercontent.com/ethereum/homebrew-ethereum/
↪ 77cce03da9f289e5a3ffe579840d3c5dc0a62717/solidity.rb
```

Gentoo Linux also provides a solidity package that can be installed using emerge:

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

7.2.6 Building from Source

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:

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

Solidity has git submodules. Ensure they are properly loaded:

```
git submodule update --init --recursive
```

Prerequisites - macOS

For macOS, ensure that you have the latest version of [Xcode installed](#). This contains the [Clang C++ compiler](#), the [Xcode IDE](#) and other Apple development tools which 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 builds require you to [install 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 will need to install the following dependencies for Windows builds of Solidity:

Software	Notes
Git for Windows	Command-line tool for retrieving source from Github.
CMake	Cross-platform build file generator.
Visual Studio 2015	C++ compiler and dev environment.

External Dependencies

We now have a “one button” script which installs all required external dependencies on macOS, Windows and on numerous Linux distros. This used to be a multi-step manual process, but is now a one-liner:

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

Or, on Windows:

```
scripts\install_deps.bat
```

Command-Line Build

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

Solidity project uses CMake to configure the build. Building Solidity is quite similar on Linux, macOS and other Unices:

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

or even easier:

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

And even for Windows:

```
mkdir build
cd build
cmake -G "Visual Studio 14 2015 Win64" ..
```

This latter set of instructions 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 **RelWithDebInfo** configuration, but all others work.

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

```
cmake --build . --config RelWithDebInfo
```

7.2.7 CMake options

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

7.2.8 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 has 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`

7.2.9 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*.

7.3 Solidity by Example

7.3.1 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.

```

pragma solidity ^0.4.11;

/// @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;

    /// Create a new ballot to choose one of `proposalNames`.
    function Ballot(bytes32[] 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) {
            // If the argument of `require` evaluates to `false`,
            // it terminates and reverts all changes to
            // the state and to Ether balances. It is often
            // a good idea to use this if functions are
            // called incorrectly. But watch out, this
            // will currently also consume all provided gas
            // (this is planned to change in the future).

```

```

    require((msg.sender == chairperson) && !voters[voter].voted && (voters[voter].
↪weight == 0));
    voters[voter].weight = 1;
}

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

    // Self-delegation is not allowed.
    require(to != msg.sender);

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

    // 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) {
    Voter storage sender = voters[msg.sender];
    require(!sender.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() constant
    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
// of the winner contained in the proposals array and then
// returns the name of the winner
function winnerName() constant
    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?

7.3.2 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 her money back. After the end of the bidding period, the contract has to be called manually for the beneficiary to receive his money - contracts cannot activate themselves.

```
pragma solidity ^0.4.11;

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

    // 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
bool ended;

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

// The following is a so-called natspec comment,
// 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`.
function SimpleAuction(
    uint _biddingTime,
    address _beneficiary
) {
    beneficiary = _beneficiary;
    auctionEnd = now + _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() 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(now <= auctionEnd);

    // If the bid is not higher, send the
    // money back.
    require(msg.value > highestBid);

    if (highestBidder != 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;
    HighestBidIncreased(msg.sender, msg.value);
}

```

```
/// Withdraw a bid that was overbid.
function withdraw() 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() {
    // 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(now >= auctionEnd); // auction did not yet end
    require(!ended); // this function has already been called

    // 2. Effects
    ended = true;
    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 her 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 he won the auction is to make her 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.

```
pragma solidity ^0.4.11;

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

    address 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(now < _time); _; }
    modifier onlyAfter(uint _time) { require(now > _time); _; }

    function BlindAuction(
        uint _biddingTime,
        uint _revealTime,
        address _beneficiary
    ) {
        beneficiary = _beneficiary;
        biddingEnd = now + _biddingTime;
        revealEnd = biddingEnd + _revealTime;
    }

    /// Place a blinded bid with `_blindedBid` = keccak256(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
    /// 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)
    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[] _values,
    bool[] _fake,
    bytes32[] _secret
)
    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++) {
        var bid = bids[msg.sender][i];
        var (value, fake, secret) =
            (_values[i], _fake[i], _secret[i]);
        if (bid.blindedBid != keccak256(value, fake, secret)) {
            // Bid was not actually revealed.
            // Do not refund deposit.
            continue;
        }
        refund += bid.deposit;
        if (!fake && bid.deposit >= value) {
            if (placeBid(msg.sender, value))
                refund -= value;
        }
        // Make it impossible for the sender to re-claim
        // the same deposit.
        bid.blindedBid = bytes32(0);
    }
    msg.sender.transfer(refund);
}

// 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 != 0) {
        // Refund the previously highest bidder.
        pendingReturns[highestBidder] += highestBid;
    }
    highestBid = value;
    highestBidder = bidder;
    return true;
}

/// Withdraw a bid that was overbid.
function withdraw() {
    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 `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()
    onlyAfter(revealEnd)
{
    require(!ended);
    AuctionEnded(highestBidder, highestBid);
    ended = true;
    beneficiary.transfer(highestBid);
}
}

```

7.3.3 Safe Remote Purchase

```

pragma solidity ^0.4.11;

contract Purchase {
    uint public value;
    address public seller;
    address public buyer;
    enum State { Created, Locked, Inactive }
    State public state;

    // 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.
    function Purchase() payable {
        seller = msg.sender;
        value = msg.value / 2;
        require((2 * value) == msg.value);
    }
}

```

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

modifier onlyBuyer() {
    require(msg.sender == buyer);
    _;
}

modifier onlySeller() {
    require(msg.sender == seller);
    _;
}

modifier inState(State _state) {
    require(state == _state);
    _;
}

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

/// Abort the purchase and reclaim the ether.
/// Can only be called by the seller before
/// the contract is locked.
function abort()
    onlySeller
    inState(State.Created)
{
    Aborted();
    state = State.Inactive;
    seller.transfer(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()
    inState(State.Created)
    condition(msg.value == (2 * value))
    payable
{
    PurchaseConfirmed();
    buyer = msg.sender;
    state = State.Locked;
}

/// Confirm that you (the buyer) received the item.
/// This will release the locked ether.
function confirmReceived()
    onlyBuyer
    inState(State.Locked)
{
    ItemReceived();
    // It is important to change the state first because

```

```
// otherwise, the contracts called using `send` below
// can call in again here.
state = State.Inactive;

// NOTE: This actually allows both the buyer and the seller to
// block the refund - the withdraw pattern should be used.

buyer.transfer(value);
seller.transfer(this.balance);
}
}
```

7.3.4 Micropayment Channel

To be written.

7.4 Solidity in Depth

This section should provide you with all you need to know about Solidity. If something is missing here, please contact us on [Gitter](#) or make a pull request on [Github](#).

7.4.1 Layout of a Solidity Source File

Source files can contain an arbitrary number of contract definitions, include directives and pragma directives.

Version Pragma

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

The version pragma is used as follows:

```
pragma solidity ^0.4.0;
```

Such a source file will not compile with a compiler earlier than version 0.4.0 and it will also not work on a compiler starting from version 0.5.0 (this second condition is added by using `^`). The idea behind this is that there will be no breaking changes until version `0.5.0`, so we can always be sure that our code will compile the way we intended it to. We do not fix the exact version of the compiler, so that bugfix releases are still possible.

It is possible to specify much more complex rules for the compiler version, the expression follows those used by [npm](#).

Importing other Source Files

Syntax and Semantics

Solidity supports import statements that are very similar to those available in JavaScript (from ES6 on), although Solidity does not know 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).

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

...creates a new global symbol `symbolName` whose members are all the global symbols from “filename”.

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

...creates new global symbols `alias` and `symbol2` which reference `symbol1` and `symbol2` from “filename”, respectively.

Another syntax is not part of ES6, but probably convenient:

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

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

Paths

In the above, `filename` is always treated as a path with `/` as directory separator, `.` 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 `x` from the same directory as the current file, use `import "./x" as x;`. If you use `import "x" as x;` instead, a different file could be referenced (in a global “include directory”).

It depends on the compiler (see below) how to actually resolve the paths. In general, the directory hierarchy does not need to strictly map onto your local filesystem, it can also map to resources discovered via e.g. ipfs, http or git.

Use in Actual Compilers

When the compiler is invoked, it is not only possible to specify how to discover the first element of a path, but it is possible to specify path prefix remappings so that e.g. `github.com/ethereum/dapp-bin/library` is remapped to `/usr/local/dapp-bin/library` and the compiler will read the files from there. If multiple remappings can be applied, the one with the longest key is tried first. This allows for a “fallback-remapping” with e.g. `"` maps to `/usr/local/include/solidity`. Furthermore, these remappings can depend on the 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), these remappings are provided as `context:prefix=target` arguments, where both the `context:` and the `=target` parts are optional (where `target` defaults to `prefix` in that case). All remapping values that are regular files are compiled (including their dependencies). This mechanism is completely backwards-compatible (as long as no filename contains `=` or `:`) and thus not a breaking change. All imports in files in or below the directory `context` that import a file that starts with `prefix` are redirected by replacing `prefix` by `target`.

So as an 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;
```

and then run the compiler as

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

As a more complex example, suppose you rely on some module that uses a very old version of dapp-bin. That old version of dapp-bin is checked out at `/usr/local/dapp-bin_old`, then you can use

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

so that all imports in `module2` point to the old version but imports in `module1` get the new version.

Note that `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, just 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 will also automatically retrieve the file over the network: You can import the iterable mapping by e.g. `import "github.com/ethereum/dapp-bin/library/iterable_mapping.sol" as it_mapping;`

Other source code providers may be added in the future.

Comments

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

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

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

Additionally, there is another type of comment called a natspec comment, for which the documentation is not yet written. 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 input parameters and two returned values.

```
pragma solidity ^0.4.0;

/** @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) returns (uint s, uint p) {  
        s = w * h;  
        p = 2 * (w + h);  
    }  
}
```

7.4.2 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*, *Structs Types* and *Enum Types*. Furthermore, contracts can inherit from other contracts.

State Variables

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

```
pragma solidity ^0.4.0;  
  
contract SimpleStorage {  
    uint storedData; // State variable  
    // ...  
}
```

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

Functions

Functions are the executable units of code within a contract.

```
pragma solidity ^0.4.0;  
  
contract SimpleAuction {  
    function bid() payable { // Function  
        // ...  
    }  
}
```

Function Calls can happen internally or externally and have different levels of visibility (*Visibility and Getters*) towards other contracts.

Function Modifiers

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

```
pragma solidity ^0.4.11;  
  
contract Purchase {  
    address public seller;  
  
    modifier onlySeller() { // Modifier
```

```

    require(msg.sender == seller);
    _;
}

function abort() onlySeller { // Modifier usage
    // ...
}
}

```

Events

Events are convenience interfaces with the EVM logging facilities.

```

pragma solidity ^0.4.0;

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

    function bid() payable {
        // ...
        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.

Structs Types

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

```

pragma solidity ^0.4.0;

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

```

Enum Types

Enums can be used to create custom types with a finite set of values (see [Enums](#) in types section).

```

pragma solidity ^0.4.0;

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

```

7.4.3 Types

Solidity is a statically typed language, which means that the type of each variable (state and local) needs to be specified (or at least known - see *Type Deduction* below) at compile-time. 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*.

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)
- Arithmetic operators: `+`, `-`, unary `-`, unary `+`, `*`, `/`, `%` (remainder), `**` (exponentiation), `<<` (left shift), `>>` (right shift)

Division always truncates (it is just compiled to the `DIV` opcode of the EVM), but it does not truncate if both operators are *literals* (or literal expressions).

Division by zero and modulus with zero throws a runtime exception.

The result of a shift operation is the type of the left operand. The expression `x << y` is equivalent to `x * 2**y`, and `x >> y` is equivalent to `x / 2**y`. This means that shifting negative numbers sign extends. Shifting by a negative amount throws a runtime exception.

Warning: The results produced by shift right of negative values of signed integer types is different from those produced by other programming languages. In Solidity, shift right maps to division so the shifted negative values are going to be rounded towards zero (truncated). In other programming languages the shift right of negative values works like division with rounding down (towards negative infinity).

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` represent the number of bits taken by the type and `N` represent 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 `ufixed128x19` and `fixed128x19`, respectively.

Operators:

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

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

`address`: Holds a 20 byte value (size of an Ethereum address). Address types also have members and serve as a base for all contracts.

Operators:

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

Note: Starting with version 0.5.0 contracts do not derive from the address type, but can still be explicitly converted to address.

Members of Addresses

- `balance` and `transfer`

For a quick reference, see [Address Related](#).

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

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

Note: If `x` is a contract address, its code (more specifically: 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`, `callcode` and `delegatecall`

Furthermore, to interface with contracts that do not adhere to the ABI, the function `call` is provided which takes an arbitrary number of arguments of any type. These arguments are padded to 32 bytes and concatenated. One exception is the case where the first argument is encoded to exactly four bytes. In this case, it is not padded to allow the use of function signatures here.

```
address nameReg = 0x72ba7d8e73fe8eb666ea66bab8116a41bfb10e2;
nameReg.call("register", "MyName");
nameReg.call(bytes4(keccak256("fun(uint256)")), a);
```

`call` returns a boolean indicating whether the invoked function terminated (`true`) or caused an EVM exception (`false`). It is not possible to access the actual data returned (for this we would need to know the encoding and size in advance).

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

```
nameReg.call.gas(1000000)("register", "MyName");
```

Similarly, the supplied Ether value can be controlled too:

```
nameReg.call.value(1 ether)("register", "MyName");
```

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

```
nameReg.call.gas(1000000).value(1 ether)("register", "MyName");
```

Note: It is not yet possible to use the gas or value modifiers on overloaded functions.

A workaround is to introduce a special case for gas and value and just re-check whether they are present at the point of overload resolution.

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. 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.

All three functions `call`, `delegatecall` and `callcode` 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: All contracts inherit the members of `address`, so it is possible to query the balance of the current contract using `this.balance`.

Note: The use of `callcode` is discouraged and will be removed in the future.

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.

Fixed-size byte arrays

`bytes1`, `bytes2`, `bytes3`, ..., `bytes32`. `byte` is an alias for `bytes1`.

Operators:

- Comparisons: `<=`, `<`, `==`, `!=`, `>=`, `>` (evaluate to `bool`)
- Bit operators: `&`, `|`, `^` (bitwise exclusive or), `~` (bitwise negation), `<<` (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 any integer type as right operand (but will return the type of the left operand), which denotes the number of bits to shift by. Shifting by a negative amount will cause a runtime exception.

Members:

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

Note: It is possible to use an array of bytes as `byte[]`, but it is wasting a lot of space, 31 bytes every element, to be exact, when passing in calls. It is better to use `bytes`.

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!

As a rule of thumb, 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 `bytes1` to `bytes32` because they are much cheaper.

Address Literals

Hexadecimal literals that pass the address checksum test, for example `0xCad3a6d3569DF655070DEd06cb7A1b2Ccd1D3AF` are of `address` type. Hexadecimal literals that are between 39 and 41 digits long and do not pass the checksum test produce a warning and are treated as regular rational number literals.

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, while the exponent cannot. Examples include `2e10`, `-2e10`, `2e-10`, `2.5e1`.

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). 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).

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.

Warning: Division on integer literals used to truncate in earlier versions, but it will now convert into a rational number, i.e. `5 / 2` is not equal to `2`, but to `2.5`.

Note: Number literal expressions are converted into a non-literal type as soon as they are used with non-literal expressions. Even though we know that the value of the expression assigned to `b` in the following example evaluates to an integer, but the partial expression `2.5 + a` does not type check so the code does not compile

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


String Literals

String literals are written with either double or single-quotes ("foo" or 'bar'). 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`.

String literals support escape characters, such as `\n`, `\xNN` and `\uNNNN`. `\xNN` takes a hex value and inserts the appropriate byte, while `\uNNNN` takes a Unicode codepoint and inserts an UTF-8 sequence.

Hexadecimal Literals

Hexadecimal Literals are prefixed with the keyword `hex` and are enclosed in double or single-quotes (`hex"001122FF"`). Their content must be a hexadecimal string and their value will be the binary representation of those values.

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 conversions check the value ranges at runtime and a failure causes an exception. Enums needs at least one member.

```
pragma solidity ^0.4.0;

contract test {
    enum ActionChoices { GoLeft, GoRight, GoStraight, SitStill }
    ActionChoices choice;
    ActionChoices constant defaultChoice = ActionChoices.GoStraight;

    function setGoStraight() {
        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 values,
    // `uint16` will be used and so on.
    function getChoice() returns (ActionChoices) {
        return choice;
    }

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

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|constant|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. In contrast, contract functions themselves are public by default, only when used as the name of a type, the default is internal.

There are two ways to access a function in the current contract: Either directly by its name, `f`, or using `this.f`. The former will result in an internal function, the latter in an external function.

If a function type variable is not initialized, calling it will result in an exception. 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`.

Additionally, public (or external) functions also have a special member called `selector`, which returns the *ABI function selector*:

```
pragma solidity ^0.4.0;  
  
contract Selector {  
    function f() returns (bytes4) {  
        return this.f.selector;  
    }  
}
```

Example that shows how to use internal function types:

```
pragma solidity ^0.4.5;  
  
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) returns (uint) f)  
        internal  
        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) returns (uint) f
  )
  internal
  returns (uint r)
  {
    r = self[0];
    for (uint i = 1; i < self.length; i++) {
      r = f(r, self[i]);
    }
  }
  function range(uint length) internal 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) returns (uint) {
    return ArrayUtils.range(l).map(square).reduce(sum);
  }
  function square(uint x) internal returns (uint) {
    return x * x;
  }
  function sum(uint x, uint y) internal returns (uint) {
    return x + y;
  }
}

```

Another example that uses external function types:

```

pragma solidity ^0.4.11;

contract Oracle {
  struct Request {
    bytes data;
    function(bytes memory) external callback;
  }
  Request[] requests;
  event NewRequest(uint);
  function query(bytes data, function(bytes memory) external callback) {
    requests.push(Request(data, callback));
    NewRequest(requests.length - 1);
  }
  function reply(uint requestID, bytes response) {
    // Here goes the check that the reply comes from a trusted source
    requests[requestID].callback(response);
  }
}

contract OracleUser {
  Oracle constant oracle = Oracle(0x1234567); // known contract
  function buySomething() {
    oracle.query("USD", this.oracleResponse);
  }
  function oracleResponse(bytes response) {

```

```
require(msg.sender == address(oracle));
// Use the data
}
}
```

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

Reference Types

Complex types, i.e. types which do not always fit into 256 bits have to be handled more carefully than the value-types we have already seen. Since copying them can be quite expensive, we have to think about whether we want them to be stored in **memory** (which is not persisting) or **storage** (where the state variables are held).

Data location

Every complex type, i.e. *arrays* and *structs*, has an additional annotation, the “data location”, about whether it is stored in memory or in storage. Depending on the context, there is always a default, but it can be overridden by appending either `storage` or `memory` to the type. The default for function parameters (including return parameters) is `memory`, the default for local variables is `storage` and the location is forced to `storage` for state variables (obviously).

There is also a third data location, `calldata`, which is a non-modifiable, non-persistent area where function arguments are stored. Function parameters (not return parameters) of external functions are forced to `calldata` and behave mostly like `memory`.

Data locations are important because they change how assignments behave: assignments between storage and memory and also to a state variable (even from other state variables) always create an independent copy. Assignments to local storage variables only assign a reference though, and this reference always points to the state variable even if the latter is changed in the meantime. On the other hand, assignments from a memory stored reference type to another memory-stored reference type do not create a copy.

```
pragma solidity ^0.4.0;

contract C {
    uint[] x; // the data location of x is storage

    // the data location of memoryArray is memory
    function f(uint[] memoryArray) {
        x = memoryArray; // works, copies the whole array to storage
        var y = x; // works, assigns a pointer, data location of y is storage
        y[7]; // fine, returns the 8th element
        y.length = 2; // 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 storageArray) internal {}
function h(uint[] memoryArray) {}
}
```

Summary

Forced data location:

- parameters (not return) of external functions: `calldata`
- state variables: `storage`

Default data location:

- parameters (also return) of functions: `memory`
- all other local variables: `storage`

Arrays

Arrays can have a compile-time fixed size or they can be dynamic. For storage arrays, the element type can be arbitrary (i.e. also other arrays, mappings or structs). For memory arrays, it cannot be a mapping and has to be an ABI type if it is an argument of a publicly-visible function.

An array of fixed size k and element type T is written as $T[k]$, an array of dynamic size as $T[]$. As an example, an array of 5 dynamic arrays of `uint` is `uint[][5]` (note that the notation is reversed when compared to some other languages). To access the second `uint` in the third dynamic array, you use `x[2][1]` (indices are zero-based and access works in the opposite way of the declaration, i.e. `x[2]` shaves off one level in the type from the right).

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

So `bytes` should always be preferred over `byte[]` because it is 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!

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

Allocating Memory Arrays

Creating arrays with variable length in memory can be done using the `new` keyword. As opposed to storage arrays, it is **not** possible to resize memory arrays by assigning to the `.length` member.

```
pragma solidity ^0.4.0;

contract C {
    function f(uint len) {
        uint[] memory a = new uint[](7);
        bytes memory b = new bytes(len);
        // Here we have a.length == 7 and b.length == len
        a[6] = 8;
    }
}
```

```
}  
}
```

Array Literals / Inline Arrays

Array literals are arrays that are written as an expression and are not assigned to a variable right away.

```
pragma solidity ^0.4.0;  
  
contract C {  
    function f() {  
        g([uint(1), 2, 3]);  
    }  
    function g(uint[3] _data) {  
        // ...  
    }  
}
```

The type of an array literal is a memory array of fixed size whose base type is the common type of the given elements. The type of `[1, 2, 3]` is `uint8[3]` memory, because the type of each of these constants is `uint8`. Because of that, it was necessary to convert the first element in the example above to `uint`. Note that currently, fixed size memory arrays cannot be assigned to dynamically-sized memory arrays, i.e. the following is not possible:

```
// This will not compile.  
  
pragma solidity ^0.4.0;  
  
contract C {  
    function f() {  
        // The next line creates a type error because uint[3] memory  
        // cannot be converted to uint[] memory.  
        uint[] x = [uint(1), 3, 4];  
    }  
}
```

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

Members

length: Arrays have a `length` member to hold their number of elements. Dynamic arrays can be resized in storage (not in memory) by changing the `.length` member. This does not happen automatically when attempting to access elements outside the current length. The size 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 can be used to append an element at the end of the array. The function returns the new length.

Warning: It is not yet possible to use arrays of arrays in external functions.

Warning: Due to limitations of the EVM, it is not possible to return dynamic content from external function calls. The function `f` in contract `C { function f() returns (uint[]) { ... } }` will return something if called from `web3.js`, but not if called from Solidity.

The only workaround for now is to use large statically-sized arrays.

```
pragma solidity ^0.4.0;

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).
    bool[2][] m_pairsOfFlags;
    // newPairs is stored in memory - the default for function arguments

    function setAllFlagPairs(bool[2][] newPairs) {
        // assignment to a storage array replaces the complete array
        m_pairsOfFlags = newPairs;
    }

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

    function changeFlagArraySize(uint newSize) {
        // if the new size is smaller, removed array elements will be cleared
        m_pairsOfFlags.length = newSize;
    }

    function clear() {
        // these clear the arrays completely
        delete m_pairsOfFlags;
        delete m_aLotOfIntegers;
        // identical effect here
        m_pairsOfFlags.length = 0;
    }

    bytes m_byteData;

    function byteArray(bytes data) {
        // byte arrays ("bytes") are different as they are stored without padding,
        // but can be treated identical to "uint8[]"
        m_byteData = data;
        m_byteData.length += 7;
        m_byteData[3] = 8;
        delete m_byteData[2];
    }

    function addFlag(bool[2] flag) returns (uint) {
        return m_pairsOfFlags.push(flag);
    }

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

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

Structs

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

```
pragma solidity ^0.4.11;

contract CrowdFunding {
    // Defines a new type with two fields.
    struct Funder {
        address addr;
        uint amount;
    }

    struct Campaign {
        address beneficiary;
        uint fundingGoal;
        uint numFunders;
        uint amount;
        mapping (uint => Funder) funders;
    }

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

    function newCampaign(address beneficiary, uint goal) returns (uint campaignID) {
        campaignID = numCampaigns++; // campaignID is return variable
        // Creates new struct and saves in storage. We leave out the mapping type.
        campaigns[campaignID] = Campaign(beneficiary, goal, 0, 0);
    }

    function contribute(uint campaignID) 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) 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 itself 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. 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 (of the default storage data location). 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`.

Mappings

Mapping types are declared as `mapping(_KeyType => _ValueType)`. Here `_KeyType` can be almost any type except for a mapping, a dynamically sized array, a contract, an enum and a struct. `_ValueType` can actually be any type, including mappings.

Mappings can be seen as [hash tables](#) which are virtually initialized 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 here, though: The key data is not actually stored in a mapping, only its `keccak256` hash 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”.

Mappings are only allowed for state variables (or as storage reference types in internal functions).

It is possible to mark mappings `public` and have Solidity create a *getter*. The `_KeyType` will become a required parameter for the getter and it will return `_ValueType`.

The `_ValueType` can be a mapping too. The getter will have one parameter for each `_KeyType`, recursively.

```
pragma solidity ^0.4.0;

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

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

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

Note: Mappings are not iterable, but it is possible to implement a data structure on top of them. For an example, see [iterable mapping](#).

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 `--`, `*=`, `/=`, `%=`, `a |=`, `&=` 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 reset. For structs, it assigns a struct with all members reset.

`delete` has no effect on whole 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.

It is important to note that `delete a` really behaves like an assignment to `a`, i.e. it stores a new object in `a`.

```
pragma solidity ^0.4.0;

contract DeleteExample {
    uint data;
    uint[] dataArray;

    function f() {
        uint x = data;
        delete x; // sets x to 0, does not affect data
        delete data; // sets data to 0, does not affect x which still holds a copy
        uint[] 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.
    }
}
```

Conversions between Elementary Types

Implicit Conversions

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). In general, an implicit conversion between value-types is possible if it makes sense semantically and no information is lost: `uint8` is convertible to `uint16` and `int128` to `int256`, but `int8` is not convertible to `uint256` (because `uint256` cannot hold e.g. `-1`). Furthermore, unsigned integers can be converted to bytes of the same or larger size, but not vice-versa. Any type that can be converted to `uint160` can also be converted to `address`.

Explicit Conversions

If the compiler does not allow implicit conversion but you know what you are doing, an explicit type conversion is sometimes possible. Note that this may give you some unexpected behaviour so be sure to test to ensure that the result is what you want! Take the following example where you are converting a negative `int8` to a `uint`:

```
int8 y = -3;
uint x = uint(y);
```

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

If a type 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
```

Type Deduction

For convenience, it is not always necessary to explicitly specify the type of a variable, the compiler automatically infers it from the type of the first expression that is assigned to the variable:

```
uint24 x = 0x123;
var y = x;
```

Here, the type of `y` will be `uint24`. Using `var` is not possible for function parameters or return parameters.

Warning: The type is only deduced from the first assignment, so the loop in the following snippet is infinite, as `i` will have the type `uint8` and the highest value of this type is smaller than 2000. `for (var i = 0; i < 2000; i++) { ... }`

7.4.4 Units and Globally Available Variables

Ether Units

A literal number can take a suffix of `wei`, `finney`, `szabo` or `ether` to convert between the subdenominations of Ether, where Ether currency numbers without a postfix are assumed to be Wei, e.g. `2 ether == 2000 finney` evaluates to `true`.

Time Units

Suffixes like `seconds`, `minutes`, `hours`, `days`, `weeks` and `years` after literal numbers can be used to convert between 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`

- 1 years == 365 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.

These suffixes cannot be applied to variables. If you want to interpret some input variable in e.g. days, you can do it in the following way:

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

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.

Block and Transaction Properties

- `block.blockhash(uint blockNumber)` returns (bytes32): hash of the given block - only works for 256 most recent blocks excluding current
- `block.coinbase(address)`: 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
- `msg.data(bytes)`: complete calldata
- `msg.gas(uint)`: remaining gas
- `msg.sender(address)`: 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
- `now(uint)`: current block timestamp (alias for `block.timestamp`)
- `tx.gasprice(uint)`: gas price of the transaction
- `tx.origin(address)`: 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`, `now` and `block.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.

Error Handling

assert(bool condition): throws if the condition is not met - to be used for internal errors.

require(bool condition): throws if the condition is not met - to be used for errors in inputs or external components.

revert(): abort execution and revert state changes

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} .

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} .

keccak256(...) returns (bytes32): compute the Ethereum-SHA-3 (Keccak-256) hash of the *(tightly packed) arguments*

sha256(...) returns (bytes32): compute the SHA-256 hash of the *(tightly packed) arguments*

sha3(...) returns (bytes32): alias to keccak256

ripemd160(...) returns (bytes20): compute RIPEMD-160 hash of the *(tightly packed) arguments*

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 (*example usage*)

In the above, “tightly packed” means that the arguments are concatenated without padding. This means that the following are all identical:

```
keccak256("ab", "c")
keccak256("abc")
keccak256(0x616263)
keccak256(6382179)
keccak256(97, 98, 99)
```

If padding is needed, explicit type conversions can be used: `keccak256("\x00\x12")` is the same as `keccak256(uint16(0x12))`.

Note that constants will be packed using the minimum number of bytes required to store them. This means that, for example, `keccak256(0) == keccak256(uint8(0))` and `keccak256(0x12345678) == keccak256(uint32(0x12345678))`.

It might be that you run into Out-of-Gas for `sha256`, `ripemd160` or `ecrecover` on a *private blockchain*. The reason for this is that those are implemented as so-called precompiled contracts and these contracts only really exist after they received the first message (although their contract code is hardcoded). Messages to non-existing contracts are more expensive and thus the execution runs into an Out-of-Gas error. A workaround for this problem is to first send e.g. 1 Wei to each of the contracts before you use them in your actual contracts. This is not an issue on the official or test net.

Address Related

`<address>.balance(uint256)`: balance of the *Address* in Wei

`<address>.transfer(uint256 amount)`: send given amount of Wei to *Address*, throws on failure

`<address>.send(uint256 amount)` returns `(bool)`: send given amount of Wei to *Address*, returns `false` on failure

`<address>.call(...)` returns `(bool)`: issue low-level CALL, returns `false` on failure

`<address>.callcode(...)` returns `(bool)`: issue low-level CALLCODE, returns `false` on failure

`<address>.delegatecall(...)` returns `(bool)`: issue low-level DELEGATECALL, returns `false` on failure

For more information, see the section on *Address*.

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: The use of `callcode` is discouraged and will be removed in the future.

Contract Related

`this` (current contract's type): the current contract, explicitly convertible to *Address*

`selfdestruct(address recipient)`: destroy the current contract, sending its funds to the given *Address*

`suicide(address recipient)`: alias to `selfdestruct`

Furthermore, all functions of the current contract are callable directly including the current function.

7.4.5 Expressions and Control Structures

Input Parameters and Output Parameters

As in Javascript, functions may take parameters as input; unlike in Javascript and C, they may also return arbitrary number of parameters as output.

Input Parameters

The input parameters are declared the same way as variables are. As an exception, unused parameters can omit the variable name. For example, suppose we want our contract to accept one kind of external calls with two integers, we would write something like:

```
pragma solidity ^0.4.0;

contract Simple {
    function taker(uint _a, uint _b) {
        // do something with _a and _b.
    }
}
```

Output Parameters

The output parameters can be declared with the same syntax after the `returns` keyword. For example, suppose we wished to return two results: the sum and the product of the two given integers, then we would write:

```
pragma solidity ^0.4.0;

contract Simple {
    function arithmetics(uint _a, uint _b) returns (uint o_sum, uint o_product) {
        o_sum = _a + _b;
        o_product = _a * _b;
    }
}
```

The names of output parameters can be omitted. The output values can also be specified using `return` statements. The `return` statements are also capable of returning multiple values, see [Returning Multiple Values](#). Return parameters are initialized to zero; if they are not explicitly set, they stay to be zero.

Input parameters and output parameters can be used as expressions in the function body. There, they are also usable in the left-hand side of assignment.

Control Structures

Most of the control structures from JavaScript are available in Solidity except for `switch` and `goto`. So there is: `if`, `else`, `while`, `do`, `for`, `break`, `continue`, `return`, `?:`, with the usual semantics known from C or JavaScript.

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.

Returning Multiple Values

When a function has multiple output parameters, `return (v0, v1, ..., vn)` can return multiple values. The number of components must be the same as the number of output parameters.

Function Calls

Internal Function Calls

Functions of the current contract can be called directly (“internally”), also recursively, as seen in this nonsensical example:

```
pragma solidity ^0.4.0;

contract C {
    function g(uint a) returns (uint ret) { return f(); }
    function f() 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 can be called internally.

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.

When calling functions of other contracts, the amount of Wei sent with the call and the gas can be specified with special options `.value()` and `.gas()`, respectively:

```
pragma solidity ^0.4.0;

contract InfoFeed {
    function info() payable returns (uint ret) { return 42; }
}

contract Consumer {
    InfoFeed feed;
    function setFeed(address addr) { feed = InfoFeed(addr); }
    function callFeed() { feed.info.value(10).gas(800)(); }
}
```

The modifier `payable` has to be used for `info`, because otherwise, the `.value()` option would not be available.

Note that the expression `InfoFeed(addr)` performs an explicit type conversion stating that “we know that the type of the contract at the given address is `InfoFeed`” and this does not execute a constructor. Explicit type conversions have to be handled with extreme caution. Never call a function on a contract where you are not sure about its type.

We could also have used `function setFeed(InfoFeed _feed) { feed = _feed; }` directly. Be careful about the fact that `feed.info.value(10).gas(800)` only (locally) sets the value and amount of gas sent with the function call and only the parentheses at the end perform the actual call.

Function calls cause exceptions if the called contract does not exist (in the sense that the account does not contain code) or 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.

Named Calls and Anonymous Function Parameters

Function call arguments can also 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.

```
pragma solidity ^0.4.0;

contract C {
    function f(uint key, uint value) {
        // ...
    }

    function g() {
        // named arguments
        f({value: 2, key: 3});
    }
}
```

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.

```
pragma solidity ^0.4.0;

contract C {
    // omitted name for parameter
    function func(uint k, uint) returns(uint) {
        return k;
    }
}
```

Creating Contracts via new

A contract can create a new contract using the `new` keyword. The full code of the contract being created has to be known in advance, so recursive creation-dependencies are not possible.

```
pragma solidity ^0.4.0;

contract D {
```

```
uint x;
function D(uint a) payable {
    x = a;
}

contract C {
    D d = new D(4); // will be executed as part of C's constructor

    function createdD(uint arg) {
        D newD = new D(arg);
    }

    function createAndEndowD(uint arg, uint amount) payable {
        // Send ether along with the creation
        D newD = (new D).value(amount)(arg);
    }
}
```

As seen in the example, it is possible to forward 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.

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. See *Order of Precedence of Operators* for more information.

Assignment

Destructuring Assignments and Returning Multiple Values

Solidity internally allows tuple types, i.e. a list of objects of potentially different types whose size is a constant at compile-time. Those tuples can be used to return multiple values at the same time and also assign them to multiple variables (or LValues in general) at the same time:

```
pragma solidity ^0.4.0;

contract C {
    uint[] data;

    function f() returns (uint, bool, uint) {
        return (7, true, 2);
    }

    function g() {
        // Declares and assigns the variables. Specifying the type explicitly is not
        ↪ possible.
        var (x, b, y) = f();
        // Assigns to a pre-existing variable.
        (x, y) = (2, 7);
        // 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).
// If the tuple ends in an empty component,
// the rest of the values are discarded.
(data.length,) = f(); // Sets the length to 7
// The same can be done on the left side.
(,data[3]) = f(); // Sets data[3] to 2
// Components can only be left out at the left-hand-side of assignments, with
// one exception:
(x,) = (1,);
// (1,) is the only way to specify a 1-component tuple, because (1) is
// equivalent to 1.
}

```

Complications for Arrays and Structs

The semantics of assignment are a bit more complicated for non-value types like arrays and structs. Assigning *to* a state variable always creates an independent copy. On the other hand, assigning to a local variable creates an independent copy only for elementary types, i.e. static types that fit into 32 bytes. If structs or arrays (including `bytes` and `string`) are assigned from a state variable to a local variable, the local variable holds a reference to the original state variable. A second assignment to the local variable does not modify the state but only changes the reference. Assignments to members (or elements) of the local variable *do* change the state.

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. Finally, for dynamically-sized arrays, `bytes` and `string`, the default value is an empty array or string.

A variable declared anywhere within a function will be in scope for the *entire function*, regardless of where it is declared. This happens because Solidity inherits its scoping rules from JavaScript. This is in contrast to many languages where variables are only scoped where they are declared until the end of the semantic block. As a result, the following code is illegal and cause the compiler to throw an error, `Identifier already declared`:

```

// This will not compile

pragma solidity ^0.4.0;

contract ScopingErrors {
    function scoping() {
        uint i = 0;

        while (i++ < 1) {
            uint same1 = 0;
        }

        while (i++ < 2) {
            uint same1 = 0; // Illegal, second declaration of same1
        }
    }

    function minimalScoping() {

```

```
    {
        uint same2 = 0;
    }

    {
        uint same2 = 0; // Illegal, second declaration of same2
    }
}

function forLoopScoping() {
    for (uint same3 = 0; same3 < 1; same3++) {
    }

    for (uint same3 = 0; same3 < 1; same3++) { // Illegal, second declaration of
↪ same3
    }
}
}
```

In addition to this, if a variable is declared, it will be initialized at the beginning of the function to its default value. As a result, the following code is legal, despite being poorly written:

```
pragma solidity ^0.4.0;

contract C {
    function foo() returns (uint) {
        // baz is implicitly initialized as 0
        uint bar = 5;
        if (true) {
            bar += baz;
        } else {
            uint baz = 10; // never executes
        }
        return bar; // returns 5
    }
}
```

Error handling: Assert, Require, Revert and Exceptions

Solidity uses state-reverting exceptions to handle errors. Such an exception will undo all changes made to the state in the current call (and all its sub-calls) and also flag an error to the caller. 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. The `require` function should be used to ensure valid conditions, such as inputs, or contract state variables are met, or to validate return values from calls to external contracts. If used properly, analysis tools can evaluate your contract to identify the conditions and function calls which will reach a failing `assert`. Properly functioning code should never reach a failing `assert` statement; if this happens there is a bug in your contract which you should fix.

There are two other ways to trigger exceptions: The `revert` function can be used to flag an error and revert the current call. In the future it might be possible to also include details about the error in a call to `revert`. The `throw` keyword can also be used as an alternative to `revert()`.

Note: From version 0.4.13 the `throw` keyword is deprecated and will be phased out in the future.

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 `callcode` – those return `false` in case of an exception instead of “bubbling up”.

Warning: The low-level `call`, `delegatecall` and `callcode` will return success if the calling account is non-existent, as part of the design of EVM. Existence must be checked prior to calling if desired.

Catching exceptions is not yet possible.

In the following example, you can see how `require` can be used to easily check conditions on inputs and how `assert` can be used for internal error checking:

```
pragma solidity ^0.4.0;

contract Sharer {
    function sendHalf(address addr) payable returns (uint balance) {
        require(msg.value % 2 == 0); // Only allow even numbers
        uint balanceBeforeTransfer = 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(this.balance == balanceBeforeTransfer - msg.value / 2);
        return this.balance;
    }
}
```

An `assert`-style exception is generated in the following situations:

1. If you access an array 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 shift by a negative amount.
5. If you convert a value too big or negative into an enum type.
6. If you call a zero-initialized variable of internal function type.
7. If you call `assert` with an argument that evaluates to false.

A `require`-style exception is generated in the following situations:

1. Calling `throw`.
2. Calling `require` with an argument that evaluates to false.
3. 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` or `callcode` is used. The low level operations never throw exceptions but indicate failures by returning `false`.
4. If you create a contract using the `new` keyword but the contract creation does not finish properly (see above for the definition of “not finish properly”).
5. If you perform an external function call targeting a contract that contains no code.
6. If your contract receives Ether via a public function without `payable` modifier (including the constructor and the fallback function).
7. If your contract receives Ether via a public getter function.

8. If a `.transfer()` fails.

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 retain the atomicity of transactions, the safest thing to do is to revert all changes and make the whole transaction (or at least call) without effect. Note that `assert`-style exceptions consume all gas available to the call, while `require`-style exceptions will not consume any gas starting from the Metropolis release.

7.4.6 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 are inaccessible.

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.

Creating contracts programatically on Ethereum is best done via using the JavaScript API [web3.js](#). As of today it has a method called `web3.eth.Contract` to facilitate contract creation.

When a contract is created, its constructor (a function with the same name as the contract) is executed once. A constructor is optional. Only one constructor is allowed, and this means overloading is not supported.

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.

```
pragma solidity ^0.4.0;

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.
    function OwnedToken(bytes32 _name) {
        // State variables are accessed via their name
        // and not via e.g. this.owner. This also applies
        // to functions and especially in the constructors,
        // you can only call them like that ("internally"),
        // because the contract itself does not exist yet.
        owner = msg.sender;
        // We do 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 check that.
    }
}
```

```

    creator = TokenCreator(msg.sender);
    name = _name;
}

function changeName(bytes32 newName) {
    // Only the creator can alter the name --
    // the comparison is possible since contracts
    // are implicitly convertible to addresses.
    if (msg.sender == address(creator))
        name = newName;
}

function transfer(address newOwner) {
    // Only the current owner can transfer the token.
    if (msg.sender != owner) return;
    // We also want to ask the creator if the transfer
    // is fine. Note that this calls a function of the
    // contract defined below. If the call fails (e.g.
    // due to out-of-gas), the execution here stops
    // immediately.
    if (creator.isTokenTransferOK(owner, newOwner))
        owner = newOwner;
}
}

contract TokenCreator {
    function createToken(bytes32 name)
        returns (OwnedToken tokenAddress)
    {
        // Create a new Token contract and return its address.
        // From the JavaScript side, the return type is simply
        // "address", as this is the closest type available in
        // the ABI.
        return new OwnedToken(name);
    }

    function changeName(OwnedToken tokenAddress, bytes32 name) {
        // Again, the external type of "tokenAddress" is
        // simply "address".
        tokenAddress.changeName(name);
    }

    function isTokenTransferOK(
        address currentOwner,
        address newOwner
    ) returns (bool ok) {
        // Check some arbitrary condition.
        address tokenAddress = msg.sender;
        return (keccak256(newOwner) & 0xff) == (bytes20(tokenAddress) & 0xff);
    }
}

```

Visibility and Getters

Since 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), there are four types of visibilities for functions and state variables.

Functions can be specified as being `external`, `public`, `internal` or `private`, where the default is `public`. For state variables, `external` is not possible and the default is `internal`.

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.

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 external observers. Making something `private` only prevents other contracts from accessing and 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.

```
pragma solidity ^0.4.0;

contract C {
    function f(uint a) private 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`.

```
// This will not compile

pragma solidity ^0.4.0;

contract C {
    uint private data;

    function f(uint a) private returns (uint b) { return a + 1; }
    function setData(uint a) { data = a; }
    function getData() public returns (uint) { return data; }
    function compute(uint a, uint b) internal returns (uint) { return a+b; }
}

contract D {
    function readData() {
        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() {
    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`. The initialization of state variables can be done at declaration.

```

pragma solidity ^0.4.0;

contract C {
    uint public data = 42;
}

contract Caller {
    C c = new C();
    function f() {
        uint local = c.data();
    }
}

```

The getter functions have external visibility. If the symbol is accessed internally (i.e. without `this.`), it is evaluated as a state variable. If it is accessed externally (i.e. with `this.`), it is evaluated as a function.

```

pragma solidity ^0.4.0;

contract C {
    uint public data;
    function x() {
        data = 3; // internal access
        uint val = this.data(); // external access
    }
}

```

The next example is a bit more complex:

```

pragma solidity ^0.4.0;

contract Complex {
    struct Data {
        uint a;
        bytes3 b;
        mapping (uint => uint) map;
    }
    mapping (uint => mapping(bool => Data[])) public data;
}

```

It will generate a function of the following form:

```

function data(uint arg1, bool arg2, uint arg3) returns (uint a, bytes3 b) {
    a = data[arg1][arg2][arg3].a;
}

```

```
b = data[arg1][arg2][arg3].b;
}
```

Note that the mapping in the struct is omitted because there is no good way to provide the key for the mapping.

Function Modifiers

Modifiers can be used to easily change the behaviour of functions. For example, they can automatically check a condition prior to executing the function. Modifiers are inheritable properties of contracts and may be overridden by derived contracts.

```
pragma solidity ^0.4.11;

contract owned {
    function owned() { owner = msg.sender; }
    address 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);
        _;
    }
}

contract mortal is owned {
    // This contract inherits the "onlyOwner"-modifier from
    // "owned" and applies it to the "close"-function, which
    // causes that calls to "close" only have an effect if
    // they are made by the stored owner.
    function close() onlyOwner {
        selfdestruct(owner);
    }
}

contract priced {
    // Modifiers can receive arguments:
    modifier costs(uint price) {
        if (msg.value >= price) {
            _;
        }
    }
}

contract Register is priced, owned {
    mapping (address => bool) registeredAddresses;
    uint price;

    function Register(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() payable costs(price) {
    registeredAddresses[msg.sender] = true;
}

function changePrice(uint _price) onlyOwner {
    price = _price;
}
}

contract Mutex {
    bool locked;
    modifier noReentrancy() {
        require(!locked);
        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() noReentrancy returns (uint) {
        require(msg.sender.call());
        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).

Constant State Variables

State variables can be declared as `constant`. In this case, they have to be assigned from an expression which is a constant at compile time. Any expression that accesses storage, blockchain data (e.g. `now`, `this.balance` or `block.number`) or execution data (`msg.gas`) or make calls to external contracts are 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 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.

The compiler does not reserve a storage slot for these variables, and every occurrence is replaced by the respective constant expression (which might be computed to a single value by the optimizer).

Not all types for constants are implemented at this time. The only supported types are value types and strings.

```
pragma solidity ^0.4.0;

contract C {
    uint constant x = 32**22 + 8;
    string constant text = "abc";
    bytes32 constant myHash = keccak256("abc");
}
```

View Functions

Functions can be declared `view` in which case they promise not to modify the state.

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.

```
pragma solidity ^0.4.16;

contract C {
    function f(uint a, uint b) view returns (uint) {
        return a * (b + 42) + now;
    }
}
```

Note: `constant` is an alias to `view`.

Note: Getter methods are marked `view`.

Warning: The compiler does not enforce yet that a `view` method is not modifying state.

Pure Functions

Functions can be declared `pure` in which case they promise not to read from or modify the state.

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 `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.

```
pragma solidity ^0.4.16;

contract C {
    function f(uint a, uint b) pure returns (uint) {
        return a * (b + 42);
    }
}
```

Warning: The compiler does not enforce yet that a `pure` method is not reading from the state.

Fallback Function

A contract can have exactly one unnamed function. This function cannot have arguments and cannot return anything. It is executed on a call to the contract if none of the other functions match the given function identifier (or if no data was supplied at all).

Furthermore, this function is executed whenever the contract receives plain Ether (without data). Additionally, in order to receive Ether, the fallback function must be marked `payable`. If no such function exists, the contract cannot receive Ether through regular transactions.

In such a context, there is usually very little gas available to the function call (to be precise, 2300 gas), so it is important to make fallback functions as cheap as possible. Note that the gas required by a transaction (as opposed to an internal call) that invokes the fallback function is much higher, because each transaction charges an additional amount of 21000 gas or more for things like signature checking.

In particular, the following operations will consume more gas than the stipend provided to a fallback function:

- Writing to storage
- Creating a contract
- Calling an external function which consumes a large amount of gas
- Sending Ether

Please ensure you test your fallback function thoroughly to ensure the execution cost is less than 2300 gas before deploying a contract.

Note: Even though the fallback function cannot have arguments, one can still use `msg.data` to retrieve any payload supplied with the call.

Warning: Contracts that receive Ether directly (without a function call, i.e. using `send` or `transfer`) but do not define a 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 fallback function.

Warning: A contract without a payable fallback 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 `this.balance` can be higher than the sum of some manual accounting implemented in a contract (i.e. having a counter updated in the fallback function).

```
pragma solidity ^0.4.0;

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.
    function() { x = 1; }
    uint x;
}

// This contract keeps all Ether sent to it with no way
// to get it back.
contract Sink {
    function() payable { }
}

contract Caller {
    function callTest(Test test) {
        test.call(0xabcd01); // hash does not exist
        // results in test.x becoming == 1.

        // The following will not compile, but even
        // if someone sends ether to that contract,
        // the transaction will fail and reject the
        // Ether.
        //test.send(2 ether);
    }
}
```

Events

Events allow the convenient usage of the EVM logging facilities, which in turn can be used to “call” JavaScript callbacks in the user interface of a dapp, which listen for these events.

Events are inheritable members of contracts. When they are called, 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 and will be incorporated into the blockchain and stay there as long as a block is accessible (forever as of Frontier and Homestead, but this might change with Serenity). Log and event data is not accessible from within contracts (not even from the contract that created them).

SPV proofs for logs are possible, so if an external entity supplies a contract with such a proof, it can check that the log actually exists inside the blockchain. But be aware that block headers have to be supplied because the contract can only see the last 256 block hashes.

Up to three parameters can receive the attribute `indexed` which will cause the respective arguments to be searched for: It is possible to filter for specific values of indexed arguments in the user interface.

If arrays (including `string` and `bytes`) are used as indexed arguments, the Keccak-256 hash of it is stored as topic instead.

The hash of the signature of the event is one of the topics except if you declared the event with `anonymous` specifier. This means that it is not possible to filter for specific anonymous events by name.

All non-indexed arguments will be stored in the data part of the log.

Note: Indexed arguments will not be stored themselves. You can only search for the values, but it is impossible to retrieve the values themselves.

```
pragma solidity ^0.4.0;

contract ClientReceipt {
    event Deposit(
        address indexed _from,
        bytes32 indexed _id,
        uint _value
    );

    function deposit(bytes32 _id) payable {
        // Any call to this function (even deeply nested) can
        // be detected from the JavaScript API by filtering
        // for `Deposit` to be called.
        Deposit(msg.sender, _id, msg.value);
    }
}
```

The use in the JavaScript API would be as follows:

```
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 will contain various information
    // including the arguments 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);
});
```

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`. `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

```
pragma solidity ^0.4.10;

contract C {
    function f() {
        bytes32 _id = 0x420042;
        log3(
            bytes32(msg.value),
            bytes32(0x50cb9fe53daa9737b786ab3646f04d0150dc50ef4e75f59509d83667ad5adb20),
            bytes32(msg.sender),
            _id
        );
    }
}
```

where the long hexadecimal number is equal to `keccak256("Deposit(address,hash256,uint256)")`, the signature of the event.

Additional Resources for Understanding Events

- [Javascript documentation](#)
- [Example usage of events](#)
- [How to access them in js](#)

Inheritance

Solidity supports multiple inheritance by copying code including polymorphism.

All function calls are virtual, which means that the most derived function is called, except when the contract name is explicitly given.

When a contract inherits from multiple contracts, only a single contract is created on the blockchain, and the code from all the base contracts is copied into the created contract.

The general inheritance system is very similar to [Python's](#), especially concerning multiple inheritance.

Details are given in the following example.

```
pragma solidity ^0.4.0;

contract owned {
    function owned() { owner = msg.sender; }
    address 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 mortal is owned {
    function kill() {
        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.
contract Config {
    function lookup(uint id) returns (address adr);
}

contract NameReg {
    function register(bytes32 name);
    function unregister();
}

// Multiple inheritance is possible. Note that "owned" is
// also a base class of "mortal", yet there is only a single
// instance of "owned" (as for virtual inheritance in C++).
contract named is owned, mortal {
    function named(bytes32 name) {
        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.
    function kill() {
        if (msg.sender == owner) {
            Config config = Config(0xd5f9d8d94886e70b06e474c3fb14fd43e2f23970);
            NameReg(config.lookup(1)).unregister();
            // It is still possible to call a specific
            // overridden function.
            mortal.kill();
        }
    }
}

// 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, mortal, named("GoldFeed") {
    function updateInfo(uint newInfo) {
        if (msg.sender == owner) info = newInfo;
    }

    function get() constant returns (uint r) { return info; }

    uint info;
}

```

Note that above, we call `mortal.kill()` to “forward” the destruction request. The way this is done is problematic,

as seen in the following example:

```
pragma solidity ^0.4.0;

contract owned {
    function owned() { owner = msg.sender; }
    address owner;
}

contract mortal is owned {
    function kill() {
        if (msg.sender == owner) selfdestruct(owner);
    }
}

contract Base1 is mortal {
    function kill() { /* do cleanup 1 */ mortal.kill(); }
}

contract Base2 is mortal {
    function kill() { /* do cleanup 2 */ mortal.kill(); }
}

contract Final is Base1, Base2 {
}
```

A call to `Final.kill()` will call `Base2.kill` as the most derived override, but this function will bypass `Base1.kill`, basically because it does not even know about `Base1`. The way around this is to use `super`:

```
pragma solidity ^0.4.0;

contract owned {
    function owned() { owner = msg.sender; }
    address owner;
}

contract mortal is owned {
    function kill() {
        if (msg.sender == owner) selfdestruct(owner);
    }
}

contract Base1 is mortal {
    function kill() { /* do cleanup 1 */ super.kill(); }
}

contract Base2 is mortal {
    function kill() { /* do cleanup 2 */ super.kill(); }
}

contract Final is Base2, Base1 {
}
```

If `Base1` 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 `Base2.kill()` (note that the final inheritance sequence is – starting with the most derived contract: `Final`, `Base1`, `Base2`, `mortal`, `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.

Arguments for Base Constructors

Derived contracts need to provide all arguments needed for the base constructors. This can be done in two ways:

```
pragma solidity ^0.4.0;

contract Base {
    uint x;
    function Base(uint _x) { x = _x; }
}

contract Derived is Base(7) {
    function Derived(uint _y) Base(_y * _y) {
    }
}
```

One way is directly in the inheritance list (`is Base(7)`). The other is in the way a modifier would be invoked as part of the header 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. If, as in this silly example, both places are used, the modifier-style argument takes precedence.

Multiple Inheritance and Linearization

Languages that allow multiple inheritance have to deal with several problems. One is the [Diamond Problem](#). Solidity follows the path of Python and uses “[C3 Linearization](#)” to force a specific order in the 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. In the following code, Solidity will give the error “Linearization of inheritance graph impossible”.

```
// This will not compile

pragma solidity ^0.4.0;

contract X {}
contract A is X {}
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.

A simple rule to remember is to specify the base classes in the order from “most base-like” to “most derived”.

Inheriting Different Kinds of Members of the Same Name

When the inheritance results in a contract with a function and a modifier of the same name, it is considered as an error. This error is produced also by an event and a modifier of the same name, and a function and an event of the same name. As an exception, a state variable getter can override a public function.

Abstract Contracts

Contract functions can lack an implementation as in the following example (note that the function declaration header is terminated by `;`):

```
pragma solidity ^0.4.0;

contract Feline {
    function utterance() returns (bytes32);
}
```

Such contracts cannot be compiled (even if they contain implemented functions alongside non-implemented functions), but they can be used as base contracts:

```
pragma solidity ^0.4.0;

contract Feline {
    function utterance() returns (bytes32);
}

contract Cat is Feline {
    function utterance() returns (bytes32) { return "miaow"; }
}
```

If a contract inherits from an abstract contract and does not implement all non-implemented functions by overriding, it will itself be abstract.

Interfaces

Interfaces are similar to abstract contracts, but they cannot have any functions implemented. There are further restrictions:

1. Cannot inherit other contracts or interfaces.
2. Cannot define constructor.
3. Cannot define variables.
4. Cannot define structs.
5. Cannot define enums.

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:

```
pragma solidity ^0.4.11;

interface Token {
    function transfer(address recipient, uint amount);
}
```

Contracts can inherit interfaces as they would inherit other contracts.

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).

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 (`L.f()` if `L` is the name of the library). Furthermore, `internal` functions of libraries are visible in all contracts, just as if the library were a base contract. Of course, calls to `internal` functions use the `internal` calling convention, which means that all `internal` types can be passed and `memory` types 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 be pulled into the calling contract, and a regular `JUMP` call will be used instead of a `DELEGATECALL`.

The following example illustrates how to use libraries (but be sure to check out *using for* for a more advanced example to implement a set).

```
pragma solidity ^0.4.11;

library Set {
    // We define a new struct datatype that will be used to
    // hold its data in the calling contract.
    struct Data { mapping(uint => bool) flags; }

    // 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)
        returns (bool)
    {
        if (self.flags[value])
            return false; // already there
        self.flags[value] = true;
        return true;
    }

    function remove(Data storage self, uint value)
        returns (bool)
    {
        if (!self.flags[value])
            return false; // not there
        self.flags[value] = false;
        return true;
    }

    function contains(Data storage self, uint value)
        returns (bool)
    {
        return self.flags[value];
    }
}
```

```

contract C {
    Set.Data knownValues;

    function register(uint value) {
        // 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, take care 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 memory types and internal functions in libraries in order to implement custom types without the overhead of external function calls:

```

pragma solidity ^0.4.0;

library BigInt {
    struct bigint {
        uint[] limbs;
    }

    function fromUint(uint x) internal returns (bigint r) {
        r.limbs = new uint[](1);
        r.limbs[0] = x;
    }

    function add(bigint _a, bigint _b) internal returns (bigint r) {
        r.limbs = new uint[(max(_a.limbs.length, _b.limbs.length))];
        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)];
            for (i = 0; i < r.limbs.length; ++i)
                newLimbs[i] = r.limbs[i];
            newLimbs[i] = carry;
            r.limbs = newLimbs;
        }
    }
}

```

```

function limb(bigint _a, uint _limb) internal returns (uint) {
    return _limb < _a.limbs.length ? _a.limbs[_limb] : 0;
}

function max(uint a, uint b) private returns (uint) {
    return a > b ? a : b;
}
}

contract C {
    using BigInt for BigInt.bigint;

    function f() {
        var x = BigInt.fromUint(7);
        var y = BigInt.fromUint(uint(-1));
        var z = x.add(y);
    }
}

```

As the compiler cannot know where the library will be deployed at, these addresses have to be filled into the final bytecode by a linker (see *Using the Commandline Compiler* for how to use the commandline compiler for linking). If the addresses are not given as arguments to the compiler, the compiled hex code will contain placeholders of the form `__Set_____` (where `Set` is the name of the library). The address can be filled manually by replacing all those 40 symbols by the hex encoding of the address of the library contract.

Restrictions for libraries in comparison to contracts:

- No state variables
- Cannot inherit nor be inherited
- Cannot receive Ether

(These might be lifted at a later point.)

Using For

The directive `using A for B;` can be used to attach library functions (from the library `A`) to any type (`B`). 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, even those where the type of the first parameter does not match the type of the object, are attached. 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 for the current scope, which is limited to a contract for now but will be lifted to the global scope later, so that by including a module, its data types including library functions are available without having to add further code.

Let us rewrite the `set` example from the *Libraries* in this way:

```

pragma solidity ^0.4.11;

// This is the same code as before, just without comments
library Set {
    struct Data { mapping(uint => bool) flags; }

    function insert(Data storage self, uint value)
        returns (bool)

```

```
{
    if (self.flags[value])
        return false; // already there
    self.flags[value] = true;
    return true;
}

function remove(Data storage self, uint value)
    returns (bool)
{
    if (!self.flags[value])
        return false; // not there
    self.flags[value] = false;
    return true;
}

function contains(Data storage self, uint value)
    returns (bool)
{
    return self.flags[value];
}
}

contract C {
    using Set for Set.Data; // this is the crucial change
    Set.Data knownValues;

    function register(uint value) {
        // Here, all variables of type Set.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:

```
pragma solidity ^0.4.0;

library Search {
    function indexOf(uint[] storage self, uint value) 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) {
        data.push(value);
    }

    function replace(uint _old, uint _new) {
        // 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 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.

7.4.7 Solidity Assembly

Solidity defines an assembly language that can also be used without Solidity. This assembly language can also be used as “inline assembly” inside Solidity source code. We start with describing how to use inline assembly and how it differs from standalone assembly and then specify assembly itself.

Inline Assembly

For more fine-grained control especially in order to enhance the language by writing libraries, it is possible to interleave Solidity statements with inline assembly in a language close to the one of the virtual machine. Due to the fact that the EVM is a stack machine, it is often hard to address the correct stack slot and provide arguments to opcodes at the correct point on the stack. Solidity’s inline assembly tries to facilitate that and other issues arising when writing manual assembly by the following features:

- functional-style opcodes: `mul(1, add(2, 3))` instead of `push1 3 push1 2 add push1 1 mul`
- assembly-local variables: `let x := add(2, 3) let y := mload(0x40) x := add(x, y)`
- access to external variables: `function f(uint x) { assembly { x := sub(x, 1) } }`
- labels: `let x := 10 repeat: x := sub(x, 1) jumpi(repeat, eq(x, 0))`
- loops: `for { let i := 0 } lt(i, x) { i := add(i, 1) } { y := mul(2, y) }`
- if statements: `if slt(x, 0) { x := sub(0, x) }`
- switch statements: `switch x case 0 { y := mul(x, 2) } default { y := 0 }`
- function calls: `function f(x) -> y { switch x case 0 { y := 1 } default { y := mul(x, f(sub(x, 1))) } }`

We now want to describe the inline assembly language in detail.

Warning: Inline assembly is a way to access the Ethereum Virtual Machine at a low level. This discards several important safety features of Solidity.

Note: TODO: Write about how scoping rules of inline assembly are a bit different and the complications that arise when for example using internal functions of libraries. Furthermore, write about the symbols defined by the compiler.

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 at all with “plain Solidity” and the idea is that assembly libraries will be used to enhance the language in such ways.

```
pragma solidity ^0.4.0;

library GetCode {
    function at(address _addr) returns (bytes 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 could also be beneficial in cases where the optimizer fails to produce efficient code. Please be aware that assembly is much more difficult to write because the compiler does not perform checks, so you should use it for complex things only if you really know what you are doing.

```
pragma solidity ^0.4.12;

library VectorSum {
    // This function is less efficient because the optimizer currently fails to
    // remove the bounds checks in array access.
    function sumSolidity(uint[] _data) returns (uint o_sum) {
        for (uint i = 0; i < _data.length; ++i)
            o_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[] _data) returns (uint o_sum) {
        for (uint i = 0; i < _data.length; ++i) {
            assembly {
                o_sum := add(o_sum, mload(add(add(_data, 0x20), mul(i, 0x20))))
            }
        }
    }

    // Same as above, but accomplish the entire code within inline assembly.
    function sumPureAsm(uint[] _data) returns (uint o_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, len) }
  lt(data, end)
  { data := add(data, 0x20) }
  {
    o_sum := add(o_sum, mload(data))
  }
}
}

```

Syntax

Assembly parses comments, literals and identifiers exactly as Solidity, so you can use the usual `//` and `/* */` comments. Inline assembly is marked by `assembly { ... }` and inside these curly braces, the following can be used (see the later sections for more details)

- literals, i.e. `0x123`, `42` or `"abc"` (strings up to 32 characters)
- opcodes (in “instruction style”), e.g. `mload` `sload` `dup1` `sstore`, for a list see below
- opcodes in functional style, e.g. `add(1, mload(0))`
- labels, e.g. `name:`
- variable declarations, e.g. `let x := 7`, `let x := add(y, 3)` or `let x` (initial value of empty (0) is assigned)
- identifiers (labels or assembly-local variables and externals if used as inline assembly), e.g. `jump(name)`, `3` `x` `add`
- assignments (in “instruction style”), e.g. `3 =: x`
- assignments in functional style, e.g. `x := add(y, 3)`
- blocks where local variables are scoped inside, e.g. `{ let x := 3 { let y := add(x, 1) } }`

Opcodes

This document does not want to be a full description of the Ethereum virtual machine, but the following list can be used as a reference of its opcodes.

If an opcode takes arguments (always from the top of the stack), they are given in parentheses. Note that the order of arguments can be seen to be reversed in non-functional style (explained below). Opcodes marked with `-` do not push an item onto the stack, those marked with `*` are special and all others push exactly one item onto the stack.

In the following, `mem[a...b]` signifies the bytes of memory starting at position `a` up to (excluding) position `b` and `storage[p]` signifies the storage contents at position `p`.

The opcodes `pushi` and `jumpdest` cannot be used directly.

In the grammar, opcodes are represented as pre-defined identifiers.

stop	-	stop execution, identical to return(0,0)
add(x, y)		$x + y$
sub(x, y)		$x - y$
mul(x, y)		$x * y$
div(x, y)		x / y
sdiv(x, y)		x / y , for signed numbers in two's complement
mod(x, y)		$x \% y$
smod(x, y)		$x \% y$, for signed numbers in two's complement
exp(x, y)		x to the power of y
not(x)		$\sim x$, every bit of x is negated
lt(x, y)		1 if $x < y$, 0 otherwise
gt(x, y)		1 if $x > y$, 0 otherwise
slt(x, y)		1 if $x < y$, 0 otherwise, for signed numbers in two's complement
sgt(x, y)		1 if $x > y$, 0 otherwise, for signed numbers in two's complement
eq(x, y)		1 if $x == y$, 0 otherwise
iszero(x)		1 if $x == 0$, 0 otherwise
and(x, y)		bitwise and of x and y
or(x, y)		bitwise or of x and y
xor(x, y)		bitwise xor of x and y
byte(n, x)		nth byte of x, where the most significant byte is the 0th byte
addmod(x, y, m)		$(x + y) \% m$ with arbitrary precision arithmetics
mulmod(x, y, m)		$(x * y) \% m$ with arbitrary precision arithmetics
signextend(i, x)		sign extend from $(i * 8 + 7)$ th bit counting from least significant
keccak256(p, n)		keccak(mem[p...(p+n)])
sha3(p, n)		keccak(mem[p...(p+n)])
jump(label)	-	jump to label / code position
jumpi(label, cond)	-	jump to label if cond is nonzero
pc		current position in code
pop(x)	-	remove the element pushed by x
dup1 ... dup16		copy ith stack slot to the top (counting from top)
swap1 ... swap16	*	swap topmost and ith stack slot below it
mload(p)		mem[p..(p+32))
mstore(p, v)	-	mem[p..(p+32)) := v
mstore8(p, v)	-	mem[p] := v & 0xff - only modifies a single byte
sload(p)		storage[p]
sstore(p, v)	-	storage[p] := v
msize		size of memory, i.e. largest accessed memory index
gas		gas still available to execution
address		address of the current contract / execution context
balance(a)		wei balance at address a
caller		call sender (excluding delegatecall)
callvalue		wei sent together with the current call
calldataload(p)		call data starting from position p (32 bytes)
calldatasize		size of call data in bytes
calldatacopy(t, f, s)	-	copy s bytes from calldata at position f to mem at position t
codesize		size of the code of the current contract / execution context
codecopy(t, f, s)	-	copy s bytes from code at position f to mem at position t

Table 7.1 – continued from previous page

<code>extcodesize(a)</code>		size of the code at address <code>a</code>
<code>extcodecopy(a, t, f, s)</code>	-	like <code>codecopy(t, f, s)</code> but take code at address <code>a</code>
<code>returndatasize</code>		size of the last <code>returndata</code>
<code>returndatacopy(t, f, s)</code>	-	copy <code>s</code> bytes from <code>returndata</code> at position <code>f</code> to mem at position <code>t</code>
<code>create(v, p, s)</code>		create new contract with code mem[<code>p..(p+s)</code>] and send <code>v</code> wei and return the new address
<code>create2(v, n, p, s)</code>		create new contract with code mem[<code>p..(p+s)</code>] at address <code>keccak256(<address> . n . keccak256(v))</code>
<code>call(g, a, v, in, insize, out, outsize)</code>		call contract at address <code>a</code> with input mem[<code>in..(in+insize)</code>] providing <code>g</code> gas and <code>v</code> wei and return mem[<code>out..(out+outsize)</code>]
<code>callcode(g, a, v, in, insize, out, outsize)</code>		identical to <code>call</code> but only use the code from <code>a</code> and stay in the context of the current contract
<code>delegatecall(g, a, in, insize, out, outsize)</code>		identical to <code>callcode</code> but also keep <code>caller</code> and <code>callvalue</code>
<code>staticcall(g, a, in, insize, out, outsize)</code>		identical to <code>call(g, a, 0, in, insize, out, outsize)</code> but do not allow state modifications
<code>return(p, s)</code>	-	end execution, return data mem[<code>p..(p+s)</code>]
<code>revert(p, s)</code>	-	end execution, revert state changes, return data mem[<code>p..(p+s)</code>]
<code>selfdestruct(a)</code>	-	end execution, destroy current contract and send funds to <code>a</code>
<code>invalid</code>	-	end execution with invalid instruction
<code>log0(p, s)</code>	-	log without topics and data mem[<code>p..(p+s)</code>]
<code>log1(p, s, t1)</code>	-	log with topic <code>t1</code> and data mem[<code>p..(p+s)</code>]
<code>log2(p, s, t1, t2)</code>	-	log with topics <code>t1, t2</code> and data mem[<code>p..(p+s)</code>]
<code>log3(p, s, t1, t2, t3)</code>	-	log with topics <code>t1, t2, t3</code> and data mem[<code>p..(p+s)</code>]
<code>log4(p, s, t1, t2, t3, t4)</code>	-	log with topics <code>t1, t2, t3, t4</code> and data mem[<code>p..(p+s)</code>]
<code>origin</code>		transaction sender
<code>gasprice</code>		gas price of the transaction
<code>blockhash(b)</code>		hash of block nr <code>b</code> - only for last 256 blocks excluding current
<code>coinbase</code>		current mining beneficiary
<code>timestamp</code>		timestamp of the current block in seconds since the epoch
<code>number</code>		current block number
<code>difficulty</code>		difficulty of the current block
<code>gaslimit</code>		block gas limit of the current block

Literals

You can use integer constants by typing them in decimal or hexadecimal notation and an appropriate `PUSHi` instruction will automatically be generated. The following creates code to add 2 and 3 resulting in 5 and then computes the bitwise and with the string "abc". Strings are stored left-aligned and cannot be longer than 32 bytes.

```
assembly { 2 3 add "abc" and }
```

Functional Style

You can type opcode after opcode in the same way they will end up in bytecode. For example adding 3 to the contents in memory at position `0x80` would be

```
3 0x80 mload add 0x80 mstore
```

As it is often hard to see what the actual arguments for certain opcodes are, Solidity inline assembly also provides a “functional style” notation where the same code would be written as follows

```
mstore(0x80, add(mload(0x80), 3))
```

Functional style expressions cannot use instructional style internally, i.e. `1 2 mstore(0x80, add)` is not valid assembly, it has to be written as `mstore(0x80, add(2, 1))`. For opcodes that do not take arguments, the

parentheses can be omitted.

Note that the order of arguments is reversed in functional-style as opposed to the instruction-style way. If you use functional-style, the first argument will end up on the stack top.

Access to External Variables and Functions

Solidity variables and other identifiers can be accessed by simply using their name. For memory variables, this will push the address and not the value onto the stack. Storage variables are different: Values in storage might not occupy a full storage slot, so 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 used `x_slot` and to retrieve the byte-offset you used `x_offset`.

In assignments (see below), we can even use local Solidity variables to assign to.

Functions external to inline assembly can also be accessed: The assembly will push their entry label (with virtual function resolution applied). The calling semantics in solidity are:

- the caller pushes return label, `arg1`, `arg2`, ..., `argn`
- the call returns with `ret1`, `ret2`, ..., `retm`

This feature is still a bit cumbersome to use, because the stack offset essentially changes during the call, and thus references to local variables will be wrong.

```
pragma solidity ^0.4.11;

contract C {
    uint b;
    function f(uint x) returns (uint r) {
        assembly {
            r := mul(x, sload(b_slot)) // ignore the offset, we know it is zero
        }
    }
}
```

Labels

Another problem in EVM assembly is that `jump` and `jumpi` use absolute addresses which can change easily. Solidity inline assembly provides labels to make the use of jumps easier. Note that labels are a low-level feature and it is possible to write efficient assembly without labels, just using assembly functions, loops, if and switch instructions (see below). The following code computes an element in the Fibonacci series.

```
{
    let n := calldataload(4)
    let a := 1
    let b := a
loop:
    jumpi(loopend, eq(n, 0))
    a add swap1
    n := sub(n, 1)
    jump(loop)
loopend:
    mstore(0, a)
    return(0, 0x20)
}
```

Please note that automatically accessing stack variables can only work if the assembler knows the current stack height. This fails to work if the jump source and target have different stack heights. It is still fine to use such jumps, but you should just not access any stack variables (even assembly variables) in that case.

Furthermore, the stack height analyser goes through the code opcode by opcode (and not according to control flow), so in the following case, the assembler will have a wrong impression about the stack height at label `two`:

```
{
  let x := 8
  jump(two)
  one:
    // Here the stack height is 2 (because we pushed x and 7),
    // but the assembler thinks it is 1 because it reads
    // from top to bottom.
    // Accessing the stack variable x here will lead to errors.
    x := 9
    jump(three)
  two:
    7 // push something onto the stack
    jump(one)
  three:
}
```

This problem can be fixed by manually adjusting the stack height for the assembler - you can provide a stack height delta that is added to the stack height just prior to the label. Note that you will not have to care about these things if you just use loops and assembly-level functions.

As an example how this can be done in extreme cases, please see the following.

```
{
  let x := 8
  jump(two)
  0 // This code is unreachable but will adjust the stack height correctly
  one:
    x := 9 // Now x can be accessed properly.
    jump(three)
    pop // Similar negative correction.
  two:
    7 // push something onto the stack
    jump(one)
  three:
    pop // We have to pop the manually pushed value here again.
}
```

Declaring Assembly-Local Variables

You can use the `let` keyword to declare variables that are only visible in inline assembly and actually only in the current `{...}`-block. What happens is that the `let` instruction will create a new stack slot that is reserved for the variable and automatically removed again when the end of the block is reached. You need to provide an initial value for the variable which can be just 0, but it can also be a complex functional-style expression.

```
pragma solidity ^0.4.0;

contract C {
  function f(uint x) returns (uint b) {
    assembly {
      let v := add(x, 1)
    }
  }
}
```

```
mstore(0x80, v)
{
    let y := add(sload(v), 1)
    b := y
} // y is "deallocated" here
b := add(b, v)
} // v is "deallocated" here
}
```

Assignments

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.

There are two kinds of assignments: functional-style and instruction-style. For functional-style assignments (`variable := value`), you need to provide a value in a functional-style expression that results in exactly one stack value and for instruction-style (`=: variable`), the value is just taken from the stack top. For both ways, the colon points to the name of the variable. The assignment is performed by replacing the variable's value on the stack by the new value.

```
{
    let v := 0 // functional-style assignment as part of variable declaration
    let g := add(v, 2)
    sload(10)
    =: v // instruction style assignment, puts the result of sload(10) into v
}
```

If

The if statement can be used for conditionally executing code. There is no “else” part, consider using “switch” (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 a very basic version of “if/else”. It takes the value of an expression and compares it to several constants. The branch corresponding to the matching constant is taken. Contrary to the error-prone behaviour of some programming languages, control flow does not continue from one case to the next. There can be a fallback or default case called `default`.

```
{
    let x := 0
    switch calldataload(4)
    case 0 {
        x := calldataload(0x24)
    }
}
```



```

default {
    x := calldataload(0x44)
}
sstore(0, div(x, 2))
}

```

The list of cases does not require curly braces, but the body of a case does require them.

Loops

Assembly supports a simple for-style loop. For-style loops have a header containing an initializing part, a condition and a post-iteration part. The condition has to be a functional-style expression, while the other two are blocks. If the initializing part declares any variables, the scope of these variables is extended into the body (including the condition and the post-iteration part).

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 written so that they behave like 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)
    }
}

```

Functions

Assembly allows the definition of low-level functions. These take their arguments (and a return PC) from the stack and also put the results onto the stack. Calling a function looks the same way as executing a functional-style opcode.

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. There is no explicit `return` statement.

If you call a function that returns multiple values, you have to assign them to a tuple using `a, b := f(x)` or `let a, b := f(x)`.

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) }
    }
}
```

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, managing jump labels, counting stack height for variable access and removing stack slots for assembly-local variables when the end of their block is reached. Especially for those two last cases, it is important to know that the assembler only counts stack height from top to bottom, not necessarily following control flow. Furthermore, operations like swap will only swap the contents of the stack but not the location of variables.

Conventions in Solidity

In contrast to EVM assembly, Solidity knows types which are narrower than 256 bits, e.g. `uint24`. In order to make them more efficient, most arithmetic operations just treat them as 256-bit numbers and the higher-order bits are only cleaned at the point where it is necessary, i.e. just 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 a very simple way: There is a “free memory pointer” at position `0x40` in memory. If you want to allocate memory, just use the memory from that point on and update the pointer accordingly.

Elements in memory arrays in Solidity always occupy multiples of 32 bytes (yes, 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 then only the array elements follow.

Warning: Statically-sized memory arrays do not have a length field, but it will be added soon to allow better convertibility between statically- and dynamically-sized arrays, so please do not rely on that.

Standalone Assembly

The assembly language described as inline assembly above can also be used standalone and in fact, the plan is to use it as an intermediate language for the Solidity compiler. In this form, it tries to achieve several goals:

1. Programs written in it should be readable, even if the code is generated by a compiler from Solidity.
2. The translation from assembly to bytecode should contain as few “surprises” as possible.
3. Control flow should be easy to detect to help in formal verification and optimization.

In order to achieve the first and last goal, assembly provides high-level constructs like `for` loops, `if` and `switch` statements and function calls. It should be possible to write assembly programs that do not make use of explicit `SWAP`, `DUP`, `JUMP` and `JUMPI` statements, 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.

The second goal is achieved by introducing a desugaring phase that only removes the higher level constructs in a very regular way and still allows inspecting the generated low-level assembly code. The only non-local operation performed by the assembler is name lookup of user-defined identifiers (functions, variables, ...), which follow very simple and regular scoping rules and cleanup of local variables from the stack.

Scoping: An identifier that is declared (label, variable, function, assembly) is only visible in the block where it was declared (including nested blocks inside the current block). It is not legal to access local variables across function borders, even if they would be in scope. Shadowing is not allowed. Local variables cannot be accessed before they were declared, but labels, functions and assemblies can. Assemblies are special blocks that are used for e.g. returning runtime code or creating contracts. No identifier from an outer assembly is visible in a sub-assembly.

If control flow passes over the end of a block, pop instructions are inserted that match the number of local variables declared in that block. Whenever a local variable is referenced, the code generator needs to know its current relative position in the stack and thus it needs to keep track of the current so-called stack height. Since all local variables are removed at the end of a block, the stack height before and after the block should be the same. If this is not the case, a warning is issued.

Why do we use higher-level constructs like `switch`, `for` and functions:

Using `switch`, `for` and functions, it should be possible to write complex code without using `jump` or `jumpi` manually. This makes it much easier to analyze the control flow, which allows for improved formal verification and optimization.

Furthermore, if manual jumps are allowed, computing the stack height is rather complicated. The position of all local variables on the stack needs to be known, otherwise neither references to local variables nor removing local variables automatically from the stack at the end of a block will work properly. The desugaring mechanism correctly inserts operations at unreachable blocks that adjust the stack height properly in case of jumps that do not have a continuing control flow.

Example:

We will follow an example compilation from Solidity to desugared assembly. We consider the runtime bytecode of the following Solidity program:

```
pragma solidity ^0.4.0;

contract C {
    function f(uint x) returns (uint y) {
        y = 1;
        for (uint i = 0; i < x; i++)
            y = 2 * y;
    }
}
```

The following assembly will be generated:

```
{
    mstore(0x40, 0x60) // store the "free memory pointer"
    // function dispatcher
    switch div(calldataload(0), exp(2, 226))
    case 0xb3de648b {
        let (r) = f(calldataload(4))
        let ret := $allocate(0x20)
        mstore(ret, r)
        return(ret, 0x20)
    }
    default { revert(0, 0) }
    // memory allocator
    function $allocate(size) -> pos {
        pos := mload(0x40)
```

```

    mstore(0x40, add(pos, size))
  }
  // the contract function
  function f(x) -> y {
    y := 1
    for { let i := 0 } lt(i, x) { i := add(i, 1) } {
      y := mul(2, y)
    }
  }
}

```

After the desugaring phase it looks as follows:

```

{
  mstore(0x40, 0x60)
  {
    let $0 := div(calldataload(0), exp(2, 226))
    jumpi($case1, eq($0, 0xb3de648b))
    jump($caseDefault)
    $case1:
    {
      // the function call - we put return label and arguments on the stack
      $ret1 calldataload(4) jump(f)
      // This is unreachable code. Opcodes are added that mirror the
      // effect of the function on the stack height: Arguments are
      // removed and return values are introduced.
      pop pop
      let r := 0
      $ret1: // the actual return point
      $ret2 0x20 jump($allocate)
      pop pop let ret := 0
      $ret2:
      mstore(ret, r)
      return(ret, 0x20)
      // although it is useless, the jump is automatically inserted,
      // since the desugaring process is a purely syntactic operation that
      // does not analyze control-flow
      jump($endswitch)
    }
    $caseDefault:
    {
      revert(0, 0)
      jump($endswitch)
    }
    $endswitch:
  }
  jump($afterFunction)
  allocate:
  {
    // we jump over the unreachable code that introduces the function arguments
    jump($start)
    let $retpos := 0 let size := 0
    $start:
    // output variables live in the same scope as the arguments and is
    // actually allocated.
    let pos := 0
    {
      pos := mload(0x40)
    }
  }
}

```

```

    mstore(0x40, add(pos, size))
  }
  // This code replaces the arguments by the return values and jumps back.
  swap1 pop swap1 jump
  // Again unreachable code that corrects stack height.
  0 0
}
f:
{
  jump($start)
  let $retpos := 0 let x := 0
  $start:
  let y := 0
  {
    let i := 0
    $for_begin:
    jumpi($for_end, iszero(lt(i, x)))
    {
      y := mul(2, y)
    }
    $for_continue:
    { i := add(i, 1) }
    jump($for_begin)
    $for_end:
  } // Here, a pop instruction will be inserted for i
  swap1 pop swap1 jump
  0 0
}
$afterFunction:
stop
}

```

Assembly happens in four stages:

1. Parsing
2. Desugaring (removes switch, for and functions)
3. Opcode stream generation
4. Bytecode generation

We will specify steps one to three in a pseudo-formal way. More formal specifications will follow.

Parsing / Grammar

The tasks of the parser are the following:

- Turn the byte stream into a token stream, discarding C++-style comments (a special comment exists for source references, but we will not explain it here).
- Turn the token stream into an AST according to the grammar below
- Register identifiers with the block they are defined in (annotation to the AST node) and note from which point on, variables can be accessed.

The assembly lexer follows the one defined by Solidity itself.

Whitespace is used to delimit tokens and it consists of the characters Space, Tab and Linefeed. Comments are regular JavaScript/C++ comments and are interpreted in the same way as Whitespace.

Grammar:

```

AssemblyBlock = '{' AssemblyItem* '}'
AssemblyItem =
  Identifier |
  AssemblyBlock |
  FunctionalAssemblyExpression |
  AssemblyLocalDefinition |
  FunctionalAssemblyAssignment |
  AssemblyAssignment |
  LabelDefinition |
  AssemblyIf |
  AssemblySwitch |
  AssemblyFunctionDefinition |
  AssemblyFor |
  'break' | 'continue' |
  SubAssembly | 'dataSize' '(' Identifier ')' |
  LinkerSymbol |
  'errorLabel' | 'bytecodeSize' |
  NumberLiteral | StringLiteral | HexLiteral
Identifier = [a-zA-Z_$] [a-zA-Z_0-9]*
FunctionalAssemblyExpression = Identifier '(' ( AssemblyItem ( ',' AssemblyItem ) * )?
↳ ')'
AssemblyLocalDefinition = 'let' IdentifierOrList ':' FunctionalAssemblyExpression
FunctionalAssemblyAssignment = IdentifierOrList ':=' FunctionalAssemblyExpression
IdentifierOrList = Identifier | '(' IdentifierList ')'
IdentifierList = Identifier ( ',' Identifier ) *
AssemblyAssignment = '=' Identifier
LabelDefinition = Identifier ':'
AssemblyIf = 'if' FunctionalAssemblyExpression AssemblyBlock
AssemblySwitch = 'switch' FunctionalAssemblyExpression AssemblyCase*
  ( 'default' AssemblyBlock )?
AssemblyCase = 'case' FunctionalAssemblyExpression AssemblyBlock
AssemblyFunctionDefinition = 'function' Identifier '(' IdentifierList? ')'
  ( '->' '(' IdentifierList ')' )? AssemblyBlock
AssemblyFor = 'for' ( AssemblyBlock | FunctionalAssemblyExpression )
  FunctionalAssemblyExpression ( AssemblyBlock | FunctionalAssemblyExpression )
↳ AssemblyBlock
SubAssembly = 'assembly' Identifier AssemblyBlock
LinkerSymbol = 'linkerSymbol' '(' StringLiteral ')'
NumberLiteral = HexNumber | DecimalNumber
HexLiteral = 'hex' ( '"' ([0-9a-fA-F]{2}) * '"' | '\' ' ' ([0-9a-fA-F]{2}) * '\' ' ' )
StringLiteral = '"' ([^"r\n\\] | '\\ ' . ) * '"'
HexNumber = '0x' [0-9a-fA-F]+
DecimalNumber = [0-9]+

```

Desugaring

An AST transformation removes for, switch and function constructs. The result is still parseable by the same parser, but it will not use certain constructs. If jumpdests are added that are only jumped to and not continued at, information about the stack content is added, unless no local variables of outer scopes are accessed or the stack height is the same as for the previous instruction.

Pseudocode:

```

desugar item: AST -> AST =
match item {

```

```

AssemblyFunctionDefinition('function' name '(' arg1, ..., argn ')' '->' ( '(' ret1, ..
↳., retm ')' body) ->
  <name>:
  {
    jump($<name>_start)
    let $retPC := 0 let argn := 0 ... let arg1 := 0
    $<name>_start:
    let ret1 := 0 ... let retm := 0
    { desugar(body) }
    swap and pop items so that only ret1, ... retm, $retPC are left on the stack
    jump
    0 (1 + n times) to compensate removal of arg1, ..., argn and $retPC
  }
AssemblyFor('for' { init } condition post body) ->
  {
    init // cannot be its own block because we want variable scope to extend into the_
↳body
    // find I such that there are no labels $forI_*
    $forI_begin:
    jumpi($forI_end, iszero(condition))
    { body }
    $forI_continue:
    { post }
    jump($forI_begin)
    $forI_end:
  }
'break' ->
  {
    // find nearest enclosing scope with label $forI_end
    pop all local variables that are defined at the current point
    but not at $forI_end
    jump($forI_end)
    0 (as many as variables were removed above)
  }
'continue' ->
  {
    // find nearest enclosing scope with label $forI_continue
    pop all local variables that are defined at the current point
    but not at $forI_continue
    jump($forI_continue)
    0 (as many as variables were removed above)
  }
AssemblySwitch(switch condition cases ( default: defaultBlock )? ) ->
  {
    // find I such that there is no $switchI* label or variable
    let $switchI_value := condition
    for each of cases match {
      case val: -> jumpi($switchI_caseJ, eq($switchI_value, val))
    }
    if default block present: ->
      { defaultBlock jump($switchI_end) }
    for each of cases match {
      case val: { body } -> $switchI_caseJ: { body jump($switchI_end) }
    }
    $switchI_end:
  }
FunctionalAssemblyExpression( identifier(arg1, arg2, ..., argn) ) ->
  {

```

```

if identifier is function <name> with n args and m ret values ->
{
  // find I such that $funcallI_* does not exist
  $funcallI_return argn ... arg2 arg1 jump(<name>)
  pop (n + 1 times)
  if the current context is let (idl, ..., idm) := f(...) ->
    let idl := 0 ... let idm := 0
    $funcallI_return:
  else ->
    0 (m times)
    $funcallI_return:
    turn the functional expression that leads to the function call
    into a statement stream
}
else -> desugar(children of node)
}
default node ->
  desugar(children of node)
}

```

Opcode Stream Generation

During opcode stream generation, we keep track of the current stack height in a counter, so that accessing stack variables by name is possible. The stack height is modified with every opcode that modifies the stack and with every label that is annotated with a stack adjustment. Every time a new local variable is introduced, it is registered together with the current stack height. If a variable is accessed (either for copying its value or for assignment), the appropriate DUP or SWAP instruction is selected depending on the difference between the current stack height and the stack height at the point the variable was introduced.

Pseudocode:

```

codegen item: AST -> opcode_stream =
match item {
AssemblyBlock({ items }) ->
  join(codegen(item) for item in items)
  if last generated opcode has continuing control flow:
    POP for all local variables registered at the block (including variables
    introduced by labels)
    warn if the stack height at this point is not the same as at the start of the_
↪block
Identifier(id) ->
  lookup id in the syntactic stack of blocks
  match type of id
  Local Variable ->
    DUPi where i = 1 + stack_height - stack_height_of_identifier(id)
  Label ->
    // reference to be resolved during bytecode generation
    PUSH<bytecode position of label>
  SubAssembly ->
    PUSH<bytecode position of subassembly data>
FunctionalAssemblyExpression(id ( arguments ) ) ->
  join(codegen(arg) for arg in arguments.reversed())
  id (which has to be an opcode, might be a function name later)
AssemblyLocalDefinition(let (idl, ..., idn) := expr) ->
  register identifiers idl, ..., idn as locals in current block at current stack_
↪height

```



```

codegen(expr) - assert that expr returns n items to the stack
FunctionalAssignment((idl, ..., idn) := expr) ->
  lookup id1, ..., idn in the syntactic stack of blocks, assert that they are
  ↪variables
  codegen(expr)
  for j = n, ..., 1:
    SWAPi where i = 1 + stack_height - stack_height_of_identifier(idj)
  POP
AssemblyAssignment(=: id) ->
  look up id in the syntactic stack of blocks, assert that it is a variable
  SWAPi where i = 1 + stack_height - stack_height_of_identifier(id)
  POP
LabelDefinition(name:) ->
  JUMPDEST
NumberLiteral(num) ->
  PUSH<num interpreted as decimal and right-aligned>
HexLiteral(lit) ->
  PUSH32<lit interpreted as hex and left-aligned>
StringLiteral(lit) ->
  PUSH32<lit utf-8 encoded and left-aligned>
SubAssembly(assembly <name> block) ->
  append codegen(block) at the end of the code
dataSize(<name>) ->
  assert that <name> is a subassembly ->
  PUSH32<size of code generated from subassembly <name>>
linkerSymbol(<lit>) ->
  PUSH32<zeros> and append position to linker table
}

```

7.4.8 Miscellaneous

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 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 that many bytes that 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).

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.

The elements of structs and arrays are stored after each other, just as if they were given explicitly.

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 an (unfilled) slot in storage at some position `p` according to the above rule (or by recursively applying this rule for mappings to mappings or arrays of arrays). For a dynamic array, this slot stores the number of elements in the array (byte arrays and strings are an exception here, see below). For a mapping, 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)`.

`bytes` and `string` store their data in the same slot where also the length is stored if they are short. 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`. If it is longer, the main slot stores `length * 2 + 1` and the data is stored as usual in `keccak256(slot)`.

So for the following contract snippet:

```
pragma solidity ^0.4.0;

contract C {
    struct s { uint a; uint b; }
    uint x;
    mapping(uint => mapping(uint => s)) data;
}
```

The position of `data[4][9].b` is at `keccak256(uint256(9) . keccak256(uint256(4) . uint256(1))) + 1`.

Layout in Memory

Solidity reserves three 256-bit slots:

- 0 - 64: scratch space for hashing methods
- 64 - 96: currently allocated memory size (aka. free memory pointer)

Scratch space can be used between statements (ie. within inline assembly).

Solidity always places new objects at the free memory pointer and memory is never freed (this might change in the future).

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 lifecycle, the pointer is not updated. The memory may or may not be zeroed out. Because of this, one shouldn't expect the free memory to be zeroed out.

Layout of Call Data

When a Solidity contract is deployed and when it is called from an account, the input data is assumed to be in the format in *the ABI specification*. The ABI specification requires arguments to be padded to multiples of 32 bytes. The internal function calls use a different convention.

Internals - 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 the 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

Internals - The Optimizer

The Solidity optimizer operates on assembly, so it can be and also is used by other languages. It splits the sequence of instructions into basic blocks at `JUMPS` and `JUMPDESTs`. Inside these blocks, the instructions are analysed and every modification to the stack, to memory or storage is recorded as an expression which consists of an instruction and a list of arguments which are essentially pointers to other expressions. The main idea is now to find expressions that are always equal (on every input) and combine them into an expression class. The optimizer first tries to find each new expression in a list of already known expressions. If this does not work, the expression is simplified according to rules like `constant + constant = sum_of_constants` or `X * 1 = X`. Since this is done recursively, we can also apply the latter rule if the second factor is a more complex expression where we know that it will always evaluate 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 actually do not know what is stored at `x` after we wrote to `y`. On the other hand, if a 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`.

At the end of this process, we know which expressions have to be on the stack in 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 a `JUMPI` is found whose condition evaluates to a constant, it is transformed to an unconditional jump.

As the last step, the code in each block is completely re-generated. A dependency graph is created from the expressions on the stack at the end of the block and every operation that is not part of this graph is essentially dropped. Now code is generated 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) and finally, 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, and thus code like

```
var x = 7;
data[7] = 9;
if (data[x] != x + 2)
    return 2;
else
    return 1;
```

is simplified to code which can also be compiled from

```
data[7] = 9;
return 1;
```

even though the instructions contained a jump in the beginning.

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.

Both kinds of source mappings use integer identifiers to refer to source files. These are regular array indices into a list of source files usually called "sourceList", which is part of the combined-json and the output of the json / npm compiler.

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` 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 and `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.

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::2;;
```

Tips and Tricks

- Use `delete` on arrays to delete all its elements.
- Use shorter types for struct elements and sort them such that short types are grouped together. This can lower the gas costs as multiple `SSTORE` operations might be combined into a single (`SSTORE` costs 5000 or 20000 gas, so this is what you want to optimise). Use the gas price estimator (with optimiser enabled) to check!
- Make your state variables public - the compiler will create *getters* for you automatically.
- If you end up checking conditions on input or state a lot at the beginning of your functions, try using *Function Modifiers*.
- If your contract has a function called `send` but you want to use the built-in `send`-function, use `address(contractVariable).send(amount)`.
- Initialise storage structs with a single assignment: `x = MyStruct({a: 1, b: 2});`

Cheatsheet

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	<code>++, --</code>
	New expression	<code>new <typename></code>
	Array subscripting	<code><array>[<index>]</code>
	Member access	<code><object>.<member></code>
	Function-like call	<code><func>(<args...>)</code>
	Parentheses	<code>(<statement>)</code>
2	Prefix increment and decrement	<code>++, --</code>
	Unary plus and minus	<code>+, -</code>
	Unary operations	<code>delete</code>
	Logical NOT	<code>!</code>
	Bitwise NOT	<code>~</code>
3	Exponentiation	<code>**</code>
4	Multiplication, division and modulo	<code>*, /, %</code>
5	Addition and subtraction	<code>+, -</code>
6	Bitwise shift operators	<code><<, >></code>
7	Bitwise AND	<code>&</code>
8	Bitwise XOR	<code>^</code>
9	Bitwise OR	<code> </code>
10	Inequality operators	<code><, >, <=, >=</code>
11	Equality operators	<code>==, !=</code>
12	Logical AND	<code>&&</code>
13	Logical OR	<code> </code>
14	Ternary operator	<code><conditional> ? <if-true> : <if-false></code>
15	Assignment operators	<code>=, =, ^=, &=, <<=, >>=, +=, -=, *=, /=, %=</code>
16	Comma operator	<code>,</code>

Global Variables

- `block.blockhash(uint blockNumber)` returns `(bytes32)`: hash of the given block - only

works for 256 most recent blocks

- `block.coinbase` (address): 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
- `msg.data` (bytes): complete calldata
- `msg.gas` (uint): remaining gas
- `msg.sender` (address): sender of the message (current call)
- `msg.value` (uint): number of wei sent with the message
- `now` (uint): current block timestamp (alias for `block.timestamp`)
- `tx.gasprice` (uint): gas price of the transaction
- `tx.origin` (address): 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)
- `revert` (): abort execution and revert state changes
- `keccak256(...)` returns (bytes32): compute the Ethereum-SHA-3 (Keccak-256) hash of the *(tightly packed) arguments*
- `sha3(...)` returns (bytes32): an alias to `keccak256`
- `sha256(...)` returns (bytes32): compute the SHA-256 hash of the *(tightly packed) arguments*
- `ripemd160(...)` returns (bytes20): compute the RIPEMD-160 hash of the *(tightly packed) arguments*
- `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}
- `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}
- `this` (current contract's type): the current contract, explicitly convertible to address
- `super`: the contract one level higher in the inheritance hierarchy
- `selfdestruct` (address recipient): destroy the current contract, sending its funds to the given address
- `suicide` (address recipient): an alias to `selfdestruct`
- `<address>.balance` (uint256): balance of the *Address* in Wei
- `<address>.send`(uint256 amount) returns (bool): send given amount of Wei to *Address*, returns false on failure
- `<address>.transfer` (uint256 amount): send given amount of Wei to *Address*, throws on failure

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

Modifiers

- `pure` for functions: Disallows modification or access of state - this is not enforced yet.
- `view` for functions: Disallows modification of state - this is not enforced yet.
- `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.
- `constant` for functions: Same as `view`.
- `anonymous` for events: Does not store event signature as topic.
- `indexed` for event parameters: Stores the parameter as topic.

Reserved Keywords

These keywords are reserved in Solidity. They might become part of the syntax in the future:

`abstract`, `after`, `case`, `catch`, `default`, `final`, `in`, `inline`, `let`, `match`, `null`, `of`, `relocatable`, `static`, `switch`, `try`, `type`, `typeof`.

Language Grammar

```
SourceUnit = (PragmaDirective | ImportDirective | ContractDefinition)*

// Pragma actually parses anything up to the trailing ';' to be fully forward-
↳compatible.
PragmaDirective = 'pragma' Identifier ([^;]+) ';'

ImportDirective = 'import' StringLiteral ('as' Identifier)? ';'
                | 'import' ('*' | Identifier) ('as' Identifier)? 'from' StringLiteral ';'
                | 'import' '{' Identifier ('as' Identifier)? (',' Identifier ('as'
↳Identifier)? )* '}' 'from' StringLiteral ';'

ContractDefinition = ( 'contract' | 'library' | 'interface' ) Identifier
                    ( 'is' InheritanceSpecifier (',' InheritanceSpecifier)* )?
                    '{' ContractPart* '}'

ContractPart = StateVariableDeclaration | UsingForDeclaration
              | StructDefinition | ModifierDefinition | FunctionDefinition |
↳EventDefinition | EnumDefinition
```

```

InheritanceSpecifier = UserDefinedTypeName ( '(' Expression ( ',' Expression ) * ')' ) ?

StateVariableDeclaration = TypeName ( 'public' | 'internal' | 'private' | 'constant' |
↳ Identifier ( '=' Expression ) ? ';'
UsingForDeclaration = 'using' Identifier 'for' ( '*' | TypeName ) ';'
StructDefinition = 'struct' Identifier '{'
    ( VariableDeclaration ';' ( VariableDeclaration ';' ) * ) ? '}'

ModifierDefinition = 'modifier' Identifier ParameterList? Block
ModifierInvocation = Identifier ( '(' ExpressionList? ')' ) ?

FunctionDefinition = 'function' Identifier? ParameterList
    ( ModifierInvocation | StateMutability | 'external' | 'public' |
↳ 'internal' | 'private' ) *
    ( 'returns' ParameterList ) ? ( ';' | Block )
EventDefinition = 'event' Identifier IndexedParameterList 'anonymous'? ';'

EnumValue = Identifier
EnumDefinition = 'enum' Identifier '{' EnumValue? ( ',' EnumValue ) * '}'

IndexedParameterList = '(' ( TypeName 'indexed'? Identifier? ( ',' TypeName 'indexed'?
↳ Identifier? ) * ) ? ')'
ParameterList = '(' ( TypeName Identifier? ( ',' TypeName
↳ Identifier? ) * ) ? ')'
TypeNameList = '(' ( TypeName ( ',' TypeName ) * ) ? ')'

// semantic restriction: mappings and structs (recursively) containing mappings
// are not allowed in argument lists
VariableDeclaration = TypeName StorageLocation? Identifier

TypeName = ElementaryTypeName
    | UserDefinedTypeName
    | Mapping
    | ArrayTypeName
    | FunctionTypeName

UserDefinedTypeName = Identifier ( '.' Identifier ) *

Mapping = 'mapping' '(' ElementaryTypeName '=>' TypeName ')'
ArrayTypeName = TypeName '[' Expression? ']'
FunctionTypeName = 'function' TypeNameList ( 'internal' | 'external' |
↳ StateMutability ) *
    ( 'returns' TypeNameList ) ?
StorageLocation = 'memory' | 'storage'
StateMutability = 'pure' | 'constant' | 'view' | 'payable'

Block = '{' Statement* '}'
Statement = IfStatement | WhileStatement | ForStatement | Block |
↳ InlineAssemblyStatement |
    ( DoWhileStatement | PlaceholderStatement | Continue | Break | Return |
    Throw | SimpleStatement ) ';'

ExpressionStatement = Expression
IfStatement = 'if' '(' Expression ')' Statement ( 'else' Statement ) ?
WhileStatement = 'while' '(' Expression ')' Statement
PlaceholderStatement = '_'
SimpleStatement = VariableDefinition | ExpressionStatement

```



```

ForStatement = 'for' '(' (SimpleStatement)? ';' (Expression)? ';' ↵
↳ (ExpressionStatement)? ')' Statement
InlineAssemblyStatement = 'assembly' StringLiteral? InlineAssemblyBlock
DoWhileStatement = 'do' Statement 'while' '(' Expression ')'
Continue = 'continue'
Break = 'break'
Return = 'return' Expression?
Throw = 'throw'
VariableDefinition = ('var' IdentifierList | VariableDeclaration) ( '=' Expression )?
IdentifierList = '(' ( Identifier? ',' )* Identifier? ')'

// Precedence by order (see github.com/ethereum/solidity/pull/732)
Expression
= Expression ('++' | '--')
| NewExpression
| IndexAccess
| MemberAccess
| FunctionCall
| '(' Expression ')'
| ('!' | '~' | 'delete' | '++' | '--' | '+' | '-') Expression
| Expression '**' Expression
| Expression ('*' | '/' | '%') Expression
| Expression ('+' | '-') Expression
| Expression ('<<' | '>>') Expression
| Expression '&' Expression
| Expression '^' Expression
| Expression '|' Expression
| Expression ('<' | '>' | '<=' | '>=') Expression
| Expression ('==' | '!=') Expression
| Expression '&&' Expression
| Expression '||' Expression
| Expression '?' Expression ':' Expression
| Expression ('=' | '|=' | '^=' | '&=' | '<<=' | '>>=' | '+=' | '-=' | '*=' | '/=' ↵
↳ '|%=') Expression
| PrimaryExpression

PrimaryExpression = BooleanLiteral
                    | NumberLiteral
                    | HexLiteral
                    | StringLiteral
                    | TupleExpression
                    | Identifier
                    | ElementaryTypeNameExpression

ExpressionList = Expression ( ',' Expression )*
NameValueList = Identifier ':' Expression ( ',' Identifier ':' Expression )*

FunctionCall = Expression '(' FunctionCallArguments ')'
FunctionCallArguments = '{' NameValueList? '}'
                    | ExpressionList?

NewExpression = 'new' TypeName
MemberAccess = Expression '.' Identifier
IndexAccess = Expression '[' Expression? ']'

BooleanLiteral = 'true' | 'false'
NumberLiteral = ( HexNumber | DecimalNumber ) ( ' ' NumberUnit)?
NumberUnit = 'wei' | 'szabo' | 'finney' | 'ether'

```

```

    | 'seconds' | 'minutes' | 'hours' | 'days' | 'weeks' | 'years'
HexLiteral = 'hex' ('' ([0-9a-fA-F]{2})* '' | '\' ([0-9a-fA-F]{2})* '\')
StringLiteral = '' ([^"r\n\\] | '\\'.)* ''
Identifier = [a-zA-Z_] [a-zA-Z_0-9]*

HexNumber = '0x' [0-9a-fA-F]+
DecimalNumber = [0-9]+ ( '.' [0-9]* )? ( [eE] [0-9]+ )?

TupleExpression = '(' ( Expression? ( ',' Expression? )* )? ')'
                | '[' ( Expression ( ',' Expression )* )? ']'

ElementaryTypeNameExpression = ElementaryTypeName

ElementaryTypeName = 'address' | 'bool' | 'string' | 'var'
                   | Int | Uint | Byte | Fixed | Ufixed

Int = 'int' | 'int8' | 'int16' | 'int24' | 'int32' | 'int40' | 'int48' | 'int56' |
↳ 'int64' | 'int72' | 'int80' | 'int88' | 'int96' | 'int104' | 'int112' | 'int120' |
↳ 'int128' | 'int136' | 'int144' | 'int152' | 'int160' | 'int168' | 'int176' | 'int184
↳ 'int192' | 'int200' | 'int208' | 'int216' | 'int224' | 'int232' | 'int240' |
↳ 'int248' | 'int256'

Uint = 'uint' | 'uint8' | 'uint16' | 'uint24' | 'uint32' | 'uint40' | 'uint48' |
↳ 'uint56' | 'uint64' | 'uint72' | 'uint80' | 'uint88' | 'uint96' | 'uint104' |
↳ 'uint112' | 'uint120' | 'uint128' | 'uint136' | 'uint144' | 'uint152' | 'uint160' |
↳ 'uint168' | 'uint176' | 'uint184' | 'uint192' | 'uint200' | 'uint208' | 'uint216' |
↳ 'uint224' | 'uint232' | 'uint240' | 'uint248' | 'uint256'

Byte = 'byte' | 'bytes' | 'bytes1' | 'bytes2' | 'bytes3' | 'bytes4' | 'bytes5' |
↳ 'bytes6' | 'bytes7' | 'bytes8' | 'bytes9' | 'bytes10' | 'bytes11' | 'bytes12' |
↳ 'bytes13' | 'bytes14' | 'bytes15' | 'bytes16' | 'bytes17' | 'bytes18' | 'bytes19' |
↳ 'bytes20' | 'bytes21' | 'bytes22' | 'bytes23' | 'bytes24' | 'bytes25' | 'bytes26' |
↳ 'bytes27' | 'bytes28' | 'bytes29' | 'bytes30' | 'bytes31' | 'bytes32'

Fixed = 'fixed' | ( 'fixed' [0-9]+ 'x' [0-9]+ )

Ufixed = 'ufixed' | ( 'ufixed' [0-9]+ 'x' [0-9]+ )

InlineAssemblyBlock = '{' AssemblyItem* '}'

AssemblyItem = Identifier | FunctionalAssemblyExpression | InlineAssemblyBlock |
↳ AssemblyLocalBinding | AssemblyAssignment | AssemblyLabel | NumberLiteral |
↳ StringLiteral | HexLiteral
AssemblyLocalBinding = 'let' Identifier ':' FunctionalAssemblyExpression
AssemblyAssignment = ( Identifier ':' FunctionalAssemblyExpression ) | ( '='
↳ Identifier )
AssemblyLabel = Identifier ':'
FunctionalAssemblyExpression = Identifier '(' AssemblyItem? ( ',' AssemblyItem )* ')'

```

7.5 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)!

7.5.1 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):

```
pragma solidity ^0.4.0;

// THIS CONTRACT CONTAINS A BUG - DO NOT USE
contract Fund {
    /// Mapping of ether shares of the contract.
    mapping(address => uint) shares;
    /// Withdraw your share.
    function withdraw() {
        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 always includes 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.

To avoid re-entrancy, you can use the Checks-Effects-Interactions pattern as outlined further below:

```
pragma solidity ^0.4.11;

contract Fund {
    /// Mapping of ether shares of the contract.
    mapping(address => uint) shares;
    /// Withdraw your share.
    function withdraw() {
        var 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 `constant` 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), the fallback function is executed. If it does not have a fallback function, the Ether will be rejected (by throwing an exception). During the execution of the fallback function, the contract can only rely on the “gas stipend” (2300 gas) being available to it at that time. This stipend is not enough to access storage in any way. To be sure that your contract can receive Ether in that way, check the gas requirements of the fallback function (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 only returns a failure code and does not automatically propagate 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.
- 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 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`, `throw` or because the operation is just 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()`, `.callcode()` and `.delegatecall()` behave in the same way.

tx.origin

Never use `tx.origin` for authorization. Let's say you have a wallet contract like this:

```
pragma solidity ^0.4.11;

// THIS CONTRACT CONTAINS A BUG - DO NOT USE
contract TxUserWallet {
    address owner;

    function TxUserWallet() {
        owner = msg.sender;
    }

    function transferTo(address dest, uint amount) {
        require(tx.origin == owner);
        dest.transfer(amount);
    }
}
```

Now someone tricks you into sending ether to the address of this attack wallet:

```
pragma solidity ^0.4.11;

interface TxUserWallet {
    function transferTo(address dest, uint amount);
}

contract TxAttackWallet {
    address owner;

    function TxAttackWallet() {
        owner = msg.sender;
    }

    function() {
        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.

Minor Details

- In `for (var i = 0; i < arrayName.length; i++) { ... }`, the type of `i` will be `uint8`, because this is the smallest type that is required to hold the value 0. If the array has more than 255 elements, the loop will not terminate.
- The `constant` keyword for functions is currently not enforced by the compiler. Furthermore, it is not enforced by the EVM, so a contract function that “claims” to be constant might still cause changes to the state.

- 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.

7.5.2 Recommendations

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.

7.5.3 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.

7.6 Using the compiler

7.6.1 Using the Commandline Compiler

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. Before you deploy your contract, activate the optimizer while compiling using `solc --optimize --bin sourceFile.sol`. 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 --asm sourceFile.sol`.

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/=/usr/local/lib/fallback file.sol
```

This essentially instructs the compiler to search for anything starting with `github.com/ethereum/dapp-bin/` under `/usr/local/lib/dapp-bin` and if it does not find the file there, it will look at `/usr/local/lib/fallback` (the empty prefix always matches). `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.

If there are multiple matches due to remappings, the one with the longest common prefix is selected.

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.

If your contracts use *libraries*, you will notice that the bytecode contains substrings of the form `__LibraryName__`. You can use `solc` as a linker meaning that it will insert the library addresses for you at those points:

Either add `--libraries "Math:0x12345678901234567890 Heap:0xabcdef0123456"` to your command to provide an address for each library or store the string in a file (one library per line) and run `solc` using `--libraries fileName`.

If `solc` is called with the option `--link`, all input files are interpreted to be unlinked binaries (hex-encoded) in the `__LibraryName__`-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.

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.

7.6.2 Compiler Input and Output JSON Description

These JSON formats are used by the compiler API as well as are available through `solc`. These are subject to change, some fields are optional (as noted), but it is aimed at to only make backwards compatible changes.

The compiler API expects a JSON formatted input and outputs the compilation result in a JSON formatted output. Comments are of course not permitted and used here only for explanatory purposes.

Input Description

```
{
  // Required: Source code language, such as "Solidity", "serpent", "l1l", "assembly",
  ↪ etc.
  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.
      "urls":
      [
        "bzzr://56ab...",
        "ipfs://Qma...",
        "file:///tmp/path/to/file.sol"
      ]
    },
    "mortal":
    {
      // Optional: keccak256 hash of the source file
      "keccak256": "0x234...",
      // Required (unless "urls" is used): literal contents of the source file
      "content": "contract mortal is owned { function kill() { if (msg.sender == ↪
  ↪owner) selfdestruct(owner); } }"
    }
  },
  // Optional
  settings:
  {
    // Optional: Sorted list of remappings
    remappings: [ "g/dir" ],
    // Optional: Optimizer settings (enabled defaults to false)
    optimizer: {
      enabled: true,
      runs: 500
    },
    // Metadata settings (optional)
    metadata: {
      // Use only literal content and not URLs (false by default)
      useLiteralContent: true
    },
    // 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.
// 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 is the contract name,
↳where empty contract name refers to the file itself,
↳while the star refers to all of the contracts.
//
// The available output types are as follows:
// abi - ABI
// ast - AST of all source files
// legacyAST - legacy AST of all source files
// devdoc - Developer documentation (natspec)
// userdoc - User documentation (natspec)
// metadata - Metadata
// ir - New assembly format before desugaring
// evm.assembly - New assembly format after desugaring
// 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.deployedBytecode* - Deployed bytecode (has the same options as evm.
↳bytecode)
// evm.methodIdentifiers - The list of function hashes
// evm.gasEstimates - Function gas estimates
// ewasm.wast - eWASM S-expressions format (not supported atm)
// ewasm.wasm - eWASM binary format (not supported atm)
//
// 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: {
    // Enable the metadata and bytecode outputs of every single contract.
    "*": {
        "*": [ "metadata", "evm.bytecode" ]
    },
    // Enable the abi and opcodes output of MyContract defined in file def.
    "def": {
        "MyContract": [ "abi", "evm.opcodes" ]
    },
    // Enable the source map output of every single contract.
    "*": {
        "*": [ "evm.sourceMap" ]
    },
    // Enable the legacy AST output of every single file.
    "*": {
        "": [ "legacyAST" ]
    }
}

```

```

    }
  }
}

```

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
      },
      // Mandatory: Error type, such as "TypeError", "InternalCompilerError",
      ↪ "Exception", etc
      type: "TypeError",
      // Mandatory: Component where the error originated, such as "general", "ewasm", ↪
      ↪ etc.
      component: "general",
      // Mandatory ("error" or "warning")
      severity: "error",
      // Mandatory
      message: "Invalid keyword"
      // Optional: the message formatted with source location
      formattedMessage: "sourceFile.sol:100: Invalid keyword"
    }
  ],
  // This contains the file-level outputs. In can be limited/filtered by the ↪
  ↪ outputSelection settings.
  sources: {
    "sourceFile.sol": {
      // Identifier (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://github.com/ethereum/wiki/wiki/Ethereum-Contract-ABI
        abi: [],
        // See the Metadata Output documentation (serialised JSON string)
        metadata: "{...}",
        // User documentation (natspec)
        userdoc: {},

```

```

// Developer documentation (natspec)
devdoc: {},
// Intermediate representation (string)
ir: "",
// 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: "",
    // 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 }
        ]
      }
    },
  },
  // The same layout as above.
  deployedBytecode: { },
  // 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: ""
}

```

```

    }
  }
}

```

7.7 Contract Metadata

The Solidity compiler automatically generates a JSON file, the contract metadata, that contains information about the current contract. It can be used to query the compiler version, the sources used, the ABI and NatSpec documentation in order to more safely interact with the contract and to verify its source code.

The compiler appends a Swarm 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.

Of course, you have to publish the metadata file to Swarm (or some other service) so that others can access it. The file can be output by using `solc --metadata` and the file will be called `ContractName_meta.json`. It will contain 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 of course also 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.2dabbd0.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..." ]
    },
    "mortal": {
      // Required: keccak256 hash of the source file
      "keccak256": "0x234...",
      // Required (unless "url" is used): literal contents of the source file
      "content": "contract mortal is owned { function kill() { if (msg.sender ==
↳owner) selfdestruct(owner); } }"
    }
  },
  // Required: Compiler settings

```

```

settings:
{
  // Required for Solidity: Sorted list of remappings
  remappings: [ ":g/dir" ],
  // Optional: Optimizer settings (enabled defaults to false)
  optimizer: {
    enabled: true,
    runs: 500
  },
  // 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: [ ... ],
}
}

```

Note: Note the ABI definition above has no fixed order. It can change with compiler versions.

Note: Since the bytecode of the resulting contract contains the metadata hash, any change to the metadata will result in a change of the bytecode. Furthermore, since the metadata includes a hash of all the sources used, a single whitespace change in any of the source codes will result in a different metadata, and subsequently a different bytecode.

7.7.1 Encoding of the Metadata Hash in the Bytecode

Because we might support other ways to retrieve the metadata file in the future, the mapping {"bzzr0": <Swarm hash>} is stored CBOR-encoded. Since 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 thus adds the following to the end of the deployed bytecode:

```
0xa1 0x65 'b' 'z' 'z' 'r' '0' 0x58 0x20 <32 bytes swarm hash> 0x00 0x29
```

So in order to retrieve the data, the end of the deployed bytecode can be checked to match that pattern and use the Swarm hash to retrieve the file.

7.7.2 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) retrieves the code of the contract, from that the 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, Mist can use the userdoc to display a confirmation message to the user whenever they interact with the contract.

7.7.3 Usage for Source Code Verification

In order to verify the compilation, sources can be retrieved from 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.

7.8 Application Binary Interface Specification

7.8.1 Basic Design

The Application Binary Interface 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. No introspection mechanism will be provided. 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. Should these cases become important they can be adequately handled as facilities built within the Ethereum ecosystem.

7.8.2 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 (SHA-3) hash of the signature of the function. The signature is defined as the canonical expression of the basic prototype, i.e. the function name with the parenthesised list of parameter types. Parameter types are split by a single comma - no spaces are used.

7.8.3 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.

7.8.4 Types

The following elementary types exist:

- `uint<M>`: unsigned integer type of M bits, $0 < M \leq 256, M \% 8 == 0$. e.g. `uint32`, `uint8`, `uint256`.
- `int<M>`: two's complement signed integer type of M bits, $0 < M \leq 256, M \% 8 == 0$.
- `address`: equivalent to `uint160`, except for the assumed interpretation and language typing.
- `uint`, `int`: synonyms for `uint256`, `int256` respectively (this shorthand not to be used for computing the function selector).
- `bool`: equivalent to `uint8` restricted to the values 0 and 1
- `fixed<M>x<N>`: signed fixed-point decimal number of M bits, $8 \leq M \leq 256, M \% 8 == 0$, and $0 < N \leq 80$, which denotes the value v as $v / (10 ** N)$.
- `ufixed<M>x<N>`: unsigned variant of `fixed<M>x<N>`.
- `fixed`, `ufixed`: synonyms for `fixed128x19`, `ufixed128x19` respectively (this shorthand not to be used for computing the function selector).
- `bytes<M>`: binary type of M bytes, $0 < M \leq 32$.
- `function`: equivalent to `bytes24`: an address, followed by a function selector

The following (fixed-size) array type exists:

- `<type> [M]`: a fixed-length array of M elements, $M > 0$, of the given type.

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 a finite non-negative number of them inside parentheses, separated by commas:

- `(T1, T2, ..., Tn)`: tuple consisting of the types `T1`, ..., `Tn`, $n \geq 0$

It is possible to form tuples of tuples, arrays of tuples and so on.

Note: Solidity supports all the types presented above with the same names with the exception of tuples. The ABI tuple type is utilised for encoding Solidity `structs`.

7.8.5 Formal Specification of the Encoding

We will now formally specify the encoding, such that it will have the following properties, which are especially useful if some arguments are nested arrays:

Properties:

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”

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 $k > 0$
- (T_1, \dots, T_k) if any T_i is dynamic for $1 \leq i \leq k$

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

- (T_1, \dots, T_k) for $k \geq 0$ and any types `T1, ..., Tk`

`enc(X) = head(X(1)) ... head(X(k-1)) tail(X(0)) ... tail(X(k-1))`

where `X(i)` is the *i*th component of the value, and `head` and `tail` are defined for `Ti` being a static type as

`head(X(i)) = enc(X(i))` and `tail(X(i)) = ""` (the empty string)

and as

`head(X(i)) = enc(len(head(X(0)) ... head(X(k-1)) tail(X(0)) ... tail(X(i-1))))`
`tail(X(i)) = enc(X(i))`

otherwise, i.e. if `Ti` is a dynamic type.

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. Its value 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[1], ..., X[k]])`

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.

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]`: `EVENT_INDEXED_ARGS[n - 1]` (`EVENT_INDEXED_ARGS` is the series of `EVENT_ARGS` that are indexed);
- `data`: `abi_serialise(EVENT_NON_INDEXED_ARGS)` (`EVENT_NON_INDEXED_ARGS` is the series of `EVENT_ARGS` that are not indexed, `abi_serialise` is the ABI serialisation function used for returning a series of typed values from a function, as described above).

7.8.10 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", or "fallback" (the *unnamed "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`, can be omitted if function doesn't return anything;
- `payable`: `true` if function accepts ether, defaults to `false`;
- `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` and `payable` (same as payable above).
- `constant`: `true` if function is either `pure` or `view`

`type` can be omitted, defaulting to "function".

Constructor and fallback function never have `name` or `outputs`. Fallback function doesn't have `inputs` either.

Sending non-zero ether to non-payable function will throw. Don't do it.

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,

```
pragma solidity ^0.4.0;

contract Test {
    function Test() { b = 0x12345678901234567890123456789012; }
    event Event(uint indexed a, bytes32 b);
    event Event2(uint indexed a, bytes32 b);
    function foo(uint a) { 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

```
contract Test {
    struct S { uint a; uint[] b; T[] c; }
    struct T { uint x; uint y; }
    function f(S s, T t, uint a) { }
}
```

would result in the JSON:

```
[
  {
    "name": "f",
    "type": "function",
    "inputs": [
```

```
{
  "name": "s",
  "type": "tuple",
  "components": [
    {
      "name": "a",
      "type": "uint256"
    },
    {
      "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": []
}
```

7.8.11 Non-standard Packed Mode

Solidity supports a non-standard packed mode where:

- no *function selector* is encoded,

- types shorter than 32 bytes are neither zero padded nor sign extended and
- dynamic types are encoded in-place and without the length.

As an example encoding `int1`, `bytes1`, `uint16`, string with values `-1`, `0x42`, `0x2424`, `"Hello, world!"` results in

```
0xff42242448656c6c6f2c20776f726c6421
^^                               int1(-1)
^^                               bytes1(0x42)
^^^                              uint16(0x2424)
^^^^                             string("Hello, world!") without a length field
```

More specifically, each statically-sized type takes as many bytes as its range has and dynamically-sized types like `string`, `bytes` or `uint[]` are encoded without their length field. This means that the encoding is ambiguous as soon as there are two dynamically-sized elements.

7.9 Joyfully Universal Language for (Inline) Assembly

JULIA is an intermediate language that can compile to various different backends (EVM 1.0, EVM 1.5 and eWASM are planned). Because of that, it is designed to be a usable common denominator of all three platforms. It can already be used for “inline assembly” inside Solidity and future versions of the Solidity compiler will even use JULIA as intermediate language. It should also be easy to build high-level optimizer stages for JULIA.

The core components of JULIA are functions, blocks, variables, literals, for-loops, if-statements, switch-statements, expressions and assignments to variables.

JULIA is typed, both variables and literals must specify the type with postfix notation. The supported types are `bool`, `u8`, `s8`, `u32`, `s32`, `u64`, `s64`, `u128`, `s128`, `u256` and `s256`.

JULIA in itself does not even provide operators. If the EVM is targeted, opcodes will be available as built-in functions, but they can be reimplemented if the backend changes. For a list of mandatory built-in functions, see the section below.

The following example program assumes that the EVM opcodes `mul`, `div` and `mod` are available either natively or as functions and computes exponentiation.

```
{
  function power(base:u256, exponent:u256) -> result:u256
  {
    switch exponent
    case 0:u256 { result := 1:u256 }
    case 1:u256 { result := base }
    default:
    {
      result := power(mul(base, base), div(exponent, 2:u256))
      switch mod(exponent, 2:u256)
      case 1:u256 { result := mul(base, result) }
    }
  }
}
```

It is also possible to implement the same function using a for-loop instead of with recursion. Here, we need the EVM opcodes `lt` (less-than) and `add` to be available.

```
{
  function power(base:u256, exponent:u256) -> result:u256
  {
    result := 1:u256
```

```

    for { let i := 0:u256 } lt(i, exponent) { i := add(i, 1:u256) }
    {
        result := mul(result, base)
    }
}

```

7.9.1 Specification of JULIA

JULIA code is described in this chapter. JULIA code is usually placed into a JULIA object, which is described in the following chapter.

Grammar:

```

Block = '{' Statement* '}'
Statement =
    Block |
    FunctionDefinition |
    VariableDeclaration |
    Assignment |
    Expression |
    Switch |
    ForLoop |
    BreakContinue
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' Block )?
Case =
    'case' Literal Block
ForLoop =
    'for' Block Expression Block Block
BreakContinue =
    'break' | 'continue'
FunctionCall =
    Identifier '(' ( Expression ( ',' Expression ) * )? ')'
Identifier = [a-zA-Z_] [a-zA-Z_0-9]*
IdentifierList = Identifier ( ',' Identifier ) *
TypeName = Identifier | BuiltinTypeName
BuiltinTypeName = 'bool' | [us] ( '8' | '32' | '64' | '128' | '256' )
TypedIdentifierList = Identifier ':' TypeName ( ',' Identifier ':' TypeName ) *
Literal =
    (NumberLiteral | StringLiteral | HexLiteral | TrueLiteral | FalseLiteral) ':' '␣'
↔ TypeName
NumberLiteral = HexNumber | DecimalNumber
HexLiteral = 'hex' ( '"' ([0-9a-fA-F]{2}) * '"' | '\' ([0-9a-fA-F]{2}) * '\' )
StringLiteral = '"' ([^"r\n\\] | '\\' .) * '"'
TrueLiteral = 'true'

```



```
FalseLiteral = 'false'
HexNumber = '0x' [0-9a-fA-F]+
DecimalNumber = [0-9]+
```

Restrictions on the Grammar

Switches must have at least one case (including the default case). If all possible values of the expression is covered, the default case should not be allowed (i.e. a switch with a `bool` expression and having both a true and false case should 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 values 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.

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 condition part of the for-loop has to evaluate to exactly one value.

Literals cannot be larger than the their type. The largest type defined is 256-bit wide.

Scoping Rules

Scopes in JULIA 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, identifiers defined in the “init” part of the for-loop (the first block) are visible in all other parts of the for-loop (but not outside of the loop). Identifiers declared in the other parts of the for loop respect the regular syntactical scoping rules. The parameters and return parameters of functions are visible in the function body and their names cannot overlap.

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 JULIA by providing an evaluation function `E` overloaded on the various nodes of the AST. Any functions can have side effects, so `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` and `continue` 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 var1, ..., varn := rhs>: VariableDeclaration) =
  E(G, L, <var1, ..., varn := rhs>: Assignment)
E(G, L, <let var1, ..., varn>: VariableDeclaration) =
  let L1 be a copy of L where L1[$vari] = 0 for i = 1, ..., n
  G, L1, regular
E(G, L, <var1, ..., varn := rhs>: Assignment) =
  let G1, L1, v1, ..., vn = E(G, L, rhs)
  let L2 be a copy of L1 where L2[$vari] = vi for i = 1, ..., n
  G, L2, regular
E(G, L, <for { i1, ..., in } condition post body>: ForLoop) =
  if n >= 1:
    let G1, L1, mode = E(G, L, i1, ..., in)
    // mode has to be regular due to the syntactic restrictions
    let G2, L2, mode = E(G1, L1, for {} condition post body)
    // mode has to be regular due to the syntactic restrictions
    let L3 be the restriction of L2 to only variables of L
    G2, L3, regular
  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
      else:
        G3, L3, mode = E(G2, L2, post)
        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, <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: HexLiteral) = G, L, hexString(l),
  where hexString decodes l from hex and left-aligns it into 32 bytes
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

```

Type Conversion Functions

JULIA has no support for implicit type conversion and therefore functions exists to provide explicit conversion. When converting a larger type to a shorter type a runtime exception can occur in case of an overflow.

The following type conversion functions must be available: - u32tobool(x:u32) -> y:bool - booltou32(x:bool) -> y:u32 - u32tou64(x:u32) -> y:u64 - u64tou32(x:u64) -> y:u32 - etc. (TBD)

Low-level Functions

The following functions must be available:

<i>Arithmetics</i>
addu256(x:u256, y:u256) -> z:u256 x + y
subu256(x:u256, y:u256) -> z:u256 x - y
mulu256(x:u256, y:u256) -> z:u256 x * y
divu256(x:u256, y:u256) -> z:u256 x / y

Continued on next page

Table 7.2 – continued from previous page

<code>divs256(x:s256, y:s256) -> z:s256</code> <code>x / y</code> , for signed numbers in two's complement
<code>modu256(x:u256, y:u256) -> z:u256</code> <code>x % y</code>
<code>mods256(x:s256, y:s256) -> z:s256</code> <code>x % y</code> , for signed numbers in two's complement
<code>signextendu256(i:u256, x:u256) -> z:u256</code> sign extend from (i*8+7)th bit counting from least significant
<code>expu256(x:u256, y:u256) -> z:u256</code> <code>x</code> to the power of <code>y</code>
<code>addmodu256(x:u256, y:u256, m:u256) -> z:u256</code> <code>(x + y) % m</code> with arbitrary precision arithmetics
<code>mulmodu256(x:u256, y:u256, m:u256) -> z:u256</code> <code>(x * y) % m</code> with arbitrary precision arithmetics
<code>ltu256(x:u256, y:u256) -> z:bool</code> 1 if <code>x < y</code> , 0 otherwise
<code>gtu256(x:u256, y:u256) -> z:bool</code> 1 if <code>x > y</code> , 0 otherwise
<code>sltu256(x:s256, y:s256) -> z:bool</code> 1 if <code>x < y</code> , 0 otherwise, for signed numbers in two's complement
<code>sgtu256(x:s256, y:s256) -> z:bool</code> 1 if <code>x > y</code> , 0 otherwise, for signed numbers in two's complement
<code>equ256(x:u256, y:u256) -> z:bool</code> 1 if <code>x == y</code> , 0 otherwise
<code>notu256(x:u256) -> z:u256</code> <code>~x</code> , every bit of <code>x</code> is negated
<code>andu256(x:u256, y:u256) -> z:u256</code> bitwise and of <code>x</code> and <code>y</code>
<code>oru256(x:u256, y:u256) -> z:u256</code> bitwise or of <code>x</code> and <code>y</code>
<code>xoru256(x:u256, y:u256) -> z:u256</code> bitwise xor of <code>x</code> and <code>y</code>
<code>shlu256(x:u256, y:u256) -> z:u256</code> logical left shift of <code>x</code> by <code>y</code>
<code>shru256(x:u256, y:u256) -> z:u256</code> logical right shift of <code>x</code> by <code>y</code>
<code>saru256(x:u256, y:u256) -> z:u256</code> arithmetic right shift of <code>x</code> by <code>y</code>
<code>byte(n:u256, x:u256) -> v:u256</code> nth byte of <code>x</code> , where the most significant byte is the 0th byte Cannot this be just replaced by <code>and256(shr256(n, x), 0xff)</code> and let it be optimised out by the EVM backend?
<i>Memory and storage</i>
<code>mload(p:u256) -> v:u256</code> <code>mem[p..(p+32))</code>
<code>mstore(p:u256, v:u256) mem[p..(p+32)) := v</code>
<code>mstore8(p:u256, v:u256) mem[p] := v & 0xff</code> - only modifies a single byte
<code>sload(p:u256) -> v:u256</code> <code>storage[p]</code>
<code>sstore(p:u256, v:u256) storage[p] := v</code>
<code>msize()</code> -> <code>size:u256</code> size of memory, i.e. largest accessed memory index, albeit due to the memory extension function, which extends by words, this will always be a multiple of 32 bytes
<i>Execution control</i>
<code>create(v:u256, p:u256, s:u256)</code> create new contract with code <code>mem[p..(p+s))</code> and send <code>v wei</code> and return the new address
<code>call(g:u256, a:u256, v:u256, in:u256, call contract at address <code>a</code> with input <code>mem[in..(in+insize))</code> <code>insize:u256</code>, <code>out:u256</code>, providing <code>g</code> gas and <code>v</code> wei and output area <code>outside:u256</code> <code>mem[out..(out+outside))</code> returning 0 on error (eg. out of gas) -> <code>r:u256</code> and 1 on success</code>
<code>callcode(g:u256, a:u256, v:u256, in:u256, identical to <code>call</code> but only use the code from <code>a</code> <code>insize:u256</code>, <code>out:u256</code>, and stay in the context of the <code>outside:u256</code>) -> <code>r:u256</code> current contract otherwise</code>
<code>delegatecall(g:u256, a:u256, in:u256, identical to <code>callcode</code>, <code>insize:u256</code>, <code>out:u256</code>, but also keep caller <code>outside:u256</code>) -> <code>r:u256</code> and <code>callvalue</code></code>
<code>stop()</code> stop execution, identical to <code>return(0,0)</code> Perhaps it would make sense retiring this as it equals to <code>return(0,0)</code> . It can be an optimisation by the EVM backend.
<code>abort()</code> abort (equals to invalid instruction on EVM)
<code>return(p:u256, s:u256)</code> end execution, return data <code>mem[p..(p+s))</code>
<code>revert(p:u256, s:u256)</code> end execution, revert state changes, return data <code>mem[p..(p+s))</code>
<code>selfdestruct(a:u256)</code> end execution, destroy current contract and send funds to <code>a</code>
Continued on next page

Table 7.2 – continued from previous page

log0(p:u256, s:u256) log without topics and data mem[p..(p+s))
log1(p:u256, s:u256, t1:u256) log with topic t1 and data mem[p..(p+s))
log2(p:u256, s:u256, t1:u256, t2:u256) log with topics t1, t2 and data mem[p..(p+s))
log3(p:u256, s:u256, t1:u256, t2:u256, t3:u256) log with topics t1, t2, t3 and data mem[p..(p+s))
log4(p:u256, s:u256, t1:u256, t2:u256, t3:u256, t4:u256) log with topics t1, t2, t3, t4 and data mem[p..(p+s))
<i>State queries</i>
blockcoinbase() -> address:u256 current mining beneficiary
blockdifficulty() -> difficulty:u256 difficulty of the current block
blockgaslimit() -> limit:u256 block gas limit of the current block
blockhash(b:u256) -> hash:u256 hash of block nr b - only for last 256 blocks excluding current
blocknumber() -> block:u256 current block number
blocktimestamp() -> timestamp:u256 timestamp of the current block in seconds since the epoch
txorigin() -> address:u256 transaction sender
txgasprice() -> price:u256 gas price of the transaction
gasleft() -> gas:u256 gas still available to execution
balance(a:u256) -> v:u256 wei balance at address a
this() -> address:u256 address of the current contract / execution context
caller() -> address:u256 call sender (excluding delegatecall)
callvalue() -> v:u256 wei sent together with the current call
calldataload(p:u256) -> v:u256 call data starting from position p (32 bytes)
calldatasize() -> v:u256 size of call data in bytes
calldatacopy(t:u256, f:u256, s:u256) copy s bytes from calldata at position f to mem at position t
codesize() -> size:u256 size of the code of the current contract / execution context
codecopy(t:u256, f:u256, s:u256) copy s bytes from code at position f to mem at position t
extcodesize(a:u256) -> size:u256 size of the code at address a
extcodecopy(a:u256, t:u256, f:u256, s:u256) like codecopy(t, f, s) but take code at address a
<i>Others</i>
discardu256(used:u256) discard value
splitu256tou64(x:u256) -> (x1:u64, x2:u64, split u256 to four u64's x3:u64, x4:u64)
combineu64tou256(x1:u64, x2:u64, x3:u64, combine four u64's into a single u256 x4:u64) -> (x:u256)
sha3(p:u256, s:u256) -> v:u256 keccak(mem[p...(p+s)))

Backends

Backends or targets are the translators from JULIA to a specific bytecode. Each of the backends can expose functions prefixed with the name of the backend. We reserve `evm_` and `ewasm_` prefixes for the two proposed backends.

Backend: EVM

The EVM target will have all the underlying EVM opcodes exposed with the `evm_` prefix.

Backend: “EVM 1.5”

TBD

Backend: eWASM

TBD

7.9.2 Specification of JULIA Object

Grammar:

```

TopLevelObject = 'object' '{' Code? ( Object | Data )* '}'
Object = 'object' StringLiteral '{' Code? ( Object | Data )* '}'
Code = 'code' Block
Data = 'data' StringLiteral HexLiteral
HexLiteral = 'hex' ('' ([0-9a-fA-F]{2})* '' | '\\' ([0-9a-fA-F]{2})* '\\')
StringLiteral = '"' ([^"r\n\\] | '\\' .)* '"'

```

Above, Block refers to Block in the JULIA code grammar explained in the previous chapter.

An example JULIA Object is shown below:

..code:

```

// Code consists of a single object. A single "code" node is the 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
object {
  code {
    let size = datasize("runtime")
    let offset = allocate(size)
    // This will turn into a memory->memory copy for eWASM and
    // a codecopy for EVM
    datacopy(dataoffset("runtime"), offset, size)
    // this is a constructor and the runtime code is returned
    return(offset, size)
  }

  data "Table2" hex"4123"

  object "runtime" {
    code {
      // runtime code

      let size = datasize("Contract2")
      let offset = allocate(size)
      // This will turn into a memory->memory copy for eWASM and
      // a codecopy for EVM
      datacopy(dataoffset("Contract2"), offset, size)
      // constructor parameter is a single number 0x1234
      mstore(add(offset, size), 0x1234)
      create(offset, add(size, 32))
    }

    // 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"
}
}
```

7.10 Style Guide

7.10.1 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.

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!

7.10.2 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:

```
contract A {  
    ...  
}  
  
contract B {  
    ...  
}  
  
contract C {  
    ...  
}
```

No:

```
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:

```
contract A {  
    function spam();  
    function ham();  
}  
  
contract B is A {  
    function spam() {  
        ...  
    }  
  
    function ham() {  
        ...  
    }  
}
```

No:

```
contract A {  
    function spam() {  
        ...  
    }  
    function ham() {  
        ...  
    }  
}
```

Source File Encoding

UTF-8 or ASCII encoding is preferred.

Imports

Import statements should always be placed at the top of the file.

Yes:

```
import "owned";

contract A {
    ...
}

contract B is owned {
    ...
}
```

No:

```
contract A {
    ...
}

import "owned";

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
- fallback function (if exists)
- external
- public
- internal
- private

Within a grouping, place the constant functions last.

Yes:

```
contract A {
    function A() {
        ...
    }

    function() {
        ...
    }

    // External functions
    // ...

    // External functions that are constant
    // ...

    // Public functions
    // ...

    // Internal functions
    // ...

    // Private functions
    // ...
}
```

No:

```
contract A {

    // External functions
    // ...

    // Private functions
    // ...

    // Public functions
    // ...

    function A() {
        ...
    }

    function() {
        ...
    }

    // 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:

```
spam(ham[1], Coin({name: "ham"}));
```

No:

```
spam( ham[ 1 ], Coin( { name: "ham" } ) );
```

Exception:

```
function singleLine() { spam(); }
```

Immediately before a comma, semicolon:

Yes:

```
function spam(uint i, Coin coin);
```

No:

```
function spam(uint i , Coin coin) ;
```

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 fallback function:

Yes:

```
function() {
    ...
}
```

No:

```
function () {
    ...
}
```

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:

```
contract Coin {
  struct Bank {
    address owner;
    uint balance;
  }
}
```

No:

```
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) returns (uint) {
    return x + 1;
}

function increment(uint x) public onlyowner returns (uint) {
    return x + 1;
}

```

No:

```

function increment(uint x) returns (uint)
{
    return x + 1;
}

```

```
}  
  
function increment(uint x) returns (uint) {  
    return x + 1;  
}  
  
function increment(uint x) returns (uint) {  
    return x + 1;  
}  
  
function increment(uint x) returns (uint) {  
    return x + 1;}  
}
```

The visibility modifiers for a function should come before any custom modifiers.

Yes:

```
function kill() public onlyowner {  
    selfdestruct(owner);  
}
```

No:

```
function kill() onlyowner public {  
    selfdestruct(owner);  
}
```

For long function declarations, it is recommended to drop each argument onto its 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:

```
function thisFunctionHasLotsOfArguments(  
    address a,  
    address b,  
    address c,  
    address d,  
    address e,  
    address f  
) {  
    doSomething();  
}
```

No:

```
function thisFunctionHasLotsOfArguments(address a, address b, address c,  
    address d, address e, address f) {  
    doSomething();  
}  
  
function thisFunctionHasLotsOfArguments(address a,  
                                         address b,  
                                         address c,  
                                         address d,  
                                         address e,  
                                         address f) {  
    doSomething();  
}
```

```

}

function thisFunctionHasLotsOfArguments(
    address a,
    address b,
    address c,
    address d,
    address e,
    address f) {
    doSomething();
}

```

If a long function declaration has modifiers, then each modifier should be dropped to it's 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();
}

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();
}
```

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:

```
contract A is B, C, D {
    function A(uint param1, uint param2, uint param3, uint param4, uint param5)
        B(param1)
        C(param2, param3)
        D(param4)
    {
        // do something with param5
    }
}
```

No:

```
contract A is B, C, D {
    function A(uint param1, uint param2, uint param3, uint param4, uint param5)
        B(param1)
        C(param2, param3)
        D(param4)
    {
        // do something with param5
    }
}

contract A is B, C, D {
    function A(uint param1, uint param2, uint param3, uint param4, uint param5)
        B(param1)
        C(param2, param3)
        D(param4) {
            // do something with param5
        }
}
```

When declaring short functions with a single statement, it is permissible to do it on a single line.

Permissible:

```
function shortFunction() { 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

TODO

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;
```

7.10.3 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 supercede 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 abbreviations in `CapWords`, capitalize all the letters of the abbreviation. Thus `HTTPServerError` is better than `HttpServerError`.

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`.

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`.

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.

General Recommendations

TODO

7.11 Common Patterns

7.11.1 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 `send` call, this is not recommended as it introduces a potential security risk. You may read more about this on the *Security Considerations* page.

This 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 usurped as the richest, you will receive the funds of the person who has gone on to become the new richest.

```
pragma solidity ^0.4.11;

contract WithdrawalContract {
    address public richest;
    uint public mostSent;

    mapping (address => uint) pendingWithdrawals;

    function WithdrawalContract() payable {
        richest = msg.sender;
        mostSent = msg.value;
    }

    function becomeRichest() payable returns (bool) {
        if (msg.value > mostSent) {
            pendingWithdrawals[richest] += msg.value;
            richest = msg.sender;
            mostSent = msg.value;
            return true;
        } else {
            return false;
        }
    }

    function withdraw() {
        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:

```
pragma solidity ^0.4.11;

contract SendContract {
    address public richest;
    uint public mostSent;

    function SendContract() payable {
        richest = msg.sender;
        mostSent = msg.value;
    }

    function becomeRichest() payable returns (bool) {
        if (msg.value > mostSent) {
            // This line can cause problems (explained below).
            richest.transfer(msg.value);
            richest = msg.sender;
            mostSent = msg.value;
            return true;
        } else {
            return false;
        }
    }
}
```

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 fallback function which fails (e.g. by using `revert()` or by just consuming more than the 2300 gas stipend). 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.

7.11.2 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 make 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.

```
pragma solidity ^0.4.11;

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 = now;

    // 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);
        // 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)
        onlyBy(owner)
    {
        owner = _newOwner;
    }

    modifier onlyAfter(uint _time) {
        require(now >= _time);
        _;
    }
}
```

```
/// Erase ownership information.
/// May only be called 6 weeks after
/// the contract has been created.
function disown()
    onlyBy(owner)
    onlyAfter(creationTime + 6 weeks)
{
    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);
    _;
    if (msg.value > _amount)
        msg.sender.send(msg.value - _amount);
}

function forceOwnerChange(address _newOwner)
    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.

7.11.3 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.

```
pragma solidity ^0.4.11;

contract StateMachine {
    enum Stages {
        AcceptingBlindedBids,
        RevealBids,
        AnotherStage,
        AreWeDoneYet,
        Finished
    }

    // This is the current stage.
    Stages public stage = Stages.AcceptingBlindedBids;

    uint public creationTime = now;

    modifier atStage(Stages _stage) {
        require(stage == _stage);
        _;
    }

    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 &&
            now >= creationTime + 10 days)
            nextStage();
        if (stage == Stages.RevealBids &&
            now >= creationTime + 12 days)
            nextStage();
        // The other stages transition by transaction
        _;
    }

    // Order of the modifiers matters here!
    function bid()
        payable
        timedTransitions
        atStage(Stages.AcceptingBlindedBids)
    {
```

```
    // We will not implement that here
  }

  function reveal()
    timedTransitions
    atStage(Stages.RevealBids)
  {
  }

  // This modifier goes to the next stage
  // after the function is done.
  modifier transitionNext()
  {
    _;
    nextStage();
  }

  function g()
    timedTransitions
    atStage(Stages.AnotherStage)
    transitionNext
  {
  }

  function h()
    timedTransitions
    atStage(Stages.AreWeDoneYet)
    transitionNext
  {
  }

  function i()
    timedTransitions
    atStage(Stages.Finished)
  {
  }
}
```

7.12 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. Currently, this is an object that can contain a boolean value `optimizer`, which means that the optimizer has to be switched on to enable the bug. If no conditions are given, assume that the bug is present.

```
[
  {
    "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"
  }
]

```

7.13 Contributing

Help is always appreciated!

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.

In particular, we need help in the following areas:

- Improving the documentation
- Responding to questions from other users on [StackExchange](#) and the [Solidity Gitter](#)

- Fixing and responding to Solidity's GitHub issues, especially those tagged as `up-for-grabs` which are meant as introductory issues for external contributors.

7.13.1 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.

7.13.2 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.

Additionally, if you are writing a new feature, please ensure you write appropriate Boost test cases and place them under `test/`.

However, if you are making a larger change, please consult with the Gitter channel, first.

Finally, please make sure you respect the [coding standards](#) 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!

7.13.3 Running the compiler tests

Solidity includes different types of tests. They are included in the application called `soltest`. Some of them require the `cpp-ethereum` client in testing mode, some others require `libz3` to be installed.

To disable the z3 tests, use `./build/test/soltest -- --no-smt` and to run a subset of the tests that do not require `cpp-ethereum`, use `./build/test/soltest -- --no-ipc`.

For all other tests, you need to install `cpp-ethereum` and run it in testing mode: `eth --test -d /tmp/testeth`.

Then you run the actual tests: `./build/test/soltest -- --ipcpath /tmp/testeth/geth.ipc`.

To run a subset of tests, filters can be used: `soltest -t TestSuite/TestName -- --ipcpath /tmp/testeth/geth.ipc`, where `TestName` can be a wildcard `*`.

Alternatively, there is a testing script at `scripts/test.sh` which executes all tests and runs `cpp-ethereum` automatically if it is in the path (but does not download it).

Travis CI even runs some additional tests (including `solc-js` and testing third party Solidity frameworks) that require compiling the Emscripten target.

7.13.4 Whiskers

Whiskers is a 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 *Inline Assembly* (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 will 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.

7.14 Frequently Asked Questions

This list was originally compiled by [fivedogit](#).

7.14.1 Basic Questions

Example contracts

There are some [contract examples](#) by [fivedogit](#) and there should be a [test contract](#) for every single feature of Solidity.

Create and publish the most basic contract possible

A quite simple contract is the [greeter](#)

Is it possible to do something on a specific block number? (e.g. publish a contract or execute a transaction)

Transactions are not guaranteed to happen on the next block or any future specific block, since it is up to the miners to include transactions and not up to the submitter of the transaction. This applies to function calls/transactions and contract creation transactions.

If you want to schedule future calls of your contract, you can use the [alarm clock](#).

What is the transaction “payload”?

This is just the bytecode “data” sent along with the request.

Is there a decompiler available?

There is no decompiler to Solidity. This is in principle possible to some degree, but for example variable names will be lost and great effort will be necessary to make it look similar to the original source code.

Bytecode can be decompiled to opcodes, a service that is provided by several blockchain explorers.

Contracts on the blockchain should have their original source code published if they are to be used by third parties.

Create a contract that can be killed and return funds

First, a word of warning: Killing contracts sounds like a good idea, because “cleaning up” is always good, but as seen above, it does not really clean up. Furthermore, if Ether is sent to removed contracts, the Ether will be forever lost.

If you want to deactivate your contracts, it is preferable to **disable** them by changing some internal state which causes all functions to throw. This will make it impossible to use the contract and ether sent to the contract will be returned automatically.

Now to answering the question: Inside a constructor, `msg.sender` is the creator. Save it. Then `selfdestruct(creator);` to kill and return funds.

example

Note that if you `import "mortal"` at the top of your contracts and declare `contract SomeContract is mortal { ...` and compile with a compiler that already has it (which includes [Remix](#)), then `kill()` is taken care of for you. Once a contract is “mortal”, then you can `contractname.kill.sendTransaction({from:eth.coinbase})`, just the same as my examples.

Store Ether in a contract

The trick is to create the contract with `{from:someaddress, value: web3.toWei(3,"ether") ...}`

See `endowment_retriever.sol`.

Use a non-constant function (req `sendTransaction`) to increment a variable in a contract

See `value_incrementer.sol`.

Get a contract to return its funds to you (not using `selfdestruct(...)`).

This example demonstrates how to send funds from a contract to an address.

See `endowment_retriever`.

Can you return an array or a `string` from a solidity function call?

Yes. See `array_receiver_and_returner.sol`.

What is problematic, though, is returning any variably-sized data (e.g. a variably-sized array like `uint[]`) from a fuction **called from within Solidity**. This is a limitation of the EVM and will be solved with the next protocol update.

Returning variably-sized data as part of an external transaction or call is fine.

Is it possible to in-line initialize an array like so: `string[] myarray = ["a", "b"];`

Yes. However it should be noted that this currently only works with statically sized memory arrays. You can even create an inline memory array in the return statement. Pretty cool, huh?

Example:

```
pragma solidity ^0.4.0;

contract C {
    function f() returns (uint8[5]) {
```

```
    string[4] memory adaArr = ["This", "is", "an", "array"];
    return ([1, 2, 3, 4, 5]);
}
}
```

Can a contract function return a struct?

Yes, but only in internal function calls.

If I return an enum, I only get integer values in web3.js. How to get the named values?

Enums are not supported by the ABI, they are just supported by Solidity. You have to do the mapping yourself for now, we might provide some help later.

Can state variables be initialized in-line?

Yes, this is possible for all types (even for structs). However, for arrays it should be noted that you must declare them as static memory arrays.

Examples:

```
pragma solidity ^0.4.0;

contract C {
    struct S {
        uint a;
        uint b;
    }

    S public x = S(1, 2);
    string name = "Ada";
    string[4] adaArr = ["This", "is", "an", "array"];
}

contract D {
    C c = new C();
}
```

How do structs work?

See [struct_and_for_loop_tester.sol](#).

How do for loops work?

Very similar to JavaScript. There is one point to watch out for, though:

If you use `for (var i = 0; i < a.length; i++) { a[i] = i; }`, then the type of `i` will be inferred only from `0`, whose type is `uint8`. This means that if `a` has more than 255 elements, your loop will not terminate because `i` can only hold values up to 255.

Better use `for (uint i = 0; i < a.length...`

See [struct_and_for_loop_tester.sol](#).

What are some examples of basic string manipulation (substring, indexOf, charAt, etc)?

There are some string utility functions at `stringUtils.sol` which will be extended in the future. In addition, Arachnid has written `solidity-stringutils`.

For now, if you want to modify a string (even when you only want to know its length), you should always convert it to a `bytes` first:

```
pragma solidity ^0.4.0;

contract C {
    string s;

    function append(byte c) {
        bytes(s).push(c);
    }

    function set(uint i, byte c) {
        bytes(s)[i] = c;
    }
}
```

Can I concatenate two strings?

You have to do it manually for now.

Why is the low-level function `.call()` less favorable than instantiating a contract with a variable (`ContractB b;`) and executing its functions (`b.doSomething();`)?

If you use actual functions, the compiler will tell you if the types or your arguments do not match, if the function does not exist or is not visible and it will do the packing of the arguments for you.

See `ping.sol` and `pong.sol`.

Is unused gas automatically refunded?

Yes and it is immediate, i.e. done as part of the transaction.

When returning a value of say `uint` type, is it possible to return an undefined or “null”-like value?

This is not possible, because all types use up the full value range.

You have the option to `throw` on error, which will also revert the whole transaction, which might be a good idea if you ran into an unexpected situation.

If you do not want to throw, you can return a pair:

```
pragma solidity ^0.4.0;

contract C {
    uint[] counters;

    function getCounter(uint index)
        returns (uint counter, bool error) {
```

```
        if (index >= counters.length)
            return (0, true);
        else
            return (counters[index], false);
    }

    function checkCounter(uint index) {
        var (counter, error) = getCounter(index);
        if (error) {
            // ...
        } else {
            // ...
        }
    }
}
```

Are comments included with deployed contracts and do they increase deployment gas?

No, everything that is not needed for execution is removed during compilation. This includes, among others, comments, variable names and type names.

What happens if you send ether along with a function call to a contract?

It gets added to the total balance of the contract, just like when you send ether when creating a contract. You can only send ether along to a function that has the `payable` modifier, otherwise an exception is thrown.

Is it possible to get a tx receipt for a transaction executed contract-to-contract?

No, a function call from one contract to another does not create its own transaction, you have to look in the overall transaction. This is also the reason why several block explorer do not show Ether sent between contracts correctly.

What is the `memory` keyword? What does it do?

The Ethereum Virtual Machine has three areas where it can store items.

The first is “storage”, where all the contract state variables reside. Every contract has its own storage and it is persistent between function calls and quite expensive to use.

The second is “memory”, this is used to hold temporary values. It is erased between (external) function calls and is cheaper to use.

The third one is the stack, which is used to hold small local variables. It is almost free to use, but can only hold a limited amount of values.

For almost all types, you cannot specify where they should be stored, because they are copied everytime they are used.

The types where the so-called storage location is important are structs and arrays. If you e.g. pass such variables in function calls, their data is not copied if it can stay in memory or stay in storage. This means that you can modify their content in the called function and these modifications will still be visible in the caller.

There are defaults for the storage location depending on which type of variable it concerns:

- state variables are always in storage
- function arguments are in memory by default

- local variables of struct, array or mapping type reference storage by default
- local variables of value type (i.e. neither array, nor struct nor mapping) are stored in the stack

Example:

```
pragma solidity ^0.4.0;

contract C {
    uint[] data1;
    uint[] data2;

    function appendOne() {
        append(data1);
    }

    function appendTwo() {
        append(data2);
    }

    function append(uint[] storage d) internal {
        d.push(1);
    }
}
```

The function `append` can work both on `data1` and `data2` and its modifications will be stored permanently. If you remove the `storage` keyword, the default is to use `memory` for function arguments. This has the effect that at the point where `append(data1)` or `append(data2)` is called, an independent copy of the state variable is created in memory and `append` operates on this copy (which does not support `.push` - but that is another issue). The modifications to this independent copy do not carry back to `data1` or `data2`.

A common mistake is to declare a local variable and assume that it will be created in memory, although it will be created in storage:

```
/// THIS CONTRACT CONTAINS AN ERROR

pragma solidity ^0.4.0;

contract C {
    uint someVariable;
    uint[] data;

    function f() {
        uint[] x;
        x.push(2);
        data = x;
    }
}
```

The type of the local variable `x` is `uint[] storage`, but since `storage` is not dynamically allocated, it has to be assigned from a state variable before it can be used. So no space in storage will be allocated for `x`, but instead it functions only as an alias for a pre-existing variable in storage.

What will happen is that the compiler interprets `x` as a storage pointer and will make it point to the storage slot 0 by default. This has the effect that `someVariable` (which resides at storage slot 0) is modified by `x.push(2)`.

The correct way to do this is the following:

```
pragma solidity ^0.4.0;
```

```
contract C {
    uint someVariable;
    uint[] data;

    function f() {
        uint[] x = data;
        x.push(2);
    }
}
```

7.14.2 Advanced Questions

How do you get a random number in a contract? (Implement a self-returning gambling contract.)

Getting randomness right is often the crucial part in a crypto project and most failures result from bad random number generators.

If you do not want it to be safe, you build something similar to the [coin flipper](#) but otherwise, rather use a contract that supplies randomness, like the [RANDAO](#).

Get return value from non-constant function from another contract

The key point is that the calling contract needs to know about the function it intends to call.

See [ping.sol](#) and [pong.sol](#).

Get contract to do something when it is first mined

Use the constructor. Anything inside it will be executed when the contract is first mined.

See [replicator.sol](#).

How do you create 2-dimensional arrays?

See [2D_array.sol](#).

Note that filling a 10x10 square of `uint8` + contract creation took more than 800,000 gas at the time of this writing. 17x17 took 2,000,000 gas. With the limit at 3.14 million... well, there's a pretty low ceiling for what you can create right now.

Note that merely "creating" the array is free, the costs are in filling it.

Note2: Optimizing storage access can pull the gas costs down considerably, because 32 `uint8` values can be stored in a single slot. The problem is that these optimizations currently do not work across loops and also have a problem with bounds checking. You might get much better results in the future, though.

What happens to a struct's mapping when copying over a struct?

This is a very interesting question. Suppose that we have a contract field set up like such:

```

struct user {
    mapping(string => string) comments;
}

function somefunction {
    user user1;
    user1.comments["Hello"] = "World";
    user user2 = user1;
}

```

In this case, the mapping of the struct being copied over into the userList is ignored as there is no “list of mapped keys”. Therefore it is not possible to find out which values should be copied over.

How do I initialize a contract with only a specific amount of wei?

Currently the approach is a little ugly, but there is little that can be done to improve it. In the case of a contract A calling a new instance of contract B, parentheses have to be used around `new B` because `B.value` would refer to a member of B called `value`. You will need to make sure that you have both contracts aware of each other’s presence and that contract B has a payable constructor. In this example:

```

pragma solidity ^0.4.0;

contract B {
    function B() payable {}
}

contract A {
    address child;

    function test() {
        child = (new B).value(10)(); //construct a new B with 10 wei
    }
}

```

Can a contract function accept a two-dimensional array?

This is not yet implemented for external calls and dynamic arrays - you can only use one level of dynamic arrays.

What is the relationship between bytes32 and string? Why is it that bytes32 somevar = "stringliteral"; works and what does the saved 32-byte hex value mean?

The type `bytes32` can hold 32 (raw) bytes. In the assignment `bytes32 somevar = "stringliteral";`, the string literal is interpreted in its raw byte form and if you inspect `somevar` and see a 32-byte hex value, this is just "stringliteral" in hex.

The type `bytes` is similar, only that it can change its length.

Finally, `string` is basically identical to `bytes` only that it is assumed to hold the UTF-8 encoding of a real string. Since `string` stores the data in UTF-8 encoding it is quite expensive to compute the number of characters in the string (the encoding of some characters takes more than a single byte). Because of that, `string s; s.length` is not yet supported and not even index access `s[2]`. But if you want to access the low-level byte encoding of the string, you can use `bytes(s).length` and `bytes(s)[2]` which will result in the number of bytes in the UTF-8 encoding of the string (not the number of characters) and the second byte (not character) of the UTF-8 encoded string, respectively.

Can a contract pass an array (static size) or string or bytes (dynamic size) to another contract?

Sure. Take care that if you cross the memory / storage boundary, independent copies will be created:

```
pragma solidity ^0.4.0;

contract C {
    uint[20] x;

    function f() {
        g(x);
        h(x);
    }

    function g(uint[20] y) internal {
        y[2] = 3;
    }

    function h(uint[20] storage y) internal {
        y[3] = 4;
    }
}
```

The call to `g(x)` will not have an effect on `x` because it needs to create an independent copy of the storage value in memory (the default storage location is memory). On the other hand, `h(x)` successfully modifies `x` because only a reference and not a copy is passed.

Sometimes, when I try to change the length of an array with `ex: arrayname.length = 7;` I get a compiler error `Value must be an lvalue. Why?`

You can resize a dynamic array in storage (i.e. an array declared at the contract level) with `arrayname.length = <some new length>;`. If you get the “lvalue” error, you are probably doing one of two things wrong.

1. You might be trying to resize an array in “memory”, or
2. You might be trying to resize a non-dynamic array.

```
int8[] memory memArr;           // Case 1
memArr.length++;                // illegal
int8[5] storageArr;            // Case 2
somearray.length++;            // legal
int8[5] storage storageArr2;    // Explicit case 2
somearray2.length++;           // legal
```

Important note: In Solidity, array dimensions are declared backwards from the way you might be used to declaring them in C or Java, but they are access as in C or Java.

For example, `int8[][5] somearray;` are 5 dynamic `int8` arrays.

The reason for this is that `T[5]` is always an array of 5 `T`'s, no matter whether `T` itself is an array or not (this is not the case in C or Java).

Is it possible to return an array of strings (`string[]`) from a Solidity function?

Not yet, as this requires two levels of dynamic arrays (`string` is a dynamic array itself).

If you issue a call for an array, it is possible to retrieve the whole array? Or must you write a helper function for that?

The automatic *getter function* for a public state variable of array type only returns individual elements. If you want to return the complete array, you have to manually write a function to do that.

What could have happened if an account has storage value(s) but no code? Example: <http://test.ether.camp/account/5f740b3a43fbb99724ce93a879805f4dc89178b5>

The last thing a constructor does is returning the code of the contract. The gas costs for this depend on the length of the code and it might be that the supplied gas is not enough. This situation is the only one where an “out of gas” exception does not revert changes to the state, i.e. in this case the initialisation of the state variables.

<https://github.com/ethereum/wiki/wiki/Subtleties>

After a successful CREATE operation’s sub-execution, if the operation returns x , $5 * \text{len}(x)$ gas is subtracted from the remaining gas before the contract is created. If the remaining gas is less than $5 * \text{len}(x)$, then no gas is subtracted, the code of the created contract becomes the empty string, but this is not treated as an exceptional condition - no reverts happen.

What does the following strange check do in the Custom Token contract?

```
require((balanceOf[_to] + _value) >= balanceOf[_to]);
```

Integers in Solidity (and most other machine-related programming languages) are restricted to a certain range. For `uint256`, this is 0 up to $2^{256} - 1$. If the result of some operation on those numbers does not fit inside this range, it is truncated. These truncations can have *serious consequences*, so code like the one above is necessary to avoid certain attacks.

More Questions?

If you have more questions or your question is not answered here, please talk to us on [gitter](#) or file an [issue](#).

A

abi, 122
abstract contract, **79**
access
 restricting, 153
account, **19**
addmod, 57, 105
address, 19, 41, 43
anonymous, 107
application binary interface, 122
array, 48, **49**
 allocating, **49**
 length, **50**
 literals, **50**
 push, **50**
asm, **85, 131**
assembly, **85, 131**
assert, 57, **64, 105**
assignment, 53, **62**
 destructuring, **62**
auction
 blind, 28
 open, 28

B

balance, 19, 41, 58, 105
ballot, 25
base
 constructor, **79**
base class, **76**
blind auction, 28
block, **18, 56, 105**
 number, 56, 105
 timestamp, 56, 105
bool, **40**
break, 59
Bugs, 156
byte array, 43
bytes, 45

bytes32, 43

C

C3 linearization, **79**
call, 41, 58
callcode, 20, 41, 58, 80
cast, **54**
coding style, 139
coin, 17
coinbase, 56, 105
commandline compiler, **115**
comment, **37**
common subexpression elimination, 103
compiler
 commandline, 115
constant, **71, 107**
constant propagation, 103
constructor, 66
 arguments, 66
continue, 59
contract, 38, **66**
 abstract, **79**
 base, **76**
 creation, **66**
 interface, **80**
contract creation, 21
contract verification, 120
contracts
 creating, 61
cryptography, 57, 105

D

data, 56, 105
days, 55
declarations, 63
default value, 63
delegatecall, 20, 41, 58, 80
delete, **53**
deriving, **76**

difficulty, 56, 105
do/while, 59

E

ecrecover, 57, 105
else, 59
enum, 38, 45
escrow, 33
ether, 55
ethereum virtual machine, **18**
event, 17, 38, **74**
evm, **18**
evmasm, **85, 131**
exception, **64**
external, 67, 106

F

fallback function, **73**
false, **40**
finney, 55
fixed, **41**
fixed point number, **41**
for, 59
function, 38
 call, 20, **59**
 external, 59
 fallback, 73
 getter, **69**
 internal, 59
 modifier, 38, **70, 153, 154**
function type, **45**

G

gas, **19, 56, 105**
gas price, **19, 56, 105**
getter
 function, **69**
goto, 59

H

hours, 55

I

if, 59
import, **35**
indexed, 107
inheritance, **76**
 multiple, **79**
inline
 arrays, **50**
installing, **21**
instruction, **20**
int, **40**
integer, **40**

interface contract, **80**
internal, 67, 106

J

julia, **131**

K

keccak256, 57, 105

L

length, 50
library, 20, **80, 83**
linearization, **79**
linker, **115**
literal, 44, 45
 address, 43
 rational, 44
 string, 44
location, 48
log, 20, **75**
lvalue, 53

M

mapping, 17, **53, 101**
memory, **19, 48**
message call, **20**
metadata, 120
minutes, 55
modifiers, 107
msg, 56, 105
mulmod, 57, 105

N

natspec, 37
new, 49, **61**
now, 56, 105
number, 56, 105

O

open auction, 28
optimizer, 103
origin, 56, 105

P

parameter, **58**
 input, 58
 output, 58
payable, 107
pragma, **35**
precedence, 105
private, 67, 106
public, 67, 106
purchase, 33

pure, 107
push, 50

R

reference type, **48**
remote purchase, 33
require, 57, **64**, 105
return, 59
revert, 57, **64**, 105
ripemd160, 57, 105

S

scoping, **63**
seconds, 55
selfdestruct, 21, 58, 105
send, 41, 58, 105
sender, 56, 105
set, 81
sha256, 57, 105
solc, **115**
source file, 35
source mappings, 104
stack, **19**
state machine, 154
state variable, 38, 101
storage, **19**, 19, 48, 101
string, 44
struct, 38, 48, **52**
style, 139
subcurrency, **16**
super, 105
switch, 59
szabo, 55

T

this, 58, 105
throw, **64**
time, 55
timestamp, 56, 105
transaction, 18, **19**
transfer, 41, 58
true, **40**
type, 39
 conversion, **54**
 deduction, **55**
 function, **45**
 reference, **48**
 struct, **52**
 value, **40**

U

ufixed, **41**
uint, **40**
using for, 81, **83**

V

value, 56, 105
value type, **40**
var, **55**
version, 35
view, 107
visibility, **67**, 106
voting, 25

W

weeks, 55
wei, 55
while, 59
withdrawal, 151

Y

years, 55