ETHEREUM: A SECURE DECENTRALISED GENERALISED TRANSACTION LEDGER
EIP-150 REVISION (a04ea02 - 2017-09-30)

DR. GAVIN WOOD
FOUNDER, ETHEREUM & ETHCORE
GAVIN@ETHCORE.IO

ABSTRACT. The blockchain paradigm when coupled with cryptographically-secured transactions has demonstrated its utility through a number of projects, not the least being Bitcoin. Each such project can be seen as a simple application on a decentralised, but singleton, compute resource. We can call this paradigm a transactional singleton machine with shared-state.

Ethereum implements this paradigm in a generalised manner. Furthermore it provides a plurality of such resources, each with a distinct state and operating code but able to interact through a message-passing framework with others. We discuss its design, implementation issues, the opportunities it provides and the future hurdles we envisage.

1. INTRODUCTION

With ubiquitous internet connections in most places of the world, global information transmission has become incredibly cheap. Technology-rooted movements like Bitcoin have demonstrated, through the power of the default, consensus mechanisms and voluntary respect of the social contract that it is possible to use the internet to make a decentralised value-transfer system, shared across the world and virtually free to use. This system can be said to be a very specialised version of a cryptographically secure, transaction-based state machine. Follow-up systems such as Namecoin adapted this original "currency application" of the technology into other applications albeit rather simplistic ones.

Ethereum is a project which attempts to build the generalised technology; technology on which all transaction-based state machine concepts may be built. Moreover it aims to provide to the end-developer a tightly integrated end-to-end system for building software on a hitherto unexplored compute paradigm in the mainstream: a trustful object messaging compute framework.

1.1. Driving Factors.
There are many goals of this project; one key goal is to facilitate transactions between consenting individuals who would otherwise have no means to trust one another. This may be due to geographical separation, interfacing difficulty, or perhaps the incompatibility, incompetence, unwillingness, expense, uncertainty, inconvenience or corruption of existing legal systems. By specifying a state-change system through a rich and unambiguous language, and furthermore architecting a system such that we can reasonably expect that an agreement will be thus enforced autonomously, we can provide a means to this end.

Dealings in this proposed system would have several attributes not often found in the real world. The incorruptibility of judgement, often difficult to find, comes naturally from a disinterested algorithmic interpreter. Transparency, or being able to see exactly how a state or judgement came about through the transaction log and rules or instructional codes, never happens perfectly in human-based systems since natural language is necessarily vague, information is often lacking, and plain old prejudices are difficult to shake.

Overall, I wish to provide a system such that users can be guaranteed that no matter with which other individuals, systems or organisations they interact, they can do so with absolute confidence in the possible outcomes and how those outcomes might come about.

1.2. Previous Work.
Buterin [2013a] first proposed the kernel of this work in late November, 2013. Though now evolved in many ways, the key functionality of a blockchain with a Turing-complete language and an effectively unlimited inter-transaction storage capability remains unchanged.

Dwork and Naor [1992] provided the first work into the usage of a cryptographic proof of computational expenditure ("proof-of-work") as a means of transmitting a value signal over the Internet. The value-signal was utilised here as a spam deterrence mechanism rather than any kind of currency, but critically demonstrated the potential for a basic data channel to carry a *strong economic signal*, allowing a receiver to make a physical assertion without having to rely upon *trust*. Back [2002] later produced a system in a similar vein.

The first example of utilising the proof-of-work as a strong economic signal to secure a currency was by Vishnumurthy et al. [2003]. In this instance, the token was used to keep peer-to-peer file trading in check, providing "consumers" with the ability to make micro-payments to "suppliers" for their services. The security model afforded by the proof-of-work was augmented with digital signatures and a ledger in order to ensure that the historical record couldn't be corrupted and that malicious actors could not spoof payment or unjustly complain about service delivery. Five years later, Nakamoto [2008] introduced another such proof-of-work-secured value token, somewhat wider in scope. The fruits of this project, Bitcoin, became the first widely adopted global decentralised transaction ledger.

Other projects built on Bitcoin's success; the alt-coins introduced numerous other currencies through alteration to the protocol. Some of the best known are Litecoin and Primecoin, discussed by Sprankel [2013]. Other projects sought to take the core value content mechanism of the

protocol and repurpose it; Aron [2012] discusses, for example, the Namecoin project which aims to provide a decentralised name-resolution system.

Other projects still aim to build upon the Bitcoin network itself, leveraging the large amount of value placed in the system and the vast amount of computation that goes into the consensus mechanism. The Mastercoin project, first proposed by Willett [2013], aims to build a richer protocol involving many additional high-level features on top of the Bitcoin protocol through utilisation of a number of auxiliary parts to the core protocol. The Coloured Coins project, proposed by Rosenfeld [2012], takes a similar but more simplified strategy, embellishing the rules of a transaction in order to break the fungibility of Bitcoin's base currency and allow the creation and tracking of tokens through a special "chroma-wallet"-protocol-aware piece of software.

Additional work has been done in the area with discarding the decentralisation foundation; Ripple, discussed by Boutellier and Heinzen [2014], has sought to create a "federated" system for currency exchange, effectively creating a new financial clearing system. It has demonstrated that high efficiency gains can be made if the decentralisation premise is discarded.

Early work on smart contracts has been done by Szabo [1997] and Miller [1997]. Around the 1990s it became clear that algorithmic enforcement of agreements could become a significant force in human cooperation. Though no specific system was proposed to implement such a system, it was proposed that the future of law would be heavily affected by such systems. In this light, Ethereum may be seen as a general implementation of such a *crypto-law* system.

For a list of terms used in this paper, refer to Appendix A.

2. THE BLOCKCHAIN PARADIGM

Ethereum, taken as a whole, can be viewed as a transaction-based state machine: we begin with a genesis state and incrementally execute transactions to morph it into some final state. It is this final state which we accept as the canonical "version" of the world of Ethereum. The state can include such information as account balances, reputations, trust arrangements, data pertaining to information of the physical world; in short, anything that can currently be represented by a computer is admissible. Transactions thus represent a valid arc between two states; the 'valid' part is important—there exist far more invalid state changes than valid state changes. Invalid state changes might, e.g., be things such as reducing an account balance without an equal and opposite increase elsewhere. A valid state transition is one which comes about through a transaction. Formally:

$$(1) \qquad \sigma_{t+1} \equiv \Upsilon(\sigma_\mathbf{t}, T)$$

where Υ is the Ethereum state transition function. In Ethereum, Υ, together with σ are considerably more powerful then any existing comparable system; Υ allows components to carry out arbitrary computation, while σ allows components to store arbitrary state between transactions.

Transactions are collated into blocks; blocks are chained together using a cryptographic hash as a means of reference. Blocks function as a journal or ledger, recording a series of transactions together with the previous block and an identifier for the final state (though blocks do not store the final state itself—that would be far too big). They also punctuate the transaction series with incentives for nodes to *mine*. This incentivisation takes place as a state-transition function, adding value to a nominated account.

Mining is the process of dedicating effort (working) to bolster one series of transactions (a block) over any other potential competitor block. It is achieved thanks to a cryptographically secure proof. This scheme is known as a proof-of-work and is discussed in detail in section 11.5.

Formally, we expand to:

$$(2) \qquad \sigma_{t+1} \equiv \Pi(\sigma_\mathbf{t}, B)$$
$$(3) \qquad B \equiv (..., (T_0, T_1, ...))$$
$$(4) \qquad \Pi(\sigma, B) \equiv \Omega(B, \Upsilon(\Upsilon(\sigma, T_0), T_1)...)$$

Where Ω is the block-finalisation state transition function (a function that rewards a nominated party); B is this block, which includes a series of transactions amongst some other components; and Π is the block-level state-transition function.

This is the basis of the blockchain paradigm, a model that forms the backbone of not only Ethereum, but all decentralised consensus-based transaction systems to date.

2.1. **Value.** In order to incentivise computation within the network, there needs to be an agreed method for transmitting value. To address this issue, Ethereum has an intrinsic currency, Ether, known also as ETH and sometimes referred to by the Old English Ð. The smallest subdenomination of Ether, and thus the one in which all integer values of the currency are counted, is the Wei. One Ether is defined as being 10^{18} Wei. There exist other subdenominations of Ether:

Multiplier	Name
10^0	Wei
10^{12}	Szabo
10^{15}	Finney
10^{18}	Ether

Throughout the present work, any reference to value, in the context of Ether, currency, a balance or a payment, should be assumed to be counted in Wei.

2.2. **Which History?** Since the system is decentralised and all parties have an opportunity to create a new block on some older pre-existing block, the resultant structure is necessarily a tree of blocks. In order to form a consensus as to which path, from root (the genesis block) to leaf (the block containing the most recent transactions) through this tree structure, known as the blockchain, there must be an agreed-upon scheme. If there is ever a disagreement between nodes as to which root-to-leaf path down the block tree is the 'best' blockchain, then a *fork* occurs.

This would mean that past a given point in time (block), multiple states of the system may coexist: some nodes believing one block to contain the canonical transactions, other nodes believing some other block to be canonical, potentially containing radically different or incompatible transactions. This is to be avoided at all costs as the uncertainty that would ensue would likely kill all confidence in the entire system.

The scheme we use in order to generate consensus is a simplified version of the GHOST protocol introduced by Sompolinsky and Zohar [2013]. This process is described in detail in section 10.

Sometimes, a path follows a new protocol from a particular height. This document describes one version of the protocol. In order to follow back the history of a path, one might need to reference multiple versions of this document.

3. CONVENTIONS

I use a number of typographical conventions for the formal notation, some of which are quite particular to the present work:

The two sets of highly structured, 'top-level', state values, are denoted with bold lowercase Greek letters. They fall into those of world-state, which are denoted $\boldsymbol{\sigma}$ (or a variant thereupon) and those of machine-state, $\boldsymbol{\mu}$.

Functions operating on highly structured values are denoted with an upper-case greek letter, e.g. Υ, the Ethereum state transition function.

For most functions, an uppercase letter is used, e.g. C, the general cost function. These may be subscripted to denote specialised variants, e.g. C_{SSTORE}, the cost function for the SSTORE operation. For specialised and possibly externally defined functions, I may format as typewriter text, e.g. the Keccak-256 hash function (as per the winning entry to the SHA-3 contest, rather than later releases (Guido Bertoni and Joan Daemen and Michaël Peeters and Gilles Van Assche and Ronny Van Keer [2017], Kec)) is denoted KEC (and generally referred to as plain Keccak). Also KEC512 is referring to the Keccak 512 hash function.

Tuples are typically denoted with an upper-case letter, e.g. T, is used to denote an Ethereum transaction. This symbol may, if accordingly defined, be subscripted to refer to an individual component, e.g. $T_{\mathbf{n}}$, denotes the nonce of said transaction. The form of the subscript is used to denote its type; e.g. uppercase subscripts refer to tuples with subscriptable components.

Scalars and fixed-size byte sequences (or, synonymously, arrays) are denoted with a normal lower-case letter, e.g. n is used in the document to denote a transaction nonce. Those with a particularly special meaning may be greek, e.g. δ, the number of items required on the stack for a given operation.

Arbitrary-length sequences are typically denoted as a bold lower-case letter, e.g. $\mathbf{o}$ is used to denote the byte sequence given as the output data of a message call. For particularly important values, a bold uppercase letter may be used.

Throughout, we assume scalars are positive integers and thus belong to the set $\mathbb{P}$. The set of all byte sequences is $\mathbb{B}$, formally defined in Appendix B. If such a set of sequences is restricted to those of a particular length, it is denoted with a subscript, thus the set of all byte sequences of length 32 is named $\mathbb{B}_{32}$ and the set of all positive integers smaller than 2^{256} is named $\mathbb{P}_{256}$. This is formally defined in section 4.3.

Square brackets are used to index into and reference individual components or subsequences of sequences, e.g. $\boldsymbol{\mu}_{\mathbf{s}}[0]$ denotes the first item on the machine's stack. For subsequences, ellipses are used to specify the intended range, to include elements at both limits, e.g. $\boldsymbol{\mu}_{\mathbf{m}}[0..31]$ denotes the first 32 items of the machine's memory.

In the case of the global state $\boldsymbol{\sigma}$, which is a sequence of accounts, themselves tuples, the square brackets are used to reference an individual account.

When considering variants of existing values, I follow the rule that within a given scope for definition, if we assume that the unmodified 'input' value be denoted by the placeholder $\square$ then the modified and utilisable value is denoted as $\square'$, and intermediate values would be $\square^*$, $\square^{**}$ &c. On very particular occasions, in order to maximise readability and only if unambiguous in meaning, I may use alpha-numeric subscripts to denote intermediate values, especially those of particular note.

When considering the use of existing functions, given a function f, the function f^* denotes a similar, element-wise version of the function mapping instead between sequences.

I define a number of useful functions throughout. One of the more common is ℓ, which evaluates to the last item in the given sequence:

$$(5) \qquad \ell(\mathbf{x}) \equiv \mathbf{x}[\|\mathbf{x}\| - 1]$$

4. BLOCKS, STATE AND TRANSACTIONS

Having introduced the basic concepts behind Ethereum, we will discuss the meaning of a transaction, a block and the state in more detail.

4.1. **World State.** The world state (*state*), is a mapping between addresses (160-bit identifiers) and account states (a data structure serialised as RLP, see Appendix B). Though not stored on the blockchain, it is assumed that the implementation will maintain this mapping in a modified Merkle Patricia tree (*trie*, see Appendix D). The trie requires a simple database backend that maintains a mapping of bytearrays to bytearrays; we name this underlying database the state database. This has a number of benefits; firstly the root node of this structure is cryptographically dependent on all internal data and as such its hash can be used as a secure identity for the entire system state. Secondly, being an immutable data structure, it allows any previous state (whose root hash is known) to be recalled by simply altering the root hash accordingly. Since we store all such root hashes in the blockchain, we are able to trivially revert to old states.

The account state comprises the following four fields:

nonce: A scalar value equal to the number of transactions sent from this address or, in the case of accounts with associated code, the number of contract-creations made by this account. For account of address a in state $\boldsymbol{\sigma}$, this would be formally denoted $\boldsymbol{\sigma}[a]_{\mathbf{n}}$.

balance: A scalar value equal to the number of Wei owned by this address. Formally denoted $\boldsymbol{\sigma}[a]_{\mathbf{b}}$.

storageRoot: A 256-bit hash of the root node of a Merkle Patricia tree that encodes the storage contents of the account (a mapping between 256-bit integer values), encoded into the trie as a mapping from the Keccak 256-bit hash of the 256-bit integer keys to the RLP-encoded 256-bit integer values. The hash is formally denoted $\boldsymbol{\sigma}[a]_{\mathbf{s}}$.

codeHash: The hash of the EVM code of this account—this is the code that gets executed should this address receive a message call; it is immutable and thus, unlike all other fields, cannot be changed after construction. All such code fragments are contained in the state database under their corresponding hashes for later retrieval. This hash is formally denoted $\boldsymbol{\sigma}[a]_{\mathbf{c}}$, and thus the code may be denoted as $\mathbf{b}$, given that $\mathtt{KEC}(\mathbf{b}) = \boldsymbol{\sigma}[a]_{\mathbf{c}}$.

Since I typically wish to refer not to the trie's root hash but to the underlying set of key/value pairs stored within, I define a convenient equivalence:

$$(6) \qquad \mathtt{TRIE}\big(L_I^*(\boldsymbol{\sigma}[a]_{\mathbf{s}})\big) \equiv \boldsymbol{\sigma}[a]_{\mathbf{s}}$$

The collapse function for the set of key/value pairs in the trie, L_I^*, is defined as the element-wise transformation of the base function L_I, given as:

$$(7) \qquad L_I\big((k, v)\big) \equiv \big(\mathtt{KEC}(k), \mathtt{RLP}(v)\big)$$

where:

$$(8) \qquad k \in \mathbb{B}_{32} \quad \wedge \quad v \in \mathbb{P}$$

It shall be understood that $\boldsymbol{\sigma}[a]_{\mathbf{s}}$ is not a 'physical' member of the account and does not contribute to its later serialisation.

If the **codeHash** field is the Keccak-256 hash of the empty string, i.e. $\boldsymbol{\sigma}[a]_{\mathbf{c}} = \mathtt{KEC}\big(()\big)$, then the node represents a simple account, sometimes referred to as a "non-contract" account.

Thus we may define a world-state collapse function L_S:

$$(9) \qquad L_S(\boldsymbol{\sigma}) \equiv \{p(a) : \boldsymbol{\sigma}[a] \neq \varnothing\}$$

where

$$(10) \quad p(a) \equiv \big(\mathtt{KEC}(a), \mathtt{RLP}((\boldsymbol{\sigma}[a]_{\mathbf{n}}, \boldsymbol{\sigma}[a]_{\mathbf{b}}, \boldsymbol{\sigma}[a]_{\mathbf{s}}, \boldsymbol{\sigma}[a]_{\mathbf{c}}))\big)$$

This function, L_S, is used alongside the trie function to provide a short identity (hash) of the world state. We assume:

$$(11) \qquad \forall a : \boldsymbol{\sigma}[a] = \varnothing \ \vee \ (a \in \mathbb{B}_{20} \ \wedge \ v(\boldsymbol{\sigma}[a]))$$

where v is the account validity function:

$$(12) \quad v(x) \equiv x_{\mathbf{n}} \in \mathbb{P}_{256} \wedge x_{\mathbf{b}} \in \mathbb{P}_{256} \wedge x_{\mathbf{s}} \in \mathbb{B}_{32} \wedge x_{\mathbf{c}} \in \mathbb{B}_{32}$$

4.2. The Transaction. A transaction (formally, T) is a single cryptographically-signed instruction constructed by an actor externally to the scope of Ethereum. While it is assumed that the ultimate external actor will be human in nature, software tools will be used in its construction and dissemination[1]. There are two types of transactions: those which result in message calls and those which result in the creation of new accounts with associated code (known informally as 'contract creation'). Both types specify a number of common fields:

nonce: A scalar value equal to the number of transactions sent by the sender; formally $T_{\mathbf{n}}$.

gasPrice: A scalar value equal to the number of Wei to be paid per unit of *gas* for all computation costs incurred as a result of the execution of this transaction; formally $T_{\mathbf{p}}$.

gasLimit: A scalar value equal to the maximum amount of gas that should be used in executing this transaction. This is paid up-front, before any computation is done and may not be increased later; formally $T_{\mathbf{g}}$.

to: The 160-bit address of the message call's recipient or, for a contract creation transaction, $\varnothing$, used here to denote the only member of $\mathbb{B}_0$; formally $T_{\mathbf{t}}$.

value: A scalar value equal to the number of Wei to be transferred to the message call's recipient or, in the case of contract creation, as an endowment to the newly created account; formally $T_{\mathbf{v}}$.

v, r, s: Values corresponding to the signature of the transaction and used to determine the sender of the transaction; formally $T_{\mathbf{w}}$, $T_{\mathbf{r}}$ and $T_{\mathbf{s}}$. This is expanded in Appendix F.

Additionally, a contract creation transaction contains:

init: An unlimited size byte array specifying the EVM-code for the account initialisation procedure, formally $T_{\mathbf{i}}$.

init is an EVM-code fragment; it returns the **body**, a second fragment of code that executes each time the account receives a message call (either through a transaction or due to the internal execution of code). **init** is executed only once at account creation and gets discarded immediately thereafter.

In contrast, a message call transaction contains:

data: An unlimited size byte array specifying the input data of the message call, formally $T_{\mathbf{d}}$.

$$(13)$$
$$L_{\mathbf{T}}(T) \equiv \begin{cases} (T_{\mathbf{n}}, T_{\mathbf{p}}, T_{\mathbf{g}}, T_{\mathbf{t}}, T_{\mathbf{v}}, T_{\mathbf{i}}, T_{\mathbf{w}}, T_{\mathbf{r}}, T_{\mathbf{s}}) & \text{if } T_{\mathbf{t}} = \varnothing \\ (T_{\mathbf{n}}, T_{\mathbf{p}}, T_{\mathbf{g}}, T_{\mathbf{t}}, T_{\mathbf{v}}, T_{\mathbf{d}}, T_{\mathbf{w}}, T_{\mathbf{r}}, T_{\mathbf{s}}) & \text{otherwise} \end{cases}$$

Here, we assume all components are interpreted by the RLP as integer values, with the exception of the arbitrary length byte arrays $T_{\mathbf{i}}$ and $T_{\mathbf{d}}$.

$$(14) \quad \begin{array}{llllll} T_{\mathbf{n}} \in \mathbb{P}_{256} & \wedge & T_{\mathbf{v}} \in \mathbb{P}_{256} & \wedge & T_{\mathbf{p}} \in \mathbb{P}_{256} & \wedge \\ T_{\mathbf{g}} \in \mathbb{P}_{256} & \wedge & T_{\mathbf{w}} \in \mathbb{P}_5 & \wedge & T_{\mathbf{r}} \in \mathbb{P}_{256} & \wedge \\ T_{\mathbf{s}} \in \mathbb{P}_{256} & \wedge & T_{\mathbf{d}} \in \mathbb{B} & \wedge & T_{\mathbf{i}} \in \mathbb{B} \end{array}$$

where

$$(15) \qquad \mathbb{P}_{\mathbf{n}} = \{P : P \in \mathbb{P} \wedge P < 2^n\}$$

The address hash $T_{\mathbf{t}}$ is slightly different: it is either a 20-byte address hash or, in the case of being a contract-creation transaction (and thus formally equal to $\varnothing$), it is the RLP empty byte sequence and thus the member of $\mathbb{B}_0$:

$$(16) \qquad T_{\mathbf{t}} \in \begin{cases} \mathbb{B}_{20} & \text{if } T_{\mathbf{t}} \neq \varnothing \\ \mathbb{B}_0 & \text{otherwise} \end{cases}$$

---

[1]Notably, such 'tools' could ultimately become so causally removed from their human-based initiation—or humans may become so causally-neutral—that there could be a point at which they rightly be considered autonomous agents. e.g. contracts may offer bounties to humans for being sent transactions to initiate their execution.

4.3. **The Block.** The block in Ethereum is the collection of relevant pieces of information (known as the block *header*), H, together with information corresponding to the comprised transactions, $\mathbf{T}$, and a set of other block headers $\mathbf{U}$ that are known to have a parent equal to the present block's parent's parent (such blocks are known as *ommers*[2]). The block header contains several pieces of information:

parentHash: The Keccak 256-bit hash of the parent block's header, in its entirety; formally $H_{\mathbf{p}}$.

ommersHash: The Keccak 256-bit hash of the ommers list portion of this block; formally $H_{\mathbf{o}}$.

beneficiary: The 160-bit address to which all fees collected from the successful mining of this block shall be transferred; formally $H_{\mathbf{c}}$.

stateRoot: The Keccak 256-bit hash of the root node of the state trie, after all transactions are executed and finalisations applied; formally $H_{\mathbf{r}}$.

transactionsRoot: The Keccak 256-bit hash of the root node of the trie structure populated with each transaction in the transactions list portion of the block; formally $H_{\mathbf{t}}$.

receiptsRoot: The Keccak 256-bit hash of the root node of the trie structure populated with the receipts of each transaction in the transactions list portion of the block; formally $H_{\mathbf{e}}$.

logsBloom: The Bloom filter composed from indexable information (logger address and log topics) contained in each log entry from the receipt of each transaction in the transactions list; formally $H_{\mathbf{b}}$.

difficulty: A scalar value corresponding to the difficulty level of this block. This can be calculated from the previous block's difficulty level and the timestamp; formally $H_{\mathbf{d}}$.

number: A scalar value equal to the number of ancestor blocks. The genesis block has a number of zero; formally $H_{\mathbf{i}}$.

gasLimit: A scalar value equal to the current limit of gas expenditure per block; formally $H_{\mathbf{l}}$.

gasUsed: A scalar value equal to the total gas used in transactions in this block; formally $H_{\mathbf{g}}$.

timestamp: A scalar value equal to the reasonable output of Unix's time() at this block's inception; formally $H_{\mathbf{s}}$.

extraData: An arbitrary byte array containing data relevant to this block. This must be 32 bytes or fewer; formally $H_{\mathbf{x}}$.

mixHash: A 256-bit hash which proves combined with the nonce that a sufficient amount of computation has been carried out on this block; formally $H_{\mathbf{m}}$.

nonce: A 64-bit hash which proves combined with the mix-hash that a sufficient amount of computation has been carried out on this block; formally $H_{\mathbf{n}}$.

The other two components in the block are simply a list of ommer block headers (of the same format as above), $B_{\mathbf{U}}$, and a series of the transactions, $B_{\mathbf{T}}$. Formally, we

can refer to a block B:

$$(17) \qquad B \equiv (B_H, B_{\mathbf{T}}, B_{\mathbf{U}})$$

4.3.1. *Transaction Receipt.* In order to encode information about a transaction concerning which it may be useful to form a zero-knowledge proof, or index and search, we encode a receipt of each transaction containing certain information from concerning its execution. Each receipt, denoted $B_{\mathbf{R}}[i]$ for the ith transaction, is placed in an index-keyed trie and the root recorded in the header as $H_{\mathbf{e}}$.

The transaction receipt, R, is a tuple of four items comprising the post-transaction state, $R_{\boldsymbol{\sigma}}$, the cumulative gas used in the block containing the transaction receipt as of immediately after the transaction has happened, $R_{\mathbf{u}}$, the set of logs created through execution of the transaction, $R_{\mathbf{l}}$ and the Bloom filter composed from information in those logs, $R_{\mathbf{b}}$:

$$(18) \qquad R \equiv (R_{\boldsymbol{\sigma}}, R_{\mathbf{u}}, R_{\mathbf{b}}, R_{\mathbf{l}})$$

The function L_R trivially prepares a transaction receipt for being transformed into an RLP-serialised byte array:

$$(19) \qquad L_R(R) \equiv (\mathtt{TRIE}(L_S(R_{\boldsymbol{\sigma}})), R_{\mathbf{u}}, R_{\mathbf{b}}, R_{\mathbf{l}})$$

thus the post-transaction state, $R_{\boldsymbol{\sigma}}$ is encoded into a trie structure, the root of which forms the first item.

We assert $R_{\mathbf{u}}$, the cumulative gas used is a positive integer and that the logs Bloom, $R_{\mathbf{b}}$, is a hash of size 2048 bits (256 bytes):

$$(20) \qquad R_{\mathbf{u}} \in \mathbb{P} \quad \wedge \quad R_{\mathbf{b}} \in \mathbb{B}_{256}$$

The log entries, $R_{\mathbf{l}}$, is a series of log entries, termed, for example, $(O_0, O_1, ...)$. A log entry, O, is a tuple of a logger's address, $O_{\mathbf{a}}$, a series of 32-bytes log topics, $O_{\mathbf{t}}$ and some number of bytes of data, $O_{\mathbf{d}}$:

$$(21) \qquad O \equiv (O_{\mathbf{a}}, (O_{\mathbf{t}0}, O_{\mathbf{t}1}, ...), O_{\mathbf{d}})$$

$$(22) \qquad O_{\mathbf{a}} \in \mathbb{B}_{20} \quad \wedge \quad \forall_{t \in O_{\mathbf{t}}} : t \in \mathbb{B}_{32} \quad \wedge \quad O_{\mathbf{d}} \in \mathbb{B}$$

We define the Bloom filter function, M, to reduce a log entry into a single 256-byte hash:

$$(23) \qquad M(O) \equiv \bigvee_{t \in \{O_{\mathbf{a}}\} \cup O_{\mathbf{t}}} \left(M_{3:2048}(t) \right)$$

where $M_{3:2048}$ is a specialised Bloom filter that sets three bits out of 2048, given an arbitrary byte sequence. It does this through taking the low-order 11 bits of each of the first three pairs of bytes in a Keccak-256 hash of the byte sequence.[3] Formally:

$$(24) \quad M_{3:2048}(\mathbf{x} : \mathbf{x} \in \mathbb{B}) \quad \equiv \quad \mathbf{y} : \mathbf{y} \in \mathbb{B}_{256} \quad \text{where:}$$

$$(25) \qquad\qquad\qquad \mathbf{y} \quad = \quad (0, 0, ..., 0) \quad \text{except:}$$

$$(26) \qquad \forall_{i \in \{0,2,4\}} \quad : \quad \mathcal{B}_{m(\mathbf{x},i)}(\mathbf{y}) = 1$$

$$(27) \qquad m(\mathbf{x}, i) \quad \equiv \quad \mathtt{KEC}(\mathbf{x})[i, i+1] \bmod 2048$$

where $\mathcal{B}$ is the bit reference function such that $\mathcal{B}_j(\mathbf{x})$ equals the bit of index j (indexed from 0) in the byte array $\mathbf{x}$.

---

[2]*ommer* is the most prevalent (which is not saying much, as it is not a well-known word) gender-neutral term to mean "sibling of parent"; see http://nonbinary.org/wiki/Gender_neutral_language#Family_Terms

[3]11 bits = 2^{2048}, and the low-order 11 bits is the modulo 2048 of the operand, which is in this case is "each of the first three pairs of bytes in a Keccak-256 hash of the byte sequence".

4.3.2. *Holistic Validity.* We can assert a block's validity if and only if it satisfies several conditions: it must be internally consistent with the ommer and transaction block hashes and the given transactions $B_{\mathbf{T}}$ (as specified in sec 11), when executed in order on the base state $\boldsymbol{\sigma}$ (derived from the final state of the parent block), result in a new state of the identity $H_{\mathbf{r}}$:

(28)
$$
\begin{aligned}
H_{\mathbf{r}} &\equiv \text{TRIE}(L_S(\Pi(\boldsymbol{\sigma}, B))) &&\wedge \\
H_{\mathbf{o}} &\equiv \text{KEC}(\text{RLP}(L_H^*(B_{\mathbf{U}}))) &&\wedge \\
H_{\mathbf{t}} &\equiv \text{TRIE}(\{\forall i < \|B_{\mathbf{T}}\|, i \in \mathbb{P} : p(i, L_{\mathbf{T}}(B_{\mathbf{T}}[i]))\}) &&\wedge \\
H_{\mathbf{e}} &\equiv \text{TRIE}(\{\forall i < \|B_{\mathbf{R}}\|, i \in \mathbb{P} : p(i, L_R(B_{\mathbf{R}}[i]))\}) &&\wedge \\
H_{\mathbf{b}} &\equiv \bigvee_{\mathbf{r} \in B_{\mathbf{R}}} (\mathbf{r_b})
\end{aligned}
$$

where $p(k, v)$ is simply the pairwise RLP transformation, in this case, the first being the index of the transaction in the block and the second being the transaction receipt:

(29) $$ p(k, v) \equiv \big(\text{RLP}(k), \text{RLP}(v)\big) $$

Furthermore:

(30) $$ \text{TRIE}(L_S(\boldsymbol{\sigma})) = P(B_H)_{H_{\mathbf{r}}} $$

Thus $\text{TRIE}(L_S(\boldsymbol{\sigma}))$ is the root node hash of the Merkle Patricia tree structure containing the key-value pairs of the state $\boldsymbol{\sigma}$ with values encoded using RLP, and $P(B_H)$ is the parent block of B, defined directly.

The values stemming from the computation of transactions, specifically the transaction receipts, $B_{\mathbf{R}}$, and that defined through the transaction's stateR-accumulation function, Π, are formalised later in section 11.4.

4.3.3. *Serialisation.* The function L_B and L_H are the preparation functions for a block and block header respectively. Much like the transaction receipt preparation function L_R, we assert the types and order of the structure for when the RLP transformation is required:

(31) $L_H(H) \equiv ($ $H_{\mathbf{p}}, H_{\mathbf{o}}, H_{\mathbf{c}}, H_{\mathbf{r}}, H_{\mathbf{t}}, H_{\mathbf{e}}, H_{\mathbf{b}}, H_{\mathbf{d}},$
$H_{\mathbf{i}}, H_{\mathbf{l}}, H_{\mathbf{g}}, H_{\mathbf{s}}, H_{\mathbf{x}}, H_{\mathbf{m}}, H_{\mathbf{n}})$

(32) $L_B(B) \equiv \big(L_H(B_H), L_{\mathbf{T}}^*(B_{\mathbf{T}}), L_H^*(B_{\mathbf{U}})\big)$

With $L_{\mathbf{T}}^*$ and L_H^* being element-wise sequence transformations, thus:

(33)
$$ f^*\big((x_0, x_1, ...)\big) \equiv \big(f(x_0), f(x_1), ...\big) \quad \text{for any function } f $$

The component types are defined thus:

(34)
$$
\begin{aligned}
&H_{\mathbf{p}} \in \mathbb{B}_{32} &&\wedge& &H_{\mathbf{o}} \in \mathbb{B}_{32} &&\wedge& &H_{\mathbf{c}} \in \mathbb{B}_{20} &&\wedge \\
&H_{\mathbf{r}} \in \mathbb{B}_{32} &&\wedge& &H_{\mathbf{t}} \in \mathbb{B}_{32} &&\wedge& &H_{\mathbf{e}} \in \mathbb{B}_{32} &&\wedge \\
&H_{\mathbf{b}} \in \mathbb{B}_{256} &&\wedge& &H_{\mathbf{d}} \in \mathbb{P} &&\wedge& &H_{\mathbf{i}} \in \mathbb{P} &&\wedge \\
&H_{\mathbf{l}} \in \mathbb{P} &&\wedge& &H_{\mathbf{g}} \in \mathbb{P} &&\wedge& &H_{\mathbf{s}} \in \mathbb{P}_{256} &&\wedge \\
&H_{\mathbf{x}} \in \mathbb{B} &&\wedge& &H_{\mathbf{m}} \in \mathbb{B}_{32} &&\wedge& &H_{\mathbf{n}} \in \mathbb{B}_8
\end{aligned}
$$

where

(35) $$ \mathbb{B}_{\mathbf{n}} = \{B : B \in \mathbb{B} \wedge \|B\| = n\} $$

We now have a rigorous specification for the construction of a formal block structure. The RLP function RLP (see Appendix B) provides the canonical method for transforming this structure into a sequence of bytes ready for transmission over the wire or storage locally.

4.3.4. *Block Header Validity.* We define $P(B_H)$ to be the parent block of B, formally:

(36) $$ P(H) \equiv B' : \text{KEC}(\text{RLP}(B'_H)) = H_{\mathbf{p}} $$

The block number is the parent's block number incremented by one:

(37) $$ H_{\mathbf{i}} \equiv P(H)_{H_{\mathbf{i}}} + 1 $$

The canonical difficulty of a block of header H is defined as $D(H)$:

(38)
$$ D(H) \equiv \begin{cases} D_0 & \text{if } H_{\mathbf{i}} = 0 \\ \max\big(D_0, P(H)_{H_{\mathbf{d}}} + x \times \varsigma_2 + \epsilon\big) & \text{otherwise} \end{cases} $$

where:

(39) $$ D_0 \equiv 131072 $$

(40) $$ x \equiv \left\lfloor \frac{P(H)_{H_{\mathbf{d}}}}{2048} \right\rfloor $$

(41) $$ \varsigma_2 \equiv \max\left(1 - \left\lfloor \frac{H_{\mathbf{s}} - P(H)_{H_{\mathbf{s}}}}{10} \right\rfloor, -99\right) $$

(42) $$ \epsilon \equiv \left\lfloor 2^{\lfloor H_{\mathbf{i}} \div 100000 \rfloor - 2} \right\rfloor $$

The canonical gas limit $H_{\mathbf{l}}$ of a block of header H must fulfil the relation:

(43) $$ H_{\mathbf{l}} < P(H)_{H_{\mathbf{l}}} + \left\lfloor \frac{P(H)_{H_{\mathbf{l}}}}{1024} \right\rfloor \quad \wedge $$

(44) $$ H_{\mathbf{l}} > P(H)_{H_{\mathbf{l}}} - \left\lfloor \frac{P(H)_{H_{\mathbf{l}}}}{1024} \right\rfloor \quad \wedge $$

(45) $$ H_{\mathbf{l}} \geqslant 125000 $$

$H_{\mathbf{s}}$ is the timestamp of block H and must fulfil the relation:

(46) $$ H_{\mathbf{s}} > P(H)_{H_{\mathbf{s}}} $$

This mechanism enforces a homeostasis in terms of the time between blocks; a smaller period between the last two blocks results in an increase in the difficulty level and thus additional computation required, lengthening the likely next period. Conversely, if the period is too large, the difficulty, and expected time to the next block, is reduced.

The nonce of a block, $H_{\mathbf{n}}$, must satisfy the relations:

(47) $$ n \leqslant \frac{2^{256}}{H_{\mathbf{d}}} \quad \wedge \quad m = H_{\mathbf{m}} $$

with $(n, m) = \text{PoW}(H_{\not{\mathbf{n}}}, H_{\mathbf{n}}, \mathbf{d})$.

Where $H_{\not{\mathbf{n}}}$ is the new block's header H, but *without* the nonce and mix-hash block components, $\mathbf{d}$ being the current DAG, a large data set needed to compute the mix-hash, and PoW is the proof-of-work function (see section 11.5): this evaluates to an array with the first item being the mix-hash, to prove that a correct DAG has been used, and the second item being a pseudo-random number cryptographically dependent on H and $\mathbf{d}$. Given an approximately uniform distribution in the range $[0, 2^{64})$, the expected time to find a solution is proportional to the difficulty, $H_{\mathbf{d}}$.

This is the foundation of the security of the blockchain and is the fundamental reason why a malicious node cannot propagate newly created blocks that would otherwise overwrite ("rewrite") history. Because the block nonce must satisfy this requirement, and because its satisfaction

depends on the contents of the block and in turn its composed transactions, creating new, valid, blocks is difficult and, over time, requires approximately the total compute power of the trustworthy portion of the mining peers.

Thus we are able to define the block header validity function $V(H)$:

$$(48) \quad V(H) \quad \equiv \quad n \leqslant \frac{2^{256}}{H_{\mathbf{d}}} \wedge m = H_{\mathbf{m}} \quad \wedge$$

$$(49) \qquad\qquad H_{\mathbf{d}} = D(H) \quad \wedge$$

$$(50) \qquad\qquad H_{\mathbf{g}} \leq H_{\mathbf{l}} \quad \wedge$$

$$(51) \qquad\qquad H_{\mathbf{l}} < P(H)_{H_{\mathbf{l}}} + \left\lfloor \frac{P(H)_{H_{\mathbf{l}}}}{1024} \right\rfloor \quad \wedge$$

$$(52) \qquad\qquad H_{\mathbf{l}} > P(H)_{H_{\mathbf{l}}} - \left\lfloor \frac{P(H)_{H_{\mathbf{l}}}}{1024} \right\rfloor \quad \wedge$$

$$(53) \qquad\qquad H_{\mathbf{l}} \geqslant 125000 \quad \wedge$$

$$(54) \qquad\qquad H_{\mathbf{s}} > P(H)_{H_{\mathbf{s}}} \quad \wedge$$

$$(55) \qquad\qquad H_{\mathbf{i}} = P(H)_{H_{\mathbf{i}}} + 1 \quad \wedge$$

$$(56) \qquad\qquad \|H_{\mathbf{x}}\| \leq 32$$

where $(n, m) = \mathtt{PoW}(H_{\cancel{n}}, H_{\mathbf{n}}, \mathbf{d})$

Noting additionally that **extraData** must be at most 32 bytes.

5. Gas and Payment

In order to avoid issues of network abuse and to sidestep the inevitable questions stemming from Turing completeness, all programmable computation in Ethereum is subject to fees. The fee schedule is specified in units of *gas* (see Appendix G for the fees associated with various computation). Thus any given fragment of programmable computation (this includes creating contracts, making message calls, utilising and accessing account storage and executing operations on the virtual machine) has a universally agreed cost in terms of gas.

Every transaction has a specific amount of gas associated with it: **gasLimit**. This is the amount of gas which is implicitly purchased from the sender's account balance. The purchase happens at the according **gasPrice**, also specified in the transaction. The transaction is considered invalid if the account balance cannot support such a purchase. It is named **gasLimit** since any unused gas at the end of the transaction is refunded (at the same rate of purchase) to the sender's account. Gas does not exist outside of the execution of a transaction. Thus for accounts with trusted code associated, a relatively high gas limit may be set and left alone.

In general, Ether used to purchase gas that is not refunded is delivered to the *beneficiary* address, the address of an account typically under the control of the miner. Transactors are free to specify any **gasPrice** that they wish, however miners are free to ignore transactions as they choose. A higher gas price on a transaction will therefore cost the sender more in terms of Ether and deliver a greater value to the miner and thus will more likely be selected for inclusion by more miners. Miners, in general, will choose to advertise the minimum gas price for which they will execute transactions and transactors will be free to canvas these prices in determining what gas price to offer. Since there will be a (weighted) distribution of minimum acceptable gas prices, transactors will necessarily have a trade-off to make between lowering the gas price

and maximising the chance that their transaction will be mined in a timely manner.

6. Transaction Execution

The execution of a transaction is the most complex part of the Ethereum protocol: it defines the state transition function Υ. It is assumed that any transactions executed first pass the initial tests of intrinsic validity. These include:

(1) The transaction is well-formed RLP, with no additional trailing bytes;
(2) the transaction signature is valid;
(3) the transaction nonce is valid (equivalent to the sender account's current nonce);
(4) the gas limit is no smaller than the intrinsic gas, g_0, used by the transaction;
(5) the sender account balance contains at least the cost, v_0, required in up-front payment.

Formally, we consider the function Υ, with T being a transaction and $\boldsymbol{\sigma}$ the state:

$$(57) \qquad\qquad \boldsymbol{\sigma}' = \Upsilon(\boldsymbol{\sigma}, T)$$

Thus $\boldsymbol{\sigma}'$ is the post-transactional state. We also define Υ^g to evaluate to the amount of gas used in the execution of a transaction and $\Upsilon^{\mathbf{l}}$ to evaluate to the transaction's accrued log items, both to be formally defined later.

6.1. **Substate.** Throughout transaction execution, we accrue certain information that is acted upon immediately following the transaction. We call this *transaction substate*, and represent it as A, which is a tuple:

$$(58) \qquad\qquad A \equiv (A_{\mathbf{s}}, A_{\mathbf{l}}, A_{\mathbf{r}})$$

The tuple contents include $A_{\mathbf{s}}$, the self-destruct set: a set of accounts that will be discarded following the transaction's completion. $A_{\mathbf{l}}$ is the log series: this is a series of archived and indexable 'checkpoints' in VM code execution that allow for contract-calls to be easily tracked by onlookers external to the Ethereum world (such as decentralised application front-ends). Finally there is $A_{\mathbf{r}}$, the refund balance, increased through using the SSTORE instruction in order to reset contract storage to zero from some non-zero value. Though not immediately refunded, it is allowed to partially offset the total execution costs.

For brevity, we define the empty substate A^0 to have no self-destructs, no logs and a zero refund balance:

$$(59) \qquad\qquad A^0 \equiv (\varnothing, (), 0)$$

6.2. **Execution.** We define intrinsic gas g_0, the amount of gas this transaction requires to be paid prior to execution, as follows:

$$(60) \qquad g_0 \equiv \sum_{i \in T_{\mathbf{i}}, T_{\mathbf{d}}} \begin{cases} G_{\text{txdatazero}} & \text{if} \quad i = 0 \\ G_{\text{txdatanonzero}} & \text{otherwise} \end{cases}$$

$$(61) \qquad\qquad + \begin{cases} G_{\text{txcreate}} & \text{if} \quad T_{\mathbf{t}} = \varnothing \\ 0 & \text{otherwise} \end{cases}$$

$$(62) \qquad\qquad + G_{\text{transaction}}$$

where $T_{\mathbf{i}}, T_{\mathbf{d}}$ means the series of bytes of the transaction's associated data and initialisation EVM-code,

depending on whether the transaction is for contract-creation or message-call. G_{txcreate} is added if the transaction is contract-creating, but not if a result of EVM-code. G is fully defined in Appendix G.

The up-front cost v_0 is calculated as:

$$(63) \qquad v_0 \equiv T_{\mathbf{g}} T_{\mathbf{p}} + T_{\mathbf{v}}$$

The validity is determined as:

$$(64) \qquad \begin{aligned} S(T) &\neq \varnothing \quad \wedge \\ \boldsymbol{\sigma}[S(T)] &\neq \varnothing \quad \wedge \\ T_{\mathbf{n}} &= \boldsymbol{\sigma}[S(T)]_{\mathbf{n}} \quad \wedge \\ g_0 &\leqslant T_{\mathbf{g}} \quad \wedge \\ v_0 &\leqslant \boldsymbol{\sigma}[S(T)]_{\mathbf{b}} \quad \wedge \\ T_{\mathbf{g}} &\leqslant B_{H\mathbf{l}} - \ell(B_{\mathbf{R}})_{\mathbf{u}} \end{aligned}$$

Note the final condition; the sum of the transaction's gas limit, $T_{\mathbf{g}}$, and the gas utilised in this block prior, given by $\ell(B_{\mathbf{R}})_{\mathbf{u}}$, must be no greater than the block's **gasLimit**, $B_{H\mathbf{l}}$.

The execution of a valid transaction begins with an irrevocable change made to the state: the nonce of the account of the sender, $S(T)$, is incremented by one and the balance is reduced by part of the up-front cost, $T_{\mathbf{g}} T_{\mathbf{p}}$. The gas available for the proceeding computation, g, is defined as $T_{\mathbf{g}} - g_0$. The computation, whether contract creation or a message call, results in an eventual state (which may legally be equivalent to the current state), the change to which is deterministic and never invalid: there can be no invalid transactions from this point.

We define the checkpoint state $\boldsymbol{\sigma}_0$:

$$(65) \qquad \boldsymbol{\sigma}_0 \equiv \boldsymbol{\sigma} \quad \text{except:}$$

$$(66) \qquad \boldsymbol{\sigma}_0[S(T)]_b \equiv \boldsymbol{\sigma}[S(T)]_b - T_g T_p$$

$$(67) \qquad \boldsymbol{\sigma}_0[S(T)]_{\mathbf{n}} \equiv \boldsymbol{\sigma}[S(T)]_{\mathbf{n}} + 1$$

Evaluating $\boldsymbol{\sigma}_P$ from $\boldsymbol{\sigma}_0$ depends on the transaction type; either contract creation or message call; we define the tuple of post-execution provisional state $\boldsymbol{\sigma}_P$, remaining gas g' and substate A:

$$(68)$$

$$(\boldsymbol{\sigma}_P, g', A) \equiv \begin{cases} \Lambda(\boldsymbol{\sigma}_0, S(T), T_{\mathbf{o}}, \\ \qquad g, T_{\mathbf{p}}, T_{\mathbf{v}}, T_{\mathbf{i}}, 0) & \text{if} \quad T_{\mathbf{t}} = \varnothing \\ \Theta_3(\boldsymbol{\sigma}_0, S(T), T_{\mathbf{o}}, \\ \qquad T_{\mathbf{t}}, T_{\mathbf{t}}, g, T_{\mathbf{p}}, T_{\mathbf{v}}, T_{\mathbf{v}}, T_{\mathbf{d}}, 0) & \text{otherwise} \end{cases}$$

where g is the amount of gas remaining after deducting the basic amount required to pay for the existence of the transaction:

$$(69) \qquad g \equiv T_{\mathbf{g}} - g_0$$

and $T_{\mathbf{o}}$ is the original transactor, which can differ from the sender in the case of a message call or contract creation not directly triggered by a transaction but coming from the execution of EVM-code.

Note we use Θ_3 to denote the fact that only the first three components of the function's value are taken; the final represents the message-call's output value (a byte array) and is unused in the context of transaction evaluation.

After the message call or contract creation is processed, the state is finalised by determining the amount to be refunded, g^* from the remaining gas, g', plus some allowance from the refund counter, to the sender at the original rate.

$$(70) \qquad g^* \equiv g' + \min\left\{ \left\lfloor \frac{T_{\mathbf{g}} - g'}{2} \right\rfloor, A_{\mathbf{r}} \right\}$$

The total refundable amount is the legitimately remaining gas g', added to $A_{\mathbf{r}}$, with the latter component being capped up to a maximum of half (rounded down) of the total amount used $T_{\mathbf{g}} - g'$.

The Ether for the gas is given to the miner, whose address is specified as the beneficiary of the present block B. So we define the pre-final state $\boldsymbol{\sigma}^*$ in terms of the provisional state $\boldsymbol{\sigma}_P$:

$$(71) \qquad \boldsymbol{\sigma}^* \equiv \boldsymbol{\sigma}_P \quad \text{except}$$

$$(72) \qquad \boldsymbol{\sigma}^*[S(T)]_{\mathbf{b}} \equiv \boldsymbol{\sigma}_P[S(T)]_{\mathbf{b}} + g^* T_{\mathbf{p}}$$

$$(73) \qquad \boldsymbol{\sigma}^*[m]_{\mathbf{b}} \equiv \boldsymbol{\sigma}_P[m]_{\mathbf{b}} + (T_{\mathbf{g}} - g^*) T_{\mathbf{p}}$$

$$(74) \qquad m \equiv B_{H\mathbf{c}}$$

The final state, $\boldsymbol{\sigma}'$, is reached after deleting all accounts that appear in the self-destruct set:

$$(75) \qquad \boldsymbol{\sigma}' \equiv \boldsymbol{\sigma}^* \quad \text{except}$$

$$(76) \qquad \forall i \in A_{\mathbf{s}} : \boldsymbol{\sigma}'[i] \equiv \varnothing$$

And finally, we specify Υ^g, the total gas used in this transaction and $\Upsilon^{\mathbf{l}}$, the logs created by this transaction:

$$(77) \qquad \Upsilon^g(\boldsymbol{\sigma}, T) \equiv T_{\mathbf{g}} - g'$$

$$(78) \qquad \Upsilon^{\mathbf{l}}(\boldsymbol{\sigma}, T) \equiv A_{\mathbf{l}}$$

These are used to help define the transaction receipt.

7. Contract Creation

There are a number of intrinsic parameters used when creating an account: sender (s), original transactor (o), available gas (g), gas price (p), endowment (v) together with an arbitrary length byte array, $\mathbf{i}$, the initialisation EVM code and finally the present depth of the message-call/contract-creation stack (e).

We define the creation function formally as the function Λ, which evaluates from these values, together with the state $\boldsymbol{\sigma}$ to the tuple containing the new state, remaining gas and accrued transaction substate $(\boldsymbol{\sigma}', g', A)$, as in section 6:

$$(79) \qquad (\boldsymbol{\sigma}', g', A) \equiv \Lambda(\boldsymbol{\sigma}, s, o, g, p, v, \mathbf{i}, e)$$

The address of the new account is defined as being the rightmost 160 bits of the Keccak hash of the RLP encoding of the structure containing only the sender and the nonce. Thus we define the resultant address for the new account a:

$$(80) \qquad a \equiv \mathcal{B}_{96..255}\Big(\text{KEC}\Big(\text{RLP}\big((s, \boldsymbol{\sigma}[s]_{\mathbf{n}} - 1) \big) \Big) \Big)$$

where KEC is the Keccak 256-bit hash function, RLP is the RLP encoding function, $\mathcal{B}_{a..b}(X)$ evaluates to binary value containing the bits of indices in the range $[a, b]$ of the binary data X and $\boldsymbol{\sigma}[x]$ is the address state of x or $\varnothing$ if none exists. Note we use one fewer than the sender's nonce value; we assert that we have incremented the sender account's nonce prior to this call, and so the value used is the sender's nonce at the beginning of the responsible transaction or VM operation.

The account's nonce is initially defined as zero, the balance as the value passed, the storage as empty and the code hash as the Keccak 256-bit hash of the empty string; the sender's balance is also reduced by the value passed. Thus the mutated state becomes $\boldsymbol{\sigma}^*$:

$$(81) \qquad \boldsymbol{\sigma}^* \equiv \boldsymbol{\sigma} \quad \text{except:}$$

$$(82) \qquad \boldsymbol{\sigma}^*[a] \quad \equiv \quad \big(0, v + v', \mathtt{TRIE}(\varnothing), \mathtt{KEC}(())\big)$$

$$(83) \qquad \boldsymbol{\sigma}^*[s]_{\mathbf{b}} \quad \equiv \quad \boldsymbol{\sigma}[s]_{\mathbf{b}} - v$$

where v' is the account's pre-existing value, in the event it was previously in existence:

$$(84) \qquad v' \equiv \begin{cases} 0 & \text{if} \quad \boldsymbol{\sigma}[a] = \varnothing \\ \boldsymbol{\sigma}[a]_{\mathbf{b}} & \text{otherwise} \end{cases}$$

Finally, the account is initialised through the execution of the initialising EVM code $\mathbf{i}$ according to the execution model (see section 9). Code execution can effect several events that are not internal to the execution state: the account's storage can be altered, further accounts can be created and further message calls can be made. As such, the code execution function Ξ evaluates to a tuple of the resultant state $\boldsymbol{\sigma}^{**}$, available gas remaining g^{**}, the accrued substate A and the body code of the account $\mathbf{o}$.

$$(85) \qquad (\boldsymbol{\sigma}^{**}, g^{**}, A, \mathbf{o}) \equiv \Xi(\boldsymbol{\sigma}^*, g, I)$$

where I contains the parameters of the execution environment as defined in section 9, that is:

$$(86) \qquad I_{\mathbf{a}} \quad \equiv \quad a$$
$$(87) \qquad I_{\mathbf{o}} \quad \equiv \quad o$$
$$(88) \qquad I_{\mathbf{p}} \quad \equiv \quad p$$
$$(89) \qquad I_{\mathbf{d}} \quad \equiv \quad ()$$
$$(90) \qquad I_{\mathbf{s}} \quad \equiv \quad s$$
$$(91) \qquad I_{\mathbf{v}} \quad \equiv \quad v$$
$$(92) \qquad I_{\mathbf{b}} \quad \equiv \quad \mathbf{i}$$
$$(93) \qquad I_{\mathbf{e}} \quad \equiv \quad e$$

$I_{\mathbf{d}}$ evaluates to the empty tuple as there is no input data to this call. I_H has no special treatment and is determined from the blockchain.

Code execution depletes gas, and gas may not go below zero, thus execution may exit before the code has come to a natural halting state. In this (and several other) exceptional cases we say an out-of-gas (OOG) exception has occurred: The evaluated state is defined as being the empty set, $\varnothing$, and the entire create operation should have no effect on the state, effectively leaving it as it was immediately prior to attempting the creation.

If the initialization code completes successfully, a final contract-creation cost is paid, the code-deposit cost, c, proportional to the size of the created contract's code:

$$(94) \qquad c \equiv G_{\text{codedeposit}} \times |\mathbf{o}|$$

If there is not enough gas remaining to pay this, i.e. $g^{**} < c$, then we also declare an out-of-gas exception.

The gas remaining will be zero in any such exceptional condition, i.e. if the creation was conducted as the reception of a transaction, then this doesn't affect payment of the intrinsic cost of contract creation; it is paid regardless. However, the value of the transaction is not transferred to the aborted contract's address when we are out-of-gas.

If such an exception does not occur, then the remaining gas is refunded to the originator and the now-altered state is allowed to persist. Thus formally, we may specify the resultant state, gas and substate as $(\boldsymbol{\sigma}', g', A)$ where:

$$(95) \qquad g' \equiv \begin{cases} 0 & \text{if} \quad F \\ g^{**} - \mathbf{c} & \text{otherwise} \end{cases}$$

$$(96) \qquad \boldsymbol{\sigma}' \equiv \begin{cases} \boldsymbol{\sigma} & \text{if} \quad F \\ \boldsymbol{\sigma}^{**} & \text{except:} \\ \quad \boldsymbol{\sigma}'[a]_{\mathbf{c}} = \mathtt{KEC}(\mathbf{o}) & \text{otherwise} \end{cases}$$

where

$$(97) \qquad F \equiv \big(\boldsymbol{\sigma}^{**} = \varnothing \ \vee \ \mathbf{g}^{**} < \mathbf{c} \ \vee \ |\mathbf{o}| > 24576\big)$$

The exception in the determination of $\boldsymbol{\sigma}'$ dictates that $\mathbf{o}$, the resultant byte sequence from the execution of the initialisation code, specifies the final body code for the newly-created account. Note that the 24576 byte limit for $\mathbf{o}$ exists because a contract creation call can trigger O(n) cost in terms of reading the code from disk, preprocessing the code for VM execution, and also adding O(n) data to the Merkle proof for the block's proof-of-validity. With higher gas limits that can be caused by dynamic gas limit rules, this is a greater concern, and is especially inconvenient with light clients verifying proofs of validity or invalidity.

Note that the intention is that the result is either a successfully created new contract with its endowment, or no new contract with no transfer of value.

7.1. **Subtleties.** Note that while the initialisation code is executing, the newly created address exists but with no intrinsic body code. Thus any message call received by it during this time causes no code to be executed. If the initialisation execution ends with a SELFDESTRUCT instruction, the matter is moot since the account will be deleted before the transaction is completed. For a normal STOP code, or if the code returned is otherwise empty, then the state is left with a zombie account, and any remaining balance will be locked into the account forever.

8. Message Call

In the case of executing a message call, several parameters are required: sender (s), transaction originator (o), recipient (r), the account whose code is to be executed (c, usually the same as recipient), available gas (g), value (v) and gas price (p) together with an arbitrary length byte array, $\mathbf{d}$, the input data of the call and finally the present depth of the message-call/contract-creation stack (e).

Aside from evaluating to a new state and transaction substate, message calls also have an extra component—the output data denoted by the byte array $\mathbf{o}$. This is ignored when executing transactions, however message calls can be initiated due to VM-code execution and in this case this information is used.

$$(98) \qquad (\boldsymbol{\sigma}', g', A, \mathbf{o}) \equiv \Theta(\boldsymbol{\sigma}, s, o, r, c, g, p, v, \tilde{v}, \mathbf{d}, e)$$

Note that we need to differentiate between the value that is to be transferred, v, from the value apparent in the execution context, $\tilde{v}$, for the DELEGATECALL instruction.

We define $\boldsymbol{\sigma}_1$, the first transitional state as the original state but with the value transferred from sender to recipient:

$$(99) \qquad \boldsymbol{\sigma}_1[r]_{\mathbf{b}} \equiv \boldsymbol{\sigma}[r]_{\mathbf{b}} + v \quad \wedge \quad \boldsymbol{\sigma}_1[s]_{\mathbf{b}} \equiv \boldsymbol{\sigma}[s]_{\mathbf{b}} - v$$

unless $s = r$.

Throughout the present work, it is assumed that if $\boldsymbol{\sigma}_1[r]$ was originally undefined, it will be created as an account with no code or state and zero balance and nonce. Thus the previous equation should be taken to mean:

$$(100) \qquad \boldsymbol{\sigma}_1 \equiv \boldsymbol{\sigma}_1' \quad \text{except:}$$

$$(101) \qquad \boldsymbol{\sigma}_1[s]_\mathbf{b} \equiv \boldsymbol{\sigma}_1'[s]_\mathbf{b} - v$$

$$(102) \qquad \text{and} \quad \boldsymbol{\sigma}_1' \equiv \boldsymbol{\sigma} \quad \text{except:}$$

$$(103) \quad \begin{cases} \boldsymbol{\sigma}_1'[r] \equiv (v, 0, \texttt{KEC}(()), \texttt{TRIE}(\varnothing)) & \text{if} \quad \boldsymbol{\sigma}[r] = \varnothing \\ \boldsymbol{\sigma}_1'[r]_\mathbf{b} \equiv \boldsymbol{\sigma}[r]_\mathbf{b} + v & \text{otherwise} \end{cases}$$

The account's associated code (identified as the fragment whose Keccak hash is $\boldsymbol{\sigma}[c]_\mathbf{c}$) is executed according to the execution model (see section 9). Just as with contract creation, if the execution halts in an exceptional fashion (i.e. due to an exhausted gas supply, stack underflow, invalid jump destination or invalid instruction), then no gas is refunded to the caller and the state is reverted to the point immediately prior to balance transfer (i.e. $\boldsymbol{\sigma}$).

$$(104) \qquad \boldsymbol{\sigma}' \equiv \begin{cases} \boldsymbol{\sigma} & \text{if} \quad \boldsymbol{\sigma}^{**} = \varnothing \\ \boldsymbol{\sigma}^{**} & \text{otherwise} \end{cases}$$

$$(105) \qquad g' \equiv \begin{cases} 0 & \text{if} \quad \boldsymbol{\sigma}^{**} = \varnothing \\ g^{**} & \text{otherwise} \end{cases}$$

$$(106) \; (\boldsymbol{\sigma}^{**}, g^{**}, A, \mathbf{o}) \equiv \begin{cases} \Xi_{\texttt{ECREC}}(\boldsymbol{\sigma}_1, g, I) & \text{if} \quad r = 1 \\ \Xi_{\texttt{SHA256}}(\boldsymbol{\sigma}_1, g, I) & \text{if} \quad r = 2 \\ \Xi_{\texttt{RIP160}}(\boldsymbol{\sigma}_1, g, I) & \text{if} \quad r = 3 \\ \Xi_{\texttt{ID}}(\boldsymbol{\sigma}_1, g, I) & \text{if} \quad r = 4 \\ \Xi(\boldsymbol{\sigma}_1, g, I) & \text{otherwise} \end{cases}$$

$$(107) \qquad I_\mathbf{a} \equiv r$$
$$(108) \qquad I_\mathbf{o} \equiv o$$
$$(109) \qquad I_\mathbf{p} \equiv p$$
$$(110) \qquad I_\mathbf{d} \equiv \mathbf{d}$$
$$(111) \qquad I_\mathbf{s} \equiv s$$
$$(112) \qquad I_\mathbf{v} \equiv \tilde{v}$$
$$(113) \qquad I_\mathbf{e} \equiv e$$
$$(114) \quad \text{Let } \texttt{KEC}(I_\mathbf{b}) = \boldsymbol{\sigma}[c]_\mathbf{c}$$

It is assumed that the client will have stored the pair $(\texttt{KEC}(I_\mathbf{b}), I_\mathbf{b})$ at some point prior in order to make the determination of $I_\mathbf{b}$ feasible.

As can be seen, there are four exceptions to the usage of the general execution framework Ξ for evaluation of the message call: these are four so-called 'precompiled' contracts, meant as a preliminary piece of architecture that may later become *native extensions*. The four contracts in addresses 1, 2, 3 and 4 execute the elliptic curve public key recovery function, the SHA2 256-bit hash scheme, the RIPEMD 160-bit hash scheme and the identity function respectively.

Their full formal definition is in Appendix E.

9. Execution Model

The execution model specifies how the system state is altered given a series of bytecode instructions and a small tuple of environmental data. This is specified through a formal model of a virtual state machine, known as the Ethereum Virtual Machine (EVM). It is a *quasi*-Turing-complete machine; the *quasi* qualification comes from the fact that the computation is intrinsically bounded through a parameter, *gas*, which limits the total amount of computation done.

9.1. **Basics.** The EVM is a simple stack-based architecture. The word size of the machine (and thus size of stack item) is 256-bit. This was chosen to facilitate the Keccak-256 hash scheme and elliptic-curve computations. The memory model is a simple word-addressed byte array. The stack has a maximum size of 1024. The machine also has an independent storage model; this is similar in concept to the memory but rather than a byte array, it is a word-addressable word array. Unlike memory, which is volatile, storage is non volatile and is maintained as part of the system state. All locations in both storage and memory are well-defined initially as zero.

The machine does not follow the standard von Neumann architecture. Rather than storing program code in generally-accessible memory or storage, it is stored separately in a virtual ROM interactable only through a specialised instruction.

The machine can have exceptional execution for several reasons, including stack underflows and invalid instructions. Like the out-of-gas exception, they do not leave state changes intact. Rather, the machine halts immediately and reports the issue to the execution agent (either the transaction processor or, recursively, the spawning execution environment) which will deal with it separately.

9.2. **Fees Overview.** Fees (denominated in gas) are charged under three distinct circumstances, all three as prerequisite to the execution of an operation. The first and most common is the fee intrinsic to the computation of the operation (see Appendix G). Secondly, gas may be deducted in order to form the payment for a subordinate message call or contract creation; this forms part of the payment for CREATE, CALL and CALLCODE. Finally, gas may be paid due to an increase in the usage of the memory.

Over an account's execution, the total fee for memory-usage payable is proportional to the smallest multiple of 32 bytes that is required such that all memory indices (whether for read or write) are included in the range. This is paid for on a just-in-time basis; as such, referencing an area of memory at least 32 bytes greater than any previously indexed memory will certainly result in an additional memory usage fee. Due to this fee it is highly unlikely addresses will ever go above 32-bit bounds. That said, implementations must be able to manage this eventuality.

Storage fees have a slightly nuanced behaviour—to incentivise minimisation of the use of storage (which corresponds directly to a larger state database on all nodes), the execution fee for an operation that clears an entry in the storage is not only waived, a qualified refund is given; in fact, this refund is effectively paid up-front since the initial usage of a storage location costs substantially more than normal usage.

See Appendix H for a rigorous definition of the EVM gas cost.

9.3. **Execution Environment.** In addition to the system state $\boldsymbol{\sigma}$, and the remaining gas for computation g, there are several pieces of important information used in the execution environment that the execution agent must provide; these are contained in the tuple I:

- $I_{\mathbf{a}}$, the address of the account which owns the code that is executing.
- $I_{\mathbf{o}}$, the sender address of the transaction that originated this execution.
- $I_{\mathbf{p}}$, the price of gas in the transaction that originated this execution.
- $I_{\mathbf{d}}$, the byte array that is the input data to this execution; if the execution agent is a transaction, this would be the transaction data.
- $I_{\mathbf{s}}$, the address of the account which caused the code to be executing; if the execution agent is a transaction, this would be the transaction sender.
- $I_{\mathbf{v}}$, the value, in Wei, passed to this account as part of the same procedure as execution; if the execution agent is a transaction, this would be the transaction value.
- $I_{\mathbf{b}}$, the byte array that is the machine code to be executed.
- I_H, the block header of the present block.
- $I_{\mathbf{e}}$, the depth of the present message-call or contract-creation (i.e. the number of CALLs or CREATEs being executed at present).

The execution model defines the function Ξ, which can compute the resultant state $\boldsymbol{\sigma}'$, the remaining gas g', the accrued substate A and the resultant output, $\mathbf{o}$, given these definitions. For the present context, we will defined it as:

$$(115) \qquad (\boldsymbol{\sigma}', g', A, \mathbf{o}) \equiv \Xi(\boldsymbol{\sigma}, g, I)$$

where we will remember that A, the accrued substate is defined as the tuple of the suicides set $\mathbf{s}$, the log series $\mathbf{l}$ and the refunds r:

$$(116) \qquad A \equiv (\mathbf{s}, \mathbf{l}, r)$$

9.4. **Execution Overview.** We must now define the Ξ function. In most practical implementations this will be modelled as an iterative progression of the pair comprising the full system state, $\boldsymbol{\sigma}$ and the machine state, $\boldsymbol{\mu}$. Formally, we define it recursively with a function X. This uses an iterator function O (which defines the result of a single cycle of the state machine) together with functions Z which determines if the present state is an exceptional halting state of the machine and H, specifying the output data of the instruction if and only if the present state is a normal halting state of the machine.

The empty sequence, denoted $()$, is not equal to the empty set, denoted $\varnothing$; this is important when interpreting the output of H, which evaluates to $\varnothing$ when execution is to continue but a series (potentially empty) when execution

should halt.

$$(117) \qquad \Xi(\boldsymbol{\sigma}, g, I) \quad \equiv \quad (\boldsymbol{\sigma}', \boldsymbol{\mu}'_{\mathbf{g}}, A, \mathbf{o})$$

$$(118) \qquad (\boldsymbol{\sigma}, \boldsymbol{\mu}', A, ..., \mathbf{o}) \quad \equiv \quad X\big((\boldsymbol{\sigma}, \boldsymbol{\mu}, A^0, I)\big)$$

$$(119) \qquad \boldsymbol{\mu}_{\mathbf{g}} \quad \equiv \quad \mathbf{g}$$

$$(120) \qquad \boldsymbol{\mu}_{pc} \quad \equiv \quad 0$$

$$(121) \qquad \boldsymbol{\mu}_{\mathbf{m}} \quad \equiv \quad (0, 0, ...)$$

$$(122) \qquad \boldsymbol{\mu}_{\mathbf{i}} \quad \equiv \quad 0$$

$$(123) \qquad \boldsymbol{\mu}_{\mathbf{s}} \quad \equiv \quad ()$$

$$(124)$$

$$X\big((\boldsymbol{\sigma}, \boldsymbol{\mu}, A, I)\big) \equiv \begin{cases} \big(\varnothing, \boldsymbol{\mu}, A^0, I, ()\big) & \text{if} \quad Z(\boldsymbol{\sigma}, \boldsymbol{\mu}, I) \\ O(\boldsymbol{\sigma}, \boldsymbol{\mu}, A, I) \cdot \mathbf{o} & \text{if} \quad \mathbf{o} \neq \varnothing \\ X\big(O(\boldsymbol{\sigma}, \boldsymbol{\mu}, A, I)\big) & \text{otherwise} \end{cases}$$

where

$$(125) \qquad \mathbf{o} \quad \equiv \quad H(\boldsymbol{\mu}, I)$$

$$(126) \qquad (a, b, c, d) \cdot e \quad \equiv \quad (a, b, c, d, e)$$

Note that, when we evaluate Ξ, we drop the fourth element I' and extract the remaining gas $\boldsymbol{\mu}'_{\mathbf{g}}$ from the resultant machine state $\boldsymbol{\mu}'$.

X is thus cycled (recursively here, but implementations are generally expected to use a simple iterative loop) until either Z becomes true indicating that the present state is exceptional and that the machine must be halted and any changes discarded or until H becomes a series (rather than the empty set) indicating that the machine has reached a controlled halt.

9.4.1. *Machine State.* The machine state $\boldsymbol{\mu}$ is defined as the tuple $(g, pc, \mathbf{m}, i, \mathbf{s})$ which are the gas available, the program counter $pc \in \mathbb{P}_{256}$, the memory contents, the active number of words in memory (counting continuously from position 0), and the stack contents. The memory contents $\boldsymbol{\mu}_{\mathbf{m}}$ are a series of zeroes of size 2^{256}.

For the ease of reading, the instruction mnemonics, written in small-caps (e.g. ADD), should be interpreted as their numeric equivalents; the full table of instructions and their specifics is given in Appendix H.

For the purposes of defining Z, H and O, we define w as the current operation to be executed:

$$(127) \qquad w \equiv \begin{cases} I_{\mathbf{b}}[\boldsymbol{\mu}_{pc}] & \text{if} \quad \boldsymbol{\mu}_{pc} < \|I_{\mathbf{b}}\| \\ \text{STOP} & \text{otherwise} \end{cases}$$

We also assume the fixed amounts of δ and α, specifying the stack items removed and added, both subscriptable on the instruction and an instruction cost function C evaluating to the full cost, in gas, of executing the given instruction.

9.4.2. *Exceptional Halting.* The exceptional halting function Z is defined as:

$$(128) \qquad Z(\boldsymbol{\sigma}, \boldsymbol{\mu}, I) \equiv \begin{aligned} &\boldsymbol{\mu}_{\mathbf{g}} < C(\boldsymbol{\sigma}, \boldsymbol{\mu}, I) \quad \vee \\ &\delta_w = \varnothing \quad \vee \\ &\|\boldsymbol{\mu}_{\mathbf{s}}\| < \delta_w \quad \vee \\ &(w \in \{\text{JUMP}, \text{JUMPI}\} \quad \wedge \\ &\quad \boldsymbol{\mu}_{\mathbf{s}}[0] \notin D(I_{\mathbf{b}})) \quad \vee \\ &\|\boldsymbol{\mu}_{\mathbf{s}}\| - \delta_w + \alpha_w > 1024 \end{aligned}$$

This states that the execution is in an exceptional halting state if there is insufficient gas, if the instruction is invalid (and therefore its δ subscript is undefined), if there are insufficient stack items, if a JUMP/JUMPI destination

is invalid or the new stack size would be larger then 1024. The astute reader will realise that this implies that no instruction can, through its execution, cause an exceptional halt.

9.4.3. *Jump Destination Validity.* We previously used D as the function to determine the set of valid jump destinations given the code that is being run. We define this as any position in the code occupied by a JUMPDEST instruction.

All such positions must be on valid instruction boundaries, rather than sitting in the data portion of PUSH operations and must appear within the explicitly defined portion of the code (rather than in the implicitly defined STOP operations that trail it).

Formally:

$$(129) \qquad D(\mathbf{c}) \equiv D_J(\mathbf{c}, 0)$$

where:

$$(130)$$
$$D_J(\mathbf{c}, i) \equiv \begin{cases} \{\} & \text{if} \quad i \geqslant |\mathbf{c}| \\ \{i\} \cup D_J(\mathbf{c}, N(i, \mathbf{c}[i])) & \text{if} \quad \mathbf{c}[i] = \text{JUMPDEST} \\ D_J(\mathbf{c}, N(i, \mathbf{c}[i])) & \text{otherwise} \end{cases}$$

where N is the next valid instruction position in the code, skipping the data of a PUSH instruction, if any:

$$(131)$$
$$N(i, w) \equiv \begin{cases} i + w - \text{PUSH1} + 2 & \text{if} \quad w \in [\text{PUSH1}, \text{PUSH32}] \\ i + 1 & \text{otherwise} \end{cases}$$

9.4.4. *Normal Halting.* The normal halting function H is defined:

$$(132)$$
$$H(\boldsymbol{\mu}, I) \equiv \begin{cases} H_{\text{RETURN}}(\boldsymbol{\mu}) & \text{if} \quad w = \text{RETURN} \\ () & \text{if} \quad w \in \text{STOP}, \text{SELFDESTRUCT} \\ \varnothing & \text{otherwise} \end{cases}$$

The data-returning halt operation, RETURN, has a special function H_{RETURN}. Note also the difference between the empty sequence and the empty set as discussed here.

9.5. **The Execution Cycle.** Stack items are added or removed from the left-most, lower-indexed portion of the series; all other items remain unchanged:

$$(133) \qquad O\big((\boldsymbol{\sigma}, \boldsymbol{\mu}, A, I)\big) \equiv (\boldsymbol{\sigma}', \boldsymbol{\mu}', A', I)$$
$$(134) \qquad \Delta \equiv \alpha_w - \delta_w$$
$$(135) \qquad \|\boldsymbol{\mu}'_{\mathbf{s}}\| \equiv \|\boldsymbol{\mu}_{\mathbf{s}}\| + \Delta$$
$$(136) \qquad \forall x \in [\alpha_w, \|\boldsymbol{\mu}'_{\mathbf{s}}\|) : \boldsymbol{\mu}'_{\mathbf{s}}[x] \equiv \boldsymbol{\mu}_{\mathbf{s}}[x - \Delta]$$

The gas is reduced by the instruction's gas cost and for most instructions, the program counter increments on each cycle, for the three exceptions, we assume a function J, subscripted by one of two instructions, which evaluates to the according value:

$$(137) \qquad \boldsymbol{\mu}'_g \equiv \boldsymbol{\mu}_g - C(\boldsymbol{\sigma}, \boldsymbol{\mu}, I)$$

$$(138) \qquad \boldsymbol{\mu}'_{pc} \equiv \begin{cases} J_{\text{JUMP}}(\boldsymbol{\mu}) & \text{if} \quad w = \text{JUMP} \\ J_{\text{JUMPI}}(\boldsymbol{\mu}) & \text{if} \quad w = \text{JUMPI} \\ N(\boldsymbol{\mu}_{pc}, w) & \text{otherwise} \end{cases}$$

In general, we assume the memory, self-destruct set and system state don't change:

$$(139) \qquad \boldsymbol{\mu}'_{\mathbf{m}} \equiv \boldsymbol{\mu}_{\mathbf{m}}$$
$$(140) \qquad \boldsymbol{\mu}'_{\mathbf{i}} \equiv \boldsymbol{\mu}_{\mathbf{i}}$$
$$(141) \qquad A' \equiv A$$
$$(142) \qquad \boldsymbol{\sigma}' \equiv \boldsymbol{\sigma}$$

However, instructions do typically alter one or several components of these values. Altered components listed by instruction are noted in Appendix H, alongside values for α and δ and a formal description of the gas requirements.

10. BLOCKTREE TO BLOCKCHAIN

The canonical blockchain is a path from root to leaf through the entire block tree. In order to have consensus over which path it is, conceptually we identify the path that has had the most computation done upon it, or, the *heaviest* path. Clearly one factor that helps determine the heaviest path is the block number of the leaf, equivalent to the number of blocks, not counting the unmined genesis block, in the path. The longer the path, the greater the total mining effort that must have been done in order to arrive at the leaf. This is akin to existing schemes, such as that employed in Bitcoin-derived protocols.

Since a block header includes the difficulty, the header alone is enough to validate the computation done. Any block contributes toward the total computation or *total difficulty* of a chain.

Thus we define the total difficulty of block B recursively as:

$$(143) \qquad B_{\mathbf{t}} \equiv B'_{\mathbf{t}} + B_{\mathbf{d}}$$
$$(144) \qquad B' \equiv P(B_H)$$

As such given a block B, $B_{\mathbf{t}}$ is its total difficulty, B' is its parent block and $B_{\mathbf{d}}$ is its difficulty.

11. BLOCK FINALISATION

The process of finalising a block involves four stages:
  (1) Validate (or, if mining, determine) ommers;
  (2) validate (or, if mining, determine) transactions;
  (3) apply rewards;
  (4) verify (or, if mining, compute a valid) state and nonce.

11.1. **Ommer Validation.** The validation of ommer headers means nothing more than verifying that each ommer header is both a valid header and satisfies the relation of Nth-generation ommer to the present block where $N \leq 6$. The maximum of ommer headers is two. Formally:

$$(145) \qquad \|B_{\mathbf{U}}\| \leqslant 2 \bigwedge_{U \in B_{\mathbf{U}}} V(U) \ \wedge \ k(U, P(B_H)_H, 6)$$

where k denotes the "is-kin" property:

$$(146)$$
$$k(U, H, n) \equiv \begin{cases} false & \text{if} \quad n = 0 \\ s(U, H) & \\ \quad \vee \ k(U, P(H)_H, n - 1) & \text{otherwise} \end{cases}$$

and s denotes the "is-sibling" property:

$$(147)$$
$$s(U, H) \equiv (P(H) = P(U) \ \wedge \ H \neq U \ \wedge \ U \notin B(H)_{\mathbf{U}})$$

where $B(H)$ is the block of the corresponding header H.

11.2. Transaction Validation. The given **gasUsed** must correspond faithfully to the transactions listed: $B_{H\mathbf{g}}$, the total gas used in the block, must be equal to the accumulated gas used according to the final transaction:

$$(148) \qquad\qquad B_{H\mathbf{g}} = \ell(\mathbf{R})_{\mathbf{u}}$$

11.3. Reward Application. The application of rewards to a block involves raising the balance of the accounts of the beneficiary address of the block and each ommer by a certain amount. We raise the block's beneficiary account by $R_{\mathbf{b}}$; for each ommer, we raise the block's beneficiary by an additional $\frac{1}{32}$ of the block reward and the beneficiary of the ommer gets rewarded depending on the block number. Formally we define the function Ω, the block-finalisation state transition function (a function that rewards a nominated party):

$$(149)\quad \Omega(B, \boldsymbol{\sigma}) \quad\equiv\quad \boldsymbol{\sigma}' : \boldsymbol{\sigma}' = \boldsymbol{\sigma} \quad \text{except:}$$

$$(150)\boldsymbol{\sigma}'[B_{H\mathbf{c}}]_{\mathbf{b}} \quad=\quad \boldsymbol{\sigma}[B_{H\mathbf{c}}]_{\mathbf{b}} + (1 + \frac{\|B_{\mathbf{U}}\|}{32})R_{\mathbf{b}}$$

$$(151)\, \forall_{U \in B_{\mathbf{U}}} :$$

$$\boldsymbol{\sigma}'[U_{\mathbf{c}}]_{\mathbf{b}} \quad=\quad \boldsymbol{\sigma}[U_{\mathbf{c}}]_{\mathbf{b}} + (1 + \frac{1}{8}(U_{\mathbf{i}} - B_{H\mathbf{i}}))R_{\mathbf{b}}$$

If there are collisions of the beneficiary addresses between ommers and the block (i.e. two ommers with the same beneficiary address or an ommer with the same beneficiary address as the present block), additions are applied cumulatively.

We define the block reward as 5 Ether:

$$(152) \qquad\qquad \text{Let} \quad R_{\mathbf{b}} = 5 \times 10^{18}$$

11.4. State & Nonce Validation. We may now define the function, Γ, that maps a block B to its initiation state:
$$(153)$$
$$\Gamma(B) \equiv \begin{cases} \boldsymbol{\sigma}_0 & \text{if} \quad P(B_H) = \varnothing \\ \boldsymbol{\sigma}_{\mathbf{i}} : \text{TRIE}(L_S(\boldsymbol{\sigma}_{\mathbf{i}})) = P(B_H)_{H_{\mathbf{r}}} & \text{otherwise} \end{cases}$$

Here, $\text{TRIE}(L_S(\boldsymbol{\sigma}_{\mathbf{i}}))$ means the hash of the root node of a trie of state $\boldsymbol{\sigma}_{\mathbf{i}}$; it is assumed that implementations will store this in the state database, which is trivial and efficient since the trie is by nature an immutable data structure. And finally we define Φ, the block transition function, which maps an incomplete block B to a complete block B':

$$(154)\, \Phi(B) \quad\equiv\quad B' : \quad B' = B^* \quad \text{except:}$$

$$(155) \quad B'_{\mathbf{n}} \quad=\quad n : \quad x \leqslant \frac{2^{256}}{H_{\mathbf{d}}}$$

$$(156) \quad B'_{\mathbf{m}} \quad=\quad m \quad \text{with } (x, m) = \text{PoW}(B^*_{H\!\!\!/}, n, \mathbf{d})$$

$$(157) \quad B^* \quad\equiv\quad B \quad \text{except:} \quad B^*_{\mathbf{r}} = r(\Pi(\Gamma(B), B))$$

With $\mathbf{d}$ being a dataset as specified in appendix J.

As specified at the beginning of the present work, Π is the state-transition function, which is defined in terms of Ω, the block finalisation function and Υ, the transaction-evaluation function, both now well-defined.

As previously detailed, $\mathbf{R}[n]_{\boldsymbol{\sigma}}$, $\mathbf{R}[n]_{\mathbf{l}}$ and $\mathbf{R}[n]_{\mathbf{u}}$ are the nth corresponding states, logs and cumulative gas used after each transaction ($\mathbf{R}[n]_{\mathbf{b}}$, the fourth component in the tuple, has already been defined in terms of the logs). The former is defined simply as the state resulting from applying the corresponding transaction to the state resulting

from the previous transaction (or the block's initial state in the case of the first such transaction):

$$(158) \qquad \mathbf{R}[n]_{\boldsymbol{\sigma}} = \begin{cases} \Gamma(B) & \text{if} \quad n < 0 \\ \Upsilon(\mathbf{R}[n-1]_{\boldsymbol{\sigma}}, B_{\mathbf{T}}[n]) & \text{otherwise} \end{cases}$$

In the case of $B_{\mathbf{R}}[n]_{\mathbf{u}}$, we take a similar approach defining each item as the gas used in evaluating the corresponding transaction summed with the previous item (or zero, if it is the first), giving us a running total:

$$(159) \quad \mathbf{R}[n]_{\mathbf{u}} = \begin{cases} 0 & \text{if} \quad n < 0 \\ \Upsilon^g(\mathbf{R}[n-1]_{\boldsymbol{\sigma}}, B_{\mathbf{T}}[n]) \\ \quad + \mathbf{R}[n-1]_{\mathbf{u}} & \text{otherwise} \end{cases}$$

For $\mathbf{R}[n]_{\mathbf{l}}$, we utilise the $\Upsilon^{\mathbf{l}}$ function that we conveniently defined in the transaction execution function.

$$(160) \qquad \mathbf{R}[n]_{\mathbf{l}} = \Upsilon^{\mathbf{l}}(\mathbf{R}[n-1]_{\boldsymbol{\sigma}}, B_{\mathbf{T}}[n])$$

Finally, we define Π as the new state given the block reward function Ω applied to the final transaction's resultant state, $\ell(B_{\mathbf{R}})_{\boldsymbol{\sigma}}$:

$$(161) \qquad\qquad \Pi(\boldsymbol{\sigma}, B) \equiv \Omega(B, \ell(\mathbf{R})_{\boldsymbol{\sigma}})$$

Thus the complete block-transition mechanism is defined, except for PoW, the proof-of-work function.

11.5. Mining Proof-of-Work. The mining proof-of-work (PoW) exists as a cryptographically secure nonce that proves beyond reasonable doubt that a particular amount of computation has been expended in the determination of some token value n. It is utilised to enforce the blockchain security by giving meaning and credence to the notion of difficulty (and, by extension, total difficulty). However, since mining new blocks comes with an attached reward, the proof-of-work not only functions as a method of securing confidence that the blockchain will remain canonical into the future, but also as a wealth distribution mechanism.

For both reasons, there are two important goals of the proof-of-work function; firstly, it should be as accessible as possible to as many people as possible. The requirement of, or reward from, specialised and uncommon hardware should be minimised. This makes the distribution model as open as possible, and, ideally, makes the act of mining a simple swap from electricity to Ether at roughly the same rate for anyone around the world.

Secondly, it should not be possible to make super-linear profits, and especially not so with a high initial barrier. Such a mechanism allows a well-funded adversary to gain a troublesome amount of the network's total mining power and as such gives them a super-linear reward (thus skewing distribution in their favour) as well as reducing the network security.

One plague of the Bitcoin world is ASICs. These are specialised pieces of compute hardware that exist only to do a single task (Smith [1997], ASI [a]). In Bitcoin's case the task is the SHA256 hash function (ASI [b]). While ASICs exist for a proof-of-work function, both goals are placed in jeopardy. Because of this, a proof-of-work function that is ASIC-resistant (i.e. difficult or economically inefficient to implement in specialised compute hardware) has been identified as the proverbial silver bullet.

Two directions exist for ASIC resistance; firstly make it sequential memory-hard, i.e. engineer the function such

that the determination of the nonce requires a lot of memory and bandwidth such that the memory cannot be used in parallel to discover multiple nonces simultaneously. The second is to make the type of computation it would need to do general-purpose; the meaning of "specialised hardware" for a general-purpose task set is, naturally, general purpose hardware and as such commodity desktop computers are likely to be pretty close to "specialised hardware" for the task. For Ethereum *Frontier* and *Homestead* we have chosen the first path.

More formally, the proof-of-work function takes the form of PoW:

(162)
$$m = H_{\mathbf{m}} \quad \wedge \quad n \leqslant \frac{2^{256}}{H_{\mathbf{d}}} \quad \text{with} \quad (m, n) = \mathtt{PoW}(H_{\slashed{n}}, H_{\mathbf{n}}, \mathbf{d})$$

Where $H_{\slashed{n}}$ is the new block's header but *without* the nonce and mix-hash components; $H_{\mathbf{n}}$ is the nonce of the block header; $\mathbf{d}$ is a large data set needed to compute the mixHash and $H_{\mathbf{d}}$ is the new block's difficulty value (i.e. the block difficulty from section 10). PoW is the proof-of-work function which evaluates to an array with the first item being the mixHash and the second item being the block nonce, a pseudo-random number cryptographically dependent on H and $\mathbf{d}$. The underlying algorithm is called Ethash and is described below.

11.5.1. *Ethash.* Ethash is the PoW algorithm for Ethereum *Frontier* and *Homestead*. It is the latest version of Dagger-Hashimoto, introduced by Buterin [2013b] and Dryja [2014], although it can no longer appropriately be called that since many of the original features of both algorithms were drastically changed with R&D from February 2015 until May 4 2015 (Christoph Jentzsch). The general route that the algorithm takes is as follows:

There exists a seed which can be computed for each block by scanning through the block headers up until that point. From the seed, one can compute a pseudorandom cache, $J_{\text{cacheinit}}$ bytes in initial size. Light clients store the cache. From the cache, we can generate a dataset, $J_{\text{datasetinit}}$ bytes in initial size, with the property that each item in the dataset depends on only a small number of items from the cache. Full clients and miners store the dataset. The dataset grows linearly with time.

Mining involves grabbing random slices of the dataset and hashing them together. Verification can be done with low memory by using the cache to regenerate the specific pieces of the dataset that you need, so you only need to store the cache. The large dataset is updated once every J_{epoch} blocks, so the vast majority of a miner's effort will be reading the dataset, not making changes to it. The mentioned parameters as well as the algorithm is explained in detail in appendix J.

12. Implementing Contracts

There are several patterns of contracts engineering that allow particular useful behaviours; two of these that I will briefly discuss are data feeds and random numbers.

12.1. **Data Feeds.** A data feed contract is one which provides a single service: it gives access to information from the external world within Ethereum. The accuracy and timeliness of this information is not guaranteed and it is the task of a secondary contract author—the contract that

utilises the data feed—to determine how much trust can be placed in any single data feed.

The general pattern involves a single contract within Ethereum which, when given a message call, replies with some timely information concerning an external phenomenon. An example might be the local temperature of New York City. This would be implemented as a contract that returned that value of some known point in storage. Of course this point in storage must be maintained with the correct such temperature, and thus the second part of the pattern would be for an external server to run an Ethereum node, and immediately on discovery of a new block, creates a new valid transaction, sent to the contract, updating said value in storage. The contract's code would accept such updates only from the identity contained on said server.

12.2. **Random Numbers.** Providing random numbers within a deterministic system is, naturally, an impossible task. However, we can approximate with pseudo-random numbers by utilising data which is generally unknowable at the time of transacting. Such data might include the block's hash, the block's timestamp and the block's beneficiary address. In order to make it hard for malicious miner to control those values, one should use the BLOCKHASH operation in order to use hashes of the previous 256 blocks as pseudo-random numbers. For a series of such numbers, a trivial solution would be to add some constant amount and hashing the result.

13. Future Directions

The state database won't be forced to maintain all past state trie structures into the future. It should maintain an age for each node and eventually discard nodes that are neither recent enough nor checkpoints. Checkpoints, or a set of nodes in the database that allow a particular block's state trie to be traversed, could be used to place a maximum limit on the amount of computation needed in order to retrieve any state throughout the blockchain.

Blockchain consolidation could be used in order to reduce the amount of blocks a client would need to download to act as a full, mining, node. A compressed archive of the trie structure at given points in time (perhaps one in every 10,000th block) could be maintained by the peer network, effectively recasting the genesis block. This would reduce the amount to be downloaded to a single archive plus a hard maximum limit of blocks.

Finally, blockchain compression could perhaps be conducted: nodes in state trie that haven't sent/received a transaction in some constant amount of blocks could be thrown out, reducing both Ether-leakage and the growth of the state database.

13.1. **Scalability.** Scalability remains an eternal concern. With a generalised state transition function, it becomes difficult to partition and parallelise transactions to apply the divide-and-conquer strategy. Unaddressed, the dynamic value-range of the system remains essentially fixed and as the average transaction value increases, the less valuable of them become ignored, being economically pointless to include in the main ledger. However, several strategies exist that may potentially be exploited to provide a considerably more scalable protocol.

Some form of hierarchical structure, achieved by either consolidating smaller lighter-weight chains into the main block or building the main block through the incremental combination and adhesion (through proof-of-work) of smaller transaction sets may allow parallelisation of transaction combination and block-building. Parallelism could also come from a prioritised set of parallel blockchains, consolidating each block and with duplicate or invalid transactions thrown out accordingly.

Finally, verifiable computation, if made generally available and efficient enough, may provide a route to allow the proof-of-work to be the verification of final state.

14. Conclusion

I have introduced, discussed and formally defined the protocol of Ethereum. Through this protocol the reader may implement a node on the Ethereum network and join others in a decentralised secure social operating system. Contracts may be authored in order to algorithmically specify and autonomously enforce rules of interaction.

15. Acknowledgements

Many thanks to Aeron Buchanan for authoring the Homestead revisions, Christoph Jentzsch for authoring the Ethash algorithm and Yoichi Hirai for doing most of the EIP-150 changes. Important maintenance, useful corrections and suggestions were provided by a number of others from the Ethereum DEV organisation and Ethereum community at large including Gustav Simonsson, Paweł Bylica, Jutta Steiner, Nick Savers, Viktor Trón, Marko Simovic, Giacomo Tazzari and, of course, Vitalik Buterin.

16. Availability

The source of this paper is maintained at `https://github.com/ethereum/yellowpaper/`. An auto-generated PDF is located at `https://ethereum.github.io/yellowpaper/paper.pdf`.

References

Application-specific integrated circuit, a. URL `https://en.wikipedia.org/wiki/Application-specific_integrated_circuit`. `https://web.archive.org/web/20170929032958/https://en.wikipedia.org/wiki/Application-specific_integrated_circuit`.

ASIC, b. URL `{https://en.bitcoin.it/wiki/ASIC}`. `https://web.archive.org/web/20170929042224/https://en.bitcoin.it/wiki/ASIC`.

Elliptic Curve Digital Signature Algorithm. URL `https://en.wikipedia.org/wiki/Elliptic_Curve_Digital_Signature_Algorithm`. `https://web.archive.org/web/20170916033830/https://en.wikipedia.org/wiki/Elliptic_Curve_Digital_Signature_Algorithm`.

SHA-3. URL `https://en.wikipedia.org/wiki/SHA-3`. `https://web.archive.org/web/20171001081722/https://en.wikipedia.org/wiki/SHA-3`.

Lattice (order). URL `https://en.wikipedia.org/wiki/Lattice_(order)`. `https://web.archive.org/save/https://en.wikipedia.org/wiki/Lattice_(order)`.

Secp256k1. URL `https://en.bitcoin.it/wiki/Secp256k1`. `https://web.archive.org/web/20170916032105/https://en.bitcoin.it/wiki/Secp256k1`.

Pierre Arnaud, Mathieu Schroeter, and Sam Le Barbare. Electrum, 2017. URL `https://www.npmjs.com/package/electrum`. `https://web.archive.org/save/https://www.npmjs.com/package/electrum`.

Jacob Aron. BitCoin software finds new life. *New Scientist*, 213(2847):20, 2012. URL `http://www.sciencedirect.com/science/article/pii/S0262407912601055`.

Adam Back. Hashcash - Amortizable Publicly Auditable Cost-Functions. 2002. URL `http://www.hashcash.org/papers/amortizable.pdf`. `https://web.archive.org/web/20170810043047/http://www.hashcash.org/papers/amortizable.pdf`.

Roman Boutellier and Mareike Heinzen. Pirates, Pioneers, Innovators and Imitators. In *Growth Through Innovation*, pages 85–96. Springer, 2014. URL `https://www.springer.com/gb/book/9783319040158`. URL available at `http://wiki.erights.org/wiki/Documentation`. `https://web.archive.org/web/20170810040208/https://www.springer.com/gb/book/9783319040158`.

Vitalik Buterin. Ethereum: A Next-Generation Smart Contract and Decentralized Application Platform. 2013a. URL `https://github.com/ethereum/wiki/wiki/White-Paper`.

Vitalik Buterin. Dagger: A Memory-Hard to Compute, Memory-Easy to Verify Scrypt Alternative. 2013b. URL `http://www.hashcash.org/papers/dagger.html`. `https://web.archive.org/web/20170810043955/http://www.hashcash.org/papers/dagger.html`. Dead original link as of 10 August 2017: `http://vitalik.ca/ethereum/dagger.html`.

Christoph Jentzsch. Commit date for ethash. URL `https://github.com/ethereum/yellowpaper/commit/77a8cf2428ce245bf6e2c39c5e652ba58a278666#commitcomment-24644869`. Not able to be archived by the Wayback Machine, the saved page doesn't load showing the cited change and comment.

Thaddeus Dryja. Hashimoto: I/O bound proof of work. 2014. URL `https://pdfs.semanticscholar.org/3b23/7cc60c1b9650e260318d33bec471b8202d5e.pdf`. `https://web.archive.org/web/20170810043640/https://pdfs.semanticscholar.org/3b23/7cc60c1b9650e260318d33bec471b8202d5e.pdf`. Dead original link as of 10 August 2017: `https://mirrorx.com/files/hashimoto.pdf`.

Cynthia Dwork and Moni Naor. Pricing via processing or combatting junk mail. In *In 12th Annual International Cryptology Conference*, pages 139–147, 1992. URL `http://www.wisdom.weizmann.ac.il/~naor/PAPERS/pvp.pdf`. `https://web.archive.org/web/20170810035254/http://www.wisdom.weizmann.ac.il/~naor/PAPERS/pvp.pdf`.

Fabian Vogelsteller and Alex Van de Sande and Everton Fraga and Ramesh Nair and Luca Zeug. [mist release 0.8.0. URL `https://github.com/ethereum/mist/releases/tag/0.8.0`. `https://web.archive.org/web/20170930071729/https://github.com/ethereum/mist/releases/tag/0.8.0`.

Phong Vo Glenn Fowler, Landon Curt Noll. Fowler–Noll–Vo hash function. 1991. URL `https://en.wikipedia.org/wiki/Fowler%E2%80%93Noll%E2%80%93Vo_hash_function#cite_note-2`.

Guido Bertoni and Joan Daemen and Michaël Peeters and Gilles Van Assche and Ronny Van Keer. KECCAK, 2017. URL `https://keccak.team/keccak.html`. Unable to be archived by the Wayback Machine.

Nils Gura, Arun Patel, Arvinderpal Wander, Hans Eberle, and Sheueling Chang Shantz. Comparing elliptic curve cryptography and RSA on 8-bit CPUs. In *Cryptographic Hardware and Embedded Systems-CHES 2004*, pages 119–132. Springer, 2004. URL `https://www.iacr.org/archive/ches2004/31560117/31560117.pdf`. `https://web.archive.org/web/20170810035057/https://www.iacr.org/archive/ches2004/31560117/31560117.pdf`.

Don Johnson, Alfred Menezes, and Scott Vanstone. The Elliptic Curve Digital Signature Algorithm (ECDSA), 2001. URL `http://cs.ucsb.edu/~koc/ccs130h/notes/ecdsa-cert.pdf`. `https://web.archive.org/save/http://cs.ucsb.edu/~koc/ccs130h/notes/ecdsa-cert.pdf`.

Sergio Demian Lerner. Strict Memory Hard Hashing Functions. 2014. URL `http://www.hashcash.org/papers/memohash.pdf`. `https://web.archive.org/web/20170810044110/http://www.hashcash.org/papers/memohash.pdf`.

Mark Miller. The Future of Law. In *paper delivered at the Extro 3 Conference (August 9)*, 1997. URL `https://drive.google.com/file/d/0BwOVXJKBgYPMS0J2VGIyWWlocms/edit?usp=sharing`.

Satoshi Nakamoto. Bitcoin: A peer-to-peer electronic cash system. *Consulted*, 1:2012, 2008. URL `http://nakamotoinstitute.org/bitcoin/`.

Meni Rosenfeld. Overview of Colored Coins. 2012. URL `https://bitcoil.co.il/BitcoinX.pdf`. Test: `https://web.archive.org/web/20170810040120/https://bitcoil.co.il/BitcoinX.pdf`.

Michael John Sebastian Smith. *Application-Specific Integrated Circuits*. Addison-Wesley, 1997. ISBN 0201500221. Available here: `https://www.slideshare.net/nhKhanhNguyn/michael-john-sebastian-smith-applicationspecific-integrated-circuitsaddisonwesley-professional-1997-1` with the archive: `https://web.archive.org/save/https://www.slideshare.net/nhKhanhNguyn/michael-john-sebastian-smith-applicationspecific-integrated-circuitsaddisonwesley-professional-1997-1`. A preview is available at `https://www.amazon.com/Application-Specific-Integrated-Circuits-Michael-Smith/dp/0321602757`, with the archive: `https://web.archive.org/web/20170929041938/https://www.amazon.com/Application-Specific-Integrated-Circuits-Michael-Smith/dp/0321602757`.

Yonatan Sompolinsky and Aviv Zohar. Accelerating Bitcoin's Transaction Processing. Fast Money Grows on Trees, Not Chains. Cryptology ePrint Archive, Report 2013/881, 2013. `http://eprint.iacr.org/2013/881`.

Simon Sprankel. Technical Basis of Digital Currencies, 2013. URL `http://www.coderblog.de/wp-content/uploads/technical-basis-of-digital-currencies.pdf`. `https://web.archive.org/web/20170810025028/http://www.coderblog.de/wp-content/uploads/technical-basis-of-digital-currencies.pdf`.

Nick Szabo. Formalizing and securing relationships on public networks. *First Monday*, 2(9), 1997. URL `http://firstmonday.org/ojs/index.php/fm/article/view/548`. `https://web.archive.org/web/20170810042659/http://firstmonday.org/ojs/index.php/fm/article/view/548`.

Vivek Vishnumurthy, Sangeeth Chandrakumar, and Emin Gün Sirer. Karma: A secure economic framework for peer-to-peer resource sharing, 2003. URL `https://www.cs.cornell.edu/people/egs/papers/karma.pdf`. `https://web.archive.org/web/20170810031834/https://www.cs.cornell.edu/people/egs/papers/karma.pdf`.

J. R. Willett. MasterCoin Complete Specification. 2013. URL `https://github.com/mastercoin-MSC/spec`. `https://web.archive.org/web/20170810035927/https://github.com/OmniLayer/spec`.

Pieter Wuille. What does the curve used in Bitcoin, secp256k1, look like?, 2014. URL `https://bitcoin.stackexchange.com/a/21911/48875`. `https://web.archive.org/web/20170916031457/https://bitcoin.stackexchange.com/questions/21907/what-does-the-curve-used-in-bitcoin-secp256k1-look-like/21911`.

Appendix A. Terminology

External Actor: A person or other entity able to interface to an Ethereum node, but external to the world of Ethereum. It can interact with Ethereum through depositing signed Transactions and inspecting the blockchain and associated state. Has one (or more) intrinsic Accounts.

Address: A 160-bit code used for identifying Accounts.

Account: Accounts have an intrinsic balance and transaction count maintained as part of the Ethereum state. They also have some (possibly empty) EVM Code and a (possibly empty) Storage State associated with them. Though homogenous, it makes sense to distinguish between two practical types of account: those with empty associated EVM Code (thus the account balance is controlled, if at all, by some external entity) and those with non-empty associated EVM Code (thus the account represents an Autonomous Object). Each Account has a single Address that identifies it.

Transaction: A piece of data, signed by an External Actor. It represents either a Message or a new Autonomous Object. Transactions are recorded into each block of the blockchain.

Autonomous Object: A notional object existent only within the hypothetical state of Ethereum. Has an intrinsic address and thus an associated account; the account will have non-empty associated EVM Code. Incorporated only as the Storage State of that account.

Storage State: The information particular to a given Account that is maintained between the times that the Account's associated EVM Code runs.

Message: Data (as a set of bytes) and Value (specified as Ether) that is passed between two Accounts, either through the deterministic operation of an Autonomous Object or the cryptographically secure signature of the Transaction.

Message Call: The act of passing a message from one Account to another. If the destination account is associated with non-empty EVM Code, then the VM will be started with the state of said Object and the Message acted upon. If the message sender is an Autonomous Object, then the Call passes any data returned from the VM operation.

Gas: The fundamental network cost unit. Paid for exclusively by Ether (as of PoC-4), which is converted freely to and from Gas as required. Gas does not exist outside of the internal Ethereum computation engine; its price is set by the Transaction and miners are free to ignore Transactions whose Gas price is too low.

Contract: Informal term used to mean both a piece of EVM Code that may be associated with an Account or an Autonomous Object.

Object: Synonym for Autonomous Object.

App: An end-user-visible application hosted in the Ethereum Browser.

Ethereum Browser: (aka Ethereum Reference Client) A cross-platform GUI of an interface similar to a simplified browser (a la Chrome) that is able to host sandboxed applications whose backend is purely on the Ethereum protocol, which is known as Mist since 8 July 2016 (Fabian Vogelsteller and Alex Van de Sande and Everton Fraga and Ramesh Nair and Luca Zeug).

Ethereum Virtual Machine: (aka EVM) The virtual machine that forms the key part of the execution model for an Account's associated EVM Code.

Ethereum Runtime Environment: (aka ERE) The environment which is provided to an Autonomous Object executing in the EVM. Includes the EVM but also the structure of the world state on which the EVM relies for certain I/O instructions including CALL & CREATE.

EVM Code: The bytecode that the EVM can natively execute. Used to formally specify the meaning and ramifications of a message to an Account.

EVM Assembly: The human-readable form of EVM-code.

LLL: The Lisp-like Low-level Language, a human-writable language used for authoring simple contracts and general low-level language toolkit for trans-compiling to.

Appendix B. Recursive Length Prefix

This is a serialisation method for encoding arbitrarily structured binary data (byte arrays).

We define the set of possible structures $\mathbb{T}$:

$$(163) \qquad \mathbb{T} \quad \equiv \quad \mathbb{L} \cup \mathbb{B}$$

$$(164) \qquad \mathbb{L} \quad \equiv \quad \{ \mathbf{t} : \mathbf{t} = (\mathbf{t}[0], \mathbf{t}[1], ...) \ \wedge \ \forall_{n < \|\mathbf{t}\|} \ \mathbf{t}[n] \in \mathbb{T} \}$$

$$(165) \qquad \mathbb{B} \quad \equiv \quad \{ \mathbf{b} : \mathbf{b} = (\mathbf{b}[0], \mathbf{b}[1], ...) \ \wedge \ \forall_{n < \|\mathbf{b}\|} \ \mathbf{b}[n] \in \mathbb{O} \}$$

Where $\mathbb{O}$ is the set of bytes. Thus $\mathbb{B}$ is the set of all sequences of bytes (otherwise known as byte-arrays, and a leaf if imagined as a tree), $\mathbb{L}$ is the set of all tree-like (sub-)structures that are not a single leaf (a branch node if imagined as a tree) and $\mathbb{T}$ is the set of all byte-arrays and such structural sequences.

We define the RLP function as $\mathtt{RLP}$ through two sub-functions, the first handling the instance when the value is a byte array, the second when it is a sequence of further values:

$$(166) \qquad \mathtt{RLP}(\mathbf{x}) \equiv \begin{cases} R_{\mathbf{b}}(\mathbf{x}) & \text{if} \quad \mathbf{x} \in \mathbb{B} \\ R_{\mathbf{l}}(\mathbf{x}) & \text{otherwise} \end{cases}$$

If the value to be serialised is a byte-array, the RLP serialisation takes one of three forms:

- If the byte-array contains solely a single byte and that single byte is less than 128, then the input is exactly equal to the output.
- If the byte-array contains fewer than 56 bytes, then the output is equal to the input prefixed by the byte equal to the length of the byte array plus 128.
- Otherwise, the output is equal to the input prefixed by the minimal-length byte-array which when interpreted as a big-endian integer is equal to the length of the input byte array, which is itself prefixed by the number of bytes required to faithfully encode this length value plus 183.

Formally, we define $R_{\mathbf{b}}$:

$$(167) \qquad R_{\mathbf{b}}(\mathbf{x}) \quad \equiv \quad \begin{cases} \mathbf{x} & \text{if} \quad \|\mathbf{x}\| = 1 \wedge \mathbf{x}[0] < 128 \\ (128 + \|\mathbf{x}\|) \cdot \mathbf{x} & \text{else if} \quad \|\mathbf{x}\| < 56 \\ \big(183 + \big\|\mathtt{BE}(\|\mathbf{x}\|)\big\|\big) \cdot \mathtt{BE}(\|\mathbf{x}\|) \cdot \mathbf{x} & \text{otherwise} \end{cases}$$

$$(168) \qquad \mathtt{BE}(x) \quad \equiv \quad (b_0, b_1, ...) : b_0 \neq 0 \wedge x = \sum_{n=0}^{n < \|\mathbf{b}\|} b_n \cdot 256^{\|\mathbf{b}\| - 1 - n}$$

$$(169) \qquad (a) \cdot (b, c) \cdot (d, e) \quad = \quad (a, b, c, d, e)$$

Thus BE is the function that expands a positive integer value to a big-endian byte array of minimal length and the dot operator performs sequence concatenation.

If instead, the value to be serialised is a sequence of other items then the RLP serialisation takes one of two forms:

- If the concatenated serialisations of each contained item is less than 56 bytes in length, then the output is equal to that concatenation prefixed by the byte equal to the length of this byte array plus 192.
- Otherwise, the output is equal to the concatenated serialisations prefixed by the minimal-length byte-array which when interpreted as a big-endian integer is equal to the length of the concatenated serialisations byte array, which is itself prefixed by the number of bytes required to faithfully encode this length value plus 247.

Thus we finish by formally defining R_l:

$$(170) \qquad R_\mathrm{l}(\mathbf{x}) \quad \equiv \quad \begin{cases} (192 + \|s(\mathbf{x})\|) \cdot s(\mathbf{x}) & \text{if} \quad \|s(\mathbf{x})\| < 56 \\ \left(247 + \big\|\mathtt{BE}(\|s(\mathbf{x})\|)\big\|\right) \cdot \mathtt{BE}(\|s(\mathbf{x})\|) \cdot s(\mathbf{x}) & \text{otherwise} \end{cases}$$

$$(171) \qquad s(\mathbf{x}) \quad \equiv \quad \mathtt{RLP}(\mathbf{x}_0) \cdot \mathtt{RLP}(\mathbf{x}_1)...$$

If RLP is used to encode a scalar, defined only as a positive integer ($\mathbb{P}$ or any x for $\mathbb{P}_x$), it must be specified as the shortest byte array such that the big-endian interpretation of it is equal. Thus the RLP of some positive integer i is defined as:

$$(172) \qquad \mathtt{RLP}(i : i \in \mathbb{P}) \equiv \mathtt{RLP}(\mathtt{BE}(i))$$

When interpreting RLP data, if an expected fragment is decoded as a scalar and leading zeroes are found in the byte sequence, clients are required to consider it non-canonical and treat it in the same manner as otherwise invalid RLP data, dismissing it completely.

There is no specific canonical encoding format for signed or floating-point values.

Appendix C. Hex-Prefix Encoding

Hex-prefix encoding is an efficient method of encoding an arbitrary number of nibbles as a byte array. It is able to store an additional flag which, when used in the context of the trie (the only context in which it is used), disambiguates between node types.

It is defined as the function HP which maps from a sequence of nibbles (represented by the set $\mathbb{Y}$) together with a boolean value to a sequence of bytes (represented by the set $\mathbb{B}$):

$$(173) \qquad \mathtt{HP}(\mathbf{x}, t) : \mathbf{x} \in \mathbb{Y} \quad \equiv \quad \begin{cases} (16 f(t), 16\mathbf{x}[0] + \mathbf{x}[1], 16\mathbf{x}[2] + \mathbf{x}[3], ...) & \text{if} \quad \|\mathbf{x}\| \text{ is even} \\ (16(f(t) + 1) + \mathbf{x}[0], 16\mathbf{x}[1] + \mathbf{x}[2], 16\mathbf{x}[3] + \mathbf{x}[4], ...) & \text{otherwise} \end{cases}$$

$$(174) \qquad f(t) \quad \equiv \quad \begin{cases} 2 & \text{if} \quad t \neq 0 \\ 0 & \text{otherwise} \end{cases}$$

Thus the high nibble of the first byte contains two flags; the lowest bit encoding the oddness of the length and the second-lowest encoding the flag t. The low nibble of the first byte is zero in the case of an even number of nibbles and the first nibble in the case of an odd number. All remaining nibbles (now an even number) fit properly into the remaining bytes.

Appendix D. Modified Merkle Patricia Tree

The modified Merkle Patricia tree (trie) provides a persistent data structure to map between arbitrary-length binary data (byte arrays). It is defined in terms of a mutable data structure to map between 256-bit binary fragments and arbitrary-length binary data, typically implemented as a database. The core of the trie, and its sole requirement in terms of the protocol specification is to provide a single value that identifies a given set of key-value pairs, which may be either a 32 byte sequence or the empty byte sequence. It is left as an implementation consideration to store and maintain the structure of the trie in a manner that allows effective and efficient realisation of the protocol.

Formally, we assume the input value $\mathfrak{I}$, a set containing pairs of byte sequences:

$$(175) \qquad \mathfrak{I} = \{(k_0 \in \mathbb{B}, v_0 \in \mathbb{B}), k_1 \in \mathbb{B}, v_1 \in \mathbb{B}), ...\}$$

When considering such a sequence, we use the common numeric subscript notation to refer to a tuple's key or value, thus:

$$(176) \qquad \forall_{I \in \mathfrak{I}} I \equiv (I_0, I_1)$$

Any series of bytes may also trivially be viewed as a series of nibbles, given an endian-specific notation; here we assume big-endian. Thus:

$$(177) \qquad y(\mathfrak{I}) \quad = \quad \{(k_0' \in \mathbb{Y}, v_0 \in \mathbb{B}), (k_1' \in \mathbb{Y}, v_1 \in \mathbb{B}), ...\}$$

$$(178) \qquad \forall_n \quad \forall_{i:i<2\| k_n\|} \quad k_n'[i] \quad \equiv \quad \begin{cases} \lfloor k_n[i \div 2] \div 16 \rfloor & \text{if } i \text{ is even} \\ k_n[\lfloor i \div 2 \rfloor] \bmod 16 & \text{otherwise} \end{cases}$$

We define the function TRIE, which evaluates to the root of the trie that represents this set when encoded in this structure:

$$(179) \qquad \mathtt{TRIE}(\mathfrak{I}) \equiv \mathtt{KEC}(c(\mathfrak{I}, 0))$$

We also assume a function n, the trie's node cap function. When composing a node, we use RLP to encode the structure. As a means of reducing storage complexity, for nodes whose composed RLP is fewer than 32 bytes, we store the RLP directly; for those larger we assert prescience of the byte array whose Keccak hash evaluates to our reference. Thus we define in terms of c, the node composition function:

$$(180) \qquad n(\mathfrak{I}, i) \equiv \begin{cases} () & \text{if} \quad \mathfrak{I} = \varnothing \\ c(\mathfrak{I}, i) & \text{if} \quad \|c(\mathfrak{I}, i)\| < 32 \\ \mathtt{KEC}(c(\mathfrak{I}, i)) & \text{otherwise} \end{cases}$$

In a manner similar to a radix tree, when the trie is traversed from root to leaf, one may build a single key-value pair. The key is accumulated through the traversal, acquiring a single nibble from each branch node (just as with a radix tree). Unlike a radix tree, in the case of multiple keys sharing the same prefix or in the case of a single key having a unique suffix, two optimising nodes are provided. Thus while traversing, one may potentially acquire multiple nibbles from each of the other two node types, extension and leaf. There are three kinds of nodes in the trie:

Leaf: A two-item structure whose first item corresponds to the nibbles in the key not already accounted for by the accumulation of keys and branches traversed from the root. The hex-prefix encoding method is used and the second parameter to the function is required to be *true*.

Extension: A two-item structure whose first item corresponds to a series of nibbles of size greater than one that are shared by at least two distinct keys past the accumulation of the keys of nibbles and the keys of branches as traversed from the root. The hex-prefix encoding method is used and the second parameter to the function is required to be *false*.

Branch: A 17-item structure whose first sixteen items correspond to each of the sixteen possible nibble values for the keys at this point in their traversal. The 17th item is used in the case of this being a terminator node and thus a key being ended at this point in its traversal.

A branch is then only used when necessary; no branch nodes may exist that contain only a single non-zero entry. We may formally define this structure with the structural composition function c:

$$(181)$$

$$c(\mathfrak{I}, i) \equiv \begin{cases} \mathtt{RLP}\Big(\big(\mathtt{HP}(I_0[i..(\|I_0\| - 1)], true), I_1 \big) \Big) & \text{if} \quad \|\mathfrak{I}\| = 1 \quad \text{where } \exists I : I \in \mathfrak{I} \\ \mathtt{RLP}\Big(\big(\mathtt{HP}(I_0[i..(j-1)], false), n(\mathfrak{I}, j) \big) \Big) & \text{if} \quad i \neq j \quad \text{where } j = \arg\max_x : \exists \mathbf{l} : \|\mathbf{l}\| = x : \forall_{I \in \mathfrak{I}} : I_0[0..(x-1)] = \mathbf{l} \\ \mathtt{RLP}\Big(\big(u(0), u(1), ..., u(15), v \big) \Big) & \text{otherwise} \quad \text{where } u(j) \equiv n(\{I : I \in \mathfrak{I} \wedge I_0[i] = j\}, i+1) \\ & \qquad\qquad\qquad\quad v = \begin{cases} I_1 & \text{if} \quad \exists I : I \in \mathfrak{I} \wedge \|I_0\| = i \\ () & \text{otherwise} \end{cases} \end{cases}$$

D.1. **Trie Database.** Thus no explicit assumptions are made concerning what data is stored and what is not, since that is an implementation-specific consideration; we simply define the identity function mapping the key-value set $\mathfrak{I}$ to a 32-byte hash and assert that only a single such hash exists for any $\mathfrak{I}$, which though not strictly true is accurate within acceptable precision given the Keccak hash's collision resistance. In reality, a sensible implementation will not fully recompute the trie root hash for each set.

A reasonable implementation will maintain a database of nodes determined from the computation of various tries or, more formally, it will memoise the function c. This strategy uses the nature of the trie to both easily recall the contents of any previous key-value set and to store multiple such sets in a very efficient manner. Due to the dependency relationship, Merkle-proofs may be constructed with an $O(\log N)$ space requirement that can demonstrate a particular leaf must exist within a trie of a given root hash.

APPENDIX E. PRECOMPILED CONTRACTS

For each precompiled contract, we make use of a template function, Ξ_{PRE}, which implements the out-of-gas checking.

$$(182) \qquad \Xi_{\mathrm{PRE}}(\boldsymbol{\sigma}, g, I) \equiv \begin{cases} (\varnothing, 0, A^0, ()) & \text{if} \quad \mathbf{g} < \mathbf{g_r} \\ (\boldsymbol{\sigma}, \mathbf{g} - \mathbf{g_r}, A^0, \mathbf{o}) & \text{otherwise} \end{cases}$$

The precompiled contracts each use these definitions and provide specifications for the $\mathbf{o}$ (the output data) and $\mathbf{g_r}$, the gas requirements.

For the elliptic curve DSA recover VM execution function, we also define $\mathbf{d}$ to be the input data, well-defined for an infinite length by appending zeroes as required. Importantly in the case of an invalid signature ($\mathtt{ECDSARECOVER}(h, v, r, s) =$

$\varnothing$), then we have no output.

$$(183) \qquad \Xi_{\texttt{ECREC}} \quad \equiv \quad \Xi_{\texttt{PRE}} \quad \text{where:}$$

$$(184) \qquad \mathbf{g_r} \quad = \quad 3000$$

$$(185) \qquad |\mathbf{o}| \quad = \quad \begin{cases} 0 & \text{if} \quad \texttt{ECDSARECOVER}(h,v,r,s) = \varnothing \\ 32 & \text{otherwise} \end{cases}$$

$$(186) \qquad \text{if} \quad |\mathbf{o}| = 32:$$

$$(187) \qquad \mathbf{o}[0..11] \quad = \quad 0$$

$$(188) \qquad \mathbf{o}[12..31] \quad = \quad \texttt{KEC}\big(\texttt{ECDSARECOVER}(h,v,r,s)\big)[12..31] \quad \text{where:}$$

$$(189) \qquad \mathbf{d}[0..(|I_\mathbf{d}| - 1)] \quad = \quad I_\mathbf{d}$$

$$(190) \qquad \mathbf{d}[|I_\mathbf{d}|..] \quad = \quad (0, 0, ...)$$

$$(191) \qquad h \quad = \quad \mathbf{d}[0..31]$$

$$(192) \qquad v \quad = \quad \mathbf{d}[32..63]$$

$$(193) \qquad r \quad = \quad \mathbf{d}[64..95]$$

$$(194) \qquad s \quad = \quad \mathbf{d}[96..127]$$

The two hash functions, RIPEMD-160 and SHA2-256 are more trivially defined as an almost pass-through operation. Their gas usage is dependent on the input data size, a factor rounded up to the nearest number of words.

$$(195) \qquad \Xi_{\texttt{SHA256}} \quad \equiv \quad \Xi_{\texttt{PRE}} \quad \text{where:}$$

$$(196) \qquad \mathbf{g_r} \quad = \quad 60 + 12 \left\lceil \frac{|I_\mathbf{d}|}{32} \right\rceil$$

$$(197) \qquad \mathbf{o}[0..31] \quad = \quad \texttt{SHA256}(I_\mathbf{d})$$

$$(198) \qquad \Xi_{\texttt{RIP160}} \quad \equiv \quad \Xi_{\texttt{PRE}} \quad \text{where:}$$

$$(199) \qquad \mathbf{g_r} \quad = \quad 600 + 120 \left\lceil \frac{|I_\mathbf{d}|}{32} \right\rceil$$

$$(200) \qquad \mathbf{o}[0..11] \quad = \quad 0$$

$$(201) \qquad \mathbf{o}[12..31] \quad = \quad \texttt{RIPEMD160}(I_\mathbf{d})$$

$$(202)$$

For the purposes here, we assume we have well-defined standard cryptographic functions for RIPEMD-160 and SHA2-256 of the form:

$$(203) \qquad \texttt{SHA256}(\mathbf{i} \in \mathbb{B}) \quad \equiv \quad o \in \mathbb{B}_{32}$$

$$(204) \qquad \texttt{RIPEMD160}(\mathbf{i} \in \mathbb{B}) \quad \equiv \quad o \in \mathbb{B}_{20}$$

Finally, the fourth contract, the identity function $\Xi_{\texttt{ID}}$ simply defines the output as the input:

$$(205) \qquad \Xi_{\texttt{ID}} \quad \equiv \quad \Xi_{\texttt{PRE}} \quad \text{where:}$$

$$(206) \qquad \mathbf{g_r} \quad = \quad 15 + 3 \left\lceil \frac{|I_\mathbf{d}|}{32} \right\rceil$$

$$(207) \qquad \mathbf{o} \quad = \quad I_\mathbf{d}$$

APPENDIX F. SIGNING TRANSACTIONS

The method of signing transactions is similar to the 'Electrum style signatures' as defined by Arnaud et al. [2017], heading "Managing styles with Radium" in the bullet point list. This method utilises the SECP-256k1 curve as described by sec and Wuille [2014], and is implemented similarly to as described by Gura et al. [2004] on p. 9 of 15, para. 3.

It is assumed that the sender has a valid private key p_r, which is a randomly selected positive integer (represented as a byte array of length 32 in big-endian form) in the range $[1, \texttt{secp256k1n} - 1]$.

We assert the functions ECDSAPUBKEY, ECDSARECOVER and ECDSASIGN. These are formally defined in the literature, e.g. by Johnson et al. [2001] and ECD. [4]

$$(208) \qquad \texttt{ECDSAPUBKEY}(p_r \in \mathbb{B}_{32}) \quad \equiv \quad p_u \in \mathbb{B}_{64}$$

$$(209) \qquad \texttt{ECDSASIGN}(e \in \mathbb{B}_{32}, p_r \in \mathbb{B}_{32}) \quad \equiv \quad (v \in \mathbb{B}_1, r \in \mathbb{B}_{32}, s \in \mathbb{B}_{32})$$

$$(210) \qquad \texttt{ECDSARECOVER}(e \in \mathbb{B}_{32}, v \in \mathbb{B}_1, r \in \mathbb{B}_{32}, s \in \mathbb{B}_{32}) \quad \equiv \quad p_u \in \mathbb{B}_{64}$$

Where p_u is the public key, assumed to be a byte array of size 64 (formed from the concatenation of two positive integers each $< 2^{256}$) and p_r is the private key, a byte array of size 32 (or a single positive integer in the aforementioned range). It is assumed that v is the 'recovery id', a 1 byte value specifying the sign and finiteness of the curve point; this value is in the range of $[27, 30]$, however we declare the upper two possibilities, representing infinite values, invalid.

---

[4]For the former citation, refer to section 6.2 for ECDSAPUBKEY, and section 7 for ECDSASIGN and ECDSARECOVER.

We declare that a signature is invalid unless all the following conditions are true:

$$(211) \qquad\qquad 0 < r < \texttt{secp256k1n}$$

$$(212) \qquad\qquad 0 < s < \texttt{secp256k1n} \div 2 + 1$$

$$(213) \qquad\qquad v \in \{27, 28\}$$

where:

$$(214) \qquad \texttt{secp256k1n} = 115792089237316195423570985008687907852837564279074904382605163141518161494337$$

For a given private key, p_r, the Ethereum address $A(p_r)$ (a 160-bit value) to which it corresponds is defined as the right most 160-bits of the Keccak hash of the corresponding ECDSA public key:

$$(215) \qquad A(p_r) = \mathcal{B}_{96..255}\big(\texttt{KEC}\big(\texttt{ECDSAPUBKEY}(p_r)\big)\big)$$

The message hash, $h(T)$, to be signed is the Keccak hash of the transaction without the latter three signature components, formally described as $T_\mathbf{r}$, $T_\mathbf{s}$ and $T_\mathbf{w}$:

$$(216) \qquad L_S(T) \quad\equiv\quad \begin{cases} (T_\mathbf{n}, T_\mathbf{p}, T_\mathbf{g}, T_\mathbf{t}, T_v, T_\mathbf{i}) & \text{if } T_\mathbf{t} = 0 \\ (T_\mathbf{n}, T_\mathbf{p}, T_\mathbf{g}, T_\mathbf{t}, T_v, T_\mathbf{d}) & \text{otherwise} \end{cases}$$

$$(217) \qquad h(T) \quad\equiv\quad \texttt{KEC}(L_S(T))$$

The signed transaction $G(T, p_r)$ is defined as:

$$(218) \qquad\qquad G(T, p_r) \equiv T \quad \text{except:}$$

$$(219) \qquad\qquad (T_\mathbf{w}, T_r, T_s) = \texttt{ECDSASIGN}(h(T), p_r)$$

We may then define the sender function S of the transaction as:

$$(220) \qquad S(T) \equiv \mathcal{B}_{96..255}\big(\texttt{KEC}\big(\texttt{ECDSARECOVER}(h(T), T_\mathbf{w}, T_r, T_s)\big)\big)$$

The assertion that the sender of a signed transaction equals the address of the signer should be self-evident:

$$(221) \qquad\qquad \forall T : \forall p_r : S(G(T, p_r)) \equiv A(p_r)$$

Appendix G. Fee Schedule

The fee schedule G is a tuple of 31 scalar values corresponding to the relative costs, in gas, of a number of abstract operations that a transaction may effect.

Name	Value	Description*
G_{zero}	0	Nothing is paid for operations of the set W_{zero}.
G_{base}	2	This is the amount of gas to pay for operations of the set W_{base}.
$G_{verylow}$	3	This is the amount of gas to pay for operations of the set $W_{verylow}$.
G_{low}	5	This is the amount of gas to pay for operations of the set W_{low}.
G_{mid}	8	This is the amount of gas to pay for operations of the set W_{mid}.
G_{high}	10	This is the amount of gas to pay for operations of the set W_{high}.
$G_{extcode}$	700	This is the amount of gas to pay for operations of the set $W_{extcode}$.
$G_{balance}$	400	This is the amount of gas to pay for a BALANCE operation.
G_{sload}	200	This is paid for an SLOAD operation.
$G_{jumpdest}$	1	This is paid for a JUMPDEST operation.
G_{sset}	20000	This is paid for an SSTORE operation when the storage value is set to non-zero from zero.
G_{sreset}	5000	This is the amount for an SSTORE operation when the storage value's zeroness remains unchanged or is set to zero.
R_{sclear}	15000	This is the refund given (added into the refund counter) when the storage value is set to zero from non-zero.
$R_{selfdestruct}$	24000	This is the refund given (added into the refund counter) for self-destructing an account.
$G_{selfdestruct}$	5000	This is the amount of gas to pay for a SELFDESTRUCT operation.
G_{create}	32000	This is paid for a CREATE operation.
$G_{codedeposit}$	200	This is paid per byte for a CREATE operation to succeed in placing code into the state.
G_{call}	700	This is paid for a CALL operation.
$G_{callvalue}$	9000	This is paid for a non-zero value transfer as part of the CALL operation.
$G_{callstipend}$	2300	This is a stipend for the called contract subtracted from $G_{callvalue}$ for a non-zero value transfer.
$G_{newaccount}$	25000	This is paid for a CALL or for a SELFDESTRUCT operation which creates an account.
G_{exp}	10	This is a partial payment for an EXP operation.
$G_{expbyte}$	50	This is a partial payment when multiplied by $\lceil \log_{256}(exponent) \rceil$ for the EXP operation.
G_{memory}	3	This is paid for every additional word when expanding memory.
$G_{txcreate}$	32000	This is paid by all contract-creating transactions after the *Homestead transition*.
$G_{txdatazero}$	4	This is paid for every zero byte of data or code for a transaction.
$G_{txdatanonzero}$	68	This is paid for every non-zero byte of data or code for a transaction.
$G_{Gtransaction}$	21000	This is paid for every transaction.
G_{log}	375	This is a partial payment for a LOG operation.
$G_{logdata}$	8	This is paid for each byte in a LOG operation's data.
$G_{logtopic}$	375	This is paid for each topic of a LOG operation.
G_{sha3}	30	This is paid for each SHA3 operation.
$G_{sha3word}$	6	This is paid for each word (rounded up) for input data to a SHA3 operation.
G_{copy}	3	This is a partial payment for *COPY operations, multiplied by the number of words copied, rounded up.
$G_{blockhash}$	20	This is a payment for a BLOCKHASH operation.

Appendix H. Virtual Machine Specification

When interpreting 256-bit binary values as integers, the representation is big-endian.

When a 256-bit machine datum is converted to and from a 160-bit address or hash, the rightwards (low-order for BE) 20 bytes are used and the left most 12 are discarded or filled with zeroes, thus the integer values (when the bytes are interpreted as big-endian) are equivalent.

H.1. **Gas Cost.** The general gas cost function, C, is defined as:

(222)

$$C(\boldsymbol{\sigma}, \boldsymbol{\mu}, I) \equiv C_{mem}(\boldsymbol{\mu_i'}) - C_{mem}(\boldsymbol{\mu_i}) + \begin{cases} C_{\text{SSTORE}}(\boldsymbol{\sigma}, \boldsymbol{\mu}) & \text{if} \quad w = \text{SSTORE} \\ G_{\text{exp}} & \text{if} \quad w = \text{EXP} \wedge \boldsymbol{\mu_s}[1] = 0 \\ G_{\text{exp}} + G_{\text{expbyte}} \times (1 + \lfloor \log_{256}(\boldsymbol{\mu_s}[1]) \rfloor) & \text{if} \quad w = \text{EXP} \wedge \boldsymbol{\mu_s}[1] > 0 \\ G_{\text{verylow}} + G_{\text{copy}} \times \lceil \boldsymbol{\mu_s}[2] \div 32 \rceil & \text{if} \quad w = \text{CALLDATACOPY} \vee \text{CODECOPY} \\ G_{\text{extcode}} + G_{\text{copy}} \times \lceil \boldsymbol{\mu_s}[3] \div 32 \rceil & \text{if} \quad w = \text{EXTCODECOPY} \\ G_{\text{log}} + G_{\text{logdata}} \times \boldsymbol{\mu_s}[1] & \text{if} \quad w = \text{LOG0} \\ G_{\text{log}} + G_{\text{logdata}} \times \boldsymbol{\mu_s}[1] + G_{\text{logtopic}} & \text{if} \quad w = \text{LOG1} \\ G_{\text{log}} + G_{\text{logdata}} \times \boldsymbol{\mu_s}[1] + 2G_{\text{logtopic}} & \text{if} \quad w = \text{LOG2} \\ G_{\text{log}} + G_{\text{logdata}} \times \boldsymbol{\mu_s}[1] + 3G_{\text{logtopic}} & \text{if} \quad w = \text{LOG3} \\ G_{\text{log}} + G_{\text{logdata}} \times \boldsymbol{\mu_s}[1] + 4G_{\text{logtopic}} & \text{if} \quad w = \text{LOG4} \\ C_{\text{CALL}}(\boldsymbol{\sigma}, \boldsymbol{\mu}) & \text{if} \quad w = \text{CALL} \vee \text{CALLCODE} \vee \text{DELEGATECALL} \\ C_{\text{SELFDESTRUCT}}(\boldsymbol{\sigma}, \boldsymbol{\mu}) & \text{if} \quad w = \text{SELFDESTRUCT} \\ G_{\text{create}} & \text{if} \quad w = \text{CREATE} \\ G_{\text{sha3}} + G_{\text{sha3word}} \lceil \mathbf{s}[1] \div 32 \rceil & \text{if} \quad w = \text{SHA3} \\ G_{\text{jumpdest}} & \text{if} \quad w = \text{JUMPDEST} \\ G_{\text{sload}} & \text{if} \quad w = \text{SLOAD} \\ G_{\text{zero}} & \text{if} \quad w \in W_{\text{zero}} \\ G_{\text{base}} & \text{if} \quad w \in W_{\text{base}} \\ G_{\text{verylow}} & \text{if} \quad w \in W_{\text{verylow}} \\ G_{\text{low}} & \text{if} \quad w \in W_{\text{low}} \\ G_{\text{mid}} & \text{if} \quad w \in W_{\text{mid}} \\ G_{\text{high}} & \text{if} \quad w \in W_{\text{high}} \\ G_{\text{extcode}} & \text{if} \quad w \in W_{\text{extcode}} \\ G_{\text{balance}} & \text{if} \quad w = \text{BALANCE} \\ G_{\text{blockhash}} & \text{if} \quad w = \text{BLOCKHASH} \end{cases}$$

(223)
$$w \equiv \begin{cases} I_{\mathbf{b}}[\boldsymbol{\mu_{pc}}] & \text{if} \quad \boldsymbol{\mu_{pc}} < \|I_{\mathbf{b}}\| \\ \text{STOP} & \text{otherwise} \end{cases}$$

where:

(224)
$$C_{mem}(a) \equiv G_{\text{memory}} \cdot a + \left\lfloor \frac{a^2}{512} \right\rfloor$$

with C_{CALL}, $C_{\text{SELFDESTRUCT}}$ and C_{SSTORE} as specified in the appropriate section below. We define the following subsets of instructions:

$W_{\text{zero}} = \{\text{STOP, RETURN}\}$

$W_{\text{base}} = \{\text{ADDRESS, ORIGIN, CALLER, CALLVALUE, CALLDATASIZE, CODESIZE, GASPRICE, COINBASE, TIMESTAMP, NUMBER, DIFFICULTY, GASLIMIT, POP, PC, MSIZE, GAS}\}$

$W_{\text{verylow}} = \{\text{ADD, SUB, NOT, LT, GT, SLT, SGT, EQ, ISZERO, AND, OR, XOR, BYTE, CALLDATALOAD, MLOAD, MSTORE, MSTORE8, PUSH*, DUP*, SWAP*}\}$

$W_{\text{low}} = \{\text{MUL, DIV, SDIV, MOD, SMOD, SIGNEXTEND}\}$

$W_{\text{mid}} = \{\text{ADDMOD, MULMOD, JUMP}\}$

$W_{\text{high}} = \{\text{JUMPI}\}$

$W_{\text{extcode}} = \{\text{EXTCODESIZE}\}$

Note the memory cost component, given as the product of G_{memory} and the maximum of 0 & the ceiling of the number of words in size that the memory must be over the current number of words, $\boldsymbol{\mu_i}$ in order that all accesses reference valid memory whether for read or write. Such accesses must be for non-zero number of bytes.

Referencing a zero length range (e.g. by attempting to pass it as the input range to a CALL) does not require memory to be extended to the beginning of the range. $\boldsymbol{\mu_i'}$ is defined as this new maximum number of words of active memory; special-cases are given where these two are not equal.

Note also that C_{mem} is the memory cost function (the expansion function being the difference between the cost before and after). It is a polynomial, with the higher-order coefficient divided and floored, and thus linear up to 724B of memory used, after which it costs substantially more.

While defining the instruction set, we defined the memory-expansion for range function, M, thus:

(225)
$$M(s, f, l) \equiv \begin{cases} s & \text{if} \quad l = 0 \\ \max(s, \lceil (f + l) \div 32 \rceil) & \text{otherwise} \end{cases}$$

Another useful function is "all but one 64th" function L defined as:

$$(226) \qquad\qquad\qquad L(n) \equiv n - \lfloor n/64 \rfloor$$

H.2. **Instruction Set.** As previously specified in section 9, these definitions take place in the final context there. In particular we assume O is the EVM state-progression function and define the terms pertaining to the next cycle's state $(\boldsymbol{\sigma}', \boldsymbol{\mu}')$ such that:

$$(227) \qquad\qquad O(\boldsymbol{\sigma}, \boldsymbol{\mu}, A, I) \equiv (\boldsymbol{\sigma}', \boldsymbol{\mu}', A', I) \quad \text{with exceptions, as noted}$$

Here given are the various exceptions to the state transition rules given in section 9 specified for each instruction, together with the additional instruction-specific definitions of J and C. For each instruction, also specified is α, the additional items placed on the stack and δ, the items removed from the stack, as defined in section 9.

<div align="center">

0s: Stop and Arithmetic Operations

All arithmetic is modulo 2^{256} unless otherwise noted. The zero-th power of zero 0^0 is defined to be one.

</div>

Value	Mnemonic	δ	α	Description
0x00	STOP	0	0	This operation halts execution, outputting the empty sequence as per equation 132.
0x01	ADD	2	1	This is the addition operation. $\boldsymbol{\mu}'_{\mathbf{s}}[0] \equiv \boldsymbol{\mu}_{\mathbf{s}}[0] + \boldsymbol{\mu}_{\mathbf{s}}[1]$
0x02	MUL	2	1	This is the multiplication operation. $\boldsymbol{\mu}'_{\mathbf{s}}[0] \equiv \boldsymbol{\mu}_{\mathbf{s}}[0] \times \boldsymbol{\mu}_{\mathbf{s}}[1]$
0x03	SUB	2	1	This is the subtraction operation. $\boldsymbol{\mu}'_{\mathbf{s}}[0] \equiv \boldsymbol{\mu}_{\mathbf{s}}[0] - \boldsymbol{\mu}_{\mathbf{s}}[1]$
0x04	DIV	2	1	This is the integer division operation. $\boldsymbol{\mu}'_{\mathbf{s}}[0] \equiv \begin{cases} 0 & \text{if } \boldsymbol{\mu}_{\mathbf{s}}[1] = 0 \\ \lfloor \boldsymbol{\mu}_{\mathbf{s}}[0] \div \boldsymbol{\mu}_{\mathbf{s}}[1] \rfloor & \text{otherwise} \end{cases}$
0x05	SDIV	2	1	This is the signed integer division operation (truncated). $\boldsymbol{\mu}'_{\mathbf{s}}[0] \equiv \begin{cases} 0 & \text{if } \boldsymbol{\mu}_{\mathbf{s}}[1] = 0 \\ -2^{255} & \text{if } \boldsymbol{\mu}_{\mathbf{s}}[0] = -2^{255} \wedge \boldsymbol{\mu}_{\mathbf{s}}[1] = -1 \\ \mathbf{sgn}(\boldsymbol{\mu}_{\mathbf{s}}[0] \div \boldsymbol{\mu}_{\mathbf{s}}[1]) \lfloor |\boldsymbol{\mu}_{\mathbf{s}}[0] \div \boldsymbol{\mu}_{\mathbf{s}}[1]| \rfloor & \text{otherwise} \end{cases}$ Where all values are treated as two's complement signed 256-bit integers. Note the overflow semantic when -2^{255} is negated.
0x06	MOD	2	1	This is the modulo remainder operation. $\boldsymbol{\mu}'_{\mathbf{s}}[0] \equiv \begin{cases} 0 & \text{if } \boldsymbol{\mu}_{\mathbf{s}}[1] = 0 \\ \boldsymbol{\mu}_{\mathbf{s}}[0] \bmod \boldsymbol{\mu}_{\mathbf{s}}[1] & \text{otherwise} \end{cases}$
0x07	SMOD	2	1	This is the signed modulo remainder operation. $\boldsymbol{\mu}'_{\mathbf{s}}[0] \equiv \begin{cases} 0 & \text{if } \boldsymbol{\mu}_{\mathbf{s}}[1] = 0 \\ \mathbf{sgn}(\boldsymbol{\mu}_{\mathbf{s}}[0])(|\boldsymbol{\mu}_{\mathbf{s}}[0]| \bmod |\boldsymbol{\mu}_{\mathbf{s}}[1]|) & \text{otherwise} \end{cases}$ Where all values are treated as two's complement signed 256-bit integers.
0x08	ADDMOD	3	1	This is the modulo addition operation. $\boldsymbol{\mu}'_{\mathbf{s}}[0] \equiv \begin{cases} 0 & \text{if } \boldsymbol{\mu}_{\mathbf{s}}[2] = 0 \\ (\boldsymbol{\mu}_{\mathbf{s}}[0] + \boldsymbol{\mu}_{\mathbf{s}}[1]) \bmod \boldsymbol{\mu}_{\mathbf{s}}[2] & \text{otherwise} \end{cases}$ All intermediate calculations of this operation are not subject to the 2^{256} modulo.
0x09	MULMOD	3	1	This is the modulo multiplication operation. $\boldsymbol{\mu}'_{\mathbf{s}}[0] \equiv \begin{cases} 0 & \text{if } \boldsymbol{\mu}_{\mathbf{s}}[2] = 0 \\ (\boldsymbol{\mu}_{\mathbf{s}}[0] \times \boldsymbol{\mu}_{\mathbf{s}}[1]) \bmod \boldsymbol{\mu}_{\mathbf{s}}[2] & \text{otherwise} \end{cases}$ All intermediate calculations of this operation are not subject to the 2^{256} modulo.
0x0a	EXP	2	1	This is the exponential operation. $\boldsymbol{\mu}'_{\mathbf{s}}[0] \equiv \boldsymbol{\mu}_{\mathbf{s}}[0]^{\boldsymbol{\mu}_{\mathbf{s}}[1]}$
0x0b	SIGNEXTEND	2	1	Extend the length of a two's complement signed integer. $\forall i \in [0..255] : \boldsymbol{\mu}'_{\mathbf{s}}[0]_{\mathbf{i}} \equiv \begin{cases} \boldsymbol{\mu}_{\mathbf{s}}[1]_{\mathbf{t}} & \text{if } i \leqslant t \quad \text{where } t = 256 - 8(\boldsymbol{\mu}_{\mathbf{s}}[0] + 1) \\ \boldsymbol{\mu}_{\mathbf{s}}[1]_{\mathbf{i}} & \text{otherwise} \end{cases}$

$\boldsymbol{\mu}_{\mathbf{s}}[x]_{\mathbf{i}}$ gives the ith bit (counting from zero) of $\boldsymbol{\mu}_{\mathbf{s}}[x]$

10s: Comparison & Bitwise Logic Operations

Value	Mnemonic	δ	α	Description
0x10	LT	2	1	This is the less-than comparison. $$\boldsymbol{\mu}'_\mathbf{s}[0] \equiv \begin{cases} 1 & \text{if} \quad \boldsymbol{\mu}_\mathbf{s}[0] < \boldsymbol{\mu}_\mathbf{s}[1] \\ 0 & \text{otherwise} \end{cases}$$
0x11	GT	2	1	This is the greater-than comparison. $$\boldsymbol{\mu}'_\mathbf{s}[0] \equiv \begin{cases} 1 & \text{if} \quad \boldsymbol{\mu}_\mathbf{s}[0] > \boldsymbol{\mu}_\mathbf{s}[1] \\ 0 & \text{otherwise} \end{cases}$$
0x12	SLT	2	1	This is the signed less-than comparison. $$\boldsymbol{\mu}'_\mathbf{s}[0] \equiv \begin{cases} 1 & \text{if} \quad \boldsymbol{\mu}_\mathbf{s}[0] < \boldsymbol{\mu}_\mathbf{s}[1] \\ 0 & \text{otherwise} \end{cases}$$ Where all values are treated as two's complement signed 256-bit integers.
0x13	SGT	2	1	This is the signed greater-than comparison. $$\boldsymbol{\mu}'_\mathbf{s}[0] \equiv \begin{cases} 1 & \text{if} \quad \boldsymbol{\mu}_\mathbf{s}[0] > \boldsymbol{\mu}_\mathbf{s}[1] \\ 0 & \text{otherwise} \end{cases}$$ Where all values are treated as two's complement signed 256-bit integers.
0x14	EQ	2	1	This is the equality comparison. $$\boldsymbol{\mu}'_\mathbf{s}[0] \equiv \begin{cases} 1 & \text{if} \quad \boldsymbol{\mu}_\mathbf{s}[0] = \boldsymbol{\mu}_\mathbf{s}[1] \\ 0 & \text{otherwise} \end{cases}$$
0x15	ISZERO	1	1	This is the logical negation operation, also called the logical complement or the NOT operation. $$\boldsymbol{\mu}'_\mathbf{s}[0] \equiv \begin{cases} 1 & \text{if} \quad \boldsymbol{\mu}_\mathbf{s}[0] = 0 \\ 0 & \text{otherwise} \end{cases}$$
0x16	AND	2	1	This is the bitwise AND operation. $\forall i \in [0..255] : \boldsymbol{\mu}'_\mathbf{s}[0]_\mathtt{i} \equiv \boldsymbol{\mu}_\mathbf{s}[0]_\mathtt{i} \wedge \boldsymbol{\mu}_\mathbf{s}[1]_\mathtt{i}$
0x17	OR	2	1	This is the bitwise OR operation. $\forall i \in [0..255] : \boldsymbol{\mu}'_\mathbf{s}[0]_\mathtt{i} \equiv \boldsymbol{\mu}_\mathbf{s}[0]_\mathtt{i} \vee \boldsymbol{\mu}_\mathbf{s}[1]_\mathtt{i}$
0x18	XOR	2	1	This is the bitwise XOR operation. $\forall i \in [0..255] : \boldsymbol{\mu}'_\mathbf{s}[0]_\mathtt{i} \equiv \boldsymbol{\mu}_\mathbf{s}[0]_\mathtt{i} \oplus \boldsymbol{\mu}_\mathbf{s}[1]_\mathtt{i}$
0x19	NOT	1	1	This is the bitwise NOT operation. $$\forall i \in [0..255] : \boldsymbol{\mu}'_\mathbf{s}[0]_\mathtt{i} \equiv \begin{cases} 1 & \text{if} \quad \boldsymbol{\mu}_\mathbf{s}[0]_\mathtt{i} = 0 \\ 0 & \text{otherwise} \end{cases}$$
0x1a	BYTE	2	1	Retrieve a single byte from a word. $$\forall i \in [0..255] : \boldsymbol{\mu}'_\mathbf{s}[0]_\mathtt{i} \equiv \begin{cases} \boldsymbol{\mu}_\mathbf{s}[1]_{(i+8\boldsymbol{\mu}_\mathbf{s}[0])} & \text{if} \quad i < 8 \wedge \boldsymbol{\mu}_\mathbf{s}[0] < 32 \\ 0 & \text{otherwise} \end{cases}$$ For the Nth byte, we count from the left (i.e. $N = 0$ would be the most significant in big endian).

20s: SHA3

Value	Mnemonic	δ	α	Description
0x20	SHA3	2	1	Compute a Keccak-256 hash. $\boldsymbol{\mu}'_\mathbf{s}[0] \equiv \texttt{Keccak}(\boldsymbol{\mu}_\mathbf{m}[\boldsymbol{\mu}_\mathbf{s}[0] \ldots (\boldsymbol{\mu}_\mathbf{s}[0] + \boldsymbol{\mu}_\mathbf{s}[1] - 1)])$ $\boldsymbol{\mu}'_\mathbf{i} \equiv M(\boldsymbol{\mu}_\mathbf{i}, \boldsymbol{\mu}_\mathbf{s}[0], \boldsymbol{\mu}_\mathbf{s}[1])$

30s: Environmental Information

Value	Mnemonic	δ	α	Description
0x30	ADDRESS	0	1	Get the address of the currently executing account. $\boldsymbol{\mu}'_{\mathbf{s}}[0] \equiv I_{\mathbf{a}}$
0x31	BALANCE	1	1	Get the balance of the given account. $\boldsymbol{\mu}'_{\mathbf{s}}[0] \equiv \begin{cases} \boldsymbol{\sigma}[\boldsymbol{\mu}_{\mathbf{s}}[0]]_{\mathbf{b}} & \text{if } \boldsymbol{\sigma}[\boldsymbol{\mu}_{\mathbf{s}}[0] \mod 2^{160}] \neq \varnothing \\ 0 & \text{otherwise} \end{cases}$
0x32	ORIGIN	0	1	Get the address that the execution originated from. $\boldsymbol{\mu}'_{\mathbf{s}}[0] \equiv I_{\mathbf{o}}$ This is the sender of the original transaction; it is never an account with non-empty associated code.
0x33	CALLER	0	1	Get the caller address. $\boldsymbol{\mu}'_{\mathbf{s}}[0] \equiv I_{\mathbf{s}}$ This is the address of the account that is directly responsible for this execution.
0x34	CALLVALUE	0	1	Get the deposited value by the instruction/transaction responsible for this execution. $\boldsymbol{\mu}'_{\mathbf{s}}[0] \equiv I_{\mathbf{v}}$
0x35	CALLDATALOAD	1	1	Get the input data of the current environment. $\boldsymbol{\mu}'_{\mathbf{s}}[0] \equiv I_{\mathbf{d}}[\boldsymbol{\mu}_{\mathbf{s}}[0] \dots (\boldsymbol{\mu}_{\mathbf{s}}[0] + 31)] \quad \text{with} \quad I_{\mathbf{d}}[x] = 0 \quad \text{if} \quad x \geqslant \|I_{\mathbf{d}}\|$ This pertains to the input data passed with the message call instruction or transaction.
0x36	CALLDATASIZE	0	1	Get the size of the input data in the current environment. $\boldsymbol{\mu}'_{\mathbf{s}}[0] \equiv \|I_{\mathbf{d}}\|$ This pertains to the input data passed with the message call instruction or transaction.
0x37	CALLDATACOPY	3	0	Copy the input data in the current environment to memory. $\forall_{i \in \{0 \dots \boldsymbol{\mu}_{\mathbf{s}}[2]-1\}} \boldsymbol{\mu}'_{\mathbf{m}}[\boldsymbol{\mu}_{\mathbf{s}}[0] + i] \equiv \begin{cases} I_{\mathbf{d}}[\boldsymbol{\mu}_{\mathbf{s}}[1] + i] & \text{if } \boldsymbol{\mu}_{\mathbf{s}}[1] + i < \|I_{\mathbf{d}}\| \\ 0 & \text{otherwise} \end{cases}$ The additions in $\boldsymbol{\mu}_{\mathbf{s}}[1] + i$ are not subject to the 2^{256} modulo. $\boldsymbol{\mu}'_{\mathbf{i}} \equiv M(\boldsymbol{\mu}_{\mathbf{i}}, \boldsymbol{\mu}_{\mathbf{s}}[0], \boldsymbol{\mu}_{\mathbf{s}}[2])$ This pertains to the input data passed with the message call instruction or transaction.
0x38	CODESIZE	0	1	Get the size of code running in the current environment. $\boldsymbol{\mu}'_{\mathbf{s}}[0] \equiv \|I_{\mathbf{b}}\|$
0x39	CODECOPY	3	0	Copy code running in the current environment to memory. $\forall_{i \in \{0 \dots \boldsymbol{\mu}_{\mathbf{s}}[2]-1\}} \boldsymbol{\mu}'_{\mathbf{m}}[\boldsymbol{\mu}_{\mathbf{s}}[0] + i] \equiv \begin{cases} I_{\mathbf{b}}[\boldsymbol{\mu}_{\mathbf{s}}[1] + i] & \text{if } \boldsymbol{\mu}_{\mathbf{s}}[1] + i < \|I_{\mathbf{b}}\| \\ \text{STOP} & \text{otherwise} \end{cases}$ $\boldsymbol{\mu}'_{\mathbf{i}} \equiv M(\boldsymbol{\mu}_{\mathbf{i}}, \boldsymbol{\mu}_{\mathbf{s}}[0], \boldsymbol{\mu}_{\mathbf{s}}[2])$ The additions in $\boldsymbol{\mu}_{\mathbf{s}}[1] + i$ are not subject to the 2^{256} modulo.
0x3a	GASPRICE	0	1	Get the price of gas in the current environment. $\boldsymbol{\mu}'_{\mathbf{s}}[0] \equiv I_{\mathbf{p}}$ This is the gas price specified by the originating transaction.
0x3b	EXTCODESIZE	1	1	Get the size of an account's code. $\boldsymbol{\mu}'_{\mathbf{s}}[0] \equiv \|\boldsymbol{\sigma}[\boldsymbol{\mu}_{\mathbf{s}}[0] \mod 2^{160}]_{\mathbf{c}}\|$
0x3c	EXTCODECOPY	4	0	Copy an account's code to memory. $\forall_{i \in \{0 \dots \boldsymbol{\mu}_{\mathbf{s}}[3]-1\}} \boldsymbol{\mu}'_{\mathbf{m}}[\boldsymbol{\mu}_{\mathbf{s}}[1] + i] \equiv \begin{cases} \mathbf{c}[\boldsymbol{\mu}_{\mathbf{s}}[2] + i] & \text{if } \boldsymbol{\mu}_{\mathbf{s}}[2] + i < \|\mathbf{c}\| \\ \text{STOP} & \text{otherwise} \end{cases}$ where $\mathbf{c} \equiv \boldsymbol{\sigma}[\boldsymbol{\mu}_{\mathbf{s}}[0] \mod 2^{160}]_{\mathbf{c}}$ $\boldsymbol{\mu}'_{\mathbf{i}} \equiv M(\boldsymbol{\mu}_{\mathbf{i}}, \boldsymbol{\mu}_{\mathbf{s}}[1], \boldsymbol{\mu}_{\mathbf{s}}[3])$ The additions in $\boldsymbol{\mu}_{\mathbf{s}}[2] + i$ are not subject to the 2^{256} modulo.

40s: Block Information

Value	Mnemonic	δ	α	Description
0x40	BLOCKHASH	1	1	Get the hash of one of the 256 most recent complete blocks. $\boldsymbol{\mu}_\mathbf{s}'[0] \equiv P(I_{H_\mathbf{p}}, \boldsymbol{\mu}_\mathbf{s}[0], 0)$ where P is the hash of a block of a particular number, up to a maximum age. 0 is left on the stack if the looked for block number is greater than the current block number or more than 256 blocks behind the current block. $$P(h, n, a) \equiv \begin{cases} 0 & \text{if} \quad n > H_\mathbf{i} \vee a = 256 \vee h = 0 \\ h & \text{if} \quad n = H_\mathbf{i} \\ P(H_\mathbf{p}, n, a+1) & \text{otherwise} \end{cases}$$ and we assert that the reason why the header H can be determined is because its hash is the parent hash in the block following it.
0x41	COINBASE	0	1	Get the block's beneficiary address. $\boldsymbol{\mu}_\mathbf{s}'[0] \equiv I_{H_\mathbf{c}}$
0x42	TIMESTAMP	0	1	Get the block's timestamp. $\boldsymbol{\mu}_\mathbf{s}'[0] \equiv I_{H_\mathbf{s}}$
0x43	NUMBER	0	1	Get the block's number. $\boldsymbol{\mu}_\mathbf{s}'[0] \equiv I_{H_\mathbf{i}}$
0x44	DIFFICULTY	0	1	Get the block's difficulty. $\boldsymbol{\mu}_\mathbf{s}'[0] \equiv I_{H_\mathbf{d}}$
0x45	GASLIMIT	0	1	Get the block's gas limit. $\boldsymbol{\mu}_\mathbf{s}'[0] \equiv I_{H_\mathbf{l}}$

50s: Stack, Memory, Storage and Flow Operations

Value	Mnemonic	δ	α	Description
0x50	POP	1	0	Remove the top, first item from the stack.
0x51	MLOAD	1	1	Load the first word from memory. $\boldsymbol{\mu'_{\mathbf{s}}}[0] \equiv \boldsymbol{\mu_{\mathbf{m}}}[\boldsymbol{\mu_{\mathbf{s}}}[0]\ldots(\boldsymbol{\mu_{\mathbf{s}}}[0]+31)]$ $\boldsymbol{\mu'_{\mathbf{i}}} \equiv \max(\boldsymbol{\mu_{\mathbf{i}}}, \lceil(\boldsymbol{\mu_{\mathbf{s}}}[0]+32) \div 32\rceil)$ The addition in the calculation of $\boldsymbol{\mu'_{\mathbf{i}}}$ is not subject to the 2^{256} modulo.
0x52	MSTORE	2	0	Save the first word (which is the first item in the stack) to memory. $\boldsymbol{\mu'_{\mathbf{m}}}[\boldsymbol{\mu_{\mathbf{s}}}[0]\ldots(\boldsymbol{\mu_{\mathbf{s}}}[0]+31)] \equiv \boldsymbol{\mu_{\mathbf{s}}}[1]$ $\boldsymbol{\mu'_{\mathbf{i}}} \equiv \max(\boldsymbol{\mu_{\mathbf{i}}}, \lceil(\boldsymbol{\mu_{\mathbf{s}}}[0]+32) \div 32\rceil)$ The addition in the calculation of $\boldsymbol{\mu'_{\mathbf{i}}}$ is not subject to the 2^{256} modulo.
0x53	MSTORE8	2	0	Save the first byte to memory. $\boldsymbol{\mu'_{\mathbf{m}}}[\boldsymbol{\mu_{\mathbf{s}}}[0]] \equiv (\boldsymbol{\mu_{\mathbf{s}}}[1] \bmod 256)$ $\boldsymbol{\mu'_{\mathbf{i}}} \equiv \max(\boldsymbol{\mu_{\mathbf{i}}}, \lceil(\boldsymbol{\mu_{\mathbf{s}}}[0]+1) \div 32\rceil)$ The addition in the calculation of $\boldsymbol{\mu'_{\mathbf{i}}}$ is not subject to the 2^{256} modulo.
0x54	SLOAD	1	1	Load the first word from storage. $\boldsymbol{\mu'_{\mathbf{s}}}[0] \equiv \boldsymbol{\sigma}[I_{\mathbf{a}}]_{\mathbf{s}}[\boldsymbol{\mu_{\mathbf{s}}}[0]]$
0x55	SSTORE	2	0	Save the first word to storage. $\boldsymbol{\sigma'}[I_{\mathbf{a}}]_{\mathbf{s}}[\boldsymbol{\mu_{\mathbf{s}}}[0]] \equiv \boldsymbol{\mu_{\mathbf{s}}}[1]$ $C_{\text{SSTORE}}(\boldsymbol{\sigma}, \boldsymbol{\mu}) \equiv \begin{cases} G_{\text{sset}} & \text{if } \boldsymbol{\mu_{\mathbf{s}}}[1] \neq 0 \wedge \boldsymbol{\sigma}[I_{\mathbf{a}}]_{\mathbf{s}}[\boldsymbol{\mu_{\mathbf{s}}}[0]] = 0 \\ G_{\text{sreset}} & \text{otherwise} \end{cases}$ $A'_r \equiv A_r + \begin{cases} R_{\text{sclear}} & \text{if } \boldsymbol{\mu_{\mathbf{s}}}[1] = 0 \wedge \boldsymbol{\sigma}[I_{\mathbf{a}}]_{\mathbf{s}}[\boldsymbol{\mu_{\mathbf{s}}}[0]] \neq 0 \\ 0 & \text{otherwise} \end{cases}$
0x56	JUMP	1	0	Alter the program counter. $J_{\text{JUMP}}(\boldsymbol{\mu}) \equiv \boldsymbol{\mu_{\mathbf{s}}}[0]$ This has the effect of writing said value to $\boldsymbol{\mu}_{pc}$. See equation 138.
0x57	JUMPI	2	0	Conditionally alter the program counter. $J_{\text{JUMPI}}(\boldsymbol{\mu}) \equiv \begin{cases} \boldsymbol{\mu_{\mathbf{s}}}[0] & \text{if } \boldsymbol{\mu_{\mathbf{s}}}[1] \neq 0 \\ \boldsymbol{\mu}_{pc} + 1 & \text{otherwise} \end{cases}$ This has the effect of writing said value to $\boldsymbol{\mu}_{pc}$. See section 138.
0x58	PC	0	1	Get the value of the program counter *prior* to the increment corresponding to this instruction. $\boldsymbol{\mu'_{\mathbf{s}}}[0] \equiv \boldsymbol{\mu}_{pc}$
0x59	MSIZE	0	1	Get the size of active memory in bytes. $\boldsymbol{\mu'_{\mathbf{s}}}[0] \equiv 32\boldsymbol{\mu}_i$
0x5a	GAS	0	1	Get the amount of available gas, including the corresponding reduction for the cost of this instruction. $\boldsymbol{\mu'_{\mathbf{s}}}[0] \equiv \boldsymbol{\mu}_g$
0x5b	JUMPDEST	0	0	Mark a valid destination for jumps. This operation has no effect on the machine state during execution.

60s & 70s: Push Operations

Value	Mnemonic	δ	α	Description
0x60	PUSH1	0	1	Place a 1 byte item on the stack. $\boldsymbol{\mu}'_{\mathbf{s}}[0] \equiv c(\boldsymbol{\mu}_{pc} + 1)$ where $\quad c(x) \equiv \begin{cases} I_{\mathbf{b}}[x] & \text{if} \quad x < \|I_{\mathbf{b}}\| \\ 0 & \text{otherwise} \end{cases}$ The bytes are read in line from the program code's bytes array. The function c ensures the bytes default to zero if they extend past the limits. The byte is right-aligned (i.e. it takes the lowest significant place and the last, highest address, which is the big-endian interpretation).
0x61	PUSH2	0	1	Place a 2-byte item on the stack. $\boldsymbol{\mu}'_{\mathbf{s}}[0] \equiv \boldsymbol{c}\big((\boldsymbol{\mu}_{pc} + 1)\ldots(\boldsymbol{\mu}_{pc} + 2)\big)$ with $\boldsymbol{c}(\boldsymbol{x}) \equiv (c(\boldsymbol{x}_0), ..., c(\boldsymbol{x}_{\|x\|-1}))$ with c as defined as above. Similarly, the bytes are right-aligned (i.e. they take the lowest significant place and the last, highest address, which is the big-endian interpretation).
$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$
0x7f	PUSH32	0	1	Place 32-byte (full word) item on stack. $\boldsymbol{\mu}'_{\mathbf{s}}[0] \equiv \boldsymbol{c}\big((\boldsymbol{\mu}_{pc} + 1)\ldots(\boldsymbol{\mu}_{pc} + 32)\big)$ where $\boldsymbol{c}$ is defined as above. Similarly, the bytes are right-aligned (i.e. they take the lowest significant place and the last, highest address, which is the big-endian interpretation).

80s: Duplication Operations

Value	Mnemonic	δ	α	Description
0x80	DUP1	1	2	Duplicate the 1st stack item. $\boldsymbol{\mu}'_{\mathbf{s}}[0] \equiv \boldsymbol{\mu}_{\mathbf{s}}[0]$
0x81	DUP2	2	3	Duplicate the 2nd stack item. $\boldsymbol{\mu}'_{\mathbf{s}}[0] \equiv \boldsymbol{\mu}_{\mathbf{s}}[1]$
$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$
0x8f	DUP16	16	17	Duplicate the 16th stack item. $\boldsymbol{\mu}'_{\mathbf{s}}[0] \equiv \boldsymbol{\mu}_{\mathbf{s}}[15]$

90s: Exchange Operations

Value	Mnemonic	δ	α	Description
0x90	SWAP1	2	2	Exchange the 1st and the 2nd stack items. $\boldsymbol{\mu}'_{\mathbf{s}}[0] \equiv \boldsymbol{\mu}_{\mathbf{s}}[1]$ $\boldsymbol{\mu}'_{\mathbf{s}}[1] \equiv \boldsymbol{\mu}_{\mathbf{s}}[0]$
0x91	SWAP2	3	3	Exchange the 1st and the 3rd stack items. $\boldsymbol{\mu}'_{\mathbf{s}}[0] \equiv \boldsymbol{\mu}_{\mathbf{s}}[2]$ $\boldsymbol{\mu}'_{\mathbf{s}}[2] \equiv \boldsymbol{\mu}_{\mathbf{s}}[0]$
$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$
0x9f	SWAP16	17	17	Exchange the 1st and the 17th stack items. $\boldsymbol{\mu}'_{\mathbf{s}}[0] \equiv \boldsymbol{\mu}_{\mathbf{s}}[16]$ $\boldsymbol{\mu}'_{\mathbf{s}}[16] \equiv \boldsymbol{\mu}_{\mathbf{s}}[0]$

a0s: Logging Operations

For all logging operations, the state change is to append an additional log entry on to the substate's log series:

$A'_\mathbf{l} \equiv A_\mathbf{l} \cdot (I_\mathbf{a}, \mathbf{t}, \boldsymbol{\mu}_\mathbf{m}[\boldsymbol{\mu}_\mathbf{s}[0] \ldots (\boldsymbol{\mu}_\mathbf{s}[0] + \boldsymbol{\mu}_\mathbf{s}[1] - 1)])$

and to update the memory consumption counter:

$\boldsymbol{\mu}'_\mathbf{i} \equiv M(\boldsymbol{\mu}_\mathbf{i}, \boldsymbol{\mu}_\mathbf{s}[0], \boldsymbol{\mu}_\mathbf{s}[1])$

The entry's topic series, $\mathbf{t}$, differs accordingly:

Value	Mnemonic	δ	α	Description
0xa0	LOG0	2	0	Append the log record with no topics. $\mathbf{t} \equiv ()$
0xa1	LOG1	3	0	Append the log record with one topic. $\mathbf{t} \equiv (\boldsymbol{\mu}_\mathbf{s}[2])$
$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$
0xa4	LOG4	6	0	Append the log record with four topics. $\mathbf{t} \equiv (\boldsymbol{\mu}_\mathbf{s}[2], \boldsymbol{\mu}_\mathbf{s}[3], \boldsymbol{\mu}_\mathbf{s}[4], \boldsymbol{\mu}_\mathbf{s}[5])$

f0s: System operations

Value	Mnemonic	δ	α	Description
0xf0	CREATE	3	1	Create a new account with associated code.

$\mathbf{i} \equiv \boldsymbol{\mu}_\mathbf{m}[\boldsymbol{\mu}_\mathbf{s}[1] \dots (\boldsymbol{\mu}_\mathbf{s}[1] + \boldsymbol{\mu}_\mathbf{s}[2] - 1)]$

$$(\boldsymbol{\sigma}', \boldsymbol{\mu}'_\mathbf{g}, A^+) \equiv \begin{cases} \Lambda(\boldsymbol{\sigma}^*, I_\mathbf{a}, I_\mathbf{o}, L(\boldsymbol{\mu}_\mathbf{g}), I_\mathbf{p}, \boldsymbol{\mu}_\mathbf{s}[0], \mathbf{i}, I_\mathbf{e} + 1) & \text{if} \quad \boldsymbol{\mu}_\mathbf{s}[0] \leqslant \boldsymbol{\sigma}[I_\mathbf{a}]_\mathbf{b} \wedge I_\mathbf{e} < 1024 \\ (\boldsymbol{\sigma}, \boldsymbol{\mu}_\mathbf{g}, \varnothing) & \text{otherwise} \end{cases}$$

$\boldsymbol{\sigma}^* \equiv \boldsymbol{\sigma}$ except $\boldsymbol{\sigma}^*[I_\mathbf{a}]_\mathbf{n} = \boldsymbol{\sigma}[I_\mathbf{a}]_\mathbf{n} + 1$

$A' \equiv A \uplus A^+$ which implies: $A'_\mathbf{s} \equiv A_\mathbf{s} \cup A^+_\mathbf{s} \quad \wedge \quad A'_\mathbf{l} \equiv A_\mathbf{l} \cdot A^+_\mathbf{l} \quad \wedge \quad A'_\mathbf{r} \equiv A_\mathbf{r} + A^+_\mathbf{r}$

$\boldsymbol{\mu}'_\mathbf{s}[0] \equiv x$

where $x = 0$ if the code execution for this operation failed due to an exceptional halting:
$Z(\boldsymbol{\sigma}^*, \boldsymbol{\mu}, I) = \top$ or $I_\mathbf{e} = 1024$
(the maximum call depth limit is reached) or $\boldsymbol{\mu}_\mathbf{s}[0] > \boldsymbol{\sigma}[I_\mathbf{a}]_\mathbf{b}$ (the balance of the caller
is too low to fulfil the value transfer); and otherwise $x = A(I_\mathbf{a}, \boldsymbol{\sigma}[I_\mathbf{a}]_\mathbf{n})$, the address of the
newly created account.

$\boldsymbol{\mu}'_\mathbf{i} \equiv M(\boldsymbol{\mu}_\mathbf{i}, \boldsymbol{\mu}_\mathbf{s}[1], \boldsymbol{\mu}_\mathbf{s}[2])$

Thus the operand order is: value, input offset, input size.

0xf1	CALL	7	1	Message-call into an account.

$\mathbf{i} \equiv \boldsymbol{\mu}_\mathbf{m}[\boldsymbol{\mu}_\mathbf{s}[3] \dots (\boldsymbol{\mu}_\mathbf{s}[3] + \boldsymbol{\mu}_\mathbf{s}[4] - 1)]$

$$(\boldsymbol{\sigma}', g', A^+, \mathbf{o}) \equiv \begin{cases} \Theta(\boldsymbol{\sigma}, I_\mathbf{a}, I_\mathbf{o}, t, t, & \text{if} \quad \boldsymbol{\mu}_\mathbf{s}[2] \leqslant \boldsymbol{\sigma}[I_\mathbf{a}]_\mathbf{b} \wedge \\ \quad C_{\text{CALLGAS}}(\boldsymbol{\mu}), I_\mathbf{p}, \boldsymbol{\mu}_\mathbf{s}[2], \boldsymbol{\mu}_\mathbf{s}[2], \mathbf{i}, I_\mathbf{e} + 1) & \quad I_\mathbf{e} < 1024 \\ (\boldsymbol{\sigma}, g, \varnothing, ()) & \text{otherwise} \end{cases}$$

$n \equiv \min(\{\boldsymbol{\mu}_\mathbf{s}[6], |\mathbf{o}|\})$

$\boldsymbol{\mu}'_\mathbf{m}[\boldsymbol{\mu}_\mathbf{s}[5] \dots (\boldsymbol{\mu}_\mathbf{s}[5] + n - 1)] = \mathbf{o}[0 \dots (n - 1)]$

$\boldsymbol{\mu}'_\mathbf{g} \equiv \boldsymbol{\mu}_\mathbf{g} + g'$

$\boldsymbol{\mu}'_\mathbf{s}[0] \equiv x$

$A' \equiv A \uplus A^+$

$t \equiv \boldsymbol{\mu}_\mathbf{s}[1] \mod 2^{160}$

where $x = 0$ if the code execution for this operation failed due to an exceptional halting
$Z(\boldsymbol{\sigma}, \boldsymbol{\mu}, I) = \top$ or if
$\boldsymbol{\mu}_\mathbf{s}[2] > \boldsymbol{\sigma}[I_\mathbf{a}]_\mathbf{b}$ (not enough funds) or $I_\mathbf{e} = 1024$ (call depth limit reached); $x = 1$
otherwise.

$\boldsymbol{\mu}'_\mathbf{i} \equiv M(M(\boldsymbol{\mu}_\mathbf{i}, \boldsymbol{\mu}_\mathbf{s}[3], \boldsymbol{\mu}_\mathbf{s}[4]), \boldsymbol{\mu}_\mathbf{s}[5], \boldsymbol{\mu}_\mathbf{s}[6])$

Thus the operand order is: gas, to, value, in offset, in size, out offset, out size.

$C_{\text{CALL}}(\boldsymbol{\sigma}, \boldsymbol{\mu}) \equiv C_{\text{GASCAP}}(\boldsymbol{\sigma}, \boldsymbol{\mu}) + C_{\text{EXTRA}}(\boldsymbol{\sigma}, \boldsymbol{\mu})$

$$C_{\text{CALLGAS}}(\boldsymbol{\sigma}, \boldsymbol{\mu}) \equiv \begin{cases} C_{\text{GASCAP}}(\boldsymbol{\sigma}, \boldsymbol{\mu}) + G_{\text{callstipend}} & \text{if} \quad \boldsymbol{\mu}_\mathbf{s}[2] \neq 0 \\ C_{\text{GASCAP}}(\boldsymbol{\sigma}, \boldsymbol{\mu}) & \text{otherwise} \end{cases}$$

$$C_{\text{GASCAP}}(\boldsymbol{\sigma}, \boldsymbol{\mu}) \equiv \begin{cases} \min\{L(\boldsymbol{\mu}_\mathbf{g} - C_{\text{EXTRA}}(\boldsymbol{\sigma}, \boldsymbol{\mu})), \boldsymbol{\mu}_\mathbf{s}[0]\} & \text{if} \quad \boldsymbol{\mu}_\mathbf{g} \geq C_{\text{EXTRA}}(\boldsymbol{\sigma}, \boldsymbol{\mu}) \\ \boldsymbol{\mu}_\mathbf{s}[0] & \text{otherwise} \end{cases}$$

$C_{\text{EXTRA}}(\boldsymbol{\sigma}, \boldsymbol{\mu}) \equiv G_{\text{call}} + C_{\text{XFER}}(\boldsymbol{\mu}) + C_{\text{NEW}}(\boldsymbol{\sigma}, \boldsymbol{\mu})$

$$C_{\text{XFER}}(\boldsymbol{\mu}) \equiv \begin{cases} G_{\text{callvalue}} & \text{if} \quad \boldsymbol{\mu}_\mathbf{s}[2] \neq 0 \\ 0 & \text{otherwise} \end{cases}$$

$$C_{\text{NEW}}(\boldsymbol{\sigma}, \boldsymbol{\mu}) \equiv \begin{cases} G_{\text{newaccount}} & \text{if} \quad \boldsymbol{\sigma}[\boldsymbol{\mu}_\mathbf{s}[1] \mod 2^{160}] = \varnothing \\ 0 & \text{otherwise} \end{cases}$$

0xf2	CALLCODE	7	1	Message-call into this account with an alternative account's code.

Exactly equivalent to CALL except:

$$(\boldsymbol{\sigma}', g', A^+, \mathbf{o}) \equiv \begin{cases} \Theta(\boldsymbol{\sigma}^*, I_\mathbf{a}, I_\mathbf{o}, I_\mathbf{a}, t, & \text{if} \quad \boldsymbol{\mu}_\mathbf{s}[2] \leqslant \boldsymbol{\sigma}[I_\mathbf{a}]_\mathbf{b} \wedge \\ \quad C_{\text{CALLGAS}}(\boldsymbol{\mu}), I_\mathbf{p}, \boldsymbol{\mu}_\mathbf{s}[2], \boldsymbol{\mu}_\mathbf{s}[2], \mathbf{i}, I_\mathbf{e} + 1) & \quad I_\mathbf{e} < 1024 \\ (\boldsymbol{\sigma}, g, \varnothing, ()) & \text{otherwise} \end{cases}$$

Note the change in the fourth parameter to the call Θ from the 2nd stack value $\boldsymbol{\mu}_\mathbf{s}[1]$
(as in CALL) to the present address $I_\mathbf{a}$. This means that the recipient is in fact the
same account as at present, simply that the code is overwritten.

0xf3	RETURN	2	0	Halt execution returning output data.

$H_{\text{RETURN}}(\boldsymbol{\mu}) \equiv \boldsymbol{\mu}_\mathbf{m}[\boldsymbol{\mu}_\mathbf{s}[0] \dots (\boldsymbol{\mu}_\mathbf{s}[0] + \boldsymbol{\mu}_\mathbf{s}[1] - 1)]$

This has the effect of halting the execution at this point with output defined.
See equation 132.

$\boldsymbol{\mu}'_\mathbf{i} \equiv M(\boldsymbol{\mu}_\mathbf{i}, \boldsymbol{\mu}_\mathbf{s}[0], \boldsymbol{\mu}_\mathbf{s}[1])$

0xf4	DELEGATECALL	6	1	Message-call into this account with an alternative account's code, but persisting the current values for *sender* and *value*. Compared with CALL, DELEGATECALL takes one fewer arguments. The omitted argument is $\boldsymbol{\mu_s}[2]$. As a result, $\boldsymbol{\mu_s}[3]$, $\boldsymbol{\mu_s}[4]$, $\boldsymbol{\mu_s}[5]$ and $\boldsymbol{\mu_s}[6]$ in the definition of CALL should respectively be replaced with $\boldsymbol{\mu_s}[2]$, $\boldsymbol{\mu_s}[3]$, $\boldsymbol{\mu_s}[4]$ and $\boldsymbol{\mu_s}[5]$. Otherwise it is exactly equivalent to CALL except:

$$(\boldsymbol{\sigma}', g', A^+, \mathbf{o}) \equiv \begin{cases} \Theta(\boldsymbol{\sigma}^*, I_{\mathbf{s}}, I_{\mathbf{o}}, I_{\mathbf{a}}, t, & \\ \quad \boldsymbol{\mu_s}[0], I_{\mathbf{P}}, 0, I_{\mathbf{v}}, \mathbf{i}, I_{\mathbf{e}} + 1) & \text{if} \quad I_{\mathbf{v}} \leqslant \boldsymbol{\sigma}[I_{\mathbf{a}}]_{\mathbf{b}} \wedge I_{\mathbf{e}} < 1024 \\ (\boldsymbol{\sigma}, g, \varnothing, ()) & \text{otherwise} \end{cases}$$

Note the changes (in addition to that of the fourth parameter) to the second and ninth parameters to the call Θ. This means that the recipient is in fact the same account as at present, simply that the code is overwritten *and* the context is almost entirely identical.

0xfe	INVALID	$\varnothing$	$\varnothing$	Designated invalid instruction.
0xff	SELFDESTRUCT	1	0	Halt execution and register account for later deletion.

$$A'_{\mathbf{s}} \equiv A_{\mathbf{s}} \cup \{I_{\mathbf{a}}\}$$
$$\boldsymbol{\sigma}'[\boldsymbol{\mu_s}[0] \mod 2^{160}]_{\mathbf{b}} \equiv \boldsymbol{\sigma}[\boldsymbol{\mu_s}[0] \mod 2^{160}]_{\mathbf{b}} + \boldsymbol{\sigma}[I_{\mathbf{a}}]_{\mathbf{b}}$$
$$\boldsymbol{\sigma}'[I_{\mathbf{a}}]_{\mathbf{b}} \equiv 0$$
$$A'_r \equiv A_r + \begin{cases} R_{\text{selfdestruct}} & \text{if} \quad I_{\mathbf{a}} \notin A_{\mathbf{s}} \\ 0 & \text{otherwise} \end{cases}$$
$$C_{\text{SELFDESTRUCT}}(\boldsymbol{\sigma}, \boldsymbol{\mu}) \equiv G_{\text{selfdestruct}} + \begin{cases} G_{\text{newaccount}} & \text{if} \quad \boldsymbol{\sigma}[\boldsymbol{\mu_s}[0] \mod 2^{160}] = \varnothing \\ 0 & \text{otherwise} \end{cases}$$

Appendix I. Genesis Block

The genesis block is 15 items, and is specified thus:

$$(228) \qquad \big((0_{256}, \texttt{KEC}(\texttt{RLP}(())), 0_{160}, stateRoot, 0, 0, 0_{2048}, 2^{17}, 0, 0, 3141592, time, 0, 0_{256}, \texttt{KEC}((42))), (), () \big)$$

Where 0_{256} refers to the parent hash, a 256-bit hash which is all zeroes; 0_{160} refers to the beneficiary address, a 160-bit hash which is all zeroes; 0_{2048} refers to the log bloom, a 2048-bit sequence of all zeros; 2^{17} refers to the difficulty; the transaction trie root, receipt trie root, gas used, block number and extradata are all 0 (listed in the order of decimal zeros in the above sequence), being equivalent to the empty byte array. The sequences of both ommers and transactions are empty and represented by (). $\texttt{KEC}((42))$ refers to the Keccak hash of a byte array of length one whose first and only byte is of value 42, used for the nonce. $\texttt{KEC}(\texttt{RLP}(()))$ value refers to the hash of the ommer lists in RLP, both empty lists.

The proof-of-concept series includes a development premine, making the state root hash some value *stateRoot*. Also *time* will be set to the initial timestamp of the genesis block. The latest documentation should be consulted for those values.

Appendix J. Ethash

J.1. **Definitions.** We employ the following definitions:

Name	Value	Description
$J_{\text{wordbytes}}$	4	This is the bytes in a word.
$J_{\text{datasetinit}}$	2^{30}	This is the bytes in the dataset at genesis.
$J_{\text{datasetgrowth}}$	2^{23}	This is the dataset growth per epoch.
$J_{\text{cacheinit}}$	2^{24}	This is the bytes in the cache at genesis.
$J_{\text{cachegrowth}}$	2^{17}	This is the cache growth per epoch.
J_{epoch}	30000	This is the blocks per epoch.
J_{mixbytes}	128	This is the mix length in bytes.
$J_{\text{hashbytes}}$	64	This is the hash length in bytes.
J_{parents}	256	This is the number of parents of each dataset element.
$J_{\text{cacherounds}}$	3	This is the number of rounds in cache production.
J_{accesses}	64	This is the number of accesses in a hashimoto loop.

J.2. **Size of dataset and cache.** The size for Ethash's cache $\mathbf{c} \in \mathbb{B}$ and dataset $\mathbf{d} \in \mathbb{B}$ depend on the epoch, which in turn depends on the block number.

$$(229) \qquad E_{\text{epoch}}(H_{\mathbf{i}}) = \left\lfloor \frac{H_{\mathbf{i}}}{J_{\text{epoch}}} \right\rfloor$$

The size of the dataset growth is $J_{\text{datasetgrowth}}$ bytes, and the size of the cache growth is $J_{\text{cachegrowth}}$ bytes, where the growth of each occurs every epoch. In order to avoid regularity leading to cyclic behavior, the size must be a prime

number. Therefore the size is reduced by a multiple of J_{mixbytes}, for the dataset, and $J_{\text{hashbytes}}$ for the cache. Let $d_{\text{size}} = \|\mathbf{d}\|$ be the size of the dataset. Which is calculated using:

$$(230) \qquad d_{\text{size}} = E_{\text{prime}}(J_{\text{datasetinit}} + J_{\text{datasetgrowth}} \cdot E_{\text{epoch}} - J_{\text{mixbytes}}, J_{\text{mixbytes}})$$

The size of the cache, c_{size}, is calculated using

$$(231) \qquad c_{\text{size}} = E_{\text{prime}}(J_{\text{cacheinit}} + J_{\text{cachegrowth}} \cdot E_{\text{epoch}} - J_{\text{hashbytes}}, J_{\text{hashbytes}})$$

$$(232) \qquad E_{\text{prime}}(x, y) = \begin{cases} x & \text{if} \quad x/y \in \mathbb{P} \\ E_{\text{prime}}((x - 1) \cdot y, y) & \text{otherwise} \end{cases}$$

J.3. Dataset generation. In order the generate the dataset we need the cache c, which is an array of bytes. It depends on the cache size c_{size} and the seed hash $s \in \mathbb{B}_{32}$.

J.3.1. *Seed hash.* The seed hash is different for every epoch. For the first epoch it is the Keccak-256 hash of a series of 32 bytes of zeros. For every other epoch it is always the Keccak-256 hash of the previous seed hash:

$$(233) \qquad s = C_{\text{seedhash}}(H_{\mathbf{i}})$$

$$(234) \qquad C_{\text{seedhash}}(H_{\mathbf{i}}) = \begin{cases} \texttt{KEC}(\mathbf{0}_{32}) & \text{if} \quad E_{\text{epoch}}(H_{\mathbf{i}}) = 0 \\ \texttt{KEC}(C_{\text{seedhash}}(H_{\mathbf{i}} - J_{\text{epoch}})) & \text{otherwise} \end{cases}$$

With $\mathbf{0}_{32}$ being 32 bytes of zeros.

J.3.2. *Cache.* The cache production process involves using the seed hash to first sequentially fill up c_{size} bytes of memory, then performing $J_{\text{cacherounds}}$ passes of the RandMemoHash algorithm created by Lerner [2014]. The initial cache $\mathbf{c}'$, being an array of arrays of single bytes, will be constructed as follows.

We define the array $\mathbf{c}_i$, consisting of 64 single bytes, as the ith element of the initial cache:

$$(235) \qquad \mathbf{c}_i = \begin{cases} \texttt{KEC512}(\mathbf{s}) & \text{if} \quad i = 0 \\ \texttt{KEC512}(\mathbf{c}_{i-1}) & \text{otherwise} \end{cases}$$

Therefore $\mathbf{c}'$ can be defined as

$$(236) \qquad \mathbf{c}'[i] = \mathbf{c}_i \quad \forall \quad i < n$$

$$(237) \qquad n = \left\lfloor \frac{c_{\text{size}}}{J_{\text{hashbytes}}} \right\rfloor$$

The cache is calculated by performing $J_{\text{cacherounds}}$ rounds of the RandMemoHash algorithm to the inital cache $\mathbf{c}'$:

$$(238) \qquad \mathbf{c} = E_{\text{cacherounds}}(\mathbf{c}', J_{\text{cacherounds}})$$

$$(239) \qquad E_{\text{cacherounds}}(\mathbf{x}, y) = \begin{cases} \mathbf{x} & \text{if} \quad y = 0 \\ E_{\text{RMH}}(\mathbf{x}) & \text{if} \quad y = 1 \\ E_{\text{cacherounds}}(E_{\text{RMH}}(\mathbf{x}), y - 1) & \text{otherwise} \end{cases}$$

Where a single round modifies each subset of the cache as follows:

$$(240) \qquad E_{\text{RMH}}(\mathbf{x}) = \big(E_{rmh}(\mathbf{x}, 0), E_{rmh}(\mathbf{x}, 1), ..., E_{rmh}(\mathbf{x}, n - 1) \big)$$

$$(241) \quad E_{rmh}(\mathbf{x}, i) = \texttt{KEC512}(\mathbf{x}'[(i - 1 + n) \mod n] \oplus \mathbf{x}'[\mathbf{x}'[i][0] \mod n])$$
$$\text{with} \quad \mathbf{x}' = \mathbf{x} \quad \text{except} \quad \mathbf{x}'[j] = E_{rmh}(\mathbf{x}, j) \quad \forall \quad j < i$$

J.3.3. *Full dataset calculation.* Essentially, we combine data from J_{parents} pseudorandomly selected cache nodes, and hash that to compute the dataset. The entire dataset is then generated by a number of items, each $J_{\text{hashbytes}}$ bytes in size:

$$(242) \qquad \mathbf{d}[i] = E_{\text{datasetitem}}(\mathbf{c}, i) \quad \forall \quad i < \left\lfloor \frac{d_{\text{size}}}{J_{\text{hashbytes}}} \right\rfloor$$

In order to calculate the single item we use an algorithm inspired by the FNV hash (Glenn Fowler [1991]) in some cases as a non-associative substitute for XOR.

$$(243) \qquad E_{\text{FNV}}(\mathbf{x}, \mathbf{y}) = (\mathbf{x} \cdot (\texttt{0x01000193} \oplus \mathbf{y})) \mod 2^{32}$$

The single item of the dataset can now be calculated as:

$$(244) \qquad E_{\text{datasetitem}}(\mathbf{c}, i) = E_{\text{parents}}(\mathbf{c}, i, -1, \varnothing)$$

$$(245) \qquad E_{\text{parents}}(\mathbf{c}, i, p, \mathbf{m}) = \begin{cases} E_{\text{parents}}(\mathbf{c}, i, p + 1, E_{\text{mix}}(\mathbf{m}, \mathbf{c}, i, p + 1)) & \text{if} \quad p < J_{\text{parents}} - 2 \\ E_{\text{mix}}(\mathbf{m}, \mathbf{c}, i, p + 1) & \text{otherwise} \end{cases}$$

$$(246) \qquad E_{\text{mix}}(\mathbf{m}, \mathbf{c}, i, p) = \begin{cases} \texttt{KEC512}(\mathbf{c}[i \mod \mathbf{c}_{\text{size}}] \oplus i) & \text{if} \quad p = 0 \\ E_{\text{FNV}}\big(\mathbf{m}, \mathbf{c}[E_{\text{FNV}}(i \oplus p, \mathbf{m}[p \mod \lfloor J_{\text{hashbytes}}/J_{\text{wordbytes}} \rfloor]) \mod \mathbf{c}_{\text{size}}]\big) & \text{otherwise} \end{cases}$$

J.4. Proof-of-work function. Essentially, we maintain a "mix" J_{mixbytes} bytes wide, and repeatedly sequentially fetch J_{mixbytes} bytes from the full dataset and use the E_{FNV} function to combine it with the mix. J_{mixbytes} bytes of sequential access are used so that each round of the algorithm always fetches a full page from RAM, minimizing translation lookaside buffer misses which ASICs would theoretically be able to avoid.

If the output of this algorithm is below the desired target, then the nonce is valid. Note that the extra application of $\texttt{KEC}$ at the end ensures that there exists an intermediate nonce which can be provided to prove that at least a small amount of work was done; this quick outer PoW verification can be used for anti-DDoS purposes. It also serves to provide statistical assurance that the result is an unbiased, 256 bit number.

The PoW-function returns an array with the compressed mix as its first item and the Keccak-256 hash of the concatenation of the compressed mix with the seed hash as the second item:
$$(247)$$
$$\texttt{PoW}(H_{\slashed{n}}, H_{\mathbf{n}}, \mathbf{d}) = \{\mathbf{m_c}(\texttt{KEC}(\texttt{RLP}(L_H(H_{\slashed{n}}))), H_{\mathbf{n}}, \mathbf{d}), \texttt{KEC}(\mathbf{s_h}(\texttt{KEC}(\texttt{RLP}(L_H(H_{\slashed{n}}))), H_{\mathbf{n}}) + \mathbf{m_c}(\texttt{KEC}(\texttt{RLP}(L_H(H_{\slashed{n}}))), H_{\mathbf{n}}, \mathbf{d}))\}$$

With $H_{\slashed{n}}$ being the hash of the header without the nonce. The compressed mix $\mathbf{m_c}$ is obtained as follows:

$$(248) \qquad \mathbf{m_c}(\mathbf{h}, \mathbf{n}, \mathbf{d}) = E_{\text{compress}}(E_{\text{accesses}}(\mathbf{d}, \sum_{i=0}^{n_{\text{mix}}} \mathbf{s_h}(\mathbf{h}, \mathbf{n}), \mathbf{s_h}(\mathbf{h}, \mathbf{n}), -1), -4)$$

The seed hash being:
$$(249) \qquad \mathbf{s_h}(\mathbf{h}, \mathbf{n}) = \texttt{KEC512}(\mathbf{h} + E_{\text{revert}}(\mathbf{n}))$$

$E_{\text{revert}}(\mathbf{n})$ returns the reverted bytes sequence of the nonce $\mathbf{n}$:
$$(250) \qquad E_{\text{revert}}(\mathbf{n})[i] = \mathbf{n}[\|\mathbf{n}\| - i]$$

We note that the "+"-operator between two byte sequences results in the concatenation of both sequences.

The dataset $\mathbf{d}$ is obtained as described in section J.3.3.

The number of replicated sequences in the mix is:
$$(251) \qquad n_{\text{mix}} = \left\lfloor \frac{J_{\text{mixbytes}}}{J_{\text{hashbytes}}} \right\rfloor$$

In order to add random dataset nodes to the mix, the E_{accesses} function is used:

$$(252) \qquad E_{\text{accesses}}(\mathbf{d}, \mathbf{m}, \mathbf{s}, i) = \begin{cases} E_{\text{mixdataset}}(\mathbf{d}, \mathbf{m}, \mathbf{s}, i) & \text{if} \quad i = J_{\text{accesses}} - 2 \\ E_{\text{accesses}}(E_{\text{mixdataset}}(\mathbf{d}, \mathbf{m}, \mathbf{s}, i), \mathbf{s}, i + 1) & \text{otherwise} \end{cases}$$

$$(253) \qquad E_{\text{mixdataset}}(\mathbf{d}, \mathbf{m}, \mathbf{s}, i) = E_{\text{FNV}}(\mathbf{m}, E_{newdata}(\mathbf{d}, \mathbf{m}, \mathbf{s}, i))$$

$E_{newdata}$ returns an array with n_{mix} elements:
$$(254)$$
$$E_{newdata}(\mathbf{d}, \mathbf{m}, \mathbf{s}, i)[j] = \mathbf{d}[E_{\text{FNV}}(i \oplus \mathbf{s}[0], \mathbf{m}[i \mod \left\lfloor \frac{J_{\text{mixbytes}}}{J_{\text{wordbytes}}} \right\rfloor]) \mod \left\lfloor \frac{d_{\text{size}}}{J_{\text{hashbytes}}} n_{\text{mix}} \right\rfloor \cdot n_{\text{mix}} + j] \quad \forall \quad j < n_{\text{mix}}$$

The mix is compressed as follows:

$$(255) \quad E_{\text{compress}}(\mathbf{m}, i) = \begin{cases} \mathbf{m} & \text{if} \quad i \geqslant \|\mathbf{m}\| - 8 \\ E_{\text{compress}}(E_{\text{FNV}}(E_{\text{FNV}}(E_{\text{FNV}}(\mathbf{m}[i+4], \mathbf{m}[i+5]), \mathbf{m}[i+6]), \mathbf{m}[i+7]), i+8) & \text{otherwise} \end{cases}$$

APPENDIX K. LIST OF MATHEMATICAL SYMBOLS

Symbol	Latex Command	Description
$\bigvee$	\bigvee	This is the least upper bound, supremum, or join of all elements operated on. Thus it is the greatest element of such elements (Lat).