

ETHEREUM: A SECURE DECENTRALISED GENERALISED TRANSACTION LEDGER

SHANGHAI VERSION f3553dd – 2024-11-04

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

ABSTRACT. The blockchain paradigm when coupled with cryptographically-secured transactions has demonstrated its utility through a number of projects, with Bitcoin being one of the most notable ones. 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 that can be 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, we 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 [2013] 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 et al. [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 current state. It is this current 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:

$$(1) \quad \sigma_{t+1} \equiv \Upsilon(\sigma_t, T)$$

where Υ is the Ethereum state transition function. In Ethereum, Υ , together with σ are considerably more powerful than 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, recording a series of

transactions together with the previous block and an identifier for the final state (though do not store the final state itself—that would be far too big).

Formally, we expand to:

$$\begin{aligned} (2) \quad \sigma_{t+1} &\equiv \Pi(\sigma_t, B) \\ (3) \quad B &\equiv (\dots, (T_0, T_1, \dots), \dots) \\ (4) \quad \Pi(\sigma, B) &\equiv \Upsilon(\Upsilon(\sigma, T_0), T_1) \dots \end{aligned}$$

Where B is this block, which includes a series of transactions amongst some other components and Π is the block-level state-transition function for transactions¹.

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 |
| 10^9 | Gwei |
| 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.

Since the *Paris* hard fork, reaching consensus on new blocks is managed by a protocol called the *Beacon Chain*. It is known as the *consensus layer* of Ethereum, and it defines the rules for determining the canonical history of Ethereum blocks. This document describes the *execution layer* of Ethereum. The execution layer defines the rules for interacting with and updating the state of the Ethereum virtual machine. The consensus layer is described in greater detail in the consensus specifications. How the consensus

¹Note that since the Shanghai fork, blocks also needs to process *withdrawal operations* in order to reach their final state. Withdrawal operations are defined in sub-section 4.3, and block final state is discussed in greater detail in section 12.

layer is used to determine the canonical state of Ethereum is discussed in section 11.

There are many versions of Ethereum, as the protocol has undergone a number of updates. These updates can be specified to occur:

- at a particular block number in the case of pre-*Paris* updates,
- after reaching a *terminal total difficulty* in the case of the *Paris* update, or
- at a particular block timestamp in the case of post-*Paris* updates.

This document describes the *Shanghai* version.

In order to follow back the history of a path, one must reference multiple versions of this document. Here are the block numbers of protocol updates on the Ethereum main network:²

| Name | First Block Number |
|-------------------------------|--------------------|
| $F_{\text{Homestead}}$ | 1150000 |
| $F_{\text{TangerineWhistle}}$ | 2463000 |
| $F_{\text{SpuriousDragon}}$ | 2675000 |
| $F_{\text{Byzantium}}$ | 4370000 |
| $F_{\text{Constantinople}}$ | 7280000 |
| $F_{\text{Petersburg}}$ | 7280000 |
| F_{Istanbul} | 9069000 |
| $F_{\text{MuirGlacier}}$ | 9200000 |
| F_{Berlin} | 12244000 |
| F_{London} | 12965000 |
| $F_{\text{ArrowGlacier}}$ | 13773000 |
| $F_{\text{GrayGlacier}}$ | 15050000 |
| F_{Paris} | 15537394 |
| F_{Shanghai} | 17034870 |

Occasionally actors do not agree on a protocol change, and a permanent fork occurs. In order to distinguish between diverged blockchains, EIP-155 by Buterin [2016] introduced the concept of chain ID, which we denote by β . For the Ethereum main network

$$(5) \quad \beta = 1$$

3. CONVENTIONS

We 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 σ (or a variant thereupon) and those of machine-state, μ .

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_{STORE} , the cost function for the `STORE` operation. For specialised and possibly externally defined functions, we may format as typewriter text, e.g. the Keccak-256 hash function (as per version 3 of the winning entry to the SHA-3 contest by Bertoni et al. [2011], rather than the final SHA-3 specification), is

denoted KEC (and generally referred to as plain Keccak). Also, KEC512 refers 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_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 non-negative integers and thus belong to the set \mathbb{N} . 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 non-negative integers smaller than 2^{256} is named \mathbb{N}_{256} . This is formally defined in section 4.4.

Square brackets are used to index into and reference individual components or subsequences of sequences, e.g. $\mu_s[0]$ 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. $\mu_m[0..31]$ denotes the first 32 items of the machine’s memory.

In the case of the global state σ , which is a sequence of accounts, themselves tuples, the square brackets are used to reference an individual account.

When considering variants of existing values, we 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, we 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. It is formally defined in section 4.4.

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

$$(6) \quad \ell(\mathbf{x}) \equiv \mathbf{x}[|\mathbf{x}| - 1]$$

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.

²Note that while the Paris, Shanghai, and every upcoming forks activated at a given block number (e.g. 15,537,394 for Paris), the trigger was not the block number, but rather reaching a specified *timestamp* (or *total difficulty* for Paris). The trigger for the *Paris* and subsequent hard fork are discussed in greater detail in section 10.

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 byte arrays to byte arrays; 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, $\sigma[a]$, 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 σ , this would be formally denoted $\sigma[a]_n$.

balance: A scalar value equal to the number of Wei owned by this address. Formally denoted $\sigma[a]_b$.

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 $\sigma[a]_s$.

codeHash: The hash of the EVM code of this account—this is the code that gets executed should this address receive a message call. All such code fragments are contained in the state database under their corresponding hashes for later retrieval. This hash is formally denoted $\sigma[a]_c$, and thus the code may be denoted as \mathbf{b} , given that $\text{KEC}(\mathbf{b}) = \sigma[a]_c$.

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

$$(7) \quad \text{TRIE}(L_I^*(\sigma[a]_s)) \equiv \sigma[a]_s$$

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:

$$(8) \quad L_I((k, v)) \equiv (\text{KEC}(k), \text{RLP}(v))$$

where:

$$(9) \quad k \in \mathbb{B}_{32} \quad \wedge \quad v \in \mathbb{N}$$

It shall be understood that $\sigma[a]_s$ 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. $\sigma[a]_c = \text{KEC}()$, 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 :

$$(10) \quad L_S(\sigma) \equiv \{p(a) : \sigma[a] \neq \emptyset\}$$

where

$$(11) \quad p(a) \equiv (\text{KEC}(a), \text{RLP}((\sigma[a]_n, \sigma[a]_b, \sigma[a]_s, \sigma[a]_c)))$$

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

$$(12) \quad \forall a : \sigma[a] = \emptyset \vee (a \in \mathbb{B}_{20} \wedge v(\sigma[a]))$$

where v is the account validity function:

$$(13) \quad v(x) \equiv x_n \in \mathbb{N}_{256} \wedge x_b \in \mathbb{N}_{256} \wedge x_s \in \mathbb{B}_{32} \wedge x_c \in \mathbb{B}_{32}$$

An account is *empty* when it has no code, zero nonce and zero balance:

$$(14) \quad \text{EMPTY}(\sigma, a) \equiv \sigma[a]_c = \text{KEC}() \wedge \sigma[a]_n = 0 \wedge \sigma[a]_b = 0$$

Even callable precompiled contracts can have an empty account state. This is because their account states do not usually contain the code describing its behavior.

An account is *dead* when its account state is non-existent or empty:

$$(15) \quad \text{DEAD}(\sigma, a) \equiv \sigma[a] = \emptyset \vee \text{EMPTY}(\sigma, a)$$

4.2. The Transaction. A transaction (formally, T) is a single cryptographically-signed instruction constructed by an actor externally to the scope of Ethereum. The sender of a transaction cannot be a contract. While it is assumed that the ultimate external actor will be human in nature, software tools will be used in its construction and dissemination³. EIP-2718 by Zoltu [2020] introduced the notion of different transaction types. As of the *London* version of the protocol, there are three transaction types: 0 (legacy), 1 (EIP-2930 by Buterin and Swende [2020b]), and 2 (EIP-1559 by Buterin et al. [2019]). Further, there are two subtypes 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’). All transaction types specify a number of common fields:

type: EIP-2718 transaction type; formally T_x .

nonce: A scalar value equal to the number of transactions sent by the sender; formally T_n .

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_g .

to: The 160-bit address of the message call’s recipient or, for a contract creation transaction, \emptyset , used here to denote the only member of \mathbb{B}_0 ; formally T_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_v .

³Notably, 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.

r, s: Values corresponding to the signature of the transaction and used to determine the sender of the transaction; formally T_r and T_s . This is expanded in Appendix F.

EIP-2930 (type 1) and EIP-1559 (type 2) transactions also have:

accessList: List of access entries to warm up; formally T_A . Each access list entry E is a tuple of an account address and a list of storage keys: $E \equiv (E_a, E_s)$.

chainId: Chain ID; formally T_c . Must be equal to the network chain ID β .

yParity: Signature Y parity; formally T_y .

Legacy transactions do not have an **accessList** ($T_A = ()$), while **chainId** and **yParity** for legacy transactions are combined into a single value:

w: A scalar value encoding Y parity and possibly chain ID; formally T_w . $T_w = 27 + T_y$ or $T_w = 2\beta + 35 + T_y$ (see EIP-155 by Buterin [2016]).

There are differences in how one’s acceptable gas price is specified in type 2 transactions versus type 0 and type 1 transactions. Type 2 transactions take better advantage of the gas market improvements introduced in EIP-1559 by explicitly limiting the *priority fee*⁴ that is paid. Type 2 transactions have the following two fields related to gas:

maxFeePerGas: A scalar value equal to the maximum 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_m .

maxPriorityFeePerGas: A scalar value equal to the maximum number of Wei to be paid to the block’s fee recipient as an incentive to include the transaction; formally T_f .

In contrast, type 0 and type 1 transactions specify gas price as a single value:

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_p .⁵

Additionally, a contract creation transaction (regardless of transaction type) contains:

init: An unlimited size byte array specifying the EVM-code for the account initialisation procedure, formally T_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_d .

Appendix F specifies the function, S , which maps transactions to the sender, and happens through the ECDSA of the SECP-256k1 curve, using the hash of the transaction (excepting the latter three signature fields) as the datum

to sign. For the present we simply assert that the sender of a given transaction T can be represented with $S(T)$.

(16)

$$L_T(T) \equiv \begin{cases} (T_n, T_p, T_g, T_t, T_v, \mathbf{p}, T_w, T_r, T_s) & \text{if } T_x = 0 \\ (T_c, T_n, T_p, T_g, T_t, T_v, \mathbf{p}, T_A, T_y, T_r, T_s) & \text{if } T_x = 1 \\ (T_c, T_n, T_f, T_m, T_g, T_t, T_v, \mathbf{p}, T_A, T_y, T_r, T_s) & \text{if } T_x = 2 \end{cases}$$

where

$$(17) \quad \mathbf{p} \equiv \begin{cases} T_i & \text{if } T_t = \emptyset \\ T_d & \text{otherwise} \end{cases}$$

Here, we assume all components are interpreted by the RLP as integer values, with the exception of the access list T_A and the arbitrary length byte arrays T_i and T_d .

(18)

$$\begin{aligned} T_x \in \{0, 1, 2\} & \wedge T_c = \beta & \wedge T_n \in \mathbb{N}_{256} & \wedge \\ T_p \in \mathbb{N}_{256} & \wedge T_g \in \mathbb{N}_{256} & \wedge T_v \in \mathbb{N}_{256} & \wedge \\ T_w \in \mathbb{N}_{256} & \wedge T_r \in \mathbb{N}_{256} & \wedge T_s \in \mathbb{N}_{256} & \wedge \\ T_y \in \mathbb{N}_1 & \wedge T_d \in \mathbb{B} & \wedge T_i \in \mathbb{B} & \wedge \\ T_m \in \mathbb{N}_{256} & \wedge T_f \in \mathbb{N}_{256} & & \end{aligned}$$

where

$$(19) \quad \mathbb{N}_n = \{P : P \in \mathbb{N} \wedge P < 2^n\}$$

The address hash T_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 \emptyset), it is the RLP empty byte sequence and thus the member of \mathbb{B}_0 :

$$(20) \quad T_t \in \begin{cases} \mathbb{B}_{20} & \text{if } T_t \neq \emptyset \\ \mathbb{B}_0 & \text{otherwise} \end{cases}$$

4.3. The Withdrawal. A withdrawal (formally, W) is a tuple of data describing a consensus layer validator’s withdrawal of some amount of its staked Ether. A withdrawal is created and validated in the consensus layer of the blockchain and then pushed to the execution layer. A withdrawal is composed of the following fields:

globalIndex: zero based incrementing withdrawal index that acts as a unique identifier for this withdrawal; formally W_g .

validatorIndex: index of consensus layer’s validator this withdrawal corresponds to; formally W_v .

recipient: the 20-byte address that will receives Ether from this withdrawal; formally W_r .

amount: a nonzero amount of Ether denominated in Gwei (10^9 Wei); formally W_a .

Withdrawal serialisation is defined as:

$$(21) \quad L_W(W) \equiv (W_g, W_v, W_r, T_a)$$

Here, we assume all components are interpreted by the RLP as integer values except for W_r which is a 20-byte address:

$$(22) \quad \begin{aligned} W_g \in \mathbb{N}_{64} & \wedge W_v \in \mathbb{N}_{64} & \wedge \\ W_r \in \mathbb{B}_{20} & \wedge W_a \in \mathbb{N}_{64} & \end{aligned}$$

⁴The *priority fee* is discussed in greater detail in sections 5 and 6.

⁵Type 0 and type 1 transactions will get the same gas price behavior as a type 2 transaction with T_m and T_f set to the value of T_p .

⁶*ommer* is a gender-neutral term to mean “sibling of parent”; see https://nonbinary.miraheze.org/wiki/Gender_neutral_language_in_English#Aunt/Uncle

4.4. 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} , a now deprecated property \mathbf{U} that prior to the *Paris* hard fork contained headers of blocks whose parents were equal to the present block's parent's parent (such blocks were known as *ommers*⁶), and, since the *Shanghai* hard fork, \mathbf{W} , a collection of validator's withdrawal pushed by the consensus layer. 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_p .

ommersHash: A 256-bit hash field that is now deprecated due to the replacement of proof of work consensus. It is now to a constant, $\text{KEC}(\text{RLP}(()))$; formally H_o .

beneficiary: The 160-bit address to which priority fees from this block are transferred; formally H_c .

stateRoot: The Keccak 256-bit hash of the root node of the state trie, after all transactions and withdrawals are executed and finalisations applied; formally H_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_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_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_b .

difficulty: A scalar field that is now deprecated due to the replacement of proof of work consensus. It is set to 0; formally H_d .

number: A scalar value equal to the number of ancestor blocks. The genesis block has a number of zero; formally H_i .

gasLimit: A scalar value equal to the current limit of gas expenditure per block; formally H_l .

gasUsed: A scalar value equal to the total gas used in transactions in this block; formally H_g .

timestamp: A scalar value equal to the reasonable output of Unix's `time()` at this block's inception; formally H_s .

extraData: An arbitrary byte array containing data relevant to this block. This must be 32 bytes or fewer; formally H_x .

prevRandao: the latest RANDAO mix⁷ of the post beacon state of the previous block; formally H_a .

nonce: A 64-bit value that is now deprecated due to the replacement of proof of work consensus. It is set to `0x00000000000000000000`; formally H_n .

baseFeePerGas: A scalar value equal to the amount of wei that is burned for each unit of gas consumed; formally H_f .

withdrawalsRoot: The Keccak 256-bit hash of the root node of the trie structure populated with each withdrawal operations pushed by the consensus layer for this block; formally H_w .

The other three components in the block are a series of transactions, $B_{\mathbf{T}}$, an empty array which was previously reserved for ommer block headers, $B_{\mathbf{U}}$, and a series of withdrawals, $B_{\mathbf{W}}$. Formally, we can refer to a block B :

$$(23) \quad B \equiv (B_H, B_{\mathbf{T}}, B_{\mathbf{U}}, B_{\mathbf{W}})$$

4.4.1. 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 its execution. Each receipt, denoted $B_{\mathbf{R}}[i]$ for the i th transaction, is placed in an index-keyed trie and the root recorded in the header as H_e .

The transaction receipt, R , is a tuple of five items comprising: the type of the transaction, R_x , the status code of the transaction, R_z , the cumulative gas used in the block containing the transaction receipt as of immediately after the transaction has happened, R_u , the set of logs created through execution of the transaction, R_l and the Bloom filter composed from information in those logs, R_b :

$$(24) \quad R \equiv (R_x, R_z, R_u, R_b, R_l)$$

R_x is equal to the type of the corresponding transaction.

The function $L_{\mathbf{R}}$ prepares a transaction receipt for being transformed into an RLP-serialised byte array:

$$(25) \quad L_{\mathbf{R}}(R) \equiv (R_z, R_u, R_b, R_l)$$

We assert that the status code R_z is a non-negative integer:

$$(26) \quad R_z \in \mathbb{N}$$

We assert that R_u , the cumulative gas used, is a non-negative integer and that the logs Bloom, R_b , is a hash of size 2048 bits (256 bytes):

$$(27) \quad R_u \in \mathbb{N} \quad \wedge \quad R_b \in \mathbb{B}_{256}$$

The sequence R_l is a series of log entries, (O_0, O_1, \dots) . A log entry, O , is a tuple of the logger's address, O_a , a possibly empty series of 32-byte log topics, O_t and some number of bytes of data, O_d :

$$(28) \quad O \equiv (O_a, (O_{t0}, O_{t1}, \dots), O_d)$$

$$(29) \quad O_a \in \mathbb{B}_{20} \quad \wedge \quad \forall x \in O_t : x \in \mathbb{B}_{32} \quad \wedge \quad O_d \in \mathbb{B}$$

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

$$(30) \quad M(O) \equiv \bigvee_{x \in \{O_a\} \cup O_t} (M_{3:2048}(x))$$

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

⁷RANDAO is a pseudorandom value generated by validators on the Ethereum consensus layer. Refer to the consensus layer specs (<https://github.com/ethereum/consensus-specs>) for more detail on RANDAO.

⁸ $2048 = 2^{11}$ (11 bits), 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."

sequence.⁸ Formally:

$$(31) M_{3:2048}(\mathbf{x} : \mathbf{x} \in \mathbb{B}) \equiv \mathbf{y} : \mathbf{y} \in \mathbb{B}_{256} \quad \text{where:}$$

$$(32) \quad \mathbf{y} = (0, 0, \dots, 0) \quad \text{except:}$$

$$(33) \quad \forall i \in \{0, 2, 4\} : B_{2047-m(\mathbf{x}, i)}(\mathbf{y}) = 1$$

$$(34) \quad m(\mathbf{x}, i) \equiv \text{KEC}(\mathbf{x})[i, i+1] \bmod 2048$$

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} . Notably, it treats \mathbf{x} as big-endian (more significant bits will have smaller indices).

4.4.2. Holistic Validity. We can assert a block's validity if and only if it satisfies several conditions: the block's ommers field B_U must be an empty array and the block's header must be consistent with the given transactions B_T and withdrawals B_W . For the header to be consistent with the transactions B_T and withdrawals B_W , **stateRoot** (H_r) must match the resultant state after executing all transactions, then all withdrawals, in order on the base state σ (as specified in section 12), and **transactionsRoot** (H_t), **receiptsRoot** (H_e), **logsBloom** (H_b) and **withdrawalsRoot** (H_w) must be correctly derived from the transactions themselves, the transaction receipts resulting from execution, the resulting logs, and the withdrawals, respectively.

$$(35) \quad \begin{aligned} B_U &\equiv () && \wedge \\ H_r &\equiv \text{TRIE}(L_S(\Pi(\sigma, B))) && \wedge \\ H_t &\equiv \text{TRIE}(\{\forall i < \|B_T\|, i \in \mathbb{N} : \\ &\quad p_T(i, B_T[i])\}) && \wedge \\ H_e &\equiv \text{TRIE}(\{\forall i < \|B_R\|, i \in \mathbb{N} : \\ &\quad p_R(i, B_R[i])\}) && \wedge \\ H_w &\equiv \text{TRIE}(\{\forall i < \|B_W\|, i \in \mathbb{N} : \\ &\quad p_W(i, B_W[i])\}) && \wedge \\ H_b &\equiv \bigvee_{r \in B_R} (r_b) \end{aligned}$$

where $p_W(k, v)$ is a pairwise RLP transformation for withdrawals:

$$(36) \quad p_W(k, W) \equiv (\text{RLP}(k), \text{RLP}(L_W(W)))$$

similarly, $p_T(k, v)$ and $p_R(k, v)$ are pairwise RLP transformations, but with a special treatment for EIP-2718 transactions:

$$(37) \quad p_T(k, T) \equiv \left(\text{RLP}(k), \begin{cases} \text{RLP}(L_T(T)) & \text{if } T_x = 0 \\ (T_x) \cdot \text{RLP}(L_T(T)) & \text{otherwise} \end{cases} \right)$$

and

$$(38) \quad p_R(k, R) \equiv \left(\text{RLP}(k), \begin{cases} \text{RLP}(L_R(R)) & \text{if } R_x = 0 \\ (R_x) \cdot \text{RLP}(L_R(R)) & \text{otherwise} \end{cases} \right)$$

(\cdot is the concatenation of byte arrays).

Furthermore:

$$(39) \quad \text{TRIE}(L_S(\sigma)) = P(B_H)_{H_r}$$

Thus $\text{TRIE}(L_S(\sigma))$ is the root node hash of the Merkle Patricia tree structure containing the key-value pairs of the state σ 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_R , and that defined through the transaction's state-accumulation function, Π , are formalised later in section 12.3.

4.4.3. Serialisation. The function L_B and L_H are the preparation functions for a block and block header respectively. We assert the types and order of the structure for when the RLP transformation is required:

$$(40) L_H(H) \equiv (H_p, H_o, H_c, H_r, H_t, H_e, H_b, H_d, \\ H_i, H_1, H_g, H_s, H_x, H_a, H_n, H_f, H_w)$$

$$(41) L_B(B) \equiv (L_H(B_H), \tilde{L}_T^*(B_T), L_H^*(B_U), L_W^*(B_W))$$

where \tilde{L}_T takes a special care of EIP-2718 transactions:

$$(42) \quad \tilde{L}_T(T) = \begin{cases} L_T(T) & \text{if } T_x = 0 \\ (T_x) \cdot \text{RLP}(L_T(T)) & \text{otherwise} \end{cases}$$

with \tilde{L}_T^* , L_H^* , and L_W^* , being element-wise sequence transformations, thus:

$$(43) \quad f^*((x_0, x_1, \dots)) \equiv (f(x_0), f(x_1), \dots) \quad \text{for any function } f$$

The component types are defined thus:

$$(44) \quad \begin{aligned} H_p &\in \mathbb{B}_{32} & \wedge & H_o &\in \mathbb{B}_{32} & \wedge & H_c &\in \mathbb{B}_{20} & \wedge \\ H_r &\in \mathbb{B}_{32} & \wedge & H_t &\in \mathbb{B}_{32} & \wedge & H_e &\in \mathbb{B}_{32} & \wedge \\ H_b &\in \mathbb{B}_{256} & \wedge & H_d &\in \mathbb{N} & \wedge & H_i &\in \mathbb{N} & \wedge \\ H_1 &\in \mathbb{N} & \wedge & H_g &\in \mathbb{N} & \wedge & H_s &\in \mathbb{N}_{256} & \wedge \\ H_x &\in \mathbb{B} & \wedge & H_a &\in \mathbb{B}_{32} & \wedge & H_n &\in \mathbb{B}_8 & \wedge \\ H_f &\in \mathbb{N} & \wedge & H_w &\in \mathbb{B}_{32} \end{aligned}$$

where

$$(45) \quad \mathbb{B}_n = \{B : B \in \mathbb{B} \wedge \|B\| = n\}$$

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.

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

$$(46) \quad P(H) \equiv B' : \text{KEC}(\text{RLP}(B'_H)) = H_p$$

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

$$(47) \quad H_i \equiv P(H)_{H_i} + 1$$

The *London* release introduced the block attribute *base fee per gas* H_f (see EIP-1559 by Buterin et al. [2019]). The base fee is the amount of wei burned per unit of gas consumed while executing transactions within the block. The value of the base fee is a function of the difference between the gas used by the parent block and the parent block's *gas target*.

The expected base fee per gas is defined as $F(H)$:

$$(48) \quad F(H) \equiv \begin{cases} 1000000000 & \text{if } H_i = F_{\text{London}} \\ P(H)_{H_f} & \text{if } P(H)_{H_g} = \tau \\ P(H)_{H_f} - \nu & \text{if } P(H)_{H_g} < \tau \\ P(H)_{H_f} + \nu & \text{if } P(H)_{H_g} > \tau \end{cases}$$

where:

$$(49) \quad \tau \equiv \lfloor \frac{P(H)_{H_1}}{\rho} \rfloor$$

$$(50) \quad \rho \equiv 2$$

$$(51) \quad \nu^* \equiv \begin{cases} \lfloor \frac{P(H)_{H_f} \times (\tau - P(H)_{H_g})}{\tau} \rfloor & \text{if } P(H)_{H_g} < \tau \\ \lfloor \frac{P(H)_{H_f} \times (P(H)_{H_g} - \tau)}{\tau} \rfloor & \text{if } P(H)_{H_g} > \tau \end{cases}$$

$$(52) \quad \nu \equiv \begin{cases} \lfloor \frac{\nu^*}{\xi} \rfloor & \text{if } P(H)_{H_g} < \tau \\ \max(\lfloor \frac{\nu^*}{\xi} \rfloor, 1) & \text{if } P(H)_{H_g} > \tau \end{cases}$$

$$(53) \quad \xi \equiv 8$$

The *gas target*, τ , is defined as the gas limit H_1 divided by the *elasticity multiplier*, ρ , a global constant set to 2. So while blocks can consume as much gas as the gas limit, the base fee is adjusted so that on average, blocks consume as much gas as the gas target. The base fee is increased in the current block when the parent block's gas usage exceeds the gas target, and, conversely, the base fee is decreased in the current block when the parent block's gas usage is less than the gas target.

The magnitude of the increase or decrease in the base fee, defined as ν , is proportional to the difference between the amount of gas the parent block consumed and the parent block's gas target. The effect on the base fee is dampened by a global constant called the *base fee max change denominator*, formally ξ , set to 8. A value of 8 entails that the base fee can increase or decrease by at most 12.5% from one block to the next.

The canonical gas limit H_1 of a block of header H must fulfil the relation:

$$(54) \quad \begin{aligned} H_1 &< P(H)_{H_{1'}} + \left\lfloor \frac{P(H)_{H_{1'}}}{1024} \right\rfloor \quad \wedge \\ H_1 &> P(H)_{H_{1'}} - \left\lfloor \frac{P(H)_{H_{1'}}}{1024} \right\rfloor \quad \wedge \\ H_1 &\geq 5000 \end{aligned}$$

where:

$$(55) \quad P(H)_{H_{1'}} \equiv \begin{cases} P(H)_{H_1} \times \rho & \text{if } H_i = F_{\text{London}} \\ P(H)_{H_1} & \text{if } H_i > F_{\text{London}} \end{cases}$$

To avoid a discontinuity in gas usage, the value of the parent block gas limit for the purpose of validating the current block's gas limit is modified at the *London* fork block by multiplying it by the *elasticity multiplier*, ρ . We call this modified value $P(H)_{H_{1'}}$. This ensures that the gas target for post-*London* blocks can be set roughly in line with the gas limit of pre-*London* blocks.

H_s is the timestamp (in Unix's time()) of block H and must fulfil the relation:

$$(56) \quad H_s > P(H)_{H_s}$$

The *Paris* hard fork changed Ethereum's consensus from proof of work to proof of stake, and thus deprecated many block header properties related to proof of work. These deprecated properties include **nonce** (H_n), **ommersHash** (H_o), **difficulty** (H_d), and **mixHash** (H_m).

mixHash has been replaced with a new field **prevRando** (H_a). The other header fields related to proof of work have been replaced with constants:

$$(57) \quad H_o \equiv \text{KEC}(\text{RLP}(()))$$

$$(58) \quad H_d \equiv 0$$

$$(59) \quad H_n \equiv 0x0000000000000000$$

The value of **prevRando** must be determined using information from the Beacon Chain. While the details of generating the *RANDAO* value on the Beacon Chain is beyond the scope of this paper, we refer to the expected *RANDAO* value for the previous block as **PREVRANDAO**(\cdot).

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

$$(60) \quad \begin{aligned} V(H) &\equiv H_g \leq H_1 \quad \wedge \\ &H_1 < P(H)_{H_{1'}} + \left\lfloor \frac{P(H)_{H_{1'}}}{1024} \right\rfloor \quad \wedge \\ &H_1 > P(H)_{H_{1'}} - \left\lfloor \frac{P(H)_{H_{1'}}}{1024} \right\rfloor \quad \wedge \\ &H_1 \geq 5000 \quad \wedge \\ &H_s > P(H)_{H_s} \quad \wedge \\ &H_i = P(H)_{H_i} + 1 \quad \wedge \\ &\|H_x\| \leq 32 \quad \wedge \\ &H_f = F(H) \quad \wedge \\ &H_o = \text{KEC}(\text{RLP}(())) \quad \wedge \\ &H_d = 0 \quad \wedge \\ &H_n = 0x0000000000000000 \quad \wedge \\ &H_a = \text{PREVRANDAO}() \end{aligned}$$

Note 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 *effective gas price* defined in section 6. 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.

Since the introduction of EIP-1559 by Buterin et al. [2019] in the *London* hard fork, every transaction included in a block must pay a *base fee*, which is specified as wei per unit of gas consumed and is constant for each transaction within a block. The ether that is paid to meet the base fee is burned (taken out of circulation). The base fee adjusts dynamically as a function of the previous

block’s gas consumption relative to its *gas target* (a value that is currently half the block’s gas limit, which can be adjusted by validators). If the previous block’s total gas consumption exceeds the gas target, this indicates excess demand for block space at the current base fee, and the base fee is increased. Conversely, if the gas consumed in the previous block is lower than the gas target, demand for block space is lower than the gas target at the current base fee, and thus the base fee is decreased. This process of adjusting the base fee should bring the average block’s gas consumption in line with the gas target. Refer to section 4.4 for greater detail on how the base fee is set.

To incentivize validators to include transactions, there is an additional fee known as a *priority fee*, which is also specified as wei per unit of gas consumed. The total fee paid by the transactor therefore is the sum of the base fee per gas and the priority fee per gas multiplied by the total gas consumed. Ether used to satisfy the priority fee is delivered to the *beneficiary* address, the address of an account typically under the control of the validator.

Transactors using type 2 transactions can specify the maximum priority fee they are willing to pay (**maxPriorityFeePerGas**) as well as the max total fee they are willing to pay (**maxFeePerGas**), inclusive of both the priority fee and the base fee. **maxFeePerGas** must be at least as high as the base fee for the transaction to be included in a block. Type 0 and type 1 transactions have only one field for specifying a gas price—**gasPrice**—which also must be at least as high as the base fee for inclusion in a block. The amount by which **gasPrice** is higher than the base fee constitutes the priority fee in the case of a type 0 or type 1 transaction.

Transactors are free to select any priority fee that they wish, however validators are free to ignore transactions as they choose. A higher priority fee on a transaction will therefore cost the sender more in terms of Ether and deliver a greater value to the validator and thus will more likely be selected for inclusion. Since there will be a (weighted) distribution of minimum acceptable priority fees, transactors will necessarily have a trade-off to make between lowering the priority fee and maximising the chance that their transaction will be included in a block 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 sender account has no contract code deployed (see EIP-3607 by Feist et al. [2021]);
- (5) the gas limit is no smaller than the intrinsic gas, g_0 , used by the transaction;
- (6) the sender account balance contains at least the cost, v_0 , required in up-front payment;
- (7) the **maxFeePerGas**, T_m , in the case of type 2 transactions, or **gasPrice**, T_p , in the case of type

0 and type 1 transactions, is greater than or equal to the block’s base fee, H_f ; and

- (8) for type 2 transactions, **maxPriorityFeePerGas**, T_f , must be no larger than **maxFeePerGas**, T_m .

Formally, we consider the function Υ , with T being a transaction and σ the state:

$$(61) \quad \sigma' = \Upsilon(\sigma, T)$$

Thus σ' is the post-transactional state. We also define Υ^g to evaluate to the amount of gas used in the execution of a transaction, Υ^l to evaluate to the transaction’s accrued log items and Υ^z to evaluate to the status code resulting from the transaction. These will 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 the *accrued transaction substate*, or *accrued substate* for short, and represent it as A , which is a tuple:

$$(62) \quad A \equiv (A_s, A_l, A_t, A_r, A_a, A_K)$$

The tuple contents include A_s , the self-destruct set: a set of accounts that will be discarded following the transaction’s completion. A_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). A_t is the set of touched accounts, of which the empty ones are deleted at the end of a transaction. A_r is 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. Finally, EIP-2929 by Buterin and Swende [2020a] introduced A_a , the set of accessed account addresses, and A_K , the set of accessed storage keys (more accurately, each element of A_K is a tuple of a 20-byte account address and a 32-byte storage slot).

We define the empty accrued substate A^0 to have no self-destructs, no logs, no touched accounts, zero refund balance, all precompiled contracts in the accessed addresses, and no accessed storage:

$$(63) \quad A^0 \equiv (\emptyset, (), \emptyset, 0, \pi, \emptyset)$$

where π is the set of all precompiled addresses.

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

$$(64) \quad g_0 \equiv \sum_{i \in T_i, T_d} \begin{cases} G_{\text{txdatazero}} & \text{if } i = 0 \\ G_{\text{txdatanonzero}} & \text{otherwise} \end{cases} + \begin{cases} G_{\text{txcreate}} + R(\|T_i\|) & \text{if } T_t = \emptyset \\ 0 & \text{otherwise} \end{cases} + G_{\text{transaction}} + \sum_{j=0}^{\|T_A\|-1} (G_{\text{accesslistaddress}} + \|T_A[j]_s\| G_{\text{accessliststorage}})$$

where T_i, T_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 it is a message call. $G_{\text{accesslistaddress}}$ and $G_{\text{accessliststorage}}$ are the costs of warming up account and storage access, respectively. G is fully defined in Appendix G.

We define the *initcode cost function*, formally R , as the amount of gas that needs to be paid for each word of the initcode prior to executing the creation code of a new contract:

$$(65) \quad R(x) \equiv G_{\text{initcodeword}} \times \lceil \frac{x}{32} \rceil$$

We define the *effective gas price*, formally p , as the amount of wei the transaction signer will pay per unit of gas consumed during the transaction's execution. It is calculated as follows:

$$(66) \quad p \equiv \begin{cases} T_p & \text{if } T_x = 0 \vee T_x = 1 \\ f + H_f & \text{if } T_x = 2 \end{cases}$$

where f is the *priority fee*—the amount of wei the block's beneficiary address will receive per unit of gas consumed during the transaction's execution. It is calculated as:

$$(67) \quad f \equiv \begin{cases} T_p - H_f & \text{if } T_x = 0 \vee T_x = 1 \\ \min(T_f, T_m - H_f) & \text{if } T_x = 2 \end{cases}$$

The up-front cost v_0 is calculated as:

$$(68) \quad v_0 \equiv \begin{cases} T_g T_p + T_v & \text{if } T_x = 0 \vee T_x = 1 \\ T_g T_m + T_v & \text{if } T_x = 2 \end{cases}$$

The validity is determined as:

$$(69) \quad \begin{aligned} S(T) &\neq \emptyset \wedge \\ \sigma[S(T)]_c &= \text{KEC}(\emptyset) \wedge \\ T_n &= \sigma[S(T)]_n \wedge \\ g_0 &\leq T_g \wedge \\ v_0 &\leq \sigma[S(T)]_b \wedge \\ m &\geq H_f \wedge \\ n &\leq 49152 \wedge \\ T_g &\leq B_{H1} - \ell(B_{\mathbf{R}})_u \end{aligned}$$

where

$$(70) \quad m \equiv \begin{cases} T_p & \text{if } T_x = 0 \vee T_x = 1 \\ T_m & \text{if } T_x = 2 \end{cases}$$

and

$$(71) \quad n \equiv \begin{cases} \|T_i\| & \text{if } T_t \neq \emptyset \\ 0 & \text{otherwise} \end{cases}$$

The penultimate condition ensures that, for create transactions, the length of the initcode is no greater than 49152 bytes. Note the final condition; the sum of the transaction's gas limit, T_g , and the gas utilised in this block prior, given by $\ell(B_{\mathbf{R}})_u$, must be no greater than the block's **gasLimit**, B_{H1} . Also, with a slight abuse of notation, we assume that $\sigma[S(T)]_c = \text{KEC}(\emptyset)$, $\sigma[S(T)]_n = 0$, and $\sigma[S(T)]_b = 0$ if $\sigma[S(T)] = \emptyset$.

For type 2 transactions, we add an additional check that **maxPriorityFeePerGas** is no larger than **maxFeePerGas**:

$$(72) \quad T_m \geq T_f$$

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, $T_g p$. The gas available for the proceeding computation, g , is defined as $T_g - g_0$. 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 σ_0 :

$$(73) \quad \sigma_0 \equiv \sigma \text{ except:}$$

$$(74) \quad \sigma_0[S(T)]_b \equiv \sigma[S(T)]_b - T_g p$$

$$(75) \quad \sigma_0[S(T)]_n \equiv \sigma[S(T)]_n + 1$$

Evaluating σ_P from σ_0 depends on the transaction type; either contract creation or message call; we define the tuple of post-execution provisional state σ_P , remaining gas g' , accrued substate A and status code z :

$$(76) \quad (\sigma_P, g', A, z) \equiv \begin{cases} \Lambda_4(\sigma_0, A^*, S(T), S(T), g, \\ p, T_v, T_i, 0, \emptyset, \top) & \text{if } T_t = \emptyset \\ \Theta_4(\sigma_0, A^*, S(T), S(T), T_t, \\ T_t, g, p, T_v, T_v, T_d, 0, \top) & \text{otherwise} \end{cases}$$

where

$$(77) \quad A^* \equiv A^0 \text{ except}$$

$$(78) \quad A_{\mathbf{K}}^* \equiv \bigcup_{E \in T_{\mathbf{A}}} \{\forall i < \|E_s\|, i \in \mathbb{N} : (E_a, E_s[i])\}$$

$$(79) \quad A_{\mathbf{a}}^* \equiv \begin{cases} a \cup T_t & \text{if } T_t \neq \emptyset \\ a & \text{otherwise} \end{cases}$$

$$(80) \quad a \equiv A_{\mathbf{a}}^0 \cup S(T) \cup H_c \cup_{E \in T_{\mathbf{A}}} \{E_a\}$$

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

$$(81) \quad g \equiv T_g - g_0$$

Note we use Θ_4 and Λ_4 to denote the fact that only the first four components of the functions' values are taken; the final represents the message-call's output value (a byte array) and is unused in the context of transaction evaluation.

Then 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.

$$(82) \quad g^* \equiv g' + \min \left\{ \left\lfloor \frac{T_g - g'}{5} \right\rfloor, A_r \right\}$$

The total refundable amount is the legitimately remaining gas g' , added to A_r , with the latter component being capped up to a maximum of one fifth⁹ (rounded down) of the total amount used $T_g - g'$. Therefore, g^* is the total gas that remains after the transaction has been executed.

⁹The max refundable proportion of gas was reduced from one half to one fifth by EIP-3529 by Buterin and Swende [2021] in the *London* release

The validator, whose address is specified as the beneficiary of the present block B , receives the gas consumed multiplied by the transaction's *priority fee per gas*, defined as f in this section. The ether that is paid by the transactor that goes toward the base fee is debited from the transactor's account but credited to no other account, so it is burned.

We define the pre-final state σ^* in terms of the provisional state σ_P :

$$(83) \quad \sigma^* \equiv \sigma_P \text{ except}$$

$$(84) \quad \sigma^*[S(T)]_b \equiv \sigma_P[S(T)]_b + g^*p$$

$$(85) \quad \sigma^*[B_{Hc}]_b \equiv \sigma_P[B_{Hc}]_b + (T_g - g^*)f$$

The final state, σ' , is reached after deleting all accounts that either appear in the self-destruct set or are touched and empty:

$$(86) \quad \sigma' \equiv \sigma^* \text{ except}$$

$$(87) \quad \forall i \in A_s : \sigma'[i] = \emptyset$$

$$(88) \quad \forall i \in A_t : \sigma'[i] = \emptyset \text{ if } \text{DEAD}(\sigma^*, i)$$

And finally, we specify Υ^g , the total gas used in this transaction Υ^1 , the logs created by this transaction and Υ^z , the status code of this transaction:

$$(89) \quad \Upsilon^g(\sigma, T) \equiv T_g - g^*$$

$$(90) \quad \Upsilon^1(\sigma, T) \equiv A_1$$

$$(91) \quad \Upsilon^z(\sigma, T) \equiv z$$

These are used to help define the transaction receipt and are also used later for state validation.

7. CONTRACT CREATION

There are a number of intrinsic parameters used when creating an account: sender (s), original transactor¹⁰ (o), available gas (g), effective gas price (p), endowment (v) together with an arbitrary length byte array, \mathbf{i} , the initialisation EVM code, the present depth of the message-call/contract-creation stack (e), the salt for new account's address (ζ) and finally the permission to make modifications to the state (w). The salt ζ might be missing ($\zeta = \emptyset$); formally,

$$(92) \quad \zeta \in \mathbb{B}_{32} \cup \mathbb{B}_0$$

If the creation was caused by CREATE2, then $\zeta \neq \emptyset$.

We define the creation function formally as the function Λ , which evaluates from these values, together with the state σ and the accrued substate A , to the tuple containing the new state, remaining gas, new accrued substate, status code and output $(\sigma', g', A', z, \mathbf{o})$:

$$(93) \quad (\sigma', g', A', z, \mathbf{o}) \equiv \Lambda(\sigma, A, s, o, g, p, v, \mathbf{i}, e, \zeta, w)$$

The address of the new account is defined as being the rightmost 160 bits of the Keccak-256 hash of the RLP encoding of the structure containing only the sender and the account nonce. For CREATE2 the rule is different and is described in EIP-1014 by Buterin [2018]. Combining the two cases, we define the resultant address for the new

account a :

$$(94) \quad a \equiv \text{ADDR}(s, \sigma[s]_n - 1, \zeta, \mathbf{i})$$

$$(95) \quad \text{ADDR}(s, n, \zeta, \mathbf{i}) \equiv \mathcal{B}_{96..255}(\text{KEC}(L_A(s, n, \zeta, \mathbf{i})))$$

$$(96) \quad L_A(s, n, \zeta, \mathbf{i}) \equiv \begin{cases} \text{RLP}((s, n)) & \text{if } \zeta = \emptyset \\ (255) \cdot s \cdot \zeta \cdot \text{KEC}(\mathbf{i}) & \text{otherwise} \end{cases}$$

where \cdot is the concatenation of byte arrays, $\mathcal{B}_{a..b}(X)$ evaluates to a binary value containing the bits of indices in the range $[a, b]$ of the binary data X , and $\sigma[x]$ is the address state of x , or \emptyset 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 address of the new account is added to the set of accessed accounts:

$$(97) \quad A^* \equiv A \text{ except } A_a^* \equiv A_a \cup \{a\}$$

The account's nonce is initially defined as one, 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 σ^* :

$$(98) \quad \sigma^* \equiv \sigma \text{ except:}$$

$$(99) \quad \sigma^*[a] = (1, v + v', \text{TRIE}(\emptyset), \text{KEC}(()))$$

$$(100) \quad \sigma^*[s] = \begin{cases} \emptyset & \text{if } \sigma[s] = \emptyset \wedge v = 0 \\ \mathbf{a}^* & \text{otherwise} \end{cases}$$

$$(101) \quad \mathbf{a}^* \equiv (\sigma[s]_n, \sigma[s]_b - v, \sigma[s]_s, \sigma[s]_c)$$

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

$$(102) \quad v' \equiv \begin{cases} 0 & \text{if } \sigma[a] = \emptyset \\ \sigma[a]_b & \text{otherwise} \end{cases}$$

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 σ^{**} , available gas remaining g^{**} , the resultant accrued substate A^{**} and the body code of the account \mathbf{o} .

$$(103) \quad (\sigma^{**}, g^{**}, A^{**}, \mathbf{o}) \equiv \Xi(\sigma^*, g, A^*, I)$$

where I contains the parameters of the execution environment, that is:

¹⁰which 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

$$(104) \quad I_a \equiv a$$

$$(105) \quad I_o \equiv o$$

$$(106) \quad I_p \equiv p$$

$$(107) \quad I_d \equiv ()$$

$$(108) \quad I_s \equiv s$$

$$(109) \quad I_v \equiv v$$

$$(110) \quad I_b \equiv \mathbf{i}$$

$$(111) \quad I_e \equiv e$$

$$(112) \quad I_w \equiv w$$

I_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, \emptyset , 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:

$$(113) \quad c \equiv G_{\text{codedeposit}} \times \|\mathbf{o}\|$$

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, thus the contract's code is not stored.

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, accrued substate and status code as (σ', g', A', z) where:

$$(114)$$

$$g' \equiv \begin{cases} 0 & \text{if } F \\ g^{**} - c & \text{otherwise} \end{cases}$$

$$(115)$$

$$\sigma' \equiv \begin{cases} \sigma & \text{if } F \vee \sigma^{**} = \emptyset \\ \sigma^{**} \text{ except:} & \\ \sigma'[a] = \emptyset & \text{if } \text{DEAD}(\sigma^{**}, a) \\ \sigma^{**} \text{ except:} & \\ \sigma'[a]_c = \text{KEC}(\mathbf{o}) & \text{otherwise} \end{cases}$$

$$(116)$$

$$A' \equiv \begin{cases} A^* & \text{if } F \vee \sigma^{**} = \emptyset \\ A^{**} & \text{otherwise} \end{cases}$$

$$(117)$$

$$z \equiv \begin{cases} 0 & \text{if } F \vee \sigma^{**} = \emptyset \\ 1 & \text{otherwise} \end{cases}$$

where

$$(118)$$

$$\begin{aligned} F &\equiv (\sigma[a] \neq \emptyset \wedge (\sigma[a]_c \neq \text{KEC}(\emptyset) \vee \sigma[a]_n \neq 0)) \vee \\ &(\sigma^{**} = \emptyset \wedge \mathbf{o} = \emptyset) \vee \\ &g^{**} < c \vee \\ &\|\mathbf{o}\| > 24576 \vee \\ &\mathbf{o}[0] = \mathbf{0xef} \end{aligned}$$

Note the last condition of F indicates that contract code cannot begin with the byte $\mathbf{0xef}$ (refer to EIP-3541 by Beregszaszi et al. [2021]),

The exception in the determination of σ' 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 intention is that the result is either a successfully created new contract with its endowment, or no new contract with no transfer of value. In addition, observe that if the execution of the initialising code reverts ($\sigma^{**} = \emptyset \wedge \mathbf{o} \neq \emptyset$), the resultant gas g' is not depleted (provided there was no other exception), but no new account is created.

7.1. Subtleties. Note that while the initialisation code is executing, the newly created address exists but with no intrinsic body code¹¹. 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

¹¹During initialization code execution, `EXTCODESIZE` on the address should return zero, which is the length of the code of the account while `CODESIZE` should return the length of the initialization code (as defined in H.2).

(c , usually the same as recipient), available gas (g), value (v) and effective gas price (p) together with an arbitrary length byte array, \mathbf{d} , the input data of the call, the present depth of the message-call/contract-creation stack (e) and finally the permission to make modifications to the state (w).

Aside from evaluating to a new state and accrued 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.

$$(119) \quad (\sigma', g', A', z, \mathbf{o}) \equiv \Theta(\sigma, A, s, o, r, c, g, p, v, \tilde{v}, \mathbf{d}, e, w)$$

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 σ_1 , the first transitional state as the original state but with the value transferred from sender to recipient:

$$(120) \quad \sigma_1[r]_{\mathbf{b}} \equiv \sigma[r]_{\mathbf{b}} + v \quad \wedge \quad \sigma_1[s]_{\mathbf{b}} \equiv \sigma[s]_{\mathbf{b}} - v$$

unless $s = r$.

Throughout the present work, it is assumed that if $\sigma_1[r]$ 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:

$$(121) \quad \sigma_1 \equiv \sigma'_1 \quad \text{except:}$$

$$(122) \quad \sigma_1[s] \equiv \begin{cases} \emptyset & \text{if } \sigma'_1[s] = \emptyset \wedge v = 0 \\ \mathbf{a}_1 & \text{otherwise} \end{cases}$$

$$(123) \quad \mathbf{a}_1 \equiv (\sigma'_1[s]_{\mathbf{n}}, \sigma'_1[s]_{\mathbf{b}} - v, \sigma'_1[s]_{\mathbf{s}}, \sigma'_1[s]_{\mathbf{c}})$$

$$(124) \quad \text{and } \sigma'_1 \equiv \sigma \quad \text{except:}$$

$$(125) \quad \begin{cases} \sigma'_1[r] \equiv (0, v, \text{TRIE}(\emptyset), \text{KEC}(())) & \text{if } \sigma[r] = \emptyset \wedge v \neq 0 \\ \sigma'_1[r] \equiv \emptyset & \text{if } \sigma[r] = \emptyset \wedge v = 0 \\ \sigma'_1[r] \equiv \mathbf{a}'_1 & \text{otherwise} \end{cases}$$

$$(126) \quad \mathbf{a}'_1 \equiv (\sigma[r]_{\mathbf{n}}, \sigma[r]_{\mathbf{b}} + v, \sigma[r]_{\mathbf{s}}, \sigma[r]_{\mathbf{c}})$$

The account's associated code (identified as the fragment whose Keccak-256 hash is $\sigma[c]_{\mathbf{c}}$) 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. σ).

$$(127) \quad \sigma' \equiv \begin{cases} \sigma & \text{if } \sigma^{**} = \emptyset \\ \sigma^{**} & \text{otherwise} \end{cases}$$

$$(128) \quad g' \equiv \begin{cases} 0 & \text{if } \sigma^{**} = \emptyset \wedge \\ & \mathbf{o} = \emptyset \\ g^{**} & \text{otherwise} \end{cases}$$

$$(129) \quad A' \equiv \begin{cases} A & \text{if } \sigma^{**} = \emptyset \\ A^{**} & \text{otherwise} \end{cases}$$

$$(130) \quad z \equiv \begin{cases} 0 & \text{if } \sigma^{**} = \emptyset \\ 1 & \text{otherwise} \end{cases}$$

$$(131) \quad (\sigma^{**}, g^{**}, A^{**}, \mathbf{o}) \equiv \Xi$$

$$(132) \quad I_{\mathbf{a}} \equiv r$$

$$(133) \quad I_{\mathbf{o}} \equiv o$$

$$(134) \quad I_{\mathbf{p}} \equiv p$$

$$(135) \quad I_{\mathbf{d}} \equiv \mathbf{d}$$

$$(136) \quad I_{\mathbf{s}} \equiv s$$

$$(137) \quad I_{\mathbf{v}} \equiv \tilde{v}$$

$$(138) \quad I_{\mathbf{e}} \equiv e$$

$$(139) \quad I_{\mathbf{w}} \equiv w$$

where

$$(140) \quad \Xi \equiv \begin{cases} \Xi_{\text{ECCREC}}(\sigma_1, g, A, I) & \text{if } c = 1 \\ \Xi_{\text{SHA256}}(\sigma_1, g, A, I) & \text{if } c = 2 \\ \Xi_{\text{RIP160}}(\sigma_1, g, A, I) & \text{if } c = 3 \\ \Xi_{\text{ID}}(\sigma_1, g, A, I) & \text{if } c = 4 \\ \Xi_{\text{EXPMOD}}(\sigma_1, g, A, I) & \text{if } c = 5 \\ \Xi_{\text{BN_ADD}}(\sigma_1, g, A, I) & \text{if } c = 6 \\ \Xi_{\text{BN_MUL}}(\sigma_1, g, A, I) & \text{if } c = 7 \\ \Xi_{\text{SNARKV}}(\sigma_1, g, A, I) & \text{if } c = 8 \\ \Xi_{\text{BLAKE2F}}(\sigma_1, g, A, I) & \text{if } c = 9 \\ \Xi(\sigma_1, g, A, I) & \text{otherwise} \end{cases}$$

and

$$(141) \quad \text{KEC}(I_{\mathbf{b}}) = \sigma[c]_{\mathbf{c}}$$

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

As can be seen, there are nine exceptions to the usage of the general execution framework Ξ for evaluation of the message call: these are so-called ‘precompiled’ contracts, meant as a preliminary piece of architecture that may later become *native extensions*. The contracts in addresses 1 to 9 execute the elliptic curve public key recovery function, the SHA2 256-bit hash scheme, the RIPEMD 160-bit hash scheme, the identity function, arbitrary precision modular exponentiation, elliptic curve addition, elliptic curve scalar multiplication, an elliptic curve pairing check, and the BLAKE2 compression function F respectively. Their full formal definition is in Appendix E. We denote the set of the addresses of the precompiled contracts by π :

$$(142) \quad \pi \equiv \{1, 2, 3, 4, 5, 6, 7, 8, 9\}$$

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 items) 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, CREATE2, 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 smallest multiple of 32 bytes that are 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 σ , the remaining gas for computation g , and the accrued substate A , 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_a , the address of the account which owns the code that is executing.
- I_o , the sender address of the transaction that originated this execution.
- I_p , the price of gas paid by the signer of the transaction that originated this execution. This is defined as the effective gas price p in section 6.
- I_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_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_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_b , the byte array that is the machine code to be executed.
- I_H , the block header of the present block.
- I_e , the depth of the present message-call or contract-creation (i.e. the number of CALLs or CREATE(2)s being executed at present).
- I_w , the permission to make modifications to the state.

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

$$(143) \quad (\sigma', g', A', \mathbf{o}) \equiv \Xi(\sigma, g, A, I)$$

where we will remember that A , the accrued substate, is defined in section 6.1.

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, σ and the machine state, μ . 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 \emptyset ; this is important when interpreting the output of H , which evaluates to \emptyset when execution is to continue but a series (potentially empty) when execution

should halt.

$$\begin{aligned}
 (144) \quad \Xi(\sigma, g, A, I) &\equiv (\sigma', \mu'_g, A', \mathbf{o}) \\
 (145) \quad (\sigma', \mu', A', \dots, \mathbf{o}) &\equiv X((\sigma, \mu, A, I)) \\
 (146) \quad \mu_g &\equiv g \\
 (147) \quad \mu_{pc} &\equiv 0 \\
 (148) \quad \mu_m &\equiv (0, 0, \dots) \\
 (149) \quad \mu_i &\equiv 0 \\
 (150) \quad \mu_s &\equiv () \\
 (151) \quad \mu_o &\equiv ()
 \end{aligned}$$

$$(152) \quad X((\sigma, \mu, A, I)) \equiv \begin{cases} (\emptyset, \mu, A, I, \emptyset) & \text{if } Z(\sigma, \mu, A, I) \\ (\emptyset, \mu', A, I, \mathbf{o}) & \text{if } w = \text{REVERT} \\ O(\sigma, \mu, A, I) \cdot \mathbf{o} & \text{if } \mathbf{o} \neq \emptyset \\ X(O(\sigma, \mu, A, I)) & \text{otherwise} \end{cases}$$

where

$$\begin{aligned}
 (153) \quad \mathbf{o} &\equiv H(\mu, I) \\
 (154) \quad (a, b, c, d) \cdot e &\equiv (a, b, c, d, e) \\
 (155) \quad \mu' &\equiv \mu \text{ except:} \\
 (156) \quad \mu'_g &\equiv \mu_g - C(\sigma, \mu, A, I)
 \end{aligned}$$

Note that, when we evaluate Ξ , we drop the fourth element I' and extract the remaining gas μ'_g from the resultant machine state μ' .

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 μ is defined as the tuple $(g, pc, \mathbf{m}, i, \mathbf{s}, \mathbf{o})$ which are the gas available, the program counter $pc \in \mathbb{N}_{256}$, the memory contents, the active number of words in memory (counting continuously from position 0), the stack contents, and the returndata buffer. The memory contents μ_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:

$$(157) \quad w \equiv \begin{cases} I_b[\mu_{pc}] & \text{if } \mu_{pc} < \|I_b\| \\ \text{STOP} & \text{otherwise} \end{cases}$$

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:

$$(158) \quad Z(\sigma, \mu, A, I) \equiv \begin{aligned} &\mu_g < C(\sigma, \mu, A, I) \quad \vee \\ &\delta_w = \emptyset \quad \vee \\ &\|\mu_s\| < \delta_w \quad \vee \\ &(w = \text{JUMP} \wedge \mu_s[0] \notin D(I_b)) \quad \vee \\ &(w = \text{JUMPI} \wedge \mu_s[1] \neq 0 \wedge \\ &\quad \mu_s[0] \notin D(I_b)) \quad \vee \\ &(w = \text{RETURNDATACOPY} \wedge \\ &\quad \mu_s[1] + \mu_s[2] > \|\mu_o\|) \quad \vee \\ &\|\mu_s\| - \delta_w + \alpha_w > 1024 \quad \vee \\ &(\neg I_w \wedge W(w, \mu)) \quad \vee \\ &(w = \text{SSTORE} \wedge \mu_g \leq G_{\text{callstipend}}) \end{aligned}$$

where

$$(159) \quad W(w, \mu) \equiv \begin{aligned} &w \in \{\text{CREATE}, \text{CREATE2}, \text{SSTORE}, \\ &\text{SELFDESTRUCT}\} \vee \\ &\text{LOG0} \leq w \wedge w \leq \text{LOG4} \quad \vee \\ &w = \text{CALL} \wedge \mu_s[2] \neq 0 \end{aligned}$$

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, the new stack size would be larger than 1024 or state modification is attempted during a static call. The astute reader will realise that this implies that no instruction can, through its execution, cause an exceptional halt. Also, the execution is in an exceptional halting state if the gas left prior to executing an SSTORE instruction is less than or equal to the call stipend $G_{\text{callstipend}}$ – see EIP-2200 by Tang [2019] for more information.

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:

$$(160) \quad D(\mathbf{c}) \equiv D_J(\mathbf{c}, 0)$$

where:

$$(161) \quad D_J(\mathbf{c}, i) \equiv \begin{cases} \{\} & \text{if } i \geq \|\mathbf{c}\| \\ \{i\} \cup D_J(\mathbf{c}, N(i, \mathbf{c}[i])) & \\ \quad \text{if } \mathbf{c}[i] = \text{JUMPDEST} & \\ D_J(\mathbf{c}, N(i, \mathbf{c}[i])) & \text{otherwise} \end{cases}$$

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

$$(162) \quad N(i, w) \equiv \begin{cases} i + w - \text{PUSH1} + 2 & \\ \quad \text{if } w \in [\text{PUSH1}, \text{PUSH32}] & \\ i + 1 & \text{otherwise} \end{cases}$$

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

$$(163) \quad H(\boldsymbol{\mu}, I) \equiv \begin{cases} H_{\text{RETURN}}(\boldsymbol{\mu}) & \text{if } w \in \{\text{RETURN}, \text{REVERT}\} \\ () & \text{if } w \in \{\text{STOP}, \text{SELFDSTRUCT}\} \\ \emptyset & \text{otherwise} \end{cases}$$

The data-returning halt operations, RETURN and REVERT, have 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:

$$(164) \quad O((\boldsymbol{\sigma}, \boldsymbol{\mu}, A, I)) \equiv (\boldsymbol{\sigma}', \boldsymbol{\mu}', A', I)$$

$$(165) \quad \Delta \equiv \alpha_w - \delta_w$$

$$(166) \quad \|\boldsymbol{\mu}'_s\| \equiv \|\boldsymbol{\mu}_s\| + \Delta$$

$$(167) \quad \forall x \in [\alpha_w, \|\boldsymbol{\mu}'_s\|] : \boldsymbol{\mu}'_s[x] \equiv \boldsymbol{\mu}_s[x - \Delta]$$

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:

$$(168) \quad \boldsymbol{\mu}'_g \equiv \boldsymbol{\mu}_g - C(\boldsymbol{\sigma}, \boldsymbol{\mu}, A, I)$$

$$(169) \quad \boldsymbol{\mu}'_{pc} \equiv \begin{cases} J_{\text{JUMP}}(\boldsymbol{\mu}) & \text{if } w = \text{JUMP} \\ J_{\text{JUMPI}}(\boldsymbol{\mu}) & \text{if } w = \text{JUMPI} \\ N(\boldsymbol{\mu}_{pc}, w) & \text{otherwise} \end{cases}$$

In general, we assume the memory, accrued substate and system state do not change:

$$(170) \quad \boldsymbol{\mu}'_m \equiv \boldsymbol{\mu}_m$$

$$(171) \quad \boldsymbol{\mu}'_i \equiv \boldsymbol{\mu}_i$$

$$(172) \quad A' \equiv A$$

$$(173) \quad \boldsymbol{\sigma}' \equiv \boldsymbol{\sigma}$$

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. TRANSITION TO PROOF OF STAKE

The *Paris* hard fork changed the underlying consensus mechanism of Ethereum from proof of work to proof of stake.

Unlike all previous hard forks of Ethereum, *Paris* was not defined to occur at a particular block height, but rather after a specified *terminal total difficulty* was reached. Total difficulty was used instead of block height to avoid a scenario in which a minority of hash power could create a malicious fork that could race to satisfy the block height requirement and claim the first proof of stake block.

Thus the *terminal block*, the last proof of work block before the *Paris* fork takes effect, is defined as having:

$$(174) \quad B_t \geq 5875000000000000000000$$

$$(175) \quad P(B_H)_t < 5875000000000000000000$$

where B_t is the total difficulty of block B and $P(B_H)_t$ is the total difficulty of its parent.

Total difficulty for a proof of work (pre-*Paris*) block B is defined recursively as:

$$(176) \quad B_t \equiv P(B_H)_t + H_d$$

where H_d is the difficulty of the current block B .

Upon reaching the terminal block, new blocks are processed by the Beacon Chain.

10.1. **Post-Paris Updates.** Because the Beacon Chain generate a new slot every 12 seconds, post-*Paris* updates can be scheduled to occur at a specific timestamp. At the execution layer, the update will then happen in the first produced block after the scheduled timestamp. For example the *Shanghai* hard fork was scheduled to occur at 2023-04-12 10:27:35 UTC, on Epoch 194,048. Validators failed to propose a block during the two first slots of this Epoch, but at slot 6,209,538 a validator finally proposed the block 17,034,870, marking the transition to *Shanghai* on the execution layer.

11. BLOCKTREE TO BLOCKCHAIN

Prior to the transition to proof of stake at the *Paris* hard fork, the canonical blockchain was defined as the block path with the greatest *total difficulty*, defined in section 10 as B_t .

After reaching the *terminal block* described in section 10, the greatest *total difficulty* rule must be removed in favor of a rule known as *LMD Ghost*.¹²

Note that in order to determine what blocks comprise the canonical Ethereum blockchain after *Paris*, one must have additional information from the Beacon Chain, which is not described herein. We denote events emitted by the Beacon Chain with the prefix POS_.

On each occurrence of a POS_FORKCHOICE_UPDATED event, starting with the first at the *transition block* described in section 10, the canonical chain is defined as the chain beginning with the genesis block and ending at the block nominated by the event as the head of the chain.

The head of the chain should be updated if and only if a POS_FORKCHOICE_UPDATED is emitted, in which case the head should set to the block specified by the event. No optimistic updates to the head of the chain should be made.

The POS_FORKCHOICE_UPDATED event additionally references a finalized block. The most recent finalized block should be set to this block.

The canonical blockchain must also contain a block with the hash and number of the *terminal block* defined in section 10.

12. BLOCK FINALISATION

The process of finalising a block involves three stages:

- (1) executing withdrawals;
- (2) validate transactions;
- (3) verify state.

¹²LMD GHOST comprises two acronyms, "Latest Message Driven", and "Greedy Heaviest-Observed Sub-Tree".

12.1. Executing Withdrawals. After processing the block’s transactions, the withdrawals are executed. A withdrawal is simply an increase of the recipient account’s balance of the specified Gwei amount. No other balances are decreased, a withdrawal is not a transfer but a creation of funds. A withdrawal operation cannot fail and has no gas cost. We define the function E as the withdrawal state transition function:

$$(177) \quad E(\sigma_w, W) \equiv \sigma_{w+1}$$

$$(178) \quad \sigma_{w+1} \equiv \sigma_w \text{ except:}$$

$$(179) \quad \sigma_{w+1}[W_r]_b \equiv \sigma_w[W_r]_b + (W_a \times 10^9)$$

Finally, we define K , the block-level state transition function for withdrawals:

$$(180) \quad K(\sigma, B) \equiv E(E(\sigma, W_0), W_1) \dots$$

12.2. Transaction Validation. The given `gasUsed` must correspond faithfully to the transactions listed: B_{Hg} , the total gas used in the block, must be equal to the accumulated gas used according to the final transaction:

$$(181) \quad B_{\text{Hg}} = \ell(\mathbf{R})_u$$

12.3. State Validation. We may now define the function, Γ , that maps a block B to its initiation state:

$$(182) \quad \Gamma(B) \equiv \begin{cases} \sigma_0 & \text{if } P(B_{\text{H}}) = \emptyset \\ \sigma_i : \text{TRIE}(L_S(\sigma_i)) = P(B_{\text{H}})_{\text{H}_r} & \text{otherwise} \end{cases}$$

Here, $\text{TRIE}(L_S(\sigma_i))$ means the hash of the root node of a trie of state σ_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' :

$$(183) \quad \Phi(B) \equiv B' : B' = B \text{ except:}$$

$$(184) \quad B'_{\text{H}_r} = \text{TRIE}(L_S(K(\Pi(\Gamma(B), B), B)))$$

As specified at the beginning of the present work, Π is the state-transition function, which is defined in terms of Υ , the transaction-evaluation function.

As previously detailed, $\mathbf{R}[n]_z$, $\mathbf{R}[n]_1$ and $\mathbf{R}[n]_u$ are the n th corresponding status code, logs and cumulative gas used after each transaction ($\mathbf{R}[n]_b$, the fourth component in the tuple, has already been defined in terms of the logs). We also define the n th state $\sigma[n]$, which 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):

$$(185) \quad \sigma[n] = \begin{cases} \Gamma(B) & \text{if } n < 0 \\ \Upsilon(\sigma[n-1], B_{\text{T}}[n]) & \text{otherwise} \end{cases}$$

In the case of $B_{\mathbf{R}}[n]_u$, 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:

$$(186) \quad \mathbf{R}[n]_u = \begin{cases} 0 & \text{if } n < 0 \\ \Upsilon^g(\sigma[n-1], B_{\text{T}}[n]) \\ \quad + \mathbf{R}[n-1]_u & \text{otherwise} \end{cases}$$

For $\mathbf{R}[n]_1$, we utilise the Υ^1 function that we conveniently defined in the transaction execution function.

$$(187) \quad \mathbf{R}[n]_1 = \Upsilon^1(\sigma[n-1], B_{\text{T}}[n])$$

We define $\mathbf{R}[n]_z$ in a similar manner.

$$(188) \quad \mathbf{R}[n]_z = \Upsilon^z(\sigma[n-1], B_{\text{T}}[n])$$

Finally, we define Π as the final transaction’s resultant state, $\ell(\sigma)$:

$$(189) \quad \Pi(\sigma, B) \equiv \ell(\sigma)$$

Thus the complete block-transition mechanism (before consensus) is defined.

13. IMPLEMENTING CONTRACTS

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

13.1. 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.

13.2. 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 and the block’s beneficiary address. In order to make it hard for malicious validators 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.

14. 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 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.

15. CONCLUSION

We 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.

16. 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.

17. 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

- 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>.
- Alex Beregszaszi, Paweł Bylica, Andrei Maiboroda, Alexey Akhunov, Christian Reitwiessner, and Martin Swende. EIP-3541: Reject new contract code starting with the 0xef byte, March 2021. URL <https://eips.ethereum.org/EIPS/eip-3541>.
- Guido Bertoni, Joan Daemen, Michal Peeters, and Gilles Van Assche. The KECCAK SHA-3 submission, 2011. URL <https://keccak.team/files/Keccak-submission-3.pdf>.
- Roman Boutellier and Mareike Heinen. Pirates, Pioneers, Innovators and Imitators. In *Growth Through Innovation*, pages 85–96. Springer, 2014. URL <https://www.springer.com/gb/book/9783319040158>.
- Vitalik Buterin. Ethereum: A Next-Generation Smart Contract and Decentralized Application Platform, 2013. URL <https://github.com/ethereum/wiki/wiki/White-Paper>.
- Vitalik Buterin. EIP-2: Homestead hard-fork changes, 2015. URL <https://eips.ethereum.org/EIPS/eip-2>.
- Vitalik Buterin. EIP-155: Simple replay attack protection, October 2016. URL <https://eips.ethereum.org/EIPS/eip-155>.
- Vitalik Buterin. EIP-1014: Skinny CREATE2, April 2018. URL <https://eips.ethereum.org/EIPS/eip-1014>.
- Vitalik Buterin and Martin Swende. EIP-2929: Gas cost increases for state access opcodes, September 2020a. URL <https://eips.ethereum.org/EIPS/eip-2929>.
- Vitalik Buterin and Martin Swende. EIP-2930: Optional access lists, August 2020b. URL <https://eips.ethereum.org/EIPS/eip-2930>.
- Vitalik Buterin and Martin Swende. EIP-3529: Reduction in refunds, April 2021. URL <https://eips.ethereum.org/EIPS/eip-3529>.
- Vitalik Buterin, Eric Conner, Rick Dudley, Matthew Slipper, Ian Norden, and Abdelhamid Bakhta. EIP-1559: Fee market change for eth 1.0 chain, 2019. URL <https://eips.ethereum.org/EIPS/eip-1559>.
- Nicolas T. Courtois, Marek Grajek, and Rahul Naik. *Optimizing SHA256 in Bitcoin Mining*, pages 131–144. Springer Berlin Heidelberg, Berlin, Heidelberg, 2014. ISBN 978-3-662-44893-9. doi: 10.1007/978-3-662-44893-9_12. URL https://doi.org/10.1007/978-3-662-44893-9_12.
- B.A. Davey and H.A. Priestley. *Introduction to lattices and order*. 2nd ed. Cambridge: Cambridge University Press, 2nd ed. edition, 2002. ISBN 0-521-78451-4/pbk.
- 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>.
- Dankrad Feist, Dmitry Khovratovich, and Marius van der Wijden. EIP-3607: Reject transactions from senders with deployed code, June 2021. URL <https://eips.ethereum.org/EIPS/eip-3607>.
- Phong Vo Glenn Fowler, Landon Curt Noll. FowlerNoll-IVo hash function, 1991. URL <http://www.isthe.com/chongo/tech/comp/fnv/index.html>.
- 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>.
- Tjaden Hess, Matt Luongo, Piotr Dyraga, and James Hancock. EIP-152: Add BLAKE2 compression function 'F' precompile, October 2016. URL <https://eips.ethereum.org/EIPS/eip-152>.
- Don Johnson, Alfred Menezes, and Scott Vanstone. The Elliptic Curve Digital Signature Algorithm (ECDSA), 2001. URL <https://web.archive.org/web/20170921160141/http://cs.ucsb.edu/~koc/ccs130h/notes/ecdsa-cert.pdf>. Accessed 21 September 2017, but the original link was inaccessible on 19 October 2017. Refer to section 6.2 for ECDSAPUBKEY, and section 7 for ECDSASIGN and ECDSARECOVER.
- Sergio Demian Lerner. Strict Memory Hard Hashing Functions, 2014. URL <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/0BwOVXJKBgYPMSOJ2VGIyWWlocms/edit?usp=sharing>.
- Satoshi Nakamoto. Bitcoin: A peer-to-peer electronic cash system, 2008. URL <http://www.bitcoin.org/bitcoin.pdf>.
- Meni Rosenfeld, Yoni Assia, Vitalik Buterin, miorhakiLior, Oded Leiba, Assaf Shomer, and Eli-ran Zach. Colored Coins Protocol Specification, 2012. URL <https://github.com/Colored-Coins/Colored-Coins-Protocol-Specification>.
- Markku-Juhani Saarinen and Jean-Philippe Aumasson. RFC 7693: The BLAKE2 cryptographic hash and message authentication code (MAC), November 2015. URL <https://tools.ietf.org/html/rfc7693>.
- Simon Sprankel. Technical Basis of Digital Currencies, 2013. URL <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>.
- Wei Tang. EIP-2200: Structured definitions for net gas metering, 2019. URL <https://eips.ethereum.org/EIPS/eip-2200>.
- Vivek Vishnumurthy, Sangeeth Chandrakumar, and Emin Gn Sirer. KARMA: A secure economic framework for peer-to-peer resource sharing, 2003. URL <https://www.cs.cornell.edu/people/egs/papers/karma.pdf>.
- J. R. Willett. MasterCoin Complete Specification, 2013. URL <https://github.com/mastercoin-MSC/spec>.
- Micah Zoltu. EIP-2718: Typed transaction envelope, June 2020. URL <https://eips.ethereum.org/EIPS/eip-2718>.

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 validators 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.
- 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} :

$$(190) \quad \mathbb{T} \equiv \mathbb{L} \uplus \mathbb{B}$$

$$(191) \quad \mathbb{L} \equiv \{\mathbf{t} : \mathbf{t} = (\mathbf{t}[0], \mathbf{t}[1], \dots) \wedge \forall n < \|\mathbf{t}\| : \mathbf{t}[n] \in \mathbb{T}\}$$

$$(192) \quad \mathbb{B} \equiv \{\mathbf{b} : \mathbf{b} = (\mathbf{b}[0], \mathbf{b}[1], \dots) \wedge \forall n < \|\mathbf{b}\| : \mathbf{b}[n] \in \mathbb{O}\}$$

Where \mathbb{O} is the set of (8-bit) 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. The disjoint union \uplus is needed only to distinguish the empty byte array $() \in \mathbb{B}$ from the empty list $() \in \mathbb{L}$, which are encoded differently as defined below; as common, we will abuse notation and leave the disjoint union indices implicit, inferable from context.

We define the RLP function as 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:

$$(193) \quad \text{RLP}(\mathbf{x}) \equiv \begin{cases} R_b(\mathbf{x}) & \text{if } \mathbf{x} \in \mathbb{B} \\ R_l(\mathbf{x}) & \text{otherwise} \end{cases}$$

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, provided that it contains fewer than 2^{64} bytes, 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.

Byte arrays containing 2^{64} or more bytes cannot be encoded. This restriction ensures that the first byte of the encoding of a byte array is always below 192, and thus it can be readily distinguished from the encodings of sequences in \mathbb{L} .

Formally, we define R_b :

$$(194) \quad R_b(\mathbf{x}) \equiv \begin{cases} \mathbf{x} & \text{if } \|\mathbf{x}\| = 1 \wedge \mathbf{x}[0] < 128 \\ (128 + \|\mathbf{x}\|) \cdot \mathbf{x} & \text{else if } \|\mathbf{x}\| < 56 \\ (183 + \|\text{BE}(\|\mathbf{x}\|)\|) \cdot \text{BE}(\|\mathbf{x}\|) \cdot \mathbf{x} & \text{else if } \|\mathbf{x}\| < 2^{64} \\ \emptyset & \text{otherwise} \end{cases}$$

$$(195) \quad \text{BE}(x) \equiv (b_0, b_1, \dots) : b_0 \neq 0 \wedge x = \sum_{n=0}^{\|\mathbf{b}\|-1} b_n \cdot 256^{\|\mathbf{b}\|-1-n}$$

$$(196) \quad (x_1, \dots, x_n) \cdot (y_1, \dots, y_m) = (x_1, \dots, x_n, y_1, \dots, y_m)$$

Thus BE is the function that expands a non-negative 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, provided that they contain fewer than 2^{64} bytes, 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.

Sequences whose concatenated serialized items contain 2^{64} or more bytes cannot be encoded. This restriction ensures that the first byte of the encoding does not exceed 255 (otherwise it would not be a byte).

Thus we finish by formally defining R_l :

$$(197) \quad R_l(\mathbf{x}) \equiv \begin{cases} (192 + \|\mathbf{s}(\mathbf{x})\|) \cdot \mathbf{s}(\mathbf{x}) & \text{if } \mathbf{s}(\mathbf{x}) \neq \emptyset \wedge \|\mathbf{s}(\mathbf{x})\| < 56 \\ (247 + \|\text{BE}(\|\mathbf{s}(\mathbf{x})\|)\|) \cdot \text{BE}(\|\mathbf{s}(\mathbf{x})\|) \cdot \mathbf{s}(\mathbf{x}) & \text{else if } \mathbf{s}(\mathbf{x}) \neq \emptyset \wedge \|\mathbf{s}(\mathbf{x})\| < 2^{64} \\ \emptyset & \text{otherwise} \end{cases}$$

$$(198) \quad \mathbf{s}(\mathbf{x}) \equiv \begin{cases} \text{RLP}(\mathbf{x}[0]) \cdot \text{RLP}(\mathbf{x}[1]) \cdot \dots & \text{if } \forall i : \text{RLP}(\mathbf{x}[i]) \neq \emptyset \\ \emptyset & \text{otherwise} \end{cases}$$

If RLP is used to encode a scalar, defined only as a non-negative integer (in \mathbb{N} , or in \mathbb{N}_x for any x), it must be encoded as the shortest byte array whose big-endian interpretation is the scalar. Thus the RLP of some non-negative integer i is defined as:

$$(199) \quad \text{RLP}(i : i \in \mathbb{N}) \equiv \text{RLP}(\text{BE}(i))$$

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}):

$$(200) \quad \text{HP}(\mathbf{x}, t) : \mathbf{x} \in \mathbb{Y} \equiv \begin{cases} (16f(t), 16\mathbf{x}[0] + \mathbf{x}[1], 16\mathbf{x}[2] + \mathbf{x}[3], \dots) & \text{if } \|\mathbf{x}\| \text{ is even} \\ (16(f(t) + 1) + \mathbf{x}[0], 16\mathbf{x}[1] + \mathbf{x}[2], 16\mathbf{x}[3] + \mathbf{x}[4], \dots) & \text{otherwise} \end{cases}$$

$$(201) \quad f(t) \equiv \begin{cases} 2 & \text{if } t \neq 0 \\ 0 & \text{otherwise} \end{cases}$$

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 \mathcal{J} , a set containing pairs of byte sequences with unique keys:

$$(202) \quad \mathcal{J} = \{(\mathbf{k}_0 \in \mathbb{B}, \mathbf{v}_0 \in \mathbb{B}), (\mathbf{k}_1 \in \mathbb{B}, \mathbf{v}_1 \in \mathbb{B}), \dots\}$$

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

$$(203) \quad \forall I \in \mathcal{J} : I \equiv (I_0, I_1)$$

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:

$$(204) \quad y(\mathcal{J}) = \{(\mathbf{k}'_0 \in \mathbb{Y}, \mathbf{v}_0 \in \mathbb{B}), (\mathbf{k}'_1 \in \mathbb{Y}, \mathbf{v}_1 \in \mathbb{B}), \dots\}$$

$$(205) \quad \forall n : \forall i < 2\|\mathbf{k}_n\| : \mathbf{k}'_n[i] \equiv \begin{cases} \lfloor \mathbf{k}_n[i \div 2] \div 16 \rfloor & \text{if } i \text{ is even} \\ \mathbf{k}_n[\lfloor i \div 2 \rfloor] \bmod 16 & \text{otherwise} \end{cases}$$

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

$$(206) \quad \text{TRIE}(\mathcal{J}) \equiv \text{KEC}(\text{RLP}(c(\mathcal{J}, 0)))$$

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, we store nodes whose composed RLP is fewer than 32 bytes directly; for those larger we assert prescience of the byte array whose Keccak-256 hash evaluates to our reference. Thus we define in terms of c , the node composition function:

$$(207) \quad n(\mathcal{J}, i) \equiv \begin{cases} () \in \mathbb{B} & \text{if } \mathcal{J} = \emptyset \\ c(\mathcal{J}, i) & \text{if } \|\text{RLP}(c(\mathcal{J}, i))\| < 32 \\ \text{KEC}(\text{RLP}(c(\mathcal{J}, i))) & \text{otherwise} \end{cases}$$

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 1.

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 0.

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 :

$$(208) \quad c(\mathcal{J}, i) \equiv \begin{cases} (\text{HP}(I_0[i..(\|I_0\| - 1)], 1), I_1) & \text{if } \|\mathcal{J}\| = 1 \text{ where } \exists I : I \in \mathcal{J} \\ (\text{HP}(I_0[i..(j - 1)], 0), n(\mathcal{J}, j)) & \text{if } i \neq j \text{ where } j = \max\{x : \exists I : \|I\| = x \wedge \forall I \in \mathcal{J} : I_0[0..(x - 1)] = \mathbf{1}\} \\ (u(0), u(1), \dots, u(15), v) & \text{otherwise where } u(j) \equiv n(\{I : I \in \mathcal{J} \wedge I_0[i] = j\}, i + 1) \\ & v = \begin{cases} I_1 & \text{if } \exists I : I \in \mathcal{J} \wedge \|I_0\| = i \\ () \in \mathbb{B} & \text{otherwise} \end{cases} \end{cases}$$

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 \mathcal{J} to a 32-byte hash and assert that only a single such hash exists for any \mathcal{J} , 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.

$$(209) \quad \Xi_{\text{PRE}}(\sigma, g, A, I) \equiv \begin{cases} (\emptyset, 0, A, ()) & \text{if } g < g_r \\ (\sigma, g - g_r, A, \mathbf{o}) & \text{otherwise} \end{cases}$$

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

We define Ξ_{ECCREC} as a precompiled contract for the elliptic curve digital signature algorithm (ECDSA) public key recovery function (ecrecover). See Appendix F for the definition of the function `ECDSARECOVER` and the constant `secp256k1n`. We also define \mathbf{d} to be the input data, well-defined for an infinite length by appending zeroes as required. In the case of an invalid signature, we return no output.

$$(210) \quad \Xi_{\text{ECCREC}} \equiv \Xi_{\text{PRE}} \text{ where:}$$

$$(211) \quad g_r = 3000$$

$$(212) \quad \|\mathbf{o}\| = \begin{cases} 0 & \text{if } v \notin \{27, 28\} \vee r = 0 \vee r \geq \text{secp256k1n} \vee s = 0 \vee s \geq \text{secp256k1n} \\ 0 & \text{if } \text{ECDSARECOVER}(h, v - 27, r, s) = \emptyset \\ 32 & \text{otherwise} \end{cases}$$

$$(213) \quad \text{if } \|\mathbf{o}\| = 32 :$$

$$(214) \quad \mathbf{o}[0..11] = 0$$

$$(215) \quad \mathbf{o}[12..31] = \text{KEC}(\text{ECDSARECOVER}(h, v - 27, r, s))[12..31] \text{ where:}$$

$$(216) \quad \mathbf{d}[0..(\|I_d\| - 1)] = I_d$$

$$(217) \quad \mathbf{d}[\|I_d\|..] = (0, 0, \dots)$$

$$(218) \quad h = \mathbf{d}[0..31]$$

$$(219) \quad v = \mathbf{d}[32..63]$$

$$(220) \quad r = \mathbf{d}[64..95]$$

$$(221) \quad s = \mathbf{d}[96..127]$$

We define Ξ_{SHA256} and $\Xi_{\text{RIPEMD160}}$ as precompiled contracts implementing the SHA2-256 and RIPEMD-160 hash functions respectively. Their gas usage is dependent on the input data size, a factor rounded up to the nearest number of words.

$$(222) \quad \Xi_{\text{SHA256}} \equiv \Xi_{\text{PRE}} \text{ where:}$$

$$(223) \quad g_r = 60 + 12 \left\lceil \frac{\|I_d\|}{32} \right\rceil$$

$$(224) \quad \mathbf{o}[0..31] = \text{SHA256}(I_d)$$

$$(225) \quad \Xi_{\text{RIPEMD160}} \equiv \Xi_{\text{PRE}} \text{ where:}$$

$$(226) \quad g_r = 600 + 120 \left\lceil \frac{\|I_d\|}{32} \right\rceil$$

$$(227) \quad \mathbf{o}[0..11] = 0$$

$$(228) \quad \mathbf{o}[12..31] = \text{RIPEMD160}(I_d)$$

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

$$(229) \quad \text{SHA256}(\mathbf{i} \in \mathbb{B}) \equiv o \in \mathbb{B}_{32}$$

$$(230) \quad \text{RIPEMD160}(\mathbf{i} \in \mathbb{B}) \equiv o \in \mathbb{B}_{20}$$

The fourth contract, the identity function Ξ_{ID} simply defines the output as the input:

$$(231) \quad \Xi_{\text{ID}} \equiv \Xi_{\text{PRE}} \text{ where:}$$

$$(232) \quad g_r = 15 + 3 \left\lceil \frac{\|I_d\|}{32} \right\rceil$$

$$(233) \quad \mathbf{o} = I_d$$

The fifth contract performs arbitrary-precision exponentiation under modulo. Here, 0^0 is taken to be one, and $x \bmod 0$ is zero for all x . The first word in the input specifies the number of bytes that the first non-negative integer B occupies. The second word in the input specifies the number of bytes that the second non-negative integer E occupies. The third word in the input specifies the number of bytes that the third non-negative integer M occupies. These three words are followed by B , E and M . The rest of the input is discarded. Whenever the input is too short, the missing bytes are considered to be zero. The output is encoded big-endian into the same format as M 's.

$$(234) \quad \Xi_{\text{EXPMOD}} \equiv \Xi_{\text{PRE}} \text{ except:}$$

$$(235) \quad g_r = \max \left(200, \left\lceil \frac{f(\max(\ell_M, \ell_B)) \max(\ell'_E, 1)}{G_{\text{quaddivisor}}} \right\rceil \right)$$

$$(236) \quad G_{\text{quaddivisor}} \equiv 3$$

$$(237) \quad f(x) \equiv \left\lceil \frac{x}{8} \right\rceil^2$$

$$(238) \quad \ell'_E = \begin{cases} 0 & \text{if } \ell_E \leq 32 \wedge E = 0 \\ \lfloor \log_2(E) \rfloor & \text{if } \ell_E \leq 32 \wedge E \neq 0 \\ 8(\ell_E - 32) + \lfloor \log_2(i[(96 + \ell_B)..(127 + \ell_B)]) \rfloor & \text{if } 32 < \ell_E \wedge i[(96 + \ell_B)..(127 + \ell_B)] \neq 0 \\ 8(\ell_E - 32) & \text{otherwise} \end{cases}$$

$$(239) \quad \mathbf{o} = (B^E \bmod M) \in \mathbb{N}_{8\ell_M}$$

$$(240) \quad \ell_B \equiv i[0..31]$$

$$(241) \quad \ell_E \equiv i[32..63]$$

$$(242) \quad \ell_M \equiv i[64..95]$$

$$(243) \quad B \equiv i[96..(95 + \ell_B)]$$

$$(244) \quad E \equiv i[(96 + \ell_B)..(95 + \ell_B + \ell_E)]$$

$$(245) \quad M \equiv i[(96 + \ell_B + \ell_E)..(95 + \ell_B + \ell_E + \ell_M)]$$

$$(246) \quad i[x] \equiv \begin{cases} I_d[x] & \text{if } x < \|I_d\| \\ 0 & \text{otherwise} \end{cases}$$

E.1. zkSNARK Related Precompiled Contracts. We choose two numbers, both of which are prime.

$$(247) \quad p \equiv 21888242871839275222246405745257275088696311157297823662689037894645226208583$$

$$(248) \quad q \equiv 21888242871839275222246405745257275088548364400416034343698204186575808495617$$

Since p is a prime number, $\{0, 1, \dots, p-1\}$ forms a field with addition and multiplication modulo p . We call this field F_p .

We define a set C_1 with

$$(249) \quad C_1 \equiv \{(X, Y) \in F_p \times F_p \mid Y^2 = X^3 + 3\} \cup \{(0, 0)\}$$

We define a binary operation $+$ on C_1 for distinct elements $(X_1, Y_1), (X_2, Y_2)$ with

$$(250) \quad \begin{aligned} (X_1, Y_1) + (X_2, Y_2) &\equiv \begin{cases} (X, Y) & \text{if } X_1 \neq X_2 \\ (0, 0) & \text{otherwise} \end{cases} \\ \lambda &\equiv \frac{Y_2 - Y_1}{X_2 - X_1} \\ X &\equiv \lambda^2 - X_1 - X_2 \\ Y &\equiv \lambda(X_1 - X) - Y_1 \end{aligned}$$

In the case where $(X_1, Y_1) = (X_2, Y_2)$, we define $+$ on C_1 with

$$(251) \quad \begin{aligned} (X_1, Y_1) + (X_2, Y_2) &\equiv \begin{cases} (X, Y) & \text{if } Y_1 \neq 0 \\ (0, 0) & \text{otherwise} \end{cases} \\ \lambda &\equiv \frac{3X_1^2}{2Y_1} \\ X &\equiv \lambda^2 - 2X_1 \\ Y &\equiv \lambda(X_1 - X) - Y_1 \end{aligned}$$

$(C_1, +)$ is known to form a group. We define scalar multiplication \cdot with

$$(252) \quad n \cdot P \equiv (0, 0) + \underbrace{P + \dots + P}_n$$

for a natural number n and a point P in C_1 .

We define P_1 to be a point $(1, 2)$ on C_1 . Let G_1 be the subgroup of $(C_1, +)$ generated by P_1 . G_1 is known to be a cyclic group of order q . For a point P in G_1 , we define $\log_{P_1}(P)$ to be the smallest natural number n satisfying $n \cdot P_1 = P$. $\log_{P_1}(P)$ is at most $q - 1$.

Let F_{p^2} be a field $F_p[i]/(i^2 + 1)$. We define a set C_2 with

$$(253) \quad C_2 \equiv \{(X, Y) \in F_{p^2} \times F_{p^2} \mid Y^2 = X^3 + 3(i + 9)^{-1}\} \cup \{(0, 0)\}$$

We define a binary operation $+$ and scalar multiplication \cdot with the same equations (250), (251) and (252). $(C_2, +)$ is also known to be a group. We define P_2 in C_2 with

$$(254) \quad \begin{aligned} P_2 &\equiv (11559732032986387107991004021392285783925812861821192530917403151452391805634 \times i \\ &\quad + 10857046999023057135944570762232829481370756359578518086990519993285655852781, \\ &\quad 4082367875863433681332203403145435568316851327593401208105741076214120093531 \times i \\ &\quad + 8495653923123431417604973247489272438418190587263600148770280649306958101930) \end{aligned}$$

We define G_2 to be the subgroup of $(C_2, +)$ generated by P_2 . G_2 is known to be the only cyclic group of order q on C_2 . For a point P in G_2 , we define $\log_{P_2}(P)$ to be the smallest natural number n satisfying $n \cdot P_2 = P$. With this definition, $\log_{P_2}(P)$ is at most $q - 1$.

Let G_T be the multiplicative abelian group underlying $F_{q^{12}}$. It is known that a non-degenerate bilinear map $e : G_1 \times G_2 \rightarrow G_T$ exists. This bilinear map is a type three pairing. There are several such bilinear maps, it does not matter which is chosen to be e . Let $P_T = e(P_1, P_2)$, a be a set of k points in G_1 , and b be a set of k points in G_2 . It follows from the definition of a pairing that the following are equivalent

$$(255) \quad \log_{P_1}(a_1) \times \log_{P_2}(b_1) + \dots + \log_{P_1}(a_k) \times \log_{P_2}(b_k) \equiv 1 \pmod{q}$$

$$(256) \quad \prod_{i=0}^k e(a_i, b_i) = P_T$$

Thus the pairing operation provides a method to verify (255).

A 32 byte number $\mathbf{x} \in \mathbf{P}_{256}$ might and might not represent an element of F_p .

$$(257) \quad \delta_p(\mathbf{x}) \equiv \begin{cases} \mathbf{x} & \text{if } \mathbf{x} < p \\ \emptyset & \text{otherwise} \end{cases}$$

E.2. **BLAKE2 Precompiled Contract.** EIP-152 by Hess et al. [2016] defines $\Xi_{\text{BLAKE2.F}}$ as a precompiled contract implementing the compression function **F** used in the BLAKE2 cryptographic hashing algorithm. The **F** compression function is specified in RFC 7693 by Saarinen and Aumasson [2015].

$$\begin{aligned}
(292) \quad & \Xi_{\text{BLAKE2.F}} \equiv \Xi_{\text{PRE}} \text{ except:} \\
(293) \quad & \Xi_{\text{BLAKE2.F}}(\sigma, g, A, I) = (\emptyset, 0, A, ()) \text{ if } \|I_d\| \neq 213 \vee f \notin \{0, 1\} \\
(294) \quad & g_r = r \\
(295) \quad & \mathbf{o} \equiv \text{LE}_8(h'_0) \cdot \dots \cdot \text{LE}_8(h'_7) \\
(296) \quad & (h'_0, \dots, h'_7) \equiv \text{F}(h, m, t_{\text{low}}, t_{\text{high}}, f) \text{ with } r \text{ rounds and } w = 64 \\
(297) \quad & \text{BE}_4(r) \equiv I_d[0..4] \\
(298) \quad & \text{LE}_8(h_0) \equiv I_d[4..12] \\
(299) \quad & \dots \\
(300) \quad & \text{LE}_8(h_7) \equiv I_d[60..68] \\
(301) \quad & \text{LE}_8(m_0) \equiv I_d[68..76] \\
(302) \quad & \dots \\
(303) \quad & \text{LE}_8(m_{15}) \equiv I_d[188..196] \\
(304) \quad & \text{LE}_8(t_{\text{low}}) \equiv I_d[196..204] \\
(305) \quad & \text{LE}_8(t_{\text{high}}) \equiv I_d[204..212] \\
(306) \quad & f \equiv I_d[212]
\end{aligned}$$

where $r \in \mathbb{B}_{32}$, $\forall i \in 0..7 : h_i \in \mathbb{B}_{64}$, $\forall i \in 0..15 : m_i \in \mathbb{B}_{64}$, $t_{\text{low}} \in \mathbb{B}_{64}$, $t_{\text{high}} \in \mathbb{B}_{64}$, $f \in \mathbb{B}_8$, BE_k is the k -byte big-endian representation—compare with(195):

$$(307) \quad \text{BE}_k(x) \equiv (b_0, b_1, \dots, b_{k-1}) : x = \sum_{n=0}^{k-1} b_n \cdot 256^{k-1-n}$$

and LE_k is the k -byte little-endian representation:

$$(308) \quad \text{LE}_k(x) \equiv (b_0, b_1, \dots, b_{k-1}) : x = \sum_{n=0}^{k-1} b_n \cdot 256^n$$

APPENDIX F. SIGNING TRANSACTIONS

Transactions are signed using recoverable ECDSA signatures. This method utilises the SECP-256k1 curve as described by Courtois et al. [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 $[1, \text{secp256k1n} - 1]$.

We assume the existence of functions **ECDSAPUBKEY**, **ECDSASIGN** and **ECDSARECOVER**. These are formally defined in the literature, e.g. by Johnson et al. [2001].

$$\begin{aligned}
(309) \quad & \text{ECDSAPUBKEY}(p_r \in \mathbb{B}_{32}) \equiv p_u \in \mathbb{B}_{64} \\
(310) \quad & \text{ECDSASIGN}(e \in \mathbb{B}_{32}, p_r \in \mathbb{B}_{32}) \equiv (v \in \mathbb{B}_1, r \in \mathbb{B}_{32}, s \in \mathbb{B}_{32}) \\
(311) \quad & \text{ECDSARECOVER}(e \in \mathbb{B}_{32}, v \in \mathbb{B}_1, r \in \mathbb{B}_{32}, s \in \mathbb{B}_{32}) \equiv p_u \in \mathbb{B}_{64}
\end{aligned}$$

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 $< 2^{256}$), p_r is the private key, a byte array of size 32 (or a single positive integer in the aforementioned range) and e is the hash of the transaction, $h(T)$. It is assumed that v is the ‘recovery identifier’. The recovery identifier is a 1 byte value specifying the parity and finiteness of the coordinates of the curve point for which r is the x-value; this value is in the range of $[0, 3]$, however we declare the upper two possibilities, representing infinite values, invalid. The value 0 represents an even y value and 1 represents an odd y value.

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

$$\begin{aligned}
(312) \quad & 0 < r < \text{secp256k1n} \\
(313) \quad & 0 < s < \text{secp256k1n} \div 2 + 1 \\
(314) \quad & v \in \{0, 1\}
\end{aligned}$$

where:

$$(315) \quad \text{secp256k1n} = 115792089237316195423570985008687907852837564279074904382605163141518161494337$$

Note that this restriction on s is more stringent than restriction 212 in the Ξ_{ECCREC} precompile; see EIP-2 by Buterin [2015] for more detail.

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 rightmost 160-bits of the Keccak-256 hash of the corresponding ECDSA public key:

$$(316) \quad A(p_r) = \mathcal{B}_{96..255}(\text{KEC}(\text{ECDSAPUBKEY}(p_r)))$$

The message hash, $h(T)$, to be signed is the Keccak-256 hash of the transaction. Four different flavours of signing schemes are available:

$$(317) \quad L_X(T) \equiv \begin{cases} (T_n, T_p, T_g, T_t, T_v, \mathbf{p}) & \text{if } T_x = 0 \wedge T_w \in \{27, 28\} \\ (T_n, T_p, T_g, T_t, T_v, \mathbf{p}, \beta, (), ()) & \text{if } T_x = 0 \wedge T_w \in \{2\beta + 35, 2\beta + 36\} \\ (T_c, T_n, T_p, T_g, T_t, T_v, \mathbf{p}, T_A) & \text{if } T_x = 1 \\ (T_c, T_n, T_i, T_m, T_g, T_t, T_v, \mathbf{p}, T_A) & \text{if } T_x = 2 \end{cases}$$

where

$$(318) \quad \mathbf{p} \equiv \begin{cases} T_i & \text{if } T_t = \emptyset \\ T_d & \text{otherwise} \end{cases}$$

$$(318) \quad h(T) \equiv \begin{cases} \text{KEC}(\text{RLP}(L_X(T))) & \text{if } T_x = 0 \\ \text{KEC}(T_x \cdot \text{RLP}(L_X(T))) & \text{otherwise} \end{cases}$$

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

$$(319) \quad G(T, p_r) \equiv T \quad \text{except:}$$

$$(320) \quad (T_y, T_r, T_s) = \text{ECDSASIGN}(h(T), p_r)$$

Reiterating from previously:

$$(321) \quad T_r = r$$

$$(322) \quad T_s = s$$

and T_w of legacy transactions is either $27 + T_y$ or $2\beta + 35 + T_y$.

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

$$(323) \quad S(T) \equiv \mathcal{B}_{96..255}(\text{KEC}(\text{ECDSARECOVER}(h(T), v, T_r, T_s)))$$

$$(324) \quad v \equiv \begin{cases} T_w - 27 & \text{if } T_x = 0 \wedge T_w \in \{27, 28\} \\ (T_w - 35) \bmod 2 & \text{if } T_x = 0 \wedge T_w \in \{2\beta + 35, 2\beta + 36\} \\ T_y & \text{if } T_x = 1 \vee T_x = 2 \end{cases}$$

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

$$(325) \quad \forall T : \forall p_r : S(G(T, p_r)) \equiv A(p_r)$$

APPENDIX G. FEE SCHEDULE

The fee schedule G is a tuple of 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 paid for operations of the set W_{zero} . |
| G_{jumpdest} | 1 | Amount of gas to pay for a JUMPDEST operation. |
| G_{base} | 2 | Amount of gas to pay for operations of the set W_{base} . |
| G_{verylow} | 3 | Amount of gas to pay for operations of the set W_{verylow} . |
| G_{low} | 5 | Amount of gas to pay for operations of the set W_{low} . |
| G_{mid} | 8 | Amount of gas to pay for operations of the set W_{mid} . |
| G_{high} | 10 | Amount of gas to pay for operations of the set W_{high} . |
| $G_{\text{warmaccess}}$ | 100 | Cost of a warm account or storage access. |
| $G_{\text{accesslistaddress}}$ | 2400 | Cost of warming up an account with the access list. |
| $G_{\text{accessliststorage}}$ | 1900 | Cost of warming up a storage with the access list. |
| $G_{\text{coldaccountaccess}}$ | 2600 | Cost of a cold account access. |
| G_{coldslod} | 2100 | Cost of a cold storage access. |
| G_{sset} | 20000 | Paid for an SSTORE operation when the storage value is set to non-zero from zero. |
| G_{sreset} | 2900 | Paid for an SSTORE operation when the storage value's zeroness remains unchanged or is set to zero. |
| R_{sclear} | 4800 | Refund given (added into refund counter) when the storage value is set to zero from non-zero. The refund amount is defined as $G_{\text{sreset}} + G_{\text{accessliststorage}}$. |
| $G_{\text{selfdestruct}}$ | 5000 | Amount of gas to pay for a SELFDESTRUCT operation. |
| G_{create} | 32000 | Paid for a CREATE operation. |
| $G_{\text{codedeposit}}$ | 200 | Paid per byte for a CREATE operation to succeed in placing code into state. |
| $G_{\text{initcodeword}}$ | 2 | Paid per word of the initcode at the beginning of a deploy transaction, a CREATE, or a CREATE2 operation. |
| $G_{\text{callvalue}}$ | 9000 | Paid for a non-zero value transfer as part of the CALL operation. |
| $G_{\text{callstipend}}$ | 2300 | A stipend for the called contract subtracted from $G_{\text{callvalue}}$ for a non-zero value transfer. |
| $G_{\text{newaccount}}$ | 25000 | Paid for a CALL or SELFDESTRUCT operation which creates an account. |
| G_{exp} | 10 | Partial payment for an EXP operation. |
| G_{expbyte} | 50 | Partial payment when multiplied by the number of bytes in the exponent for the EXP operation. |
| G_{memory} | 3 | Paid for every additional word when expanding memory. |
| G_{txcreate} | 32000 | Paid by all contract-creating transactions after the <i>Homestead</i> transition. |
| $G_{\text{txdatazero}}$ | 4 | Paid for every zero byte of data or code for a transaction. |
| $G_{\text{txdatanonzero}}$ | 16 | Paid for every non-zero byte of data or code for a transaction. |
| $G_{\text{transaction}}$ | 21000 | Paid for every transaction. |
| G_{log} | 375 | Partial payment for a LOG operation. |
| G_{logdata} | 8 | Paid for each byte in a LOG operation's data. |
| G_{logtopic} | 375 | Paid for each topic of a LOG operation. |
| $G_{\text{keccak256}}$ | 30 | Paid for each KECCAK256 operation. |
| $G_{\text{keccak256word}}$ | 6 | Paid for each word (rounded up) for input data to a KECCAK256 operation. |
| G_{copy} | 3 | Partial payment for *COPY operations, multiplied by words copied, rounded up. |
| $G_{\text{blockhash}}$ | 20 | Payment for each 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 leftmost 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:

$$(326) \quad C(\sigma, \mu, A, I) \equiv C_{\text{mem}}(\mu'_i) - C_{\text{mem}}(\mu_i) + \begin{cases} C_{\text{SSTORE}}(\sigma, \mu, A, I) & \text{if } w = \text{SSTORE} \\ G_{\text{exp}} & \text{if } w = \text{EXP} \wedge \mu_s[1] = 0 \\ G_{\text{exp}} + G_{\text{expbyte}} \times (1 + \lceil \log_{256}(\mu_s[1]) \rceil) & \text{if } w = \text{EXP} \wedge \mu_s[1] > 0 \\ G_{\text{verylow}} + G_{\text{copy}} \times \lceil \mu_s[2] \div 32 \rceil & \text{if } w \in W_{\text{copy}} \\ C_{\text{aaccess}}(\mu_s[0] \bmod 2^{160}, A) + G_{\text{copy}} \times \lceil \mu_s[3] \div 32 \rceil & \text{if } w = \text{EXTCODECOPY} \\ C_{\text{aaccess}}(\mu_s[0] \bmod 2^{160}, A) & \text{if } w \in W_{\text{extaccount}} \\ G_{\text{log}} + G_{\text{logdata}} \times \mu_s[1] & \text{if } w = \text{LOG0} \\ G_{\text{log}} + G_{\text{logdata}} \times \mu_s[1] + G_{\text{logtopic}} & \text{if } w = \text{LOG1} \\ G_{\text{log}} + G_{\text{logdata}} \times \mu_s[1] + 2G_{\text{logtopic}} & \text{if } w = \text{LOG2} \\ G_{\text{log}} + G_{\text{logdata}} \times \mu_s[1] + 3G_{\text{logtopic}} & \text{if } w = \text{LOG3} \\ G_{\text{log}} + G_{\text{logdata}} \times \mu_s[1] + 4G_{\text{logtopic}} & \text{if } w = \text{LOG4} \\ C_{\text{CALL}}(\sigma, \mu, A) & \text{if } w \in W_{\text{call}} \\ C_{\text{SELFDESTRUCT}}(\sigma, \mu) & \text{if } w = \text{SELFDESTRUCT} \\ G_{\text{create}} + R(\mu_s[2]) & \text{if } w = \text{CREATE} \\ G_{\text{create}} + G_{\text{keccak256word}} \times \lceil \mu_s[2] \div 32 \rceil + R(\mu_s[2]) & \text{if } w = \text{CREATE2} \\ G_{\text{keccak256}} + G_{\text{keccak256word}} \times \lceil \mu_s[1] \div 32 \rceil & \text{if } w = \text{KECCAK256} \\ G_{\text{jumpdest}} & \text{if } w = \text{JUMPDEST} \\ C_{\text{SLOAD}}(\mu, A, I) & \text{if } w = \text{SLOAD} \\ G_{\text{zero}} & \text{if } w \in W_{\text{zero}} \\ G_{\text{base}} & \text{if } w \in W_{\text{base}} \\ G_{\text{verylow}} & \text{if } w \in W_{\text{verylow}} \\ G_{\text{low}} & \text{if } w \in W_{\text{low}} \\ G_{\text{mid}} & \text{if } w \in W_{\text{mid}} \\ G_{\text{high}} & \text{if } w \in W_{\text{high}} \\ G_{\text{blockhash}} & \text{if } w = \text{BLOCKHASH} \end{cases}$$

$$(327) \quad w \equiv \begin{cases} I_b[\mu_{\text{pc}}] & \text{if } \mu_{\text{pc}} < \|I_b\| \\ \text{STOP} & \text{otherwise} \end{cases}$$

where:

$$(328) \quad C_{\text{mem}}(a) \equiv G_{\text{memory}} \cdot a + \left\lfloor \frac{a^2}{512} \right\rfloor$$

$$(329) \quad C_{\text{aaccess}}(x, A) \equiv \begin{cases} G_{\text{warmaccess}} & \text{if } x \in A_{\mathbf{a}} \\ G_{\text{coldaccountaccess}} & \text{otherwise} \end{cases}$$

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

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

$$W_{\text{base}} = \{\text{ADDRESS, ORIGIN, CALLER, CALLVALUE, CALLDATASIZE, CODESIZE, GASPRICE, COINBASE, TIMESTAMP, NUMBER, PREVRANDAO, GASLIMIT, CHAINID, RETURNDATASIZE, POP, PC, MSIZE, GAS, BASEFEE, PUSH0}\}$$

$$W_{\text{verylow}} = \{\text{ADD, SUB, NOT, LT, GT, SLT, SGT, EQ, ISZERO, AND, OR, XOR, BYTE, SHL, SHR, SAR, CALLDATALOAD, MLOAD, MSTORE, MSTORE8, PUSH1, ..., PUSH32, DUP*, SWAP*}\}$$

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

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

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

$$W_{\text{copy}} = \{\text{CALLDATACOPY, CODECOPY, RETURNDATACOPY}\}$$

$$W_{\text{call}} = \{\text{CALL, CALLCODE, DELEGATECALL, STATICCALL}\}$$

$$W_{\text{extaccount}} = \{\text{BALANCE, EXTCODESIZE, EXTCODEHASH}\}$$

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, μ_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. μ'_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 704B of memory used, after which it costs substantially more.

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

$$(330) \quad M(s, f, l) \equiv \begin{cases} s & \text{if } l = 0 \\ \max(s, \lceil (f + l) \div 32 \rceil) & \text{otherwise} \end{cases}$$

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

$$(331) \quad L(n) \equiv n - \lfloor n/64 \rfloor$$

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 (σ', μ') such that:

$$(332) \quad O(\sigma, \mu, A, I) \equiv (\sigma', \mu', A', I) \quad \text{with exceptions, as noted}$$

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 stack, as defined in section 9.

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.

| Value | Mnemonic | δ | α | Description |
|-------|------------|----------|----------|---|
| 0x00 | STOP | 0 | 0 | Halts execution. |
| 0x01 | ADD | 2 | 1 | Addition operation. $\mu'_s[0] \equiv \mu_s[0] + \mu_s[1]$ |
| 0x02 | MUL | 2 | 1 | Multiplication operation. $\mu'_s[0] \equiv \mu_s[0] \times \mu_s[1]$ |
| 0x03 | SUB | 2 | 1 | Subtraction operation. $\mu'_s[0] \equiv \mu_s[0] - \mu_s[1]$ |
| 0x04 | DIV | 2 | 1 | Integer division operation. $\mu'_s[0] \equiv \begin{cases} 0 & \text{if } \mu_s[1] = 0 \\ \lfloor \mu_s[0] \div \mu_s[1] \rfloor & \text{otherwise} \end{cases}$ |
| 0x05 | SDIV | 2 | 1 | Signed integer division operation (truncated). $\mu'_s[0] \equiv \begin{cases} 0 & \text{if } \mu_s[1] = 0 \\ -2^{255} & \text{if } \mu_s[0] = -2^{255} \wedge \mu_s[1] = -1 \\ \text{sgn}(\mu_s[0] \div \mu_s[1]) \lfloor \mu_s[0] \div \mu_s[1] \rfloor & \text{otherwise} \end{cases}$ 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 | Modulo remainder operation. $\mu'_s[0] \equiv \begin{cases} 0 & \text{if } \mu_s[1] = 0 \\ \mu_s[0] \bmod \mu_s[1] & \text{otherwise} \end{cases}$ |
| 0x07 | SMOD | 2 | 1 | Signed modulo remainder operation. $\mu'_s[0] \equiv \begin{cases} 0 & \text{if } \mu_s[1] = 0 \\ \text{sgn}(\mu_s[0]) (\lfloor \mu_s[0] \rfloor \bmod \lfloor \mu_s[1] \rfloor) & \text{otherwise} \end{cases}$ Where all values are treated as two's complement signed 256-bit integers. |
| 0x08 | ADDMOD | 3 | 1 | Modulo addition operation. $\mu'_s[0] \equiv \begin{cases} 0 & \text{if } \mu_s[2] = 0 \\ (\mu_s[0] + \mu_s[1]) \bmod \mu_s[2] & \text{otherwise} \end{cases}$ All intermediate calculations of this operation are not subject to the 2^{256} modulo. |
| 0x09 | MULMOD | 3 | 1 | Modulo multiplication operation. $\mu'_s[0] \equiv \begin{cases} 0 & \text{if } \mu_s[2] = 0 \\ (\mu_s[0] \times \mu_s[1]) \bmod \mu_s[2] & \text{otherwise} \end{cases}$ All intermediate calculations of this operation are not subject to the 2^{256} modulo. |
| 0x0a | EXP | 2 | 1 | Exponential operation. $\mu'_s[0] \equiv \mu_s[0]^{\mu_s[1]}$ |
| 0x0b | SIGNEXTEND | 2 | 1 | Extend length of two's complement signed integer. $\forall i \in [0..255] : \mu'_s[0]_i \equiv \begin{cases} \mu_s[1]_t & \text{if } i \leq t \text{ where } t = 256 - 8(\mu_s[0] + 1) \\ \mu_s[1]_i & \text{otherwise} \end{cases}$ |

$\mu_s[x]_i$ gives the i th bit (counting from zero) of $\mu_s[x]$

| 10s: Comparison & Bitwise Logic Operations | | | | |
|---|----------|----------|----------|--|
| Value | Mnemonic | δ | α | Description |
| 0x10 | LT | 2 | 1 | Less-than comparison. $\mu'_s[0] \equiv \begin{cases} 1 & \text{if } \mu_s[0] < \mu_s[1] \\ 0 & \text{otherwise} \end{cases}$ |
| 0x11 | GT | 2 | 1 | Greater-than comparison. $\mu'_s[0] \equiv \begin{cases} 1 & \text{if } \mu_s[0] > \mu_s[1] \\ 0 & \text{otherwise} \end{cases}$ |
| 0x12 | SLT | 2 | 1 | Signed less-than comparison. $\mu'_s[0] \equiv \begin{cases} 1 & \text{if } \mu_s[0] < \mu_s[1] \\ 0 & \text{otherwise} \end{cases}$ Where all values are treated as two's complement signed 256-bit integers. |
| 0x13 | SGT | 2 | 1 | Signed greater-than comparison. $\mu'_s[0] \equiv \begin{cases} 1 & \text{if } \mu_s[0] > \mu_s[1] \\ 0 & \text{otherwise} \end{cases}$ Where all values are treated as two's complement signed 256-bit integers. |
| 0x14 | EQ | 2 | 1 | Equality comparison. $\mu'_s[0] \equiv \begin{cases} 1 & \text{if } \mu_s[0] = \mu_s[1] \\ 0 & \text{otherwise} \end{cases}$ |
| 0x15 | ISZERO | 1 | 1 | Simple not operator. $\mu'_s[0] \equiv \begin{cases} 1 & \text{if } \mu_s[0] = 0 \\ 0 & \text{otherwise} \end{cases}$ |
| 0x16 | AND | 2 | 1 | Bitwise AND operation. $\forall i \in [0..255] : \mu'_s[0]_i \equiv \mu_s[0]_i \wedge \mu_s[1]_i$ |
| 0x17 | OR | 2 | 1 | Bitwise OR operation. $\forall i \in [0..255] : \mu'_s[0]_i \equiv \mu_s[0]_i \vee \mu_s[1]_i$ |
| 0x18 | XOR | 2 | 1 | Bitwise XOR operation. $\forall i \in [0..255] : \mu'_s[0]_i \equiv \mu_s[0]_i \oplus \mu_s[1]_i$ |
| 0x19 | NOT | 1 | 1 | Bitwise NOT operation. $\forall i \in [0..255] : \mu'_s[0]_i \equiv \begin{cases} 1 & \text{if } \mu_s[0]_i = 0 \\ 0 & \text{otherwise} \end{cases}$ |
| 0x1a | BYTE | 2 | 1 | Retrieve single byte from word. $\forall i \in [0..255] : \mu'_s[0]_i \equiv \begin{cases} \mu_s[1]_{(i-248+8\mu_s[0])} & \text{if } i \geq 248 \wedge \mu_s[0] < 32 \\ 0 & \text{otherwise} \end{cases}$ For the Nth byte, we count from the left (i.e. N=0 would be the most significant in big endian). |
| 0x1b | SHL | 2 | 1 | Left shift operation. $\mu'_s[0] \equiv (\mu_s[1] \times 2^{\mu_s[0]}) \bmod 2^{256}$ |
| 0x1c | SHR | 2 | 1 | Logical right shift operation. $\mu'_s[0] \equiv \lfloor \mu_s[1] \div 2^{\mu_s[0]} \rfloor$ |
| 0x1d | SAR | 2 | 1 | Arithmetic (signed) right shift operation. $\mu'_s[0] \equiv \lfloor \mu_s[1] \div 2^{\mu_s[0]} \rfloor$ Where $\mu'_s[0]$ and $\mu_s[1]$ are treated as two's complement signed 256-bit integers, while $\mu_s[0]$ is treated as unsigned. |

20s: KECCAK256

| Value | Mnemonic | δ | α | Description |
|-------|-----------|----------|----------|--|
| 0x20 | KECCAK256 | 2 | 1 | Compute Keccak-256 hash. $\mu'_s[0] \equiv \text{KEC}(\mu_m[\mu_s[0] \dots (\mu_s[0] + \mu_s[1] - 1)])$ $\mu'_i \equiv M(\mu_i, \mu_s[0], \mu_s[1])$ |

| 30s: Environmental Information | | | | |
|---------------------------------------|--------------|----------|----------|--|
| Value | Mnemonic | δ | α | Description |
| 0x30 | ADDRESS | 0 | 1 | Get address of currently executing account. $\mu'_s[0] \equiv I_a$ |
| 0x31 | BALANCE | 1 | 1 | Get balance of the given account. $\mu'_s[0] \equiv \begin{cases} \sigma[\mu_s[0] \bmod 2^{160}]_b & \text{if } \sigma[\mu_s[0] \bmod 2^{160}] \neq \emptyset \\ 0 & \text{otherwise} \end{cases}$ $A'_a \equiv A_a \cup \{\mu_s[0] \bmod 2^{160}\}$ |
| 0x32 | ORIGIN | 0 | 1 | Get execution origination address. $\mu'_s[0] \equiv I_o$ This is the sender of original transaction; it is never an account with non-empty associated code. |
| 0x33 | CALLER | 0 | 1 | Get caller address. $\mu'_s[0] \equiv I_s$ This is the address of the account that is directly responsible for this execution. |
| 0x34 | CALLVALUE | 0 | 1 | Get deposited value by the instruction/transaction responsible for this execution. $\mu'_s[0] \equiv I_v$ |
| 0x35 | CALLDATALOAD | 1 | 1 | Get input data of current environment. $\mu'_s[0] \equiv I_d[\mu_s[0] \dots (\mu_s[0] + 31)]$ with $I_d[x] = 0$ if $x \geq \ I_d\ $ This pertains to the input data passed with the message call instruction or transaction. |
| 0x36 | CALLDATASIZE | 0 | 1 | Get size of input data in current environment. $\mu'_s[0] \equiv \ I_d\ $ This pertains to the input data passed with the message call instruction or transaction. |
| 0x37 | CALLDATACOPY | 3 | 0 | Copy input data in current environment to memory. $\forall i \in \{0 \dots \mu_s[2] - 1\} : \mu'_m[\mu_s[0] + i] \equiv \begin{cases} I_d[\mu_s[1] + i] & \text{if } \mu_s[1] + i < \ I_d\ \\ 0 & \text{otherwise} \end{cases}$ The additions in $\mu_s[1] + i$ are not subject to the 2^{256} modulo. $\mu'_i \equiv M(\mu_i, \mu_s[0], \mu_s[2])$ This pertains to the input data passed with the message call instruction or transaction. |
| 0x38 | CODESIZE | 0 | 1 | Get size of code running in current environment. $\mu'_s[0] \equiv \ I_b\ $ |
| 0x39 | CODECOPY | 3 | 0 | Copy code running in current environment to memory. $\forall i \in \{0 \dots \mu_s[2] - 1\} : \mu'_m[\mu_s[0] + i] \equiv \begin{cases} I_b[\mu_s[1] + i] & \text{if } \mu_s[1] + i < \ I_b\ \\ \text{STOP} & \text{otherwise} \end{cases}$ $\mu'_i \equiv M(\mu_i, \mu_s[0], \mu_s[2])$ The additions in $\mu_s[1] + i$ are not subject to the 2^{256} modulo. |
| 0x3a | GASPRICE | 0 | 1 | Get price of gas in current environment. This is the <i>effective gas price</i> defined in section 6. Note that as of the <i>London</i> hard fork, this value no longer represents what is received by the validator, but rather just what is paid by the sender. $\mu'_s[0] \equiv I_p$ |
| 0x3b | EXTCODESIZE | 1 | 1 | Get size of an account's code. $\mu'_s[0] \equiv \begin{cases} \ b\ & \text{if } \sigma[\mu_s[0] \bmod 2^{160}] \neq \emptyset \\ 0 & \text{otherwise} \end{cases}$ where $\text{KEC}(b) \equiv \sigma[\mu_s[0] \bmod 2^{160}]_c$ $A'_a \equiv A_a \cup \{\mu_s[0] \bmod 2^{160}\}$ |
| 0x3c | EXTCODECOPY | 4 | 0 | Copy an account's code to memory. $\forall i \in \{0 \dots \mu_s[3] - 1\} : \mu'_m[\mu_s[1] + i] \equiv \begin{cases} b[\mu_s[2] + i] & \text{if } \mu_s[2] + i < \ b\ \\ \text{STOP} & \text{otherwise} \end{cases}$ where $\text{KEC}(b) \equiv \sigma[\mu_s[0] \bmod 2^{160}]_c$ We assume $b \equiv ()$ if $\sigma[\mu_s[0] \bmod 2^{160}] = \emptyset$. $\mu'_i \equiv M(\mu_i, \mu_s[1], \mu_s[3])$ The additions in $\mu_s[2] + i$ are not subject to the 2^{256} modulo. $A'_a \equiv A_a \cup \{\mu_s[0] \bmod 2^{160}\}$ |

| | | | | |
|------|----------------|---|---|--|
| 0x3d | RETURNDATASIZE | 0 | 1 | Get size of output data from the previous call from the current environment. $\mu'_s[0] \equiv \ \mu_o\ $ |
| 0x3e | RETURNDATACOPY | 3 | 0 | Copy output data from the previous call to memory. $\forall i \in \{0 \dots \mu_s[2] - 1\} : \mu'_m[\mu_s[0] + i] \equiv \begin{cases} \mu_o[\mu_s[1] + i] & \text{if } \mu_s[1] + i < \ \mu_o\ \\ 0 & \text{otherwise} \end{cases}$ The additions in $\mu_s[1] + i$ are not subject to the 2^{256} modulo. $\mu'_i \equiv M(\mu_i, \mu_s[0], \mu_s[2])$ |
| 0x3f | EXTCODEHASH | 1 | 1 | Get hash of an account's code. $\mu'_s[0] \equiv \begin{cases} 0 & \text{if } \text{DEAD}(\sigma, \mu_s[0] \bmod 2^{160}) \\ \sigma[\mu_s[0] \bmod 2^{160}]_c & \text{otherwise} \end{cases}$ $A'_a \equiv A_a \cup \{\mu_s[0] \bmod 2^{160}\}$ |

40s: Block Information

| Value | Mnemonic | δ | α | Description |
|-------|-------------|----------|----------|--|
| 0x40 | BLOCKHASH | 1 | 1 | Get the hash of one of the 256 most recent complete blocks. $\mu'_s[0] \equiv P(I_{H_p}, \mu_s[0], 0)$ 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 or equal to the current block number or more than 256 blocks behind the current block. $P(h, n, a) \equiv \begin{cases} 0 & \text{if } n > H_i \vee a = 256 \vee h = 0 \\ h & \text{if } n = H_i \\ P(H_p, n, a + 1) & \text{otherwise} \end{cases}$ and we assert the header H can be determined from its hash h unless h is zero (as is the case for the parent hash of the genesis block). |
| 0x41 | COINBASE | 0 | 1 | Get the current block's beneficiary address. $\mu'_s[0] \equiv I_{H_c}$ |
| 0x42 | TIMESTAMP | 0 | 1 | Get the current block's timestamp. $\mu'_s[0] \equiv I_{H_s}$ |
| 0x43 | NUMBER | 0 | 1 | Get the current block's number. $\mu'_s[0] \equiv I_{H_i}$ |
| 0x44 | PREVRANDAO | 0 | 1 | Get the latest RANDAO mix of the post beacon state of the previous block. $\mu'_s[0] \equiv I_{H_a}$ |
| 0x45 | GASLIMIT | 0 | 1 | Get the current block's gas limit. $\mu'_s[0] \equiv I_{H_l}$ |
| 0x46 | CHAINID | 0 | 1 | Get the chain ID. $\mu'_s[0] \equiv \beta$ |
| 0x47 | SELFBALANCE | 0 | 1 | Get balance of currently executing account. $\mu'_s[0] \equiv \sigma[I_a]_b$ |
| 0x48 | BASEFEE | 0 | 1 | Get the current block's base fee. $\mu'_s[0] \equiv I_{H_f}$ |

50s: Stack, Memory, Storage and Flow Operations

| Value | Mnemonic | δ | α | Description |
|-------|----------|----------|----------|---|
| 0x50 | POP | 1 | 0 | Remove item from stack. |
| 0x51 | MLOAD | 1 | 1 | Load word from memory. $\mu'_s[0] \equiv \mu_m[\mu_s[0] \dots (\mu_s[0] + 31)]$ $\mu'_i \equiv \max(\mu_i, \lceil (\mu_s[0] + 32) \div 32 \rceil)$ The addition in the calculation of μ'_i is not subject to the 2^{256} modulo. |
| 0x52 | MSTORE | 2 | 0 | Save word to memory. $\mu'_m[\mu_s[0] \dots (\mu_s[0] + 31)] \equiv \mu_s[1]$ $\mu'_i \equiv \max(\mu_i, \lceil (\mu_s[0] + 32) \div 32 \rceil)$ The addition in the calculation of μ'_i is not subject to the 2^{256} modulo. |
| 0x53 | MSTORE8 | 2 | 0 | Save byte to memory. $\mu'_m[\mu_s[0]] \equiv (\mu_s[1] \bmod 256)$ $\mu'_i \equiv \max(\mu_i, \lceil (\mu_s[0] + 1) \div 32 \rceil)$ The addition in the calculation of μ'_i is not subject to the 2^{256} modulo. |
| 0x54 | SLOAD | 1 | 1 | Load word from storage. $\mu'_s[0] \equiv \sigma[I_a]_s[\mu_s[0]]$ $A'_K \equiv A_K \cup \{(I_a, \mu_s[0])\}$ $C_{SLOAD}(\mu, A, I) \equiv \begin{cases} G_{warmaccess} & \text{if } (I_a, \mu_s[0]) \in A_K \\ G_{coldload} & \text{otherwise} \end{cases}$ |
| 0x55 | SSTORE | 2 | 0 | Save word to storage. $\sigma'[I_a]_s[\mu_s[0]] \equiv \mu_s[1]$ $A'_K \equiv A_K \cup \{(I_a, \mu_s[0])\}$ $C_{SSTORE}(\sigma, \mu)$ and A'_r are specified by EIP-2200 as follows. We remind the reader that the checkpoint (“original”) state σ_0 is the state if the current transaction were to revert. Let $v_0 = \sigma_0[I_a]_s[\mu_s[0]]$ be the original value of the storage slot. Let $v = \sigma[I_a]_s[\mu_s[0]]$ be the current value. Let $v' = \mu_s[1]$ be the new value. Then: $C_{SSTORE}(\sigma, \mu, A, I) \equiv \begin{cases} 0 & \text{if } (I_a, \mu_s[0]) \in A_K \\ G_{coldload} & \text{otherwise} \end{cases}$ $+ \begin{cases} G_{warmaccess} & \text{if } v = v' \vee v_0 \neq v \\ G_{sset} & \text{if } v \neq v' \wedge v_0 = v \wedge v_0 = 0 \\ G_{sreset} & \text{if } v \neq v' \wedge v_0 = v \wedge v_0 \neq 0 \end{cases}$ $A'_r \equiv A_r + \begin{cases} R_{sclear} & \text{if } v \neq v' \wedge v_0 = v \wedge v' = 0 \\ r_{dirtyclear} + r_{dirtyreset} & \text{if } v \neq v' \wedge v_0 \neq v \\ 0 & \text{otherwise} \end{cases}$ where $r_{dirtyclear} \equiv \begin{cases} -R_{sclear} & \text{if } v_0 \neq 0 \wedge v = 0 \\ R_{sclear} & \text{if } v_0 \neq 0 \wedge v' = 0 \\ 0 & \text{otherwise} \end{cases}$ $r_{dirtyreset} \equiv \begin{cases} G_{sset} - G_{warmaccess} & \text{if } v_0 = v' \wedge v_0 = 0 \\ G_{sreset} - G_{warmaccess} & \text{if } v_0 = v' \wedge v_0 \neq 0 \\ 0 & \text{otherwise} \end{cases}$ |
| 0x56 | JUMP | 1 | 0 | Alter the program counter. $J_{JUMP}(\mu) \equiv \mu_s[0]$ This has the effect of writing said value to μ_{pc} . See section 9. |
| 0x57 | JUMPI | 2 | 0 | Conditionally alter the program counter. $J_{JUMPI}(\mu) \equiv \begin{cases} \mu_s[0] & \text{if } \mu_s[1] \neq 0 \\ \mu_{pc} + 1 & \text{otherwise} \end{cases}$ This has the effect of writing said value to μ_{pc} . See section 9. |
| 0x58 | PC | 0 | 1 | Get the value of the program counter <i>prior</i> to the increment corresponding to this instruction. $\mu'_s[0] \equiv \mu_{pc}$ |

| | | | | |
|------|----------|---|---|--|
| 0x59 | MSIZE | 0 | 1 | Get the size of active memory in bytes. $\mu'_s[0] \equiv 32\mu_i$ |
| 0x5a | GAS | 0 | 1 | Get the amount of available gas, including the corresponding reduction for the cost of this instruction. $\mu'_s[0] \equiv \mu_g$ |
| 0x5b | JUMPDEST | 0 | 0 | Mark a valid destination for jumps. This operation has no effect on machine state during execution. |

5f, 60s & 70s: Push Operations

| Value | Mnemonic | δ | α | Description |
|-------|----------|----------|----------|--|
| 0x5f | PUSH0 | 0 | 1 | Place 0 on the stack. $\mu'_s[0] \equiv 0$ |
| 0x60 | PUSH1 | 0 | 1 | Place 1 byte item on stack. $\mu'_s[0] \equiv c(\mu_{pc} + 1)$ where $c(x) \equiv \begin{cases} I_b[x] & \text{if } x < \ I_b\ \\ 0 & \text{otherwise} \end{cases}$ 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 (takes the lowest significant place in big endian). |
| 0x61 | PUSH2 | 0 | 1 | Place 2-byte item on stack. $\mu'_s[0] \equiv c((\mu_{pc} + 1) \dots (\mu_{pc} + 2))$ with $c(\mathbf{x}) \equiv (c(x_0), \dots, c(x_{\ \mathbf{x}\ -1}))$ with c as defined as above. The bytes are right-aligned (takes the lowest significant place in big endian). |
| ⋮ | ⋮ | ⋮ | ⋮ | ⋮ |
| 0x7f | PUSH32 | 0 | 1 | Place 32-byte (full word) item on stack. $\mu'_s[0] \equiv c((\mu_{pc} + 1) \dots (\mu_{pc} + 32))$ where c is defined as above. The bytes are right-aligned (takes the lowest significant place in big endian). |

80s: Duplication Operations

| Value | Mnemonic | δ | α | Description |
|-------|----------|----------|----------|--|
| 0x80 | DUP1 | 1 | 2 | Duplicate 1st stack item. $\mu'_s[0] \equiv \mu_s[0]$ |
| 0x81 | DUP2 | 2 | 3 | Duplicate 2nd stack item. $\mu'_s[0] \equiv \mu_s[1]$ |
| ⋮ | ⋮ | ⋮ | ⋮ | ⋮ |
| 0x8f | DUP16 | 16 | 17 | Duplicate 16th stack item. $\mu'_s[0] \equiv \mu_s[15]$ |

90s: Exchange Operations

| Value | Mnemonic | δ | α | Description |
|-------|----------|----------|----------|--|
| 0x90 | SWAP1 | 2 | 2 | Exchange 1st and 2nd stack items. $\mu'_s[0] \equiv \mu_s[1]$ $\mu'_s[1] \equiv \mu_s[0]$ |
| 0x91 | SWAP2 | 3 | 3 | Exchange 1st and 3rd stack items. $\mu'_s[0] \equiv \mu_s[2]$ $\mu'_s[2] \equiv \mu_s[0]$ |
| ⋮ | ⋮ | ⋮ | ⋮ | ⋮ |
| 0x9f | SWAP16 | 17 | 17 | Exchange 1st and 17th stack items. $\mu'_s[0] \equiv \mu_s[16]$ $\mu'_s[16] \equiv \mu_s[0]$ |

a0s: Logging Operations

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

$$A'_1 \equiv A_1 \cdot (I_a, \mathbf{t}, \boldsymbol{\mu}_m[\boldsymbol{\mu}_s[0] \dots (\boldsymbol{\mu}_s[0] + \boldsymbol{\mu}_s[1] - 1)])$$

and to update the memory consumption counter:

$$\boldsymbol{\mu}'_i \equiv M(\boldsymbol{\mu}_i, \boldsymbol{\mu}_s[0], \boldsymbol{\mu}_s[1])$$

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

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

f0s: System operations

| Value | Mnemonic | δ | α | Description |
|-------|----------|----------|----------|--|
| 0xf0 | CREATE | 3 | 1 | <p>Create a new account with associated code. $\mathbf{i} \equiv \mu_{\mathbf{m}}[\mu_{\mathbf{s}}[1] \dots (\mu_{\mathbf{s}}[1] + \mu_{\mathbf{s}}[2] - 1)]$ $\zeta \equiv \emptyset$</p> $(\sigma', g', A', z, \mathbf{o}) \equiv \begin{cases} \Lambda(\sigma^*, A, I_a, I_o, L(\mu_g), I_p, \mu_{\mathbf{s}}[0], \mathbf{i}, I_e + 1, \zeta, I_w) & \text{if } \mu_{\mathbf{s}}[0] \leq \sigma[I_a]_b \wedge \\ & I_e < 1024 \wedge \ \mathbf{i}\ \leq 49152 \\ (\sigma, L(\mu_g), A, 0, ()) & \text{otherwise} \end{cases}$ <p>$\sigma^* \equiv \sigma$ except $\sigma^*[I_a]_n = \sigma[I_a]_n + 1$ $\mu'_g \equiv \mu_g - L(\mu_g) + g'$ $\mu'_{\mathbf{s}}[0] \equiv x$ where $x = 0$ if $z = 0$, i.e., the contract creation process failed, or $\ \mathbf{i}\ > 49152$ (initcode is longer than 49152 bytes), or $I_e = 1024$ (the maximum call depth limit is reached), or $\mu_{\mathbf{s}}[0] > \sigma[I_a]_b$ (balance of the caller is too low to fulfil the value transfer); and otherwise $x = \text{ADDR}(I_a, \sigma[I_a]_n, \zeta, \mathbf{i})$, the address of the newly created account (95). $\mu'_i \equiv M(\mu_i, \mu_{\mathbf{s}}[1], \mu_{\mathbf{s}}[2])$ $\mu'_o \equiv \begin{cases} () & \text{if } z = 1 \\ \mathbf{o} & \text{otherwise} \end{cases}$</p> <p>Thus the operand order is: value, input offset, input size.</p> |
| 0xf1 | CALL | 7 | 1 | <p>Message-call into an account. $\mathbf{i} \equiv \mu_{\mathbf{m}}[\mu_{\mathbf{s}}[3] \dots (\mu_{\mathbf{s}}[3] + \mu_{\mathbf{s}}[4] - 1)]$</p> $(\sigma', g', A', x, \mathbf{o}) \equiv \begin{cases} \Theta(\sigma, A^*, I_a, I_o, t, t, C_{\text{CALLGAS}}(\sigma, \mu, A), & \text{if } \mu_{\mathbf{s}}[2] \leq \sigma[I_a]_b \wedge \\ I_p, \mu_{\mathbf{s}}[2], \mu_{\mathbf{s}}[2], \mathbf{i}, I_e + 1, I_w) & I_e < 1024 \\ (\sigma, C_{\text{CALLGAS}}(\sigma, \mu, A), A^*, 0, ()) & \text{otherwise} \end{cases}$ <p>$n \equiv \min(\{\mu_{\mathbf{s}}[6], \ \mathbf{o}\ \})$ $\mu'_{\mathbf{m}}[\mu_{\mathbf{s}}[5] \dots (\mu_{\mathbf{s}}[5] + n - 1)] = \mathbf{o}[0 \dots (n - 1)]$ $\mu'_o = \mathbf{o}$ $\mu'_g \equiv \mu_g - C_{\text{CALLGAS}}(\sigma, \mu, A) + g'$ $\mu'_{\mathbf{s}}[0] \equiv x$ $A^* \equiv A$ except $A^*_a \equiv A_a \cup \{t\}$ $t \equiv \mu_{\mathbf{s}}[1] \bmod 2^{160}$ $\mu'_i \equiv M(M(\mu_i, \mu_{\mathbf{s}}[3], \mu_{\mathbf{s}}[4]), \mu_{\mathbf{s}}[5], \mu_{\mathbf{s}}[6])$ where $x = 0$ if the code execution for this operation failed, or if $\mu_{\mathbf{s}}[2] > \sigma[I_a]_b$ (not enough funds) or $I_e = 1024$ (call depth limit reached); $x = 1$ otherwise. Thus the operand order is: gas, to, value, in offset, in size, out offset, out size. $C_{\text{CALL}}(\sigma, \mu, A) \equiv C_{\text{GASCAP}}(\sigma, \mu, A) + C_{\text{EXTRA}}(\sigma, \mu, A)$ $C_{\text{CALLGAS}}(\sigma, \mu, A) \equiv \begin{cases} C_{\text{GASCAP}}(\sigma, \mu, A) + G_{\text{callstipend}} & \text{if } \mu_{\mathbf{s}}[2] \neq 0 \\ C_{\text{GASCAP}}(\sigma, \mu, A) & \text{otherwise} \end{cases}$ $C_{\text{GASCAP}}(\sigma, \mu, A) \equiv \begin{cases} \min\{L(\mu_g - C_{\text{EXTRA}}(\sigma, \mu, A)), \mu_{\mathbf{s}}[0]\} & \text{if } \mu_g \geq C_{\text{EXTRA}}(\sigma, \mu, A) \\ \mu_{\mathbf{s}}[0] & \text{otherwise} \end{cases}$ $C_{\text{EXTRA}}(\sigma, \mu, A) \equiv C_{\text{aaccess}}(t, A) + C_{\text{XFER}}(\mu) + C_{\text{NEW}}(\sigma, \mu)$ $C_{\text{XFER}}(\mu) \equiv \begin{cases} G_{\text{callvalue}} & \text{if } \mu_{\mathbf{s}}[2] \neq 0 \\ 0 & \text{otherwise} \end{cases}$ $C_{\text{NEW}}(\sigma, \mu) \equiv \begin{cases} G_{\text{newaccount}} & \text{if } \text{DEAD}(\sigma, t) \wedge \mu_{\mathbf{s}}[2] \neq 0 \\ 0 & \text{otherwise} \end{cases}$</p> |
| 0xf2 | CALLCODE | 7 | 1 | <p>Message-call into this account with an alternative account's code. Exactly equivalent to CALL except:</p> $(\sigma', g', A', x, \mathbf{o}) \equiv \begin{cases} \Theta(\sigma, A^*, I_a, I_o, I_a, t, C_{\text{CALLGAS}}(\sigma, \mu, A), & \text{if } \mu_{\mathbf{s}}[2] \leq \sigma[I_a]_b \wedge \\ I_p, \mu_{\mathbf{s}}[2], \mu_{\mathbf{s}}[2], \mathbf{i}, I_e + 1, I_w) & I_e < 1024 \\ (\sigma, C_{\text{CALLGAS}}(\sigma, \mu, A), A^*, 0, ()) & \text{otherwise} \end{cases}$ <p>Note the change in the fourth parameter to the call Θ from the 2nd stack value $\mu_{\mathbf{s}}[1]$ (as in CALL) to the present address I_a. This means that the recipient is in fact the same account as at present, simply that the code is overwritten.</p> |
| 0xf3 | RETURN | 2 | 0 | <p>Halt execution returning output data. $H_{\text{RETURN}}(\mu) \equiv \mu_{\mathbf{m}}[\mu_{\mathbf{s}}[0] \dots (\mu_{\mathbf{s}}[0] + \mu_{\mathbf{s}}[1] - 1)]$ This has the effect of halting the execution at this point with output defined. See section 9. $\mu'_i \equiv M(\mu_i, \mu_{\mathbf{s}}[0], \mu_{\mathbf{s}}[1])$</p> |

| | | | | |
|------|--------------|-------------|-------------|---|
| 0xf4 | DELEGATECALL | 6 | 1 | <p>Message-call into this account with an alternative account's code, but persisting the current values for <i>sender</i> and <i>value</i>.</p> <p>Compared with CALL, DELEGATECALL takes one fewer arguments. The omitted argument is $\mu_s[2]$. As a result, $\mu_s[3]$, $\mu_s[4]$, $\mu_s[5]$ and $\mu_s[6]$ in the definition of CALL should respectively be replaced with $\mu_s[2]$, $\mu_s[3]$, $\mu_s[4]$ and $\mu_s[5]$. Otherwise it is equivalent to CALL except:</p> $(\sigma', g', A', x, \mathbf{o}) \equiv \begin{cases} \Theta(\sigma, A^*, I_s, I_o, I_a, t, C_{\text{CALLGAS}}(\sigma, \mu, A), & \text{if } I_e < 1024 \\ I_p, 0, I_v, \mathbf{i}, I_e + 1, I_w) & \\ (\sigma, C_{\text{CALLGAS}}(\sigma, \mu, A), A^*, 0, ()) & \text{otherwise} \end{cases}$ <p>Note the changes (in addition to that of the fourth parameter) to the second and ninth parameters to the call Θ.</p> <p>This means that the recipient is in fact the same account as at present, simply that the code is overwritten <i>and</i> the context is almost entirely identical.</p> |
| 0xf5 | CREATE2 | 4 | 1 | <p>Create a new account with associated code.</p> <p>Exactly equivalent to CREATE except: The salt $\zeta \equiv \mu_s[3]$.</p> |
| 0xfa | STATICCALL | 6 | 1 | <p>Static message-call into an account.</p> <p>Exactly equivalent to CALL except: The argument $\mu_s[2]$ is replaced with 0. The deeper argument $\mu_s[3]$, $\mu_s[4]$, $\mu_s[5]$ and $\mu_s[6]$ are respectively replaced with $\mu_s[2]$, $\mu_s[3]$, $\mu_s[4]$ and $\mu_s[5]$. The last argument of Θ is \perp.</p> |
| 0xfd | REVERT | 2 | 0 | <p>Halt execution reverting state changes but returning data and remaining gas.</p> $H_{\text{RETURN}}(\mu) \equiv \mu_m[\mu_s[0] \dots (\mu_s[0] + \mu_s[1] - 1)]$ <p>The effect of this operation is described in (152).</p> <p>For the gas calculation, we use the memory expansion function, $\mu'_i \equiv M(\mu_i, \mu_s[0], \mu_s[1])$</p> |
| 0xfe | INVALID | \emptyset | \emptyset | Designated invalid instruction. |
| 0xff | SELFDESTRUCT | 1 | 0 | <p>Halt execution and register account for later deletion. Readers should be warned, since the <i>Shanghai</i> update, this opcode has been marked as deprecated. Its behavior might be subject to breaking change in an upcoming update.</p> $A'_s \equiv A_s \cup \{I_a\}$ $A'_a \equiv A_a \cup \{r\}$ $\sigma'[r] \equiv \begin{cases} \emptyset & \text{if } \sigma[r] = \emptyset \wedge \sigma[I_a]_b = 0 \\ (\sigma[r]_n, \sigma[r]_b + \sigma[I_a]_b, \sigma[r]_s, \sigma[r]_c) & \text{if } r \neq I_a \\ (\sigma[r]_n, 0, \sigma[r]_s, \sigma[r]_c) & \text{otherwise} \end{cases}$ <p>where $r = \mu_s[0] \bmod 2^{160}$</p> $\sigma'[I_a]_b = 0$ $C_{\text{SELFDESTRUCT}}(\sigma, \mu) \equiv G_{\text{selfdestruct}} + \begin{cases} 0 & \text{if } r \in A_a \\ G_{\text{coldaccountaccess}} & \text{otherwise} \end{cases} + \begin{cases} G_{\text{newaccount}} & \text{if } \text{DEAD}(\sigma, r) \wedge \sigma[I_a]_b \neq 0 \\ 0 & \text{otherwise} \end{cases}$ |

APPENDIX I. GENESIS BLOCK

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

$$(333) \quad ((0_{256}, \text{KEC}(\text{RLP}(()))), 0_{160}, \text{stateRoot}, 0, 0, 0_{2048}, 2^{34}, 0, 0, 3141592, \text{time}, 0, 0_{256}, \text{KEC}((42))), (), ())$$

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, 2048-bit of all zeros; 2^{34} refers to the difficulty; the transaction trie root, receipt trie root, gas used, block number and extradata are both 0, being equivalent to the empty byte array. The sequences of both ommers and transactions are empty and represented by $()$. $\text{KEC}((42))$ refers to the Keccak-256 hash of a byte array of length one whose first and only byte is of value 42, used for the nonce. $\text{KEC}(\text{RLP}(()))$ value refers to the hash of the ommer list in RLP, both empty lists.

The proof-of-concept series include 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. **Deprecation.** This section is kept for historical purposes, but the Ethash algorithm is no longer used for consensus as of the *Paris* hard fork.

J.2. **Definitions.** We employ the following definitions:

| Name | Value | Description |
|----------------------------|----------|--|
| $J_{\text{wordbytes}}$ | 4 | Bytes in word. |
| $J_{\text{datasetinit}}$ | 2^{30} | Bytes in dataset at genesis. |
| $J_{\text{datasetgrowth}}$ | 2^{23} | Dataset growth per epoch. |
| $J_{\text{cacheinit}}$ | 2^{24} | Bytes in cache at genesis. |
| $J_{\text{cachegrowth}}$ | 2^{17} | Cache growth per epoch. |
| J_{epoch} | 30000 | Blocks per epoch. |
| J_{mixbytes} | 128 | mix length in bytes. |
| $J_{\text{hashbytes}}$ | 64 | Hash length in bytes. |
| J_{parents} | 256 | Number of parents of each dataset element. |
| $J_{\text{cacherrounds}}$ | 3 | Number of rounds in cache production. |
| J_{accesses} | 64 | Number of accesses in hashimoto loop. |

J.3. **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.

$$(334) \quad E_{\text{epoch}}(H_i) = \left\lfloor \frac{H_i}{J_{\text{epoch}}} \right\rfloor$$

The size of the dataset growth by $J_{\text{datasetgrowth}}$ bytes, and the size of the cache by $J_{\text{cachegrowth}}$ bytes, 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

$$(335) \quad d_{\text{size}} = E_{\text{prime}}(J_{\text{datasetinit}} + J_{\text{datasetgrowth}} \cdot E_{\text{epoch}} - J_{\text{mixbytes}}, J_{\text{mixbytes}})$$

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

$$(336) \quad c_{\text{size}} = E_{\text{prime}}(J_{\text{cacheinit}} + J_{\text{cachegrowth}} \cdot E_{\text{epoch}} - J_{\text{hashbytes}}, J_{\text{hashbytes}})$$

$$(337) \quad E_{\text{prime}}(x, y) = \begin{cases} x & \text{if } x/y \in \mathbb{N} \\ E_{\text{prime}}(x - 2 \cdot y, y) & \text{otherwise} \end{cases}$$

J.4. **Dataset generation.** In order to generate the dataset we need the cache \mathbf{c} , which is an array of bytes. It depends on the cache size c_{size} and the seed hash $\mathbf{s} \in \mathbb{B}_{32}$.

J.4.1. *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:

$$(338) \quad \mathbf{s} = C_{\text{seedhash}}(H_i)$$

$$(339) \quad C_{\text{seedhash}}(H_i) = \begin{cases} \mathbf{0}_{32} & \text{if } E_{\text{epoch}}(H_i) = 0 \\ \text{KEC}(C_{\text{seedhash}}(H_i - J_{\text{epoch}})) & \text{otherwise} \end{cases}$$

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

J.4.2. *Cache.* The cache production process involves using the seed hash to first sequentially filling up c_{size} bytes of memory, then performing $J_{\text{cacherrounds}}$ 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 i th element of the initial cache:

$$(340) \quad \mathbf{c}_i = \begin{cases} \text{KEC512}(\mathbf{s}) & \text{if } i = 0 \\ \text{KEC512}(\mathbf{c}_{i-1}) & \text{otherwise} \end{cases}$$

Therefore \mathbf{c}' can be defined as

$$(341) \quad \mathbf{c}'[i] = \mathbf{c}_i \quad \forall \quad i < n$$

$$(342) \quad n = \left\lfloor \frac{c_{\text{size}}}{J_{\text{hashbytes}}} \right\rfloor$$

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

$$(343) \quad \mathbf{c} = E_{\text{cacherrounds}}(\mathbf{c}', J_{\text{cacherrounds}})$$

$$(344) \quad E_{\text{cachrounds}}(\mathbf{x}, y) = \begin{cases} \mathbf{x} & \text{if } y = 0 \\ E_{\text{RMH}}(\mathbf{x}) & \text{if } y = 1 \\ E_{\text{cachrounds}}(E_{\text{RMH}}(\mathbf{x}), y - 1) & \text{otherwise} \end{cases}$$

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

$$(345) \quad E_{\text{RMH}}(\mathbf{x}) = (E_{\text{rmh}}(\mathbf{x}, 0), E_{\text{rmh}}(\mathbf{x}, 1), \dots, E_{\text{rmh}}(\mathbf{x}, n - 1))$$

$$(346) \quad E_{\text{rmh}}(\mathbf{x}, i) = \text{KEC512}(\mathbf{x}'[(i - 1 + n) \bmod n] \oplus \mathbf{x}'[\mathbf{x}'[i][0] \bmod n])$$

with $\mathbf{x}' = \mathbf{x}$ except $\mathbf{x}'[j] = E_{\text{rmh}}(\mathbf{x}, j) \quad \forall j < i$

J.4.3. 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:

$$(347) \quad \mathbf{d}[i] = E_{\text{datasetitem}}(\mathbf{c}, i) \quad \forall i < \left\lfloor \frac{d_{\text{size}}}{J_{\text{hashbytes}}} \right\rfloor$$

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.

$$(348) \quad E_{\text{FNV}}(\mathbf{x}, \mathbf{y}) = (\mathbf{x} \cdot (0\text{x}01000193 \oplus \mathbf{y})) \bmod 2^{32}$$

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

$$(349) \quad E_{\text{datasetitem}}(\mathbf{c}, i) = E_{\text{parents}}(\mathbf{c}, i, -1, \emptyset)$$

$$(350) \quad E_{\text{parents}}(\mathbf{c}, i, p, \mathbf{m}) = \begin{cases} E_{\text{parents}}(\mathbf{c}, i, p + 1, E_{\text{mix}}(\mathbf{m}, \mathbf{c}, i, p + 1)) & \text{if } p < J_{\text{parents}} - 2 \\ E_{\text{mix}}(\mathbf{m}, \mathbf{c}, i, p + 1) & \text{otherwise} \end{cases}$$

$$(351) \quad E_{\text{mix}}(\mathbf{m}, \mathbf{c}, i, p) = \begin{cases} \text{KEC512}(\mathbf{c}[i \bmod c_{\text{size}}] \oplus i) & \text{if } p = 0 \\ E_{\text{FNV}}(\mathbf{m}, \mathbf{c}[E_{\text{FNV}}(i \oplus p, \mathbf{m}[p \bmod \lfloor J_{\text{hashbytes}}/J_{\text{wordbytes}} \rfloor])] \bmod c_{\text{size}}) & \text{otherwise} \end{cases}$$

J.5. 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 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:

$$(352) \quad \text{PoW}(H_{\mathcal{H}}, H_n, \mathbf{d}) = \{\mathbf{m}_c(\text{KEC}(\text{RLP}(L_{\mathcal{H}}(H_{\mathcal{H}}))), H_n, \mathbf{d}), \text{KEC}(\mathbf{s}_h(\text{KEC}(\text{RLP}(L_{\mathcal{H}}(H_{\mathcal{H}}))), H_n) + \mathbf{m}_c(\text{KEC}(\text{RLP}(L_{\mathcal{H}}(H_{\mathcal{H}}))), H_n, \mathbf{d}))\}$$

With $H_{\mathcal{H}}$ being the hash of the header without the nonce. The compressed mix \mathbf{m}_c is obtained as follows:

$$(353) \quad \mathbf{m}_c(\mathbf{h}, \mathbf{n}, \mathbf{d}) = E_{\text{compress}}(E_{\text{accesses}}(\mathbf{d}, \sum_{i=0}^{n_{\text{mix}}} \mathbf{s}_h(\mathbf{h}, \mathbf{n}), \mathbf{s}_h(\mathbf{h}, \mathbf{n}), -1), -4)$$

The seed hash being:

$$(354) \quad \mathbf{s}_h(\mathbf{h}, \mathbf{n}) = \text{KEC512}(\mathbf{h} + E_{\text{revert}}(\mathbf{n}))$$

$E_{\text{revert}}(\mathbf{n})$ returns the reverted bytes sequence of the nonce \mathbf{n} :

$$(355) \quad E_{\text{revert}}(\mathbf{n})[i] = \mathbf{n}[|\mathbf{n}| - i]$$

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.4.3.

The number of replicated sequences in the mix is:

$$(356) \quad n_{\text{mix}} = \left\lfloor \frac{J_{\text{mixbytes}}}{J_{\text{hashbytes}}} \right\rfloor$$

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

$$(357) \quad E_{\text{accesses}}(\mathbf{d}, \mathbf{m}, \mathbf{s}, i) = \begin{cases} E_{\text{mixdataset}}(\mathbf{d}, \mathbf{m}, \mathbf{s}, i) & \text{if } i = J_{\text{accesses}} - 2 \\ E_{\text{accesses}}(E_{\text{mixdataset}}(\mathbf{d}, \mathbf{m}, \mathbf{s}, i), \mathbf{s}, i + 1) & \text{otherwise} \end{cases}$$

$$(358) \quad E_{\text{mixdataset}}(\mathbf{d}, \mathbf{m}, \mathbf{s}, i) = E_{\text{FNV}}(\mathbf{m}, E_{\text{newdata}}(\mathbf{d}, \mathbf{m}, \mathbf{s}, i))$$

E_{newdata} returns an array with n_{mix} elements:

$$(359) \quad E_{\text{newdata}}(\mathbf{d}, \mathbf{m}, \mathbf{s}, i)[j] = \mathbf{d}[E_{\text{FNV}}(i \oplus \mathbf{s}[0], \mathbf{m}[i \bmod \left\lfloor \frac{J_{\text{mixbytes}}}{J_{\text{wordbytes}}} \right\rfloor})] \bmod \left\lfloor \frac{d_{\text{size}}/J_{\text{hashbytes}}}{n_{\text{mix}}} \right\rfloor \cdot n_{\text{mix}} + j] \quad \forall \quad j < n_{\text{mix}}$$

The mix is compressed as follows:

$$(360) \quad E_{\text{compress}}(\mathbf{m}, i) = \begin{cases} \mathbf{m} & \text{if } i \geq \|\mathbf{m}\| - 8 \\ E_{\text{compress}}(E_{\text{FNV}}(E_{\text{FNV}}(E_{\text{FNV}}(\mathbf{m}[i+4], \mathbf{m}[i+5]), \mathbf{m}[i+6]), \mathbf{m}[i+7]), i+8) & \text{otherwise} \end{cases}$$

APPENDIX K. ANOMALIES ON THE MAIN NETWORK

K.1. Deletion of an Account Despite Out-of-gas. At block 2675119, in the transaction 0xcf416c536ec1a19ed1fb89e4ec7ffb3cf73aa413b3aa9b77d60e4fd81a4296ba, an account at address 0x03 was called and an out-of-gas occurred during the call. Against the equation (209), this added 0x03 in the set of touched addresses, and this transaction turned $\sigma[0x03]$ into \emptyset .

APPENDIX L. LIST OF MATHEMATICAL SYMBOLS

| Symbol | Latex Command | Description |
|-----------|----------------------|--|
| \bigvee | <code>\bigvee</code> | This is the least upper bound, supremum, or join of all elements operated on. Thus it is the greatest element of such elements (Davey and Priestley [2002]). |