

Enso

A general-purpose virtual machine

Bruno França
Trinkler Software
company@trinkler.software

Version 3 – March 4, 2019


Abstract—In this paper we introduce Enso, a virtual machine designed to be used as general-purpose state transition function in blockchains. This design allows the blockchain application logic to be coded into the state, instead of into the state transition function, making it much more flexible and easier to modify. A byproduct is reducing the frequency of forks, concerted or not. Finally, we discuss how Enso can easily be implemented as a runtime module on Parity’s Substrate.

I. INTRODUCTION

Despite the diversity and complexity in the blockchain world today, all blockchains are just distributed virtual machines.

Independently of the function for which it is used, a blockchain is just a virtual machine running some code distributed across several computers.

The components of such a distributed virtual machine can be visualized as a stack:




Starting at the bottom, we have the networking. This basically includes all the tools that the computer needs to connect to the other nodes and to receive, transmit and relay data across the network. It is a standard component for any peer-to-peer network.

Then we have the consensus. This allows all nodes to reach an agreement on what blocks are accepted. There is a huge variety of consensus protocols, but they all have the same objective: making all honest nodes reach an agreement on some input data.

Lastly, we have the state and the state transition function (STF). State represents, of course, the current state of the virtual machine. The state, for most users, is the only component that they interact with. The state transition function is the function that takes as input the previous state and a block, and produces a new state as output. So, it is what defines the rules for modifying the state. Contrary to the networking and consensus, these two components no longer interact with, or

have any notion of, other nodes. Together they just form a simple virtual machine.

The way in which each of these components is modified, or upgraded, varies from blockchain to blockchain. The following image illustrates how legacy chains, like Bitcoin [1] or Ethereum [2] [3], are modified:



The state is always modified with extrinsics¹. This is how most people will interact with the blockchain and it is very straightforward.

Next we have the state transition function and the consensus. Both can be modified with forks, either soft forks or hard forks. The problem with forking is that forks are difficult to organize and can cause the chain to split into two networks (ex: Bitcoin and Bitcoin Cash, or Ethereum and Ethereum Classic).

Finally we have the network. Changes in the network normally are not contentious since they can happen with gradual user adoption. For example, some nodes in the Bitcoin network can start accepting transactions via Tor and relaying them to miners. There is no need for any type of consensus on whether transactions should be accepted via Tor. Any node that wishes to do so can make that change and make that service available to the network right away. So, networking updates happen with organic user adoption.

Recently, blockchains appeared which claim to be upgradable by *governance*. We are only interested in the technical aspect of upgrading a blockchain and not in the social or political aspect (ex: if voting should be proportional to stake, etc). The most famous examples in this category are probably Tezos [4] and Parity’s Substrate (Substrate is not actually a blockchain but rather a blockchain development kit). With this type of chain we have the following stack:

¹We use same definition for transaction as the one taken by Parity. There are extrinsics, which are any input to the state transition function, and both transactions and inherents are mutually exclusive types of extrinsics. Transactions are extrinsics that are propagated through the network and signed. Inherents are neither propagated nor signed. An example of an inherent is a timestamp.

This work is licensed under the Creative Commons Attribution-ShareAlike 4.0 International License.

State	←	Extrinsics
STF	←	Swaps
Consensus	←	Swaps
Networking	←	Adoption

Note that the state and networking are still modified in the same fashion as legacy chains, the difference is that the state transition function (STF) and the consensus use *module swaps* instead of forks.

In these blockchains, there is a shell that contains only the networking and some other auxiliary code. Then there are two slots (one for the STF, another for the consensus) that accept arbitrary code blocks, called *modules*.

The idea is that a blockchain can be modified by simply swapping one module for another. If the blockchain also has governance functionality coded into it, it is possible for the nodes to reach an agreement on whether a module should be swapped or not.

Governance and module swaps are different components. Governance only allows a network to reach an agreement on whether the blockchain is to be updated or not. The actual method of doing that update is the module swapping. Objectively, the main difference between forks and swaps is that swaps are concerted, while forks are not. Otherwise the process is similar. The miners (or validator, or authorities) all agree to, at a specific time, stop using one module and start using the next one.

The breakthrough on the part of Tezos was realizing that legacy blockchains have their governance *off-chain* and that instead the governance should be placed *on-chain* into the STF.

Another way of analyzing this stack is seeing what functionality is coded into each component. As we have said, a blockchain is a distributed virtual machine but, of course, each blockchain is running a specific application (also commonly called the *runtime*). For example, Bitcoin is a distributed virtual machine running a ledger application. We can visualize where each part of the application is located in the stack:

App data	≡	State
App logic	≡	STF
		Consensus
		Networking

The application logic and data are separated, with the logic being hardcoded into the STF and the data being stored into the state.

This separation of *code* and *data* is how early computers worked. Programs were hardcoded into the hardware using plugboards and then data would be fed into the computer using some storage medium (for example, punch cards). This was time-consuming, requiring the plugboard to be rewired whenever a new program was to be run in the computer. Then along came the stored-program computer, which treats code and data equally. The paradigm change was thinking of the hardware as a general-purpose shell without any specific functionality predefined into it. Now programs could be fed into the computer just like regular data.

Enso applies the stored-program computer concept to blockchains. It adopts a very simple general-purpose virtual machine as the state transition function so that the application logic can be put into the state. This allows for both the application logic and data to be modified with extrinsics, while the consensus and the virtual machine itself are still modified with forks or module swaps.

App logic and data	≡	State	←	Extrinsics
Virtual machine	≡	STF	←	Forks/Swaps
		Consensus	←	Forks/Swaps
		Networking	←	Adoption

By treating the logic and the data equally, blockchains can be built with a greater degree of flexibility. And extrinsics allow for simpler and more granular modifications to the application logic. Changes to the consensus or to the virtual machine still need to be done using the module swaps, but we expect that consensus and virtual machine updates will be fairly rare, at least when compared with changes to the application logic.

II. SPECIFICATION

In this section we will delve into several aspects of the Enso virtual machine. Beginning by formally defining what a blockchain is.

A blockchain is an ordered series of states produced by the repeated application of a state transition function. Let us denote the state as S and a block as B , then the state transition function can be defined as:

$$STF(B_t, S_t) \rightarrow S_{t+1}$$

A block in this context is just an input to the state transition function, having no further structure. The consensus is not *explicitly* defined in this framework, but rather *implicitly* because the consensus is necessary to reach an agreement over a set of inputs to the state transition function and its order.

The current state is then defined as the result of applying the state transition function repeatedly using as inputs the *genesis* state S_0 and the blocks on which a consensus was reached.

$$S_n = STF(B_{n-1}, \dots STF(B_1, STF(B_0, S_0)) \dots)$$

Now that we defined what a blockchain is, we can discuss what a blockchain that uses Enso looks like. For Enso everything is either an *object* or an *event*.

A. Objects

An object is an entity composed of the following components:

- *ID*: An unique identifier of the object. It can be any string, as long as it is unique.
- *Code*: A block of code containing functions that can be called by other objects.
- *Storage*: A data structure containing arbitrary information and that can be read and modified by the code.

Note that objects, because of their general purpose code and storage, can support any kind of logic. Also, objects are dynamic and can have any part of them modified at any time.

Using Enso, the state of a blockchain will be simply the set of all objects. So, all of the application logic and data will be coded into the objects.

B. Events

Events, put simply, are *asynchronous* function calls from one object to another. By asynchronous we mean that when an object creates an event, the object will not be blocked while it waits for the event to be processed. Instead it will keep working and accepting other events.

This asynchronicity is necessary because we want the virtual machine to keep running even if a function call fails.

An event e takes the form $e(\text{priority}, \text{ID_to}, \text{function_call}, \text{parameters})$, where:

- priority : The priority of the event, used for the event queue.
- ID_to : The ID of the object that will receive the event.
- function_call : The name of the function that will be called.
- parameters : A list of parameters that will be passed to the function.

Now that we know that the state is composed of objects, and that objects communicate between themselves with events, we can describe the inner workings of the Enso virtual machine. It is composed of three components: the *super object*, the *events queue* and the *input rules*. We will now specify each one of them.

C. Super Object

The super object is a special object that has the following properties:

- It is an unique object, there can only be one super object in the state.
- Its ID and code are static, they cannot be changed or modified in any way.
- It is the only object that can modify other objects.

The name '*super object*' is a reference to the *super user* in Linux systems, and is meant to signify that it has complete control over the state.

The super object code is composed of the following functions:

- $\text{Create}(\text{ID}, \text{code}, \text{storage})$: Creates a new object with the given ID, code and storage.
- $\text{Delete}(\text{ID})$: Deletes object with the given ID.
- $\text{Check}(\text{ID})$: Checks if any object exists with the given ID and returns the answer.
- $\text{Request_object}(\text{ID})$: Returns the code and storage of the object with the given ID.
- $\text{Change_ID}(\text{ID}, \text{new_ID})$: Changes the ID of the object with the given ID to new_ID .
- $\text{Change_code}(\text{ID}, \text{new_code})$: Changes the code of the object with the given ID to new_code .

- $\text{Change_storage}(\text{ID}, \text{new_storage})$: Changes the storage of the object with the given ID to new_storage .
- $\text{Set_input}(\text{ID})$: Sets the input object to the object with given ID.

These functions can be used to arbitrarily change the state. The last function can change the input object of the state, i.e., the object that all inputs are initially fed into.

D. Events queue

The events queue is just a simple priority queue for events. So, events are added to the queue and then events with a higher priority are processed before events with a lower priority. Events with the same priority are processed on a *first in first out* basis.

When an event is processed, it is taken out of the queue and the function function_call of the object ID_to is called with the input parameters parameters . If, for some reason, the function cannot be called, for example if the object doesn't exist or there is no function with that name, the event is simply ignored.

E. Input rules

The input rules determine how extrinsics are fed into the Enso virtual machine and processed, thus resulting into a new state.

First, there is always a single *input object*, which is the object that is designated by the super object to receive extrinsics. Extrinsics in Enso are events just like any other, but with the caveat that they must have $\text{ID_to}=\text{input_object_ID}$. Any object can be the input object, the only requirement is that it designated as such by the super object.

Second, a block is just an ordered list of extrinsics. When the virtual machine receives a new block, it just needs to process those extrinsics in order and in the end it will have reached a new state.

So, to process a new block, a node follows these rules:

- 1) Receive an ordered list of extrinsics.
- 2) Add the first extrinsic to the event queue.
- 3) The resulting event will be sent to the input object. The input object then may create more events, that in turn will also be processed and create more events, and so forth, until no more events are created.
- 4) Add the next extrinsic to the event queue and keep repeating steps 3 and 4 until there are no more extrinsics.
- 5) When there are no more extrinsics, or events, to be processed, the virtual machine has reached a new state.

III. IMPLEMENTATION

The objective of the last section was to describe the components of the Enso virtual machine and its operation in an abstract form. In this section we want to give some details on how Enso can be implemented in practice.

A. Objects

The ID of an object is only required to be unique and it is desirable to have human-readable IDs. Therefore, the ID format can be defined by choosing some alphabet (ex: lower-case letters, upper-case letters and numbers) and some maximum length. An idea would be to use email address standards since the users are already familiar with these.

The code of an object can, in theory, be written in any language as long as it supports function calls of some kind. In practice, we also want it to be fast to execute and easy to write. For these reasons, we believe the code for objects should be in WebAssembly, since its speed rivals that of native machine code and supports compilation from a variety of high-level languages like C, C++, and Rust.

The storage of an object does not need to be strongly defined because normally it will only be accessed by the code of the object itself. So, different objects can have different data structures as storage without any issues arising. The best option here then is to have the storage of an object be a binary blob, which is just a chunk of binary data.

B. Kernel

Even though both the application data and logic will coexist in the state, in most cases there will be a clear separation between both. By this we mean that, when developing a new blockchain with Enso, it is useful to divide the objects into two categories:

- **Kernel objects:** All objects created at the genesis block and during updates to the blockchain. These are objects that define the rules for how the runtime works and are unique objects, meaning that there is only instance of each object type.
- **User objects:** All objects created by the users. There is a predetermined list of object templates, from which users can create new user objects.

The idea is that the entire application logic will be coded into the kernel objects, for example we may have one object to handle the incoming extrinsics, another that implements some governance method, etc. While the user objects will be regular objects like accounts or smart contracts, effectively representing the application data.

C. Substrate

Since Enso only aims to substitute the state transition function of most blockchains, there is no need to develop an entire blockchain around it. Instead, we intend to implement Enso as a runtime module for Parity's Substrate Core. Substrate Core has an `execute_block` function that is general enough to support Enso.

IV. CONCLUSION

This paper analyzes the organization of a blockchain code by thinking of it as a distributed virtual machine. We then argue that a distributed virtual machine needs to act analogously to a stored-program computer, which implies that the application logic needs to be coded into the state and that the

state transition function needs to be a general-purpose virtual machine. Using this method it is possible to have blockchains that are simpler and faster to develop and operate, and that require far fewer forks.

REFERENCES

- [1] S. Nakamoto *et al.*, "Bitcoin: A peer-to-peer electronic cash system," 2008. [Online]. Available: <https://bitcoin.org/bitcoin.pdf>
- [2] V. Buterin *et al.*, "A next-generation smart contract and decentralized application platform," *white paper*, 2014. [Online]. Available: https://cryptorating.eu/whitepapers/Ethereum/Ethereum_white_paper.pdf
- [3] G. Wood, "Ethereum: A secure decentralised generalised transaction ledger," *Ethereum project yellow paper*, vol. 151, pp. 1–32, 2014. [Online]. Available: <http://gavwood.com/paper.pdf>
- [4] L. Goodman, "Tezos: a self-amending crypto-ledger," *URL: https://www.tezos.com/static/papers/white_paper.pdf*, 2014. [Online]. Available: https://everipedia-storage.s3.amazonaws.com/NewlinkFiles/16739988/8d4d1f1b-8/white_paper.pdf