

HyperCycle: A Lightweight Agent-System-Based Ledgerless Blockchain Architecture, Supporting Secure, Efficient and Scalable AI Microservices, Agnostic and Interoperable

Ben Goertzel¹, Toufi Saliba², Dann Toliver² Anton Kolonin¹, Sergei Rodionov¹

¹ SingularityNET Foundation

² TODA Network
April 19, 2023

Abstract

A novel true p2p ledgerless blockchain implementation of Toda/IP is presented, oriented toward secure, yet inexpensive, high-speed large-scale on-chain execution of microservices, especially (but not exclusively) those carrying out AI-related functions. This HyperCycle "network" based blockchain can leverage several smart contract language (and underlying UTxO model) and TODA/IP p2p communication protocol for ownerships, the TODA asset model, and SingularityNET's Proof of Reputation system. We call it Layer Zero Sub Plus Plus, Layer 0++. A HyperCycle network consists of a population of autonomous agents each owning their own transaction history and accumulating their own reputation, grouped together in rings in which they collaborate to execute consensus mechanisms enabling execution of nano transactions based on cryptographic proofs of computation performed in a p2p setting with zero third party dependencies. HyperCycle agents are constructed as specific TODA files/packets (Sato-Servers), complete with the TODA transaction and cycle trie mech-

anisms for secure strongly decentralized management of file and network history (the cycle trie being the inspiration for the name "HyperCycle"). Consensus mechanisms involved in HyperCycle rings may be ledgerless or e.g. based on hierarchically sharded ledgers, depending on the particular requirements. Any ledger can be optionally used but it never depends on the ledgers. Some of the usage can be for broadcasting proofs while others for long term global memory, and for ledger-based long-term backup storage of a selected fraction of on-HyperCycle transactions or aggregate of transactions. Brief discussion is given on the particulars of HyperCycle customization to key application areas like swarm AI, rating and reward in media networks, decentralized payments and computer processing, and public/private chain interoperability. Integration of the OpenCog Hyperon frameworks' MeTTa programming language into Plutus is also envisioned, as a strategy for enabling broad usability via the creation of MeTTa-based DSLs focused on smart contract programming in specific vertical areas. Such DSLs may straightforwardly be presented to the user as low or no code frameworks auto-generated from the underlying MeTTa/-Plutus code, effectively leveraging the HyperCycle framework behind the scenes.

Contents

1	Introduction	3
2	Background: Any Chain such as Cardano, Hydra, TODA/IP, Proof of Reputation	9
2.1	Any UTxO Transaction Model	9
2.2	Hydra: Cardano's Interoperability Solution	11
	2.2.1 Head Protocol Transactions	12
	2.2.2 Hydra Tail Protocol	14
2.3	The TODA/IP ledgerless blockchain	15
	2.3.1 TODA/IP Transaction Processing: Decentralized and Localized	17
	2.3.2 TODA: Interoperable Proof Structures	19
	2.3.3 Key TODA Data Structures and the TODA Cycle	20
2.4	Proof of Reputation	21

3	HyperCycle	23
3.1	Core HyperCycle Architecture Concepts	24
3.2	HyperCycle Consensus Mechanisms	25
3.2.1	Lightweight Rings	25
3.2.2	Hierarchical Sharding Enhanced Consensus Mechanisms	26
3.2.3	Further Notes on Security, Reliability and Performance Tradeoffs	27
3.3	Accelerated Smart Contract Execution	32
3.4	The MeTTa AI Language As A Plutus Framework	34
3.4.1	Smart Contract DSLs for HyperCycle: A Path to Broad Usability	34
4	A Few Promising Early Applications	35
4.1	Swarm AI: Evolutionary Learning, Algorithmic Chemistry, En- semble ML	36
4.2	Ratings and Rewards in Online Networks	36
4.3	Decentralizing Payments and Processing Power	37
4.4	Adaptively Blending Public and Private Chains	37
5	Conclusion: Catalyzing the Future Of AI and Blockchain Ecosystem	38

1 Introduction

In the time since the Bitcoin whitepaper was released in 2008, a variety of different blockchain technologies and networks have been created, leveraging similar to Bitcoin underlying algorithms and data structures and in some cases significant evolutions thereof. New concepts such as smart contracts, DAOs and multi-layer networks have been added to the picture which adds more dependencies in what's supposed to p2p. It has also become clear that there may not be end up being One Blockchain Network To Rule Them All; rather, we may be facing an ecosystem comprising a decentralized network of decentralized networks, the various networks each architected with underlying mechanisms that combine roughly the same underlying math and computer science in ways specialized to suit particular classes of use-cases.

The HyperCycle design could leverage any existing blockchains and their layers as HyperCycle is by construct interoperable. It truly sits inside the net-

work packets and operates on that level below any and all existing siloed AIs, blockchains and ledgers. HyperCycle integrates smartcontracts executions using an extremely secure and deterministic VM, yet unprecedentedly efficient. This concepts with core algorithms and structures from the TODA asset model [**TODAPrimer**] and the TODA/IP ledgerless blockchain [**TODAWP**] , and with the Proof of Reputation concept [**aluko2021proof**] developed for use within the SingularityNET AI/blockchain network.

Part of the motivation for the HyperCycle design is the requirements of the SingularityNET protocol

[**GoertzelSNet2019**]

[**goertzel2017singularitynet**]

and [**montes2019distributed**].

SingularityNET is a decentralized platform for communication and coordination between AI agents and their users, originally implemented on Ethereum and now in the midst of a port to Any Ledger Blockchain such as Cardano. The Cardano infrastructure still imposes sufficient overhead to drive SingularityNET AI agent design toward patterns in which individual agents carry out massive and complex AI processes internally at a modular level and globally open to every single connected AI, using the blockchain only for relatively infrequent communications. Via radically decreasing the cost and increasing the speed of blockchain transactions, HyperCycle allows design patterns involving lighter-weight AI agents interoperating more frequently and therefore resulting in emergence that improves the totality of the intelligence especially when taken at an exponential level. Of course these improvements will come with some tradeoffs, which are acceptable in the SingularityNET application context but would be unappealing in some other applications. In this paper we may have mentions of certain ledger based blockchains more than others, that doesn't mean the implementation is mutually exclusive. In fact every utilization of existing blockchains has been thought out mainly to enhance their existing ecosystems and maximize the synergistic emergence without having to turn any existing ledger based blockchain into dependency layer. At the time we wrote this paper, we felt the HyperCycle Implementation of Toda/IP can have several deployments some more focused on existing ledger based blockchains, like Bitcoin, Cardano, Ethereum, Cosmos, Polkadot, Algorand, Avalanche, Avalanche (to name a few) while a major deployment that will always be agnostic yet can still be interoperable with other deployments and can of course be pluggable accordingly.

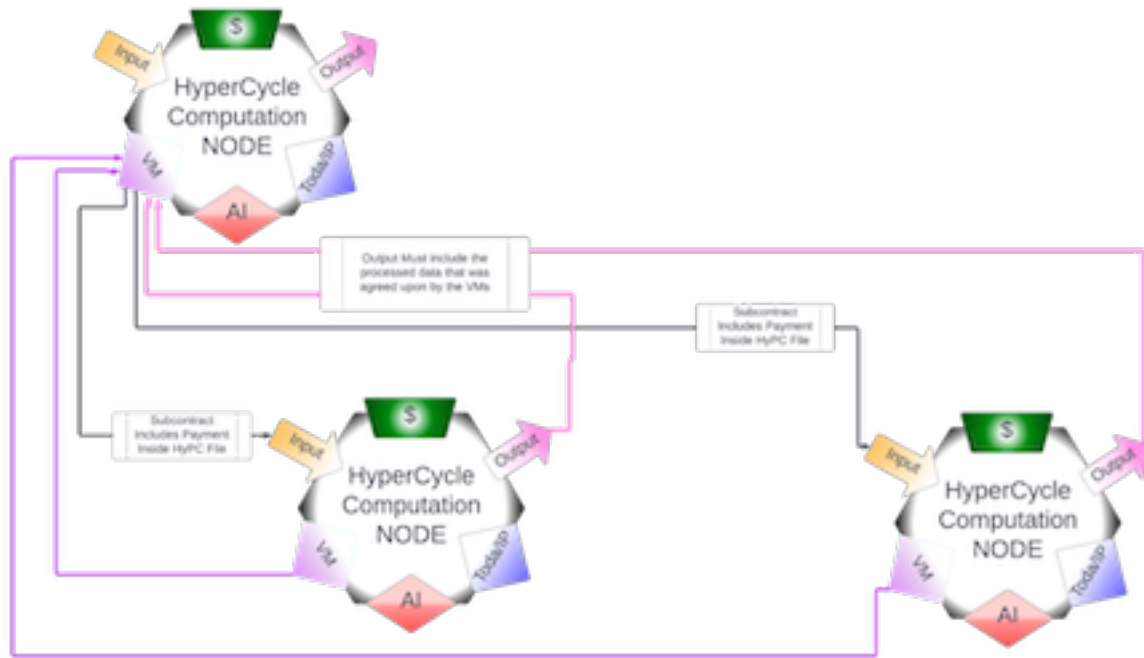


Figure 1: HyperCycle Computation Nodes Interacting

Below is an image showcasing a single HyperCycle Computation Node subcontracting 2 other nodes simultaneously and zero dependencies. Yet it is crucial to be reminded that no finality ever reached without having something global being impacted. We cover that in details in Toda/IP.

One important class of use-cases for which HyperCycle is intended is where:

- You have a large number of transactions that need to happen in a short period of time such as HPC centers (a single HPC center can demand over 26 million micro transactions per second to have finalities in milliseconds, highest security and near zero cost)
- You don't necessarily need a rapidly checkable long-term record of all aspects of all of them
- It may be OK if some correctness-checking occurs after the fact rather than in real-time ¹

¹In one of the TODA asset model, which is heavily leveraged within HyperCycle, you

While not the only class of situations HyperCycle can be applied to, it's the particular sort of use-case that has been most firmly in mind while developing the design, and it is envisioned as playing a large role in shaping the initial implementation.

The HyperCycle design fulfills these requirements (as well as those of other use-cases) via leveraging

- Solidity, Plutus smart contracts, Hydra and other decentralized and centralized chains
- TODA/IP's ability to coordinate a large number of secure blockchain transactions without the necessity of maintaining any large or small replicated ledgers. Effectively zero-dependencies
- The SingularityNET weighted liquid rank reputation system's capability to assess the contribution of an agent to a network and thus the reliability of that agent to assist with various network functions.

Although there is no transaction fees for the AI Compute transactions, but there are royalties split up automatically and deterministically to ensure all parties are properly incentivized. Royalties and exchange of value between nodes can enable any currency to be loaded inside those HyPC, including but limited to USD, EURO, WON, AGIX, ADA, ETH, USDC, Ava, Algo, etc.

- HyPC tokens, native to HyperCycle network
- AGIX tokens to pay for the AI functionality required to maintain an uncorrupted reputation system.
- ETH, BTC, ADA or others to pay for the long-term storage of partial snapshots of the HyperCycle network state.

can "perform" any transaction you like, any time you want, even ridiculous completely made up ones, but anyone looking at that transaction can completely verify its integrity, based only on the information in the transaction's proof of provenance. So no bad transaction can provide evidence that it is a good transaction, and every good transaction has standalone evidence proving its veracity to anyone who takes time to look. However, the transaction may still take time to reach finality globally, but in the meantime and within the cycle and its reach to other cycle, the finality is immediate, which is unprecedented at this time this document was last updated. The question then becomes, in a given application, whether to accept a partially complete transaction or insist on a fully complete one before proceeding.

Rewards for helping with consensus in HyperCycle network will be given in HyPC tokens.

A HyperCycle network (as loosely depicted in Figure 2) comprises a collection of "rings", each containing a set of HyperCycle agents, each of which contains and controls a record of its own transaction history, and each of which is able to enter into various sorts of relationships with other agents (e.g. sending or receiving transactions, or participating in validating transactions). Unlike in standard ledger based blockchains, and drawing key structures and methods from TODA and TODA/IP (such as the TODA transaction trie and the cycle trie that gave HyperCycle its name), in HyperCycle individual agents and their purposeful interactions are the core of the system – no large-scale replicated ledgers are needed, and where they are used they are best considered efficiency-oriented augmentations to the metadata that agents retain regarding their own histories and properties.

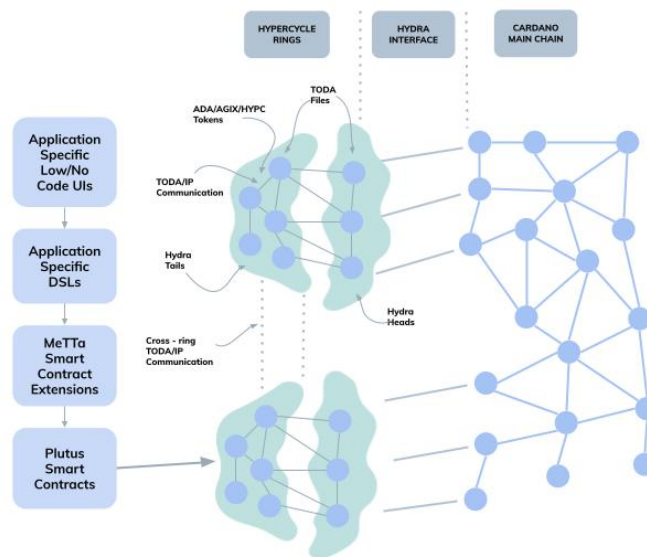


Figure 2: Conceptual illustration of two HyperCycle rings and their connection to the Any Ledger Blockchain mainchain.

Each ring in the network performs its own consensus judgments, and is connecting to any other nodes in the network or outside, such as ETH, BTC, ADA etc. Each ring contains a certain number of consistent validator nodes (agents) that are heavily involved in validating transactions in that

ring, and other nodes that can be opportunistically pulled into the validation process based on various criteria. The validators in each ring act like the participants in a multiparty state channel in a Hydra head. Participants in a HyperCycle ring who are non-validators or opportunistic validators will be like the participants in a Hydra "tail" , though the particular design of the HyperCycle "tail protocol" may be different than the generic Hydra tail protocol (which is not yet publicly well specified).

The SingularityNET reputation system, deployed among the nodes involved in HyperCycle rings, provides a dynamically updated reputation measure to each node. Proof of Reputation means that the choice of validators within HyperCycle rings will be biased toward higher-reputation network nodes (who then get fees in HyperCycle tokens for helping with consensus decisions)... Reliance on reputation in this way helps avoid various sorts of attacks in a relatively straightforward and elegant way. To seed the reputation system we could e.g. give initial reputation boosts to AGIX or ADA stakers, and also to early purchasers of HyperCycle tokens.

Different HyperCycle rings specialized for different purposes may use different consensus mechanisms, tuned e.g. for different balances of security guarantees versus speed and cost of processing. Purely ledgerless consensus may be used where looser security at the protocol level is permissible (with certain types of security dependent on any deposits made by participants on the any chain, for instance) and fast cheap processing is critical. Where maximal protocol level security is required then custom hierarchically sharded ledgers may be deployed at the ring or ring-set level, adding both cost and certain sorts of reliability.

Section ?? briefly reviews a few example applications ideally suited for HyperCycle – swarm AI, rating and reward in media networks, decentralized payments and computer processing, and public/private chain interoperability – and roughly indicates the sort of consensus mechanism and HyperCycle/any-mainchain interaction likely to be needed in each case.

The HyperCycle design also supports flexible public/private deployment. Given a fully implemented HyperCycle, SingularityNET applications (along with many other sorts of applications) will be deployable with subnetworks on both public and private HyperCycle rings.

Finally, all this underlying complexity can nevertheless be presented to application developers in a simple and usable manner, appropriate to the domain areas in which their applications are operating. Toward this end, integration of the OpenCog Hyperon [**Hyperon**] framework's MeTTa pro-

programming language [MeTTaSpec] into Plutus is also envisioned, as a strategy for enabling broad usability via the creation of MeTTa-based DSLs focused on smart contract programming in specific vertical areas. Such DSLs may straightforwardly be presented to the user as low or no code frameworks auto-generated from the underlying MeTTa/Plutus code, effectively leveraging the HyperCycle framework behind the scenes.

2 Background: Any Chain such as Cardano, Hydra, TODA/IP, Proof of Reputation

In this section we review key aspects of the existing blockchain networks and designs that are critical to the HyperCycle design. This review may seem like a bit of a wild ride, given that the ingredients being drawn together use somewhat different terminologies and core data structures and processes, and are largely coming from different directions conceptually. However, Section 3 will pull all the pieces together.

2.1 Any UTxO Transaction Model

2

HyperCycle involves a unique combination of mechanisms enabling efficient, lightweight decentralized coordination of software processes; however, key to the design is its not dependent nor piggybacking on the Cardano blockchain, as it already has an extremely solid foundation for all aspects of secure, decentralized software coordination. However, integrations can provide certain functionality that may very well be needed. HyperCycle can also leverage the UTxO model such as Cardano blockchain and also the Plutus smart contract language and numerous other related features, however substituting a different consensus mechanism and a different system for data management.

Cardano's Extended UTxO (EUTxO) model [EUTXO] extends the standard UTxO model underlying Bitcoin in a manner explicitly designed to support smart contracts operating according to functional programming principles.

²This section largely contains information drawn from [EUTXO] and [hydra]

Transactions in a UTxO ledger contain a set of inputs and outputs, where outputs lock an amount of cryptocurrency, such that only authorized inputs of subsequent transactions can connect and consume those funds. The set of outputs that are dangling (unconnected) pending validation, at any given point in time, are the unspent transaction outputs (UTxOs) . In addition to the locked currency, each output also comes with a validator predicate; and each input comes with a redeemer value. To determine whether a given input of the currently validated transaction is permitted to connect to a currently dangling and unspent output, the software determines whether the validator predicate the output applies to the redeemer.

The EUTxO model extends the UTxO model in two ways:

- Instead of restricting the validator/redeemer mechanism to dealing with signatures, addresses in the EUTxO model can contain arbitrary logic in the form of scripts. When a node validates a transaction, the transaction will look up the script provided by the output's address and will execute the script if the transaction is allowed to use the output as an input.
- Outputs in EUTxO can carry essentially arbitrary data in addition to an address and value. This allows scripts to carry, for instance, information regarding the state of long-running smart contracts.

The script associated with a transaction can access the data being carried by the output, the transaction being validated, and some additional pieces of data called redeemers, which the transaction provides for every input; based on this it can carry out complex logic to determine whether the given input is permissible or not. The fact that the validator can inspect the entire validated transaction enables validators to enforce that contract invariants are maintained across entire chains of transactions.

Key to the EUTxO model is that the success or failure of transaction validation depends only on the transaction itself and its inputs, and not on anything else on the blockchain. This works naturally with the "pure functional programming" nature of the Haskell language which underlies the Plutus smart contract system; in a pure functional language like Haskell, computation is broken down into functions whose outputs depend only on their inputs, without side-effect as are rampant in imperative language. The functional nature of Plutus is the central aspect making Cardano smart contracts tractably formally verifiable.

The purely functional nature of EUTxO validation also means the validity of a transaction can be checked off-chain, before the transaction is sent to the blockchain. A transaction can still fail if some other transaction concurrently consumes an input that the transaction is expecting; however, if all inputs are still present throughout the period of execution, the transaction is mathematically guaranteed to succeed.

The differences between this UTxO-based model and the account-based models used in Ethereum and other similarly structured blockchains necessitate significantly different design patterns on the DApp level. EUTxO makes it much more tractable and straightforward to create DApps with provable guarantees regarding their behavior; and EUTxO also provides a superior foundation for powerful extensions to the core Cardano model like HyperCycle. However, it does require a way of thinking about transactions and data representation that feels initially unfamiliar to developers who got their start with account-based blockchains. EUTxO smart contracts can be formally modeled using a type of state machine called a *constraint emitting machine* (CEM). Based on Mealy machines, each one consists of a set of states S_c , a set of inputs $i = I_c$, a predicate $final_c : S_c \rightarrow Bool$ identifying final states, and a step relation $s_? \rightarrow (s, tx)$, which takes a state s on an input i to a successor state $s?$ under the requirements that the constraints tx are satisfied. Cardano implements CEMs on a EUTxO ledger (the mainchain) by representing a sequence of CEM states as a sequence of transactions, each of which has a state-machine input i_c and a state-machine output o_c , where the latter is locked by a validator $?_c$, implementing the step relation (exceptions being the initial and final state, which have no state-machine input or output respectively).

2.2 Hydra: Cardano’s Interoperability Solution

3

Hydra [**hydra**] is the algorithmic and software framework enabling the Cardano mainchain to communicate efficiently and elegantly with other blockchains. It supports implementation of what would informally be called sidechains, and also of proxies to other fully autonomous and separate blockchains.

When Hydra is used to connect to another blockchain network, the ”head parties” are a distinguished set of high-performance and high-availability par-

³This section draws mainly on [**hydra**] and [**EnterHydra**]

ticipants in this other network. "Tail parties" are other participants in this other network, who may be unreliable in terms of being online at any given time or being computationally capable of carrying out operations critical for decentralized network functioning. The Hydra architecture can be divided accordingly into the head protocol, the tail protocol, and the cross-head-and-tail communication protocol. These components are supported by additional underlying protocols for routing, reconfiguration and virtualization.

The head protocol is the only portion of Hydra currently developed to a state of reasonable maturity. It allows the heads parties in a blockchain network connected to the Cardano mainchain to rapidly process large numbers of transactions with minimal storage requirements by way of a multiparty state channel. The tail protocol, which has not yet been publicly described in a detailed way, enables the network heads to provide scalability for large numbers of additional network participants who may use the system from low-power devices, such as mobile phones, and who may be offline for extended periods of time. The cross-head-and-tail communication protocol leverages virtualization to allow heads and tails to communicate without going through the medium of the Cardano mainchain.

Hydra's state channels are isomorphic to Cardano in the sense that they make use of the same transaction format and contract code as the underlying Cardano blockchain. This means smart contracts can be directly moved back and forth between Hydra state channels and the Cardano blockchain – so that Hydra state channels effectively yield parallel, off-chain siblings of transactions on the Cardano blockchain.

2.2.1 Head Protocol Transactions

In terms of transaction management, what is happening behind the scenes in using the Hydra head protocol to connect a network to Cardano mainchain is: The head parties involved cooperate to commit a set of UTxOs comprising the initial head state, which these head parties then evolve by handling smart contracts and transactions among themselves without mainchain interaction. Transaction validation, including script execution, proceeds according to the exact same rules as onchain, leveraging the exact same validation code. In case of disputes or in case some party wishes to terminate the offchain protocol, the head parties decommit the current state of the head back to the blockchain - which then necessarily results in an updated blockchain state that is consistent with the offchain protocol evolution on the initially commit-

ted UTxO set. The mainchain doesn't need to know the detailed transaction history that led to the committed state, it only needs to know the committed state. Further, the decommit process is designed such that, when the latest state in the head is very large, the head state can be decommitted via parallel decommitment of small chunks. UTxOs can also be added to and removed from a running head without closing it.

To get a Hydra head started, any party may take the role of an initiator and ask a set of parties to participate. Each party then establishes pairwise authenticated channels to all other parties in the head. This public-key material is used both for the authentication of head-related onchain transactions that are restricted to head members and for multisignature-based event confirmation in the head. The initiator then submits an initial transaction to the mainchain that contains the head parameters. This automatically initializes a state machine for the head instance that manages the transfer of UTxOs between mainchain and head. Each head member then attaches a commit transaction, which locks on the mainchain the UTxOs that the party wants to commit to the head.

From this point on, the head enters into an open state. The head members continue running the offchain head protocol, which evolves the initial UTxO set independently of the mainchain. In the case that some head members fail to post a commit transaction, the head can be aborted by going directly from initial to final.

Once a head is in the closed state, the underlying state machine grants parties a contestation period, during which each head party has one chance to contest the closure by providing the certificate for a newer head UTxO set. Contesting leads back to the state closed. After the contestation period has elapsed, the state machine may proceed to the final state.

What's happening on the mainchain while a Hydra head is evolving offchain is that the mainchain Hydra protocol:

1. Upon head initiation, locks the mainchain UTxOs committed to the head
2. Keeps these locks and waits while the head is active
3. Facilitates the settlement of the final head state back to the mainchain after the head is closed.

The end result is a dynamic that replaces the initial head UTxO set by the final head UTxO set on the mainchain in a manner that respects but does not persist the complete set of head transactions.

Some careful record-keeping is required here behind the scenes, in order to ensure that each head member posts exactly one commit transaction and that the open transaction faithfully collects all commit transactions. To enable this, a single non-fungible "participation token" is issued to each head member, connoting a capability and obligation to participate in the head protocol.. This token must flow through the commit transaction of the respective head member and to be valid the open transaction must collect the full set of participation tokens.

2.2.2 Hydra Tail Protocol

The Hydra head protocol has many powerful use-cases on its own. For instance, a natural strategy for porting the SingularityNET marketplace to Cardano is to map the multiparty escrow contract in the Ethereum-based SingularityNET implementation into a Hydra head. The SingularityNET agents participating in the multiparty token and API access transaction are then the head parties and the head instance remains active until the escrow is cleared and the multiparty transaction is done.

However, if one wants to associate a sizeable external network with Cardano using Hydra as the interface in a "sidechain" like design pattern, it's necessary to also handle the situation where there are many participants in the sidechain that are intermittently online or have unreliable or minimal computational power. This situation leads to a variety of security challenges, as there are various scams a network participant can carry out more effectively when not required to be consistently online to participate in transactions. The Hydra Tail protocol [**EnterHydra**] has two key mechanisms for addressing this:

- requiring participants to put collateral on the mainchain, which will be lost if fraud is attempted
- instantiating a Challenge-Response-Protocol on the mainchain, with which clients can dispute claims by any participant

As will be elaborated in Section 3 below, this is one aspect of the Hydra tail protocol that is improved in HyperCycle, via leveraging Proof of Reputation

and TODA/IP mechanisms. Given HyperCycle’s particular requirements and the underspecification of the Hydra tail protocol, the approach we’ve taken in the HyperCycle design is to create a custom version of the Hydra tail protocol that suits HyperCycle’s needs, incorporating mechanisms and ideas beyond those that have previous occurred in the Cardano sphere.

2.3 The TODA/IP ledgerless blockchain

TODA/IP [TODAWP] is a protocol for secure communication and coordination in decentralized networks, which relies on the same underlying data structures and encryption mechanisms as standard functioning cryptography that is running this planet including the innovations by blockchain platforms, but combines these in a radically different way to enable tremendously greater scalability.

The most obvious distinction between TODA/IP and other blockchain systems is that TODA/IP is fundamentally *ledgerless* – there is no essential dependence on any ledger be it centralized or decentralized, removing the dependencies introduce unprecedented advantages. Rather, each record managed by the network independently records its own transaction history while any movement of ownerships do impact a global hash that is carefully and deterministically constructed in a fully decentralized and distributed setting. One could then say that the ledger exists implicitly within the network of file transaction records, to be partially spidered by various processes on an as-needed basis. A crude metaphor to traditional blockchain architectures would be to say that in TODA/IP the ledger is ”sharded all the way down” (though of course the underlying mechanisms are not precisely that). However, if any ledger succeeds and getting sharded all the way down to the smallest possible set of data, they are likely to coincide with where Toda/IP has started and therefore be interoperable

TODA/IP significantly bypasses well known problems w/ replicated ledgers, such as

- Coordination overhead associated with distributing information from a large network to a number of replicated ledgers
- Communications overhead required for sending, receiving, and processing messages (e.g. according to gossip protocols which, while resilient, are known for their considerable communications overhead)

- Decreased performance with scale. Replicated ledgers require that every fully participating node must process every single transaction which occurs. This means that as more transactions occur, and more nodes participate, overall transaction times suffer severely. This type of protocol offers security, neutrality and censorship resistance at the cost of scalability.

Basically, in typical replicated ledger blockchain systems, every peer receives all transaction data, all hash reference values, and all block headers. Every node needs to be supplied with enough information to be able to re-create the entire chain. This is how blockchain explorers are able to provide browsing through the history of major blockchains.

Sharding palliates all these problematic factors, but only partially. To make sharding work really well requires hierarchical sharding architectures which give algorithmically more satisfactory performance but at cost of adding significant complexity and overhead. TODA/IP is in some ways similar in spirit to sharding, but keeps things simpler and lower-overhead by adopting a more fully decentralized approach, in which each local data-chunk is responsible for its own historical information (or in some cases that of its close neighbors). This changes the nature of the communication protocols a fair bit, but removes the main source of complexity and cost in ledger-based blockchains, which is the maintenance of a population of ledgers each storing broad-based information about the network as opposed to simply a population of inter-transacting agents each with their own local data and interests.

A network that is ledgerless in its foundational operation can still make use of ledgers when appropriate. For instance if one wants a rapidly searchable backup of a certain subset of network transactions, storing all these in a replicated ledger makes a lot of sense. HyperCycle makes use of the Cardano mainchain for this purpose, along with other purposes like help maintaining certain sorts of security guarantees. As will be elaborated in Section 3, HyperCycle runs most of its transactions independently of Cardano mainchain, leveraging TODA/IP and Proof of Reputation mechanisms along with Cardano's EUTxO model and smart contracts to achieve scalable secure decentralized ledgerless computing. But then a judiciously chosen subset of HyperCycle transactions are periodically snapshotted and pushed to Cardano mainchain where they will be stored in Cardano's replicated ledger for backup, and rapid search and access.

It is also possible to use TODA/IP together with hierarchical sharding, when one wants to achieve particular sorts of balance between security guarantees and performance. In this case the overhead of sharding is still there, but is significantly less than in cases where a sharded ledger is required to do all the work.

2.3.1 TODA/IP Transaction Processing: Decentralized and Localized

The core structural element of TODA/IP is its association of individual records with their own localized ledgers. This makes these records semi-autonomous agents in a sense that is not remotely the case for entries stored in the ledger of a traditional distributed/replicated ledger-based blockchain.

The core dynamical aspect of TODA/IP, correspondingly, is the way secure transaction processing occurs in a manner that is wholly decentralized and is also "localized" in the sense that it doesn't require access to any overall ledger of all transactions that have ever occurred in the blockchain, but requires only interaction of a small set of parties who are closer to the transaction at issue. In this dynamics, records are acting as "nodes" in a blockchain network – specifically, each record may correspond to a number of different nodes, each acting in a different portion of the tree.

Let us now briefly unpack how this works, in a simple case of a ring with a "plain vanilla" fully-replicated-ledger-less TODA/IP consensus system.

Given a record R owned by wallet A , that record can be sent to wallet B through the generation of a transaction request, which is signed by wallet A , then signed by wallet B , then distributed and signed by a set of validators. A cycle in TODA/IP consists of a round of transactions requests and ensuing validations (for those requested transactions that are valid).

But how are the validators chosen? In the simplest TODA/IP implementations, basically, any active device that is geographically dispersed may be chosen as a validator. A pseudorandom function is used to select devices in the network for participation in a computation validating a given transaction, based on their geographic dynamic disbursement. Devices cannot apply more computation than the modest amount required for the pseudorandom selection of work that is assigned to them in each block of time ? an amount of work that cannot be known in advance, but is fixed immutably the moment it is assigned.

Because there is no way to achieve extra privileges in the network via

extra work, there is no intrinsic need for “mining farms” or similar mechanisms. Also, while this approach enables economical incentives for devices to be on the network, it requires no disincentive for not participating in the computation. The amount of work involved here is independent of the size of the network; and the work is spread out through the system, rather than concentrated as it would be in typical replicated ledger-based blockchains.

Each validator provides four important functions:

- Confirming the validity of the transaction (structural soundness and proof correctness)
- Prevent sending a packet twice in this cycle
- Help build the consensus proofs for the transaction
- Provide matching proofs for A and B

One of the validators will also be chosen to create a new record containing the signed transaction validation message the validation process returns to the recipient. This new file will then be sent from the chosen validator to itself, resulting in the appending of this file to the validator’s record.

After the transaction is done, the Merkle root of the data structure constructed as the transaction proceeds is a cryptographically secure representation of the entire transaction.

It makes logical sense for a TODA/IP-based network to charge modest transaction fees to cover the network, electricity and depreciation cost of carrying out the work involved in transaction validation. These fees would be naturally split among sender, recipient and validators, in a manner that may be different for different TODA/IP implementations. There are obvious variations such as preferentially assigning the validator role (and the modest revenue obtainable therefrom) to nodes that have provided rapid validation in the past. This leads on to the integration of TODA/IP with reputation systems that we will discuss in Section 3 below and that forms a key aspect of the HyperCycle design.

The nodes in a TODA/IP network may be divided into “rings”, each of which groups a certain set of records, and which are hierarchically connected. A single record can then live in multiple concentric rings. Different rings can have different consensus mechanisms for approving transactions. Transactions made to the record in different rings can then be merged consistently within the record.

A TODA/IP system may hold a collection of rings, what we might call a "ring-set", can be configured to all share a common "TODATree" (fully saturated and balanced BST) structure, which provides a universal file address system that may roughly be thought of as a decentralized analogue to IPV6. In the context of a TODATree, each record has a unique address number that it retains over its lifetime, defined via the branch of the TODATree it occupies. A highest level TODATree of height 256 has been posited to supervene over all TODA/IP implementations, with each of the 2^{96} branches at level 160 associated with a particular TODA/IP implementation. HyperCycle will then involve a TODA/IP ring-set associated with a particular TODATree level 160 subtree.

2.3.2 TODA: Interoperable Proof Structures

While TODA/IP is a full ledgerless blockchain system, TODA simply describes a data structure.⁴ The core of the TODA design is the "TODA File", which essentially is a digital data file that comes along with a per-file ledger attached as metadata. The binding between a file's internal data and its ledger allows a file to behave like a "unique digital object" (a sort of NFT) rather than a more abstract, interchangeable symbolic token.

Every transaction that a file is involved with causes a corresponding record to be attached to the file's associated ledger. These transaction records contain among other information the addresses of the other parties involved in the various transactions. The functioning of TODA is founded on the way these transaction records enable each file to have a single canonical well formed proof of provenance.

The transaction detail describing a particular transaction for a particular file specifies several aspects such as:

- the destination address: who will become the new owner
- a metadata hash, which allows users to attach arbitrary metadata to the transaction
- an encumbrance, which is a special restriction on how the transaction will occur

⁴This section is based on [TODAPrimer]; later designs supersede this but are generally compatible with it.

- a signature on this data, to certify that the current owner really wishes the transaction to occur

The transaction detail becomes an intrinsic part of the proof of a file. By only requiring the hash of the metadata to be included in a transaction, the size of proofs is kept under control, and users are given the power to choose not to share the metadata of a transaction with someone else.

These qualities provide TODA files with the ability to move transparently between rings, regardless of the consensus mechanisms or networks involved, provided only that the capacity to create corresponding structures exists. This allows records that make use of these data structures to interoperate natively across unrelated systems.

2.3.3 Key TODA Data Structures and the TODA Cycle

Operation on a TODA file depend centrally on two Merkle trie data structures: The File Trie and the Cycle Trie. Explicating how these tries work is the best way to go to the next level of detail in the description of TODA operations.

The file trie is a Merkle trie which associates file IDs with transaction detail hashes for transactions involving that file. A transaction proof is the Merkle proof associating a single file detail with a file ID. The transaction details associated with a file's ID by its owner's file trie are the only valid representation of operations on that file.

The file trie for a given address in a given cycle contains all of the operations performed over files owned by that address in that cycle. Files owned by that address that are not present in that file trie for that cycle are said to have null proofs. Because the transaction details referenced in the transaction proofs associated with a file are the only valid representation of operations on a file, a file with a null proof is demonstrably not operated on within the context of the file trie.

Each file trie is contained within a cycle trie. The cycle trie can be built in a highly distributed way using e.g. TODA/IP, as we describe below, but it can also be built in a fully centralized setting, or something in between. The values contained in this cycle trie are nothing other than the file trie Merkle roots corresponding to the transactions occurring in the cycle. Each file only requires the small slice of the overall cycle trie that contains it.

As the cycle proceeds, the Merkle trie is built containing proofs for each transaction. For a given transaction its two proofs (one from the sender, one

from the receiver) will be known only to the sender and the receiver, but all are part of the cycle trie.

Not all nodes in a ring, in general, will be available to participate in a given cycle. However, cycles in any ring that a file belongs to are relevant to that file whether the node corresponding to that file is transacting in that cycle or not. In order to be able to continually prove that your file's ledgers are complete and not missing any information, you need to collect a bit of data every cycle that contains proof that your address did not contribute to the cycle. Luckily, these trimmed proofs tend to be especially short.

A protocol-level way to deal with these zero-transaction proofs is to say that the nodes assembling the cycle trie during a given cycle are responsible for maintaining proofs of their non-participating neighboring nodes (where a "neighbor" is a node that a given node has transacted with before). As many neighboring nodes who didn't contribute to a cycle will generally share the same proof, this is a relatively trivial set of data to maintain. There are also other routes to managing these proofs which are more at the service level, i.e. specific rings may maintain specialized ledgers or other mechanisms for managing these proofs that in some contexts may provide greater efficiency than relying on neighboring nodes. For instance, with sufficient trickery it may be possible to store "proofs of null proofs" together with cycle roots without actually storing whole cycle tries; there is some research yet to be done in this direction.

2.4 Proof of Reputation

5

The final key ingredient woven into the HyperCycle design is the "Proof of Reputation" consensus mechanism that was incubated within SingularityNET by a team led by Anton Kolonin and fleshed out fully in collaboration with Oladotun Aluko [[aluko2021proof](#)].

SingularityNET's "weighted liquid rank" reputation system

[[kolonin2018reputation](#)] is a general mechanism for assigning reputation to members of a decentralized network, integrating implicit and explicit ratings, and incorporating various subtleties such as adjusting the effect of member A's impact on member B's reputation based on member A's reputation. The length of a network member's history, the amount of currency they

⁵This section is based on [[aluko2021proof](#)]

have spent, the predictiveness of their own ratings, and many other factors can play into reputation calculations. The framework also includes a role for machine learning driven "reputation integrity analysis" agents to identify potential reputation fraud. Several simulations of the use of weighted liquid rank in various situations have been implemented and evaluated, e.g. reputation in ecommerce marketplaces [[kolonin2019reputation](#)].

The core idea of Proof of Reputation (PoR) is to use weighted liquid rank based reputations as the basis of a consensus mechanism for blockchain networks. The PoR framework uses the interaction of nodes in a network over time to determine the amount of reputation associated with each node in the network. A node's reputation is calculated by blending together a normalized set of ratings and the corresponding reputation values of the node providing the rating at a given point in time, rather than simply the value of the direct rating given by other nodes. A node's behavior also directly influences its overall reputation value by appropriate formulae. The PoR mechanism then uses the reputation to determine a set of consensus nodes responsible for maintaining the network's shared state. Reputation values are ongoingly updated as interactions between nodes in the network progress over time.

A few of the core principles underlying this scheme are:

- . The reputation value computed for a node is based on the reputation value of the node providing the rating – this is the "liquid" aspect
- The temporal scoping of reputation, according to which reputation values accumulated by members in the longer past contribute less to the current reputation value (while at the same time allowing agent history to contribute appropriately to reputation)
- The openness of all reputation values to all members of the community so that audits can be performed.

In accordance with principles of openness and decentralization, PoR stores node reputation values on a sidechain rather than in a centralized database. In a HyperCycle context, this ends up meaning that each HyperCycle network spawns a side HyperCycle network managing the reputation values for the parent HyperCycle network. This doesn't spawn an endless recursion because the nodes in the child network (containing reputation values) have the same owners as nodes in the parent network, so the consensus mechanism in the child mirrors that in the parent.

At the start of every PoR consensus round, consensus group members need to be selected and added into a consensus group. Members of the consensus group are chosen from the nodes with the highest reputation values, e.g. with collective reputation scores that exceed 50% of the total reputation values of the network. A leader from the group is then selected, who serves the functions of:

- Packaging all valid transactions from the list of pending transactions to a block
- Calculating the new reputation values for all network nodes using data from transactions in the transaction list
- Broadcasting the commit message to the consensus group

To verify a transaction, the consensus group waits until a certain minimum number of consensus members has sent a confirmatory message. This process constitutes a consensus group vote, in which each node has a weighted vote proportional to its reputation value from the previous round.

This general mechanism is consistent with a variety of different underlying blockchain architectures, e.g. it could be used with account-based, UTxO or EUTxO frameworks. HyperCycle layers PoR based consensus into a system based on EUTxO and TODA, resulting in a non-replicated-ledger-dependent blockchain framework supporting powerful functional-programming smart contracts but incorporating the particular security that PoR brings against various sorts of attacks.

3 HyperCycle

HyperCycle is a novel blockchain architecture which is formed by assembling key pieces from the existing blockchain networks and designs reviewed above: Cardano (EUTXO/Plutus/Hydra), TODA/IP, TODA, and PoR. The "assembly" process involved is not entirely straightforward, however, and requires some modifications and extensions to all of the ingredients. The result of this process is a blockchain with unprecedented capability for handling high-speed, large-scale on-chain agent interactions such as is required for on-chain deployment of population-based AI algorithms, tokenomics-powered interactive media, and numerous other applications.

The strongly modular design of HyperCycle and all its components also makes it relatively straightforward to extend various portions of the system, for instance the smart contract language. The MeTTa (Meta Type Talk) language [MeTTaSpec] under development in the context of the OpenCog Hyperon AGI project [Hyperon] has a number of properties that make it favorable for the development of application-specific smart-contract DSLs. Toward this end we envision making MeTTa accessible as a framework within Haskell (as it will be made available within e.g. python and Julia for AI scripting purposes), so that e.g. finance-focused, medicine-focused or evolutionary-learning-focused smart contract languages for on-chain execution could be scripted in MeTTa and then run via Plutus. It will then be possible to create generic tools for mapping MeTTa DSLs into appropriate low or no code development UIs, resulting in a HyperCycle framework with a high degree of usability for developers working in niches for which DSLs have been created.

3.1 Core HyperCycle Architecture Concepts

The HyperCycle network is a population of "agents" which correspond to TODA files and which are able to execute transactions in a p2p setting without depending on any third party. However, using third parties like Cardano as a service but not dependency can have certain features that folks may like. Using Cardano's EUTXO mechanisms (for instance they are able to run Plutus scripts). Each agent has an owner, and one owner may control multiple agents. The agents are organized into TODA-style rings, and a given TODA file may correspond to multiple agents living within different rings. As usual with TODA, each ring may have its own particular consensus mechanism.

There are no fees for HyperCycle transactions, however it can incorporate royalties for ever successful AI computation. Having said that the HyPC token has the capability to carry values inside it such as ETH, USD, ADA, AGIX and others.

Each HyperCycle ring is independent from any ledger based blockchain including the Cardano mainchain via Hydra. Each ring has its own policy determining which of its internal transactions gets pushed back in snapshots to Cardano mainchain, and how often these snapshots are pushed back.

Agent reputations and agent-owner reputations are calculated dynamically via a weighted liquid rank reputation system that operates at the overall

HyperCycle network level. Three dimensions of reputation are maintained for each agent: Reliability, cooperativeness and honesty. An agent with poor uptime or slow processing may be rated unreliable even if it has never done anything fraudulent or problematic. An agent with strong technical properties may still refuse to participate in constructing proofs, which will merit it a low cooperativeness rating, but this is still important to distinguish from a fraudster agent which will get a low honesty rating.

Use of reputation and social connectivity dynamics should in many cases enable dynamic evolution of HyperCycle rings, and in rings that use hierarchical sharding (to be discussed below), potentially of shards as well.

Each ring, in each cycle of its operation, has certain cycle head parties and certain cycle tail parties, in the rough sense of the Hydra protocol. Agents in the ring with high reliability, cooperativeness and honesty may be thought of as essentially the same as conventional Hydra head parties. Any member of the ring who wants to participate actively in consensus during the next cycle can broadcast a message saying so, and will then be considered as a cycle head for the next cycle.

As compared to how Hydra currently operates, for use with HyperCycle we will need Hydra to be a little more flexible in terms of how consensus is achieved. Currently Hydra requires all head parties to explicitly agree in order to confirm a transaction, in the manner of a standard multiparty state channel. In HyperCycle, instead, things will proceed more like in TODA/IP, but with reputation-based weighting. Validators for each transaction proposed within a cycle will be chosen via a pseudorandom function from among the agents in the ring who have proposed themselves as cycle head parties. Preference will be given to those proposer agents with stronger reputations.

This modification to Hydra proposed for use with HyperCycle has similar intention to the Hydra tail protocol, but is likely different in some particulars from the tail protocol as currently being pursued within the Cardano community. Collaborative development in this regard will likely be very productive.

3.2 HyperCycle Consensus Mechanisms

3.2.1 Lightweight Rings

The simplest sort of HyperCycle ring is what we call a "lightweight ring," which is wholly ledgerless in the manner of the "simple TODA/IP ring"

roughly described in Section 2.3 above. This sort of ring lacks certain sorts of strong security guarantees, as is discussed in Section 3.2.3 below. However this is compensated significantly via security at the service level. Agent-owners participating in a lightweight ring are required to post a certain amount of ADA and AGIX as deposits on Cardano mainchain. If dishonest or uncooperative behavior on the part of an agent is discovered, this can result in loss of all or part of the owner's deposit.

The "lightweight ring" consensus mechanism is not adequate for rings involving individual transactions of huge financial value, but is well suited for rings oriented toward micro-transactions. Examples would be:

- AI agents that are participating in population-based AI algorithms like genetic algorithms, genetic programming, ensemble machine learning, swarm AI
- Small-scale tokenomic transactions such as rating articles or comments on a tokenomics-enabled media site
- Small-scale financial micropayments for any purpose

3.2.2 Hierarchical Sharding Enhanced Consensus Mechanisms

For cases where a greater level of protocol-level security and reliability is required, rings with alternative consensus mechanisms are needed. For instance one may create a "hierarchical sharding enhanced ring" (HS enhanced ring), in which the basic localized transactions between agents are supplemented by a hierarchically sharded ledger which contains proofs corresponding to the ring.

This HyperCycle ring level HS ledger is distinct from the Cardano mainchain ledger, and most of the time most of the information it contains won't end up getting pushed to Cardano mainchain, because it will make more sense to send mainchain snapshots of actual transactions that happened rather than snapshots of , for instance, zero-transaction proofs for numerous agents in each cycle.

The HS ledger may be used in various different ways, yielding multiple variations of the HS consensus mechanism. One option is to use the HS ledger purely as a backup, in case the more fully decentralized peer-to-peer approach used in the lightweight consensus mechanism fails. This allows guarantees beyond what is possible with the lightweight consensus, but in

the case where the agents in the ring are mostly effective and cooperative, it only modestly increases operational cost and speed.

Another option is to use the HS ledger more aggressively. E.g. if zero-transaction proofs are stored in the HS ledger then there's an option to omit them from file metadata, thus keeping proofs smaller. However this has risk to significantly increase the time cost of doing proofs because of the need for agents to download relevant portions of the HS ledger every time they do proofs. The severity of this issue will depend on the habitual distributions of the inter-agent transactions in the network – in particular, on the degree to which the transactions among the parties in the ring fall into a hierarchical statistical pattern that can be reflected in the hierarchical sharding structure.

Broadly speaking, the use of the HS ledger provides greater security against various sorts of attacks, and thus an HS ring is going to be appropriate for agents carrying out larger, less micro-level/disposable transactions than those appropriate for a lightweight ring. However this added security will come at an added performance cost – which can be reduced via clever design and efficient implementation, but will often remain a significant factor.

One general principle we see in these aspects of the HS consensus mechanism is that the TODA design providing core elements to HyperCycle does not obviate the key tradeoffs underlying all blockchain design, such as the performance versus security tradeoff – but does provide a framework in which various different tradeoff choices appropriate for different applications can be elegantly inserted, via use of ring-specific consensus mechanisms associated with their own auxiliary data structures and dynamics building on the core HyperCycle mechanisms.

3.2.3 Further Notes on Security, Reliability and Performance Tradeoffs

In this section we dig a little more deeply into the technical tradeoffs involved in the choice between the lightweight and HS ring consensus mechanisms described above. The discussion may also shed a bit of a broader light on some of the tradeoffs involved in choosing a lightweight ring versus a heavier ring for a particular application.

Incentivizing Agent Participation in Consensus Firstly, one key question underlying the operation of a ring under the lightweight consensus mechanism is: What incentivizes agents to contribute the knowledge they possess,

which is necessary for other agents' in the proof construction process they need to undertake to construct the Cycle Trie? In some cases this is not a relevant question, e.g. if all the agents in the ring are strongly externally incentivized to contribute to the overall outcome of the process the ring is carrying out (e.g. AI agents collaborating to solve a problem of mutual interest to all the agent owners). But in other cases the rapid cooperation of agents in providing distributed data to feed proofs to form the cycle trie could be a potential critical bottleneck for lightweight consensus operation.

The basic answers to this question are: payment and reputation. The network design involves modest transaction fees and some of these can be directed to agents providing necessary proof data for cycle trie formation. Also, maintaining acceptable ongoing reputation will be necessary for agent-owners to continue participating in HyperCycle networks. While low-reputation agent-owners may always try to re-enter the network using a different identity, there will be machine learning driven "reputation police" mechanisms designed to spot this sort of behavior. Further, agents with higher reputation are more likely to be chosen to participate in consensus and thus are more likely to profit from their participation in the network ongoingly. One can also juice up the benefits of having high reputation even further by making transaction fees lower for high reputation agents, so that agents have a direct financial incentive for playing nice.

Another possible mechanism, to be introduced with great care and only in particularly appropriate sorts of applications, is to allow sufficiently high reputation agents to add transactions to the cycle trie even if they don't have all the necessary proof data at hand. This is not a route one would want to take in a network hosting large financial transactions, but in a ring comprising AI agents acting cooperatively to solve problems together on-chain, it may be perfectly acceptable.

Payment and reputation are "soft incentives" rather than hard protocol guarantees, so in the lightweight consensus scenario there is still a possibility that some agents will simply refuse to provide data that they hold locally, and eat the reputation loss. In this case, the Hydra conflict resolution mechanism will be invoked, and the agent-owners involved will lose all or part of the deposits they have made on the Cardano mainchain. But there's a possibility that some transactions will not be able to happen.

Costs and Benefits of HS Enhanced Consensus The HS consensus mechanism sketched above is, in a sense, a hybrid between the fully ledgerless TODA/IP system and a more traditional ledger-based blockchain with hierarchical sharding (see e.g. [chen2019thinkey] for a fleshed-out description of hierarchical sharding in the Thinkey blockchain, but many other examples exist). In this case we have information which, ideally, everybody in a ring should have (with consensus): the cycle roots of the TODA/IP system. In the HS approach we increase the odds of achieving this ideal by increasing this common information a bit beyond what's there in the lightweight consensus mechanism. All head parties in a ring still construct the data structures for the record. But along with building the cycle tree, the non-zero transaction packets are added to the hierarchically sharded ledger. This system has almost the same level of privacy as in the lightweight consensus protocol (from common information it will be impossible to infer who owns items). But it automatically solves the two requirements mentioned above, and dramatically reduces the proof sizes (since proofs for zero transactions will not be required in proofs of provenance).

The big downside of the HS approach is that, in order to guarantee its ability the needed proof for any file received, an agent needs to be able to download the relevant portion of the HS ledger. In the worst case it needs to download the whole ledger, which could become huge as the system grows. On the other hand, if the transactions in the network have a strong tendency to be localized, so that most transactions occur between agents in the same small shard, then this may be only a moderate-sized problem. Still, at least in unfavorable scenarios, the HS consensus mechanism loses the lightweight consensus mechanism's favorable property of accelerating rather than decelerating as the number of network participants increases.

Two Strong (But Optional) Requirements One way to think about the value added by the HS approach (or other similar mechanisms) is to look at the following two requirements, both of which are clearly desirable for maximal network security:

1. If a user adds a non zero transaction then this transaction can make its way to the final ?cycle root? (cycle root for which we have consensus) *if and only if* user has merkle proof for this transaction. It means that it is not possible by design to fail (because of network error or intentionally) to send the element of merkle proof to the user.

2. Users should always be able to reliably recover all cycle roots and all merkel proof for all zero transactions.

In a "low security margin" ring, such as one using the lightweight consensus mechanism outlined above, we can have nonzero probability of failure for these two properties. On the other hand, it is difficult to see how one would satisfy the second requirement in particular without doing something roughly "decentralized ledger like" – i.e. without having decentralized consensus about proofs of zero transactions, and making these available for everybody. The first requirement is in some way trickier but is also less complex to address in ledger-like scenarios.

However, it is not a given that these requirements must necessarily be met in all rings via guarantees at the protocol level. Maximizing security guarantees generally also maximizes cost (computational and otherwise), which limits the scope of potential use cases. Bitcoin is currently primarily concerned with the use-case of transacting large quantities of speculative assets, and for this purpose a decentralized ledger based approach is relatively apropos. But other sorts of use-case, e.g. in agent-based AI or social media or micropayments, have dramatically different sets of requirements which point toward different trade-off choices, and some of these choices may sensibly involve giving up some protocol-level security guarantees in pursuit of performance or other beneficial aspects.

Defusing the Power of Forking It's also worth noting that essentially every TODA/IP-based system will possess, by design, certain sorts of security that elude all modern decentralized ledger-based systems – for instance, security against the chaos and economic perversity habitually incurred by blockchain forks. All too often the fallback solution to a major security issue in a ledger-based blockchain is "we can always fork the chain." This has worked disturbingly well for various cryptocurrencies on a financial basis. For instance, the results of the Bitcoin forks leading to Bitcoin Cash and Bitcoin SV, and the Ethereum fork leading to Ethereum Classic, were broadly profitable: Both pre and post fork tokens retained significant financial value, so that the forks increased the market caps of the underlying blockchains. However, in other cases the ability to fork in this way is highly undesirable; for instance real-world assets like loyalty points, fiat, video game assets, etc. The fundamental uniqueness of TODA/IP records obviates the power of this sort of forking. In a HyperCycle network with HS consensus, forking the HS

ledger still doesn't fork the records underlying the HyperCycle agents. There can of course be software upgrades but there is no structural bias for these to lead to perverse forking situations as seem to routinely occur with today's standard distributed ledger-based blockchains.

File Uniqueness as a Key to Understanding Tradeoffs HyperCycle inherits from TODA the centrally important invariant that *each file must have a single canonical well formed proof of provenance*. The key questions, which we have been discussing in this section as regards lightweight versus HS consensus mechanisms are ones like: How quickly can that proof be obtained? How complex is obtaining it? How much of a security margin is there for my files? What recourse is there if something blows up (either through lack of availability or through widespread equivocation)?

The answer to each of those questions is ultimately "it depends", because, as with TODA, there is a goal of having HyperCycle workable in a variety of different use-cases with different practical requirements, while maintaining canonical uniqueness for the underlying files. This requires flexibility because the security margin you want for high value assets always costs a fair bit in terms of finance and/or performance, and continues to cost over time (nothing is free), whereas for efficient micro-transactions you have to accept significantly lower security margins to get the efficiency required. What we can ask for is not a single detailed blockchain design that addresses all possible requirements with the exact same set of algorithms, but rather a flexibly modularized blockchain design that can be elegantly customized to meet a broad set of requirements – for instance via diverse pluggable consensus mechanisms, as in the case of HyperCycle.

File Uniqueness Across Multiple Rings These tradeoff-related issues arise in particularly subtle ways when considering the dynamics of TODA/IP-based systems across multiple rings. The uniqueness of a file in a multi-ring network is guaranteed only if we have consensus about cycle roots, in an appropriately defined sense. In the simplest case, if there is agreement about all cycle roots across all rings in the network, we can pass files between rings unproblematically. However, this strong situation is not the only one of interest. For instance, if I have a file which was in another ring A for some period of time, then to appropriately carry out proofs regarding A I require cycle roots from ring A for this period of time, but not necessarily beyond

this period of time. The proof data from A over the relevant period of time is included in the file's proof of provenance, and serves to link the cycle roots of A into the greater network of rings.

3.3 Accelerated Smart Contract Execution

Handling of smart contract execution in HyperCycle is a separate topic that we will touch only lightly here. Rather than relying on generic consensus mechanisms, greater efficiency can be achieved by use of specialized mechanisms for validating smart contract output in a secure and decentralized way. In this manner we can achieve multiple orders of magnitude speedup over existing blockchains.

In the standard Ethereum-blockchain approach, each smart contract is run on every full node in the network, thus providing consensus on what the output of the contract should be according to the network's consensus algorithm. This is secure (or at least as secure as the network's consensus algorithm) but obvious extremely inefficient. If 2000 nodes need to approve to give consensus on a smart contract's output, then the smart contract is basically running at least 2000x slower than if it were just a program running on a single machine.

A few workarounds for this smart contract execution bottleneck have been proposed, most famously zero-knowledge rollups, which are part of the Ethereum roadmap, and were being pursued for Cardano by the now-abandoned Orbis project. However the simplest (not that simple) form of zk-rollup based system relies on a small number of centralized provers, and decentralized alternatives become increasingly complex relying on special marketplaces of nodes providing zk proofs.

The HyperCycle framework supports considerably simpler and lower-overhead methods of radically accelerating smart contract execution. Basically, pseudorandom choice and a few other tricks are used to enable each smart contract to be run only on a small set of nodes rather than the whole array. This kind of flexible mechanism becomes especially convenient in a system that has no globally shared or replicated ledger and correspondingly no universal list of nodes who would be called on to validate all smart contracts.

In the HyperCycle approach, when one node has a smart contract it wants to run, it uses a specially designed pseudorandom function to select a small number (say $N=3$) of other nodes to run the smart contract as well. In the

simplest case of a smart contract that just gives a single output and doesn't report interim results: These other N nodes each run the contract and report an encrypted version of the result of the execution ... if the results from the N nodes are all proved equivalent then the result is decrypted and delivered to the original requestor node, and payment is delivered to the N nodes.

In the case of a smart contract that produces interim results, and does potentially expensive processing in-between the results, then the process of comparison among the N nodes doing execution may be done at appropriate check-points during smart contract execution, as regards interim results. If one is working with a pre-existing smart contract interpreter (such as the Plutus interpreter) then enabling this variation will generally require some minor tweaks to interpreter.

This approach can work across various different VMs, and can handle smart contracts running on VM A which then spawn smart contracts running on VM B , and more complex scenarios.

For most smart contracts, whose output is drawn from a reasonably large space of possibilities, a choice like $N = 3$ will reduce the odds that all N executor nodes fallaciously return the same result very very close to zero.

Roughly speaking the result of this approach is that to run a smart contract in a provably almost-surely-correct way only takes say $3x$ (plus some not-that-huge overhead) as much compute effort as running comparable code on just one machine. Which is a very modest slowdown compared to running smart contracts on existing popular blockchain networks (orders of magnitude less).

In cases where aren't too many operating devices able to run the smart contract, or where the smart contract is being validated only within a relatively small ring of participants, and therefore guessing who's being selected becomes easier by attackers, additional obfuscation mechanisms may be valuable. To provide extra defense against collusion by the pseudorandomly chosen nodes, one can transform the functions (contracts) being sent to each node so that each node is executing a slightly different contract, and the nodes then have no tractable way to connect their assignments with those of other nodes. This additional complexity, however, is not necessary if one is dealing with small, quickly-executed smart contracts that can easily be run on any pseudorandomly selected machine in a large ring of HyperCycle nodes.

3.4 The MeTTa AI Language As A Plutus Framework

Cardano’s Plutus smart contract language provides an extremely powerful framework for developing software processes for on-chain execution; and part of this power is the flexibility that it provides to create smart contract languages with different syntax and semantic and implement and deploy these using Plutus as a base layer.

MeTTa (Meta Type Talk) [**MeTTaSpec**] is a highly flexible AI programming language, developed as part of the OpenCog Hyperon AGI framework [**Hyperon**]. It features a native meta-type system designed to ease creation of custom application and algorithm specific type systems, including systems centered on gradual, dependent and probabilistic types. It provides a natural means for software agents to describe each others’ properties to each other (and for this reason will likely be inserted as a replacement for Idris2 in the AI-DSL under development by Cardano and SingularityNET for structuring and guiding interaction of SingularityNET agents [**AI-DSL**]).

MeTTa will be deployed within the Hyperon project as a framework accessible via python, Julia and possibly other languages. Making it accessible within Haskell as well would have the effect of enabling the scripting of Plutus smart contracts in a combination of Haskell and MeTTa, which could have numerous advantages given the greater richness of MeTTa’s semantics – both advantages for smart contracts involving AI operations more concisely expressed using MeTTa’s abstract type constructs, and also for applications suited to leverage MeTTa’s facilities for creating DSLs.

3.4.1 Smart Contract DSLs for HyperCycle: A Path to Broad Usability

Smart contract coding can be a subtle and sensitive matter, and it’s preferable to have a framework in which most application developers don’t need to think much about the nitty-gritty of smart contracts and can focus on their application. In the Ethereum world, the closest thing that has emerged to usability is a ”cut and paste culture” in which most developers work by copying others’ Solidity code and tweaking it to suit their needs. However, this is not how things work in a software ecosystem powered by tools with reasonable composability. Haskell provides powerful composability but has the disadvantage of being relatively opaque to most developers without background in functional programming or related mathematics.

An alternate approach to achieving usable smart contract development is the creation of domain specific languages presenting developers with only those blockchain-related functions they really need in their domain of application. Cardano’s Marlowe DSL-for-decentralized-finance pioneered this approach [Marlowe], but also highlights some relevant caveats. While Marlowe is extremely elegant it seems not to provide out-of-the-box the most convenient approach for achieving scalability in modern DeFi applications.

So, for example, a Cardano DeFi application like SundaeSwap ended up directly leveraging Plutus (which is generic) rather than Marlowe (which is finance-specific but not tailored for the particular sort of financial operations at the heart of SundaeSwap) [SundaeSwap]. One thing this illustrates is the need for a framework enabling rapid updating and modification of DSLs within an extremely flexible base language. This course was not open for Marlowe when it was initially created as its launch preceded that of Plutus, but it is important to consider in a HyperCycle context.

MeTTa provides an elegant approach for the creation of application and algorithm specific DSLs, similar to the use of Idris for DSL creation [brady2013idris] [brady2020idris] but with greater flexibility. The simple and generic nature of the process of DSL creation in MeTTa opens the door for tools that auto-generate low-code or no-code UI frameworks for MeTTa-specified DSLs. This arguably is the most (and perhaps only) viable approach for making smart contract development adequately simple and usable for the average product developer without sacrificing security, robustness or capability.

Depending on the DSL, in many cases it may be possible to create low-code or no-code frameworks that implicitly conduct authoring in the DSL. The abstract nature of MeTTa lends itself well to the templating of this sort of process, e.g. to the creation of general-purpose tools for mapping sufficiently simple application-specific DSLs into corresponding low/no-code frameworks. In this approach the only people who need to deal with the technicalities of Plutus and MeTTa are those who are writing DSLs for new domains, or creating custom applications sufficiently recondite that they don’t correspond to any of the DSLs in the library.

4 A Few Promising Early Applications

We now give a few brief notes about a few key HyperCycle application areas that we are eager to explore once a sufficiently mature version of HyperCycle

is available. We note this is somewhat of a haphazard sampling of the possibility space, reflecting our current landscape of pursuits which may omit projects and verticals that will be at the top of our minds by the time HyperCycle is ready for scalable usage. Also, of course, much of the beauty of launching an extremely flexible and general purpose platform is that others can then utilize it in ways and domains one never conceived.

4.1 Swarm AI: Evolutionary Learning, Algorithmic Chemistry, Ensemble ML

One major use-case driving the design of HyperCycle is the desire to run population-based AI methods on-chain – for instance, run a genetic algorithm where each population member is an on-chain agent and crossover and mutation are on-chain transactions. ”Algorithmic chemistry” applications as explored in e.g. [Cogistry] [buliga2020artificial] present similar issues and opportunities.

There is also great potential in the area of decentralized ensemble machine learning – building model ensembles confronting a common data-analysis or reinforcement learning problem, in a setting where each model in the ensemble may be owned and managed by a different party. To carry this out in a transparent and fully decentralized way, one would like as much as possible of the process of passing around datasets, assessing and merging results, running tests and so forth to be done on-chain.

These are applications where for almost all problems a lightweight consensus approach will be preferable – maximal real-time security is not key. If an attack or failure causes some interim AI results to be lost, this is unfortunate but rarely tragic. DSLs here will serve as AI algorithm configuration languages enabling developers to quickly experiment with different approaches on different datasets in different applications.

4.2 Ratings and Rewards in Online Networks

Managing ratings and rewards in online networks is another case where lightweight consensus is desirable. In this case it is important that attacks or system failures not cause participants’ ratings or rewards to get permanently lost, but it’s OK if making participants whole after infrequent unfortunate incidents takes some time and depends on a slow-paced conflict resolution

mechanism. Having maximally strong real-time security guarantees is less important than having a system with very low cost of operation, and which is quick and effective under all normal circumstances. It is also easy to envision a low or no code DSL enabling developers to quickly specify the customization of the weighted liquid rank reputation framework and various related reward tokens in their applications.

4.3 Decentralizing Payments and Processing Power

Two somewhat subtler use-cases are decentralized payments processing, and the tokenized decentralized management of processing power done in the NuNet platform ⁶.

In each of these cases, one has a spectrum of transactions, some of which are small in magnitude (micropayments on an online media platform, utilization of small amounts of background processing on an individual's smartphone) and naturally match with a lightweight consensus protocol, others of which are more serious (larger payments, utilization of large chunks of processing such as supercomputer time or significant portions of a corporate computing network) and would merit the added protocol-level security guarantees of something like HS enhanced consensus.

These applications would appear to benefit from multiple rings with different consensus mechanisms, and the ability for some agents to participate in transactions in both rings. So for instance if I pay a few US cents equivalent to a blogger based on the time I spent reading their article, this goes through a lightweight consensus based ring; but if I pay a larger chunk of funds to an AI server farm to train a big ML model, this should go through an HS enhanced ring.

DSL-wise, the Marlowe framework would seem to have many positive lessons to teach here, and one envisions a Marlowe-like language embedded in Plutus via MeTTa leveraging unique HyperCycle transactional features based on TODA and reputation.

4.4 Adaptively Blending Public and Private Chains

TODAQ, one of the companies in the TODA ecosystem, has pioneered the use of centralized TODA rings for enterprise applications. A private TODA ring

⁶<http://nunet.io>

of this nature can interoperate naturally with public TODA rings running on decentralized consensus mechanisms like TODA/IP, due to the elastic nature of the TODA design in which the sovereignty of the individual data file is key.

A similar approach can be taken for enterprise deployments of HyperCycle. In fact there are multiple approaches to be pursued for different enterprise contexts.

One strategy is to develop a HyperCycle ring with outright centralized control and data management: Basically going beyond the HS enhanced consensus and just using a centralized database as a HyperCycle ledger, with only as much replication as is needed for practical efficiency (no need to use replication to assure decentralized control in this case).

Another strategy, opposite in some ways, is to recognize that within the bounds of a particular enterprise, it may be that lightweight consensus mechanisms can safely be used, because it's safe to assume that all the agents operating in the HyperCycle ring are going to be honest and cooperative with respect to each other. This approach can enable greater efficiency within the walled garden of an enterprise than can ever be possible in the "wild west" of a public blockchain.

The choice between these two strategies depends on the degree of centralization and internal trust in the particular enterprise in question. The HyperCycle infrastructure supports a variety of custom possibilities beyond the two specific extreme strategies indicated here. Broadly, the point to be made in this regard is that the flexibility to define rings with custom consensus mechanisms, and then transact agents and files across these rings, allows hybrid public/private blockchain deployments to be created to meet the needs of essentially any enterprise. After an initial batch of enterprise deployments has been successfully executed, as a side-effect set of template "enterprise consensus mechanisms" will have been created, which will handle the vast majority of enterprise situations.

5 Conclusion: Catalyzing the Future Of AI and Blockchain Ecosystem

While Blockchain technologies are still considered in their early days, and while it would be a mistake to assume that the technologies that have be-

come popular today represent mature solutions that compellingly solve the problems they address, it would also be a mistake not to utilize a scientific method in evolving systems everywhere and not just in academia. Science is cumulative and so is computer science. Some leading blockchain frameworks do embody algorithms, structures and processes that appear solid enough to persist into the decentralized networks of the future – EUTxO and Plutus being two examples we’ve reviewed here. However, there are also major aspects of current blockchains, such as the default use of distributed replicated ledgers, that appear to us to owe their current dominance largely to historical accident rather than fundamental superiority or appropriateness.

We have articulated here a novel blockchain design, HyperCycle, which we believe has the power and the flexible customizability to serve as the basis for the decentralized networks of the future, at least in the particular arena of large-scale systems of software agents providing microservices to themselves and external consumers. We refer to HyperCycle as in Layer Zero Plus Plus Layer 0++ mainly because it is driven by the motivation to service the entire ecosystem of microservices including AI as we know whenever AI service is executed it generates value and that value needs p2p to be exchanged without having to ever depend on a third party that itself may claim to be decentralized. The goal here is to present the highest security at the lowest possible network level to minimize friction which introduces unprecedented capabilities. We are putting core aspects of Latest proven cryptography including Bitcoin, Cardano, Ethereum , ZK proofs together with essential ideas, structures and processes from the TODA/IP blockchain and the Proof of Reputation and OpenCog Hyperon frameworks, we believe a radically different future for blockchain technology can be laid out with relatively modest (though still far from trivial) engineering effort.

Our most immediate use-cases for HyperCycle are the SingularityNET decentralized AI network, and the various spinoff projects emergent from the SingularityNET ecosystem (e.g. NuNet for decentralized processing power, Mindplex for decentralized media, Rejuve for decentralized medical data and research, SophiaVerse for decentralized metaverses, SingularityDAO for DeFi, etc.). However, HyperCycle’s scope extends far beyond these particular application and we foresee a very broad range of utilization once HyperCycle’s engineering and deployment are complete. Indeed we believe HyperCycle technology has the potential to play a key role in projecting blockchain beyond the niches it currently serves and into the dominant role in the global tech and financial ecosystem that various pundits have long foreseen.

mybibliography.bib