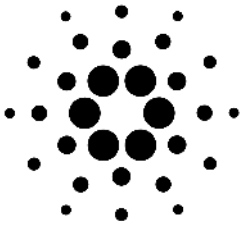

WHY WE ARE BUILDING CARDANO

A Subjective Approach

CHARLES HOSKINSON
<Charles.Hoskinson@iohk.io>
<C3A6 5E46 7B54 77DF 3C4C 9790 4D22 B3CA 5B32 FF66>

1. Introduction

Motivation

Cardano is a project that began in 2015 as an effort to change the way cryptocurrencies are designed and developed. The overall focus beyond a particular set of innovations is to provide a more balanced and sustainable ecosystem that better accounts for the needs of its users as well as other systems seeking integration.

In the spirit of many open source projects, Cardano did not begin with a comprehensive roadmap or even an authoritative white paper. Rather it embraced a collection of design principles, engineering best practices and avenues for exploration. These include the following:

- Separation of accounting and computation into different layers
- Implementation of core components in highly modular functional code
- Small groups of academics and developers competing with peer reviewed research
- Heavy use of interdisciplinary teams including early use of InfoSec experts
- Fast iteration between white papers, implementation and new research required to correct issues discovered during review
- Building in the ability to upgrade post-deployed systems without destroying the network
- Development of a decentralized funding mechanism for future work

- A long-term view on improving the design of cryptocurrencies so they can work on mobile devices with a reasonable and secure user experience
- Bringing stakeholders closer to the operations and maintenance of their cryptocurrency
- Acknowledging the need to account for multiple assets in the same ledger
- Abstracting transactions to include optional metadata in order to better conform to the needs of legacy systems
- Learning from the nearly 1,000 altcoins by embracing features that make sense
- Adopt a standards-driven process inspired by the Internet Engineering Task Force using a dedicated foundation to lock down the final protocol design
- Explore the social elements of commerce
- Find a healthy middle ground for regulators to interact with commerce without compromising some core principles inherited from Bitcoin

From this unstructured set of ideas, the principals working on Cardano began both to explore cryptocurrency literature and to build a toolset of abstractions. The output of this research is IOHK's extensive library of papers, numerous survey results such as this recent scripting language overview as well as an Ontology of Smart Contracts, and the Scorex project. Lessons yielded an appreciation for the cryptocurrency industry's unusual and at times counterproductive growth.

First, unlike successful protocols such as TCP/IP, there is little layering in the design of cryptocurrencies. There has been a desire to preserve a single notion of consensus around facts and events recorded in a single ledger, regardless of whether it makes sense.

For example, Ethereum has encumbered enormous complexity attempting to become a universal world computer, but suffers from trivial concerns potentially destroying the system's ability to operate as a store of value. Should everyone's program be a first class citizen regardless of its economic value, cost to maintain, or regulatory consequences?

Second, there is little appreciation for prior results in mainstream cryptographic research. For example, Bitshares' delegated Proof of Stake could have easily and reliably generated random numbers using coin tossing with guaranteed output delivery, which is a technique known since the 1980s (see the seminal paper by Rabin and Ben-Or).

Third, most altcoins (with a few notable exceptions such as Tezos) have not made any accommodation for future updates. The ability to successfully push a soft or hard fork is pivotal to the long-term success of any cryptocurrency.

As a corollary, enterprise users cannot commit millions of dollars worth of resources to protocols where the roadmap and actors behind them are ephemeral, petty or radicalized. There

needs to be an efficient process through which social consensus can form around a vision for evolving the underlying protocol. If this process is enormously burdensome, fragmentation could break the community apart.

Finally, money is ultimately a social phenomenon. In the effort to anonymize and disintermediate central actors, Bitcoin and its contemporaries have also discarded the need for stable identities, metadata and reputation in commercial transactions. Adding these data through centralized solutions removes the auditability, global availability and immutability — which is the entire point of using a blockchain.

Legacy financial systems such as those composed of SWIFT, FIX and ACH are rich in transactional metadata. It is not enough to know how much value moved between accounts, regulation often requires the attribution of actors involved, compliance information, reporting suspicious activity, and other records and actions. In some cases, the metadata is more important than the transaction.

Hence, it seems reasonable to infer that the manipulation of metadata could be as harmful as counterfeiting currency or rewriting transaction history. Making no accommodation for actors who want to voluntarily include these fields seems counterproductive to mainstream adoption and consumer protection.

Sojourn's End

The aggregation of our principled exploration of the cryptocurrency space is two collections of protocols. Respectively, a provably secure Proof-of-Stake [1][2] based cryptocurrency called the Cardano Settlement Layer (CSL) and a set of protocols called the Cardano Computation Layer (CCL).

Our design emphasis is to accommodate the social aspects of cryptocurrencies, build in layers by separating the accounting of value from complex computation and address the needs of regulators within the scope of several immutable principles[1]. Furthermore, where it is sensible, we attempt to vet proposed protocols through peer review and check code against formal specifications.

[1] See Regulation section for list

Proof of Stake

Using proof of stake for a cryptocurrency is a hotly debated design choice, however because it adds a mechanism to introduce secure voting, has more capacity to scale, and permits more exotic incentive schemes, we decided to embrace it.

Our proof of stake protocol is called Ouroboros and it has been designed by an extremely talented team of cryptographers from five academic institutions[2] led by Professor Aggelos Kiayias of the University of Edinburgh. The core innovation it brings beyond being proven secure using a rigorous cryptographic model is a modular and flexible design that allows for the composition of many protocols to enhance functionality.

This modularity allows for features such as delegation, sidechains, subscribable checkpoints, better data structures for light clients, different forms of random number generation and even different synchronization assumptions. As a network develops from having thousands to millions and even billions of users, the requirements of its consensus algorithm will also change. Thus, it is vital to have enough flexibility to accommodate these changes and thereby future-proof the heart of a cryptocurrency.

Social Elements of Money

Cryptocurrencies are a prime example of the social component of money. When restricting analysis solely to technology, there is little difference between Bitcoin and Litecoin and even less so between Ethereum and Ethereum Classic. Yet, both Litecoin and Ethereum Classic maintain large market capitalizations and robust, dynamic communities as well as their own social mandates.

It can be argued that a large part of the value of a cryptocurrency is derived from its community, the way it uses the currency, and its level of engagement in the currency's evolution. Furthering the thought, currencies such as Dash have even integrated systems directly into the protocol to engage their community in deciding what should be a priority to develop and fund.

[2] University of Connecticut, University of Athens, University of Edinburgh, Aarhus University, Tokyo Institute of Technology

The vast diversity of cryptocurrencies also provides evidence for their social elements. Disagreements about philosophy, monetary policy, or even just between the core developers lead to fragmentation and forks. Yet unlike their cryptocurrency counterparts, fiat currencies of superpowers tend to survive political shifts and local disagreements without a currency crisis or mass exodus.

Therefore, it seems that there are elements of legacy systems that are missing from the cryptocurrency industry. We argue — and have inculcated into the Cardano roadmap – that users of a protocol need incentives to understand the social contract behind their protocol and have the freedom to propose changes in a productive way. This freedom extends to every aspect of a value exchange system, from deciding how markets should be regulated to which projects should be funded. Yet it cannot be brokered through centralized actors nor require some special credential that could be co-opted by a well funded minority.

Cardano will implement a system of overlay protocols built on top of CSL to accommodate the needs of its users.

First, regardless of the success of a crowdsale to bootstrap development, funds will eventually dissipate. Hence, Cardano will include a decentralized trust[3] funded from monotonically decreasing inflation and transaction fees.

Any user should be eligible to request funds from the trust by a ballot system and the stakeholders of CSL vote on who becomes a beneficiary. The process creates a productive feedback loop seen in other cryptocurrencies with treasury/trust systems, such as Dash, by starting a conversation about who should and should not be funded.

Funding discussions force a relation of long and short term goals, the cryptocurrency's social contract, priorities and the belief in value creation with particular proposals. This conversation means that the community is constantly evaluating and debating its beliefs against possible roadmaps.

Second, our hope is that Cardano will eventually include a formal, blockchain based system to propose and vote on both soft and hard forks. Bitcoin with its block size debate, Ethereum with the DAO fork, and many other cryptocurrencies besides have endured long standing and, in frequent cases, unresolved arguments over the technical and moral direction of the codebase.

It can and should be argued that many of these disagreements, and the fracturing of the community that results when action is taken, are a direct result of a lack of formal processes for debating change.

[3] This is also known as a treasury system

Where does one go to convince Bitcoin users to adopt Segregated Witness? How should the core developers of Ethereum measure community sentiment for bailing out the DAO? If the community fractures, is the cryptocurrency damaged beyond repair?

In the worst cases, moral authority to act could simply devolve to whoever has the developers, infrastructural relationships and money, not the best wishes of the vast majority of the community. Furthermore, if a large portion of the community is inaccessible or disengaged due to bad incentives[4], then how can one truly know if their acts are legitimate?

Proposed cryptocurrencies such as Tezos provide an interesting model to examine where a cryptocurrency protocol is treated like a constitution containing three sections (Transaction, Consensus and Network) with a set of formal rules and process to update the constitution. Yet there remains much work to be done with incentives and over how exactly to model and change a cryptocurrency with a formal language.

The use of formal methods, machine understandable specifications and merging a treasury with this process for financial incentives are being explored as possible avenues for inspiration. Ultimately, just the ability to propose a protocol change in a transparent, censorship free way with blockchain based voting should improve the process, even if more elegant solutions cannot be designed.

Designing in Layers – Cardano Settlement Layer

When designing great protocols and languages, one should not look to the future, but rather to the past. History provides a litany of examples of great ideas that are perfect on paper, yet somehow have not survived, such as the Open Systems Interconnection standards. History also provides happy accidents that have endured from TCP/IP to JavaScript.

Some principles extracted from a historical view are the following:

1. You cannot predict the future so build in wiggle room
2. Complexity is nice on paper, but simplicity usually wins
3. Too many cooks spoil the broth
4. Once a standard is set it will probably stick around, regardless of whether it is suboptimal

[4] See rational ignorance

5. Bad ideas can actually evolve into pretty good ones if there is a will

Cardano is a financial system that accepts its social nature. There will be a tremendous need for flexibility and the ability to address arbitrary complexity in a particular user's transaction. If successful, there will be a need for tremendous computational, storage and network resources to accommodate millions of concurrent transactions.

Yet we do not have a digital, decentralized Robin Hood to take from the rich nodes and give to the poor ones in order to achieve a fair network. Nor do we have the luxury of trusting human beneficence to altruistically sacrifice for the greater good of the network. Therefore, Cardano's design borrows from TCP/IP the concept of separation of concerns.

Blockchains are ultimately databases ordering facts and events with guarantees about timestamps and immutability. In the context of money, they order ownership of assets. Adding complex computation by storing and executing programs is an orthogonal concept. Do we want to know how much value went from Alice to Bob, or do we want to get involved in figuring out the whole story behind the transaction and deciding how much to send?

It is incredibly tempting to choose the latter as Ethereum has done because it is more flexible, but it violates the design principles above. Figuring out the story means that a single protocol has to be able to understand arbitrary events, script arbitrary transactions, permit arbitration in cases of fraud and even potentially reverse transactions when new information is made available.

Then one has to make difficult design decisions about what metadata to store for each transaction. What elements of the story behind Alice and Bob's transaction are relevant? Are they relevant forever? When can we throw away some data? Does doing so violate the law in some countries?

Furthermore, some computation is private in nature. For example, when calculating the average salary of workers in an office, we would not necessarily want to leak how much each person makes. But what if every computation is publicly known? What if this publicity biases execution order to harm outcome?

Thus, we have chosen the position that the accounting of value should be separated from the story behind why the value was moved. In other words, separation of value from computation. This separation does not mean that Cardano will not support smart contracts. On the contrary, by making the separation explicit, it permits significantly more flexibility in the design, use, privacy and execution of smart contracts.

The value ledger is called the Cardano Settlement Layer (CSL). As the purpose is to account for value, the roadmap has the following goals:

1. Support two sets of scripting languages, one to move value and another to enhance overlay protocol support
2. Provide support for KMZ sidechains[5] to link to other ledgers
3. Support multiple types of signature including quantum resistant signatures for higher security
4. Support multiple user issued assets
5. Achieve true scalability, meaning as more users join, the capabilities of the system increase

Scripting

Starting with the scripting language, transactions between addresses in a ledger require some form of a script to execute and be proven valid. Ideally, one would not want Eve to access Alice's money, nor would one want a poorly designed script to accidently send value to a dead address making the funds irretrievable.

Systems such as Bitcoin provide an extremely inflexible and draconian scripting language that is difficult to program bespoke transactions in, and to read and understand. Yet the general programmability of languages such as Solidity introduce an extraordinary amount of complexity into the system and are useful to only a much smaller set of actors.

Therefore, we have chosen to design a new language called Simon[6] in honor of its creator Simon Thompson and the creator of the concepts that inspired it, Simon Peyton Jones. Simon is a domain-specific language that is based upon Composing contracts: an adventure in financial engineering.

The principal idea is that financial transactions are generally composed from a collection of foundational elements[7]. If one assembles a financial periodic table of elements, then one can provide support for an arbitrarily large set of compound transactions that will cover most, if not all, common transaction types without requiring general programmability.

[5] Coming soon in a paper from Kiayias, Zindros and Miller
[6] Specifics will be released in an upcoming specification. The full language will be supported in the Shelley CSL release planned for Q4 of 2017
[7] Project ACTUS has an in-depth elaboration

The primary advantage is that security and execution can be extremely well understood. Proofs can be written to show correctness of templates and exhaust the execution space of problematic transaction events, such as the creation of new money out of thin air or transaction malleability. Second, one can leave in extensions to add more elements by way of soft forks if new functionality is required.

That said, there will always be a need to connect CSL to overlay protocols, legacy financial systems, and special purpose servers. Thus we have developed Plutus as both a general purpose smart contract language and also a special purpose DSL for interoperability.

Plutus is a typed functional language based on concepts from Haskell, which can be used to write custom transaction scripts. For CSL, it will be used for complex transactions required to add support for other layers we need to connect, such as our sidechains scheme.

Sidechains

With respect to sidechains, Cardano will support a new protocol developed by Kiayias, Miller and Zindros (KMZ sidechains) based upon prior results from proofs of proofs of work. The particular design is beyond the scope of this paper; however, the concept allows for the secure and non-interactive movement of funds from CSL to any Cardano Computation Layer or other blockchain supporting the protocol.

KMZ sidechains are the key to encapsulating complexity. Ledgers with regulatory requirements, private operations, robust scripting languages and other special concerns are effectively black boxes to CSL, yet the CSL user will gain certain guarantees about accounting and the ability to recall funds once computation is complete.

Signatures

In order to securely move value from Alice to Bob, Alice needs to prove she has the right to move the funds. The most direct and reliable way of accomplishing this task is to use a public key signature scheme where funds are connected to a public key and Alice controls an associated private key.

There are hundreds of possible schemes with different security parameters and assumptions. Some rely upon mathematical problems connected to elliptic curves, whereas others are connected to exotic concepts using lattices.

The abstract goal is always the same. There exists a hard problem that cannot be solved unless someone has a secret piece of knowledge. The holder of this piece of knowledge is said to be the owner of the keypair and should be the only entity that has the ability to use it.

There are two groups of concerns a cryptocurrency faces with choosing a signature scheme. First, there is the long-term security durability of the scheme itself. Some cryptographic schemes used in the 1970s and 1980s such as DES have been broken. The period over which the scheme should be expected to survive must be decided upon.

Second, there are many enterprises, governments and other institutions that have preferred, or in some cases, mandated the use of a particular scheme. For example, the NSA maintains the Suite B protocol set. There are standards from ISO and even W3C workgroups on cryptography.

If a cryptocurrency chooses a single signature scheme, it is forced to accept that the scheme could be broken at some point in the future and at least one entity cannot use the cryptocurrency due to legal or industry restrictions. Yet a cryptocurrency cannot support every signature scheme as this would require every client to understand and validate each scheme.

For Cardano, we decided to start with using elliptic curve cryptography, the Ed25519 curve in particular. We also decided to enhance the existing libraries by adding support for HD wallets using Dr Dmitry Khovratovich and Jason Law's Specification[8].

This said, Cardano will support more signature schemes in the future. In particular, we are interested in integrating BLISS-B to add quantum computer resistant signatures to our system. We are also interested in adding SECP256k1 to enhance interoperability with legacy cryptocurrencies such as Bitcoin.

Cardano has been designed with special extensions that will allow us to add more signature schemes through a soft fork. They will be added as needed and during major updates planned in the roadmap[9].

User Issued Assets (UIAs)

Early in Bitcoin's history, protocols were quickly developed to allow users to issue assets that piggybacked on Bitcoin's accounting system in order to track multiple currencies concurrently.

[8] This is the documentation for Cardano's HD Wallet Implementation. We believe Cardano is the first cryptocurrency to support Ed25519 HD Wallets

[9] See cardanoroadmap.com

These protocols were not natively supported by the Bitcoin protocol, but implemented through clever hacks.

In the case of Bitcoin overlays such as Colored Coins and Mastercoin (now called Omni), light clients are forced to rely on trusted servers. Also transaction fees still have to be paid in bitcoins. These properties combined with the single pipeline for transaction approval make Bitcoin suboptimal for multi-asset accounting.

In the Ethereum case using the ERC20 standard, there is more feature richness. However, transaction fees still require ether. Furthermore, the Ethereum network is having difficulty scaling to the needs of all the issued ERC20 tokens.

The fundamental problem can be broken into three parts: resources, incentives and concern. With respect to resources, adding an entirely new currency to the same ledger means one has two independent UTXO (unspent transaction inputs) sets sharing the bandwidth, mempool and block space. Consensus nodes responsible for embedding transactions of these currencies need an incentive for doing so. And not every user of a cryptocurrency will or should care about a particular entity's currency.

Given these problems, the benefits are tremendous as the primary token of a multiasset ledger can effectively serve as a bridge currency allowing for decentralized market making. Special purpose assets could be issued to provide additional utility such as value stable assets like Tether or MakerDAO that are useful for lending and remittance applications.

Given the challenges, Cardano has adopted a pragmatic approach to multiasset accounting. Building in stages, the first challenge is designing the necessary infrastructure to support the demands of thousands of UIAs. Namely the following advancements are necessary:

1. Special purpose authenticated data structures to permit the tracking of a very large UTXO state
2. The ability to have a distributed mempool to hold a huge set of pending transactions
3. Blockchain partitioning and checkpoints to permit a huge global blockchain
4. An incentive scheme that rewards consensus nodes for including different sets of transactions
5. A subscription mechanic that allows users to decide which currencies they want to track
6. Strong security guarantees that UIAs enjoy similar security as the native asset
7. Support for decentralized market making to improve liquidity between UIA and the primary token

Our preliminary efforts for finding the right authenticated data structure have resulted in a new type of AVL+ Tree jointly developed by Leo Reyzin, IOHK and Waves. More research is required, but it is a foundational advancement that will be included in a later version of Cardano.

A distributed mempool could be implemented using Stanford University's RAMCloud protocol. Experiments will begin in Q3 of 2017 to study its integration into Cardano's consensus layer.

The remaining topics are interconnected and covered by ongoing research. We expect — subject to research results — to include a protocol into Cardano for UIAs during the Basho of CSL release in 2018.

Scalability

Distributed systems are composed of a set of computers (nodes) agreeing to run a protocol or suite of protocols to accomplish a common goal. This goal could be sharing a file as defined by the BitTorrent protocol or folding a protein using Folding@Home.

The most effective protocols gain resources as nodes join the network. A file hosted by BitTorrent, for example, can be downloaded much faster on average if many peers are concurrently downloading it. The speed increases because the peers provide resources while also consuming them. This characteristic is what one typically means when stating a distributed system scales.

The challenge with the design of all current cryptocurrencies is that they actually are not designed to be scalable. Blockchains, for example, are usually an append-only linked list of blocks. The security and availability of a blockchain protocol relies upon many nodes possessing a full copy of the blockchain data. Thus, a single byte of data must be replicated among N nodes. Additional nodes do not provide additional resources.

This result is the same for transaction processing and the gossiping of messages throughout the system. Adding more nodes to the consensus system does not provide additional transaction processing power. It just means more resources have to be spent to do the same job. More network relaying meaning more nodes have to pass the same messages to keep the whole network in synchronization with the most current block.

Given this topology, cryptocurrencies cannot scale to a global network on par with legacy financial systems. In contrast, legacy infrastructure is scalable and has orders of magnitude for more processing and storage power. Adding a specific point, Bitcoin is a very small network relative to its payment peers, yet struggles to manage its current load.

Our scalability goals for Cardano are greatly aided by our consensus algorithm. Ouroboros permits a decentralized way to elect a quorum of consensus nodes, which in turn can run more traditional protocols developed over the last 20 years to accommodate the needs of large infrastructure providers such as Google and Facebook[10].

For example, the election of a quorum for an epoch means we have a trusted set of nodes to maintain the ledger for a specific time period. It is trivial to elect multiple quorums concurrently and partition transactions to different quorums.

Similar techniques could be applied for network propagation and also sharding the blockchain itself into unique partitions. In our current roadmap, scaling methods will be applied to Ouroboros starting in 2018 and continue to be a focus in 2019 and 2020.

Cardano Computation Layer

As mentioned previously, there are two components of a transaction: the mechanism to send and record the flow of tokens and the reasons as well as conditions behind moving tokens. The latter can be arbitrarily complex and involve terabytes of data, multiple signatures and special events occurring. The latter can also be remarkably simple with a single signature pushing value to another address.

The challenge behind modeling the reasons and conditions of value flow is that they are immensely personal to the entities involved in the most unpredictable of ways. Lessons from contract law paint an even more problematic picture where the actors themselves might not even be aware that the transaction does not match commercial reality. We generally call this phenomenon “the semantic gap”[11].

Why should one build a cryptocurrency chasing an endless layer of complexity and abstraction? It seems Sisyphean in nature and naive in practice. Furthermore, each abstraction embraced has both legal and security consequences.

For example, there are numerous activities online that are universally deemed illegal or scorned such as the trafficking of child pornography or the selling of state secrets. By deploying robust

[10] There are also other independently research protocols attempting to achieve the same end such as Elastico and Bitcoin-NG

[11] Loi Luu et al. discuss this gap in their recent paper on Making Smart Contracts Smarter

decentralized infrastructure, one is now providing a channel for this activity to occur with the same censorship resistance that normal commercial transactions enjoy. It is legally unclear if the consensus nodes of the network — which have the incentive to become more federated over time to promote efficiency — would be held accountable for the content they host.

Prosecution of Tor operators, the brutal treatment of Silk Road's operator and the lack of overall legal clarity behind legal protections of protocol participants leaves an uncertain road. There is no lack of imagination of what else a sufficiently advanced cryptocurrency could enable (see the Ring of Gyges). Is it reasonable to force all users of a cryptocurrency to endorse or at least enable the worst acts and conduct of the web?

Unfortunately, there are no clear answers that provide insight to a cryptocurrency designer. It is more about picking a position and defending its merit. The advantage that both Cardano and Bitcoin have is that we have chosen to separate concerns to layers. With Bitcoin, there is Rootstock. With Cardano, there is the Cardano Computation Layer.

The kinds of complex behavior that would enable the acts elaborated previously cannot run on CSL. They require the ability to run programs written in a Turing complete language and some form of gas economics to meter computation. They also require consensus nodes willing to include the transactions in their blocks.

Thus, a functionality restriction could reasonably protect users. So far, most established governments have not taken the position that the use or maintenance of a cryptocurrency is an illegal act. Hence, the vast majority of users should be comfortable maintaining a ledger that is comparable in capability with a digital payment system.

When one wants to extend capability, there are two possibilities. It is enabled by a private collective of likeminded individuals and ephemeral in nature (for example, a poker game). Or, it is enabled by a ledger of comparable capabilities as Ethereum. In both cases, we have chosen outsourcing the events to another protocol.

In the case of a private, ephemeral event, it is reasonable to avoid the blockchain paradigm entirely, but rather restrict efforts towards a library of special purpose MPC protocols that can be invoked when desired by a group of likeminded participants. The computations and activities are coordinated in a private network and reference CSL only as a trusted bulletin board and a message passing channel when necessary.

The key insight in this case is that there is consent, encapsulation of liability and privacy. CSL is being used as a digital commons for users to meet and communicate — like a park would host a private event — but does not provide any special accommodations or facilitation. Furthermore,

the use of special purpose MPC will enable low latency interaction without the need for blockchain bloat. Thus, it improves the scale of the system.

Cardano's research efforts towards this library are centralized at our Tokyo Tech laboratory with some assistance from scientists abroad. We call the library "Tartaglia" after a fellow mathematician as well as contemporary of Cardano and expect the first iteration to be available in Q1 of 2018.

In the second case, one needs a blockchain with a virtual machine, a set of consensus nodes and a mechanism to enable communication between the two chains. We have begun the process of rigorously formalizing the Ethereum Virtual Machine using the K-framework[12] in partnership with a team from the University of Illinois.

The result of this analysis will inform the most optimal way to design a replicated and eventually distributed virtual machine[13] with clear operational semantics and strong guarantees of correct implementation from the specification. In other words, the VM actually does what the code tells it to do with the security risks minimized.

There are still unresolved questions about the gas economics proposed by Ethereum and how it relates to work such as Jan Hoffmann et al's resource aware ML and the broader study of resource estimation for computation. We are also curious about the level of language independence of the virtual machine. For example, the Ethereum project has expressed desire for transition from their current VM to Web Assembly.

The next effort is in developing a reasonable programming language to express stateful contracts that will be called as services by decentralized applications. For this task, we have chosen both the approach of supporting the legacy smart contract language Solidity for low assurance applications and developing a new language called Plutus for higher assurance applications requiring formal verification.

Like the solidity based Zeppelin project, IOHK will also develop a reference library of Plutus code for application developers to use in their projects. We will also develop a specialized set of tools for formal verification inspired by work from UCSD's Liquid Haskell project.

In terms of consensus, Ouroboros was designed in a sufficiently modular fashion to support smart contract evaluation. Hence, both CSL and CCL will share the same consensus algorithm.

[12] Invented by Professor Grigore Rosu et. al., K is a universal framework for language independent machine executable semantics. Prior to our work, it has been used to model C, Java and JavaScript
[13] Meaning that different consensus nodes run different smart contracts. Also known as state sharding

The difference is that Ouroboros can be confirmed to permit both permissioned and permissionless ledgers via token distribution.

With CSL, Ada has been distributed by a token generating event to purchasers throughout Asia who will eventually resell on a secondary market. This means that CSL's consensus algorithm is controlled by a diverse and increasingly more decentralized set of actors or their delegated assigns. With CCL, it is possible to create a special purpose token held by delegates of that ledger who could be regulated entities, thereby creating a permissioned ledger.

The flexibility of this approach allows for different instances of CCL to materialize with different rules about the evaluation of transactions. For example, gambling activities could be restricted unless KYC/AML data is present simply by blacklisting non-attributed transactions.

Our final design focus is on adding trusted hardware security modules (HSM) to our protocol stack. These are two enormous advantages when introducing these capabilities into the protocol. First, HSMs provide massive boosts in performance[14] without introducing security concerns beyond trusting the vendor. Second, through the use of Sealed Glass Proofs (SGP), HSMs can provide assurances that data can be verified and then destroyed without being copied or leaked to malicious outsiders.

Focusing on the second point, SGPs could have a revolutionary impact upon compliance. Ordinarily, when a consumer provides personally identifiable information (PII) to authenticate identity or prove the right to participate, this information is handed to a trusted third party with the hope it will not act maliciously. This activity is intrinsically centralized, the data provider loses control over their PII and is also subject to various regulations based on jurisdiction.

The ability to select a set of trusted attestors and then warehouse PII in a hardware enclave means that any actor with a sufficiently capable HSM will be able to verify facts about an actor in an unforgeable way without the verifier knowing the identity of the actor. For example, Bob is not an US citizen. Alice is an accredited investor. James is a US taxpayer and one should send taxable profits to account X.

Cardano's HSM strategy will be to attempt implemented specialized protocols over the next two years using Intel SGX and ARM Trustzone. Both modules are built into billions of consumer devices from laptops to cellphones and require no additional effort on the consumer side to use. Both are also heavily vetted, well designed and based upon years of iteration from some of the largest and best funded hardware security teams.

[14] See http://hackingdistributed.com/2016/12/22/scaling-bitcoin-with-secure-hardware/ from Cornell University

Regulation

The harsh reality of all modern financial systems is that as they scale, they accumulate a need, or at least a desire, for regulation. This outcome is generally the result of recurrent collapses due to the negligence of some actor or cabal of actors in a marketplace.

For example, the Knickerbocker Crisis of 1907 resulted in the creation of the Federal Reserve System in 1913 as a lender of last resort. Another example is the excesses of the 1920s in the United States that resulted in a terrible financial collapse, the Great Depression. This collapse yielded the creation of the Securities Exchange Commission in 1934 in order to prevent a similar event or at least hold bad actors accountable.

One can reasonably debate the need for, scope and efficacy of regulation, but one cannot deny its existence and the zeal with which major governments have enforced it. However, the challenge all regulators face as the world globalizes and cash becomes digital is two-pronged.

First, which set of regulations should be supreme when dealing with a collection of jurisdictions? The antiquated notion of Westphalian sovereignty melts when a single transaction can touch three dozen countries in under a minute. Should it simply be whomever wields the most geopolitical influence?

Second, improvements in privacy technology have created a digital arms race where it will become increasingly more difficult to even understand who has participated in a transaction, much less who owns a particular store of value. In a world where millions of dollars of assets can be controlled with nothing more than a secretly held 12-word mnemonic[15], how do you enforce effective regulation?

Like all financial systems, the Cardano protocol must have an opinion in its design over what is fair and reasonable. We have chosen to divide between individual rights and the rights of a marketplace.

Individuals should always have sole access to their funds without coercion or civil asset forfeiture. This right has to be enforced because not all governments can be trusted not to abuse their sovereign power for the personal gain of corrupt politicians, as seen in Venezuela and Zimbabwe. Cryptocurrencies have to be engineered to the lowest common denominator.

[15] See BIP39 https://github.com/bitcoin/bips/blob/master/bip-0039.mediawiki

Second, history should never be tampered with. Blockchains provide a promise of immutability. Introducing the power to roll back history or alter the official record introduces too much temptation to change the past in order to benefit a particular actor or actors.

Third, the flow of value should be unrestricted. Capital controls and other artificial walls diminish human rights. Outside of the futility of attempting to enforce them[16], in a global economy with many citizens in the least developed nations traveling outside of their jurisdiction to find a living wage, restricting capital flows usually ends up harming the poorest in the world.

These principles stated, markets are distinctly different from individuals. While the designers of Cardano believe in individual rights, we also believe that markets have the right to openly state their terms and conditions, and if an individual agrees to do business within this market, then they must be held to those standards for the sake of integrity of the entire system.

The challenge has always been cost and practicality of enforcement. Small, multijurisdictional transactions are simply too expensive in legacy systems to provide high assurance of recourse in the event of fraud or a commercial dispute. When one sends their wire transfer to the Nigerian Prince[17], it is usually too expensive to try to get one's funds back.

For Cardano, we feel we can innovate on three levels. First, through the use of smart contracts the terms and conditions of commercial relationships can be better controlled. If all assets are digital and can be solely expressed on CSL, strong guarantees of fraud-free commerce can be gained.

Second, the use of HSMs to provide an identity space where PII is not leaked but yet used to authenticate and credential actors should provide a global reputation system and allow for much lower cost regulated activities to be conducted, such as online gaming with automated tax compliance or decentralized exchanges.

Finally, in Cardano's roadmap is the creation of a modular regulation DAO that can be customized to interact with user written smart contracts in order to add mutability, consumer protection and arbitration. The scope of this project will be outlined in a later paper.

[16] As an example of a countermeasure to capital flow, see the Hawala Banking System
[17] See Advance-fee Scam

What is the Point of All of It?

Cardano has been a marathon project involving feedback from hundreds of the brightest minds inside and outside of the cryptocurrency industry. It involves tireless iteration, the active use of peer review, and shameless theft of great ideas when uncovered.

The remaining sections each cover a particular aspect of focus we have decided is a core component of our project. Some were selected due to a desire to improve the overall best practices of the space whereas others are specific to Cardano's evolution.

While no project can cover every goal or satisfy every user, our hope is to provide a vision for what a self-evolving financial stack should look like for jurisdictions that lack them. The ultimate reality of cryptocurrencies is not that they will disrupt the existing legacy financial systems. Legacy financial systems are always capable of absorbing change and maintaining their form and function.

Rather one ought to look to places where it is simply too expensive to deploy the existing banking system, where many live on less than a few dollars a day, have no stable identity and credit is impossible to find.

In these places, the power to bundle a payment system, property rights, identity, credit and risk protection into a single application running on a cell phone is not just useful, it is life changing. The reason we are building Cardano is that we feel we have a legitimate shot at delivering — or at least advancing — this vision for the developing world.

Even in failure, if we can change the way cryptocurrencies are designed, evolved and funded, then there is a great accomplishment.

2. Science and Engineering

The Art of Iteration

Cryptocurrencies are protocols implemented as software. Protocols are simply intelligent conversations between participants. Software is ultimately the manipulation of data given some

goal. Yet the difference between solid, reliable software as well as useful, secure protocols and their converse is completely human.

Good software needs accountability, clear business requirements, repeatable processes, thorough testing and tireless iteration. Good software also needs reasonably talented developers with enough domain specific knowledge to properly design a system that can fully resolve whatever problem they are trying to solve.

As for useful and secure protocols, especially ones involving cryptography and distributed systems, they start in a more academic and standards driven process. Peer review, endless debates and a firm concept of trade offs are necessary to ensure a protocol is useful. Yet these alone are not sufficient, protocols need to be implemented and tested by real life use.

The unique challenge in the cryptocurrency industry is that two completely different philosophies are mangled together without a proper Hegelian synthesis. Our thesis is a “move fast and break things” startup mentality driven by youth, greed and passion. The antithesis is a slow, methodical and academically oriented approach motivated by a desire to solidify the innovations of our space into a nice niche enjoying ample funding and prestige.

The result is that many cryptocurrencies are either entirely specified on a white paper only relevant to a CV or just by hastily written code. None of the current top ten[18] cryptocurrencies by market capitalization are based upon a peer reviewed protocol. None of the current ten top cryptocurrencies were implemented from a formal specification[19].

Yet billions of dollars of value are at stake. Once deployed, a cryptocurrency is exceedingly difficult to change. How does a user know they are using a secure system? How does a user know that the marketing claims are legitimate? What if the proposed protocol can never achieve the claims?

This lack of synthesis and respect for process is one of the primary reasons IOHK wanted to build Cardano. Our hope was to develop a reference project that would serve as an example of how to do things in a more effective, sane and honest way.

The goal is not to propose a totally new way of developing software and protocols, but rather to acknowledge that great software and protocols already exist and we can mimic the conditions that led to their creation. Second, to make these conditions publicly known and open source if possible so that they can be imitated for the benefit of the entire field.

[18] See www.coinmarketcap.com for a comprehensive listing by market capitalization
[19] Ethereum has a semi-formal specification known as the Yellow Paper; however, the EVM semantics are not fully specified nor are sufficient for a full implementation of the protocol.

Facts and Opinions

The other concern is over where facts end and opinion begins. There are hundreds of programming languages, dozens of development paradigms and more than one philosophy on project management. The academic world is riddled with its own challenges stemming from its distance from business concerns and practicality.

For Cardano, we first attempted to capture obvious deficiencies that can be universally agreed to be useful from an engineering perspective. For example, cryptography and distributed systems are both extraordinarily involved topics with far too many examples of how naive hands can make horrific mistakes. Therefore, any protocol requiring insight from these domains needs to be designed by an acknowledged expert and be submitted for review by other experts.

Ouroboros is our first case study of this area. It was designed by a team of cryptographers with a large, diverse and publicly verifiable publication history. It was built according to the standard cryptography process, with security assumptions, an adversarial model and proofs. These proofs were checked by submission to conferences[20] and also independently by computer proofs written in Isabelle by a team at the University of Cambridge[21].

Yet this work alone provides no guarantees of usefulness — just a rigorous check of a security model given some assumptions. For usefulness, one needs to implement and test the protocol. Our developers have done so in both Haskell and also Rust. This work revealed that more effort needed to be focused on the synchronization model, which led to the creation of Ouroboros Praos.

This art of iteration is what produces great protocols, with each step leading to new lessons and a requirement to re-verify the correctness of prior step[22]. It is costly, time consuming, and at times truly tedious, yet it is required to ensure a protocol is correctly designed.

Protocols — especially ones to be used by billions of people — are not short lived and rapidly evolving. Rather they are intended to be followed for years to decades. It seems entirely reasonable that, prior to burdening the world with a new financial system we all have to live with for the next 100 years, we want to demand some tedium and rigor from its designers.

[20] Accepted Paper Number 71 of the IACR's Annual Crypto Conference in California
[21] By Kawin Worrasangasilpa under the supervision of Professor Lawrence Paulson
[22] Following a tangent for a sake of levity, one should watch Professor Halmos's discussion about how to write a math textbook

Functional Sins

Moving into more opinionated territory, the tools, languages and methodologies used in software development are more artifacts of religious providence than objective reality. Source code is like written prose. Everyone has an opinion of what is good — and what is being communicated is, at times, less important than how it is communicated.

We must commit the sin of choosing a side accepting that it will be wrong in at least one person's eyes. However, there is at least a large corpus of justification behind our choice.

The protocols making Cardano possible are being implemented in Haskell. The user interface has been encapsulated in a fork of Electron that we are calling Daedalus. We have chosen to use the web architectural model where possible, and for our database, we opted for a key-value paradigm using RocksDB.

From a component level, this abstraction means that maintenance is far simpler, better technology can be substituted later with little effort, and that our stack is partly tied to the development efforts of Github and Facebook.

Using a WebGUI allows us to leverage React and develop front end features using tools understood by hundreds of thousands of JavaScript developers. Using a web architecture means that components can be treated as services and the security model is sensible.

Choosing Haskell for protocol development was the most difficult choice. Even in the functional world, there are ample choices. On the more flexible and impure side, there are languages like Clojure, Scala and F#, which benefit from the enormous libraries of Java and the .Net ecosystems while preserving some of the best aspects of functional programming.

There are more academically oriented languages such as Agda and Idris that have a close connection to techniques that would allow for strong verification of correctness. Yet they lack reasonable libraries and have a subpar development experience.

For Cardano, the choice came down to Ocaml and Haskell. Ocaml is a wonderful language with a great community, good tooling, reasonable development experience and a great legacy in the formal verification space through Coq[23]. So why did we choose Haskell?

[23] Adding to this point, IOHK actually does have a project being implemented in Ocaml called Qeditas that we inherited from the pseudonymous Bill White

Why Haskell?

The protocols that compose Cardano are distributed, bundled with cryptography and require a high degree of fault tolerance. On the best days, there will still be Byzantine actors, malformed messages and faulty clients unintentionally causing some form of havok on the network.

First, we wanted a language that enjoys a strong type system where we could easily use tools such as Quickcheck and more elaborate techniques such as Refinement Types while having a reasonable expectation of fault tolerance. An Erlang style OTP model satisfies the latter whereas languages like Haskell and Ocaml satisfy the former.

With the introduction of Cloud Haskell, Haskell gained many of Erlang's advantages while not surrendering its own. Furthermore, Haskell's modularity and composability has allowed us to use a lighter weight bespoke library called Time Warp for Cardano.

Second, Haskell's libraries have evolved greatly over the last few years thanks to extensive work of commercial entities like Galois, FP Complete and Well-Typed. As a consequence, Haskell can be used to write production applications.[24]

Third, PureScript's rapid evolution has provided a much needed bridge to the JavaScript world akin to what Clojurescript has given Clojure. We expect PureScript will be especially important when it comes to getting Cardano to work in a browser and developing mobile wallets.

Fourth, with respect to dependency resolution, Haskell in the last several years has enjoyed a significant social and technological effort led by technologists like Michael Snoyman through a platform called stackage that is both easy to use and well supported by FP Complete.

Fifth, beyond adequate dependency resolution, we aim for our software builds to be reproducible. In other words, with the same configuration values and dependency versions it should produce exactly the same build artifacts. Through stackage, we have been using NixOps to achieve reproducibility with great success.

Finally, the talent pool of developers specializing in Haskell is reasonably large — compared to its peers — and quite well-trained with the right mix of academic and industry credentials. It also acts as a competency filter as it is uncommon to find experienced Haskell developers without detailed knowledge of computer science.

[24] Bryan O'Sullivan provides a nice talk about Haskell's industrial use here.

Formal Specification and Verification

A significant strength of developing a protocol using a provably correct security model is that it provides a guaranteed limit of adversarial power. One is given a contract that as long as the protocol is followed and the proofs are correct, the adversary cannot violate the security properties claimed.

Deeper reflection makes the prior assertion even more significant. Adversaries can be arbitrarily intelligent and capable. To say they are defeated solely through a mathematical model is extraordinary. And, of course, it is not entirely true.

Reality introduces factors and circumstances that prevent the utopia of pure security and correct behavior from existing. Implementations can be wrong. Hardware can introduce attack vectors previously unconsidered. The security model might be insufficient and not conform to real life use.

A judgement call is needed about how much specification, rigor and checking is demanded for a protocol. For example, endeavors like the SeL4 Microkernel project are a prime example of an all out assault on ambiguity requiring almost 200,000 lines of Isabelle code to verify less than 10,000 lines of C code. Yet an operating system kernel is critical infrastructure that could be a serious security vulnerability if not properly implemented.

Should all cryptographic software require the same Herculean effort? Or can one choose a less vigorous path that produces equivalent outcomes? Also does it matter if the protocol is perfectly implemented if the environment it runs in is notoriously vulnerable such as on Windows XP?

For Cardano, we have chosen the following compromise. First, due to the complex nature of the domains of cryptography and distributed computing, proofs tend to be very subtle, long, complicated and sometimes quite technical. This implies that human driven checking can be tedious and error-prone. Therefore, we believe that every significant proof presented in a white paper written to cover core infrastructure needs to be machine checked.

Second, to verify Haskell code so it correctly corresponds to our white papers, we can choose between two popular options: interfacing with SMT provers via LiquidHaskell and using Isabelle/HOL.

SMT (satisfiability modulo theories) solvers deal with the problem of finding functional parameters that satisfy an equation or inequation, or alternatively showing that such parameters do not exist. As discussed by De Moura and Bjørner, use cases of SMT are various, but the key point is that these techniques are both powerful and can dramatically reduce bugs and semantic errors.

Isabelle/HOL, on the other hand, is a more expressive and diverse tool which can be used to both specify and verify implementation. Isabelle is a generic theorem solver working with higher-order logic constructs, capable of representing sets and other mathematical objects to be used in proofs. Isabelle itself integrates with Z3 SMT prover to work with problems involving such constraints.

Both approaches provide value and therefore we have decided to embrace them both in stages. Human written proofs will be encoded in Isabelle to check their correctness thereby satisfying our machine checking requirement. And we intend on gradually adding Liquid Haskell to all production code in Cardano's implementation throughout 2017 and 2018.

As a final point, formal verification is only as good as the specification one is verifying from and the toolsets available. One of the primary reasons for choosing Haskell is that it provides the right balance of practicality and theory. Specification derived from white papers looks a lot like Haskell code, and connecting the two is considerably easier than doing so with an imperative language.

There is still enormous difficulty in capturing a proper specification and also updating the specification when changes such as upgrades, bug fixes and other concerns need to be made; however, this reality does not in any way diminish the overall value. If one is going to trouble of building a foundation upon provable security, then the implementation should be what was actually proposed on paper.

Transparency

A final question when discussing the science and engineering of developing a cryptocurrency is how to address transparency. Design decisions are not Boolean and ethereal, coming to developers in dreams and then suddenly becoming canon. They are derived from experience, debate and lessons learned from earlier mistakes.

The challenge is that a totally transparent development process could influence discussion to become more theatrical than evidence based. Egos, attempts to win over a community, and fear of sounding stupid could force conversations to become sterile and counterproductive.

Furthermore, outsiders could attempt to co-opt the conversation in an effort to force their particular tangent to become the only relevant topic. Everyone has a sacred cow.

So how does one balance the need for a transparent development process, which is owed to the community that has entrusted progress to a set of core developers, with the need for freedom of expression without fear?

With Cardano, we have decided to embrace a standards driven process with directed oversight. The community needs to know that the science and the code are well thought out, checked and actually solve the things that developers claim they do. To this end, peer review should completely satisfy the science component as it has been designed specifically for this purpose and has given us the modern world.

For code, this topic is a bit more opinionated. For Cardano, we have elected to entrust the Cardano Foundation to serve as a final auditor of IOHK's work. In particular, they are entrusted with the following duties:

1. Regular review of the source code contained in the Cardano Github to check for quality, test coverage, proper comments and completeness
2. Review of all Cardano documentation for correctness and usefulness
3. Verifying the claims that the protocols produced by the scientists are fully implemented

To accomplish this task, IOHK will submit regular and timely reports to the Foundation – and its assigns – to review. The Foundation in turn will release a development oversight report to the Cardano community on at least a quarterly basis.

This first effort is intended to start a broader conversation about how a decentralized project achieves accountability. Development oversight from a trusted third party is a powerful tool to ensure that developers are on track, but it is not sufficient to completely guarantee that the project will always deliver.

For this reason, after the treasury is integrated into CSL, the Foundation will encourage additional development teams to construct alternative clients based upon the formal specifications developed jointly with IOHK. Development diversity has been a great technique used by the Ethereum project to avoid a monoculture forming around a single set of ideas or developers.

With respect to specifications, there is a wealth of knowledge to be gained from the standards process followed by the WC3 and the IETF. Ultimately, each protocol Cardano integrates requires a specification that is independent of academic work or source code. Rather it needs to be in a suitable format such as an RFC.

One of the Cardano Foundation's core tenets is to act as standards body specifically for the Cardano protocols and to host conversations to update, add or change standards relevant to Cardano. If the internet (a product of standards) through IETF can reach consensus about what core protocols shall be used, then it is entirely reasonable to assume that a dedicated body could facilitate the same outcome.

As a closing note, it is interesting to explore moving these discussions to a decentralized entity hosted on a blockchain. This concept is called a decentralized autonomous organization (DAO) and preliminary work is underway in this area. IOHK will develop a reference DAO model for entities interfacing with Cardano to use if desired and it is the Cardano Foundation's prerogative to decide whether to embrace it under their standards mandate.

3. Interoperability

The Grand Myopia

Finance and the broader idea of commerce is ultimately a human endeavor. There exist elegant languages, extremely precise tools to capture intent, and endless mazes of techniques to achieve recourse in the event of bad outcomes as well as thousands of years of laws seeking equity in trade. In fact some of the earliest forms of writing were commercial contracts.

Yet the human element cannot be eschewed regardless of the disintermediation to logic, machines or governmental sentinels entrusted with terrible powers. Therein lies the grand myopia of cryptocurrencies. They are mostly divorced from human reality.

People make mistakes. People change their minds. People do not always fully understand the business relationships they are agreeing to enter. People get misled and defrauded. Circumstances change on an individual and state level that require unique solutions. Belaboring this point, most contracts contain force majeure clauses.

However, cryptocurrencies seek to toss out human understanding, compassion and judgement in exchange for an uncaring digital judge perfectly bound to a constitution without consideration to fairness or outcome. Given that humans have always tried and will continue to attempt to change rules to selfish ends, it is refreshing to actually have a system that cannot be corrupted.

But what happens when a user needs to blend these new systems with traditional financial systems? What happens when one needs to live in the human world? For example, property rights such as land registration live entirely in the physical world. Even tokenizing the land still requires some acknowledgement of the incumbent jurisdiction.

To provide another point, a bar of gold cannot move itself. The digital judge can command its movement, but cannot force it without humans to accommodate. Hence a digital ledger can drift from reality.

Thus a protocol designer needs to decide how much human reality should be permitted in his cryptocurrency. The more flexibility, the less fidelity to the absolute one should expect. The more consumer protection, the more mechanisms have to exist to provide rollbacks, refunds and editing of history.

This section and the next on regulation covers Cardano's pragmatic approach to the topic. In terms of interoperability, there are two broad groups to discuss. First, interoperability with legacy financial systems (the non-cryptocurrency world). Second, interoperability with other cryptocurrencies.

Legacy

Fintech is not composed of a single standard or even a common language. There is tremendous diversity in approaches, the entities responsible for settlement and clearing, business processes, and other domains involved in the accounting, transformation and movement of value.

It is unreasonable to suggest that, simply because one technology is superior, the rest of the ecosystem will somehow admit defeat and upgrade. For example, many people still use Windows XP 16 years after the initial release. This sad state of affairs is equivalent to someone using the original Macintosh released in 1984 in the year 2000.

Consumer behavior aside, businesses are generally even slower in their upgrade cycle. Many banks still use back ends written in Cobol. Once infrastructure is known to work and meets business requirements, there is usually little incentive to upgrade or refine software and protocols for a consumer's benefit outside of compliance or security concerns.

For Cardano, we first have to establish what would a legacy bridge even entail? What systems, standards, entities and protocols should we target to ensure there is a reasonable certainty of interoperability? Can these bridges be federated or decentralized? Or like exchanges will they become central points of failure for hackers, malicious owners or overzealous regulators?

There are three concerns that have to be addressed. First, the representation of information and belief in its accuracy. Second, representation of value and its associated ownership. Third, representation of entities and, a particular user's alongside the aggregate level of trust in such entities.

To be useful, information and value need to freely flow between the legacy financial world and Cardano. Then outcomes need to be established and recorded to build reputation and grounds for recourse. Yet such things are mostly scoped in nature to the actors involved. To encode them on a blockchain would make them global and permanent.

Furthermore, value cannot always freely flow in the legacy world. Embargos, sanctions, capital controls and judicial action could freeze assets. To be interoperable, one cannot create an always open escape valve for value to leak.

Finally, the brand and reputation of entities is one of the cornerstones of commercial relationships. Billions of dollars are spent yearly on marketing campaigns to establish, maintain and repair brands. If libelous, false or misleading claims are made about a person or entity, then they have the right to seek legal recourse. Yet blockchains attempt to permanently preserve history.

Like our choice of programming language, there is no ideal solution for Cardano to resolve these concerns in a ubiquitously correct way. Rather, we have to yield to supported opinion again.

With respect to the flow of information, this flow is known as a trusted data feed. It has a source and content. Sources have some notion of credibility and incentive to deceive or maintain honesty. Content can be arbitrarily encoded.

Given that we intend on supporting trusted hardware in our protocol stack, we have chosen to explore adding support for Professor Ari Juel et al.'s Town Crier Protocol. Assuming the existence of a credible set of data sources, Town Crier permits the secure scraping of web content for use in smart contracts and other applications.

A bootstrap list of sources will be provided by Emurgo, IOHK and the Cardano Foundation. Later this list will be replaced by a community curated list using mechanics derived from Cardano's treasury system. Our hope is that a reputation system can materialize around good data feeds, thereby creating a positive feedback loop to gradually improve reliability and fidelity.

The representation of value is a more complex topic. Unlike information — where once the veracity, timeliness and completeness are established, protocols can behave in a reliable and deterministic way — value is more delicate.

Once tokenized, value should behave like a unique object. Information can be copied and passed around, but a token representing ownership of something (say a vehicle title) cannot be cloned and traded on two different ledgers. This act would effectively destroy the integrity of the system.

The challenge in legacy interoperability when dealing with tokenized value is that trust assumptions, reliability and auditability change as tokens flow between ledgers. For example, if Bob owns some Bitcoin and then deposits them on an exchange, then Bob now has the exchange's representation of his Bitcoin on their ledger. In the case of MtGOX, their ledger did not conform to reality, causing the users to lose everything.

The problem is further complicated by the need for legacy systems to recognize tokens living in a cryptocurrency. As mentioned previously, businesses are historically resistant to upgrading their software and supporting new protocols. This situation makes it difficult to see a clear solution.

For Cardano, our best hope is to provide an option for users to attach a rich supply of metadata to their transactions and then wait for industry standards to emerge to hook into. Some progress has been made with the Interledger workgroup, efforts like R3Cev and international mandates to upgrade old financial protocols.

However, the larger challenge remains of quantifying and qualifying value sent from a legacy system to a cryptocurrency ledger. For example if Bob is a bank owner and issues a dollar backed token, then he can always build a bridge to send his tokens to a ledger like Cardano as a user issued asset.

While Cardano would track ownership precisely and provide all the features we have come to love such as timestamping and auditability, no cryptocurrency can make Bob an honest banker. He always has the option of running a fractional reserve bank by not backing all of his dollar

tokens with real dollars. This fraud cannot be detected by a cryptocurrency unless the dollar itself was a token accounted by a digital ledger[25].

Finally, the representation of entities online is a classical network problem dating back to early days of the internet. Universities, businesses, government departments and any arbitrary users need to establish their identity at some point.

To this end, pragmatic yet centralized solutions like the web's Public Key Infrastructure and ICANN's DNS system have been implemented. Given that we enjoy the modern web, these solutions are both scalable and practical. But they do not answer a more commercially oriented question of reliability, trustworthiness and other meta characteristics necessary for determining if one wants to do business with the entity.

Multi-sided marketplace hosts like eBay have constructed a business model on providing some of this metadata alongside a framework to complete transactions. Judgements about the quality of content, events and businesses are often deeply influenced solely by online ratings from trusted sources[26].

The part of this point relevant to Cardano is a question of centralization of reputation. One of our goals for Cardano is to provide a financial stack for the developing world. A key to this effort is the ability to establish trust with actors one has never met.

If a single entity or a consortium of entities control who is labeled good or bad, not an organic process derived from actual interactions in the community as a whole, then these entities could arbitrarily blacklist anyone for any perceived sin. This power is against our values as a project and defeats the broader point of using a cryptocurrency.

Fortunately, the same mechanisms used in voting for treasury ballots, adding sources to a list of trusted data feeds and forking a protocol can be reused to establish a reputation space. It is an open area of research and our hope is to provide an overlay protocol for a decentralized reputation web of trust in 2018-2019 after more foundational elements have been settled.

Cryptocurrency Interoperability

[25] For digital ledgers on the other hand, proof of reserve has been proposed as a clever way of keeping cryptocurrency only exchanges honest.
[26] These rates even impact the creation of content itself. See this interest story on how Rotten Tomatoes has impacted the movie industry.

Moving from the legacy world to distributed digital ledgers, interoperability becomes far simpler. Each ledger has a network protocol, standards of communication and security assumptions about its respective consensus algorithm. These in turn can be easily quantified.

Movement of information is established by connecting to the foreign network and translating its messages. Movement of value can be done through a relay system, atomic cross chain trading or through a clever sidechains scheme. As there is not a centralized operator, one representation of entities restricts more to a metadiscussion of trust in developers, miners or some other powerbroker.

For Cardano, we are integrating a new sidechain protocol developed by Kiayias, Miller and Zindros. It provides a non-interactive way of safely moving value between two chains that support the protocol. This mechanism will be the primary way value will flow between CSL and a CCL layer.

For other cryptocurrencies, federated bridges should form as Cardano grows in value and user base. To help accelerate this growth, Cardano SL supports a restricted version of Plutus for interoperability scripts. New transactions will be added in the Shelley and later releases of CSL specifically to address these needs.

The Maze of Daedalus

The points on interoperability come from a global perspective. Specialized protocols, new transaction types, systems to assess credibility and the flow of information cannot be scoped to just a single gatekeeper or user. Rather they must be readily available to anyone without censorship or tolls.

Yet what happens when Cardano does not support a protocol, transaction or application that a user cannot live without? Should we just be out of scope? The web faced a similar concern during the 1990s.

Ironically, the web provides two different solutions that can be replicated with cryptocurrencies. The introduction of JavaScript provided programmability to any website to add arbitrary features. The introduction of browser plugins and extensions added custom capabilities for users willing to install them. Both approaches gave us the modern web alongside all its security horrors.

Ethereum adopted the former approach by allowing users to embed subprotocols on the Ethereum blockchain as smart contracts. Cardano supports this feature through the CCL paradigm. But what about custom extensions?

An elucidating example would be a cryptocurrency trader. Imagine a decentralized marketplace, called DM, that supports a set of different cryptocurrencies. A trader wants to automate his strategies acting on DM.

In a fragmented ecosystem, the trader would have to install dozens of clients for each cryptocurrency and then write custom software to talk to each client in order to coordinate automated trades. If one client updates, then it could break the bespoke software. Furthermore, what if the trader wants to sell the software?

Inspired from the web model of extensions, if the interface to various cryptocurrencies can be pulled into a web stack, then the trader's task becomes dramatically easier. A universal interface can be established. Installation is one click. Distribution of software can be modeled after the Chrome web store.

For Cardano, we have decided to experiment with this paradigm by deploying our reference wallet's front end on Electron. It is an open source project maintained by Github that combines both Node and Chrome together. Cardano's build of Electron is called Daedalus.

The first generation of Daedalus[27] will act as an HD wallet with support for many of the expected accounting and security features that are industry standards, such as spending passwords and BIP39. In later generations Daedalus will develop into an application framework with a store, universal integration APIs and an SDK.

The key innovations are ease of development by allowing programmers to use JavaScript, HTML5 and CSS3 to build their applications and a unified bridge for cross application communication. Complex behavior such as cryptography, managing a distributed network and database mechanics can be abstracted away thereby letting the developer focus solely on user experience and their application's core logic.

As Daedalus is intended to be a universal framework, its roadmap and evolution is somewhat independent of Cardano's. During 2017 they are tightly coupled, but later Cardano will be just another application for a Daedalus user. We also intend on exploring extremely unique features such as a universal key management service running solely in Intel SGX.

[27] Which is already available at daedaluswallet.io

Ultimately, as protocol designers, we cannot support all needs. Our hope is that the flexibility that Daedalus will provide combined with stateful smart contracts running on CCL will satisfy those left out by our design decisions. We also hope that better standards can emerge to encourage all cryptocurrencies to enjoy better interoperability and security.

4. Regulation

The False Dichotomy

As mercurial and arcane as regulation can often be, one can metaphorically infer an elegant narrative loop of the corrupt and their prosecutors seeking justice. Regulations are the toolkit of the lawbringer. But like all tools, they might be crude, old or simply misused.

Cryptocurrencies have not changed the human condition or the narrative loop. There will always be scams, bad actors and terrible outcomes despite the best of intentions. While cryptocurrencies can remove human judgement, they cannot remove human behavior.

A cryptocurrency designer has to take a position on what toolkit he will offer the regulator to correct bad events. The unique challenge cryptocurrencies face is that they are a product of regulatory and monetary failure[28].

Culturally, many in cryptocurrencies consider government action to be corrupt, inept or ineffective. Therefore, they have little respect, patience or desire to endorse a special backdoor for a regulator or lawman to right wrongs. This act would be anathema to the entire purpose of cryptocurrencies.

On the other hand, counting exchange failures and historic events, more than 10 percent of Bitcoin has been lost or stolen since the protocol started on January 3rd, 2009. As of June 30th, 2017, the value lost or stolen comes to a little over $4 billion. And this figure does not account for Bitcoin and other tokens lost to scams and poorly formed ICOs.

[28] In fact Satoshi embedded in the Bitcoin Genesis Block the following headline taken from The Times: *The Times 03/Jan/2009 Chancellor on brink of second bailout for banks*

Then there is the issue of privacy. On a macro scale, value flows through specialized channels that are regulated, rich in metadata and actively monitored by law enforcement, governments and international regulators. It is a well understood game with leakage occurring only on the cash side of affairs, which has been gradually diminishing as the world moves to digital money.[29]

The paradigm if cryptocurrencies did not exist would seem to be a world that increasingly treats financial privacy like social media content. There is none and one cannot opt out. Hence we have a dilemma yielding an apparent dichotomy.

A cryptocurrency designer can surrender principles and yield to whatever demands their local jurisdiction places upon their code, thereby compromising the privacy and integrity of their users. Or he can adopt a more principled, but anarchistic, philosophy that divorces itself from current best practices and laws.

For Cardano, we feel this narrative is a false dichotomy brought on by a lack of imagination. The reality is that most users are not concerned about rules existing for markets. They are usually concerned about sudden changes in the rules to benefit one or more actors. They are worried about a lack of transparency over who gets special privileges.

We need to distinguish between individual and market rights. Given that cryptocurrencies have a global reach, rights needs to be as user oriented as possible.

Privacy should be reasonable and at the user's control, not a gatekeeper. The flow of value should be unrestricted. Value should not be subject to sudden forfeiture without consent.

From a market perspective, the marketplace needs to be transparent about the use of data, how funds will be handled within and everyone needs to play by the same set of rules. Furthermore, once the user has consented, then they cannot suddenly change their mind due to inconvenience. Counterparties need certainty as well.

But how exactly does one move from the abstract to an actual system? What should something practical and legal look like? We have broken our solution into three categories: metadata, authentication and compliance as well as marketplace DAOs.

[29] The reader should consider picking up a copy of David Wolman's The End of Money. It covers the international movement towards cash disappearing.

Metadata

The act of something can often be less interesting than the metadata surrounding it. For example, driving from Denver to Boulder is an act. Driving from Denver to Boulder in a Ferrari 488 at an average of 120 MPH is metadata. Certainly this infers a different experience than in a Toyota Prius at an average of 30 MPH.

Financial transactions are no different. The context surrounding them is extraordinarily important to economists, tax authorities, law enforcement, businesses and other entities. Sadly in our current fiat based system, most consumers never see how rich in metadata their transactions are or who they are shared with[30].

For Cardano, we acknowledge that users could need or are legally required to share transactional metadata with certain actors like tax authorities. But we believe this sharing has to be at the user's consent.

We also believe that blockchain systems have tremendous power to eliminate fraud, waste and abuse by providing auditability, timestamping and immutability. Thus some metadata should be posted to the Cardano blockchain.

The hard part is finding a correct balance that does not condemn our blockchain to substantial bloat. Given this concern, we have chosen a pragmatic approach.

First, Daedalus will support over the next 12 months a large array of features to label transactions and financial activity. These metadata can be exported and shared on demand with whoever the user deems necessary. Furthermore, the data can be operated on by three party applications for domain specific purposes (for example, tax accounting).

Second, we are exploring adding support for special addresses that can include hashes and encrypted fields. This structure would permit a user to post metadata on our blockchain without publicly revealing it. But if she wants to share the data, it would carry all the auditability, immutability and timestamp surety that a transaction enjoys.

[30] On a more macro scale, author Juan Zarate writes about how this data is used by the US Treasury Department in the war on terrorism in Treasury's War. It provides a comprehensive view into how the current structure of global financial markets can be used for geopolitical ends.

We have already deployed an address structure that contains an attribute field. It is currently being used to store an encrypted copy of HD wallet trees structure for fast wallet recovery (see HD Wallet documentation). Later versions will generalize this construction.

Authentication and Compliance

Closely connected to transactions are the topics of the right to make transactions and the ownership of funds. For example, while there might be sufficient funds to buy something (for example alcohol), there could be restrictions on its purchase (age requirements).

Ownership and origin of funds are typically providence of know your customer regulations. When a money service business like a bank or exchange opens an account for a new customer, it is usually required to collect basic facts about the customer and where he acquired his funds from.

The technological challenge is that in the process of submitting this legally required information, the user sending it has no guarantee how it will be used, stored and if it will ever be destroyed. Compliance information is commercially valuable. It could be stolen for identity theft or resold where regulations permit.

For Cardano, we want to innovate as much as possible. On the software side of protocols, there is little to provide a guarantee that the receiver of compliance information will behave within a scope of conduct. However, on the hardware side of protocols, using trusted hardware, one can leverage Intel SGX and other HSMs to enforce certain policies.

Thus we are exploring using Sealed Glass Proofs alongside a sharing policy to permit the safe transmission of compliance information to a verifier who in turn is forced to comply with the policies it was transmitted under. We believe that both uniform standards could emerge and also that this method will reduce risk to verifiers by preventing the loss of customer data from hackers.

As a corollary to this effort, the layered model we propose for Cardano separating value from computation also can benefit from this approach. If the computation layer is run by regulated entities (say exchanges or casinos), then they would need to conduct compliance checks and potentially enforce tax policy on users.

Using SGPs, the user can send funds alongside personally identifiable information without concern that it will leak into the broader internet or be preserved by the consensus nodes of the

computation layer. Furthermore, the computation layer would gain certainty that all users transacting are authenticated and legitimate.

This paradigm also allows for customer portability between regulated entities. Exchanges could transfer balances and accounts for customers instantly through these safe channels and also — where policies permit — share data with regulators.

We expect our first beta test of this technology to be conducted in mid-2018 with an aim towards Cardano integration in late-2018 to early 2019 pending research results. This timeline also assumes the ability to collaborate with ARM and Intel in order to get code signed to run on their hardware[31].

Marketplace DAOs

The two previous sections covered the generation and movement of information assuming the existence of some external system. To ensure legacy interoperability, these features will always be necessary, but they do not address blockchain based regulation.

Smart contracts enable a completely new kind of commercial system where relationships are deterministic, self-enforcing and free of ambiguity. They can in turn be used to create rules for marketplaces including arbitrarily complex structures such as arbitration, event driven refunds, and revelation of facts given special conditions.

We call these smart contract enforced structures Marketplace DAOs. They do not require special protocol support nor mutability to be embedded in the ledger. In fact, they can be totally constructed using a collection of interdependent smart contracts.

The architectural concept is to design a collection of commercial templates inspired from contract law and business best practices. These templates can be wired into a developer's smart contract to enforce specific standards upon the marketplace.

For example, say a developer wants to issue an ERC20 token on CCL to conduct a crowdsale. A Marketplace DAO could be established specifically for crowdsales and its terms and conditions parameterized or even enforced by volunteer or legal standards. Things such as refunds, reallocation of funds or freezing of payment could be inherited in the developer's ERC20 contract.

[31] See Intel SGX Commercial License Policy

This effort allows us to have a macro discussion about how a marketplace should be controlled in order to ensure consumer protection. Second, we can discuss how to model transactions in a way to automatically ensure legal protection and rights within specific jurisdictions, such as New Hampshire.

Working with the Cardano Foundation, IOHK and other entities, the Cardano project will create a reference library of Marketplace DAOs for smart contract developers to use. Our hope is that insurance and regulatory markets can form around these DAOs and that they will be self-evolving based upon outcomes.

5. Sustainability

An immersion into the cryptocurrency area yields many conceptual contradictions. Cryptocurrencies are designed to be difficult to change, but, like all technology, they need to change to address design flaws and advancements. Blockchains are intended to prevent centralization, yet require strong actors to lead changes or maintain the code.

Perhaps the most frustrating experience comes when there are clear deficiencies that most stakeholders agree need to be corrected, yet consensus cannot emerge on the path forwards.

Bitcoin's block size debate has now been an active issue for more than two years. Daily, transactions totalling over a billion dollars are pending because the network is at peak capacity.

If changing a simple parameter — even in the presence of temporary solutions — cannot be coordinated, then how can enterprises and governments feel comfortable investing billions of dollars into building infrastructure on top of these systems? For that matter, how can any business gamble on the strategic risk of integrating accountability-free protocols that cannot make rational design upgrades?

Looking back into history, the evolution of the internet has followed a similar pattern with even simple changes like the transition from IPv4 to IPv6 taking decades to realize. Yet there is a strong contrast between blockchain technology and the internet in that they follow a very different style of custodianship.

The internet was a military project that grew out of DARPA into academic circles with strong government backing and a well-defined set of initial custodians. The internet grew under non-commercial conditions without the machinations of corporate influence attempting to

monopolize the network. In fact, e-commerce violated the NSF AUP until it was repealed in 1992.

By the time businesses had the luxury of commercializing the internet, there was already a strong set of standards, principles and evangelistic adherents. This did not stop companies like AOL and Microsoft from trying to build wall gardens and creating proprietary technology like ActiveX. This foundation has not stopped next generation actors such as Google from pushing their own agendas given their enormous user bases and capitalizations.

With swarms of rent seeking[32] actors from traders to miners, cryptocurrencies are the ultimate commercially motivated ecosystems. Given this foundation, evolution of the custodianship of cryptocurrencies has resulted in optimization around self-interest.

For example, validationless mining is starting to occur more frequently as it improves a miner's profit margin, yet this completely disregards the entire purpose and utility of mining. Mining centralization has already occurred with just a handful of actors in control of the majority of Bitcoin's hash power.

Like the internet, cryptocurrencies require consensus to change. But when such rapid centralization of power to a handful of brokers occurs, what happens when change is not convenient to them?

Unlike the internet, the bootstrapping of most cryptocurrencies is not done through altruistically non-commercial or academic means. From inception, some group seeks to make gains and there are power brokers assigned to help ensure those gains.

Founding centralization is a reality that each cryptocurrency must face in its evolution. We cannot fully escape it, but should at least try to design around gradual decentralization.

For Cardano, we thought carefully about what factors promote centralization and what techniques could be applied to encourage our protocol to gradually become public infrastructure like the web.

We fully admit that total decentralization is both impossible and perhaps even counterproductive. Yet certain factors can be encouraged to produce a more balanced system.

First, while centralized custodianship of crowdsale funds allows for agile and rapid development of the protocol during the early days, eventually funding has to diversify and the speed of

[32] See link for more information on this term

development needs to retire to a more systematic and deliberate pace. Following this point, funding needs to avoid cultural, linguistic and geographic bias.

Second, as the community becomes more informed about the underlying nature of the cryptocurrency's technology, decisions about the roadmap cannot be centralized to a set of core developers or foundation. There needs to be a blockchain based method for proposing, vetting, and enacting changes to the protocol.

Third, the incentives behind maintaining the Cardano SL blockchain have to be directly aligned with the aggregate desires of all users. We cannot permit a cabal of specialized actors to emerge who are independent of the will of the greater community.

For the first principle, we have chosen to integrate a treasury system into Cardano. For the second, we will deploy a formal process to propose Cardano Improvement Proposals through a system coordinated by CSL itself. For the third, we believe Ouroboros provides an elegant solution.

More detail could be provided on the above topics, but they are extensive in their own right and beyond the scope of a survey paper. Mechanism design is one of the most intricate and interdependent academic fields with incomplete theory and no solid canonical model to stand on.

Rather our science driven approach described in section two serves us well here. IOHK's Veritas team is working in partnership with a group of researchers from Lancaster University under the direction of Professor Bingsheng Zhang to develop Cardano's reference treasury model. With the aim of integration in 2018, we expect a dedicated peer reviewed publication by the end of 2017.

For formal description and vetting of changes to a cryptocurrency protocol, this topic is the least understood as it requires both ontological notions as well as a mechanism to incentivize broad participation. Perhaps some form of representative democratic process could emerge or use of liquid feedback to provide more rational voting.

We expect research in this direction to consume most of IOHK's formal involvement in the development of Cardano[33]. As a starting point, we will deploy alongside the reference treasury model several mechanisms to capture consent. Further study is required for a definitive solution.

[33] IOHK is retained to build Cardano until the end of 2020

Finally, work to improve incentives for Ouroboros is being supervised by Professor Elias Koutsoupias of the University of Oxford. After the cryptographic foundations of Ouroboros are solidified alongside all required scalability work, a broader study of bonds, penalties and exotic incentives will be added to the reference protocol.

6. Conclusion

A cryptocurrency is more than the sum of its protocols, source code and utility. It is ultimately a social system that inspires, enables and connects people. Frustrated by the many half measures, failures and broken promises of past protocols, we set out to build something better.

This process is not simple nor have we ever believed it can finish. Social protocols continue indefinitely changing as people and society change. To be useful, we want to trap the power of evolution and port it into Cardano.

Evolution is not guided by a single hand or a grand design. It is a process of serendipity inspired by endless mistakes and problems. Cardano seeks to be the digital embodiment of this process — fit enough to be able to survive the markets of today and adaptive enough to evolve to meet the needs of the future.

The previous sections capture a brief view into how we have been approaching this goal. We have diligently tried to recognize cognitive biases, learn from history and follow a rigorous process. We have tried to balance the need for rapid development with formal methods that traditionally cannot move quickly.

It has been an extraordinary privilege to embark on this journey. In the past two years, we have already developed a provably secure proof-of-stake protocol, recruited a small army of Haskell developers and made Cardano's development the concern of many talented scientists.

As we move from the laboratory to a deployed system in the wild, there will be growing pains, but our hope is that Cardano's future could be summarized in a single anthropomorphized sentence. Cardano is a pragmatic dreamer that learns from its elders, is a good citizen in its community, and always finds a way to pay its bills.

We cannot know the future, but we are glad to be trying to make it a better one for everyone. Thanks for reading.