



Consensus-as-a-service: a brief report on the emergence of permissioned, distributed ledger systems

By Tim Swanson


Published: April 6, 2015



CENTRALIZED



DECENTRALIZED



DISTRIBUTED



Contents

Acknowledgements and disclosures..... 3

Glossary..... 4

Introduction 6

Report highlights..... 6

Section 1: It begins and ends with the miners..... 8

Section 2: Bypassing the Maginot Line 8

Section 3: Why is this important?..... 10

Section 4: What does any of this have to do with distributed ledgers?..... 12

Section 5: Visions of the future..... 13

Section 6: What are smart contracts and what are they good for? 14

Section 7: Use-Cases 16

Section 8: Building a neutral network of value interchange..... 18

Section 9: The characteristics of a distributed ledger 21

Section 10: What about a fusion of permissioned and permissionless? 26

List of companies 29

Comparison Matrix 38

Section 11: Conclusions 42

About the Author 45

Appendix A: other notable efforts 46

Appendix B: making sense of cents 48

Endnotes 58

Acknowledgements and disclosures

I would like to thank the following individuals for their time and constructive feedback during the formation of this report:

Anton Bolotinsky, Arthur Breitman, Richard Brown, Vitalik Buterin, Preston Byrne, Alexandre Callea, Miles Cowan, Raffael Danielli, Pinar Emirdag, Daniel Feichtinger, Andrew Geyl (Organofcorti), Stuart Hoegner, Dave Hudson, Jeremy Johnson, Petri Kajander, Pavel Kravchenko, Joey Krug, Jo Lang, Mikkel Larsen, Jonathan Levin, Alec Liu, Zaki Manian, Todd McDonald, Dan O'Prey, Taulant Ramabaja, Raja Ramachandran, Meher Roy, George Samman, Robert Sams, Peter Shiau, Ryan Straus, Srinivasan Sriram, Peter Todd, Steve Waldman, John Whelan, Mark Williams, Ryan Zagone, Vlad Zamfir and Sidney Zhang

Unless stated otherwise, the views expressed within this report are mine alone and quotes and/or citations herein are not endorsements of this report by the individuals above or organizations they represent.

Financial disclosure: this research was self-funded and conducted as part of a continual effort to understand the fluid landscape involving the confluence of virtual currencies, consensus systems, economics and fintech which is forming its own field dubbed "cryptoeconomics." I did not receive any financial compensation or remuneration for the following report. I hold no coins or tokens that are mentioned by projects within or without. Some of these findings were initially presented to the Sim Kee Boon Institute in Singapore, where I am a research fellow and to R3 CEV where I am also an advisor. I am an equity-holding advisor to both Hyperledger and Clearmatics, but neither was given preferential treatment or provided any additional compensation for their inclusion.



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

Glossary

A **consensus mechanism** is the process in which a majority (or in some cases all) of network validators come to agreement on the state of a ledger. It is a set of rules and procedures that allows maintaining coherent set of facts between multiple participating nodes. In the case of Bitcoin, the “longest chain” – the chain with the most proof-of-work – is considered to be the valid ledger. There are multiple alternative consensus mechanisms which have been developed over the past three decades. For instance, the Dijkstra Prize is an annual award for academic research on distributed computing. Some of the ideas and innovations from these proceedings have influenced both distributed databases which require fault tolerance (e.g., Paxos from Leslie Lamport) and distributed economic consensus methods. The primary distinction between the former and latter is that of maturity in formalization, analysis and implementation.

A **cryptocurrency system** is a network that utilizes cryptography to secure transactions in a verifiable database that cannot be changed without being noticed. For the purposes of this report, a cryptocurrency system assumes that transactions are transmitted in a peer-to-peer fashion. Traditionally one way to undermine a peer-to-peer network is by creating large amounts of pseudonymous identities in order to gain a disproportional amount of influence (or votes). This is called a Sybil attack. Bitcoin was purposefully designed to make it expensive to attack the network in this manner. It establishes an ordering for transactions through a 'proof-of-work' process - a capital-intensive Sybil protection mechanism, which makes it costly for an attacker seeking retrospectively to impose a new transaction ordering, which could enable them to double-spend transactions. Bitcoin includes a native token (a currency application) to fund network security. Bitcoin is the first public, decentralized cryptocurrency system that used this validation method. A proposed alternative to proof-of-work, dubbed 'proof-of-stake' is sometimes used in cryptocurrency systems, but is not in the scope of this report.¹

A **distributed ledger system**, is a network that fits into a new platform category. It typically utilizes cryptocurrency-inspired technology and perhaps even part of the Bitcoin or Ethereum network itself, to verify or store votes (e.g., hashes). While some of the platforms use tokens, they are intended more as receipts and not necessarily as commodities or currencies in and of themselves.² The Bitcoin blockchain is very commonly characterized as a distributed ledger, yet for the purposes of this report, distributed ledgers are those reliant on legal institutions and as such, a final commonality is the **permissioned** identity system which is defined below.³

A **miner** is a colloquial term for a node that originally both validated and selected transactions and consequently submitted “proof-of-work” to other nodes on the

network. This set of tasks has since been divorced over the past four years: whereas miners still validate and select transactions, **hashing** (or **hashers**) focus solely on solving the “mid-state” in a proof-of-work puzzle.⁴ That is to say, hashing machines are not miners, as they do not validate or select transactions to be placed into blocks. Other literature defines a miner as a Dynamic Membership Multi-Party Signature (DMMS).⁵

A **mining farm** is an amalgamation of multiple hashing machines. Technically these should be called hashing farms but instead they are commonly referred to as mining farms. These farms can range from a hobbyist that owns a few machines at home to several thousand hashing systems located in a professionally managed data warehouse facility.

A **mining pool** is an entity that manages and operates a transaction processing node. In practice, multiple hashing farms typically submit their proof-of-work to the pool which then sends block headers back to the farms who then begin to generate values for the “puzzle.” This is a continuous, automated process.

An **on-chain** transaction is one which users settle transactions on the public blockchain. For instance, in the first couple years of its existence, Bitcoin users sent bitcoins to one another directly through the blockchain as there were no external intermediaries or custodians available. Over the past several years, transactions that occur off of the blockchain, on the edges, have become increasingly popular as it allows for faster clearing. **Off-chain** is the euphemism used to describe this activity. Hosted wallets, exchanges and many other services now exist to provide off-chain services and are managed by their own internal accounting records (e.g. exchanges perform buy/sell operations off the chain via their own private database).

A **permissioned** system is one in which identity for users is whitelisted (or blacklisted) through some type of KYB or KYC procedure; it is the common method of managing identity in traditional finance.⁶⁷ In contrast, a **permissionless** system is one in which identity of participants is either pseudonymous or even anonymous. Bitcoin was originally designed with permissionless parameters although as of this writing many of the on-ramps and off-ramps for Bitcoin are increasingly permission-based.

A **smart contract** is a simple rules engine; cryptographically assured business logic that has the ability to execute and move value.

Turing-complete is a math concept and for the purposes of this report relates to programmability. A Turing-complete programming language means that the language can be used to simulate any other computer language (not just its own) – it is a set of instructions that includes conditions, loops, and read/write memory. As an example, scripting language in Bitcoin is not Turing-complete because there are no loops. The original Bitcoin software implementation released in 2009 includes a scripting system called Script that it was intentionally not Turing-complete.

Report highlights

- Distributed ledgers and cryptocurrency systems are fundamentally different.
- The key difference involves how transactions are validated: Bitcoin uses pseudonymous and anonymous nodes to validate transactions whereas distributed ledgers require legal identities – permissioned nodes to validate transactions.
- Consequently, distributed ledgers are able to legally host off-chain assets due to their authenticated, permissioned approach to validation. Bitcoin and other permissionless systems cannot.
- Eight alternative projects are looked at. Each are at different levels of development and the overview includes a brief description of the value-add they are attempting to provide.
- Several use-cases for smart contracts and distributed ledgers are identified and comparisons between different types of distributed ledger technologies are made.
- Interoperability between systems, between new and legacy systems, is found to be an important intersection (e.g., where do smart contracts begin and end).
- There are a variety of trade-offs between permissioned and permissionless systems involving speed, cost reduction, censorship, reversibility and finality. And due to their gated approach, permissioned systems as a whole are capable of clearing and settling assets faster and are cheaper to maintain than capital-intensive permissionless systems.
- Based on this research, it is unlikely that financial service providers like banks will have a need for cryptocurrency systems primarily because systems such as Bitcoin use anonymous validators and are unable to be a legally official register of assets due to the network's vulnerability of transaction reversal by anonymous attack.
- In comparison, by design, permissioned, distributed ledger systems are more congruent with the existing banking system and therefore provide more utility to financial institutions.

Introduction

The purpose of this short report is to describe the divergence between ‘permissionless’ cryptocurrency systems (such as Bitcoin, Ethereum, Peercoin) and ‘permissioned’ distributed ledger systems (such as Ripple, Hyperledger). Consequently, we will briefly explore some of the use-cases that distributed ledger systems could play, specifically in the financial services industry.

This document assumes that the reader is already familiar with Bitcoin (the blockchain) and bitcoin (the commodity, currency or asset). If you are unfamiliar with these concepts, then some of the vocabulary, concepts and analogies may not make sense.

But before we get into the characteristics of distributed ledger systems, we will look at what has led to bifurcation. Appendix B delves deeper into the process called “mining” which is arguably the most important – yet often overlooked – part of the Bitcoin network.

Section 1: It begins and ends with the miners

Over the past year there has been a measurable increase in interest related to cryptocurrency systems by enterprises, financial institutions and governmental organizations.⁸ Many banks, for instance, have internal teams exploring, poking and probing this technology. Yet in conversations with these researchers and decision makers, they typically eschew “coins.” This position is typically stated as, “I like the blockchain but not bitcoin.”

Ignoring the task of on-boarding the existing community to such an idea, it is unlikely that any new type of code could be integrated or added to Bitcoin itself to remove bitcoin (the coin) from the blockchain. The coin is an integral part of the network’s incentive mechanism to maintain its security; the two have an existential symbiotic relationship.

But that is not to say you could not start from a fresh “mulligan,” taking part of the toolkit – some of the cryptographic primitives and concepts – and start over with something tailored to specific use-cases. For instance, if all parties to a transaction are known, do you really need to have the level of “proof-of-work” (POW) or any type of proof-of-work, as is used by Bitcoin, for this operation? Probably not.

Recall that one of the core assumptions made in the design of Bitcoin is that because the history book of transactions (the ledger where votes are collected) resides on a public, untrusted network, participants (or voters) would be comprised of both good and bad pseudonymous actors (or attackers). Proof-of-work is all about defending against a Sybil attack, which by definition is not an attack vector you have if validators are known.

In other words, unlike a ‘permissioned’ system that whitelists (or blacklists) users based on a centrally managed identity, Bitcoin’s designer attempted to create a ‘permissionless’ system to accommodate pseudonymous actors. And the way Bitcoin distinguishes between friend from foe (or a valid vote from an invalid vote) is to require each participant to expend work on a math problem that on average takes about 10 minutes to solve.

In practice, the pseudonymous voters in this system – dubbed miners – have adopted a specialized role that is technically termed as “hashing.” As described at length in Appendix B, in 2009 when Bitcoin was launched, all individuals were both miners (e.g., validating and selecting transactions) and hashers (generating proof-of-work for the “scratch-off-puzzle”).⁹ Beginning in December 2010, with the advent of pools – also detailed below – a type of “outsourcing” took place as the roles separated into two discrete groups.¹⁰

Similarly over the past several years, the on-ramps and off-ramps for Bitcoin have become increasingly permissioned-based, adding additional costs to a purposefully expensive network.

Section 2: Bypassing the Maginot Line

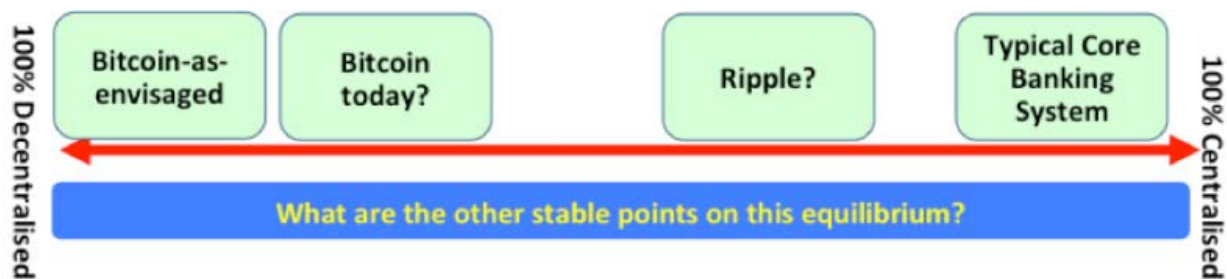
In theory, the capital outlay to change the majority-rules voting record (the consensus of the ledger) is the cost to control 51% of the voting.¹¹ In economic terms, the *maximum* costs to

“brute force” and manipulate the voting power with “hashrate” alone (or proof-of-work-rate) would read as follows: $0.51 \times MC$ where MC is the marginal cost for creating these votes.

While any number of Bitcoin promoters conclude that the network is secure from military budgets of state-funded actors, in practice, due to out-of-band attacks (i.e., “rubber-hose cryptanalysis”) the real costs are much lower. That is a topic for another report.¹²

But what if you were building a network in which each node was known, identified and had a real-life reputation to uphold?

Richard Brown, an Executive Architect, Industry Innovation for Banking and Financial Markets at IBM, has been wondering about the same question. Below is his visualization of the spectrum between 100% decentralized to 100% centralized.¹³



Source: Richard Brown

Brown poignantly asks: what are the other stable points on this equilibrium? Are there only binary choices or can the rules and reliability of consensus be finely tuned on a granular level?

For instance, if Alice manages Bank of Bob (BoB) in the United States and wants to transmit some kind of value to other banks she has multiple choices today. Alice can use ACH, Fedwire or even some kind of cryptocurrency. Each has advantages and disadvantages.

Rather than explaining how ACH and Fedwire work, let us quickly look at one of the advantages and disadvantages of plugging into the existing Bitcoin network.

For an ordinary commercial bank like the kind operated by Alice, ignoring the compliance and regulatory issues surrounding using and holding a bitcoin, one purported advantage is that a transaction could be sent to any party with a bitcoin wallet – such as a bank located on the other side of the country or the world in about an hour (assuming six confirmations) – all for a nominal fee which many Bitcoin adopters claim is a matter of cents.

Yet the disadvantage for Alice, in the long run, is the large unseen marginal cost of protecting these increasingly permissioned transactions. Recall that the corresponding bank she sends value to is known – there are no “untrusted” actors on either side of the transaction. So if they

are all known, why use pseudonymous or anonymous validators? Why create potentially new risks?

In July 2014, the European Banking Authority published a report highlighting precisely these concerns surrounding anonymous validators:¹⁴

The risks include the fact that a VC [Virtual Currency] scheme can be created, and then its function subsequently changed, by anyone, and in the case of decentralised schemes, such as Bitcoins, by anyone with a sufficient share of computational power; that payer and payee can remain anonymous; that [Virtual Currency] schemes do not respect jurisdictional boundaries and may therefore undermine financial sanctions and seizure of assets; and that market participants lack sound corporate governance arrangements.

A more expansive list of details related to the process of Bitcoin mining can be found in Appendix B.

Are these issues that a bank wants to worry about?

Section 3: Why is this important?

Again, because all parties are known in the bank transfer example above, Alice has no need for the expensive Sybil protection utility provided by Bitcoin. What does that mean for the future?

In the long run, the Bitcoin network will require either significantly higher transaction fees to end users or lower security thresholds from miners (or both). Neither of these options are attractive to Alice. She wants to be able to depend on a network for more than a decade without having to worry about cryptoeconomic game theory issues such as block reward halvings (i.e., every four years the reward miners receive splits in half).¹⁵

From an economic incentive perspective, how do we know that public-facing proof-of-work cryptocurrency systems like Bitcoin (as they are today) is not congruent with the existing commercial banking infrastructure?¹⁶




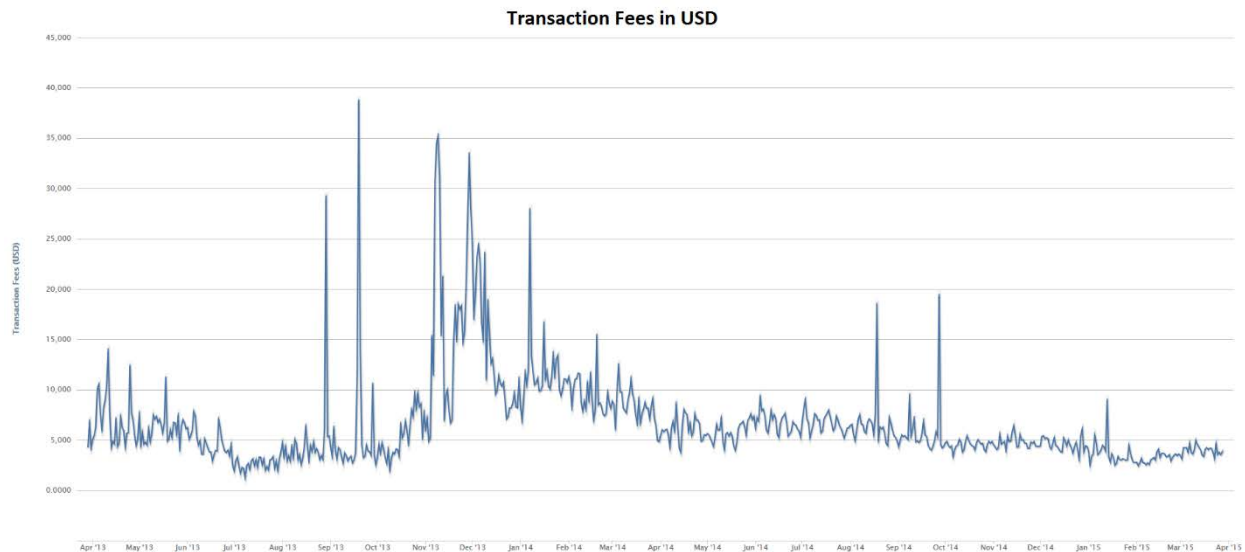
Image credit: Kerem Kaskaloglu

Above is a visual representation that depicts the assumed block reward-to-fee transition made by Satoshi Nakamoto back in October 2008.¹⁷ In section 6 of the original white paper, Nakamoto explained that:

The incentive can also be funded with transaction fees. If the output value of a transaction is less than its input value, the difference is a transaction fee that is added to the incentive value of the block containing the transaction. Once a predetermined number of coins have entered circulation, the incentive can transition entirely to transaction fees and be completely inflation free.

What has happened in practice?

Based on the theory above, we should already be observing a sharp increase in fees to replace the seigniorage-based subsidy.



Source: Blockchain.info/charts

Instead, what we see above is that over the past two years, despite a marked increase in merchant adoption and public awareness, fees to miners as measured in USD, remain flat.

Yet there is no economic law that states that Bitcoin “has to” have enough security. Theoretically, it is entirely possible that the seigniorage will drop to zero, fees will stay constant, and the network will just end up insecure and continually forked. That is the nature of public goods; they do not have to be satisfied.¹⁸

If the inflation reward goes “too low,” maybe what that means is that the incentive mechanics were not properly modeled and designed, and a system with, say, perpetual inflation as a fixed per-year security fee will become popular and replace Bitcoin. After all, if there is no incentive to pay fees other than competition for the limited blockchain space available; why would users pay it?¹⁹ Yet there may very well be value in a decentralized store-of-value with at least some way of moving that value around. However, we are unlikely going to find out for sure until the currently “high” Bitcoin inflation rate drops and hard choices have to be made.²⁰

In short, so far users prefer not to pay fees for Sybil protection to transmit value. In fact, most growth appears on the edges within trusted third parties called “hosted wallets” (hosted

wallets may be depository institutions).²¹ While the trend may change in the future, it appears that the first publicly available commodity-based censorship-resistance-as-a-service (CRaaS) has limited appeal.²² As a result, as of this writing roughly 99.49% of the revenue a miner receives is still from block rewards (seigniorage). Whether or not this is sustainable for cryptocurrency systems like Bitcoin is beyond the scope of this report.²³

Section 4: What does any of this have to do with distributed ledgers?

Rather than throwing the “baby out with the bathwater,” is there a way of using distributed consensus mechanisms to transmit value transparently and securely without expensive proof-of-work methods? Or is Bitcoin the only way to accomplish this objective?²⁴

If some of the underlying Sybil protection is unneeded, what can enterprises, financial institutions and governments do with this toolkit?

Once again, this is something Richard Brown has illustrated in the following diagram:²⁵

		Who do I trust to maintain a truthful record?			
		A central authority	A group of known actors	A group of actors, some known	Nobody
What is the universe of “things” I need people to agree on?	Ownership of on-platform assets	Central Bank, Commercial Bank		Ripple (XRP)	Bitcoin
	Ownership of off-platform assets	Custodian Bank	Hyperledger	Ripple (Gateways)	Colored Coins, Counterparty
	Obligations and rights arising from an agreement	Clearing House	Eris	Ripple (Codius)	Ethereum

Source: Richard Brown

As we can see, there are a variety of ways to look at who maintains a record and what needs to be agreed on.

However, after some discussion with both Richard Brown and Robert Sams, co-founder of Clearmatics, the X-axis could probably be modified in this model.²⁶

Why? In practice, Hyperledger, Eris and Ripple should all be in the “nobody” category. The core principle of distributed consensus is that Bob does not need to trust the validators to create a truthful record. If Alice builds a distributed consensus network on a model of known validators and has a mechanism for punishing those who do not follow the protocol, you actually minimize trust even more than you do on a network designed around anonymous validators (such as proof-of-work).

But this of course relies on some off-chain mechanisms for the authentication and enforcement and that largely compromises the censorship-resistance.²⁷ Therefore Bob and Alice have a

choice as to which variable they want to maximize: censorship resistance or irreversibility. They have to pick one because they cannot maximize both. This is discussed at length later.

Section 5: Visions of the future

Over the past two years, a number of new distributed consensus systems have been proposed and a few have been built out into proofs-of-concept with the backing of venture capital.

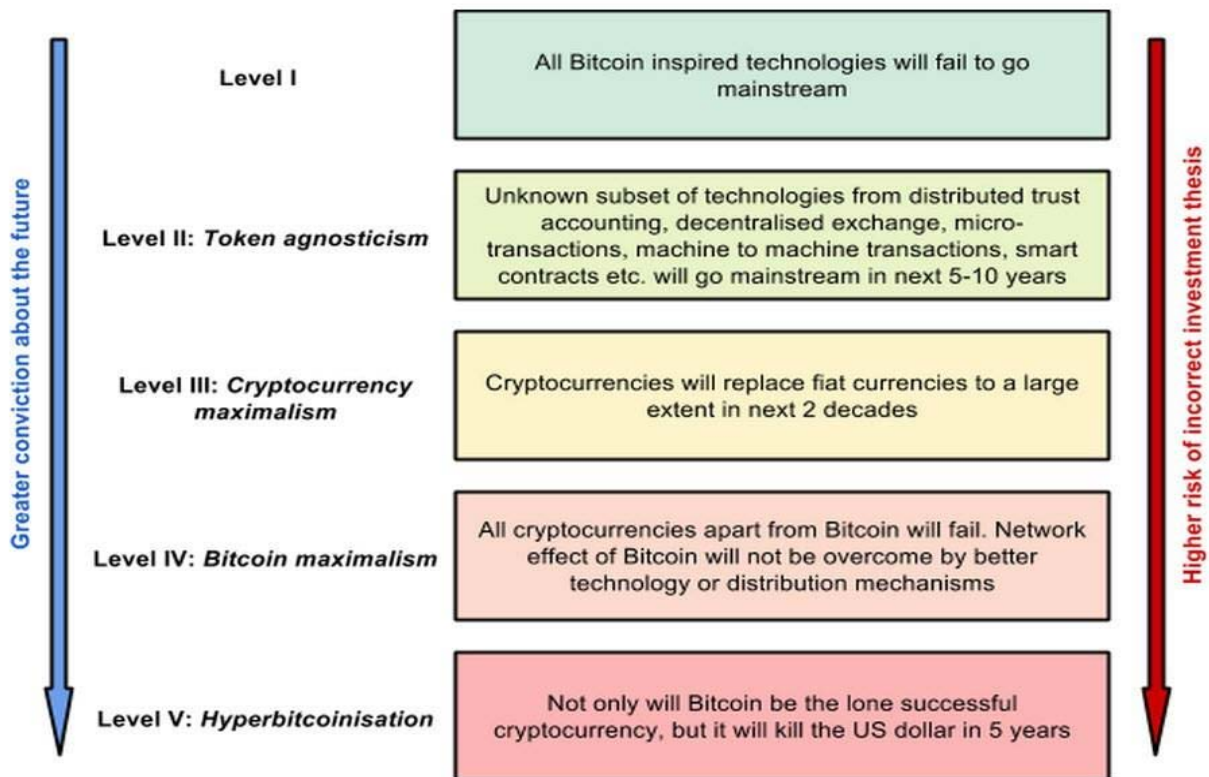
Some of these systems use tokens (like Ripple) and others do not use tokens (like Hyperledger).²⁸

Yet some of those with tokens are being built without the expectation or even without the intention of making these coins available for purchase to retail customers.

Why not? Because the developers are not necessarily trying to create another currency or commodity. Rather, the tokens primary uses are as verifiable cryptographic receipts *internally* between permissioned parties, as a way to prove that certain events happened at certain times for the parties involved as well as for outside compliance and auditing agencies.

But according to social media posts, aren't tokens supposed to eventually absorb everything?

Earlier this year, a classification of the "tokenize everything" meme was visualized in two charts by Meher Roy, an engineer at Novartis:²⁹



Source: Meher Roy

The first chart (above) visualizes the “conviction” or enthusiasm an individual may have towards a cryptocurrency-based tokenized economy.

To date, most of the deployment of venture capital has involved the funding of companies and projects based in Level III, Level IV and Level V scenarios.

Is there another way to look at this?

Belief / bet	Platform opportunities	Incremental risk	Advantages
Level I	Not Applicable	Not Applicable	Not Applicable
Token agnosticism	Hyperledger, Eris, Codius, Ripple / Stellar	<ul style="list-style-type: none"> -Lack of solutions for Identity and Private Key management -Regulatory uncertainty resulting from end-users controlling transactions -Platform specific flaws like weak consensus algorithm 	<ul style="list-style-type: none"> -Applicable to all assets including fiat money, shares and cryptocurrencies -Can replicate all applications pioneered by cryptocurrency community -Relative compatibility with existing regulations
Cryptocurrency maximalism	Bitcoin, Ethereum, Tendermint, Pebble, Ripple / Stellar (partially) etc.	<ul style="list-style-type: none"> -Societal inertia to new forms of value, needs massive network effect -System that possesses sound monetary policy and consensus method, fast transaction speed and is scalable appears late. -Associated political ideologies prevent mainstream growth 	<ul style="list-style-type: none"> -Market segment dissatisfied with conventional banking system is a ready market. -Significant public interest for the time being.
Bitcoin maximalism	Sidechains	<ul style="list-style-type: none"> -New technologies that improve on network maintenance cost, transaction speed and scalability outcompete Bitcoin 	<ul style="list-style-type: none"> -Significant first mover advantage for Bitcoin
Hyperbitcoinisation	Not Applicable	<ul style="list-style-type: none"> -Opinion proves to be a delusion 	<ul style="list-style-type: none"> -None

Source: Meher Roy

The chart above, also from Roy, illustrates the risks, advantages and opportunities of the same five levels.

For the purposes of this report we will look primarily at Levels II and a few in Level III. Whether or not Level IV or Level V occurs is also beyond the scope of this report.

The reason this is included is to provide a general, visual categorization of where investment has gone based on certain assumptions that may or may not occur.

Irrespective of what platform is adopted by market participants, it is likely that “smart contracts” are part of it.

Section 6: What are smart contracts and what are they good for?

Every business, institution and organization has different business needs. In order to determine whether or not integrating a new network is beneficial to your firm, the total costs of ownership must be clearly understood. And without looking at continuously changing business

requirements it cannot be said *a priori* whether or not your company will actually benefit from any of the technologies listed in this report.

However, there may be a couple of new tools that these distributed ledger systems (and even cryptocurrency systems) may be able to provide to a broad array of financial institutions.

For instance, “smart contracts” are perhaps the most often cited examples to try to explain the future utility of these networks.

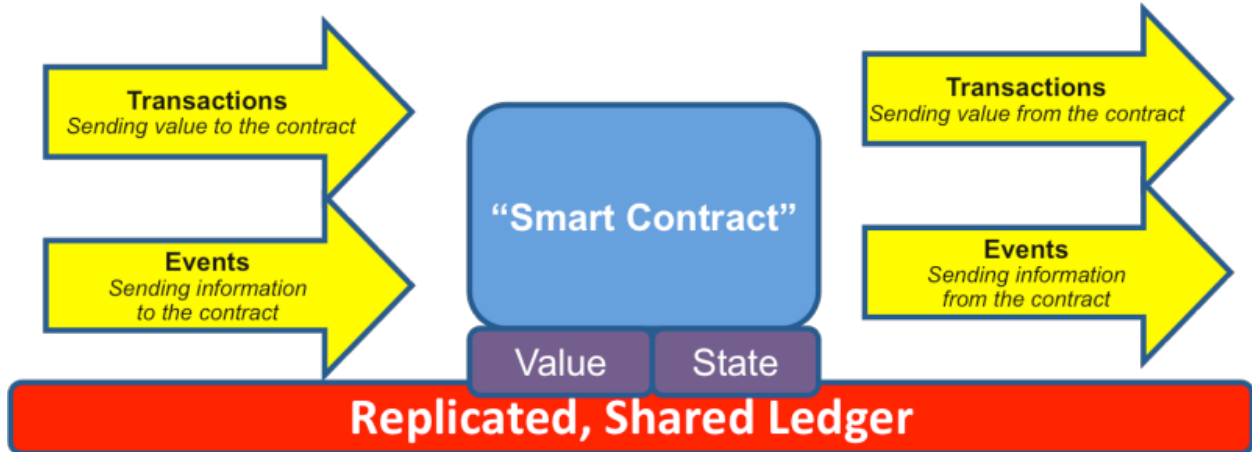
What is a “smart contract?” While there is still no consensus among the community yet, last year I previously used this definition:

Smart contracts are computer protocols that facilitate, verify, execute and enforce the terms of a commercial agreement.³⁰

I think this should be superseded by a newer, clearer definition by Richard Brown who defined it as follows:

A smart-contract is an event-driven program, with state, which runs on a replicated, shared ledger and which can take custody over assets on that ledger.³¹

One way to think of it is as a programmable calculator that can receive inputs – execute code – then provide an output. And because it resides on a distributed ledger, it is difficult for any one party to necessarily modify (abuse) the program.³²



Source: Richard Brown

Above is a visual aid of this mental model from Richard Brown.

While the proposed use-cases so far have ranged from the possibly ill-conceived to the seemingly logical, the fact that this program resides on a shared ledger likely makes it at least a new tool for managing financial controls. Is further segregation of financial controls important to your company?

Or as Brown explained:

And now you have something really interesting: neither of you have to go to the effort of reimplementing the terms of the contract in your own systems: you both know that this single piece of code satisfies *both* your purposes. And because it is running on this shared, replicated ledger and using it as its source of information, you can both be sure that whatever the program outputs will be the same for both of you. [...]

It's as if this program isn't just a computer program: it's an actor in its own right. It responds to the receipt of information, it can receive and store value – and it can send out information and send out value.

In theory, they can be valuable tools at both an interoperability level and application level (e.g., representation of a financial contract) depending on specific instances. Whether or not the legal system within a jurisdiction views smart contracts as an actual legal contract (as opposed to a secure multiparty financial computation) is a matter of debate and it will likely take several years, if resolution is possible before a more definitive conclusion can be made.³³

What can these computational contracts do?

Section 7: Use-Cases

While there have been hundreds of proposed use-cases over the past year in a variety of forums, conferences and social media channels, based on a number of conversations and research with individuals such as Mikkel Larsen, a managing director at a bank, in terms of what a “smart contract” can do, the ones probably most relevant to financial institutions include:

- Cross Border Settlement / B2B international transfers
 - Improving the SWIFT and correspondent banking network;
 - Can use consensus-as-a-service / blockchain-as-a-service to securely, transparently move value in seconds or minutes;³⁴
 - The biggest challenges however are local pools of liquidity, settlement with market makers and compliance in each jurisdiction;
- Central clearing (e.g., derivative clearing)
 - Prime case for “multi-party payments” as well as netting and clearing
 - Could be on a distributed ledger but if participants “fail” then the operator will likely move to centralize the credit risk which was the purpose and function of CCHs and CCPs in the first place;
 - Complying with existing laws such as Sarbanes-Oxley and Dodd-Frank are a continual challenge;
- Mortgages
 - The use-case is the ability to have a financial vehicle that can be used equally by many parties and “self-execute”;
 - It need not be blockchain or distributed ledger if a single bank is trusted;
 - It becomes increasingly useful only when banks (the lenders) or a new 3rd party is not trusted to fairly register (e.g., installment payments);

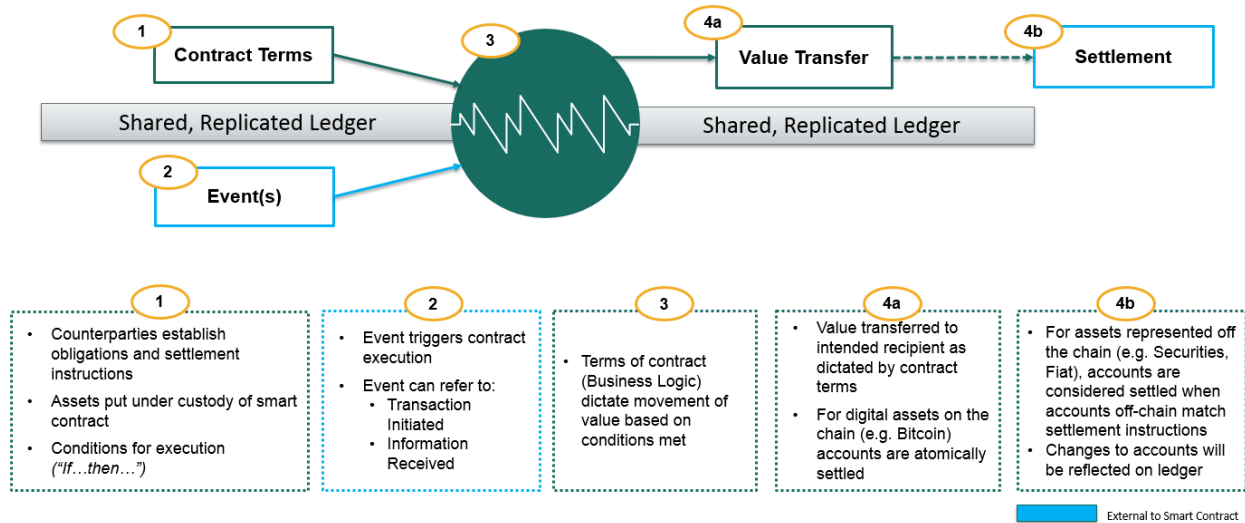
- It may be more relevant for CDOs;
- CDO/CLO/CMO/ABS
 - Smart contracts based on assumption that banks are not to be trusted to pass on all cash flows received in the “waterfall”;
 - Alternatively, building competing platforms where you set up “smart contract” (special purpose vehicle) that automatically pays through waterfall;
 - The largest problem is on enforcement of loans in case of non-payment;
- Collateralized / Guaranteed Lending
 - A bank, borrower and potentially a 3rd party providing collateral or guarantee;
 - Though without identity, credit checks or credit worthiness, the promise of decentralization may not do much;
- Letter of Credit / Bill of Lading / Trade finance
 - Multiple parties involved, trust is low, cost is high;
 - Incumbents are strong, little incentive to change, requires central changing (with “crossing the chasm” problem) and most importantly: multiple jurisdictions;
- Crowd Funding
 - Borrowers may request money on multiple platforms but also make investment fungible;
 - Challenges involve legal constraints: if securities are issued (e.g., Howey test) then SEC regulations may apply

While there are a multitude of other purported use-cases including the managing of escrow accounts and microlending there are at least three dependencies that cannot be overlooked for all of them:

- Immediate readiness
- Core to bank business lines vs ancillary
- Legal enforceability

To be frank, as of this writing there are not many smart contracts – perhaps not even one that can immediately fulfill the core business lines of banks and pass legal scrutiny in a developed country. Perhaps that will change over the course of this year.

What does this look like altogether?



Source: Jo Lang / R3 CEV

Above is a visual aid from Jo Lang of R3 CEV.³⁵ It is a flow chart for applying business logic with smart contracts.

Again, the list of potential use-cases will likely fill volumes by the close of the decade. But that does not mean that these distributed ledgers or cryptocurrency systems that provide support for “smart contracts” are of any use to your company or organization. Maybe they are all Rube Goldberg experiments, a solution to a problem that does not exist or a solution that is unnecessarily complex.

As each firm has different needs, can a distributed ledger or smart contract fulfill any of those needs?

And how can these ledgers, cryptocurrency or otherwise, be tied together on a global basis?

Section 8: Building a neutral network of value interchange

Over the past two years we have also continually heard that Bitcoin is the TCP/IP of money.³⁶

This is false. Bitcoin is not a protocol, it is a network of validators first and foremost. Not only does Bitcoin use TCP itself to operate, but TCP is an actual neutral agnostic protocol. Bitcoin is not. In order to use the Bitcoin network a user must pre-pay for a certain amount of tokens. In effect, even though the currency is “neutral,” there is vendor lock-in. With “the Internet” there is no equivalent to tcpipcoin or emailcoin. If a user is no longer satisfied with Gmail they can use a different service or even operate their own. Gmail (Google) does not own the protocol.³⁷

In fact, unlike TCP or SMTP, there are special interest groups that effectively control the Bitcoin network through token issuance and ownership. And one of the main reasons why there are now over 550 forks of the Bitcoin codebase (commonly known as “altcoins”) in operation is that once a subset of coin holders and miners becomes solidified as the “in” group, it motivates

those in the “out” group to fork the code, spin up a new chain and distribute new coins to their own special interest group (usually some small clique).³⁸

In contrast, the internet packets generated in TCP are done so dynamically on demand. There is currently no vendor lock-in or special interest group that can or does own any percentage of them not just because of their infinite creation rate but because the fundamental utility is moving amalgams of data around globally, not holding onto individual packets.³⁹ The Bitcoin network is not open like the internet is; successful Bitcoin companies directly increase the wealth of bitcoin holders leading to a chicken-egg public goods problem: why invest in a Bitcoin company at all when you can free-ride off of the utility and value others create?⁴⁰

A similar, albeit imperfect facsimile to existing user behavior in Bitcoin (where 70%+ of bitcoins are stagnant) was observed with the gated “gold rush” for web domains (TLDs) by which people bought and squatted on large numbers of domains with the hopes that eventually they could sell them for yachts and islands.⁴¹⁴²⁴³

In addition, neither SMTP nor TCP try to do everything. At a basic level they just create traffic rules for the delivery of packets. Furthermore, with TCP, there is no internal value system or attempt to secure value exchange (this is what Pactum is attempting to solve). Whereas Bitcoin tries to be the jack-of-all-trades: the network, security apparatus, currency, and so forth. And the result is a little similar to a “clown car”: the Bitcoin network is capable of fitting in lots of clowns – metacoins and clowncoins – but it eventually becomes top-heavy and neither safe, effective nor efficient.⁴⁴

So if Bitcoin is not the low level foundational system that its proponents claim, then what could a neutral network of value interchange look like?

Last fall, Meher Roy created the model below for a neutral “Internet of Money” based on the existing OSI framework.⁴⁵

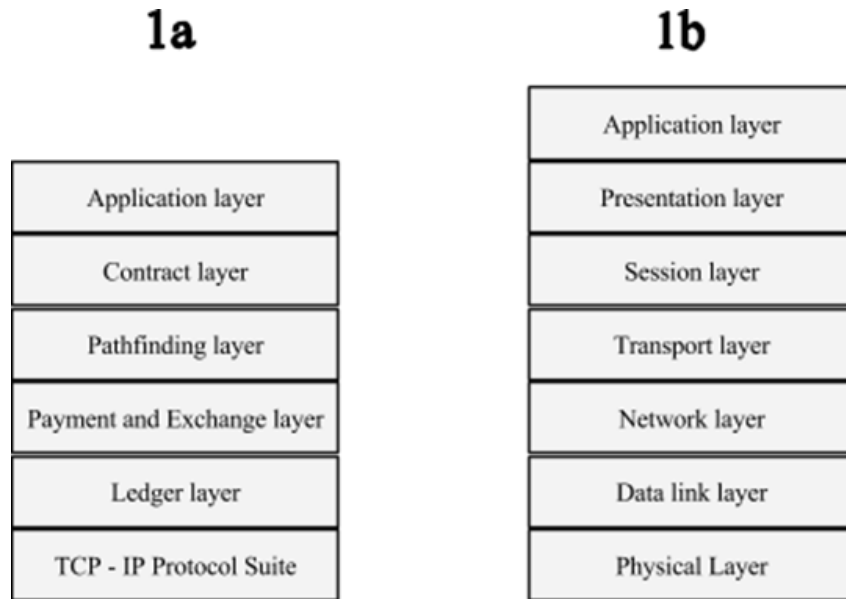


Figure 1a. Framework for the Internet of Money

Figure 1b. OSI 7-layered framework for comparison

Source: Meher Roy

The core idea behind his proposal is that a majority of financial operations can be efficiently executed by leveraging smart contracts. In his white paper, financial operations include:

1. **Asset Exchange:** An exchange of assets between Alice and Bob can be brokered by two smart contracts executed on two different ledgers tracking the assets. This setup reduces the functions of clearing and settlement into a smart contract mediated fast transaction (~2 seconds). There is no requirement for Alice and Bob to trust each other.
2. **Peer to peer payments:** A payment of USD from Alice to Bob, where Alice and Bob are customers of two different banks, requires the ACH or Fedwire systems today. Smart contract enabled distributed ledgers offer a third option: Any party, such as Carol (unrelated to the banks) can create a channel between ledgers and facilitate the payment. The transaction structure can be conceived as Alice paying Carol on one ledger, and Carol paying Bob on the other ledger. Carol's role is fully automated using a smart contract, and she can earn transaction fees as per her choosing. A market for peer-to-peer payment channels can form.
3. **Fedwire like Real Time Gross Settlement:** An RTGS transaction can be brokered by two smart contracts if the Receiving Financial institution and the Federal Reserve utilize smart contract enabled distributed ledgers. This could reduce the function of messaging, clearing and settlement into a singular smart contract based transaction.

4. ACH like Deferred Net Settlement: Similar to the point 3 above, except that inter-bank debt obligations are netted and settled later.

Smart contract enabled protocols for value transfers between ledgers are one candidate for the Internet Protocol (IP) equivalent of money. The Internet Protocol serves to mediate the exchange of data between two geographically separate Local Area Networks. Similarly, operations described above serve to mediate the transfer or exchange of value between two distributed ledgers.

Once value transfer protocols are in place, more complex inter-nation value transfers, such as payments originating in Citibank USD and terminating in UBS Swiss Francs, can be realized by a chain of operations mentioned above. The most efficient chains or *paths* can be computed using automated services.

This idea leaves intact all innovative characteristics of cryptocurrencies. Any application pioneered by the cryptocurrency community can be implemented for the banking system by programs utilizing the value interchange mechanism.

This also dovetails into two different ideas: the first involves a tradeoff between the “subjectivity” of how to qualitatively assess the health and robustness of a distributed network versus exploitability of a consensus method but this is a topic for future research.⁴⁶ It also hooks into the recent discussion surrounding a hypothetical “Fedcoin” (e.g., a distributed ledger operated by a central bank).⁴⁷

How does this neutral protocol tie in with distributed ledgers?

Section 9: The characteristics of a distributed ledger

Based on the research conducted, what are the current attributes of a distributed ledger?

For starters: a distributed ledger of off-chain assets cannot be both censorship-resistant *and* authoritative. Why not? Because of the interaction with the existing legal system, one that is not going to disappear, legal enforcement of contracts vis-à-vis identity is paramount. Recall that once a user goes “off-chain,” legal code, moves from ‘dry’ computer code to ‘wet’ human jurisprudence.⁴⁸

For instance, the attributes of a permissionless blockchain (e.g., Bitcoin) as they exist today include:⁴⁹

- Censorship resistant
- Reversals possible (though originally intended to be irreversible)
- Arguably only suitable for on-chain assets

Steve Waldman, a software developer, independently labels this type of blockchain as “antidiscretionary.” According to his definition, these blockchains have three attributes:⁵⁰

- Nodes / members lack well-defined identity

- Forks are technical glitches to be resolved mechanically, as fast as possible
- Ideally, all nodes or members face incentives to behave “correctly”, such that the behavior of the community is understood by and predictable to outside entities (“users”) who interact with the community

According to his definition, both Ethereum and Bitcoin are examples of antidiscretionary blockchains.

But what is the point of building a financial product such as a smart contract that tracks off-chain assets if it cannot be legally enforced in the real-world?

This report uses Robert Sams’s model, which proposes that a permissioned blockchain (e.g., distributed ledger) is as follows:

- Legally accountable validators (per EBA recommendations)
- No reversals – settlement finality
- Suitable for off-chain assets such as securities, fiat currency, titles




Figure 1: Permissioned blockchain

According to Waldman, the flipside to his “antidiscretionary” blockchain is a discretionary blockchain (“soylent” blockchains because people are involved). And the characteristics of discretionary blockchains in his model:

- Nodes represent identifiable members
- Forks represent disagreement. They must be resolved, but may persist a while (e.g., eventual consistency)
- Outsiders that interact with the community may tolerate a degree of temporary uncertainty and be offered a mechanism to try to force consensus (e.g. quorum and mutisig endorsement)

While there will likely be new proposed definitions⁵¹ and characteristics based on Sams's model and reinforced by Waldman's proposal, the following projects were selected as part of this report: Clearmatics, CryptoCorp, Eris, Hyperledger, Ripple Labs, Tembusu, Tezos and Tillit.

What do these 8 projects all share in common?

Again, while the definition could shift over time, I think that Robert Sams's mental model best describes the overarching four core, common themes:

- They each use or are an independent Blockchain (Permissioned not Permissionless)
- Have a built-in or companion distributed Virtual Machine (Turing-complete)
- Smart contracts govern off-chain assets
- Network achieves settlement finality

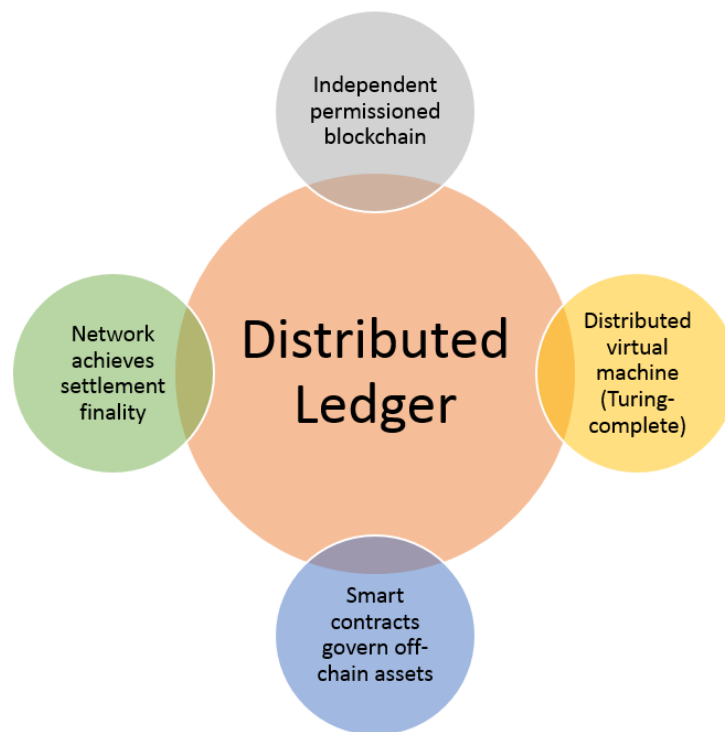


Figure 2: Components of a Distributed Ledger

Recall that due to regulatory requirements, irrevocability and finality (or sealed recording) is a *feature* of major payment systems such as the Bank of England (e.g., CHAPS), SEPA, Fedwire and Bank for International Settlements.⁵²⁵³ Conversely, it is the *potential* of reversibility in Bitcoin (and any protocol with anonymous validation) without any legal recourse or accountability that makes it a poor solution for off-chain assets.⁵⁴

Why does a distributed ledger need to have a Turing-complete virtual machine? In Sams's view, it is because projects like Clearmatics are interested in using the blockchain to host distributed ledgers of financial assets and applications.

For example, each of these off-chain asset types have a certain complexity such as corporate actions (e.g. stock splits), dividends or coupon payments. It is not a simple matter of “X units of Y held by Z.” Without Turing-complete scripting, a user would have to hard-code the logic of each asset into the blockchain's scripting.

In addition, ledger entries themselves are not simply a matter of “transfer X units of Y from Z to W.” There is complexity there too: a ledger entry is usually the performance of some contractual obligation. It dilutes the value of a blockchain-based data layer if Alice and Bob need a sundry of off-chain controls and software to ensure ledger entries are legally legitimate. In fact, one of the key value propositions of a blockchain is the encoding of the contractual obligations on the same chain that hosts the data-layer. Arguably, that needs Turing-complete scripting too.⁵⁵

Why not just use a simple database?⁵⁶

It is worth emphasizing that the shared property of the ledgers (with the question of the ownership of the records to remain open) is one of the core advantages and features of a distributed ledger as it helps avoid replication errors and delays.⁵⁷

One additional, helpful mental model that Sams’s uses to describe the pitfalls of a centralized ledger, such as a database, is comprised of “3 sins”:⁵⁸

- 1) Sin of Commission – forgery of transaction
- 2) Sin of Omission – censorship of transaction
- 3) Sin of Deletion – reversal of transaction

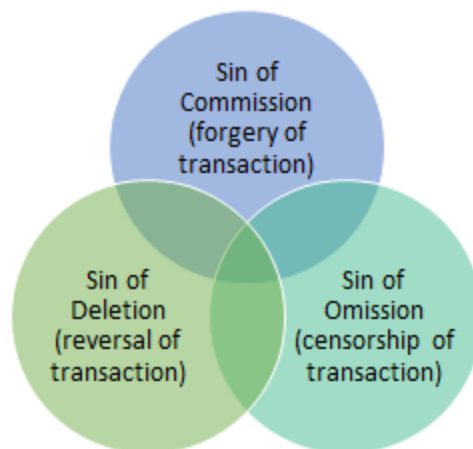


Figure 3: Pitfalls of a centralized ledger

Eliminating the first step is arguably the easiest: cryptography and secure key management already deal with #1 to prevent record forgery. It is #2 (sin of omission) and #3 (sin of deletion) where priorities differ depending upon use-cases. In fact, the “trustless” meme crooned over the past several years has likely served as a counterproductive distraction because what is really important is censorship resistance through minimizing the risk of #2.

Cryptocurrency systems prioritize mitigation of #2 over #3, whereas any system of off-chain property titles will have to prioritize #3 over #2.⁵⁹⁶⁰ And consequently, existing legal systems will likely never recognize a system of property titles that can be reversed by anonymous validators.⁶¹

Conversely, it is not hard to minimize #3 to levels of practical impossibility with a consensus protocol that solves the Sybil attack through authentication of validators rather than “proof-of-work.” As we will see, Ripple was first to try this attempt, yet others are taking the authenticated approach in new directions as well.⁶²

While it may be unpopular with “early adopters” such as some of the original cypherpunks, minimizing reversibility comes at the price of tolerating some censorship imposed by governments who can hold transparent validators accountable for the transactions they validate.

According to Steve Waldman, these two networks are fundamentally different tools with different purposes:⁶³

- Antidiscretionary blockchains prioritize values of predictability and authority
- Discretionary blockchains prioritize participation, representation, and flexibility

In his view, “fundamentally, an antidiscretionary blockchain is a technique for deploying a long-running, predictable software application on top of a community of people whose role is merely to verify. A discretionary blockchain is a technique for reifying and composing the ever-changing will of a community in the form of a distributed software application.”

What does this trade-off look like?

	Permissioned validator	Permissionless validator
Tracking on-chain assets	Sensible in some use-cases	Optimal
Tracking off-chain assets	Optimal	Does not work

Again, as shown above, this bifurcation is not saying one approach is better than the other, rather, they each solve different problems.

Yet it should be noted that because all of the 8 distributed ledger projects listed in this report use trusted validators, they can securely clear and settle assets much faster than Bitcoin, usually in a matter of seconds instead of minutes. They also do not encounter the public goods and scaling challenges that Bitcoin and Bitcoin-like networks are currently facing, such as increasing block sizes while simultaneously trying to maintain a semblance of node decentralization.⁶⁴

What about off-chain agreements?

If the goal is to have participants come to agreement over who owns off-platform assets (such as property titles), reversibility by an anonymous validator is intolerable. This is one of the reasons “sidechains” for financial settlements – while technically novel – tends towards the possibility of a reversal by anonymous (or DMMS) validators which makes this an impractical choice and thus was not included in the 8 cases above.⁶⁵

However, agreements of on-platform assets (such as a cryptocurrency) is different. These can and likely will continue to have a “code-as-law” jurisprudence that tolerates the occasional reversal as the price to pay for censorship resistance. For example, recall that in March 2013, Bitcoin underwent a 24 block reorg.⁶⁶ When the bitcoin “community” chose one side of a fork in March, anyone who ended up worse off for the maneuver had no one to hold accountable or demand restitution from. With permissioned blockchains, there are entities who can be held responsible, scrutinized and held legally accountable. It is this accountability of reversals, rather than its nonexistence, that renders permissioned blockchains more suitable for higher-stakes activity.⁶⁷

Section 10: What about a fusion of permissioned and permissionless?

There are many technological solutions to achieve the authenticating of user transactions via a set of authorized authenticators. Yet the question remains: what do we achieve by doing this? Why authorize users but not authenticators?

As a hypothetical, if we put such a scheme on top of a metacoin, or as a sidechain, or some other permissionless “internet-of-value” platform, we may not gain the benefits of the permissionless foundation such as censorship resistance and mitigation of the sins-of-omission yet may gain all of its drawbacks such as much more costly to prevent transaction reversals (sins-of-deletion). It is arguably the worst of both worlds.

In other words: one really needs to ask if it makes any sense to require the authentication of users but not validators. By authenticating validators you can dramatically lower the probability of reversal risk and do so at a fraction of the cost of a proof-of-work based system.⁶⁸

But if a company or organization can effectively approve or deny nodes, can't they also control consensus?

It is true that an authenticated validator set-up can deny any transactions that it disagrees with. That is an implication of the thesis that blockchains can either minimize censorship or reversal, but not both.

What is untrue is that the authenticator (let's call him “the witness”) thereby controls the consensus and can arbitrarily change the rules in its favor.

If the consensus protocol states that:

1) Validators must build blocks that are descendants of the last block signed by “the witness”

2) The witness can only sign blocks that are descendants of the last block signed by “the witness”

Then neither the witness nor a subset of the validators can reverse a transaction that is at or behind the last block signed by the witness (i.e., “settlement finality”).⁶⁹

One of the reasons this may not have been explored over the past six years is because “the witness” in this protocol sketch above cannot be anonymous: we have to trust him to sign only valid histories. This is not a problem if the blockchain is sponsored by some legal entity with a public reputation and accountable to the law. But the protocol does not work for an anarchic chain such as Bitcoin (as-is).

The other implication is that the witness can censor transactions. The prevailing notion within the cryptocurrency space fallaciously conflates this property with the power of a database administrator.

Why not just release an API to interact with the creator's cryptographically signed database entries instead?

To work with the database analogy, in the sketch above, “the witness” only has the power to determine which “INSERT” statements get executed (because he can refuse to sign a branch that contains transactions he does not like); but he can never execute a “DELETE” statement. And this is why the permissioned blockchain concept goes much further than what the cryptocurrency space was thinking about in the 1990's, where transaction servers ran Chaumian cash.⁷⁰ Those always assumed trusting someone to not delete history or double-spend.

Which, then, is a ledger of real-world property titles? Ideally a user would want a ledger to be such that nobody has the power to reverse history. The censorship “feature” then is a non-issue, because legal systems might not recognize a ledger that makes entries for title transfers from thieves to sanctioned counterparts (e.g., *nemo dat*).⁷¹ Yet, censorship prevention has to be there anyway to keep the distributed ledger in sync with what the law considers an authoritative record of ownership.

Thus this distributed ledger is farther ahead of database technology because there is no database administrator with the power to issue a command for “DELETE.”

And again, contrary to popular belief, there are no central ledgers in the financial system that records ownership of all the world's financial assets.⁷² In point of fact, existing ledgers are

spread out among thousands of different entities and are held in sync by very robust but slow and expensive reconciliations and financial controls.

There has also been conflation between “centralized” with “institutional.” Distributed ledgers are reliant on legal institutions that is not the same thing as being centralized. It also bears mentioning of the differentiation between “political” and “technical” centralization and decentralization. A single entity can hold the keys to the business logic of the state change, but the state change itself can happen in a completely distributed fashion. The other way around can also be true.⁷³

Consequently, a distributed ledger can replace the patchwork of existing ledgers above with a more efficient decentralized setup that is both faster and cheaper. And it becomes increasingly more robust when run on a Turing blockchain: it makes assets and financial contracts programmable, with automated enforcement of terms.

The next section will include a brief overview of each of these projects.

List of companies

clearmatics

Clearmatics is developing a new clearing and settlement platform for OTC financial markets, bringing together custodians, dealers, trading venues, buy-side firms and data providers onto a single platform. Here members can settle securities trades and automate the valuation and margining of derivatives and other financial contracts using Decentralised Clearing Network ("DCN") technology.

Their platform offers transparency, counterparty risk mitigation and settlement protocols that optimize collateral usage. Built on a model of Turing-complete scripting, Clearmatics has the flexibility to design application protocols that both integrate existing intermediaries as well as alternative, smart-contract based disintermediations like bi-lateral delivery-versus-payment.

Each DCN is Clearmatics technology managed by a membership-based legal structure ("DCN Co"), capitalized and owned by a consortium of important market participants. The DCN technology is designed such that assets represented on one DCN can be easily moved to another DCN, creating a larger "meta network" of many different DCN's dedicated to different markets.

Clearmatics aims to make the DCN the new and open standard in financial clearing and settlement (two fully integrated processes on its model), eventually replacing the patchwork of proprietary custodial accounting systems such as SWIFT messaging.

Clearmatics is based on the Ethereum Virtual Machine, specialized for financial and fiduciary computations and it uses a new consensus protocol designed to achieve finality of settlement and eligibility as a Designated Settlement System. They work closely with the Ethereum project and Eris Industries to encourage common coding standards among the different use-cases of Turing-blockchain technology.

But Clearmatics is not just about the automation that DCN's provide to existing markets. Rather in their view, it is the enabling aspect of making assets and financial contracts programmable. New derivatives specifications can be designed in programming code by end-users, and asset cash flows can be stripped, repackaged and traded "using code instead of lawyers." This includes what they term permissionless innovation in the design of new derivatives contracts (an alternative to ISDA), and a streamlined process for the origination of new securities.

The world today is comprised of dematerialized assets and derivatives contracts with fully computable terms (encoded in legal documents rather than a programming language). Clearmatics is trying to enable the settlement of these to be controlled by distributed ledgers and code run on an industry standard, distributed virtual machine.

Clearmatics was founded in January 2015 by a team of financial professionals and are based in London. On the web: Clearmatics.com

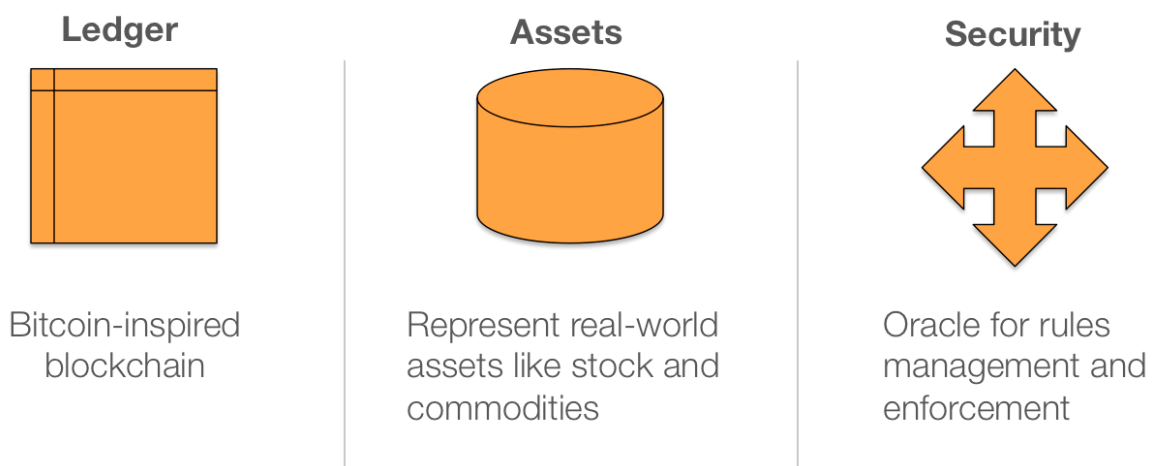


CryptoCorp Inc., offers an integrated software stack for developing blockchain-based ledger systems for traditional assets. By offering an integrated software stack, CryptoCorp helps financial institutions and other enterprises to quickly develop applications, pilot projects and prototypes without requiring deep cryptocurrency expertise.

The integrated software stack is comprised of three components: a blockchain-based ledger, an asset issuer and an oracle. The blockchain ledger is based on Bitcoin Core, but is modified for high transaction volume and enterprise deployment. The ledger can be deployed in isolation from the public Internet and is not associated with the public Bitcoin network. The asset issuer allows real world assets to be represented in the ledger such as equities, currencies and commodities. Oracles and multi-signature technology offer security and auditability - both critical requirements for any financial or enterprise application.

All three components are designed to work together. Applications can be built using different libraries, including Python, Java, and Ruby, and the components' functions are accessed via APIs. Developers can focus on an application layer on top of this stack to serve a specific function or service without needing to worry about the blockchain "plumbing." While the integrated stack is essential to most financial services functions, other components and plugins are available for additional functionality.

CryptoCorp is currently working with customers to build custom applications and services on this integrated software stack. On the web at: CryptoCorp.com and point of contact, Peter Shiau at peter@cryptocorp.com.





Eris Industries is building tools that allow developers to solve big data-driven problems:

- ErisDB: their FOSS blockchain template, is fully programmable and controllable. ErisDB is designed to give rise to independent instantiations which individual developers, corporations or other platform operators secure and control themselves.
- ErisServer: their FOSS client program/oracle machine, which runs locally on a user's computer and allows ErisDBs - or any other database - to link into any other data structure and present a harmonized application in a web browser. Together with ErisDB, it means a serverless, in their words: "internet of blockchains."

Eris believes that in order to be viable in a commercial setting, a blockchain software solution must be:

- flexible, upgradable, and reversible
- fully controllable
- infinitely repeatable (millions of separate instances if necessary)

In their view, today's blockchain platform operators do not permit this: they either create standalone cryptocurrencies or compel vendor lock-in, on their own servers, to address a limited set of use-cases. And it is their opinion that such blockchain/consensus databases are difficult to change. They cannot be tailored to meet the specific needs of specific applications. In their view, "full trustlessness" or "true decentralization" – though prized in the cryptocurrency community – is not a commercial problem in other industries. Based on their research, Institutions, governments and corporates (especially banks) need control and flexibility.

Eris also likens blockchains as database software – it will do what operators tell them. Eris views blockchain databases as rulebooks for data-driven interactions and Bitcoin as being merely one application and that consequently, more are possible with an ErisDB backbone and the modular ErisServer client.

Where Bitcoin is a rule-book for clearing and settlement (and its structure collapses settlement logic from three days to a matter of minutes), with ErisDB, a similar mode of thinking can be applied to any other process, thus giving rise to cost savings and other automation efficiencies for ErisDB operators.

They recently built a technology demonstrator, called 2Gather, to illustrate how YouTube might be done on a blockchain. The aim of 2Gather is to show that, instead of relying on Google, Uber, AirBnB or Spotify, users can now administer these applications themselves.

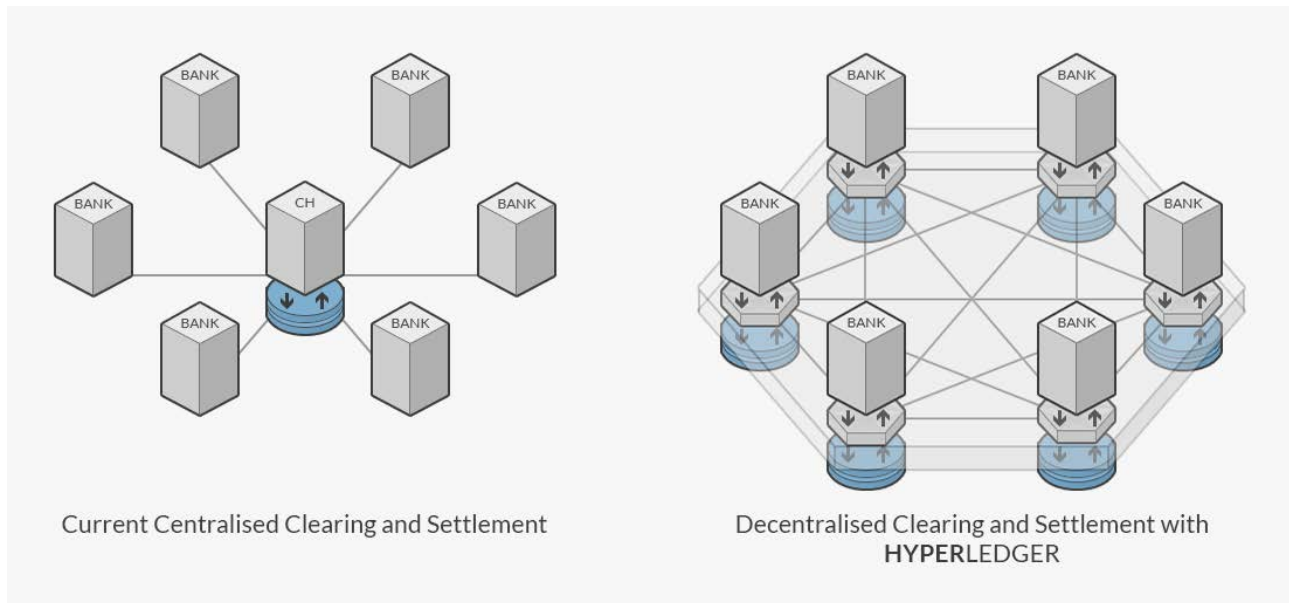
Launched in December 2014, Eris Industries is headquartered in London and its team launched the first smart contract-driven blockchain developer platform to market.

They advise a variety of companies and financial institutions on how to put distributed computing technology to use in their organizations. And their seed investor is Anthemis Group, S.A. On the web: ErisIndustries.com

Hyperledger has taken a different approach than many other blockchain systems by removing what they consider unnecessary features and provide just shared replicated ledgers. In their view, this core technology, the distributed ledger, can then be integrated into existing systems to allow banks and financial institutions to settle in real time, mitigating risk and the need for expensive reconciliation without the need for a central party.



Designed with identity and compliance in mind, Hyperledger enables operators to know all the participants in a network while still being interoperable with other pools. Hyperledger uses a consensus algorithm, Practical Byzantine fault tolerance (PBFT), which has existed for more than 15 years and is purportedly capable of handling tens of thousands of transactions per second per pool, without the need for capital intensive “proof-of-work” mining.



Unlike Bitcoin and even some distributed ledger systems, Hyperledger does not have an in-built cryptocurrency. According to them this means: less regulatory risk, less technical overhead, no volatility and truly asset agnostic rails.

Another feature of Hyperledger is that it is not one single, international, public ledger. Why? In their view financial institutions need to:

- Create multiple ledgers for different asset classes
- Keep balances and trades private
- Know who is operating the nodes and which jurisdiction they are in
- Control who can open accounts on their ledgers

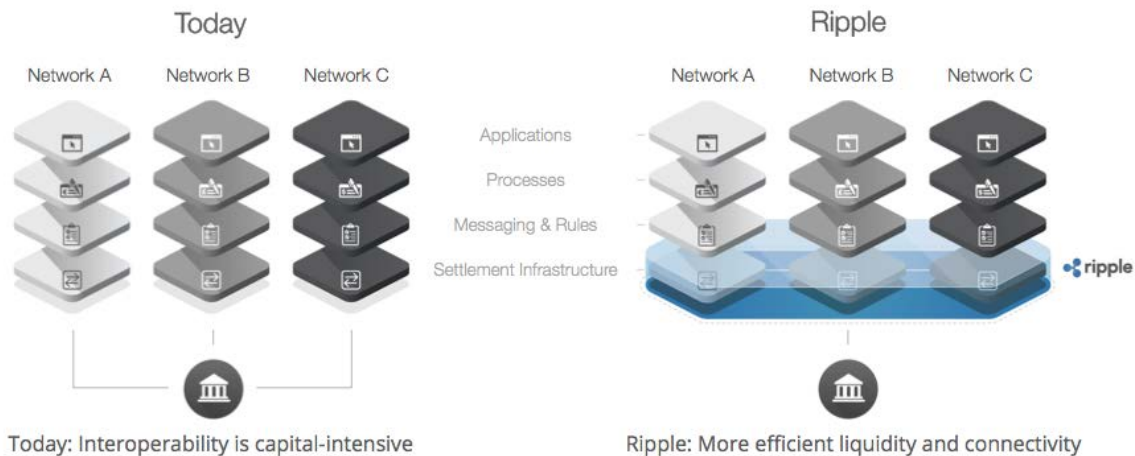
Founded in 2014 and based in San Francisco, Hyper are the commercial entity supporting the open source Hyperledger Value Transfer Protocol. They are working with financial institutions, consortia of banks and start-ups to deliver solutions to mitigate settlement risk, reduce costs and delays across multiple financial instruments; from FX, Interest Rate Swaps, Securities, to correspondent banking. They won the Bank Innovation Challenge and are semi-finalists in the SWIFT Innotribe Startup Challenge later this year. Contact team@hyperledger.com for more information.

Ripple, rather than circumventing today’s systems, makes existing currencies more efficient by maximizing liquidity and interoperability. The vision for Ripple is to create interoperability at the core of the payment stack—the settlement layer. This enables all other existing layers to remain intact and be more interoperable as well.



Make existing currencies more efficient

Maximizes liquidity and interoperability



Ripple was not designed to replace the existing payment ecosystem but improve its foundations. For instance, the side effects of correspondent banking—how international payments work today—are delays, high costs, and counterparty risk. Transfers typically take 2-4 days and the process and fees are often opaque.

Ripple is an alternative to correspondent banking that enables near-instant, bilateral, straight-through processing. Transfers typically take 3-6 seconds and the process and fees are transparent. And validating transactions does not involve “proof-of-work” that systems such as Bitcoin do.

Regarding domestic clearing, most central banks settle periodically in net batches via a real time gross settlement system. Ripple enables real time settlement -- on a gross or net basis, all day, every day. Their technology enables point-to-point settlement, lowering the cost and risk of transactions.

Ripple’s current target users include banks, central banks, and other financial institutions such as payment networks. Ripple Labs, the parent company, has publicly announced three banks—Fidor Bank, CBW Bank and Cross River Bank—and one payment network, Earthport, one of the largest in the world. Ripple Labs is also actively engaged with global regulators including the New York State Department of Financial Services, the Australian Senate, and the CSBS, among others.

The roughly 100 member team is based in San Francisco’s financial district and has received funding from Andreessen Horowitz, Google Ventures, Core Innovation Capital, Lightspeed Venture Partners, IDG Capital Partners and others. On the web at: RippleLabs.com.

Tembusu Systems

Tembusu see themselves as integrators and are attempting to provide financial services to the unbanked and underbanked through several approaches. They are currently using a fork of the Ripple network and codebase instead of Bitcoin which they view as slow and increasingly centralized, yet without the benefits of centralization. According to them, they are building a distributed transfer network and for transparency and redundancy, thus distributed ledgers and not cryptocurrency systems, made more sense.

Tembusu has a working transaction system based on “Proof of Identity” and are attempting to cover the whole on-ramping spectrum, from the nodes to mobile wallets and all the entry points for other businesses to join. They also provide APIs, reference implementations and white label applications. The plan is to start two pilots, one in Singapore and one in small country in Western Europe. These are both open to a limited number of real users.

The first pilot involves remittance from Singapore to Philippines and Thailand because based on their research there is a large amount of remittance flowing that direction, and the smaller the amount, the larger the cost. Thus they are targeting the “long tail” of the market, the millions of people who remit less than \$1,000 per month. It is planned to just be one-way for now so that they can have very tight control on identities and money laundering. The second pilot is a test network in a Western European country, with one thousand users, integrated with point-of-sale systems and cash in - cash out stations. Through this they wants to start experimenting with a completely cashless society.

They are also looking to track “real assets” on the ledger and are working with a company who wishes to track gold and other metals on the blockchain. The partnering company provides a certificate that they have deposited X kg of gold and that it is authenticated and earmarked. Afterwards the company, in this scenario, will be able to mint X kg of gold “coins” that people can buy on spot prices.

While this is already available through a number of other companies -- what Tembusu is trying to do is provide better reach for those in developing countries. For instance, a user on their system can buy the gold, keep it as an investment and send it to their family in another country, where they can cash it out on spot price. In their view, using an asset-backed coin for remittance has the advantage that an end-user can forego volatility and the foreign exchange conversion charges as well. This ties in with another product they have developed: using these metals as a collateral against a microloan. Tembusu is building a system that enables users to receive low interest loans: it is based on collateral in their accounts (e.g., gold) and also through a peer rating system, based on a real world social network.

Founded in 2014, the team of 22 employees is based in Singapore and has received \$1.5 million in seed funding. On the web: Tembusu.sg



Tezos is a distributed, cryptographic ledger for the OTC derivative market.

It currently costs more than \$25 million to create a new OTC derivative product and back-office processes are even costlier and more inefficient, particularly when the contracts are not vanilla. In an effort to lower costs and allow a variety of new OTC products to flourish, Tezos' model is for financial institutions to adopt a legal covenant whereby they agree to manage contracts and liabilities in an unambiguous smart contract language on their secure, distributed ledger.

The founder's vision is that distributed ledgers must be thought of as networks, not merely as software. Thus, Tezos provides a built-in governance model to maintain the competitive edge of its network. This governance model:

- Gives stakeholders full choice over the technological enhancements to the network.
- Solves important collective action problems, such as agreeing upon standards or trading centers.
- Allows for integrating new features into the protocol as first class citizens, which preserves scalability and composability. Tezos considers that this is a major advantage over solutions which implement new features *within* smart contracts, and not at the protocol level.

Tezos feels that their team's knowledge of financial markets allows them to get a few critical points right. In particular, financial institutions generally do not want to reveal their positions or even the existence of an open interest. In Tezos' view, most other ledgers require contract terms to be public by design. Thus, they provide what they call "trustless off-chain contract arbitration," which protects the privacy of the parties to the contracts (while retaining auditability for the regulators) and greatly improves scalability.

In addition, Tezos considers that margining and cross-collateralization are crucial aspect of OTC derivative clearing, and they are thus built within the protocol.

The project is one year into development, has a working prototype and was built from a ground up (it is not based upon Bitcoin or Ethereum, but is its own implementation). The type of development practice common for consumer-oriented tech companies (hack and fix) is not suited for the financial industry, as demonstrated by the visible failures of many new businesses in that space. The Tezos team has implemented its network client entirely in OCaml, a functional programming language, with an eye towards security and formal verification.

Tezos built its own code base from the ground. They believe that this gives them a better solution than could have been obtained by forking a preexisting project which was not designed to handle the complex requirements of financial markets.

Founded in 2014, the project is currently being led by a team of financial professionals based in New York City as well as two developers based in France. On the web: Tezos.com



Tillit is building a financial asset servicing and settlement infrastructure solution (called the “Platform”) powered by the Ripple decentralized ledger and tailored to the direct lending market.

The Platform comprises a suite of tightly integrated components, including an API and hosted interfaces for a wide variety of direct lending industry participants, including originators, investors, custodians, servicers and advisors. By integrating the company's business logic layer with the Ripple ledger, the Platform enables these participants to adopt asset servicing and settlement processes and technologies, thereby introducing distributed ledger technology to their respective services and product offerings.

The principal features of the Platform include:

- Programmable money flows to automate complex servicing functions
 - Direct borrower-lender flows
 - Multi-party structured products (i.e., securitizations)
 - Apportionment of income streams to other Platform Participants (i.e., fees, escrows, reinvestment)
- Integrated foreign exchange functionality
- Superior settlement capabilities
 - Atomic transactions (no risk of incomplete settlement)
 - Final settlement of executed trades and other transactions in seconds
- Support for custom data analysis and reporting

According to Tillit, existing participants can leverage the Platform to streamline their existing business processes, enhance their service offerings, reduce risk and enhance regulatory compliance. New entrants to the direct lending market can adopt Tillit’s Platform as a technology stack speeding the time to market for originators who are seeking to access the burgeoning direct lending funding model in their own verticals.

In their view, all participants will benefit from the liquidity and efficiency enhancements such as:

- Real-time settlement of trades enables true peer-to-peer trading by eliminating counterparty risk without the need to resort to multilateral clearing mechanisms.
- Faster movement of funds between and among income streams and multiple originator platforms reduces current fragmentation that impedes timely allocation of capital.
- Smarter and more cost-effective servicing improves the efficiency of these products and renders them more amenable to sophisticated liquidity tools like securitizations.

Tillit Inc., a Delaware corporation headquartered in New York, was formed in 2014 to develop advanced products at the intersection of the digital asset and financial industries. The company is comprised of veterans of the legal, payments, and finance industries. On the web: Tillitinc.com

Comparison Matrix

Below is a simple layout of the technology in use at these companies. It bears mentioning that there may be some changes in the future.

Table 1: Distributed ledgers with *permissioned* validators


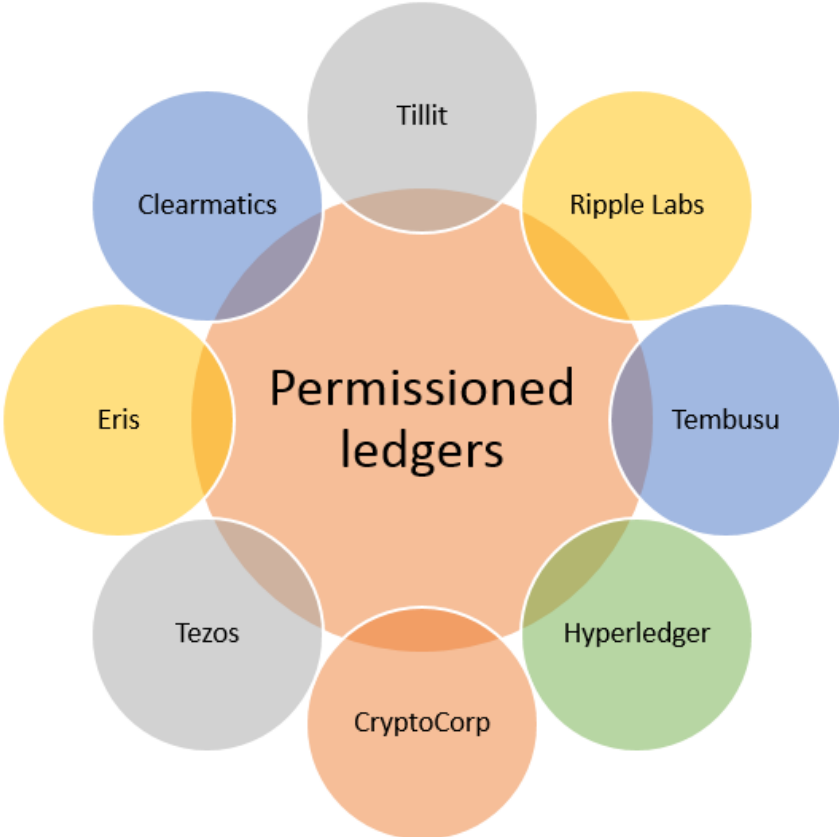
	Ledgers with an internal token	Ledgers without an internal token
Ledgers based on Ripple	Ripple (XRP) Tillit	Tembusu
Ledgers based on Hyperledger		Hyperledger (PBFT)
Ledgers based on a Bitcoin-like blockchain	CryptoCorp (fork of Bitcoin, based on a colored coin implementation)	
Ledgers based on an Ethereum-like blockchain	Eris (with “Junk”)	Eris (without “Junk”) Clearmatics (see EU Electronic Money Directive)
Ledgers based on a Tezos-like Blockchain	Tezos (consensus agnostic)	Tezos (consensus agnostic)

For comparison, listed in Table 2 are several notable projects that are being developed within the cryptocurrency ecosystem. Several of these are discussed in Appendix A.

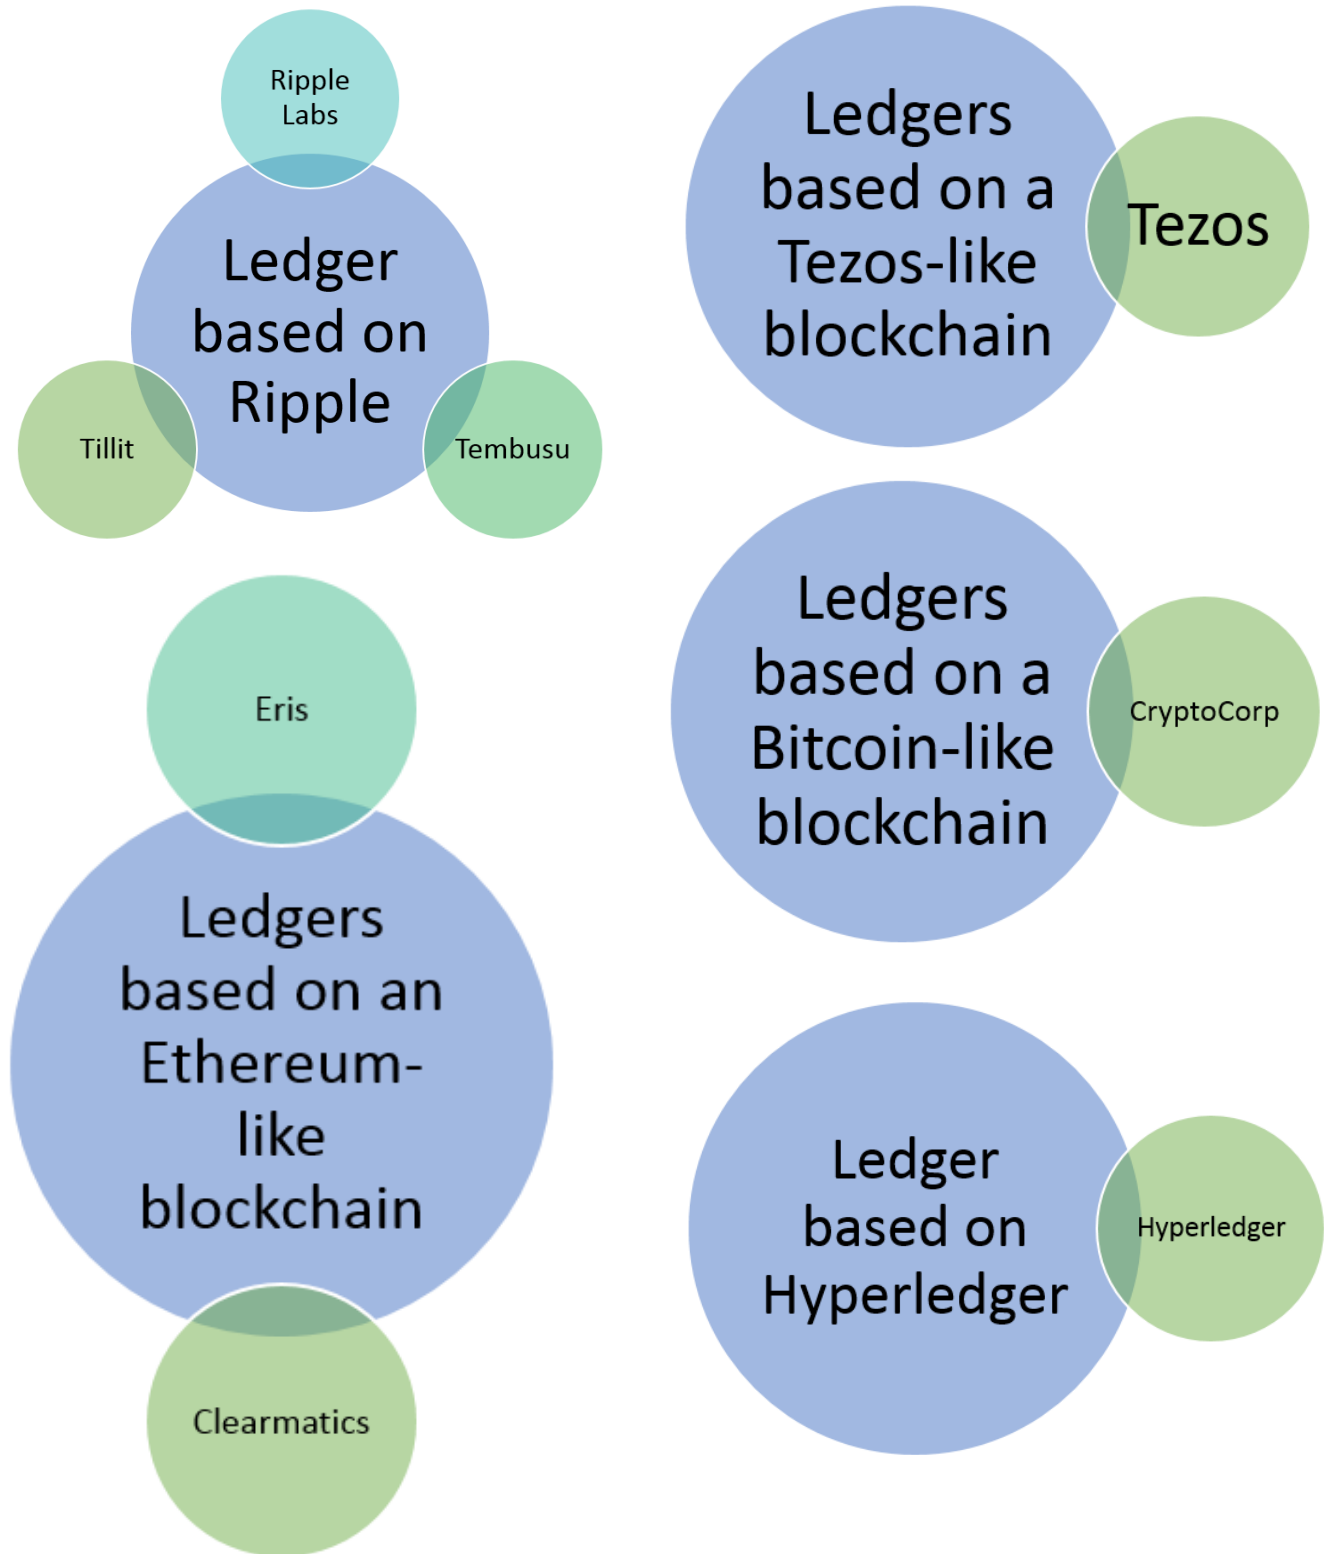
Table 2: Blockchain platforms

Services that rely on pseudonymous or anonymous validators	Blockstream (2-way peg with Bitcoin) Symbiont (“middleware,” can use BTC & XCP) Augur (Truthcoin) Colu, Chromaway, CoinPrism (based on colored coins, watermarked BTC) Ethereum (Ether) Pebble Assets (pebbles) SmartContracts.com (Bitcoin) Tendermint (“mints”) Namecoin (merge mined with Bitcoin)
Ledgers that rely on permissioned validators but not currently focused on financial services	SKUChain (non-Blockstream sidechain) Peernova
Ledger based on Hyperledger, Delegated Proof of Stake (Bitshares), & Open Transactions	Pactum (tokenless)



Permissioned and permissionless ledgers



Permissioned distributed ledgers and underlying consensus method



Permissioned ledgers with and without tokens



Section 11: Conclusions

The financial industry has a number of needs such as transparency, risk management, capital management (e.g., ability to track mobility for value) and compliance requirements that have grown in scope over the past several years.⁷⁴

Due to their authorized validators and cryptographic auditability, is it possible for distributed ledger systems to fulfill these requirements in a timelier, more cost efficient manner than centralized alternatives?

Maybe. But why bother inventing other systems? Won't Bitcoin win out because of its first mover advantage?

Two months ago Brian Armstrong, the CEO of Coinbase, said:⁷⁵

Ripple, Stellar, and Altcoins are all a distraction. Bitcoin is way too far ahead. We should be focused on bitcoin and sidechains.

This is empirically untrue. If Bitcoin was "too far ahead," then axiomatically no one would be working on all these other projects as they would clearly see this trend and focus on one platform. Due to Bitcoin's limited usage and current niche appeal other project development continues apace. In fact, the overall field as described in this report, could be much wider as it relates to actual distributed ledgers and not just cryptocurrency systems.

This would be akin to saying, Reddit, Slashdot and 4chan are all a distraction. GeoCities is way too far ahead and that we should be focused on GeoCities. Or that Windows, MacOS and Linux are all a distraction. V6 Unix is way too far ahead and we should be focused on Unix.

First movers do not necessarily mean last movers.⁷⁶ In 1899, Benz was the world's largest car manufacturer. Diners Club invented the charge card in 1950 but was later acquired by a competitor, Discover. MySpace was the largest social network in 2003 and Blackberry once the market share leader in smart phones.

While the jury is still out on what will ultimately happen with Bitcoin (and altcoins), this report explored the alternative choices that have germinated from its moving parts. Perhaps they will all succeed, fail or fall somewhere in between. We cannot know without rolling up our sleeves, trying and making the sausage. But one thing is for certain: it cannot be said *a priori* that any one is too far ahead. Or that the financial world will adopt any of this technology.

Does that mean that both cryptocurrency systems and distributed ledger systems will likely continue to co-exist independently, not necessarily interdependently, of one another?

Probably.

Why is this?

Because in the real world of finance, all participants are already authenticated and entities like validators and transmitters require legal identities. Off-chain assets in the real-world already map onto off-chain jurisprudence (and vice-versa).

Compare this to the way the edges of the Bitcoin ecosystem has evolved over the past six years: it has grafted together the structural costs of two different worlds without adding any of their overall benefits. If this trend continues, Bitcoin and efforts like sidechains – while novel technical achievements – may end up in a highly expensive, legally ambiguous purgatory: a peculiar hydra.⁷⁷ After all, developers have to solve the title transfer problem – securities are not IOU's.⁷⁸ That does not mean that there is no value or utility in these cryptocurrency systems, rather it is simply not their competitive advantage to have one foot in one world and one in another.

“2.0”

Should we call this effort Bitcoin 2.0 or Blockchain 2.0?

Last spring, the term “Bitcoin 2.0” was an oft-used umbrella term that became increasingly nebulous throughout 2014. Originally it described “hacked-in” platforms that either sat on top of Bitcoin (such as an embedded consensus system like Counterparty or Mastercoin) or that modified and enhanced the technology to support features such as “smart contracts.”

As Robert Sams has explained in several presentations:

The blockchains on which these 2.0 projects are based have a distributed consensus model based on anonymous and permissionless validators, which likely render such blockchains unsuitable as candidates for a “legally official registry for off-chain assets.” Cryptocurrency, being as it is “on-chain” value, is really governed by a jurisprudence of “protocol-as-law,” where ownership means possession of a private key, regardless of whether that key is possessed by a sanctioned counterpart, a thief, jaded lover, *ad infinitum*.

So what appears as a new and powerful technological solution to asset registries, “Bitcoin 2.0” is instead an ideological thesis devoted to bringing back bearer securities in cryptographic form. This may be desirable as a political programme for making property titles resistant to corporate and state censorship. But a network that doesn't provide a one-to-one correspondence between what the network and the law say is who-owns-what is a network that can't exist without the very legacy settlement framework that it seeks to replace, for the latter will remain the authoritative record of ownership. “Bitcoin 2.0” is useless as a solution for financial settlements in cash, securities, and other on-chain property titles.

Sams raises several germane questions: how much does trust cost? How much is pseudonymity or anonymity worth?

In practice bearer assets were difficult, cumbersome and expensive to secure, hence one of the reasons why an entire security apparatus was built around them over the past 500 years through the form of trusted third parties and permissioned identity management systems.⁷⁹⁸⁰ It is called the “traditional financial system.” And in spurts, the Bitcoin ecosystem is quickly recreating all of these same intermediaries yet with a very expensive quasi censorship-resistant network as its rails. It bears the costs of both worlds without delivering the benefits of either.

For instance, in the last four years we have seen several dozen high profile cases of individuals and companies whose bitcoins were lost, stolen or accidentally destroyed due to improper operational security. By one account there are more than a million bitcoins that are no longer with their legal owner.⁸¹⁸² Consequently, in terms of venture funding, the 2nd largest vertical that has received funds over the past 18 months is hosted wallet companies (“depository institutions”) such as Xapo and Coinbase which provide cold storage (“vaults”) and some type of insurance.⁸³ While there is clearly a market need (and demand) for these services, these entrants are increasingly taking on the same roles, vulnerabilities and cost structures as the very institutions they were purportedly to replace yet without the same financial assurance, oversight or controls.⁸⁴

Does this mean that Bitcoin is bad, irreparably broken or destined for the dustbin?

Again, it is not so much that permissioned ledgers are superior or inferior to Bitcoin, it is that they each solve different needs. Bitcoin, circa 2009, was a then-optimal solution for cypherpunks. Hence why it was first posted on a cryptography and cypherpunk mailing list and not an economics or finance mailing list.⁸⁵⁸⁶ Its design parameters require a certain level of decentralization in order to prevent and mitigate the risks of censorship of pseudonymous participants which are not relevant features for modern financial intermediaries or participants with real-world reputations.

It also bears mentioning that the diminutive usefulness of permissionless systems for participants in the permissioned traditional financial system, on the part of Bitcoin, was not some kind of unanticipated shortcoming or design flaw, but a result of intentional choices by these systems’ designers who were quite clearly reacting to aspects of permissioned systems that they disliked (see section 1 of the Nakamoto white paper). It is thus not so surprising that systems designed to circumvent the role of traditional financial institutions would not be as useful to traditional financial institutions.⁸⁷⁸⁸

Permissioned finance is different than permissionless, and each organization should look at which network best supports its requirements. Perhaps, as some Bitcoin enthusiasts suggest, this is all akin to Highlander or Lord of the Rings: there can only be “one chain” to rule them all and that chain is Bitcoin. While it cannot be known *a priori*, this narrative may not prove true for cryptocurrency systems and is most unlikely for distributed ledger networks as well.

Since permissioned, distributed ledgers are congruent with the existing banking system the questions are: what value and utility can these systems provide to you and your organization?

About the Author



Tim Swanson is a graduate of Texas A&M University and worked in East Asia for more than six years. He is currently based in the San Francisco Bay Area and is a research fellow at the [Sim Kee Boon Institute](#) at Singapore Management University.

He is the author of three books including: [Great Wall of Numbers: Business Opportunities & Challenges in China](#), [Great Chain of Numbers: A Guide to Smart Contracts, Smart Property and Trustless Asset Management](#), and most recently [The Anatomy of a Money-like Informational Commodity: A Study of Bitcoin](#)

He can be reached at: tswanson@gmail.com

Keep up to date with further information at: OfNumbers.com and Twitter: [@ofnumbers](https://twitter.com/ofnumbers)

Appendix A: other notable efforts

There are several other notable efforts that while do not currently fit the definition of a “permissioned distributed ledger,” are also creating noteworthy technology.

Blockstream – Blockstream was first announced in March 2014 and by October the team closed \$21 million in venture funding. The project involves the usage of a cryptographic bridge called a “sidechain” that validates data from other blockchains with the Bitcoin network and enables bitcoins – and other digital assets – to be transferred across these “sidechained” blockchains. The platform accomplishes this task by using a “two-way peg,” a process in which bitcoins are frozen and held in stasis on one chain and recreated on another chain at a deterministic rate. The goal is to allow bitcoins to be used as a universal coin between multiple chains, chains that each have unique features (e.g., capable of issuing and tracking smart contracts). The team is comprised of a number of Bitcoin “core developers” and is expected to release a “federated peg” version in Spring 2015.⁸⁹ On the web at: Blockstream.com

Augur - Augur is a decentralized prediction market (PM) and is based on Truthcoin. Since they rely on the “wisdom of crowds,” prediction markets are typically used as forecasting tools and Augur is trying to develop a PM platform for the general public. For example, prediction markets allow their users to buy and sell shares in the outcome of an event and the current price of a “yes” share is the market's estimate that the event will occur. After the event occurs, Bob and Alice need a way to report what actually happened (e.g. after the 2000 election, did Bush or Gore win?). All current prediction markets have a centralized party who does this. Conversely, Augur uses a distributed oracle system which reports on event outcomes (and “reporters” get a portion of trading fees from markets). Using Augur, Bob and Alice can create a question about anything for both prediction market use or to simply ask its “oracles” (trusted data feed) to get external data into the blockchain. A beta is scheduled to be released in Spring 2015. On the web at: Augur.net

SKUChain – SKUChain was founded in September 2014 and as its name suggests, it is taking parts of a blockchain and integrating it with existing supply chains. Its aim is to give manufacturers and distributors upstream provenance and downstream integrity in their supply chain. This helps them with problems such as blocked cashflow, chargeback loss, inventory stock management as well as theft and counterfeit related issues. By using a blockchain they are able to bring the underlying cryptographic primitives to inventory management by treating sku's like a “currency,” allowing them to be minted, sub-divided and transferred on an transparent ledger. A public beta is expected in Q3 2015. On the web at: SKUChain.com

Ethereum - Ethereum was announced in January 2014 and raised 31,529 bitcoins in a subsequent crowdsale. It is attempting to build a new web server platform designed to deliver a decentralized Web. It is integrating encrypted support for: user authentication, payments, P2P cloud storage, instant messaging, and reputation & trust ratings. It is also tying a Turing-

complete programming language within an independent blockchain. Ethereum is maintained and promoted by the Ethereum Commons/Foundation. The team is also responsible for developer relations with three efforts: “ethacademy” (education), “mix” (the dedicated solidity IDE), and “mist” the browser with native Ethereum support. Launch windows have moved and the most recent estimate is for late Spring 2015. On the web at: Ethereum.org

Pactum - Pactum was founded in December 2014 and creates a standardized pathfinding and interaction system for both on-chain contracts as well as Thelonious-style genesis block contracts to interact with one another. Pactum is a layered technology stack, very similar in structure to the Internet Protocol Suite and is not a single product; it is rather a stack of protocols and technologies including three things: 1) an open Blockchain and Distributed State Machine set of standards; 2) a software suite which implements and adheres to these standards; 3) a specific opt-in blockchain with the single purpose of flexible disintermediated P2P Credit and Money creation. In terms of specific software layers, it is comprised of: a networking module; Key Value Store (blockchain) module; Contract Execution Environment (split into the 1) constitutional and 2) contractual part); and the Notarization module (i.e. the Open Transaction Module). On the web at: Pactum.io

Symbiont – Symbiont was founded in January 2015 and announced in March 2015. It is a financial technology company building the first issuance and trading platform for SMART SECURITIES™ on blockchain technology. With its technology stack, Symbiont aims to create a crypto-financial platform, automate corporate actions, provide peer-to-peer settlement and escrow service, and eliminate counterparty risk, thereby decreasing costs and improving the security and efficiency of traditional financial markets. Its core team is comprised of developers that also created the Counterparty virtual currency platform. On the web at: Symbiont.io

Vennd – Vennd is a financial technology startup based in Sydney and one of 8 startups to be accepted into the Australian accelerator StartMate. Their goal is to make it easy to create, manage and exchange digital assets and currencies on a blockchain. Vennd is blockchain agnostic and can create transparently-auditable “micro” exchanges between blockchains. Furthermore, Vennd acts as a bridge enabling transfer of tokens between blockchains. Vennd is in private beta with selected with plans for public release in mid-2015. On the web: Vennd.io

Appendix B: making sense of cents

The real driver behind the concentration of mining as it relates to the economics of proof-of-work as done in Bitcoin can be framed such that: \$/kWh.⁹⁰ Or, how much it costs to dissipate a joule.

So why then is there a fee to “miners” at all and why is this fee measured in a matter of cents?

The fee is actually more akin to a donation since users are still not required to pay a fee (e.g., you can manually send a fee-less transaction).⁹¹⁹² Where does this fee (or donation) go to?

It goes to a third party in between both Alice and Bob; it is called a miner. Several years ago there were thousands of geographically dispersed miners which reduced the possibilities of a censorable, single-point-of-failure and because of this *ad hoc* “load balancing,” it prevented any one of the miners to act as a centralized trusted third party (e.g., none of the miners single handedly controlled enough “hashing” power to generate more than 51% of the votes).

Due to a number of reasons that are outside the scope of this report, capital – like in most industries – has aggregated and centralized into larger assemblages of miners called mining farms.⁹³⁹⁴ And as mentioned above, due to specialization and division of labor, these conglomerations of miners are now more accurately described as “hashers,” as the sole utility they provide is “proof-of-work” and not actual transaction verification or selection that a miner was also supposed to do.

What does that mean?

These farms of hashers (e.g., hash farms or “mining farms”) in turn collectively pool their votes and voting power together – much like coworkers do with lottery pools – forming another entity called a mining pool (sometimes referred to as ‘pooled mining’), organizations which have now been around for about four years. And in return for submitting these votes (or “work”) to pools, pool operators provide income to farmers in the form of bitcoins.

In point of fact, it is one small machine (e.g., such as a Raspberry Pi) at each of these pools that actually verifies and selects valid votes and not the rows of specialized computers called ASICs (application-specific integrated circuit).⁹⁵ And empirically, if a miner is defined as the machine that actually selects and validates a transaction, as of this writing there are only around a dozen miners on the Bitcoin network.

If it is not thousands anymore, how many miners are there then?

As of this writing, there are roughly a half dozen miners (or pools) whose collective hashing power (voting power) represents more than 60% of the votes of the network.⁹⁶ While it is arguable whether outright collusion or cartelization can or will take place, network participants must now *trust* that these pools act a certain way (e.g., do not double-spend or censor votes).

Yet it is important to know a little more about how and why these pools of capital have arisen. Mining farms have real costs to operate which will be described below. And in short, while today the nominal fee to send a transaction is purportedly a few cents, the actual cost is at least two orders of magnitude higher once seigniorage of the block reward (a “batch” reward) is factored in.⁹⁷

The costs of hashing

Why is Bitcoin much more expensive than what is commonly stated on social media?

Because, maintaining a decentralized network has a different cost structure than operating a centralized one.

Furthermore, because validators on the network were pseudonymous by default, the only way to know which nodes are not spoofing or forging identities (known as a Sybil attack) is through a Sybil protection scheme called proof-of-work which requires real costs be borne by someone – after all real resources such as electricity are consumed to do so.⁹⁸

In Bitcoin, miners (which are the sole labor force) are currently rewarded 25 bitcoins roughly every 10 minutes. This is the labor force’s wage for providing Sybil protection (e.g., censorship mitigation) effectively filtering the invalid votes from the valid votes.⁹⁹ In theory and in practice, the network operates on inflation – these additional coins add to and dilute the existing money supply (or “narrow money stock”) until its internal trust fund is fully exhausted next century.¹⁰⁰

Yet there is no need for Sybil protection in a centralized system as actors are known. And while there are other costs related to single-point of failures in centralized models that is also beyond the scope of this report.

Cost structures

Can this decentralized cost structure change?

Again, as noted above, while there are any number of ways to write new features for Bitcoin, ultimately the question boils down to is whether or not the labor force (the miners) will install the update and protect the network. Is it in their interest to do so?

For example, at the MIT Bitcoin Expo held in March 2015, Andreas Antonopoulos stated that:¹⁰¹

Right now, Bitcoin is a featherweight in terms of international terms. Bitcoin is a toddler with only a \$3 billion "valuation". There are many more corporations that have more money. There are countries with more money. By comparison, Bitcoin does not compete as a national currency. However, Bitcoin is currently running a global level security infrastructure which means that Bitcoin is resistant to international class computing attacks against it. We have bought a world class security infrastructure to run a featherweight class currency. Does that look inefficient? Yeah, it really does. We are

securing from multinationals, conglomerates, nation states, etc. Bitcoin is alive despite the fact that it is being relentlessly attacked every day and every minute of every hour. The arms race created by the incentive structure of mining has escalated to the point of delivering world class, international class security for what is still a featherweight currency.

If you look at it, you may say it's sort of expensive for a \$2 or \$4 billion currency. Well, that assumes that the currency won't grow. That's a fundamental misunderstanding, that mining has to scale as the adoption of the currency scales. We already have world class security from the mining. We don't need a single petahash of more mining to scale Bitcoin to \$100 billion or \$1 trillion. We could run a \$1 trillion Bitcoin network on the current 400 petahashes of mining that we have right now, I don't remember the exact number.

We have bought ourselves a big vault, and it was expensive, and we opened up that giant warehouse vault and we parked a tricycle in it. A kid's tricycle. One time that vault could hold 100 Ferraris, and right now it is holding a tricycle and it looks silly, but that's okay. We can grow into it.

And this is important and some people don't understand. When we grow into it, then Bitcoin is not inefficient. If Bitcoin is supporting a national-scale currency in the order of \$100 billion or \$200 billion or even \$500 billion dollars, suddenly Bitcoin is the most efficient payment network and low-cost currency that has ever been built by man. Suddenly it is the most eco-friendly currency on the planet. The carbon footprint on a per transaction basis is minuscule. Right now, it's not. But the beauty about this is that Bitcoin is two completely independent economic systems, with only a single link-- price.

This is wrong on many theoretic levels but most importantly, it is false empirically. If proof-of-work mining ever becomes more efficient (e.g., more energy efficient chips), more hashing hardware is simply added. This is called the "Red Queen Effect" – in which all participants have to run faster just to stay in place – and Bitcoin is not immune to it. In other words, if everyone eventually got the best hashing chip, then everyone will use it and we are back to square one. And for this reason, the costs of operating the network scale directly with the value of the token.

Why? The entire threat model of Bitcoin was purposefully designed to make it purposefully expensive to attack and change votes – securing Bitcoin was "inefficient" on purpose.¹⁰² Thus if the costs of maintaining the network decline, then it also means that it is cheaper to attack the network.

Why do costs scale?

Ceteris paribus, in the long run, rational hashers will only destroy as much capital as an actual bitcoin is worth.¹⁰³ So if a bitcoin is worth \$300 they will only spend up to that point, otherwise

it would be cheaper to just buy coins from the market and turn off the machines. If a bitcoin reached \$2,000 in value, the same behavior would take place: miners would destroy as much capital to reap the seigniorage (the spread between the marginal value and marginal cost) all the way up until they are expending the equivalent of \$2,000 in energy. And so on up through \$1 million, \$1 billion or even \$1 trillion per coin.

“Efficient” chips

Why would they add more hardware even if chips are more efficient?

Each hashing farm and participant has their own cost structure that are different from one another. Due to geographical arbitrage, some locations and climates provide either cheaper energy costs, labor costs, and/or better weather for cooling the equipment. If a more energy efficient chip is created, one that is 100% more efficient than the previous one, a miner would simply double their installed hashing capacity which then effectively consumes the same amount of electricity.

In other words: the Bitcoin network is currently the world’s most expensive vault that increases in costs as the value of bitcoin appreciates – and contra Antonopoulos – the vault only becomes less expensive if and *only* if the value of the coin decreases. We see this borne out empirically as a wide array of hashing farms and mining manufacturers deploy additional capital stock during “times of plenty” or conversely go bankrupt (such as Alydian and CoinTerra) – as their margins were no longer competitive.¹⁰⁴

In contrast, government-owned mints operate such that the marginal value (face value) of fiat is greater than the marginal cost of creating the fiat ($MV > MC$). This differential, this spread is called seigniorage. Bitcoin as-is, cannot become a more “efficient” seigniorage mechanism because miners effectively bid up the marginal cost of minting bitcoins such that the marginal value equals the marginal cost ($MV = MC$) in the long run. Or rather, in the long run it costs a bitcoin to make a bitcoin.

For instance, if Bitcoin became a \$100 billion network in the future, with a 14 million coin money supply this would equate to about \$7,140 per coin. By August 2016 (when the next block reward halving occurs), this amounts to \$89,000 per block or \$6.42 million per day.

For Antonopoulos’s version of history to unfold, it would mean that other market observers, other engineers and venture funds would *not* see this ‘money minting’ opportunity. It would mean that currently sidelined miners would ignore and *not* look at the aggregate capital base – 400 petahashes/second (PH/s) deployed right now.

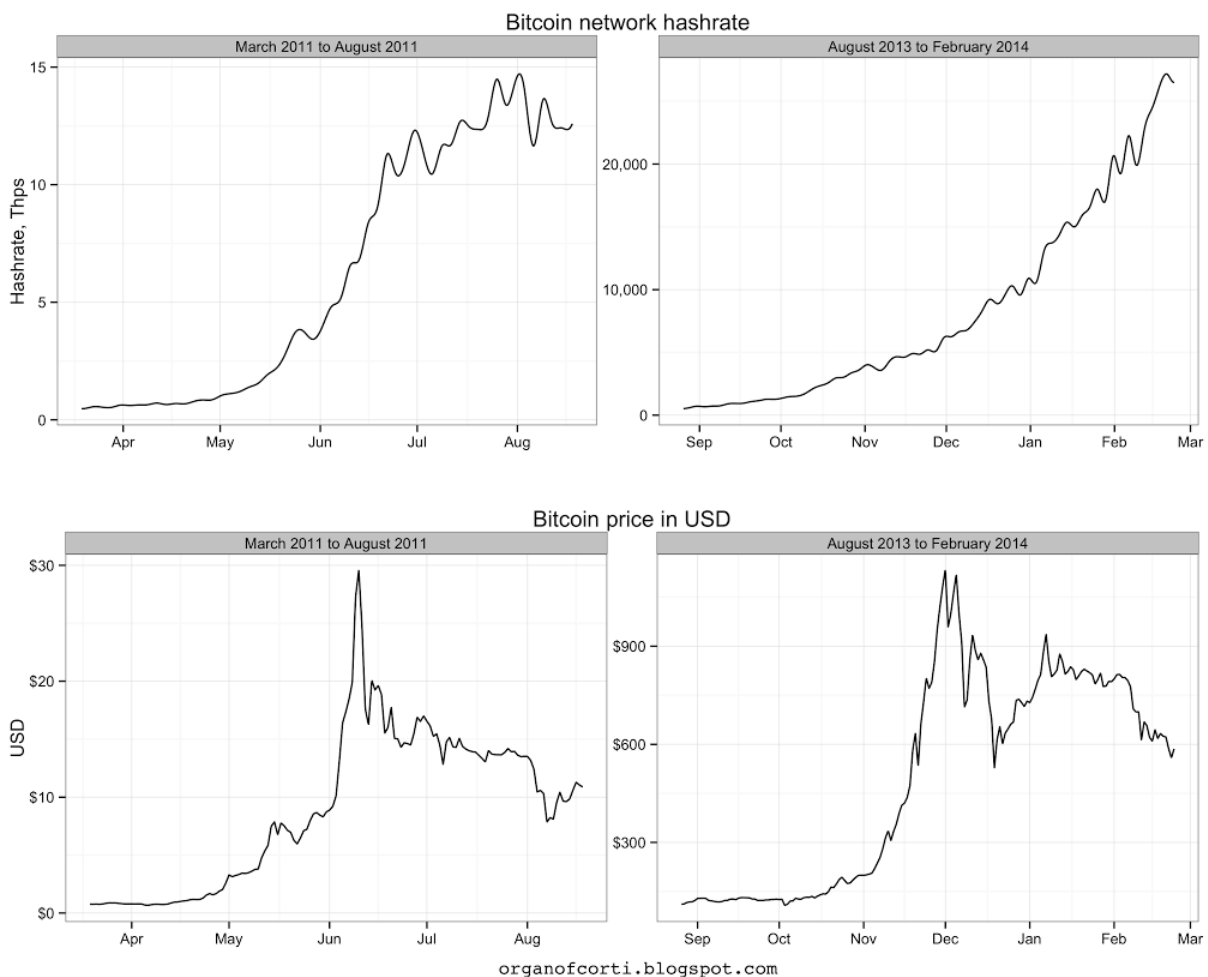
In reality, what would and does happen is that these participants would likely say, “I think I will deploy capital – 400 PH/s of my own – because that will earn me \$6.42 million per day.” Yet for every amount of hashrate added to the network, every two weeks the network adjusts the difficult rating in direct proportion to it so that there is a bi-weekly zeroing that takes place, more on that later.

We have seen this ramp-up empirically occur at least twice in the past (although some participants ramp up and down on a monthly basis).

For instance, between September 25, 2013 and January 24, 2014 the market price of bitcoin dramatically rose from \$123.42 per coin to \$825.12, an increase of 6.68x. Seeing this opportunity, between September 25, 2013 and January 24, 2014, miners added enormous hashrate, increasing the collective network hashrate by 15x.¹⁰⁵ During one month alone, October 2013, the colossal amount of ASIC mining gear that was coming online effectively “squeezed” the Poisson process – blocks were being generated at roughly every 7 minutes as opposed to the normal 10 minutes – generating an *additional* 62 blocks (1,550 bitcoins) per day.¹⁰⁶¹⁰⁷

The previous record in hashrate increase was two-and-a-half years earlier during a bubble in the spring of 2011. Between April 18, 2011 and July 19, 2011 the price rose from \$1.15 to \$15.59, an increase of 13.55x.¹⁰⁸ Simultaneously, the network hashrate increased from 661 GH/s to 11,188 GH/s, for an increase of 16.9x.

What does this look like?



As shown above, during a bull run, professional miners may be rational economic actors, bringing online additional capital as the price increases and (sometimes) reducing capital expenditures when the market value decreases.¹⁰⁹¹¹⁰ If another “black swan” event occurred in the future, by which the market value of a coin rapidly increased, it can be expected that mining power will as well. And this means increasing capital expenditures which ultimately involves dissipating exergy in some form or fashion.

What does this translate to in actual resources?

Consequently, over the past 18 months miners have been awarded approximately \$787 million in mining rewards to secure a relatively minute amount of commercial activity.¹¹¹ One estimate is that for every \$1 spent on security, roughly \$1 in actual commerce moved across the network in 2014.¹¹² While this retail commerce trend may change in the future, the fundamental costs of securing it do not.

Thus, a \$1 trillion Bitcoin network does not end up secured by what is deemed necessary at conferences, but instead by what is profitable by miners. And in practice, miners will compete to squeeze out the last slivers of that seigniorage spread until the gap between marginal value equals their marginal cost. So in theory, a \$1 trillion network will likely cost roughly \$1 trillion to run, there is no free lunch.¹¹³

Yet, a \$1 trillion Bitcoin network is not actually a \$1 trillion Bitcoin mining operation. The actual network deployment might be several order of magnitudes larger (or smaller) in terms of the actual value in transaction processing. This is beyond the scope of this report yet it is important to highlight that the \$1 trillion used in this example is the combination of the block reward and the transaction fees: neither of these have any relation to the value of what's transacted but that is also beyond the scope of this report.¹¹⁴

However, if Bitcoin or a proof-of-work-based cryptocurrency like Bitcoin did achieve Level V of Meher Roy's scale – that is to say if hyperbitcoinization took place – the network would become a proverbial monetary abyss due to its incentive mechanisms. Since the collective value of the tokens ends up equaling the amount of exergy being displaced, participants create what is dubbed “malignant computation.”¹¹⁵ Again, by design it has to be inefficient; and the larger the bitcoin market place value is, the more proof-of-work participants will eventually produce. Since this level of consumption is required then, under the most extreme scenario of hyperbitcoinization, the network would become one of the largest energy consumers on the globe as it has to continually displace capital, otherwise it becomes vulnerable to attack.¹¹⁶

Embedded chips

Can't embedded energy efficient chips flip this dynamic on its head? In August 2014 I published a book in which one chapter described an idea that has floated around from time to time: building what amounts to a legal, corporately owned botnet (in China).¹¹⁷ The main idea, is

that ASIC chips will be manufactured in bulk and installed into specific “always on” devices such as routers and distributed across China.

Why do this? For at least two reasons: at scale, the largest cost for any mining farm are the operating costs, namely electricity and administrative overhead. What if you could outsource and externalize these to consumers?

In theory, a firm will try to not only design and ship an energy efficient chip, but have it installed into devices in such a time so that the chip is still competitive when it finally reaches the homes of users. One reason this was proposed for China was that the hardest part of getting a hold of cryptocurrency in the first place are the numerous frictions involved in the foreign exchange process (e.g., passing through the KYC/AML requirements on exchanges, or buying from shady dealers on the street). So since consumers would be running the hardware anyways – so goes the theory – they not only secure the network but also receive cryptocurrency in the process.

There are challenges but these dovetail with embedded systems in general.

For example, in January 2015, IBM published a whitepaper describing ADEPT – a proof-of-work blockchain based on Ethereum that could integrate with “Internet of Things” (IoT) devices.¹¹⁸ Their goal was more focused towards securing machine-to-machine transactions but as described below, it could also introduce new hurdles that bear mentioning.

The problems for embedded devices today are not really related to monetizing data but rather on: device discovery, power management, heat dissipation or even securing the infrastructure itself. For instance, once hackers learn that there are lightweight (non-full node) cryptocurrency wallets embedded within washing machines (like the Samsung W9000 used in the IBM ADEPT proof of concept), they could create malware (such as Stuxnet or Flame) or conduct Sybil attacks on the node network itself if it was worth the cost.¹¹⁹¹²⁰

Is the hardware worth compromising?

For instance, three years ago ARM announced a line of processors that would be integrated within fridges and even light bulbs.¹²¹ And some of the new “smart appliances” shown at CES this past January have hardware that could conceivably mine for “Ether” or other altcoins; the new GE ChillHub refrigerator even integrated a Raspberry Pi.¹²² And while a typical IoT device will include an ARM Cortex-M processor running at less than 100MHz and 256 KB of RAM, script-based coins like Litecoin were mineable on ARM chips.¹²³¹²⁴

So to answer the question, based on this generation of hardware, it is unclear whether it is “worth” hacking.

This is not to say that IBM or Samsung plan to roll off Ethereum-based mining appliances, but one assumption for this IoT model is that assumes consumers are not price sensitive.

If manufacturers did install altcoin mining chips, how would consumers react? Recall that for every additional unit of hashrate that is added to the network, a corresponding amount of difficulty is also added such that there is continual counterbalancing zeroing effect. What this means is that even if these ARM chips are relatively energy efficient, the gain is temporary as the network rebalances and mining is no longer profitable.

And how does this intersect with the idea from China? In contrast to general purpose computing of ARM chips, ASICs by design have one sole function and ASIC miners in particular are depreciating capital goods that become “economic dead weight” once a certain difficulty rating is reached. Thus at some point, users in China would be spending more on electricity than they will generate in bitcoins (e.g., mining at a loss).

A second challenge, as noted by Stuxnet and Flame above, is that of security vulnerabilities. As we have seen with “traditional” bitcoin mining, hackers are willing and able to exploit vulnerabilities in internet infrastructure. One example in the spring of 2014, a hacker used a spoofed Border Gateway Protocol (BGP) to redirect mining pools to a hacker-controlled router in Canada.¹²⁵ Over the course of several months the hacker proceeded to siphon off what amounted to \$90,000 worth of cryptocurrency including bitcoin (there were several other virtual currencies being mined on the pools).

In practice, if a hacker learns that hardware designed for mining “ether” or other altcoins is embedded in numerous devices, they may be incentivized to once again externalize the operating costs onto the consumer as they previously did during the CPU and GPU-era of bitcoin mining which was plagued by botnets (e.g., hijacking computers via malware distributed by illicit botnet operators).¹²⁶

One last constraint to consider with ADEPT from IBM and Samsung as a proof of concept, is that in practice Bob, a retail customer, probably does not need a proof-of-work-based blockchain such as Ethereum to act as a “payment rail” between himself and Alice, the utility owner, because all parties are probably already known.

If IoT is really just about washing machine-to-washing machine payments and specifically arbitraging energy prices, a manufacturer could just as easily build into an appliance a ‘secure element’ that stores a prepaid debit amount and connects via existing payment rails.¹²⁷ Furthermore, it is difficult to see the value proposition for the logistics: to go from fiat to cryptocurrency and then back into fiat. This conversion process is an unnecessary friction and probably a solution looking for a currently non-existent problem for IoT. In addition, distributed ledgers such as Eris, Hyperledger and Ripple could also conceivably provide a similar cost-effective method as Alice does not need to use expensive proof-of-work to do that.

It bears mentioning that the explanation above is solely related to the proof of concept and should not reflect upon what either IBM or Samsung may use in the future. Furthermore, a

proof-of-work blockchain could still be used as a type of secure data store (e.g., proof-of-existence) or timestamping mechanism for IoT.

In addition, IoT may not be so much about inter-device payments as much as it could be about a standardized and distributed way of doing event driven activities across devices.¹²⁸ For instance, if Bob's alarm clock rings (e.g., an event) – Bob's shower responds by heating up the water; once heated (e.g., an event) – then something else happens and so on down the line. Yet it is important to distinguish between “automation” and secure computational contracting systems.¹²⁹ A user may not need complex Ethereum-style contracts for this kind of home automation but rather demand traditional automation system that do not involve vendor or cloud lock-in.¹³⁰

Closing remarks on mining

Recall that in trusted networks, operators will find and replace “inefficient” systems to reduce the operating costs to a bare minimum – even going as far as replacing AC-based power supplies with DC-based.¹³¹ In Bitcoin, the opposite actually occurs: mining farms intentionally try to add as much capital outlays in an effort to chase after the fixed block reward.¹³²

In the short-run any number of factors (difficulty rating drop, bitcoin value increased, cheaper energy rates) can contribute to the success (or failure) of individual mining efforts. Some farms in Sweden, Finland, Washington state and China are capable of generating what amounts to “bumper” coins due to their efficient scaling efforts, cheaper energy and cooler climates.

Similarly, as in all manufacturing segments, some mining manufacturers will likely be able to design and integrate chips more competitively through faster logistical turnarounds than others in this space. Thus despite the concerns and challenges raised above, a number of ventures will still likely be able to profit in the short-run.

Yet each jurisdiction has different incentives (and disincentives) for miners. In China, for example, there are several factors that have led to its ascent in aggregate hashing – as of this writing the top four global pools reside in China and collectively provide 53% of the network hashrate.¹³³

Why is this the case? After interviewing several China-based miners, there are at least five reasons:

1. Sunken costs: mining farms owners put the money down a while ago (sunk costs) and the capacity is finally coming online; owe debts so need to generate revenue even if it is at a loss.
2. Converting RMB to USD at no limitations (e.g., capital controls)
 - Individuals are limited by law to the equivalent of \$50,000 per year in cross border transfers

- Miners may do this even at a loss because it may be cheaper than converting RMB to USD
3. Believe the price of coins will go up, “but there won't be any more coins”
 - “Fear of missing out” (FOMO)
 - In their view, it makes some sense due to lack of transparency and financial controls at China-based exchanges, and doesn't leave paper trail to authorities
 4. Tax reasons: Bob can justify buying a bunch of computer “related parts” and report this without a problem to his boss or the government, but Bob can't receive permission to directly buy bitcoins
 5. Relatively cheap land / labor and the factories assembling most of the mining systems themselves are located in China, giving Chinese miners advantage in terms of lead-time

Whether or not these factors continue into the future is unclear due to the legal uncertainty of Bitcoin's status on the mainland.

There are a number of other topics related to mining and incentives that could be explored in future research but were beyond the scope of this report. For instance, according to Vitalik Buterin, creator of Ethereum:¹³⁴

It's worth differentiating between three levels of cryptocurrency token inclusion:

1. no network token at all
2. network token exists, but its use is primarily for transaction fees and protocol-required deposits and is otherwise not heavily emphasized
3. the network token is The Whole Point™

Anything involving a permissionless set of actors essentially requires game-theoretic arguments to show security, and Bitcoin's game theory is far from perfect.¹³⁵ I think there's a fundamental inclusiveness-trust tradeoff frontier in many places in society, and we can view the different categories of crypto projects as trying to advance different parts of the frontier.

How mining actually works in practice is a very intriguing and an arguably underappreciated aspect of cryptocurrency systems and their progeny. Yet the underlying incentives and punishments for certain mining activity has not changed the fundamental laws of economics. Thus, this is an open research space that is likely ripe for modeling by academic economic and financial professionals.

Endnotes

¹ While there have been numerous “proof-of-stake” (POS) systems created, as of this writing, in practice none of them have remained fully decentralized. Some novel attempts include delegated proof-of-stake (DPOS) from Invictus (it uses 101 “delegates,” which are the same as validators), [Tendermint](#), Tezos and [NeuCoin](#). Tezos is based on a Tezos-like blockchain which is consensus agnostic. Its [white paper](#) described a proof-of-stake algorithm, but the consensus mechanism can be dynamically picked by the stakeholders. Another novel idea that is being integrated into some of these POS systems is similar to [fidelity bonding](#).

² As for dollar value of “coins” – that may only matter in proof-of-work and proof-of-stake systems. It may not matter at all in a permissioned system where the coins may become irrelevant if antispam is not an issue. Depending on the customer, it could be an advantage or disadvantage – decision makers at financial service providers will likely look on a case-by-case basis.

³ Another fundamental difference between Bitcoin and distributed ledger systems is that security model is different (as described by the “permissioned” nomenclature): security of users’ private keys are crucial in Bitcoin, but in distributed ledger-based systems Bob can manage risk of theft via the freezing and refunding of funds.

⁴ See [How Bitcoin Hashing Works](#) and [On Mining](#) by Vitalik Buterin

⁵ [Enabling Blockchain Innovations with Pegged Sidechains](#) by Back, *et al.*

⁶ KYB stands for Know-Your-Business and KYC stands for Know-Your-Customer. The topic of a “legal entity identifier” (LEI) is beyond the scope of this report.

⁷ In theory, Ripple would meet the criteria for “permissionless.” In practice however, they currently whitelist and blacklist validators so it is technically *permissioned* as are the other projects on top of it (as there is no defined KYC/KYB procedure (yet) the process is *de facto* because a node is only listened to if other nodes trust it, the process requires a public identity).

⁸ See [Intel Hints at Bitcoin Play With Crypto Researcher Hire](#) from *CoinDesk* and [Researchers Biometrics and Crypto Currency & Value Chains](#) from ING

⁹ [Nonoutsourcable Scratch-Off Puzzles to Discourage Bitcoin Mining Coalitions](#) by Miller, *et al.*

¹⁰ Launched in December 2010, Slush’s pool was the first public bitcoin mining pool and is still in existence today, as [BitcoinCZ](#).

¹¹ While there may be a cost if the attack succeeds, in some circumstances it may also cost nothing for the attacker because they could collect the mining rewards. There are a variety of different hypothetical situations, but again, this is beyond the scope of this report. See [Downplaying statistically possible double-spending risks](#)

¹² [Can Bitcoin’s internal economy securely grow relative to its outputs?](#) by Tim Swanson

¹³ [The “Unbundling of Trust”: how to identify good cryptocurrency opportunities?](#) by Richard Brown

¹⁴ [EBA Opinion on ‘virtual currencies’](#) from the European Banking Authority. I would like to thank Greg Simon for bringing this to “source of funds” KYC issue to my attention last summer. See also: [Coinbase Seeks ‘Invasive’ Details on US Bitcoin Mining Operations](#) from *CoinDesk*

¹⁵ Vlad Zamfir has created two overviews of this new field: [Formalizing Cryptoeconomics](#) and [What is Cryptoeconomics?](#)

¹⁶ Some of these projects, such as CryptoCorp and SKUChain do utilize “proof-of-work” even though all parties are known, but these are private or semi-private networks and likely will not face the exact same issues as Bitcoin does today with regard to block rewards and transaction fees.

¹⁷ Illustration from [Near Zero Bitcoin Transaction Fees Cannot Last Forever](#) by Kerem Kaskaloglu

¹⁸ I would like to thank Vitalik Buterin for this insight.

¹⁹ See [Creating a decentralised payment network: A study of Bitcoin](#) by Jonathan Levin and [The Bitcoin mining game](#) by Nicolas Houy

²⁰ I would like to thank Peter Todd for this insight.

²¹ See [“The Law of Bitcoin”](#) edited by Stuart Hoegner (forthcoming)

²² If Cathy uses Coinbase to send a bitcoin to her friend David he uses Xapo, because both users are identified, it removes some of the core advantages of a permissionless blockchain (pseudonymity) and adds new costs (KYC/AML). In the long run, as fees to miners are *supposed* to increase, it is unclear what the advantage is to use the network in this manner.

²³ Empirically, perhaps the only caveat is delegated proof-of-work (DPOW) from SKUChain, but this uses trusted nodes and SKUChain white-lists the miners.

²⁴ Employing a public key infrastructure (PKI) system for documentation would be a step in a more secure direction. GPG signatures or digital signatures from smart cards on credit/debit card transactions would make the system security robust against attack from “carders” and potential wire fraud or check fraud.

²⁵ [A simple model to make sense of the proliferation of distributed ledger, smart contract and cryptocurrency projects](#) by Richard Brown

²⁶ Personal correspondence March 17, 2015

²⁷ One reviewer explained that, “Alternatively: any data/state with counterparty risk has its own centralized blockchain. All data/state between blockchains has relative liquid trading value. This means that counterparty risk can be accounted for based on relative market value of the underlying consensus/state tokens. A blockchain does not need to be a shared ledger, nor does it need to have a distributed consensus. It can be completely centralized as long as its data/state is externally verifiable and all data is immutable.”

²⁸ Stellar, originally a fork of Ripple, was not included in this report due to the fact that its consensus currently relies on one validator. However, it is expected that David Mazieres, chief scientist at Stellar Development Foundation, will announce and publish a new consensus protocol later in April 2015. See also, [Internet-Level Agreement with the Stellar Consensus Protocol, Why the Stellar Forking Issue Does Not Affect Ripple](#) by Stefan Thomas and [Safety, liveness and fault tolerance—the consensus choices](#) by Joyce Kim

²⁹ [A model to make sense of beliefs and associated Crypto-finance platforms](#) by Meher Roy

³⁰ See the glossary in [Great Chain of Numbers](#)

³¹ [A Simple Model for Smart Contracts](#) by Richard Brown

³² According to Robert Sams, “this is actually a very subtle issue. We may be unable to say that a smart contract “takes custody” over assets on the distributed ledger without re-defining what the term custody means. In law custody implies some that there is some legally accountable party who holds the assets in “safe-keeping,” thus how could a smart contract be that party – at least if it’s a consensus-computed smart contract? On the other hand, if the smart contract is computed centrally on someone’s server, then whoever is running the machine will be a custodian in the eyes of the law; the owner of the machine is the party taking custody. And this is one of the key arguments as to why the industry needs a Turing-complete blockchain (as opposed to oracles), because smart contracts computed by regulated custodians is just custodial software (e.g. Sunguard has been building these for years). Consensus-computation is the key innovation here as far as the data-layer use-case is concerned. But smart contracts don’t “take custody” – custodians do.”

³³ It is also recommended that interested readers peruse [Formalizing and Securing Relationships on Public Networks](#) by Nick Szabo

³⁴ The first person I am aware of that used the term “consensus-as-a-service” is Juan Benet-Bitez, founder of [IPFS](#) and [Filecoin](#).

³⁵ Special thanks to Jo Lang and the [R3 CEV](#) team for providing this.

³⁶ One notable, superfluous example is Joel Monegro’s “[Bitcoin Appstack](#).” In both theory and practice, Bitcoin is *unnneeded* in that “bottom” layer and it is unclear why it is used since specific neutral protocols could be designed to do the same function just like real protocols have been developed in the past.

³⁷ Dan O’Prey, co-founder of Hyperledger, explains this conundrum: “Whenever SPAM comes up, people seem to be quick to jump to the easy solution of “if we charge a little bit then it will go away” without thinking of the damaging consequences this could have to adoption and legitimate use-cases. Let’s take email for example, charging 1c an email would be pretty negligible, but what happens to newsletters? Alert systems? Notifications? Now imagine it’s not 1 cent but 1 emailcoin. Would email have overcome the early barriers to adoption to become the universal communication system it is if you had to purchase a new currency just to use it? The solution to SPAM is better software; blacklisting, greylisting, rate limits, etc. Gmail is free but I get 0% SPAM in my inbox. Snail mail costs a stamp but nearly 100% percent of my letterbox is unsolicited.”

³⁸ As of this writing there are 544 currencies and 51 assets listed on [Coinmarketcap](#). See also [Map of Coins](#) and [Coingecko](#)

³⁹ One caveat is that TCP in practice does have gatekeepers (ISPs) that a user needs permission from to access the wider internet for it to be more useful. Yet in practice, Bitcoin has trended towards this direction too, with hosted wallets and a cornucopia of trusted intermediaries on the edges.

⁴⁰ There are several “public goods” challenges for Bitcoin, this specific one was covered on p. 34 in [The Anatomy of a Money-like Informational Commodity](#) by Tim Swanson.

⁴¹ [Approximately 70% of all bitcoins have not moved in 6 or more months](#) by Tim Swanson

⁴² Another imperfect analogy (though very germane) is that Bob Kahn and Vint Cerf (“Fathers of the internet”) did not “mine” 1 million internet packets back in the mid-1970s, holding with the hopes that eventually they could buy yachts and islands with their packets. Following this admittedly imperfect analogy (since packets cannot be mined), the fact that Satoshi Nakamoto [allegedly](#) mined 1 million bitcoins has provoked a number of views. One notion includes the observation that Nakamoto only advertised Bitcoin on one obscure mailing list and then preceded to mine basically for an entire year without advertising it again and without doing any effort at all to do public relations activities. He could have run a testnet. He could have also said: “Mulligan, hey, we now have an active community, it’s been a year now, so we are going to reset it.”

⁴³ Incidentally this also happens to Namecoin (a decentralized blockchain-based DNS network), but that is beyond the scope of this as well. See [On Domain Name Squatters in Namecoin](#) by Indolering. Regarding the ownership of domain names see [Kremen v Cohen](#) 9th Circuit Court of Appeals

⁴⁴ [Can Bitcoin’s internal economy securely grow relative to its outputs?](#) and [Will colored coin extensibility throw a wrench into the automated information security costs of Bitcoin?](#) by Tim Swanson

⁴⁵ See [An architecture for the Internet of Money](#) by Meher Roy and Beyond Bitcoin, episode #27: [An Architecture For The Internet Of Money](#) at *Let’s Talk Bitcoin*.

⁴⁶ See [The Subjectivity / Exploitability Tradeoff](#) and [Proof of Stake: How I Learned to Love Weak Subjectivity](#) by Vitalik Buterin. See also the [Beyond Bitcoin Show MUSIC DEV Hangout interview](#) with Taulant Ramabaja (starting at 18:22m).

⁴⁷ See [Fedcoin](#) by JP Koning, [Fedcoin: On the Desirability of a Government Cryptocurrency](#) by David Andolfatto, [A Central Bank “cryptocurrency”? An interesting idea, but maybe not for the reason we think](#) by Richard Brown and [Which Fedcoin?](#) by Robert Sams

⁴⁸ [Wet code and dry](#) by Nick Szabo

⁴⁹ Vlad Zamfir’s definition of a distributed cryptoeconomic consensus protocol is three programs: 1) A fork-choice rule; 2) A consensus incentive mechanism; 3) A consensus strategy. In his words, “Everyone uses the same fork-choice rule, the incentive mechanism distributes payoffs by altering the state of a ledger, and the strategy at least games the incentive mechanism. My formalization is treating distributed consensus generally, rather than assuming that consensus is on a ledger.”

⁵⁰ [Soylent Blockchains](#) by Steve Waldman

⁵¹ One common strawman that some Bitcoin advocates use to portray permissioned distributed ledgers: it is a centralized Bitcoin blockchain that records exchanges of metacoins. Yet, if that is truly being built then a user could just as well use a database with PGP signed inserts. A blockchain data structure does not add anything on this model. If custodians are not involved, then this is just a shared database, not a registry of securities titles. If this is to be a prototype of a “settlement” solution, it has to show how the entries get their legal status. Cryptocurrency systems have no (current) ability to handle or manage legal, real-world assets. As an aside, Alex Mizrahi (“killerstorm”) independently [described](#) a similar definition. In his word, “Blockchain essentially boils down to this: full audit trail, available to everyone; cryptographic integrity checks; full audit trail is independently verified on each node. Why not just use a conventional database? Maybe because they don’t have cryptographic integrity checks, aren’t Byzantine fault tolerant, etc. Note that a blockchain doesn’t need to be based on PoW-based consensus (mining). Essentially you can any BFT consensus mechanism to gain features mentioned above. [...] One major advantage over “conventional databases” is that the system can keep working even if 40 bank nodes are down or were hijacked. You can’t do that with using Oracle DBMS, can you? Some DMBSs offer multi-master replication, but they can’t deal with Byzantine failures as they don’t have a thorough cryptographic verification. (SSL is not enough.) What else? They can offer some advanced features like conditional transfers and fully automated escrow at no extra cost. Why? You can certainly implement all kinds of advanced features on top of conventional databases, but then you need to deal with race conditions. Which is kinda hard in a distributed environment (i.e. each bank has its own database).”

⁵² Every CHAPS payment is unconditional, irrevocable and guaranteed, see [Accuity resource center](#). See also [A Guide to the Bank of England’s Real Time Gross Settlement System](#) and [Payment systems in the euro area](#) from the Bank for International Settlements:

Regarding electronic money institutions (ELMIs), the European Commission has interpreted that they are not covered by the definition of credit institutions in the [Settlement Finality Directive] SFD. Consequently, ELMIs would not be “institutions” within the meaning of the SFD and could not be regarded as “participants” in systems designated under the SFD. As a result, the participation of ELMIs in TARGET would not be advisable, as this could jeopardise the irrevocability and finality of payments processed through the system. Therefore, the Governing Council has decided that ELMIs are not eligible participants in TARGET as long as uncertainties remain with regard to their protection under the SFD.

⁵³ Nothing short of a court order changes that. While state adjudication of disputes will continue to take place, it does not follow that every participant running the payment system should have the power to reverse transactions for the same reason that repo men should not have unfettered access to Bob’s home when he gets behind on his credit card bills.

⁵⁴ One common misconception in the Bitcoin community is that that reversals are common in disputes such as in “flash crashes” or credit card systems. However, in the case of credit cards, this is a false equivalency. Credit card confirmations take place almost instantly and in the event that a dispute arises, require a second transaction to effectively cancel out the first (the same logistical occurrence happens with other RTGS). And exchanges reversing trades are reversals of trades *pre-settlement*. Recall that trading has two parts: the deal (via exchange or dealer) and subsequently, settlement. The deal is just a contract and can be nullified for a variety of reasons. Settlement is transfer of property. For publically traded assets, a trade is not reversible post-settlement in the way that, say, the purchase of a stolen car is, or an asset sold where seller did not disclose to buyer some lien on the property. This is enshrined in UCC article 8 and 9 in the United States and Settlement Finality Directive in the European Union. In point of fact, global capital markets really cannot work without finality. Post-Lehman Brothers, there is a global initiative to clarify it even further so that bankruptcy events do not inject uncertainty into the post-trade space. A distributed ledger solution that does not conform to this legal concept will have limited utility.

⁵⁵ I would like to thank Robert Sams and the Clearmatics team for clarifying this.

⁵⁶ There are a number of effective distributed database platforms including Hadoop, Chubby and HyperDex and others from large vendors such as IBM and Oracle who continually update and implement new features for their own platforms. Similarly both Amazon and Google (though AWS and Compute, respectively) provide distributed database services to customers. These are solutions to other existing problems. There are several reasons why a distributed ledger would be adopted over a distributed database for settlement and clearing of financial instruments but those are beyond the scope of this report which was comparing distributed ledgers with cryptocurrency systems.

⁵⁷ One reviewer suggested that, “In point of fact, a database can be turned into a blockchain. To do so it needs: a starting state as a genesis block, all data and all changes to be immutable, all changes to be appended with their corresponding hashes as well as all previous hashes (potentially a Merkle tree), changes saved only as Deltas, not as deletions to the old state. So with a secondary business logic layer and a bit of simple cryptography, a database can become a blockchain.”

⁵⁸ [Blockchain Finance](#) by Robert Sams

⁵⁹ For example: in January 2015, Ripple Labs released a new open source database called [NuDB](#) that only has insert and read functions instead of insert, read, update, delete. Even at the code level a user cannot make changes to the database.

⁶⁰ There is a slightly different balancing act involving finality within the literature regarding “[safety-liveness tradeoff](#)” and [CAP theorem](#).

⁶¹ As Robert Sams hypothesizes: “Whatever proof-of-work or proof-of-stake system you have, there will exist another system design that will protect against reversal better if you have transparent validators subject to reputation and, perhaps, legal recourse. [...] A consensus network based on authenticating the real identities of validators is an alternative to the PoW/PoS solution to the Sybil problem and could go much further in mitigating #3. But this will almost certainly come at the price of censoring transactions, as validators will be held accountable for the content of the transactions they verify, at least if the value being transferred over the network is off-chain property.”

⁶² It should be noted that while Ripple as a concept originated in 2004 from [Ryan Fugger](#), for the purposes of this report, Ripple is synonymous with the Ripple network and Ripple Labs, its current corporate sponsor.

⁶³ [Soylent Blockchains](#) by Steve Waldman

⁶⁴ See [What is the blockchain hard fork “missile crisis?”](#) and [Proof of Activity: Extending Bitcoin's Proof of Work via Proof of Stake](#) from Bentov, *et. al.* In addition, one of the continual challenges within Bitcoin currently is the principal-agent problem. There are a number of special interest groups such as miners, payment processors, hosted wallets and even “core” developers themselves that create an agency dilemma. With discussions of increasing block sizes and sidechains, there is arguably a conflict of interest between the principals (users of the network) and the agents (whose organizations they work for benefit from certain changes in code). Because it is a “public good” and the most common way of “resolving” such conflicts typically involve long debates on social media, it is unclear at this time how these changes will be resolved.

⁶⁵ Pseudonymous, permissionless systems can, in theory, manage on-blockchain, bearer instruments as long as traditional institutions merely tolerate or ignore them. Management of off-blockchain institutions requires active endorsement of traditional institutions, which is less likely to be extended to permissionless systems. Based on Appendix A of the Blockstream [whitepaper](#), the team could also use federated pegs – which in theory could also include trusted validators – but as of this writing it is unknown if this will be attempted. If trusted validators are used then Blockstream could be included in the list of distributed ledger projects.

⁶⁶ See: Greg Maxwell’s [detailed explanation](#) of the March 2013 event as well as a contemporary write-up, [Bitcoin Network Shaken by Blockchain Fork](#) by Vitalik Buterin. One alternative viewpoint, from the social engineering perspective, [comes from](#) Patrick McKenzie. During this time frame, merchants were asked to not process transactions until the chains converged, see [Alert: chain fork caused by pre-0.8 clients dealing badly with large blocks](#)

⁶⁷ I would like to thank Steve Waldman for this insight.

⁶⁸ Over the past several months, various entrepreneurs and designers including Vitalik Buterin, Peter Todd and Preston Byrne have proposed a simpler mental model of blockchains: that of a database. Robert Sams has contributed to this zeitgeist, explaining: “As a ledger is just persistent storage, and ledger entries are usually dictated by agreements, I’m not so bullish on consensus networks that are not Turing-complete, at least in financial contexts. Here, machine state rather than “ledger” is what should be the ground floor of the technology stack. Real-world ledgers are usually inseparable from agreements (think nominee/beneficial owner, hypothecation, etc.). Partitioning that complexity into a ledger layer and separate contract layer is inelegant and complicated in implementation.” Another reviewer suggested that, “Distributed ledger systems are the next step of financial IT infrastructure. It is only question when it will happen, since it is very costly to modify something that works and well tested. The question whether they will be a *public* distributed ledger.”

⁶⁹ I would like to thank Robert Sams for providing this example.

⁷⁰ David Chaum pioneered the field of “ecash” with the creation of DigiCash in 1990. See [A Quick History of Cryptocurrencies BBTC — Before Bitcoin](#) by Ken Griffith

⁷¹ One lawyer who reviewed this report thought that courts will recognize systems like Bitcoin as somewhat flawed and possibly incapable of achieving negotiability, but still capable of transferring value. See also [“The Law of Bitcoin”](#) edited by Stuart Hoegner (forthcoming)

⁷² Depository Trust Company (DTC) is just one of many institutions maintaining a ledger with assets. See [The solution to Wall Street’s 1960s paperwork crisis could also save bitcoin](#) from Quartz and [Bitcoin’s lien problem](#) from *Financial Times*

⁷³ I would like to thank Taulant Ramabaja for this insight.

⁷⁴ For instance, in the US, the SEC adopted Rule 613 to build a Consolidated Audit Trail (CAT). One reviewer suggested that, “CAT is about a very low latency, fragmented market - its original motives might not be a very good fit for this. In general financial services care more about capital management - ability to track mobility for value.” See [CAT NMS Plan](#)

⁷⁵ See his [tweet](#) from February 23, 2015.

⁷⁶ This view is called “Bitcoin maximalism.” See [On Bitcoin Maximalism, and Currency and Platform Network Effects](#) by Vitalik Buterin and [Sidechains without Pegging](#) by Jae Kwon

⁷⁷ There have been multiple proposals for “sidechains,” with the effort led by Blockstream as perhaps the most well-known. In Appendix A of their [whitepaper](#) the team discusses the use of a “federated pegs” as an intermediary solution between now and if/when a “hard fork” into Bitcoin core takes place. In a recent [interview](#) with *Epicenter Bitcoin*, two of Blockstream’s developers, Greg Maxwell and Adam Back, also briefly discuss federated pegs. In theory, these federated servers could also probably be used as trusted oracles or even trusted validating nodes for independent chains and perhaps even distributed ledger systems as well.

⁷⁸ The mental model of equating securities with IOU's is likely unsupportable from a legal perspective as well. In the cryptocurrency ecosystem, there has been a lack of treatment and analysis of custody and regulations around asset liability, responsibility that would be impacted and how, by use of ledger technology. This should be at the core of an analysis from the beginning because it is going to be a real challenge to try to retrofit it later. For instance, non-compliant (as in lacking KYC/AML) metacoins could end up trading at a premium to the compliant ones (e.g., they become a popular medium-of-exchange for illicit trade), encouraging everyone to jump off the issuer's identity scheme. What legal implications would that hold for the issuer? The issuer could get shut down but it is the very expectation of that and the resulting scarcity of the metacoin that causes them to trade at a premium. While speculative, AML authorities will not look kindly on the issuer in either scenario. *Ex ante* control by issuer is hard to avoid on issuer-IOU models. Or as one reviewer explained, "In finance, most securities that clear through normal channels are final after a certain point. And as titles change hands with high velocity, reversal would require breaking an entire chain of transactions and the very possibility of that creates intolerable uncertainty for financial markets which is why it does not happen. And it is why there is currently no central administrator with the power to do such things; nor is there a central database, but rather many distributed among CSDs, custodians and clearing houses. The idea of a central administrator who can amend a central database of securities titles arguably ranks up there with the quizzical notion that securities titles can exist on a proof-of-work blockchain."

⁷⁹ Modern day finance traces part of its lineage to the Medici Bank. See also, [The Ascent of Money](#) by Niall Ferguson

⁸⁰ A lawyer that reviewed this manuscript explained: "Perhaps the only "absolute" bearer asset is cash itself, yet even then it is not quite absolute (e.g., *nemo dat*). In contrast, there is no "legally official" registry. The job of property systems is to associate the who's with the what's. There is no infallible magic bullet. It is merely a question of best evidence. So, possession/control is a pretty crude form of evidence but often nobody has better evidence. Registration is pretty good evidence but it can still be overcome. Think about a piece of artwork that Bob consigns to a gallery or that he registers. Or a title to his house. No matter what the title search says, Bob can never really know somebody won't come out of the woodwork with better evidence of ownership. The question is really: how much protection does the law provide to an innocent purchaser for a particular type of property in a particular situation? This is still an open question with bitcoin."

⁸¹ Tabulating publicly reported bitcoins that were lost, stolen, seized, scammed and accidentally destroyed between August 2010 and March 2014 amounts to 966,531 bitcoins. See p. 196 in [The Anatomy of a Money-like Informational Commodity](#) by Tim Swanson. See also: [Bitcoin Self-Defense, Part I: Wallet Protection](#) by Vitalik Buterin

⁸² The inability to enforce a contract and retrieve losses in the event of fraud is not just a challenge for Bitcoin, but other cryptocurrency systems such as Dogecoin. For instance, Dogeparty asset "[DOGEDIGGERS](#)" was used by someone mid-November 2014 to sell shares in their "mining operation." The individual(s) behind it managed to extract a few million dogecoin before people caught on and started asking questions, identifying it as a scam and put an end to it -- the social media sites that the scammers were using to make the scam look legitimate were [taken down](#). Restitution, if there is any, will take place off-chain where contract enforcement actually exists. See also [Meet Moolah, the company that has Dogecoin by the collar](#) from *The Daily Dot*

⁸³ As of this writing, it is still unclear, whether the BitLicense is establishing a facilitator regime (like money transmission or custody) or an intermediary regime (like deposit taking). Does the BitLicense permit the acceptance of deposits by licensees? If not, then the question remains whether organizations like NYDFS and DOJ considers hosted wallet services to constitute deposit taking. If so, BitLicensees would presumably not be able to avail themselves to the securities exemption that is available to banks and other deposit takers. A deposit is a debt owed by the depository to the customer (depositor). Does holding oneself out as a depository qualify as a securities offering? If so, would licensees qualify for the bank exemption to the securities laws? See [Distributed Oversight: Custodians and Intermediaries](#)

⁸⁴ [The Last Straw for Bitcoin](#) by Ryan Straus

⁸⁵ Satoshi Nakamoto announced the release of both the [white paper](#) and [code base](#) on the Metzdown mailing list. This lends to the theory that Nakamoto was likely a member of the cryptographic/cypherpunk community and not an economist or financial professional.

⁸⁶ In contrast, what we have today is "Bitcoin In Name Only" (BINO). See [What we have today is not Bitcoin but BINO](#) and [Moving Beyond BINO Beta](#)

⁸⁷ I would like to thank Steve Waldman for highlighting this point. Furthermore, predicting a time frame as measured in months or years (e.g. by July 2020) for the various changes discussed that could occur is difficult due to a long, uncertain sales cycle (e.g., education, compliance, testing).

⁸⁸ One common question is: is it fair to compare systems designed for existing financial institutions with systems that were designed with the hope that unidentified end users, globally, will adopt the platform? This fundamentally boils down to what Gabriel Weinburg calls “[the product trap](#)” and is also referred to as “[marketing myopia](#)” which states that “businesses will do better in the end if they concentrate on meeting customers’ needs rather than on selling products.” Most entrepreneurial activity in the cryptocurrency space has thus far been focused on creating supply irrespective of what consumers actually need or demand. In contrast, permissioned distributed ledgers are designed with specific customer requirements in mind. See also: [Cryptoeconomics for beginners and experts alike](#) and [Is the adoption of blockchains and consensus ledgers a foregone conclusion?](#) by Tim Swanson

⁸⁹ Federated pegs can exist as both trusted and semi-trusted nodes and consequently could, in theory, operate independently within the distributed ledger category.

⁹⁰ I would like to thank Robert Sams for this key insight. See p. 76 in [The Anatomy of a Money-like Informational Commodity](#) by Tim Swanson

⁹¹ I believe it was Kerem Kaskaloglu who first labeled them “donations” in [Near Zero Bitcoin Transaction Fees Cannot Last Forever](#). See also [The Collective Action Problem of Mining Fees](#) by Tim Swanson

⁹² Vitalik Buterin noted that while “this is true technically, but in many protocols not economically (eg. in Ethereum gas limits will be very often whacked against and so a zero-fee transaction will just not be accepted by miners because a miner that does accept it would be paying opportunity cost).” In addition he suggested that future research could look into “the necessity of transaction fees, particularly for ledgers where anyone can become a user, they are necessary, because otherwise you fall to DDoS. But for ledgers where the set of users is kept tightly closed, you can have no fees, and manually detect and deal with “abuse” or impose per-user limits.”

⁹³ Dave Hudson has several good articles explaining the incentive to pool capital as it increases certainty and lowers the risk. That is to say, because of the Poisson process (variance), the more hashrate a pool has, the higher the probability of a payout. See [The Gambler’s Guide To Bitcoin Mining](#).

⁹⁴ [After the Bitcoin Gold Rush](#) from *The New Republic*

⁹⁵ [The myth of a cheaper Bitcoin network: a note about transaction processing, currency conversion and Bitcoinland](#) by Tim Swanson

⁹⁶ In some sense these pools act as *ad hoc* trusted nodes within the network. That is to say, since these pools (and many of the farms underlying them) are known – their management and location are known – there is little reason to use the level of proof-of-work that they are currently doing. A case could be made that because of this trend towards an expensive oligopoly, non-mining participants could just as easily and securely use any of the other 8 projects discussed. See [\[ANN\] High-speed Bitcoin Relay Network](#) by Matt Corallo and [The Future of Bitcoin: Corporate Mines and Network Peering?](#) from *Data Center Knowledge*

⁹⁷ For instance, a “block” – which in the Bitcoin world is a bundle of transaction outputs – is for all intents and purposes the equivalent to “batching” in traditional finance.

⁹⁸ See [The Marginal Cost of Cryptocurrency](#) and [Some Crypto Quibbles with Threadneedle Street](#) by Robert Sams

⁹⁹ [Economic Aspects of Bitcoin and Other Decentralized Public-Ledger Currency Platforms](#) by David Evans

¹⁰⁰ [Why Bitcoin does not have a market cap](#) by Jonathan Levin

¹⁰¹ See the transcript of his talk: “[Decentralization through game theory](#)”

¹⁰² [Cost? Trust? Something else? What’s the killer-app for Block Chain Technology?](#) by Richard Brown

¹⁰³ There are five exceptions this described on p.46 in [The Anatomy of a Money-like Informational Commodity](#) by Tim Swanson

¹⁰⁴ [Bankrupt Bitcoin Mining Company Alydian to Sell 218TH/s of Mining Power](#) from *CoinDesk* and [Bitcoin Mining Firm CoinTerra Files for Chapter 7 Bankruptcy](#) from *CoinDesk*

¹⁰⁵ See [Difficulty History](#) from Bitcoin Wisdom. See also: [Finding 2016 Blocks](#) and [Hash Rate Headaches](#) by Dave Hudson

¹⁰⁶ I would like to thank Kevin Zhou of Buttercoin for pointing this out.

¹⁰⁷ One reviewer explained this phenomenon as potentially one-off: “Prior to 2013 there had been a significant lull in hash rates so Moore’s Law-based increases were inevitable. The ASICs were so much better than their predecessors though that they practically obsoleted them overnight. There’s no way that we could ever really see

that 7 minute block times again unless there's a fundamental technology breakthrough. What's actually very interesting is that the ASIC vendors allowed 7 minute block times at all. Those very high block rates seem like a huge economic mistake because they pushed the difficulty too fast. I suspect that the soundest strategy for ASIC vendors would have been to sit back and throttle supply of hardware. Instead, they actually fell victim to "get rich quick" mindsets and the consequences were trivially predictable. By adopting a "pre-sale" concept the early ASIC vendors got funds to build things but then almost guaranteed to need to rush large amounts of hash rate out to the network. I suspect that in retrospect they actually caused a prematurely large hash rate, while more cautious sales would probably have kept them all running with huge gross margins even now."

¹⁰⁸ See [Bitcoin Difficulty Adjustments](#)

¹⁰⁹ Thanks to Andrew Geyl (Organ of Corti) for the custom charts.

¹¹⁰ Yet it is unclear how many miners are rational versus simply "were lucky." In point of fact, the number of bankruptcies may be less if miners were less exuberantly optimistic about bitcoin mining. For instance, between June 2011 and April 2013, the price of a bitcoin dropped from \$30 to around \$2, but the network hashrate (and the network difficulty) did not drop significantly. Historically hashrate will increase with an increase in price but not decrease with a decrease in price because of the optimism of many miners. This exaggerates the costs of maintaining the network (e.g., more than a few people mining at a loss).

¹¹¹ This calculation comes from the following: 18 months of coin supply (approximately 1.97 million bitcoins) x \$400 (the weighted token value of bitcoin over the past year-and-a-half).

¹¹² Despite the quintupling of merchants that accept bitcoin for payments during 2014, the aggregate amount of bitcoins processed via payment processors collectively remained flat in 2014. See [Are there changes in the volume of retail transactions through Bitpay this past year?](#), [Will colored coin extensibility throw a wrench into the automated information security costs of Bitcoin?](#) and [A brief history of Bitcoin "wallet" growth](#)

¹¹³ In practice, the arithmetic to measure the total costs is: [number of bitcoins mined per day + fees] x [price of bitcoin] to get the daily cost which is then scaled to months and years.

¹¹⁴ [Slicing data: what comprises blockchain transactions?](#) by Tim Swanson

¹¹⁵ [Malignant computation](#) from *O'Reilly*

¹¹⁶ I would like to thank Petri Kajander for providing this analogy.

¹¹⁷ See p. 55 in [The Anatomy of a Money-like Informational Commodity](#) by Tim Swanson

¹¹⁸ [IBM ADEPT Practitioner Perspective](#)

¹¹⁹ IBM partnered with Samsung and used the W9000 washing machine as a test bed for the proof of concept. See [IBM Reveals Proof of Concept for Blockchain-Powered Internet of Things](#) from *CoinDesk*

¹²⁰ [Stuxnet](#) or [Flame](#) are two examples of malware targeted at sabotaging industrial output by specifically targeting embedded systems (e.g., programmable logic controllers).

¹²¹ [ARM chip to power connected fridges and clever lighting](#) from *TechRadar*

¹²² [A washer in your washer and smart-home sprawl at CES 2015](#) from *cnet*

¹²³ [What is the Internet of Things?](#) from Android Authority

¹²⁴ For Litecoin mining, see list of "Other" at bottom. In addition, see the 2013 performance [comparison](#) of ARM chips mining on the Bitcoin network.

¹²⁵ [Hacker Redirects Traffic From 19 Internet Providers to Steal Bitcoins](#) from *Wired* and [Hacker hijacks ISPs, steals \\$83,000 from Bitcoin mining pools](#) from *ZDNet*

¹²⁶ See [Gaming Company to Pay \\$1 Million for Secretly Using Customer Computers for Bitcoin Mining](#) from *CoinDesk*, [Slow Computer? uTorrent 'Epic Scale' Bitcoin Mining Software Is Slowing Down Computers Everywhere](#) from *International Business Times* and [The ZeroAccess Botnet – Mining and Fraud for Massive Financial Gain](#) by James Wyke

¹²⁷ [What is a Secure Element \(SE\)?](#) by Ganeshji Marwaha

¹²⁸ That is not to say that there is no utility to be found in the field of "Internet of Things." IBM is planning to spend \$3 billion over the next 4 years on its new IoT unit, some of which is allocated to sifting through "big data" for weather patterns. See [IBM's latest big bet: \\$3 billion on the Internet of things](#) from *Fortune*

¹²⁹ See also comments from Stefan Thomas in ["Practical and Social Implications of Cryptocurrency"](#) at *Cryptoeconomicon*.

¹³⁰ I would like to thank Taulant Ramabaja and Steve Waldman for this insight. Smart property (via proplets and "IoT") could soon blur the boundaries between on- and off-blockchain assets and it is currently unclear that legal

authorities will recognize experiments (e.g. safes that open only with a signature from a current on-blockchain owner).

¹³¹ [The next big thing for data centers: DC power](#) from *GigaOm*

¹³² This is not to say that Bitcoin miners are not “efficiently” inefficient. Large professional farms minimize as much capital expenditures that is not involving in the hashing process, some do not even use computer cases as this is considered economic dead weight.

¹³³ This includes: F2Pool (Discus Fish), AntPool, BW.com and BTC China pool.

¹³⁴ Personal correspondence, March 14, 2015

¹³⁵ Three presentations which cover this: [Tradeoffs in Cryptocurrency](#), [Introduction to Cryptoeconomics](#), and [Blockchains Are A Friggin Database Technology](#) – all by Vitalik Buterin