Debunking Blockchain

The case for centrally administered, but highly distributed, financial utilities

Douglas Jackson, CEO, Global Standard

Abstract. This paper addresses two related topics. First is a critical examination of blockchain and DLT (Distributed Ledger Technology) premises, implementations and performance. While blockchain/DLT continues to be the banner under which new initiatives are introduced, the unmistakable direction of later generation systems is to repudiate all of the original imperatives of blockchain technology – eliminating blocks, cipher block chaining and massive redundancy of records. The anonymity of cryptocurrencies, however, while continuing to be downplayed for credulous regulators, is being hardened. More broadly, with respect to both public unpermissioned and private permissioned networks, the core shibboleth and legacy feature of DLT—the avoidance of centralized administration—continues to be a source of unnecessary complexity arising from the need for "Byzantine" consensus protocols. Having explored consequences of a movement that started with a particular technology-centric solution and then backfilled to conjure a rationale, part two presents a "clean sheet of paper" alternative proposal designed to address real world economic problems which have become more acute during the decade-long blockchain/DLT detour.

Table of Contents

Blockchain, blockchain, blockchain!

The future will be built on the blockchain.

This affirmation has emerged as an article of faith almost as universally held as belief in the inevitability of death and taxes. However, as with any other revealed truth, articulation of more particular confessions exposes schisms, some of them basic, others more arcane and only disputed among the adeptus.

The question of whether unbacked cryptocurrencies such as Bitcoin, the original incarnation of blockchain, should be ascribed condign merit represents the most established controversy. Some institutions defer such value judgments in favor of cashing in on the current Bitcoin enthusiasm with the sale of relics and indulgences; futures trading and derivative offerings such as ETFs. However, the increasingly prevalent doctrine is that, while unbacked cryptocurrencies may be false messiahs, "the viability of the underlying blockchain or distributed ledger technology (DLT)" has been "demonstrated". As further observed by Bech and Garratt, "Venture capitalists and financial institutions are investing heavily in DLT projects that seek to provide new financial services as well as deliver old ones more efficiently. Bloggers, central bankers and academics are predicting transformative or disruptive implications for payments, banks and the financial system at large".

This paper nails the following heretical theses to the door:

1. So-called "Distributed Ledger Technology" embodies a peculiar and aberrant take on the well-established principles of distributed computing that results in absurd inefficiencies that unnecessarily increase costs, hinder scalability and result in pernicious side effects. Blockchain/DLT systems may achieve efficiency and scale only by deviating from canonical blockchain dogmas in the direction of the already well-understood network engineering principles that inform genuinely well-engineered distributed systems.

2. Unbacked cryptocurrencies are play money wholly unsuited to serve as reserve assets or media of settlement. Moreover, existing and proposed blockchain/DLT arrangements are also unsuitable to be used as platforms for the issuance, distribution, circulation, redemption and de-issuance of more sound and useful implementations of Digital Base Money, whether issued by central banks or private sector institutions.
3. Most if not all economically useful use cases that have been proposed as arguments favoring blockchain/DLT technical solutions can be addressed more effectively with systems implementing other technologies.
4. While the technology-related aspects of a system are of critical importance, using a particular technical solution as the starting premise or foundation for a system is to put the cart before the horse. Design and implementation of a coherent systematic solution, especially one providing for money and payments, requires the integration of well-conceived institutional arrangements, sound monetary principles and sustainable business models or other economic incentives for participants.
5. A system must provide for intermediated payments but direct end user access to Digital Base Money—which circulates efficiently via a global P2P payment system—is more essential to a transformative beneficial solution. Private permissioned networks that fail to provide for this and focus only on reinforcing the traditional role of banks as obligatory financial intermediaries in payments are an evolutionary dead end.
6. A well-conceived, coherent and comprehensive system is proposed that achieves not only the economic benefits claimed by DLT advocates but also affords an unprecedented capability of harnessing collective wisdom to achieve an automatically self-adjusting global system for metering money supply and fostering sustainable government fiscal policies and practices.

Premises matter

Logical arguments, even if valid, lead to unsound inferences if based on faulty premises.

The currently ubiquitous belief in the superiority of "decentralized" transaction systems based on blockchain/DLT over other arrangements rests on unexamined premises as first advanced by Craig Wright (aka Satoshi Nakamoto) in "Bitcoin: A Peer-to-Peer Electronic Cash System"[1]. As will be shown, one characteristic of Wright's worldview and promotional approach, and identifiable in all subsequent blockchain polemics, was to assume the superiority of a particular technology-centric solution and then backfill to conjure up a rationale for it. Invariably, the rationale consists of vague use cases consisting either of an abstraction drawn from game theory (but lacking practical real-world benefit[2]) or, even more typically, relying on straw man arguments which are quickly glossed over without close examination. The fatal problem with an approach that is grounded in faulty premises is that even though it may eventually evolve in a direction of greater rigor and sophistication, conserved legacy aspects of the starting point

[1] The historical and philosophic origins of Bitcoin and what has become the blockchain movement are explored below in Appendix: "The cool kids' new clothes".

[2] Most discussions of smart contracts fall into this "lacking practical benefit" category. Similarly, the claimed "Turing complete" attribute of the Ethereum Virtual Machine or languages for smart contract specification such as Solidity leads to a riot of impenetrable disputatious nonsense. "Turing complete", in this context, is the modern equivalent of Don Draper's inspired characterization of the basis of superiority for the tobacco in Lucky Strikes – "It's toasted!"

prevent the more refined elaborations from ever attaining the coherence and elegant simplicity of a clean sheet of paper approach built on a sound foundation.

Goals are as important as premises. Appropriate goal definition emerges from clear articulation of the problems to be addressed and possible benefits to be realized. Any other approach is certain to result in wasteful churn and incoherent schemes. Starting a journey without a particular destination in mind leads to a meandering path virtually certain not to lead to an optimal outcome. For example, the mindset evident in the recent ICO craze, echoing the same ethos as the late-90’s dotcom bubble, regards raising a lot of money for a venture (and an enticing “exit strategy”) as the primary goal and indicator of success while a coherent plan for generating broadly beneficial economic outcomes (or even operating profits) scarcely enters into consideration.

Core fallacies of DLT

As noted, a more thorough exploration of the origins and premises of blockchain/DLT is presented in the Appendix. To proceed, however, two interrelated elements that comprise the crux of blockchain/DLT dogma warrant preliminary discussion:

- Trusted third parties,
- Distributed networks.

Trust issues

Blockchain/DLT strategies place high importance on avoiding reliance on trusted third parties. For this reason, three obligatory elements are embedded in all canonical blockchain/DLT systems.

- Cipher block chaining,
- Transactions aggregated into blocks,
- All nodes validate all additions to the blockchain and keep a copy of the entire transaction ledger[3].

Cipher block chaining

First/second generation blockchain models such as Bitcoin, Ethereum and Ripple implemented cipher block chaining as a means of assuring an immutable sequential record of transactions or other recorded items of data.

The concept dates back to the 1970s. A 1976 patent application by researchers at Bell Labs described “cipher block chaining” as a technique to assure fidelity of signal transmissions through multiple intermediary nodes that might be under the control of independent, untrusted—perhaps even competing or malign—third parties. The idea was to prevent content from being altered/corrupted as it was relayed. The solution relied on cryptographic techniques.

The gist was to start with the content of a message and combine it with some pertinent information (metadata) such as the time of transmission, final destination coordinates and information regarding the cryptographic protocol being used. This composite was then put through some sort of cryptographic

[3] Blockchain apologists are currently in a frenzy to recant and repudiate these original imperatives. This is discussed below in “New directions with blockchain/DLT”.

function resulting in a unique output—in all modern communications this is known as a "hash function"—that, practically speaking[4], could only be generated by inputting exactly the same data.

The next node then would have this composite that included both the original message and the metadata that had been bundled in. More metadata would be appended at each hop such as a unique identifier of the intermediate node and the timestamp of retransmission, with the concatenation of prior hash and new metadata re-hashed to be passed on to the next node.

When the message arrived at its destination it was possible to then authenticate the message by proving that this string of intermediate steps had not introduced any alteration of the content.

All modern email systems implement cipher block chaining or a similar but enhanced variation. For example, the header of a recently received email includes this TLS cipher suite:

TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA384_P256

In this string, "SHA384" (Secure Hash Algorithm, in this version producing a 256 bit hash output and providing a 128 bit encryption key) indicates the message authentication algorithm used to generate a hash, while "CBC" means cipher block chaining was the block cipher mode used to assure fidelity of the content.

Cipher block chaining and its variants is an excellent solution for securing a particular message or thread of messages. But if message authentication were felt to require chained-in inclusion of all previous messages ever sent via a particular system, including the communications of complete strangers, it would constitute a rather cumbersome requirement.

Blocks

All blockchain protocols, since they are predicated on the notion that system participants can't be trusted and an unknown number of nodes might be under the control of malicious actors, implement complex consensus models as discussed below. Since the consensus process of first generation systems—cryptocurrencies circulating via public blockchains—was/is slow and often unwieldy, transactions were grouped together into blocks, enabling a potentially large number of transactions to be committed in a single batch.

As noted above and emphasized throughout this document, assuring an immutable sequential record is the imperative for any system designed to create a record of financial transactions. Blockchain strategies chain blocks into a sequence by requiring each block to contain a hash of (or pointer to) the previous block such that each state of the system contains a valid representation of all previous states.

But each block contains multiple transactions. How do you assure the immutability of each component transaction? You could concatenate all the details of each transaction into a combined (and potentially

[4] "Practically speaking" means that it is actually possible for different inputs to generate the same hash output. This is due to the lower entropy of the hash output – the fact that it is limited to a result that can be expressed in a specified number of bits (in the case of SHA256, 256 bits—typically a string of 64 hexadecimal characters, comprising 2^{256} possibilities) —compared to a truly infinite set of possibilities for the input information. The likelihood of two different inputs generating the same hash, a so-called "hash collision", is statistically infinitesimal. A partial hash collision, in contrast, can be as unlikely or likely as desired, depending on how many characters in the same positions are required to be identical. This is the basis of another core component of "mined" cryptocurrencies such as Bitcoin, the "hashcash" protocol.

very long) string (of characters, such as letters and numbers) and generate a hash representing the whole blob. But that would make the subsequent process of locating an individual transaction—and validating that is has not been altered—impracticably cumbersome. Therefore, to facilitate the process of finding/validating specific transactions or other component data, most blockchain protocols use some variation of a Merkle Tree logic in organizing the data structure of a block.

The use of blocks incurs some side effects. Blocks are sequential and sequential order is indeed more meaningful than a timestamp per se as a strategy for preventing double spends – though a timestamp is pretty handy and does come into play at the block level. Transactions within a block, however, are not sequential. While a Merkle Tree introduces considerable efficiencies it implies no temporal sequence of the components. For example, what if a block contains payments from two competing bidders for an asset sold under auction rules that hinge on whose payment arrives first? Or, in a non-payment scenario, suppose the data items are competing patent filings or other scenarios that involve staking a claim with a system that determines priority on a first-filed basis?

There is another, bigger, problem. As (perhaps over-) stated by promoters of a hybrid blockchain/database scheme described below, "the blockchain essentially has no querying abilities". Stated more modestly, data structures assembled by a "trusted committer" afford greater efficiency of query performance.

As will be described in the technical overview of a proposed superior alternative system, query efficiency becomes even less of a potential operational bottleneck in a design pattern that integrates the concept of "read models". A read model effectively pre-queries data in a fashion that "denormalizes" it to an optimal read format and facilitates its forward caching to enable extremely rapid interrogation.

Massively excessive redundancy (and complexities relating to privacy)

Blockchain/DLT protocols embody a concept of "distributed" that is at variance with well-conceived principles of distributed computing networks. As described by Vitalik Buterin, the (trusted, even revered) mastermind of the Ethereum system, "Currently, in all blockchain protocols each node stores all states (account balances, contract code and storage, etc.) and processes all transactions".

The idea is that if every participant has her own copy of the entire record of transactions, any exploit designed to alter, delete or insert a back-dated item of data would be evident.

One problem with this approach is that over time the continuous accretion of new data can cause a blockchain to become massively unwieldy. The Bitcoin blockchain has passed the 160 gigabyte mark. To initialize a full Bitcoin node (without even attempting to engage in mining) on a typical laptop or home computer may require devoting its entire computing capacity to the task for several days.

The problem is compounded by the requirement that every full node maintain a complete copy. In addition to the wastefulness of a massively excessive number of redundant copies, this introduces complexities with respect to privacy. On the one hand, the fact that everyone can see the transactions or other activity of all other participants can be spun as a good thing – the ultimate in auditability and transparency. On the other hand, it is often appropriate for one's transactions to not be broadcast to and viewable by the whole world. So other blockchain strategies may go overboard in the other direction; all nodes hold a demonstrably authentic copy of everything but only designated participants – typically the parties to the transaction – can see the actual content of the transaction. This approach, the angle employed by systems such as ZCash or Monero for competitive advantage, is optimal for money laundering, terrorist finance and other criminal activity. It also is likely to make such systems a modality of choice for government agencies engaged in unauthorized black ops.

The DLT constraint that every full node process all transactions and store all data requires a level of redundancy that could only make sense in a value system akin to religious dogmatism. Even non-public blockchain systems in which, for instance, participant banks each operate a node face exponentially worsening complexity with respect to consensus models when the number of nodes reaches about twenty.

As will be discussed below, design of a well-engineered distributed architecture system determines how many copies of data should be maintained based on considerations of consistency and availability. At a minimum there are backup copies to avoid data loss and assure continuity of service in the event of equipment failures and network partitions. The question of how many copies boils down to an engineering question, the answer to which (for any large system) would never be "every node must maintain a full copy".

Alternative approaches

Safeguards against data alteration

To review, the major rationale for the blockchain/DLT sine quibus non—linking transactions in a chain of blocks and arranging for each node to have a full copy of everything—is more or less to preclude someone from monkeying with the transaction history and balances, that is, representations of system state. In the following section, I will relate why that concern is a red herring. Briefly, it would not be in the best interest of a financial provider to commit data tampering or to offer and rely upon a system that could be compromised in such a fashion. But, supposing certain users of a system felt that a tamper-proof system was of such paramount importance that they wanted belt and suspenders safeguards, is a combination of blockchain and everyone-gets-their-own-copy the best way of achieving such certainty?

No. There are dramatically less resource-wasting strategies that could be employed.

For starters, with most users of a system—such as the many millions of people worldwide who use online or mobile banking systems—these concerns are not even on their radar. They might not particularly like their bank but they assume their bank isn't going to lose track of or make off with their money. A resource-costly universal system may not be warranted if, say, only 5% of people (probably a high number) regard it is an issue.

Of this putative 5%, it is also possible that special assurances aren't needed all the time. There may have been a particularly large deposit they would prefer to document with extraordinary rigor or perhaps they feel the need to prove they made a certain payment before some sort of deadline. It is a stretch to imagine that even 10% of the transactions of this 5% fall into this non-routine category requiring more rigorous documentation. Does it make sense to implement a resource-costly universal requirement for such a use case when for > 99.5% of transactions there is no demand for it?

An alternative approach would be for a provider of transaction services to offer a utility enabling customers to order/obtain a real-time report that is digitally signed by the provider using a suitably strong cryptographic protocol. Public key cryptographic capabilities are easily implemented enabling the provider to implement an automated system for signing a report with the provider's private key in a fashion that enables it to be authenticated against their published public key (see here for an example).

Such a snapshot in the possession of the user, documenting the state of certain transactions or account balances at a particular timestamp (and containing the report parameters such as range of dates, transaction value or data fields to include) constitutes proof positive that would be determinative in any subsequent dispute. This process could be automated to run on a schedule or on an ad hoc basis, with either option supporting customer-specified parameters to define the range of transactions or other account information to be certified.

It is uncertain if commercial demand for such strong cryptography-authenticated reports exists or would actually emerge. For most people—and this applies even to people providing documents for tax or other audits—printing out a completely unauthenticated report from third-party-maintained transaction registers seems to be good enough. For that matter, even in matters as important as contract litigation, it is not unusual for the parties to print out unsecured email messages that could easily be faked as documentation of written understandings! While this is a ridiculous state of affairs, the provision of cryptographically authenticated reports provided by a suitably trusted third party would be a definitive and much more resource-efficient alternative than blockchain/DLT arrangements.

A more rigorous methodology for verifying immutability might entail implementation of a Verifiable Data Structure in the form of a verifiable log-derived map (VLDM). The potential utility of integrating a VLDM is being evaluated in relation to the proposed alternative system described below, which is based on an append-only Event Store model.

Transparency/auditability

The companion argument for a blockchain, a full copy of which is maintained and held by each node, is transparency and assured auditability. But the one-size-fits-all nature of public blockchain schemes, with either too much disclosure and lack of privacy, or too much privacy/anonymity, facilitating criminal usage, is a poor solution.

In June 1999, the e-gold® system deployed a "Publicly Viewable Balance" feature[5] in support of concurrent efforts to encourage banks to offer e-gold deposit facilities backed by reserves of actual e-gold. This feature enabled a user to specify that the e-gold system would automatically display the real-time balances (there were four e-metal® Currencies) upon request, by anybody, of that user's account. It would be technically trivial to again implement such a feature, extending it to include parameter-specified transactions. For example, a charitable institution or NGO could elect to designate not only account balances as publicly viewable but also all transaction details, funding and/or disbursements, in accordance with a requesting user's specified date range or minimum value threshold, including counterparty identity.

The trusted third party issue is a red herring

In the United States, CHIPS is the principal clearinghouse for large value USD-denominated payments. With institutional origins dating back to the 19th century, CHIPS has evolved sophisticated protocols enabling intraday settlement of payment messages with finality, commonly achieved in seconds to minutes. It has eliminated (payment) default risk via a well-conceived prefunding mechanism and the

[5] The Publicly Viewable Feature was first implemented for DigiGold™, a digital cash based on public-key cryptography and implemented via smart contracts. DigiGold, while long forgotten, was the first digital cash to ever attract commercial usage, unlike the widely known but essentially never used real money version of Digicash. Unlike later "cryptocurrencies", DigiGold was backed by a 100% reserve of e-gold. Circulation and reserves exceeded 4kg for the e-gold-backed AUG and 50kg for the e-silver-backed AGG.

direct % cost of processing a large value payment through CHIPS is negligible. The operational performance of CHIPS is unblemished. Even with the less robust protocols prevailing nearly two decades ago, and despite elements of the system being collocated practically at ground zero, CHIPS withstood the 9/11 disruptions with scarcely a hiccup. Since that time, CHIPS has been successively hardened and improved through multiple generations of upgrades and refinements and continues to roll out new features and capabilities.

Multiple other financial utilities around the world offer comparably robust, low cost, reliable services.

There is a commonality that can be abstracted with respect to how these financial utilities relate to their participant financial institutions. Banks that are direct participants each implement major league information technology systems, the core modules of which, in nearly all cases, are provided by software and systems providers who specialize in banking systems. A dozen participant banks may use systems provided by a dozen (or more) different providers, although some providers have captured global market share more effectively than others and may as a result serve hundreds of banks. But each provider builds their systems as to interface with the increasingly standards-driven interfaces implemented by the centrally administered financial utilities.

There are aspects common to all of these arrangements that afford possibilities for radical advances—addressed by the alternative system presented below—but it is ridiculous to suggest that their possible shortcomings have anything to do with them being “trusted third parties”.

In contrast, the “decentralized” systems that are discussed below in relation to private blockchains or proposed permissioned networks can be represented by a different abstraction. All participants must operate their own, complicated (and therefore less efficient) nodes, in a “Byzantine” network of mutually untrusting participants, all implementing The System provided by the winner of the technology wars, as if such a winner would be any more likely to emerge than in the current competitive market for provision of banking software[6].

The trusted third party boogeyman alleged by Blockchain/DLT advocates is a straw man. The tradition of ascribing vague horrors perpetrated by (or visited upon) trusted third parties dates back to Wright’s original Bitcoin whitepaper and continues every day. Two facets can be identified:

- Misdeeds on the part of the trusted third party itself,
- Casting any form of centralized functionality or responsibility as a “single point of failure”.

Misdeeds of trusted third parties

What vile deeds might a trusted third party commit to the detriment of those foolish enough to depend on it? Generally speaking, the idea is that records of past events may be altered, purposely deleted or lost, or backdated records inserted as to appear as if certain events or transactions occurred prior to their actual time of occurrence.

[6] The proposed alternative system fits with the existing pattern in the sense that it provides new centrally administered financial utilities, offering extended capabilities that transcend the (almost unrecognized) arbitrary limits imposed by legacy institutional constraints. The idea is then for competitive providers of banking software to make sure their modules interface well, which will not be challenging since the proposed alternative supports existing standardized interfaces such as those used for specifying payment messages.

Wright qua Nakamoto cites a single example of a harm he ascribes to "financial institutions serving as trusted third parties to process electronic payments" – payment repudiation. To my knowledge, no one has ever challenged his glib attribution of cause – his implication that payment repudiation is a matter of a financial institution altering or deleting a record. But payment repudiation simply does not stem from alteration of records. The idea that any financial intermediary or utility using software and systems even remotely conforming to universal modern standards amends financial records in such a fashion is ludicrous.

Payment repudiation arises from institutional arrangements and transaction models. If a payment modality is designed to enforce non-repudiation it will be engineered accordingly. e-gold®, for example, guaranteed non-repudiation of payments. Even with the relatively primitive technologies employed at the time, not a single instance of payment reversal ever occurred (out of over 95 million cumulative e-gold Spends processed), even in cases where it turned out that e-gold's affiliated exchange service, OmniPay® had been defrauded by a counterparty.

Credit card payments are subject to chargebacks. But this is due to their credit nature; the system is designed such that merchant recipients (rather than financial or operational intermediaries) bear the direct costs and risks stemming from fraud or certain categories of payer non-performance.

To this day, blockchain advocates suggest that non-blockchain data persistence mechanisms—typically depicted as old-timey SQL-type relational databases—implement the full suite of operations commonly referred to as CRUD; Create, Read, Update, Delete. This, however, is counterfactual – at least with respect to systems engineered to standards that would make them suitable for use by a financial institution or utility. Real world systems conform to long-established accounting conventions; committed records are not altered but rather, if amendment is required (for example if an erroneous entry has been committed) correcting journal entries are added so as to afford a fully auditable record[7].

As with blockchain, virtually every data persistence mechanism employed in accounting-type systems only appends records. As will be further detailed below in the description of the proposed alternative system, an append-only mode of data persistence not only assures data integrity and auditability but also affords optimization of efficiency.

Single point of failure

Blockchain/DLT advocates reflexively conflate any element of centralized functionality in a system with "single point of failure". This meme has become so pervasive that even sophisticated analysts who should know better have adopted this convention[8]. There are two aspects to this ascription:

- A system with a single point of failure may malfunction such that there is service interruption or a potentially catastrophic loss or corruption of data.

[7] Actually, even with SQL Server, the relational database that is the archetypical boogeyman for straw man use cases where records might be altered or inserted after the fact, a utility called SQL Server Audit can make even UPDATE writes auditable – though it is hard to imagine any credible vendor of financial transaction software giving any process UPDATE or INSERT permissions in a production system.

[8] For example, Accenture's recent report on the Bank of Singapore's Project Ubin Phase 2, repeatedly cites potential "single point of failure" vulnerabilities in the various "decentralized" schemes evaluated.

- External coercion—actions on the part of private sector criminals or government authorities—may result in system dysfunction or cessation.

The simple fact of the matter is that all systems, whether blockchain/DLT or based on other strategies must be designed and implemented with these risks in mind. Mitigation of these and other risks is always an imperative for system engineering requirements. But it is also very much a consideration that must inform institutional arrangements. The ethos that was foundational in the design of the proposed alternative system was to integrate institutional arrangements, transaction and business/revenue model and technical implementation solutions as to systematically identify and mitigate any aspects that might require anyone to trust any person or entity with responsibilities relating to system integrity. Safeguards to thwart malfeasance on the part of incumbents assigned to critical roles also serve to mitigate risks stemming from error or external coercion.

As will be examined in the Appendix, the abhorrence of trusted third parties in blockchain/DLT doctrines largely stems from the reverence accorded by blockchain/DLT clergy to game theory. But Quixotic efforts to construct a system free of trusted third parties tend to lead to the result that no one is responsible and liable when something goes wrong. And with any real-world system, something always goes wrong.

Cryptocurrency "fails"

With respect to assuring an immutable sequential record, there are a lot of ways to skin the cat and a particularly robust alternative system is presented below. The reality, however, is that some of the most ridiculously monumental fails with respect to system integrity and performance have occurred with blockchain/DLT systems. Each can be construed as an instance where system integrity was dependent on someone somewhere who had effectively been entrusted with a responsibility.

Buterin (May 2016) related with relish the following screw-ups that occurred with Bitcoin:

- "In 2010, an attacker managed to give themselves 186 billion BTC by exploiting an integer overflow vulnerability. This was fixed, but at the cost of reverting half a day's worth of transactions.
- In 2013, the blockchain forked because of a bug that existed in one version of the software but not another version, leading to part of the network rejecting a chain that was accepted as dominant by the other part. The split was resolved after 6 hours.
- In 2015, roughly six blocks were reverted because a Bitcoin mining pool was mining invalid blocks without verifying them."

The Ethereum community can hardly boast a better track record.

- In May 2016, a virtual entity—the DAO (Distributed Autonomous Organization)—was established, constituted as a smart contract on the Ethereum network, having raised 11.5 million ETH (about 14% of all Ether then in circulation), equivalent at the time to $186 million, from anonymous investors in the most successful (largest) crowdfunding campaign ever recorded. Less than a month later, an attacker exploited bugs in the smart contract code and made off with a third of the money. The Ethereum organization responded by repudiating the stolen Ether, about 5% of all Ether, via a "hard fork".
- In July 2017, hackers stole a total $32 million worth of Ether from three wallet addresses holding balances in the popular "multi-sig" Parity wallet. A team of white hat hackers then went into

action, essentially hacking all Parity wallets in order to safeguard another $208 million’ worth. This introduced yet another bug enabling a well-meaning novice hacker to accidentally take control of the apparently not-so-smart contract and irrecoverably “kill” $300 million worth of value.

In none of these cases was anyone held financially liable.

Disasters involving exchange services are another illustration of a systematic deficiency of public blockchain schemes. Such systems are “unpermissioned” in the sense that no centralized administrative entity/capability organic to the payment system exists to impose or enforce rules on would-be providers of exchange services (or any other category of user).

High profile failures in which customers holding balances on account with providers of exchange include:

- Mt. Gox – a Japanese service which at the time of its failure in February 2014 was reportedly handling 70% of all Bitcoin currency exchange activity worldwide. Approximately 850,000 bitcoins belonging to customers and the company—an amount valued at more than $450 million at the time—were reported missing and likely stolen.
- In January 2018, Coincheck, another Japanese cryptocurrency exchange service, claimed to be hacked resulting in the loss of $534 million worth of “NEM”, another unbacked cryptocurrency.

There is an additional inconvenient truth relating to the trusted-third-party-avoidance fetish of blockchain advocates.

Blockchain/DLT systems require as much “trust” as legacy systems

Deployed blockchain/DLT systems are riddled through and through with trusted third parties though, in many instances, the (often unrecognized and unknown) third parties on which elements of system integrity depend are wholly unaccountable.

Core programming teams – The earliest adopters of blockchain products were tech savvy people, the sort who contribute code to open source consortiums and prefer to download source code from Github as to compile applications such as wallets themselves. Later participants, attracted in the case of Bitcoin, Ethereum or (perhaps most incredibly) Ripple, by the prospect of overnight riches, are in a whole ‘nuther category. Some may exert themselves to wade through technical material with the goal of being able to impress others with their facility for using the latest jargon. But most are entirely reliant on the elite programmers who undertake to maintain and advance the code base. Sometimes, however, as noted, disastrous bugs compromise system integrity. Even when bugs and exploitable vulnerabilities are contained, participants rely on core teams to find solutions to baked-in problems such as limitations to scalability leading periodically to prolonged latency or spikes in transaction fees. Teams must be trusted to resolve schisms that crop up continuously such as the prolonged recent factional controversies with respect to Segwit2x or the sorts of clashes that led to agree-to-disagree hard forks (Bitcoin Classic, Bitcoin Cash, Bitcoin Gold) and even rogue variants (the latest apparently being Bitcoin Clashic). With Ethereum, the faithful await for Vitalik Buterin to eventually descend from the mountaintop and pronounce his blessing on the prophesied conversion from Ethereum’s Proof-of-Work consensus protocol (i.e. “mining”) to the “Casper” Proof-of-Stake consensus mechanism, trusting that it won’t mess everything up. Ripple is no different, with its Interledger Protocol team making it up as they go along, reminiscent of the economic models that come in and out of vogue with central bankers exercising discretionary authority over monetary policies.

Web-based wallets – Reminiscent of the ill-conceived and ill-fated Digicash discussed in the premises and origins appendix, normal people cannot be bothered to download and install computer programs (and probably have no business doing so, given their inability to protect themselves from malware). So a higher and higher proportion of blockchain end users access such systems just as they have always accessed trusted third parties such as banks – they log in to some website that holds customer-owned value on account, perhaps via a "web-based wallet". Many exchange services also hold, or claim to hold, customer-prefunded real money balances as to facilitate rapid execution of customer trades without exposing the exchange itself to the risks of payment repudiation endemic to the payment systems used to convey conventional money. This inexorable tendency, first manifest with Bitcoin, also is applicable to Ethereum and a proliferating plethora of "alt-coins".

Smart contract mongers – Having noted the fate of the DAO, at latest count, some 90,000+ different ERC20 token contracts that circulate via the Ethereum network are now on offer. Would-be holders or purchasers must trust some authoritative sounding reviewer, who may be nothing but a shill, to tell them if any particular code-defined token is itself worthy of trust.

ICO principals - The idea of an ICO is "I will create a specified quantity of a token from thin air, millions or billions of them, for which I have no liability to buy any of them back or to hold assets earmarked to assure that a secondary market for them exists. I will keep a large portion of them and sell the others to the public in exchange for consideration that you must trust will be used for a specified clever purpose but which I am free to use for whatever purpose I choose, certainly including my own compensation." One of the more legendarily successful such offerings was Ripple. In 2012, the Ripple founders created 100 billion units of "XRP", keeping 20 billion for themselves and awarding the remainder to the Ripple company to dole out to the public in dribs and drabs as market demand affords. So far, the folks metering the distribution of this most-fiat of any currency ever conceived have managed to avoid killing the golden goose. For a brief interval in early January 2018, Ripple co-founder Chris Larsen was worth nearly $17 billion on paper, making him one of the wealthiest people on the planet.

Validators – The Ripple system relies on company-designated "validators" and is therefore not regarded by DLT purists as a decentralized system.

Centralized functions in schemes for decentralized clearing and settlement of intermediated payments – The recent report of Singapore's Project Ubin Phase 2 detailed three decentralized solutions evaluated in a pilot study. Two of the three (Hyperledger Fabric and Quorum) required certain functions to be performed by a centrally administered process and all provided full visibility of all transactions to the sponsoring government monetary authority.

Issuers of asset-based media – The idealized notion of a system that eliminates any role for trusted third parties goes out the window in the case of an issuer purporting to assume liability for representations that a particular token is backed by real world assets held in readiness as to assure performance of obligations. Assets reside in custodial arrangements that require them to be titled to someone or something. This of course is a good thing but it conflicts with the imperative to engineer-out the need to trust anyone. As discussed below, the insistence on avoiding trusted third parties was at the root of the monetary fallacies baked into Bitcoin and other unbacked cryptocurrencies.

Tether/Bitfinex

Tether/Bitfinex is a special case that illustrates that the blockchain community itself both relies heavily on (egregiously) untrustworthy third parties and is at particular risk of collapse due to a single point of failure.

As this note is being written, it is likely to be revealed that Tether, and its co-owned accomplice Bitfinex, have long been engaged in massive fraud that will result in customer losses in the hundreds of millions or even billions (USD-equivalent). A cursory look (by this author) at Tether's website in December 2017 revealed blatant discrepancies between the claims presented on the most commonly viewed pages ("100% backed by actual fiat currency assets in our reserve account… redeemable") and the actual Terms of Service ("Tethers are not money and are not monetary instruments. They are also not stored value or currency. There is no contractual right or other right or legal claim against us to redeem or exchange your Tethers for money. We do not guarantee any right of redemption or exchange of Tethers by us for money.") Though obviously a company engaged in the business of money transmitting, the Tether site made no reference to a license, anywhere (or a determination that license was not required), and listed neither physical address nor service of process information.

It is a sign of the flakiness and speculative mania that underpins blockchain enthusiasm that over $2.5 billion of value (up from a claimed $10 million in January 2017) is more or less entrusted/ascribed to such a patently crooked operation. For example, despite glaring indications that Tether's "USDT" "stablecoin" is likely backed by at most a few pennies to the dollar, the largest trading pairs on numerous high-volume exchanges involve trading speculative cryptos such as Bitcoin and Ethereum (and hundreds of others) against USDT. Looking just at Binance, a total of sixteen cryptocurrencies trade against USDT with a combined 24-hour trading volume (June 13, 2018) in excess of $750 million. [The remaining 337 pairs are other cryptocurrencies between traded against each other]. Looking at all exchanges combined, 2.2% of Bitcoin's "market cap" trades daily, while the daily trading churn for Ethereum is even higher at 2.8%, the vast majority of both being traded against other cryptos and Tether.

But what if you wanted to exchange your USDT for real money? Tether won't (and can't) redeem it. So the few souls who do must seek out one of the two exchanges that support trading between Tether and actual US dollars. This does not entail redemption but rather is a matter of relying on other punters foolish or reckless enough to offer real money for USDT tokens. This, however, is relatively uncommon as the vast majority of cryptocurrency speculators are content with the paper gains numbers the crypto house of cards entails. Only about $1.1 million worth of USDT (0.04% of USDT's "market cap") trades against actual USD daily, almost all of it on Kraken (which offers 2X leverage on USDT trading.

Distributed network architectures

As asserted above, the concept of "distributed" as the term is applied to so-called "Distributed Ledger Technology" is at variance with well-established principles of distributed computing or, more specifically, distributed computing systems implemented by means of networked computers.

A distributed computing system means a network of autonomous computational entities (computers or nodes), each of which has its own local memory and communicates with other nodes via a messaging system to achieve coordinated action. Why does this matter?

Suppose a robust strategy for assuring an immutable sequential record of transactions is implemented in a fashion that generates a single instance of a ledger persisted on a single computer. Such an arrangement would be vulnerable to loss or non-availability… because computers fail. Similarly, if the computer was remotely located, i.e. somewhere where it could only be accessed via a network such as the internet, problems with the network—a "network partition" in software engineering terms—could make the system inaccessible.

For this reason, all but the most trivial systems that involve processing information and persisting data (recording data for later retrieval) are implemented on systems involving multiple computers, often arrayed with some degree of geographic separation between the various nodes – a distributed network.

Distributed networks have been around for decades and millions of man hours of creative effort have gone into advances and refinements designed to prevent distributed data being corrupted, lost or unavailable. Critical systems must be fault tolerant (performing reliably despite computer malfunctions or outright destruction) and may need to build in capabilities for fail-over. For example, were a data center in the US to be vaporized in a nuclear attack, back-up facilities in another continent might kick in in a fashion such that even in-process transactions either complete or fail gracefully.

A useful abstraction for evaluating requirements and determining suitable designs for robust large scale networked systems is the CAP Theorem, also referred to as "Brewer's Theorem" for Eric Brewer, the computer scientist and network architecture expert who first articulated it in 1999. In particular, the CAP Theorem can be used as a framework for consideration of strategies to assure that data is persisted as an immutable sequential record that is highly available to system participants.

CAP Theorem

The CAP Theorem holds that there are tensions and inevitable tradeoffs between the imperatives of Consistency and Availability in every network architecture strategy that affords Partition Tolerance.

Consistency—the "C" in the CAP Theorem—means that every "read" operation (e.g. a database query) receives the most recent "write" (i.e. recorded entry of a particular item of data), or an error. In other words, every node gives the same answer if asked the same question. For example, a process for spending money that strictly prohibits overdraft payments must have access to the most up-to-date data regarding the payer's available balance.

Availability means that every read request receives a (non-error) response – without guarantee that it contains the most recent write. A never-ending spinning "wait" indicator on the screen is a bad thing.

Partition tolerance means that the system continues to operate despite an arbitrary number of messages being dropped (or delayed) by the network via which the nodes communicate. A somewhat related situation would be if some nodes involved in a transaction fail in the middle of it.

With blockchain/DLT (and all other networked transaction systems) the imperative of consistency imposes the corollary need for a "single source of truth" (SSOT), some locus with the true current version, including the most recently appended record, such that were a discrepancy between the SSOT vs. data on another node/computer to occur, the SSOT is right and the other(s) must be brought into agreement. Consistency means that all nodes have the true current version, each as good as any other in terms of serving as the SSOT.

This is a complicated consideration, all the more so because the concept of consistency has evolved.

Consistency as used in the ACID[9] concept

Originally, consistency was a concern with respect to database design, especially as relational database management systems (RDBMS) based on SQL (Structured Query Language) such as Oracle and the later SQL Server product line of Microsoft became industry standards. A single database required safeguards to avoid internal inconsistency. For example[10], a poor design for a table that contains elements of an address might contain a field for "country" that allows free form entry of a string of alphanumeric characters. One data entry operator might enter "UK", another "U.K.", another "United Kingdom", another UJ (they missed the K key) etc. A later query searching for all customers from "UK" would find some but not others.

A principle for designing data models was "normalization" (relevant to later discussion of the technological advantages of the alternative system I will be proposing). Normalization meant dividing up a data structure into a lot of small tables such that the potential for recording conflicting duplicate values for any data item (or violating other constraints) was eliminated. Instead of a free form field for country, a separate "Country" table might be created with a unique numerical identifier for each record (COUNTRY_ID), a COUNTRY_DISPLAY_NAME field/column, which in the example case might contain "UK", and a COUNTRY_CODE field containing the unique ISO-ALPHA3 Code which, in the case of the United Kingdom of Great Britain and Northern Ireland, is GBR. Then the user interface for data entry into any table involving address information could simply provide a dropdown list of COUNTRY_DISPLAY_NAMEs and selecting one would prompt the system to record the unique and unambiguous COUNTRY_ID in the relevant field.

Atomicity

An attribute related to consistency is "atomicity", the quality that all parts of a transaction must complete for it to be committed and failure of any part must lead to a rollback of all parts. For example, a payment involves a decrement of payer balance and an increment of recipient balance. Were either part to fail and the other to be recorded it would lead to either an increase or decrease in the overall amount of money in circulation via the system – a discrepancy representing a failure of consistency. Decrement of the balance of a payer, without increment of recipient balance, means that some money has been lost in the ether. Increment of a recipient balance without decrement of the payer balance means that money has been created out of thin air and sets the stage for a payer being able to spend the money again, a "double-spend" in the parlance of digital cash enthusiasts.

In primitive single database server RDBMS systems, atomicity might be assured my means of temporary "locks" to prevent any permanent change of the state of payer or recipient records that would affect balance until it was certain that both limbs, the decrement and increment, would commit concurrently in

[9] ACID was the acronym for Atomicity, Consistency, Isolation and Durability. We will return to the notion of atomicity in our discussion of "eventual consistency".

[10] This example takes some liberties in that the conflicting data would arise from poor application level design (an amateurish data model) whereas the original sense of consistency in the ACID notion was a matter of database programming to assure, as the Wikipedia entry phrases it, "data written to the database must be valid according to all defined rules, including constraints, cascades, triggers, and any combination thereof". Plus, I wanted to queue up an introduction to normalization/denormalization rather than digressing further into technical topics related to consistency in RDBMS, such as referential integrity.

a single atomic transaction. A big problem with such arrangements was that poorly conceived lock arrangements could slow everything down by requiring a transaction to complete before another could be processed.

Consistency in a system of networked computers

While it was tricky to assure that a single database would remain consistent, keeping multiple computers in synch was a more complex challenge, especially when the computers were geographically separated. Whether on a single machine or a networked set of computers the problem largely boiled down to messages, whether connecting two processes on the same computer or coordinating the actions of multiple separate machines. In the above example, the algorithm for achieving atomicity was "concurrent", requiring the system to wait until messages had arrived, been processed in a proper sequence and confirming messages sent and received in the other direction.

As will be discussed below, both as it relates to blockchain/DLT and the proposed alternative system, a major advance was to relax the concurrency constraint, allowing a system to grind away on multiple transactions at the same time despite the likelihood of messages arriving out of order, in duplicate (or triplicate etc.) or not arriving at all. The key is two-fold, "serialization" and "eventual consistency".

Serialization

Serialization means that messages can be sorted out somehow without requiring them to be concurrent and putting the system on hold until a transaction has been committed with strict atomicity. The changes of state end up being applied in a sequential fashion so that all later transactions are informed by an accurate sense of system state arising from earlier transactions. This is the "sequential" aspect of the "sequential immutable record" I keep referring to. It is why the Nakamoto paper places emphasis on what he calls a "timestamp server", particularly a "peer-to-peer distributed timestamp server to generate computational proof of the chronological order of transactions".

Eventual consistency

The key that enables such serialization in scenarios where messages are received asynchronously—at different times, in no particular order, and sometimes duplicated or missing altogether—is an arrangement for achieving "eventual consistency". The strict atomicity of a transaction processed via a concurrent algorithm can be relaxed a bit:

- the increment of recipient balance doesn't have to be applied at the same split second as the payer's decrement provided all increments and decrements of every balance are eventually applied in proper order, and/or,
- not all servers have to be instantaneously consistent provided that momentary inconsistencies do not result in permanent errors that could otherwise result in double spending or money being lost in the ether.

Where the rubber meets the road is the absurd overhead and, in the case of public blockchain systems such as Bitcoin, the slowness that results from their peculiar take on distributed (or decentralized) systems, the "every node processes everything and stores everything" conceit. This, for example, is why Bitcoin takes more than 1,000 times as long as e-gold did over fifteen years ago to settle a payment with finality. Eventual consistency that leads to a transaction taking several seconds to finalize is one thing. A payment that can't be relied upon until it is six blocks deep into the blockchain—taking an hour with

Bitcoin, if you're lucky and also paid an exorbitant fee—would be pretty silly when in line at the checkout counter.

But it gets worse. Keeping multiple computers in synch entails additional considerations;

- Consensus,
- Propagation.

Consensus

A need for consensus comes into play in two contexts, one of them (obviously the second one) peculiar to (public) blockchain/DLT world:

- Deciding whether and how a transaction is to be committed/persisted,
- Determining whether a committed transaction was actually bogus after all.

Consensus in a managed network

Returning to the primitive RDBMS epoch, the first level of logic for determining whether a transaction should be committed was contained in the process for achieving strict atomicity. That's fine but if at that point the only record is on a single hard drive on a single computer, a hard drive failure wipes it out. So there must be one or more duplicate copies. But what if the process for creating backup copies has not yet completed when the initial instance goes missing due to hardware failure?

This led to strategies for a so-called two-stage commit. Atomicity that entailed a single process creating a single copy was insufficient so the requirement for atomicity was extended. Not until one or more additional copies are created, and confirmation of this (still provisional) duplication received, does the transaction truly commit. A failure of this intra-transaction duplication would cause the entire transaction to roll back.

So what if a cluster of five machines is part of this two-phase commitment process; how many of them need to successfully queue up the still provisional transaction and report back to the controlling process before it is ok to commit – all? a majority? This is where consensus comes into play.

With a centrally administered system, whatever consensus rules are required (which, in centrally administered systems, tend to be simple) are determined, implemented and executed in an unambiguous logic which can be tweaked and optimized to a state of high efficiency and reliability. Typically, if consensus comes into play at all, it is a matter of assuring the system recognizes where the buck stops in terms of what constitutes the SSOT in accordance with the designed balance of consistency and availability taking the certainty of occasional but unpredictable network partitions into account.

Consensus with blockchain/DLT

It's not that simple with blockchain/DLT.

With a blockchain, whether public or private, where all full nodes are created equal and no one is in charge, which node (or nodes) does the initial write, i.e. gets to append a block to the blockchain, and how does the network reach agreement that the write was legit? Or (as is the case with certain transactions on private blockchains such as clearing and settlement arrangements for a group of banks) how do the participants reach consensus on what the transaction should even look like before it is committed by any or all of them?

Proof of Work

Bitcoin, notoriously, pairs the determination of which node gets to write a new block of transactions to the blockchain with its celebrated "mining" function. The consensus model utilized is called Proof-of-Work (PoW). Numerous nodes compete to perform innumerable otherwise useless computations in a race to stumble across a "partial hash collision" of specified statistically-predictable rarity. The result is an increasingly scandalous consumption of electricity, not to mention a thus-far unending arms race of capital expenditure – likely exceeding $3 billion USD in sunk costs for rapidly obsolete mining equipment that, being "application-specific", cannot be re-tasked to any economically useful purpose except perhaps as a space heater for use in a cold climate. Or a door stop.

Currently, per digiconomist.net, Bitcoin mining consumes an estimated 70 TWh per annum, greater than the total electricity consumption of the Czech Republic.

The Ethereum network attracts less scrutiny in terms of electricity consumption but is no slouch either, estimated at 20 TWh, approaching that of Nigeria with its population of 195 million people.

Proof of Stake

For several years, Vitalik Buterin, the Prime Mover of Ethereum, has been advocating/planning a change from Proof of Work to a Proof of Stake consensus model code-named "Casper". The rationale for this massive conversion is partly framed as a strategy to reduce the exorbitant ecological costs estimated above, aptly if dismissively characterized as "it kills trees". The more compelling impetus (based on word count and tone) is framed as a matter of realizing the Cypherpunk philosophy of thwarting the impertinence of "state level actors" who might be so presumptuous as to "attack" the "castles" erected by their betters, i.e. the Cypherpunks and their acolytes. In the Appendix I will explore this ethos which—unlike the various vague straw men economic arguments offered as justification for blockchain/DLT—has always been the prime directive underlying every financial cryptography-based scheme over the past 25 years.

Achieving consensus via Ethereum's proposed PoS model means "a set of validators [anyone holding a balance of ether] take turns proposing and voting on the next block, and the weight of each validator's vote depends on the size of its deposit (i.e. stake)". The stakes that have been advanced are locked up pending consensus. Validators that choose what turns out to be the consensus as to whether or not any given block should be part of the chain are rewarded, and/or validators that had advanced their stake in favor of a naughty block may be punished (by forfeiting their stake).

Other consensus models

The blockchain/DLT literature is replete with references to the Byzantine Generals Problem, a favorite concern of game theory aficionados – a set of which Cypherpunks comprise a proper subset. In this predicament, "a group of generals of the Byzantine army [are] camped with their troops around an enemy city. Communicating only by messenger, the generals must agree upon a common battle plan. However, one or more of them may be traitors who will try to confuse the others. The problem is to find an algorithm to ensure that the loyal generals will reach agreement." The situation is further complicated due to the failure of some messages to arrive in a timely fashion or at all.

The Byzantine Generals Problem is the epitome of untrusted nodes and network partitions. The consensus model used in Ripple is described as a novel algorithm that does not require synchronous communication and "maintains robustness in the face of Byzantine failures".

Probabilistic Finality

Suppose a node or syndication of nodes has won the contest bestowing the right to add the next block to a blockchain[11]. Is the transaction final at that point? Well, no.

As explained on the Ripple website, "In the process of determining the final, authoritative version of the ledger, there may be multiple temporary internal versions. Such internal versions will happen in [the DLT version of] distributed systems because not all nodes will receive transactions in the same order".

The original solution to this, pioneered by Bitcoin and adopted in multiple other blockchain schemes, makes the final('ish – see below) determination of which version of the blockchain is legitimate and which bogus on the basis of length. If two (or more) mining nodes had won the PoW lottery at about the same time, both/each would have written a block to the blockchain and, very likely, other blocks would have subsequently been appended to both/each version(s), creating multiple candidate blockchains.

Now suppose a transaction has been written to the blockchain that entails a decrement to a particular wallet. The person controlling the paying wallet may have also tried to spend the same money simultaneously (a "double spend") to some other recipient wallet address (perhaps controlled by the payer herself) and a block containing this transaction was appended at almost the same moment by a different miner creating a different version of the blockchain. Which version is correct?

The correct version is the one in which the money was spent the first time, and whichever contains the attempt to re-spend the same money is bogus. But which was first? [Recall the importance of a sequential record.] The logic is that whichever chain subsequently became longer, where a larger number of additional blocks have subsequently been appended, contains the legitimate spend.

How many blocks deep – i.e. how many additional blocks should have subsequently been appended – before a block can be considered final and authoritative? The answer is a matter of statistical probability. As Buterin explains, once a block is six blocks deep the statistical likelihood of it being reversed is so remote as to warrant regarding it as final. This is what probabilistic finality means.

How long does it take for a transaction to reach six block depth? With Bitcoin, average block time is about 10 minutes so one hour may suffice. With Ethereum, a new block is appended every 10-19 seconds enabling reasonable confidence in a mere one or two minutes.

Buterin acknowledges that public blockchains-involving an unknown number of widely dispersed nodes- can never match the lower latency (delay) of permissioned and possibly even collocated networks, which of course would be private blockchains. The possibility of anyone considering a permissioned system that does not involve blockchain is self-evidently too ridiculous to seriously countenance and is dismissed with the usual fleeting straw man jibes. A link is provided referencing the bumbling security arrangements of the central bank of Bangladesh as definitive proof that "there is no system in the world that offers truly 100% settlement finality in the literal sense of the term" and further cements the case by noting that even a "paper registry" could be corrupted by a hooligan altering 'every "1" to make it look like a "9"'.

[11] With unbacked—including all "mined" —cryptocurrencies, the winning node is rewarded with transaction fees and, if the protocol entails mining, some amount of newly created (unbacked) money.

Thesis: the need for consensus is a systematic flaw

The complexities and costs of all the described consensus models, plus consensus protocols used in private blockchains as touched on below can be avoided by means of a centralized administrative entity, orchestrating the activities of the potentially thousands of nodes in a well-engineered network architecture.

The reality is that the architects of most major current blockchain/DLT models recognize this and are actively revising their systems to implement permissioned (or what amounts to permissioned) validation mechanisms.

As noted above, Ripple already relies on company-designated "validators" and is therefore not regarded by DLT purists as a decentralized system.

Casper, Ethereum's promised PoS schema with its "pre-registered" validators (node owners that have posted a stake), is an obvious step toward a permissioned model.

Certain pending Bitcoin hard forks, notably the long-awaited Lightning Network, will implement a Delegated Proof of Stake (DPoS) model that may integrate a "reputation system" for which "a panel of trusted parties has to be established".

Yet, as noted in the next section, even as these systems morph into permissioned consensus models, the fact that the consensus of multiple participants is required at all assures decreased efficiency relative to a well-engineered system that does not have any baked-in requirements for consensus.

Consensus in private blockchain systems

In the introduction to the consensus discussion above I made passing reference to private blockchain arrangements enabling a group of banks to perform clearing and settlement of payments in a decentralized fashion.

A recent report relating findings of Project Ubin, Phase 2—an ambitious pilot study conducted under the auspices of the Monetary Authority of Singapore—provides the most detailed technical description of such an arrangement yet published. Appendix 2, appended below, is my analysis of this report.

To quickly summarize, the unnecessary complexity of a system built in accordance with the imperative to operate in a decentralized fashion—where "decentralized" means that all nodes are controlled by independent participants and no central administrative entity is designated with the responsibility and ability to coordinate the actions of all participant nodes—results in inefficiencies that call to mind a Rube-Goldberg machine. This is addressed below in the general discussion of private blockchain/DLT schemes.

Practically any well-engineered system—a notable example being CHIPS, operated by the largest clearing house for domestic USD payments—built as a centrally administered platform is superior to systems predicated on the imperatives of blockchain/DLT. No consensus required (except for the relatively trivial technical process when/if a two-stage commit is involved). Just process the transactions and record them to a well-secured immutable sequential record.

Propagation and consistency

It is evident that a single computer (or a single storage device in a computer such as a hard drive) may fail and it is a good idea to have data stored several places. All modern systems therefore maintain multiple copies of data, in echelons ranging from multiple hard drives controlled by a single computer, to a cluster of computers that are collocated with each comprising a node connected by a local network, to computers/nodes that are geographically dispersed (commonly with a cluster at each location) and connected by a wide area network that may be the internet.

We have noted the importance of maintaining an appropriate degree of consistency among nodes via system processes to implement, and synchronize in accordance with, an unambiguous SSOT. And we have observed that in most modern systems an eventual consistency logic solves a lot of problems provided that the latency for attaining eventual consistency is not too prolonged.

We have also touched on the fact that networks run into all sorts of problems that result in partitions such that some nodes in the networked system may become unable to communicate with other nodes.

My point now is to start highlighting the contrast between blockchain/DLT networks, whether public or private—which, per blockchain/DLT doctrines are not supposed to be centrally administered but rather all participant nodes are created equal—and distributed systems that are centrally administered.

Partition attacks

Discussion of vulnerabilities that could corrupt the Bitcoin network or other public blockchains relying on PoW often focus on Sybil attacks in which an attacker gains control of > 50% of the network's computing power enabling it to corrupt the entire blockchain.

But blockchains also are subject to problems arising from network partitions. A significant network partition with blockchain/DLT leads to hard forks. This is bad enough if it happens due to natural causes but is also a vulnerability for malicious exploits.

Suppose some portion of the Bitcoin network is cut off from the rest of the world. For example, China could elect to implement countermeasures to thwart people using cryptocurrency systems to evade capital controls. Instead of using the vaunted "Great Firewall of China" to only block access to exchange services and ICO mongers, it could potentially be tweaked to intercept all cross-border blockchain-related traffic. The isolated portion of the network would carry on as usual, appending blocks approved by the majority of computing power of still-visible and reachable nodes but the in-country blockchain would immediately diverge or "fork" from the blockchain the rest of the world sees.

In 2015, Ethan Heilman of Boston University and his collaborators detailed a different threat involving network partitions, an "eclipse attack" in which rather than hijacking the entire blockchain, an attacker isolates a targeted node or group of nodes, partitioning it off—blocking or "eclipsing" it—from the blockchain the rest of the world sees. This could be done by "an adversary controlling a sufficient number of IP addresses to monopolize all connections to and from a victim bitcoin node." This is not a remote theoretical issue, as the authors empirically confirmed "while botnet attacks require far fewer IP addresses, there are hundreds of organizations that have sufficient IP resources to launch eclipse attacks".

An alternative category of attack strategies designed to partition (or slow) the network for malicious purposes is "routing attacks". A study performed under the auspices of ETH Zurich found that " >90% of Bitcoin nodes are vulnerable to BGP hijacks". The researchers built their own "virtual AS [Autonomous System] with full BGP connectivity using Transit Portal" and launched an attack against six Bitcoin nodes under their own control, concluding "Intercepting Bitcoin traffic using BGP hijack is fast and effective: all the traffic was flowing through the attacker within 90 seconds. Results computed while performing an actual BGP hijack against our own Bitcoin nodes".

Bitcoin, being the most prominent blockchain, has received the most attention with respect to exploits that entail partitioning but Ethereum is also vulnerable. Another study out of ETH Zurich, "Ethereum Eclipse Attacks", described three vulnerabilities, one of them a "consensus critical vulnerability" via which "an attacker with limited capabilities can easily sustain an attack on the whole Ethereum Network".

Interestingly, a team at MIT CSAIL (Computer Science and Artificial Intelligence Laboratory) headed by Silvio Micali[12] has announced a blockchain protocol called Algorand which is being commercially implemented by a partnership with Toda […].

Algorand is described as immune to Sybil attacks and malicious partitions. Normal operation achieves liveness (finality in < 1 minute) via a consensus protocol dependent on a "strong synchrony" condition[13]. Algorand resolves consistency vs. availability tensions in favor of consistency. Inability to achieve and propagate consensus (thereby achieving consistency across the entire network) within a specified time interval ("(e.g., at most 1 day or 1 week") requires the system to stop appending blocks pending a recovery period "(e.g., a few hours or a day") throughout which the system's "strong synchrony" requirement must again be met and maintained. Coordination of recovery is also dependent on all users being in rough synch with respect to time by means of NTP (Network Time Protocol) or something similar.

In all cases, achieving the right balance between consistency and availability is a challenge for networked systems. In the following section I will introduce some of the advantages afforded by a system orchestrated by a centralized administrative entity vested with responsibility to assure reliable performance and liable when something goes wrong.

Managed distributed network resources

Central administration of resources deployed on a distributed network enables both an initial and ongoing optimization of network configuration, providing the ability to not only monitor network health/performance but to also make timely adjustments.

As will be discussed in relation to scalability, the advent and subsequent explosive growth and evolution of cloud computing has directly enabled and catalyzed collateral proliferation of hitherto inconceivable capabilities supporting design, deployment, operation and maintenance/enhancement of distributed computing strategies. Three developments that have emerged in concert with this phenomenon warrant comment as each illustrate the advantages of someone being responsible for and empowered to manage a system.

Devops

Devops has been described as "a methodology for helping organizations build teams and build software. DevOps is a culture, movement or practice that emphasizes the collaboration and communication of both software developers and other IT professionals [network architects, system administrators, product

[12] Micali, in collaboration with another renowned cryptographer, Ron Rivest, had previously launched a company that implemented a micropayments protocol called Peppercoin. Peppercoin raised $15 million, failed to attract significant usage and went out of business in 2007.

[13] Liveness (finality in <1 minute) rests on 'a "strong synchrony" assumption that most honest users (e.g., 95%) can send messages that will be received by most other honest users (e.g., 95%) within a known time bound'.

"owners" involved with business aspects of a product or service] while automating the process of software delivery and infrastructure changes.

Note the reference to "organizations". Contrast this to the ad hoc coalitions that seek to advance the art of public blockchain/DLT systems. Intractable disputes, sometimes involving years of impassioned wrangling, lead to hard forks as factions lack a responsible adjudicator with whom the buck stops. And when catastrophic misadventures occur, such as with Ethereum's Parity bug, no one is liable for the consequences.

Continuous Delivery

"Continuous Delivery is the ability to get changes of all types—including new features, configuration changes, bug fixes and experiments—into production, or into the hands of users, safely and quickly in a sustainable way". Lacking an organization with defined lines of authority, open source collaborative projects are at risk of someone releasing a change into the wild that has not been properly tested, resulting in adverse unanticipated consequences such as the DAO debacle.

Infrastructure as Code (and Configuration as Code)

Large distributed systems may involve thousands of nodes scattered across different continents. Fluctuations such as peaks in demand may warrant timely interventions to bring additional resources online. Localized problems such as equipment failures or network partitions often require mitigation.

In response to this reality a new industry has emerged of vendors competing to provide utilities enabling management of complex network architectures by programmatic means. Contrast this with the recurring crises with Bitcoin, where the inability to scale leads to recurring crises where getting a payment to go through might require payment of exorbitant fees and still end up taking hours or days.

Build or buy

An enterprise undertaken on a commercial basis by identified principals who are responsible for actions and outcomes does not have to re-invent the wheel for the entire code base. A system may take advantage of multiple technologies provided by third party vendors addressing the specialized functions enumerated above (or a multitude of others). This can occur to some degree with open source projects that are not funded by a sponsoring organization, integrating modular components that do not require payment of licensing/subscription charges. The responsibility aspect may come into play again, however, as component elements integrated into the commercial product may be warrantied and supported while freeware imposes no liability on anyone.

Performance issues with blockchain/DLT

Blockchain/DLT implementations have significant performance issues that are intractable unless protocols are modified to make them less blockchain'ey. The most glaring deficiencies arise in regard to:

- Efficiency of data access/retrieval,
- Scalability.

Efficiency of data access/retrieval

As noted above, query performance in systems with a "trusted committer" is better than in systems lacking one[14]. The problems with data access/retrieval with blockchain/DLT become more pronounced when large volumes of data are stored in a persistence structure, that is, when you approach the domain of "big data".

There has been relatively little attention to this in the various accolades that pose as critical evaluations of blockchain/DLT. But the issue is this – for large amounts of data, adherence to the requirement for cipher block chaining degrades performance of data access and retrieval processes. The only way to overcome this is to refactor data into a storage model and structure unencumbered by the cipher block chaining constraint.

One company that has focused on this aspect is BigchainDB GmbH, headquartered in Berlin. Not wishing to bite the blockchain-enthusiasm hand that feeds them, the company is quite coy in its documentation, never quite saying directly that the way they achieve reasonable query (and write) efficiency is by putting data into a real database and ditching the whole cipher block chaining routine for the stored blocks which contain the transactions[15]. The solution is promoted as "taking a database and adding blockchain features" which include the critical characteristics—decentralized, immutable and the ability to register and transfer assets—of virtually any modern transaction system but skipping the defining element of blockchain, the cipher block chaining of transactions/blocks.

This is not to say the BigchainDB solution is bad – it is likely highly superior to the futility of relying on actual blockchain data structures for big data.

Scalability

In the past decade, revolutionary advances have been achieved in technologies supporting distributed transaction systems. During roughly the same interval that blockchain/DLT strategies have dominated discourse with respect to new approaches to money and payments, the giants of the Internet (companies such as Google, Amazon and Facebook) were forced to face and transcend growth crises stemming from previous impediments to massive scalability. These advances in scalability both exploited and engendered concurrent advances in distributed cloud computing arrangements.

Blockchain/DLT systems embody legacy principles such that the only way for them to take advantage of these revolutionary advances is to deviate from their core imperatives, in other words to become less

[14] Ostrovsky et al acknowledge a "trusted committer" can achieve query response latency of "O(N logd−1 N), as opposed to the O(N logd N) time taken by the [various Merkle Tree-organized] protocols of this paper". The "trusted" system used for comparison was a cryptographically "authenticated data structure" abstracted by Goodrich resembling the alternative approach to making transaction data tamper resistant described above. The trusted comparison system implements what Goodrich terms a "naïve approach" in which "Queries to collection S are handled by a single trusted server ST that gives a digitally signed response to every query". Better yet, performance-wise, if a digitally signed report is only requested for <0.5% of transactions.

[15] Cipher block chaining comes into play with the BigchainDB scheme only with a separate data storage structure containing "votes". Overall, other than the chaining of votes, the concept is reminiscent of the digital cash protocol implemented by DigiGold around the turn of the century in which quantities of an asset defined in a digitally signed contract (using public key cryptography) were conveyed via digitally signed transaction instructions. The technology powering Digigold also foreshadowed the technological foundation of R3 Corda.

blockchain'ey. The proposed alternative system described below, in contrast, was designed from the ground up as an implementation that integrates these advances which can be categorized as:

- Virtualization strategies,
- Three-axis scalability architectures.

Virtualization

Cloud services offered the advantage that a large number of computers belonging to a third party provider could effectively be rented, as a service, eliminating a capital cost that could only be depreciated on an amortized basis, thereby converting it to an expense. This also meant that instead of a company bearing the sysadmin overhead of making sure that each server and network appliance (firewalls, load balancers) was properly configured, secured and kept current with respect to updates and patches, such maintenance headaches became someone else's problem. A cloud services provider's failure to meet service level agreements, leading to downtime or in some cases security issues exposed the provider to liability, serving to partially indemnify the entity making use of cloud services.

Virtual machines (VMs)

But renting entire machines was costly. Virtualization strategies were therefore implemented enabling the provider to host multiple servers—potentially servicing multiple customers—on a single machine, passing cost savings on to the customers. A decade ago, the most common virtualization strategy was to create "virtual machines" which used only a portion of the computer's processing and storage capacity. Virtual machines for several customers could be hosted on a single machine, isolated from each other so as not pose a security or performance risk, preventing a faulty application of one customer—perhaps so buggy as to cause their virtual machine to crash—from adversely impacting the virtual machines of other customers on the same server.

Containers

More recently, VMs have been giving way to "containers" as the preferred virtualization strategy for massively distributed deployments (or more modest systems). A container is "lightweight" relative to a VM. Whereas each VM requires its own "kernel" and operating system (OS), containers "run at a layer on top of the host OS and they share the OS kernel" resulting in "much lower overhead than VMs and a much smaller footprint".

Containers are loaded with "images" of selected resources – optimized lightweight versions of other applications such as database systems. This facilitates synergies when combined with other techniques such as the "microservices architecture" discussed below.

Containers are also better suited for new network management tools that offer efficiencies exceeding those of industrial automation and control systems used in semiconductor or automobile manufacture (since there are no mechanical components to worry about). Infrastructure as Code management utilities, which were relatively novel as late as 2015, have rapidly evolved in concert with adoption of container strategies giving rise to the emergence of previously unimaginable Configuration as Code capabilities.

Of course the very idea of a managed network, with its implication of centralized administration, is contrary to the "decentralized" dogma of blockchain/DLT. To preview an assertion that will be developed and defended as we proceed; while all the other core conceits of blockchain/DLT – blocks, chaining, every-node-does everything – will continue to be abandoned, all that will remain is the same wrongheaded imperative that has been a consistent flaw in all Cypherpunk monetary and payment schemes since that community first formed in the 90s. The core sine qua non of Cypherpunk philosophy-imbued digital cash and payment mechanisms is the avoidance of client-server arrangements, at any cost, in favor of having multiple network participants (ideally all) control and maintain their own node, no matter how complex and expensive.

But first a survey of other recent advances enabling massive scalability is in order.

Scale cube

The trope of a "scale cube" was first popularized by Abbott and Fisher, network architecture consultants and authors of "The Art of Scalability" a comprehensive introduction to organizational, process and architectural principles enabling design, deployment, operation, maintenance and advancements of distributed and massively scalable computing systems.

X-axis – multiple nodes, each doing the same thing

In the three dimensional scale cube visualization, the X-axis represents scaling strictly by duplicative means. In a network under centralized administration this might entail running multiple clones behind a load balancer. Each clone does the same thing and may have either access to, or maintain it own copy of, the same, perhaps all, data. This arrangement provides redundancy, protecting against data loss and, in a centrally administered arrangement, may facilitate scaling. If there are N machines (such as webservers), workload can be divided up (e.g. by the load balancer) such that each machine handles 1/N of the load.

Growth of a canonical blockhain/DLT network is limited to X-axis scaling but the notion of every node processing everything and holding a complete copy of all data defeats the possibility of dividing workload. Workload is instead multiplied.

Y-axis –microservices

Y-axis scaling entails breaking up a "monolithic application", decomposing it into smaller specialized "services", perhaps going so far as a "microservices architecture". This approach affords multiple advantages such as the ability to break up a large and highly complex code base enabling separate teams to work with more tractable components. It also facilitates a granular approach for identifying and isolating problems such as might occur with a new software release that turns out to not run as smoothly as during testing. The ability to implement Y-axis scaling can be integrated with other axes so that for instance a service being used more heavily in particular data centers can be ramped up via X-axis addition of additional container-housed instances.

As noted above, the interplay of microservices and containers affords additional flexibility and efficiencies enabling, for example, use of multiple different data persistence strategies in the same overall application. A component service processing customer-provided inputs such as a membership application form or shopping cart may use a main memory database to hold transient information such as a partially filled-out form, enabling an online user to pick up where she left off if an online session is interrupted rather than starting over. The containers where this service runs would hold an image of the essential portions of the main memory database application but not images of irrelevant resources.

Document, relational or graph databases may be better suited for other portions of a system organized as microservices each running in separate containers.

Private blockchain schemes, discussed below, focus on where the money is, that is, on selling expensive systems to banks. All of these competitive providers are scrambling to implement microservices into their systems. R3 Corda emphasizes "Nodes [author's note: all belonging to and/or operated under the administrative control of the same, typically bank, customer] are logically structured as a series of microservices and have the potential in future to be run on separate machines[16]". As of 2017, the "bleeding edge" initiatives of Hyperledger Fabric, a consortium backed by IBM and Ripple, entail refactoring to use containers and microservices. Quorum, JP Morgan's "enterprise-focused version of Ethereum" may be a little behind on the microservices front but other vendors in the Ethereum ecosystem are vying to fill in this gap.

Z-axis – sharding

Sharding is an extension of "horizontal partitioning" – breaking up a table in a database into smaller tables. With horizontal partitioning, the data model of each resulting table is the same (e.g. the same columns/fields) and the partitioning is a matter of splitting up the rows leading to fewer rows per table. The original concept contemplated a single database, generally on a single machine, where queries that needed data from more than one of the now-separated tables required a relatively resource-intensive UNION operator.

Sharding extends the concept by enabling data stores to be split among multiple nodes, each with responsibility for only a subset of data. Just as with X-axis scaling, load—in this case quantities of data as well as read and write operations—is distributed.

Leading vendors of database systems, whether relational, document, main memory or supporting multiple database types offer native support for sharding, some enabling advanced capabilities such as resharding a production database without interruption or degradation of service.

A system for persisting data as an immutable sequential record must implement sharding in order to transcend limitations of scalability. But sharding is absolutely contrary to the core tenet of a decentralized network of mutually untrusting nodes – that to prevent any node from monkeying with the data, all nodes must validate and maintain a copy of all transactions. The Sharding FAQ on the Ethereum wiki provides an exquisite snapshot of the squirming currently going on throughout all of blockchaindom trying to deal with this reality. Developing solutions to this and other problems is the dynamic that absolutely assures that blockchain/DLT must morph into something that does not involve, or minimizes the role of, blockchain per se. It will become some other thing, though likely still referred to as blockchain.

New directions with blockchain/DLT

Off-blockchain transactions

Given the slowness, inability to scale and other inefficiencies of Bitcoin and other first generation blockchain protocols, one strategy is to execute many or most transactions "off-blockchain".

Existing off-blockchain workarounds

Blockchain/DLT began with Bitcoin. As Bitcoin began to attract participants and exponentially increasing—typically sympathetic or outright enthusiastic—media coverage, other blockchain schemes began to proliferate, primarily focused on "cryptocurrency" alternatives. Participation in a blockchain-based

[16] The whole point of microservices as a scalability strategy is to distribute processing among multiple nodes. i.e. "separate machines".

cryptocurrency in these early days required installation and provisioning of a wallet which, especially initially, meant climbing a significant learning curve and a lot of messing around. Early adopters needed to either be techies or to seek out the assistance of one.

As the exchange rate of Bitcoin began its meteoric rise, vaulting over the lofty milestone of $10 per unit in the summer of 2012, the broader public began to trickle in (later to flood), enticed by the prospect of overnight riches. The broader public, however, is disinclined to engage in inconvenient exertions so a host of entrepreneurial ventures sprang up, offering more convenient options, especially for the vast majority of people who were only seeking to cash in on rising exchange rates via speculation in currency markets.

People are accustomed to interacting with online transaction capabilities by logging into a website with their browser; the whole downloading, installing, provisioning and operation of wallet software (i.e. a node) on their own computer was a no-go. Online wallets became a thing but the preferred option for the majority of people was simply to establish an account with a provider of exchange services, trusting them to handle all the complicated bits.

Having ditched the core trust-abhorring imperative that would warrant each participant operating their own node, the blockchain-based transaction mechanism was the remaining hassle to circumvent. Waiting at least an hour to know if a Bitcoin payment could be relied on was as unacceptable for currency trading as it was for merchants trying to use Bitcoin as a means of being paid. Thus did “off-blockchain” transactions begin to become the dominant mode for most cryptocurrency transactions.

These first generation off-blockchain arrangements basically entailed some entity such as an exchange service or platform holding customer value on account and, similar to a bank, keeping a separate set of books with deposit liabilities backed by assets, usually specified as being actual Bitcoin in actual Bitcoin wallets[17]. Trades, involving prefunded account balances, could be matched, cleared and settled immediately by book entry effected by whatever system had been implemented by the provider.

Cryptocurrency arrangements had become systems as riddled with trust as any legacy system but these new trusted intermediaries commonly neglected to provide the contractual protections, transparency and other safeguards of well-designed and/or regulated financial intermediaries.

Alarmed by the flakiness of such arrangements and also censorious of unregenerate people failing to conform to the Cypherpunk imperative for decentralization[18]—which mandates that participants operate their own nodes—new protocols were proposed. These new protocols, the Lightning Network being the first to attract attention, would provide off-blockchain transactions in a more systematic way, providing several benefits:

- The “severe counterparty risk” arising from “extreme cases” of “privileged, trusted parties”, particularly “centralized trusted custodians who have full custody of customers’ funds” would be

[17] Exchange services that hold customer value on account also hold deposit balances denominated and payable in conventional money. While various arrangements might be implemented to assure that the cryptocurrency deposits are backed by a 100% reserve of the native cryptocurrency, arrangements for securing customer deposits of conventional money have attracted hardly any scrutiny.

[18] In keeping with the overtones of religious zeal of the DLT community, its preeminent prophet, Vitalik Buterin recently expressed “I definitely hope centralized exchanges go burn in hell as much as possible.”

avoided (except for the aspect of also holding deposits denominated and payable in conventional money),

- The already painfully apparent inability of Bitcoin or any other blockchain scheme to scale would be circumvented,
- Rather than taking minutes, hours or, in some cases, days for a transaction to be acceptably final, transactions (at least small ones) could be almost instant,
- The need for cumbersome (literally "Byzantine") consensus mechanisms, or for bundling transactions into blocks, or for chaining transactions together in an immutable record or, in many cases, any record at all would be avoided.

It had been revealed that "Real blockchain transactions" don't need to be broadcast to all nodes or persisted in an immutable sequential record. For the vast majority of routine transactions "it's not necessary for all other nodes in the bitcoin network to know about that transaction. It is instead preferable to only have the bare minimum of information on the blockchain".

This new paradigm for off-blockchain transactions is called "payment channels".

Payment Channels

Ripple was actually first to implement "Payment Channels", described as enabling "small, rapid off-ledger payments of XRP that can be later reconciled with the consensus ledger"[19]. Bitcoin's long-awaited Lightning Network[20], based on payment channels and initially framed as a micropayments solution, made its debut in 2018, achieving a total capacity by July of around 100 BTC. The Raiden Network, also initially emphasizing micropayments, describes itself as "Ethereum's version of Bitcoin's Lightning Network".

As with the blockchain itself, payment channels schemes can be superficially summarized with marketing-oriented buzzwords but developing a grasp of the actual nuts and bolts is similarly challenging. For example, an analogy is drawn by the explainers of various channel schemes with the TCP/IP protocols/stack that make the internet work (though "Lightning is Bitcoin's TCP/IP stack" drew the prompt rebuttal "No, Lightning isn't Bitcoin's TCP/IP stack").

"Channels" are the building blocks out of which payment channels are formed. With the Lightning or Raiden model, a channel is an on-ledger mutually prefunded contractual binding of two nodes, where the prefunded amount collateralizes transfers that traverse that channel. The Interledger Protocol (ILP), in contrast, does not specify funding arrangements beyond requiring that "Sender and connector have funds on some shared ledger or settlement system to rebalance their accounts with one another". ILP more closely resembles SWIFT in the sense that successful processing of a payment message results in contractual obligations to pay between nodes/intermediaries but movement of funds[21] to satisfy the obligations is a separate matter. In all channel protocols, payments traverse the network from payer to the designated recipient via routes that may involve multiple hops, i.e. traverse multiple component channels. Routes are discovered via a routing protocol involving "channel graphs" which are analogous to

[19] Ripple added Payment Channels to XRP Ledger to enable XRP integration into the then-nascent Interledger Protocol.

[20] Not to be confused with Lightning Bitcoin, yet another "hard fork"-introduced Bitcoin variant that just possibly may be seeking to cash in on the enthusiasm for the much-anticipated Lightning Network.

[21] With SWIFT, the funding component is called "cover", while ILP refers to "Funding and rebalancing".

the routing tables used by routers for routing packets through the internet (hence the TCP/IP reference cited above).

Promoters continue to use the "chain" metaphor to describe the ephemeral path by which a Lightning Network transaction reaches its destination but the chain evaporates whenever any component channel is closed – the only means for one of the nodes bound in that channel to recover actual Bitcoin that can be spent on-blockchain.

Transactions sent via the Lightning Network are not recorded to the blockchain. In fact, other than a log or transaction register in the sender's or end-recipient's wallet, no copy (even a backup) is necessarily recorded anywhere. The only transactions written to the blockchain are those required to open or close a channel[22]. As a result, if Lightning Network catches on (and works), it is possible that subsequently only a small subset of Bitcoin payments would ever end up committed to the associated blockchain. No persistent record of, or (perhaps) even means to indirectly infer, the most basic elements (payer/recipient address or amount) of the (anticipated) vast majority of payments would exist.

There are several reasons for the emphasis on micropayments[23] but a major one stems from the fact that the amount of value that may traverse any component channel is limited by the prefunding committed to its defining contract (which in Lightning Network is) called a Hashed Timelock Contract (HTLC)[24]. Another motive is anonymity. A third motive, touched on below, is fulfillment of the decades-enduring Cypherpunk dream of streaming payments, similar to how content such as video is streamed over the internet.

While little discussed so far in popular media (and, as per usual, scarcely registering on the radar of regulators), use of payment channels such as Lightning Network may provide an unprecedented level of anonymity relative to their underlying cryptocurrency[25]. Anonymity will be discussed below but, in partial summary, first generation systems such as Bitcoin, advocated single-transaction usage of a new wallet address for every incoming payment. This practice, combined with other measures, enabled a motivated/disciplined user of Bitcoin (such as a criminal or terrorist) to achieve strong anonymity. With Lightning Network, anonymity (lack of auditability or any permanent record) is built in and further reinforced by use of "onion-style routing" as used with Tor. The degree of anonymity that will actually emerge depends on how the Lightning Network evolves. If, as the Lightning Network architects prefer, an almost innumerable multitude of channels proliferates—all supporting modest value traverses—the system could approach Tor-like anonymity. Alternatively, should well-endowed nodes/channels supporting large transactions dominate, as has occurred with concentration of Bitcoin mining power, their de facto "hub" nature could be used to influence the routes selected by other nodes, facilitating efforts

[22] Closing an on-ledger channel commits a transaction to the blockchain in the amount of the net result of all transactions received, made or otherwise passed from/to or across that channel and is recorded as if coming from (if the net is positive) the directly-channel-linked node.

[23] The term "micropayments" has been used over the past two decades to describe a wide range of values but with blockchain based payment channels it implies infinitesimal (even-"sub satoshi") values.

[24] Raiden does not use a time-based logic for its hash-locked smart contracts that create a channel, but rather employs cryptographic "balance proofs".

[25] Counteracting the tendency toward anonymity would be the incentive to use a particular wallet address, once connected in a channel, for multiple transactions, both for oneself and as a conduit for others, in order to avoid the potentially steep transaction fee for the on-blockchain transaction required to open or close a channel. Additionally, there would be a strong tendency for centralized "hubs" to emerge.

to analyze traffic patterns and infer value flows[26]. Innumerable streams of teeny micropayments would be the most anonymous outcome.

While micropayments as micropayments are viewed as a worthwhile use case, it is intended that they play a role in more economically useful payments. For this to occur, a more commercially sensible (normal size) payment might be broken into a stream of sequential micropayments[27]. For both the payer and recipient of payment, however, the maximum value that can be sent or received is limited by the combined prefunded sums directly committed by the end node and its directly channel-connected, network-adjacent nodes.

It is not yet clear how payment channel schemes will play out though that has not deterred intrepid early adopters from plunging in to the deep end of the pool. One intensively discussed analysis (attracting comments by some of the biggest names in the blockchain development community) claims to offer "Mathematical Proof That the Lightning Network Cannot Be a Decentralized Bitcoin Scaling Solution". This was followed in short order by the also hotly-debated but well argued "Continued Discussion on why Lightning Network Cannot Scale".

While not involving payment channels per se, the Interledger Protocol (ILP)—a scheme for payments potentially spanning multiple ledgers—also entails breaking payments into tiny chunks and routing the packets through a path involving multiple nodes (called "connectors"). The ILP is discussed below in conjunction with, and in contrast to, the proposed alternative approach for intermediated payments

To preview a proposition I will examine in further detail in the Appendix, in many ways the sound and fury of ongoing payment channels debates reflects the conflict between one of the most fundamental ramifications of Cypherpunk prescriptions…and practical reality. As I first pointed out in April 1997 to the consternation and scorn of a mailing group frequented by Cypherpunks, Cypherpunk schemes embody an abhorrence of client-server arrangements[28]. As with the payment channels topic, for the full beauty and benefit of their visions to be realized, a multitude of individuals must undertake the hassle and inconvenience of setting up, operating and maintaining their own nodes, thus achieving the "decentralized" ideal. In reality, most people will never care enough to undergo such exertions. Cryptocurrencies have attracted a lot of discussion, enthusiasm and dabbling but that has been almost entirely an artifact of the speculative impulses aroused by the cryptocurrency bubble that emerged as an adventitious consequence of: (a) the 2008-09 exit of e-gold from the market, and, (b) unprecedented alternative monetary policies undertaken by central banks worldwide in response to the GFC which have led to a global money glut and bubbles in every class of traded assets.

There may be a connection between this speculative bubble phenomenon of cryptocurrencies and another aspect touted as an advantage of payment channels. It is widely claimed that, due to the much

[26] Vulnerabilities that might enable a "well funded adversary to manipulate the availability of channels on intermediary systems to influence or control route selection" are intensively discussed here.

[27] In the case of a payment streamed as a (potentially very large) number of micropayments, it is not entirely clear how the recipient knows when a lump sum payment is complete. This may be further complicated if a route is interrupted during lump sum payment such that packets for the same payment end up arriving via different routes that entail a different final channel.

[28] This is also reflected in the Cypherpunk conception of "smart contracts" as obligatorily involving code that executes somewhere other than on a server. I will return to this issue in discussion of the alternative system elaborated below, one of the elements of which is "server-based smart contracts".

lower computational, storage and bandwidth overhead involved with operation of a node in a payment channels network, payment processing fees will be vanishingly small or zero. It should be recalled, however, that a favorite argument in favor of Bitcoin in comparison to legacy payment systems used to be the dramatic cost savings that purportedly could be realized. That all fell out of bed somewhere along the line and now a favorite argument for payment channels cites the exorbitant costs, far exceeding those of any precedent system, sometimes charged to settle even a tiny Bitcoin spend (as for instance required to instantiate a Lightning Network payment channel).

It is possible that any low fees that may obtain at first would be a temporary artifact stemming from two factors. There is a community of enthusiastic blockchain participants flush with realized and/or paper gains stemming from the cryptocurrency bubble. This serves to subsidize other efforts to get in on the ground floor of the next big thing. The second factor that arises in relation to the first one is the Ponzi scheme-like ICO phenomenon that has enabled numerous lavishly funded companies to subsidize the costs of developing and bootstrapping new blockchain-oriented initiatives. But at some point, especially if the cryptocurrency bubble bursts, participants might actually be forced to reevaluate and possibly abandon such activities if no sustainable revenue stream emerges sufficient to cover both R&D efforts and operating costs. This will be examined more closely below in the discussion of “institutional arrangements”.

Directed Acyclic Graphs (DAGs)

A graph is a data representation (that lends itself to making cool graphical images) involving nodes and edges. Edges can be thought of as lines connecting nodes. An edge can be bidirectional – think of a line with arrow heads at each end (or no arrow head at all). If an edge is “directed”, however, it means it is unidirectional; there is a from and a to. So a Directed Acyclic Graph (DAG) can be thought of as a connected series of nodes in which the direction for traversing a path, whether from one node to its immediately adjacent one, or involving multiple nodes arrayed in a chain or ramifying into a tree structure, is evident.

In double entry bookkeeping, accounts can be represented as nodes. The edges that connect accounts are directional by virtue of the fact they always entail one or more debits and one or more credits (in accordance with a convention requiring that the debits and credits of a journal entry are supposed to perfectly offset each other). In this sense, DAGs can be thought of as a fancy-sounding abstraction of the core basis of every accounting system since the invention of double entry bookkeeping circa 1458.

The Lightning Network[29] and similar protocols for off-blockchain transactions are implementations of a DAG logic. Transactions routed through the Lightning Network normally exhibit a non-branching chain structure but it would be similarly possible for transactions in a system to take the form of a tree, ramifying and/or converging as system state advances.

The Lightning and Raiden Networks are implemented as a layer overlying a blockchain-based system. But DAGs may also be used as an alternative to cipher block chaining for a cryptocurrency as they too are a means of binding a transaction to all previous transactions. Instead of embedding a hash representing the immediately previous and, via an iterative nesting logic, all preceding system states, a transaction may contain a pointer explicitly identifying the immediately preceding transaction(s), enabling reconstruction

[29] When Poon and Dryja’s first published their Lightning Network whitepaper in 2015, it contained a single instance of the word “graph”. “Directed Acyclic Graph” jargon had not yet acquired its later cachet among the blockchain cognoscenti.

of how each current state came about. The DAG logic also potentially lends itself to a persistence structure that is sequential on a transaction basis, rather than as a sequence of blocks.

This has led to a proliferation of blockchain-like systems, and projects for systems, based on DAGs instead of cipher block chaining. A sampling includes Dagcoin, Byteball, IOTA (emphasizing the "Internet of Things (IOT) with solutions that implement the Tangle), RChain (which has issued to itself a billion RHOC, an ERC20 token, later to be converted to a native token, tentatively referred to as REV) while companies such as DAGlabs offer to develop systems based on SPECTRE and PHANTOM.

Each contender for DAG primacy describes how their system will afford massive scalability, near-instant transactions and lower transactions costs. All DAG-circulated tokens, at least so far, are unbacked and being issued via an ICO model. Each system, like all DLT schemes, is predicated on participants operating their own nodes, requiring some sort of consensus mechanism to deal with the whole Byzantine razzmatazz, and with replication of at least some substantial portion of a ledger to every participant node.

Public vs Private blockchains

The original blockchain schemes were, and largely remain, "public blockchains". This means that anyone may download software, either source code or in already compiled form, enabling operation of a node/wallet without requiring anyone's permission.

With private blockchains, in contrast, the ability to directly participate by operating and controlling a node is "permissioned".

In accordance with this definition (recognizing that some people prefer a different definition), a permissioned superstructure can be built on a public blockchain, typically with smart contracts, but unpermissioned direct participation in a private blockchain is not supposed to happen.

Public – Unpermissioned (anonymous) blockchains

Anonymity

The most important element of the permissioned vs unpermissioned axis is the potential for anonymity baked-in to unpermissioned blockchains (or successor decentralized networks).

Just over a decade ago, the mere accusation that a payment mechanism afforded anonymity, even if unfounded, was considered amply sufficient grounds to portray it as criminal-friendly and warranting criminal prosecution. How times have changed!

Regulators, not wishing to be unhip, and likely pressured by politicians clamoring for Fintech jobs and investment to pour into their jurisdiction, listen attentively to panels exclusively packed with blockchain shills as they look for excuses to take a light touch to regulating such entities. Accordingly, they accept without demur bogus representations that publicly auditable blockchains afford no haven in which would-be criminals can hide. This claim is examined below.

Even government central bankers (who not long ago would also have bristled at unrestrained characterization of their money as "fiat", viewing use of such terminology by people-not-in-the-club as vaguely pejorative) evaluate the potential anonymity of a system as a feature rather than as a show-stopping red flag. Not uncommonly the anonymity of paper cash (previously a baneful horror requiring never ending laws to criminalize cash-oriented business models) is cited as an acceptable precedent, framing the anonymity of cryptocurrencies as similar to that familiar reality.

But anonymous digital currencies are not comparable to paper cash. Anonymous digital currencies enable financial and other crimes with industrial efficiency. Comparing their anonymity to paper cash is like comparing a "1 million rounds per minute" Metal Storm gun with a flintlock dueling pistol.

Criminal exploits involving paper cash expose their perpetrators to all manner of physical risk and inconveniences. Robberies, even if involving multiple victims, are each individual events. The would-be victim of a robbery might fight back or have a gun. The perp or confederate who goes to pick up a ransom payment may, respectively, be arrested on the spot, or, make off with the money herself. Muling money across a border or distributing it to smurfs entails physical logistics.

An exploit involving anonymous digital cash, in contrast, can be tested, tweaked and refined at one's leisure without leaving the comfort and safety of your grandmother's basement. And when the zero-day or moment arrives, the exploit may be unleashed in a manner that affects thousands or millions of wallets (or in the case of ransomware, computer systems) scattered around the globe.

Anonymity is baked in

In a Forbes editorial October 2012, "Bitcoin Prevents Monetary Tyranny", Founding Director of the Bitcoin Foundation[30] Jon Matonis argued "money in a free society should not be used for the purposes of identity and asset tracking. Banks and governments may be concerned with that goal, but it is not the role of our money."

Bitcoin's anonymity, by design (per the Bitcoin wiki), rests on Bitcoin's "officially encouraged practice of using a new [wallet] address for every transaction", minimizing the risk that "any of the addresses in a transaction's past or future can be tied to an actual identity". Elsewhere (in the FAQ), the documentation again emphasizes "Bitcoin addresses are designed to be used exactly once only, for a single transaction." and goes on to describe reuse of a wallet address for more than one transaction, (which would detract from anonymity) as bad practice that "harms both yourself and other unrelated third parties".

Bitcoin developers, perhaps concerned about reports of decreasing popularity with criminals, have been pushing ahead with initiatives such as Dandelion and Mimblewimble, both advanced as affording more robust anonymity.

Ethereum, similarly striving to "increase…usage and value", presses ahead in developing capabilities to make Ethereum "more anonymous". Vitalik Buterin educates blog visitors on the most up to date advances, explaining, for example how the "anonymity-enabling properties" of "zero-knowledge Succinct Transparent ARguments of Knowledge (zk-STARKs)" exceed those of the soon-to-be-passé zero-knowledge Succinct Non-interactive ARguments of Knowledge (zk-SNARKs) used by its strongly anonymous competitor (and collaborator) Zcash (discussed below).

Services to aid in anonymizing

A new industry has emerged as a host of vendors and services vie to "ensure anonymity with cryptocurrency transactions". McAFee's helpful overview of cryptocurrency anonymity resources catalogs tumblers/mixers such as CoinShuffle, JoinMarket, SharedCoin, Jumblr, SharedCoins, Darkwallet and

[30] As Bitcoin prices have skyrocketed and the incentive to tone down rhetoric in order to keep governments at bay has increased, references to "anonymity" in the Bitcoin catechism are being replaced with the euphemism "financial privacy". The Bitcoin Foundation itself has labored mightily to reinvent itself since the days when two of four founding board members— "people who matter to Bitcoin's future"—were Charlie Shrem and Mark Karpeles.

PrivateSend. Protocols and technologies such as CoinJoin, TumbleBit and CryptoNote are also explained to aid bewildered novice freedom fighters and criminals to harden their ability to thwart government snooping.

Disciplined use of Tor encouraged

Tor, an acronym of "The Onion Router", is free software (such as the Tor Browser) designed to make it harder to track people or resources online using their IP number. The Tor Project, launched December 2006 as a Massachusetts-based 501(c)(3) research-education nonprofit organization is primarily responsible for maintaining software for the Tor anonymity network.

Usage of Tor, the Dark Web and Cryptocurrencies have all grown in a mutually beneficial synergy. Since mere usage of the Tor browser is insufficient to achieve criminal-grade anonymity, helpful guides explain the "strong discipline" needed to optimize one's stealth.

As noted above, the Lightning Network and other payment channel schemes implement "onion routing", thereby integrating the core anonymity mechanism that is the basis of Tor.

Anonymity-enhanced cryptos

Monero

Monero is a Cryptonote algorithm based cryptocurrency, optimized to be anonymous and untraceable. It is mined via a Proof of Work mechanism which, similar to Ethereum, discourages use of ASICs (Algorithm-Specific Integrated Circuits) in favor of graphic cards. Due to its strong anonymity, this relative ease of mining using conventional hardware facilitates exploits whereby hackers mine Monero using large networks of other people's computers, allegedly including thousands of U.S., Canadian, U.K. and Australian government websites. Massive Monero-mining botnets increasingly have their own catchy names like Smominru, the Jenkins miner and the newly forming PyCryptoMiner.

The principals of Monero are apparently ok with Monero being used in relation to capital crimes, striving to support a "level of privacy…able to protect users in a court of law and, in extreme cases, from the death penalty".

Since Monero has not yet attained the lofty overvaluation accorded to Bitcoin, a hybrid caper is to receive proceeds of crime in Bitcoin and trade them for Monero, as occurred with the WannaCry ransom proceeds after the news cycle had moved on to other matters.

Zcash

Europol's 2017 IOCTA Report cited Bitcoin as "being the currency of choice in criminal markets, and as payment for cyber-related extortion attempts, such as from ransomware or DDoS attack" but noted "other cryptocurrencies such as Monero, Ethereum and ZCash are also gaining popularity within the digital underground". ZCash, though celebrated for its anonymity, has reportedly failed to gain as much traction in the criminal community as Monero. This may be due to skittishness stemming from the fact that US Federal authorities seized Alphabay, once "the dark web's most popular black market" just one month after it had announced plans to accept ZCash as its currency of choice. ZCash purportedly also does not port as well to "disposable burner phones, a favorite of criminals eager to stay anonymous". ZCash excels in the Burning Man-like mystique department however, as enshrined in its now annual "burning ceremony".

Traceability/transparency is a smoke screen

There are two faces to cryptocurrencies, the anonymous one that makes them so useful to criminals, and the reassuring theater offered to regulators and politicians seeking justification for their so-far virtually hands-off approach to reining them in. This light touch rests primarily on three fictions:

- The fact that everyone can see transaction registers gives would-be criminals no place to hide,
- The anonymity of public blockchains can be penetrated,
- Imposing identity verification requirements on exchange services thwarts criminal gambits to make a clean getaway with proceeds of crime.

Public auditability of limited forensic value

With the WannaCry ransomware attack noted above, the typically blockchain-sympathetic media spin was that the demands for payment to publicly viewable wallet addresses made it "far easier for the security community and law enforcement to track any attempt to anonymously cash out WannaCry profits". A thread on bitcointalk.org showcased the plight faced by the pitiable attackers, posting links to three of the four Bitcoin wallet addresses reported.

An updated peek at these addresses demonstrates the limits to this sort of traceability. As you can see for yourself (click the links and then scroll down to the red bits on 2017-08-03), the criminals waited about 10 weeks, until the Bitcoin price had gone up another $1,000/coin, and then "anonymously cashed out", leaving no breadcrumbs for "the security community and law enforcement":

https://blockchain.info/address/12t9YDPgwueZ9NyMgw519p7AA8isjr6SMw

https://blockchain.info/address/115p7UMMngoj1pMvkpHijcRdfJNXj6LrLn

https://blockchain.info/address/13AM4VW2dhxYgXeQepoHkHSQuy6NgaEb94

The fact that a criminal transaction is displayed on a global public billboard, in and of itself, is worse than useless in terms of forensic value if it cannot be associated with a person. It is analogous to posting video of a beheading to Youtube or Facebook; if done in a fashion that provides no useful forensic data it only flaunts the cleverness of the perp vs the bumbling investigatory efforts of official authorities.

The reality with unpermissioned cryptocurrencies is that a disciplined criminal can be as anonymous as they feel the situation requires.

The anonymity arms race

Media apologists for cryptocurrencies circulate stories such as "Why criminals can't hide behind Bitcoin" in which self-promoting law enforcement agencies and academics utter reassuring claims that "investigators can follow the money". Such articles and the studies they cite are then offered in testimony and panel discussions before regulatory bodies as arguments for special dispensations for blockchain/DLT initiatives, freeing them from the prudential and AML requirements other financial institutions must adhere to.

For example, testimony submitted to the New York Dept. of Financial Services cited Koshy and Biryukov to support a contention that "steps can be taken to de-anonymize [wallet] addresses involved in suspicious [Bitcoin] transactions". But such measures (and those described by Meiklejohn) would be ineffective against disciplined criminals, the only category that matters, practicing basic hygiene such as (in the case of Bitcoin) changing the wallet address and client octet, and adding some delay, after each transaction. Additionally, the Biryukov method depends on first forcing the Bitcoin network to disconnect from Tor, a phase they noted would be "quite noticeable" to the sort of criminals who matter.

As noted above, efforts to provide anonymity, and efforts to penetrate/defeat it comprise an ongoing arms race and, quite frankly, the Cypherpunk community has law enforcement hopelessly outgunned. Anonymity is, and always has been, a cherished core value for Cypherpunks, worthy of their most devoted efforts. These efforts are proliferating and accelerating while the methods of law enforcement can scarcely be propagated to investigative staff before they are obsolete. The only perps who will be caught will be the sloppy/lazy ones, such as the cases endlessly recycled in the presentations of Justice Department and law enforcement officials who show the flag at blockchain trade shows and conventions.

Exchange Portals as the Achille's Heel of would-be criminals

Regulators worldwide appear to believe that regulating the provision of exchange services and portals—venues where cryptocurrencies can be exchanged for real money—will be an effective means of "forcing" cryptocurrency users to "reveal their identities". In the United States, such activities fall within the rubric of Money Transmitting Businesses, a sub-category of Money Services Businesses, the regulation of which is largely a matter for individual states. Most of these efforts focus on requirements for identity verification and transaction monitoring.

There are several reasons why this is a specious fantasy:

- Exchange services/portals cannot determine source of funds
- Detection precedes investigation
- Jurisdictional havens
- There is no effective way to prevent unauthorized provision of exchange services
- Criminals will have less and less need to ever exchange cryptocurrencies for real money.

Determining source of funds

An attempt to evaluate whether an identified individual's or business's cryptocurrency usage is consistent with a legitimate business or personal purpose requires analysis of incoming and outgoing payments to/from that entity. Such an analysis first requires identification of each account or wallet address controlled by or operated for the benefit of the person or entity. However, the anonymity baked into each cryptocurrency makes it impossible to build out a model of the constellation of accounts or wallet addresses controlled by a determined bad actor. The problem is compounded by the prevailing practice of each bad actor using multiple cryptocurrencies to facilitate further obfuscation – part of the "layering" process of money laundering.

A bad actor can safely bind a wallet address to his real world identity-verified customer profile on an exchange service because inflows to that address may be at the end of a bewilderingly complex sequence of obfuscating detours. No exchange service routinely looks multiple hops upstream to evaluate the provenance of value in the addresses linked to customer profiles and trying to do so would be futile.

Achieving an adequate understanding of usage patterns would require integrating data derived from P2P transactions internal to a/each cryptocurrency system with interactions at the boundaries of the system(s), that is, with exchange activity. While some cryptocurrencies expose transactions to public view, no entity is tasked with responsibility to monitor all activity in a systematic way and none have the capability of requiring exchange services to participate in an integrated monitoring system. Additionally, as noted above, second and third generation systems implementing payment channels enable transactions that generate no persistent record.

Detection precedes investigation

Lacking understanding of the business model or source of funds of a customer, stemming in part from lack of visibility of their overall transaction activity, the likelihood of any particular exchange service or portal detecting indices suggestive of criminal activity through its own transaction monitoring is minimal. About the only illicit pattern that can be detected by an individual exchange service is structuring – breaking up transactions into chunks in order to avoid reporting requirements triggered by value thresholds.

Investigations into crimes involving cryptocurrency start by someone looking into a particular account or wallet address used or specified in a criminal transaction or payment instruction. Not even the most rank amateur would use an address associated with their real world identity for overt criminal activity. With any competent disciplined criminal, efforts to follow the money virtually never lead to an address bound to the real world identity of the perp. The odds of a regulated exchange that conscientiously complies with identity verification requirements receiving a subpoena that resulted from successfully following the money are already vanishingly small and will become more remote. About the only subpoenas likely to yield actionable data will be fishing expeditions trolling for tax evasion on the part of unsophisticated speculators.

Jurisdictional boundaries and havens

US states

As noted, regulatory regimes that vary state by state enable a degree of regulatory arbitrage, compounded by the avidity of some states to attract "Fintech" companies. An exchange service operating from a light-touch jurisdiction may also not bother to implement measures to prevent customers from less accommodating jurisdictions from using their system.

International

Sorting out a money trail becomes even more difficult when flows cross international jurisdictional boundaries. Law enforcement, aware of bureaucratic inefficiencies, is typically hesitant to directly subpoena a company in another jurisdiction and domestic court orders have no force beyond their jurisdiction. Requests routed through official channels may take days or weeks to process, affording a substantial head start to criminal proceeds that move at internet speed and may well have moved on through additional hops to or through one or more additional jurisdictions. The result is that, unless a case involves high value or other priorities, investigation by law enforcement is not pursued at all once it is evident proceeds of crime have crossed a border.

The futility of official efforts chasing proceeds of crime across borders is exponentially amplified with anonymous cryptocurrency. Looking again at the (momentarily) high profile Wannacry ransomware case, the few people still monitoring the Bitcoin wallet addresses ten weeks after the global attacks observed impotently as the money was sent to shapeshift.io and changelly.com and exchanged for Monero. Both Shapeshift and Changelly, in a sense colluding with law enforcement officials who needed to pretend their efforts might be effective, hastened to plant self-serving stories regarding cooperation, describing their closing-the-barn-door-after-the-horse-is-gone measures such as blocking the specific Bitcoin wallet addresses used for the identified exchange transactions.

This episode also highlights two additional factors that illustrate the fatuousness of claims that regulating exchange services can mitigate criminal activity conducted with anonymous cryptocurrencies. The proliferation of cryptocurrencies and hardening of their anonymity means that there are no effective methods to:

- Expel bad actors, or to prevent their immediate matriculation,
- Prevent unauthorized provision of exchange services.

Blocking, freezing and expelling

Blocking, blacklisting or reporting known bad wallet addresses, as in the Shapeshift and Changelly whitewash efforts, is pointless since with public blockchains it is a simple matter to create innumerable single use addresses. Proceeds of crime in a wallet address blacklisted by some or all exchange services can readily be moved through mixing utilities and presented via a fresh new address.

Freezing a wallet address, as has been done in the aftermath of multiple Ethereum disasters, requires inducing a hard fork to isolate the value. An analogy would be recalling and re-issuing all US dollars if it was known that bad guys had gained a criminal windfall, perhaps knocking off a Federal Reserve bank as recently glamorized in cinema.

If by some means the entire constellation of bad wallet addresses controlled by a specified criminal entity could be identified (an impossibility with a competent disciplined criminal), and globally blacklisted (another impossibility without a hard fork), it would still be impossible to expel the person from further participation. Nothing but apprehension/arrest could inhibit—'prevent' would be too strong a claim—their immediate rematriculation.

Preventing unauthorized provision of exchange services

The proliferation of automated crypto-to-crypto exchange services illustrates the futility of regulating known exchanges – the ones frequented by relatively non-criminal users[31].

The relatively low barrier to entry for offering automated crypto-to-crypto exchange services facilitates a pop-up model, whereby a service—announced slightly in advance or in real-time to a list of email and/or chat addresses—can appear briefly at a dark web IP-only, non-DNS site.

This business model also affords the additional benefit of facilitating criminal-on-criminal exploits, such as an illicit exchange ripping off its customers, since criminals are particularly unlikely to approach law enforcement for assistance.

Who needs government-issued money anyway?

Currently, the overwhelming proportion of cryptocurrency transactions entail speculative activity by people hoping to get rich quick. But at some point it is possible that some of the thousands of brands of tokens being cranked out start to attract usage as money for the purchase and sale of legitimate goods and services. Widespread and accelerating efforts[32] to enable inline exchange ("cross-blockchain") of cryptocurrencies in the course of a payment transaction will tend to enable people to hold value even in

[31] The term 'relatively non-criminal users' takes cognizance of the near-universal desire of cryptocurrency speculators to avoid reporting their windfalls to tax authorities.

[32] Inline exchange refers to cross-ledger payment channels, natively supported at the protocol level, making use of innumerable ad-hoc cross-token "connector" or "channel" nodes, eliminating the need to seek out a standalone exchange service.

the form of obscure unpopular tokens provided some connector(s) undertake to offer a cross between such orphans and more liquid alternatives. Should cryptocurrencies survive and this scenario play out, this will lead to less and less need for any cryptocurrency user to ever visit a regulated exchange service to swap their tokens for government-issued brands of money.

Alternatively, emergence of services such as the Interledger Ledger Protocol would enable relatively efficient performance of the final "integration" phase of money laundering as cryptocurrency balances in supported ledgers, that have already been laundered through the layering/obfuscation mechanisms touched on above, re-enter participant banks.

Similarly, the "placement" phase of legacy money laundering schemes will effectively be deprecated as the convenience of generating criminal revenues as sums of cryptocurrency promises to render the untidiness of street crime conducted with paper cash obsolete.

KYC as theater

As with measures purported to battle money laundering and terrorist finance mandated for existing regulated mainstream financial institutions, requirements for so-called "know-your-customer" measures impose enormous aggregate inconvenience on the vast majority of relatively law-abiding people while posing no significant impediment to disciplined bad actors. For them, defeating such measures is just part of the cost of doing business. But unlike all precedent systems, concerted, sophisticated, worldwide, even *competitive* effort is exerted to harden the anonymity possibilities afforded by cryptocurrencies.

Catalog of criminality facilitated by anonymity

Anonymous cryptocurrencies facilitate efficiencies with respect to established categories of criminal activity and enable innovative new exploits that heretofore would have been impracticable.

Tax evasion

There is nothing new of course about tax evasion. But the proliferation of speculative activity involving Bitcoin, other cryptocurrencies and the explosion of tokens offered in "ICOs" (Initial Coin Offerings) is leading tens of thousands of people who previously lacked a reason to seek out evasion mechanisms to avidly explore the possibilities. Many, especially millennials, are indignantly discovering for the first time the multitude of existing measures, especially those pertaining to US taxpayers, for authorities to detect unreported income. Such measures include the filing of 1099's by exchange services, FBAR requirements for US taxpayers to list overseas financial accounts and FATCA rules requiring foreign financial institution cooperation.

This rude awakening and resultant indignation is in turn driving new developments in cryptocurrency ecosystems. The efficacy of some of these new directions may be illusory. For example, part of what is driving the proliferation of crypto-to-crypto exchange activity is the belief that such transactions may be tax-deferred "like-kind exchanges" as specified in IRC Section 1031. It is not clear if this was ever a valid tactic but, whether or not it shielded such transactions in the past from capital gains tax, it will not in the future.

The explosion of new measures to harden anonymity and to provide for offshore crypto-to-crypto exchanges discussed above is due in part to tax-related concerns. But entities that claim to implement adequate identity verification, and are therefore likely to eventually be compelled to abide with FATCA requirements, will lose market share to more stealthy offerings.

Huobi is an interesting case study. Huobi.com, originally headquartered in Beijing, is listed as one of the largest cryptocurrency exchanges in the world, despite its apparent proclivity for gambling with customer money for its own gain. Due to FATCA concerns, Huobi has moved to prevent US taxpayers from using its service, though evading those restrictions is reportedly not unduly difficult. So to accommodate US-based speculators, Huobi established huobi.pro – a service enabling "global traders" to engage in strictly crypto-to-crypto exchanges. This new subsidiary is incorporated in the Seychelles, an archipelago off the coast of southeastern Africa in the Indian Ocean.

The FATF (Financial Action Task Force), which heads up international cooperative efforts to combat money laundering, illegal arms trading and terrorist finance, helpfully notes that the Seychelles are a member of—implying their AML efforts are monitored by—the Eastern and Southern Africa Anti-Money Laundering Group (ESAAMLG). The ESAAMLG reportedly last evaluated the effectiveness of the Seychelles AML efforts in 2008, but repeated attempts to pull up the report online were unsuccessful, as were attempts to access the ESAAMLG homepage.

Money Laundering – placement, layering and integration

Money Laundering may be defined as "the process of taking the proceeds of criminal activity and making them appear legal". More specifically, "Money laundering refers to a financial transaction scheme that aims to conceal the identity, source, and destination of illicitly-obtained money". In most taxonomies, the money laundering process can be broken down into three stages, placement, layering and integration. The first stage, placement, is construed as taking the proceeds of the illegal activity that garners the money and placing it "in the launderer's hands".

This conventional definition fails to encompass the reality of money laundering as practiced with cryptocurrencies. Rather than laundering money (presumed to be conventional money) generated by external illegal activity, the proceeds of crime are directly obtained in the form of cryptocurrency.

Stealing cryptocurrency

As noted, the notion of money laundering contemplates moving the proceeds of crime but is not traditionally defined as to encompass the crimes that gave rise to the proceeds. But stealing cryptocurrency, to date, appears to be more lucrative than bothering with cryptocurrency as a means of laundering criminal proceeds for illicit profits generated in the form of good old fashioned cash and bank deposits.

As noted above, anonymity enables financial crime with the efficiency of an industrial process. One of the most lucrative exploits is simply to steal cryptocurrency. As far back as 2016, such thefts were becoming commonplace.

Stealing cryptocurrency has become even more of an issue with ICOs, as reports indicate $1.5 million per month is stolen in relation to ICOs alone.

Ransom and extortion

Ransom demands payable in anonymous cryptocurrencies are not limited to ransomware. Ransom demands for anonymous digital cash in relation to real kidnappings will almost certainly displace the traditional preference for paper cash as meatspace criminals up their crypto game.

Sale and purchase of criminal goods

Another major category of crime where proceeds are generated directly in the form of cryptocurrency is the sale and purchase of criminal goods such as illegal drugs, child pornography[33] and stolen credit card and identity data.

Sale of criminal services -murder for hire

The Cypherpunk dream of Assassination Politics first articulated by Jim Bell—frequently referenced on the Cypherpunks mailing list as "AP"[34]—can finally be realized with bounties paid via systems which afford strong anonymity such as Monero and ZCash.

Terrorist finance

Cryptocurrencies are gaining increased usage as a preferred means of terrorist finance and for evading international sanctions such as those promulgated by the US Treasury Office of Foreign Assets Control (OFAC).

Unpermissioned Central Bank Digital Currency (CBDC)

In 2016, Sarah Meiklejohn—cited above as one of the academics claiming that the transparency of cryptocurrencies enables law enforcement to follow the money—also proposed that government central banks issue their own "pseudonymous" cryptocurrencies. The paper included a proposal for an "RSCoin"[35] which, "similar to…existing cryptocurrencies" would enable users to "create new pseudonyms" and could perhaps further "ensure privacy for transactions" by adapting "existing cryptographic techniques such as those employed by Zerocoin". [Zerocoin was the precursor to Zcash, discussed above].

This helped stimulate a flurry of central bank projects worldwide to evaluate the potential for central bank cryptocurrencies, sometimes also referred to as Central Bank Digital Currencies (CBDCs).

As of 2018, most of the enthusiasm for central bank-issued tokens has subsided as authorities have realized that the anonymity afforded by an "fully permissionless" system would lead to "prohibitive societal costs" while the administrative overhead/costs of effective identity verification and other elements of a permissioned system would strain central bank budgets.

Similarly, the risk that the availability of central bank money in a system enabling efficient P2P payments could induce or worsen a run on bank deposits appears to be a fatal deterrent. Central bank schemes for issuing their own digital currency—for anything more than transactions involving trivial sums—will probably remain dormant until renewed financial crisis leads to a scramble for new monetary stimulus in the form of helicopter money.

An additional artifact affecting CBDCs or other possible central bank enhancements to payment mechanisms is their invariably domestic focus. No central bank wants other central banks encroaching on

[33] As with murder for hire, discussed below, some Cypherpunks have argued for decriminalizing the purchase and possession of child pornography.
[34] A search of the raw Cyherpunks mailing list archives for just the year 1997 using the term AP (preceded and followed by a space so as to reduce false positives such as where ap is part of a word) yields 641 hits.
[35] The RSCoin proposal credited the Bank of England's 2015 One Bank Research Agenda which, in the then prevailing enthusiasm for Bitcoin, puzzled whether "a permissioned or permissionless system [would] be more appropriate".

its own domestic monetary domain and similarly hesitates to overtly offer a payment medium or mechanism designed to optimize the efficiency of international payments.

Private – Permissioned decentralized networks

Private blockchain schemes, and related permissioned/closed networks such as R3 Corda, are constituted in accordance with the logic of the "Willy Sutton rule", focusing on banks because "that's where the money is". As such, they tend to perpetuate institutional arrangements that interpose banks into payment mechanisms as obligatory intermediaries. In these systems, broad direct public access to P2P payments and other transactions takes a back seat or does not enter into the picture at all.

Private blockchains do not need to worry so much about traitorous Byzantine Generals; all participants are regarded as sufficiently reputable as to permit participation. But just as scenes where the Three Stooges try to all go through a doorway at the same time, arranging decentralized non-hierarchical coordination can be painfully inefficient. And with the insistence on avoiding use of a centrally administered financial utility, no matter how well conceived and implemented, even reputable institutions must approach each other as "untrusted" in a decentralized system that requires consensus.

It is precisely this need for consensus that is the Achille's Heel of private DLT schemes. Per the Fischer, Lynch, Paterson (FLP) Theorem[36], all consensus models are constrained by conflicts between Safety (Consistency), Liveness (Availability) and Fault Tolerance (tolerance of nodes dropping out whether due to failure or network partition). Systems that scale well cannot achieve consistency/finality without prolonged latency. Systems that can achieve near-immediate finality become unusable due to messaging and computational overhead when the number of nodes exceeds about twenty.

The related fundamental problem with this prevailing concept of a decentralized system is exactly the same as with the Cypherpunk schemes of the 90's, notably Digicash. Decentralized schemes require participants to operate a complex node. But rather than bear the overhead cost and hassle of deploying and maintaining a node, most people and entities are better served by a client server arrangement, or, in the case of financial intermediaries, by accessing a centrally administered financial utility via standardized interfaces.

The most glaring flaws of legacy systems do not stem from technology but are the result of flawed institutional arrangements. The obligatory interposition of banks as financial intermediaries for all remote payments is one of the core problems. The lack of an alternative global money issued in accordance with simple rules, as opposed to discretionary monetary policies, is another. While technology matters a great deal, the proposed solution outlined below holds institutional arrangements and monetary principles to be of greater importance than technical considerations.

Monetary premises and principles

As will be discussed in the appendix, the premises underlying unbacked cryptocurrencies (and other tokens sometimes regarded as embodiments of monetary value) entail monetary fallacies, dating back to failed Cypherpunk initiatives of the 1990's.

Cypherpunk logic holds that a token may be ascribed monetary value on the basis of

- Arbitrary scarcity,

[36] Note the similarity of FLP to the CAP Theorem discussed above.

- Safeguards to preclude counterfeiting, such as via a double-spending exploit.

Scarcity vs Monetary Liability

Scarcity-as-value stands in contrast with an alternative logic—Monetary Liability—that informs most forms of issued money, such as the moneys issued by government Monetary Authorities. The idea of monetary liability is to regard money that has been issued and outstanding as a financial liability, represented on a balance sheet[37] and offset by a portfolio of assets[38] held in readiness enabling the issuer to buy back any or all of the money in circulation. This prudential practice is observed not just by monetary authorities, whose monetary liabilities serve as Base Money, but also applies to banks and other depository institutions, the monetary liabilities of which constitute Deposits, the principal component of the Broad Money supply. Why does this matter?

The problem with unbacked cryptocurrencies is that when demand for a specific brand of money[39] declines, its exchange rate relative to all other brands of money (and everything else in the world) also declines. There is no issuer with the wherewithal (that would be provided by a trove of liquid assets) to buy in the excess quantities of the now-unwanted money. Moreover, in the event of a decline, were a cascading panic to erupt, nothing prevents the market value of the money dropping to zero[40].

This is in marked contrast with the adjustment mechanism built into moneys that are accounted as the liabilities of a solvent issuer. A decline in demand for the specific money does not necessarily affect its exchange rate to any significant degree. Instead, that money supply—the entire monetary balance sheet for that issuer—may shrink. Depending on the policies governing monetary collateral and other aspects of operational logic, the adjustment can occur automatically with zero latency—as with the proposal detailed below—or it may entail long delays stemming from the time lags and reflexivity built into discretionary monetary regimes.

The money supply for each unbacked cryptocurrency is set to an arbitrary fixed amount. It may all be pre-issued and doled out to the public in dribs and drabs, as described above with Ripple. Alternatively, a large portion may be pre-mined as a means of rewarding the developers and other insiders, with the remainder scheduled to be created by a mining process as with Ether and Bitcoin. What is lacking, however, is any orderly mechanism for shrinking the money supply.

Exchange market considerations

With the alternative proposal detailed below, the money supply is created and distributed into circulation or, if market conditions warrant, recalled and extinguished via a currency exchange mechanism. This is mediated by independent financial institutions providing currency exchange services—between system-created and conventional money—on a competitive basis. This industry, which first came into being circa

[37] For example, the Fed's weekly H.4.1 Statistical Release or, a historical private sector example, e-gold's Examiner Page.

[38] Examples of such asset portfolio reports, again, for the Fed and for e-gold.

[39] Demand for a specific brand of money should not be confused with the more complex theoretical concept of demand for money in general (as opposed to demand for other assets such as investments), normally discussed in terms of "liquidity preference".

[40] A temporary floor might form if there is a sufficient balance of short positions. To realize profits from a short position it is necessary to cover it by buying assets to replace what was borrowed. This too becomes trickier however if no market exists to buy back the assets as with a delisted stock or a collapse of cryptocurrency exchanges.

2000-2001 (created and fostered by the e-gold founders) has exploded into a multi-billion dollar sector with the cryptocurrency bubble. There is, however, a huge difference in soundness and sustainability of exchange service provision as conducted with unbacked cryptocurrencies as opposed to what is possible with well-constituted, redeemable moneys.

From time to time, a retail provider of exchange services encounters an imbalance between customers seeking to buy and customers wishing to sell. If the exchange provider itself is the counterparty to customers, in other words acting as a market maker, such a shift can lead to an imbalance in the dealer's own inventories of traded currencies. If corrective measures[41] are insufficient to bring trading balances back into trim, the erstwhile market maker may be forced to resort to an external market as a price taker (paying a spread instead of charging a spread). The problem with unbacked cryptocurrencies is that in a serious downdraft of demand, all providers are affected and there is no external market – no entity is able to provide a market of last resort. To do this would require someone to have the right to return the now-unwanted money to the issuer for redemption[42].

If and when a cascading panic bursts the cryptocurrency bubble, market-making dealers—even if at some point they refuse to accept additional sell orders—will find themselves drowning in unwanted excess cryptocurrency balances (which they cannot unload in a no-bid situation) and facing depletion of their real money balances. They will then default, unable to pay out real money to customers who had tried to sell them cryptocurrency. If the carnage is bad enough, a number of them will likely lapse into a terminal phase of outright larceny, absconding with all remaining real money balances while exclaiming "we wuz hacked!".

Not all cryptocurrency exchange facilities entail the provider acting as market maker. Many involve platforms for speculators to trade with each other. But default would still be likely to occur in many such cases. This would happen because all such trading platforms require prefunding in order to prevent customer repudiation of funding obligations. In essence, the customers trade deposits issued by the platform administrator and purportedly backed by a 100% reserve of…something. Sensing that the jig was up, the principals of many such exchange platforms would abscond with whatever real money prefunding remained. This would be especially likely with stealthy exchange services where the location or identity of the principals is uncertain.

With the alternative described below, as with its e-gold predecessor, providers of retail exchange services, if confronted with an otherwise intractable imbalance between customer buy and sell orders may always resort to an external market enabling them to bring their own trading balances back into trim. Designated "Primary Dealers" are vested with contractual responsibility to make a market in all market conditions. The ability of Primary Dealers to meet such obligations rests on their right to require the Issuer to perform redemption on demand, which in turn rests on a 100% reserve.

[41] In relatively normal market conditions, the provider may compensate by adjusting proferred bid and ask rates, or imposing limits on order size.

[42] The situation with government central banks that issue irredeemable currencies in accordance with discretionary monetary policies would be similar but more complicated.

Sloppy rhetoric

A popular trope favored by cryptocurrency apologists[43] for characterizing the money issued by a central bank such as the Federal Reserve is to describe it as "created out of thin air". But this is unhelpful and misleading. The Base Money supply issued by the Federal Reserve or other central banks is created out of government bonds, other remunerative assets such as securities and direct loans, and reserves, which may be gold, SDR's or foreign currency balances held at foreign central banks. The US Treasury bonds and Mortgage-backed Securities (MBS) that comprise the bulk of assets held by the Federal Reserve command a higher market price than air of any density. Mined cryptocurrencies, in contrast, cannot even be claimed to achieve the thin air standard as their creation process requires the destruction and degradation, respectively, of economically valuable resources and the environment.

Market cap

A more serious misuse of terminology is the practice of referring to the "market cap" (market capitalization) of various unbacked cryptocurrencies or of all cryptocurrencies in aggregate.

The market cap of a particular instrument is calculated by multiplying the market price of a single unit by the total number of units outstanding. This metric was originally devised as an aid for evaluating the priciness of a company's stock. In this setting, market cap is a meaningful number, commonly considered relative to book value – total assets minus total liabilities. Market cap, and the price-to-book ratio, may be relatively high or low, reflecting expectations concerning the prospects of a company for future growth and profitability. The comparison to book value anchors such expectations to value that might be realized in an adverse outcome such as liquidation or a forced restructuring.

A bear raid on a company—dumping shares en masse in an effort to drive down their market price—would encounter limitations as value investors perceived the opportunity to acquire the net corporate assets on the cheap. Similarly, a crisis involving the stock market—even a complete suspension of trading lasting for months, as occurred with the outbreak of WW1 —would not make shares in a viable company worthless./

No such floor exists with unbacked cryptocurrency. As this is being written, the "market cap" of Bitcoin alone is quoted as exceeding 180 billion USD. It is exceedingly unlikely, however, that so much as 1 billion USD, < 0.6% of this amount, could be sold quickly (for real money) without driving the price to zero (quite possibly taking out all other unbacked cryptocurrencies—more than 1,900 of them at latest count—with it).

Play money vs. Real Money

In 2001, Laurence Meyer, at that time a Fed Governor, gave a lecture outlining his views regarding "The Future of Money and Monetary Policy". He offered the following interesting assertion:

[43] Weirdly, this rhetoric is also a favorite of gold bugs and other advocates for rules-based government monetary regimes.

> *"...central banks, at least in developed economies, issue currency and provide clearing services, at least in part, because their services offer features, such as freedom from default risk and finality of settlement, that private providers cannot match."*

The relative incapacity of private providers is refuted below. But his pronouncement of the imperatives of assuring: a) freedom from default risk, and, b) finality of settlement (of base money transfers) comprises a basis for formulating a definition of real money[44].

Real money can be thought of as a brand of money, the base money of which is suitable to be held by financial institutions as a reserve asset—held against like-denominated broad money obligations such as bank deposits—and used as a medium of settlement. The imperatives Meyer identified comprise the essential criteria for such suitability.

But how can freedom from (issuer) default risk be assured? This is achieved by the issuer backing all monetary liabilities at all times by liquid assets held in readiness and in sufficient quantity as to assure the issuer's ability to buy back all that has been spent into circulation.

The sophism embodied in unbacked cryptocurrency is that default cannot occur because there is no liability, no obligation to perform. The problem with this however, as already noted, is the ever-present risk of a cascading panic rendering the money worthless.

Back in the 1990's, the holy grail of Cypherpunks was anonymous digital cash implemented using "blind signatures" as worked out and patented by Dr. David Chaum. Digicash bv, sporting the logo "numbers that are money" was formed to implement digital cash (affording payer, but not recipient, anonymity) on a commercial basis. The company and its advocates cited previous testing with an unbacked token called "Cyberbucks". Cyberbucks, as will be related in the appendix, was the direct precursor to the unbacked cryptocurrencies that made their debut fourteen years later. Tellingly, company press releases described Cyberbucks as "a 'monopoly(TM) money'". A more general term for unbacked cryptocurrencies (which also avoids possible genericization of Hasbro's trademark) would be "play money".

Ethereum and Ripple each involve a native currency

Banks and institutional investors have taken an interest in the blockchain/DLT phenomenon and one would be hard pressed to find a major bank that has not announced participation in some sort of blockchain/DLT or blockchain/DLT-inspired "decentralized" scheme. A theme evident in many such discussions is that DLT technology may have significant potential even if the unbacked cryptocurrencies that gave rise to them are a passing fad.

For several years, Ethereum and Ripple have sought to engage with institutions, particularly banks, offering platform solutions to support permissioned network DLT applications.

[44] Meyer's assertion also neatly abstracts the defining functions of a monetary authority:

- issuance of base money, and
- provision of a platform for settling base money transfers.

There are, however, a few flies in this ointment. Both the Ethereum Foundation and Ripple depend on selling ETH or XRP respectively to fund operations[45]. But if participation in a private network DLT scheme meant a bank was obliged to hold balances of these unbacked cryptocurrencies, as with Ripple's xRapid, it would accord poorly with their prudential obligations as fiduciaries. Ripple, accordingly places emphasis on XCurrent, its ILP-based offering which does not obligatorily involve XRP or even a shared ledger.

Similarly, solutions marketed by other private-permissioned DLT vendors such as R3 or Hyperledger do not entail holding unbacked cryptocurrency balances.

This movement away from obligatory usage of ETH or XRP poses a problem for the Ethereum Foundation and Ripple (the company), both of which were/are lavishly funded by sale of unbacked tokens, huge proportions of which were kept for their own financing and remuneration, to the public. Were the crypotocurrency bubble to collapse, the ability of both companies to sustain their ginormous burn rate for R&D, both internal and through funding of grants, would evaporate. Ripple currently manages to convince banks to pay undisclosed but presumably hefty sums for software but in the setting of an across the board cryptocurrency panic and collapse, it is easy to imagine banks losing interest. The fact that cryptocurrency has been just one of multiple asset class bubbles fueled by an unprecedented decade-long money glut also raises the likelihood that when the day of reckoning comes, it will be in conjunction with a financial crisis that once again brings banks to death's door and causes them to gut R&D budgets.

A suitable medium needed

Outside of the community of those directly engaged in or otherwise cashing in on the rampant speculation that (so far) elevates the paper value of unbacked cryptocurrencies and ICO tokens, many recognize the need for a more sound form of digital currency. Numerous schemes have therefore been advanced for issuance of asset-backed tokens, most commonly piggy-backed on the Ethereum or Ripple blockchains or platforms. The Ethereum model for this usually entails implementation as an ERC20 token, a smart contract format enabling specification of monetary liability. The original Ripple model, now deprecated, contemplated "Gateways"—financial institutions, either banks or exchange services—creating similar monetary tokens.

Financially viable issuance of a sound asset-backed currency on Ethereum or Ripple impracticable

There is a problem, however, with issuing a sound asset-backed digital currency via a blockchain scheme, which will be further examined below in the discussion of institutional arrangements. With public blockchains, it boils down to revenue models. How does a financial institution, or anyone else, generate revenue from transactions in which they play only a third party role[46]? In particular, how might an issuer of monetary liabilities profit from such activity?

There are only a limited number of categories whereby an issuer might generate operating revenues (exclusive of selling unbacked tokens):

- Treasury management/operations;
- Fees;

[45] A class-action suit filed May 2018 alleging XRP to constitute an unregistered security describes the Ripple business model as a "never ending ICO", noting ongoing sales of XRP ("over $342.8 million …in the last year alone") to be Ripple's "primary source of income".

[46] By a "third party role" I mean neither as the payer nor the recipient of money or other consideration tendered in exchange for goods or services.

- Exchange spreads.

Treasury management/operations

Banks, including government central banks and currency boards, make money from holding remunerative assets which may generate interest income or capital gains in excess of the funding cost of their deposit liabilities. But we are talking about an issuer of money, that is, an entity whose liabilities are monetary liabilities, that is, liabilities suitable to be held and used by others as money. The problem with this—discussed below in the context of eliminating the risk of issuer default—is that remunerative assets, in all normal circumstances[47], carry a degree of risk. If the market value of the assets held against outstanding monetary liabilities drops below the notional value of those monetary liabilities, the issuer is insolvent – unable to buy back all the money it has issued into circulation. So an issuer, unless backstopped by a sovereign entity such as a national government, has no business generating revenue via the asset portfolio backing its monetary liabilities.

Fees

Excluding fees assessed in relation to (what amounts to) currency exchange, addressed below, the only other categories of fees an issuer might tap as revenue would be transaction fees, such as for processing a transfer from payer to recipient, or some sort of account maintenance fee.

The difficulty that comes into play when piggy-backing on Ethereum or original Ripple is that some entity other than the issuer—a miner or validator—is awarded transfer fees.

Account maintenance fees, per se, are also not amenable to capture because wallets differ from bank accounts – where the issuer is also the entity that processes and keeps track of the transactions and balances of money holders. The only variant of an account maintenance fee that might benefit an issuer in a DLT arrangement would be a smart contract that defines a wasting asset, where the wallet balance is designed to decrement with time, thereby reducing the outstanding liabilities of the issuer. Historically however, people have never chosen money (in bearer form) that evaporates with the passage of time—sometimes referred to as "demurrage" —if they have ready access to money of constant notional amount.

Revenues from provision of currency exchange services

A currency exchange transaction can be abstracted to encompass an exchange whereby an entity accepts one brand or type of money in exchange for another. While such a transaction may be conducted in a fashion where the amounts of money conveyed in both transfers are calculated using the same exchange rate, in most commercial settings a provider buys money at a lower price (exchange rate) than it sells it. This bid price vs. asking price spread could be disclosed as a fee. Alternatively, or in addition, a transaction fee may be levied by the provider on some other basis such as a flat fee.

In a DLT mechanism such as the Interledger Protocol now being championed by Ripple, certain participants (Connectors) may generate revenue on an ongoing basis by providing such exchange functionality. But this is different from an issuer who is only involved in transactions that entail enlarging or shrinking the quantity of their liabilities in circulation. They may only get a single crack at exchange-type revenues, by charging a markup for emission of new money or, occasionally buying back money at a

[47] An exception—which I hold to be bizarrely abnormal, never countenanced by serious economists prior to the past decade as anything other than a clever thought experiment—is interest paid by a central bank on reserve balances.

discount. This, however is an insufficient source of operating revenues, especially if they issue money for which the reserve assets backing it entail carrying costs.

This combination of factors is why, to date, despite frequent announcements of entities intending to issue a sound asset-backed money via a public blockchain, none are likely to do so in a fashion capable of generating operating revenues sufficient to achieve positive cash flow.

Money issued for use in a private DLT scheme

With a private-permissioned DLT arrangement, whether involving a blockchain per se or some other "decentralized" scheme, it is even sillier for a participant—including the government monetary authority itself—to issue asset-backed digital cryptotokens residing on a distributed ledger.

Most private DLT schemes involve a consortium of banks undertaking to engage in clearing and settlement of intermediated payments. As a rule, none of the participant banks would relish the prospect of holding settlement balances consisting of deposit liabilities of one of its peers/competitors. Similarly, a participant with ambitions to be top dog, would face the same constraints preventing profitable performance of the issuer role as outlined above.

Utility Settlement Coin

For this reason, a consortium of banks was launched in 2015 (to considerable fanfare) that would use various conventionally-denominated flavors of a "Utility Settlement Coin (USC)", each backed by a 100% reserve of the corresponding central bank money. As currently conceived however, the idea serves only to illustrate how even the most sophisticated financial institutions in the world seem to lack basic judgement when under the influence of blockchain enthusiasm.

Just as the notion of Central Bank Digital Currencies (for anything other than very low value transactions) discussed above is likely to be quietly shelved, the prospect of any USC being made directly available to the public as an option for holding balances or as a medium for P2P payments is nil. Banks as we know them may face extinction when efficient remote payments involving a sound transaction medium can be effected with no need for an obligatory bank intermediary; they are unlikely to volunteer for creative destruction.

If USC is only held by banks to be used as a medium of settlement it raises the question: why bother? Currently, processes exist for swapping the various forms of central bank money required by commercial banks – paper cash for ATM and over-the-counter withdrawals, and electronic reserve/settlement balances held as deposits at the central bank. Why add a third form, when account balances at the central bank can already support efficient clearing and settlement, with centralized administration enabling more efficient performance of multilateral netting than any decentralized kluge?

As will be detailed below, if there is to be a new medium created to facilitate clearing/settlement of intermediated payments, it would make more sense if the medium did double duty, also serving as a reserve for immediate funding of customer withdrawals. But the very idea of withdrawals in the form of the new medium requires that the general public would be able to hold balances external to the banking system. The ability on the part of the general public (subject to appropriate permissioning) to hold such balances logically pairs with a system enabling the medium to also circulate external to the banking system via P2P transfers.

An additional limitation to USC, as with every other monetary and payments initiative being pursued by central or commercial banks, is that it fails to address the need for an efficient global payment system based on a suitable international reserve and settlement medium.

Project Jasper

A recent illustration of the contrived contortions undertaken in order to use cryptotokens in a decentralized settlement scheme was the Jasper Project undertaken by the Bank of Canada. Jasper was a two-phase exploration of the use of DLT for clearing and settlement of interbank payments via RTGS as well as with protocols that entail netting. A prerequisite step was the pledge of conventional balances at the central bank in exchange for cryptographically signed Digital Depository Receipts (DDRs). This third form of central bank money was then redeemed at the end of each day[48] (back to conventional balances) in a process requiring the usual multi-step digital cash rigamarole to avoid double spending.

The predictable conclusion of this multi-million dollar, multi-year exercise was "It will be challenging for any DLT-based system to process payments more efficiently than the LVTS"[49] but held the door open for using DLT "If a DLT-based settlement system is able to reduce back-office reconciliation efforts". This is like saying "if everyone uses the same product/service, and also stays in synch with any upgrades that aren't backwards compatible, the inefficiencies that arise from different folks using different products/services (such as those offered by competitors of the anointed solution provider) would be eliminated". In other words, if a group of banks all elected to use any particular reasonably up-to-date and commonly agreed-upon software system both to interface with each other and for their back office operations it stands to reason that reconciliation would be a smoother process or that no reconciliation per se would be needed in the first place[50]. But unlike software vendors offering systems engineered to be free of the cumbersome restraints of a decentralized, consensus-requiring protocol, R3 and its DLT competitors manage to garner millions in what amount to grants from taxpayer funded institutions.

International money and payments

Legacy monetary and payment systems by and large are organized in accordance with political boundaries. This compartmentalization into multiple domestically focused domains leads to inefficiencies and other negative side effects.

Global currency

In the aftermath of the GGC, multiple entities, notably China and the Vatican, advocated the creation of a global currency. These proposals could be grouped in two categories.

Global Central Bank

For many, the term "global currency" connotes a global money that would supersede and replace existing national brands of moneys, more or less as was done with the formation of the Eurosystem but on a global basis. Just as the euro enabled efficiencies for cross-border transactions within the eurozone, the idea is

[48] DDRs could be redeemed intraday as well. The report noted that in an actual deployment DDRs could be held for longer or shorter intervals as desired by the participant banks. The daily pledge-generate and then redeem cycle was arbitrarily used "to avoid the need to apportion interest against DDRs rather than the underlying cash deposits".

[49] LVTS refers to Canada's existing Large Value Transfer System.

[50] It must also be noted that as institutions worldwide implement ISO 20022 messaging standards, the expanded format of ISO 20022 provides for transmission of reconciliation data, more or less with "smart contract"-like (XML) name-value pairs.

that a global currency would facilitate trade and movements of capital. For example, in 2011 the Vatican issued an encyclical that explicitly called for formation of a global central bank, which it described as a step toward a “global polity”.

SDRs

Another variation that is discussed from time to time is to expand the role of Special Drawing Rights (SDRs) – an international reserve asset created by the IMF in 1969 to supplement its member countries’ official reserves. Inspired by the “Bancor” and International Currency Union concepts articulated by Keynes in the 1940s, SDRs were intended to serve not only as an international reserve asset but as a medium of settlement for international capital flows. Keynes’s centralized global control of capital movements never came to fruition, a state of affairs many attributed to US influence around the time of the Bretton Woods Agreement.

Relatively few people have even heard of SDRs. Only a much smaller subset of this group would be able to offer a coherent explanation of what they are, and even fewer have shown any inclination to price goods, services or debt instruments using SDRs as the unit of account. This partly stems from the fact that SDRs do not circulate as money in the sense of a medium one might use to make or receive payments. Per the IMF’s SDR Factsheet: "The SDR is neither a currency nor a claim on the IMF. Rather, it is a potential claim on the freely usable currencies of IMF members. SDRs can be exchanged for these currencies.”. In practice, a country with an imbalance of payments can sell some or all of their allotment of SDRs to another country in exchange for hard currency that is actually usable for buying imports.

China, around the time of the global financial crisis, was the principal advocate of an increased role for SDRs and for inclusion of the Chinese renminbi (RMB) as one of the currency basket components commensurate with the increasing global prominence of the Chinese economy. Subsequently, in 2009, an increased allotment of SDRs was imposed on IMF member countries, expanding their supply nearly ten-fold, from 21.6 billion to 204.2 billion. In 2016, the indexed basket of currencies used to define and value the SDR was altered, with the addition of RMB.

SDRs are in fact more ubiquitous than gold as a reserve asset on government central bank balance sheets although their aggregate value of about 290 billion USD is dwarfed by the 1.4 trillion USD market value of central bank gold holdings.

But while SDRs are used occasionally for transactions between government central banks they have never attracted commercial usage. No commercial bank deposits are denominated in SDRs and they are not used as a pricing, invoicing or accounting unit by commercial firms.

Prospects of government-issued global currency

The US dollar is sometimes referred to as a global currency because USD is the most commonly used invoicing unit for international trade and shows up the most frequently in currency exchange pairs. It is also the third most widely held reserve asset on the books of government central banks (after SDRs and gold).

But extensive global usage of USD (or EUR, GBP, JPY or RMB) does not mean that it is a global currency. Like any other government-issued brand of money, USD is administered pursuant to domestic US policy objectives. Moreover, dominance of any particular national currency as an international reserve asset imposes burdens both on the country that issues it and on trading partners. Global usage of USD bolsters

US national prestige but inevitably leads to chronic current account deficits to offset capital account surpluses, a phenomenon known as the Triffin Effect. Trading partners on the other hand are affected by what D'Estaing termed the "exorbitant privilege" of the US, effectively subsidizing a higher level of debt-fueled consumption in the US.

There is another problem that makes the idea of global usage of any national currency a bad idea. Governments have a poor record when it comes to abiding by obligations, whether monetary or treaty-defined. Part of the reason SDRs were created was that the US abrogated its obligations under the Bretton Woods Agreement for USD to serve as a proxy for gold, suitable for other countries to hold as a substitute for gold. Governments, at least when it comes to their monetary obligations cannot be bound by contract.

Private sector issuance

International payments

There are two general categories of remote payments[51], whether domestic or international: payments that convey directly from payer to payee (P2P)[52], and payments that require the obligatory involvement of two or more financial intermediaries (typically banks or something like a bank), at least one of which incurs liabilities to another which must be satisfied, typically within one business day.

P2P

For payment to convey via a P2P mechanism requires both payer and recipient to use the same brand of money and directly participate in the same payment system.

Clearing and settlement of intermediated payments

Wang'lian principles/practice wrt interop

P2P

Global reserve and settlement medium

Case – Oilcoin

Speculative investment vs. price stability

Money is not an investment

Targeting price stability is a fool's errand

ILP affords little in the way of utility

Innumerable "currencies"

The fact that an ever mounting avalanche of new coins appear is celebrated, even framed as "asset based sharding", a remedy for scalability limitation.

The ICO phenomenon

[51] "Remote payments" means payments that do not require payer and payee to be in the same location or entail physical transfer of bearer instruments such as paper cash.

[52] A bank may fulfill a payment instruction in which both payer and payee are its own depositors either by a book entry process involving only its own accounting system, or by uploading the payment message to a clearinghouse like any other payment. The former would qualify as a P2P payment. The situation is analogous to postal mail where both sender and addressee both share the same postal code. There could be an internal sorting and routing process at the local post office although more likely the letter would be sent to a central facility for processing.

Achieving freedom from default risk

Finality of settlement

Problems and solutions

There are serious economic problems worldwide arising from longstanding and seemingly intractable flaws embedded in monetary and financial arrangements as well as payment systems. These problems manifest both cause and effect linkages to unsustainable government fiscal practices and give rise to both periodic financial disruptions and an increasing degree of financial inequality.

Efforts to solve or amelioriate these problems, all of them global to some extent, would best be undertaken in accordance with a coherent, comprehensive and defensible logical framework, a private industry example of which would be a well-conceived business plan.

A ubiquitous deficiency undermining blockchain initiatives which has resulted in malinvestment—massive capital expenditure, unconscionable operating costs, speculative excess, wasteful churn and incoherent schemes—is the lack of definition of goals derived from clear articulation of the problems to be addressed and possible benefits to be realized.

A better approach starts with an analysis of problems and deficiencies of antecedent, existing and proposed systems for money, payments and credit as well as potential benefits that may be realized from emergence of new systems. When problems are identified, a systematic analysis of their origin and root causes may facilitate discovery of strategies to definitively address deeply embedded legacy artifacts that, while they may have been the best option in an earlier age, have persisted via path dependent processes to endure as foundational flaws. This aspect is essential because the most urgent subset of the problem domain is the seemingly intractable systemic risk stemming from existing arrangements.

Formulation of a solution

With the problem domain precisely defined, then and only then can strategies be formulated and evaluated. This process of requirements development does not begin with evaluation of appropriate technologies but rather focuses initially on core principles and institutional arrangements. Core principles encapsulate both positive imperatives and constraints, examples of which, respectively, might be broad global access and avoidance of coercive imposition. Institutional arrangements, in turn, are a matter of roles – what sort of entities perform which categories of functions. Coherent institutional arrangements must align responsibilities, costs and risks to revenue streams and other incentives.

DLT advocates have not taken this approach. Global Standard, in contrast, has, attacking the problem domain with a "clean sheet of paper" approach. The idea was to emulate the analysis that might be undertaken by a sophisticated extra-terrestrial cognizant of but unencumbered by path dependent legacies. The resultant thesis is also accompanied by a gap analysis and highly practical transition pathway designed to minimize disruption and create attractive incentives for established institutions to matriculate to the Global Standard system and facilitate its implementation and emergence.

The problem domain

Money

Payment systems

Systemic risk

Unsustainable government fiscal practices

Core principles for a solution

Its fundamentally about the money

Rules rather than discretionary policies

Demarcation between money and credit

Market-based, non-coercive

Self-sustaining revenue model required

Global

Must support both P2P and intermediated payments

An identity framework, either external or internally provided, is essential

Central administrative control

Channelize usage patterns to prevent abuse and facilitate detection of illicit pattern

A system implemented by the author had evolved the most sophisticated detection, investigation, interdiction etc capabilitis of any financial institutuion. But this did not prevent it being subject to unremitting reputation attacks casting the system as anonymous untraceable and inaccessible to law enforcement. Prevention of illicit usage is required.

Make it easy and convenient to use system for a legitimate business or personal purpose but make it inconvenient to do naughty things.

Highly granular compliance

Avoid need to trust

Separation of Roles

Some sort of money is useful even for systems designed for other purposes, such as systems of record

Goals

Better money - Alternative global currencies implemented as Digital Base Moneys

Must achieve freedom from issuer default risk equal to or exceeding government monetary authorities

Better payments

Must achieve immediate settlement, with finality

P2P

Intermediated

Financial inclusion and levelling playing fields

Addressing financialization and concentration of wealth

Harness collective wisdom enabling automatic self-adjustment with high granularity and minimal latency

Direct Agent-Effected Modulation (DAEMON)

Entrainment effect on discretionary monetary and fiscal policies

Efforts to define the appropriate scope of State activities is a fool’s errand. The goal of this system is to assure that whatever range of activities are undertaken by a government this domain is right-sized by the constraint of sustainability.

Institutional arrangements

During the dot-com bubble, the universal metric for success was the ability to raise funding from investors. The imperative to implement a coherent revenue model affording good prospects of generating actual operating profits took a back seat. But at some point a company should be able to stand on its own two feet and generate profits for the investors, as opposed to an exit strategy predicated on flipping the equity to a greater fool, thereby capturing capital gains.

An analysis of institutional arrangements should consider roles – which entities have are required and enabled to perform which functions. Do revenues generated in performance of the role adequately compensate the direct costs and risks? If not does some other role generate excess revenue that might rationally be applied to subsidizing the unprofitable role?

System Provider

Issuer(s)

Primary Dealers

Participant financial institutions

Depository Institutions

Exchange Providers

GSACH

Transaction and business/revenue model

System design and technical implementation

HDLT™ (Highly Distributed Ledger Technology)

Note http://plasma.io/plasma.pdf Buterin tosses off unexamined assertion “it’s understood that a gross settlement system has difficulty scaling”. Not true for conventional RTGS.

Appendix - The Cool Kids’ New Clothes

Ptolemaic cosmology was grounded on the premise that the Earth is stationary and at the centre of the universe and that planets and celestial bodies orbit around the earth more or less like the moon (which actually does). Given the benefit of later understandings, most schoolchildren find it hard to imagine how anyone could have ever been so dumb as to subscribe to such nonsense. Modern pedagogy even provides a simple formula to explain the once-universal embrace of these now discredited theories – they were reinforced by similarly outmoded religious beliefs. Yet a closer examination of the Ptolemaic refinements of earlier geocentric models quickly reveals considerable sophistication. It is as hard to wade through even the Wikipedia introduction to epicycles, equants and deferents as it would be for a normal person to digest Vitalik Buterin’s discussion of sharding on the Github Ethereum wiki.

The Hans Christian Anderson tale of the Emperor’s New Clothes adroitly illustrates the very real dynamic faced by a would-be dissenter to a widely held and seemingly authoritative meme. If so many really smart people unanimously embrace certain principles, especially when couched in impenetrable but plausible sounding jargon, one expressing a contrary view risks the criticism of “you just don’t get it” or, worse yet, of being not only genuinely wrong but also on the wrong side of an epochal paradigm shift.

[develop case that an edifice built on flawed premise encounters the same ramifications of complexity as a sequence initiated by a lie. Embedded flaws lead to unintended and perhaps unanticipated consequences, each leading to further interventions in an ever increasing web of complexity. This is seen very vividly in basic blockchain, payment channels, banking as we know it and discretionary monetary policies.]

[note with e-gold there were ever more effective efforts to thwart exploits by people acting on false belief the system was anonymous or criminal friendly. This is the opposite of the accelerating efforts to strengthen the anonymity of cryptocurrencies]

[

In the case of blockchain/DLT enthusiasm, instead of an impressionable emperor manipulated by an unscrupulous hustler, the instigators themselves –the self-styled "Cypherpunks" of the 1990's – are also the yet-to-be exposed objects of admiration.

I propose to examine the origins and premises of the blockchain movement on the principle that deductions arising from invalid premises, similar to the Ptolemaic example, are not sound. The unexamined premises directly

The worldview of these "cool kids", and its impact on the blockchain/DLT phenomenon is examined in the section "Premises and Origins. An underexamined element that played a significant role in the formulation and introduction of Bitcoin was the interactions between the Cypherpunk community and the author of this paper, dating back to 1997.

The discussion that follows examines the institutional arrangements, transaction models and technological foundations of DLT pursuant to a polemic agenda – to argue the superiority of what I call the Global Standard System and its technological foundation, HDLT (Highly Distributed Ledger Technology).

Don't forget to liken to Keynes and his generation and maybe to Bloomsbury in particular.]