Public Distributed Ledger Networks

Market Assessment

pwc

Disclaimer

The information, statements, statistics and commentary contained in this report have been prepared by PwC from material provided by Hashgraph Consortium and publicly available material. PwC may, at its absolute discretion and without any obligation to do so, update, amend or supplement this document.

PwC does not express an opinion as to the accuracy or completeness of the information provided, the assumptions made by the parties that provided the information or any conclusions reached by those parties. PwC disclaims any and all liability arising from actions taken in response to this report. PwC disclaims any and all liability for any investment or strategic decisions made as a consequence of information contained in this report.

PwC, its employees, and any persons associated with the preparation of the enclosed documents are in no way responsible for any errors or omissions in the enclosed document resulting from any inaccuracy, misdescription or incompleteness of information provided or from assumptions made or opinions reached by the parties that provided information.

Note: Some of the authors of this report have invested in Hedera tokens. To maintain full objectivity, the review and sign off of the content was completed with PwC global blockchain leaders who do not have any Hashgraph based investments.

This report is not for public disclosure. Hashgraph have agreed within the terms of engaging PwC for this assessment to issue this report only to agreed upon parties (private PwC permissioned distribution of this report).

Table of Contents

Executive Summary

Fast forward two years and the reality of micro payment data, generated from many thousands of networked devices, will be voluminous. Current payment networks are in no way designed to be able to accommodate this new level of service demand.

During the same two years we are trading borderless digital securities with the current clearing and settlement platforms unable to cost effectively handle the change in service requirements. Across developing nations, telco based marketplaces have handed 1.5 billion unbanked humans access to digital financial services.

New at scale technology services are required and so the question is, are blockchains or DLTs going to solve for this new demand?

Speed of commits to immutable record is key to the service extensibility and therefore success of any of the new networks. Having this as a working assumption we conclude that proof-of-work blockchain-based networks are going to be limited, mainly due to their consensus method (as featured later), in where they take and keep market share. This leads us to focus on the more recent market entrants and conclude that using directed acyclic graphs (DAGs) and not blockchains will take the lead in this space as they have, in theory, a superior consensus type for managing scale, throughput, latency, and in some cases security.

Notable Events accelerating as more investment pours in

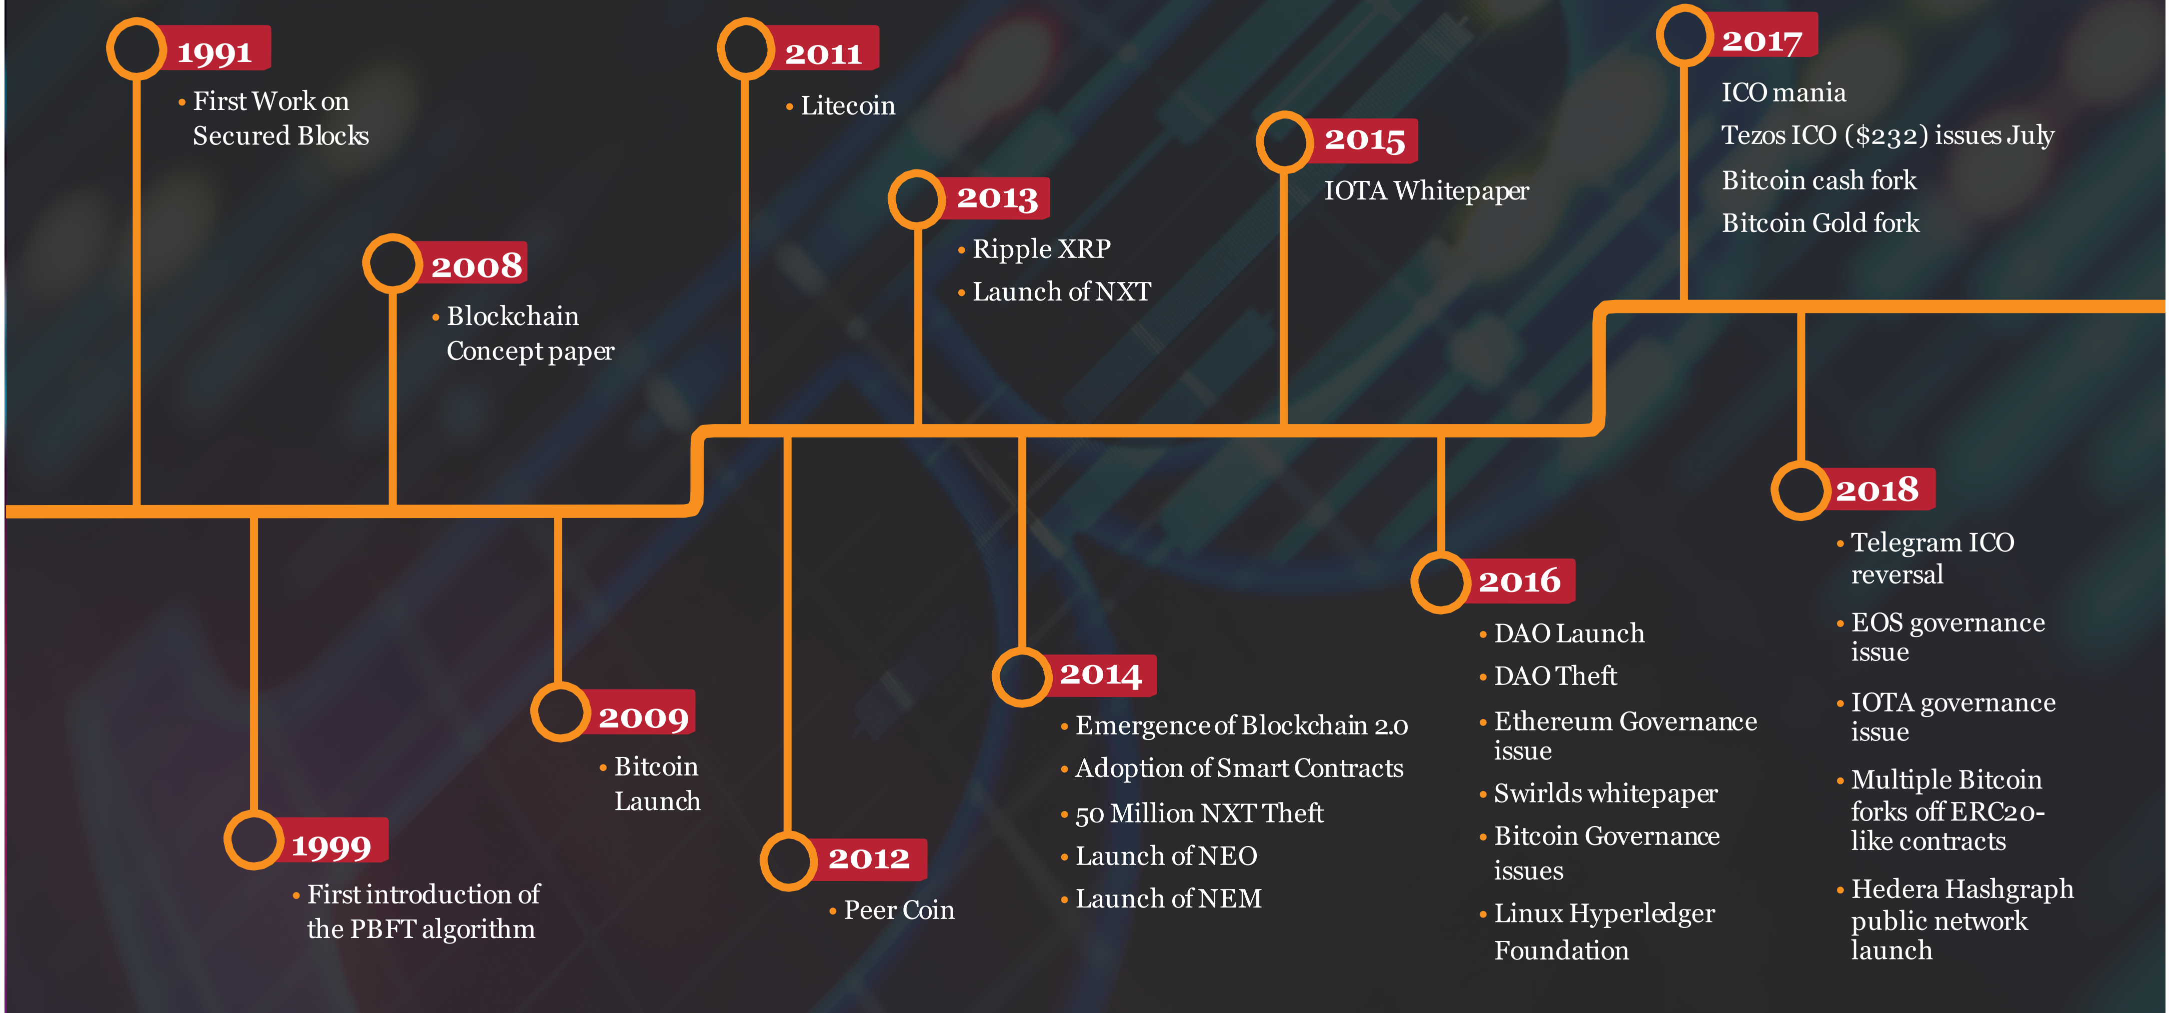

1991
- First Work on Secured Blocks

1999
- First introduction of the PBFT algorithm

2008
- Blockchain Concept paper

2009
- Bitcoin Launch

2011
- Litecoin

2012
- Peer Coin

2013
- Ripple XRP
- Launch of NXT

2014
- Emergence of Blockchain 2.0
- Adoption of Smart Contracts
- 50 Million NXT Theft
- Launch of NEO
- Launch of NEM

2015

IOTA Whitepaper

2016
- DAO Launch
- DAO Theft
- Ethereum Governance issue
- Swirlds whitepaper
- Bitcoin Governance issues
- Linux Hyperledger Foundation

2017

ICO mania

Tezos ICO ($232) issues July

Bitcoin cash fork

Bitcoin Gold fork

2018
- Telegram ICO reversal
- EOS governance issue
- IOTA governance issue
- Multiple Bitcoin forks off ERC20-like contracts
- Hedera Hashgraph public network launch

Technologies

Public DLT networks hold the promise to revolutionise commerce. They will provide for the seamless creation and execution of complex financial contracts at a scale that runs nations, whilst also making practical extremely small payments at tremendously high volumes characteristic of the Internet of Things (IoT). All this will be possible in a robust, decentralised, and trusted manner, without intermediaries and often without any human involvement at all.

Historically, parties to commercial agreements and exchanges have maintained their own independent records; frequently this leads to discrepancies and the need for a time-consuming reconciliation process. In many cases a trusted third-party may operate as an intermediary to facilitate the transactions, but in turn extract a fee. All such overhead represents a barrier to efficient trade. In contrast, a trusted, dynamic, real-time consensual view of mutual transactions would eliminate much of this overhead.

Until DLTs entered the fray, the likely implementation of this would have been a cloud-hosted database with appropriate business logic and access permissions; such systems have been built by consortia, where the economic and political drivers are overwhelming and the operating cost can be borne by the enterprises involved, but the technical and legal costs and complexities for such single-point solutions are often prohibitive.

In contrast, nascent public DLT networks provide an infrastructure and in some cases framework for the construction of trusted distributed applications, driving down both implementation and operational costs.

In this report we examine eight public networks and their various associated technologies.

We distinguish public from private (“permissioned”) DLT networks; private networks are deployments of DLT technology where the participating entities are known to each other in advance and form a relatively stable closed user group, either divisions within an individual company or a consortium, and although in some cases the same underlying technology may be used in both private and public deployments there are typically additional concerns with a public deployment, which faces a more dynamic and high-risk environment. Leading private (enterprise) DLT technologies include IBM Hyperledger Fabric, R3 Corda, Digital Asset, and Ripple.

Whereas purely private DLT networks are relatively simple to define by their closed nature, there are degrees of openness that may be associated with networks that have the ambition to be a public utility.

At the most open end of the spectrum are networks like Bitcoin and Ethereum, where there is no control over who joins and they are close to an anarchy, albeit with behaviour incentivised through game theory and individual gain; these represent a strong appeal in some quarters and an equally strong antithesis in others. These are truly “permissionless”.

In a more middle ground are networks that are open to join, but where some of the functionality is retained to a more limited group of nodes; this may be a transitory situation to guarantee stability as the network grows, or a more permanent feature.

Further on in the spectrum are networks where membership is nominally open to all, but is controlled by a governing body of some form. One could argue that this is closer to the consortium model, but the key point from our perspective is that the intent is to allow anyone to operate a node so long as they meet basic fitness criteria.

We have therefore chosen to include in this report networks on the whole spectrum above with pubic utility ambition.

Technologies

Blockchains, as the name suggests, are ledgers composed of a chain of blocks, reaching right back to the first (genesis) block and each block containing a number of transactions. New transactions are broadcast to the network by a submitting node and remain pending until selected for inclusion into a new block by a miner node. New blocks are produced at intervals by miner nodes, attached to the current head of the chain with a backward-linking hash of the previous block and broadcast to the network.

DAGs data structure is a mesh of events, each event composed of one or more transactions and linked to two prior events (by hashes of those events, similarly to a blockchain). Prior in this instance means previously in time at the event-issuing node, which due to network latency and simultaneity concerns may not necessarily be the same ordering as seen by other nodes, which implies the need for an additional process to extract a consensus order. This mesh forms a Directed (backwards in local time) Acyclic (no loops) Graph (a mathematical construct consisting of vertices (=events) and edges (=backward pointing hashes).

Focusing of public DLT networks, we further refine the grouping into blockchains and Directed Acyclic Graph-based (DAG) technologies, both so-named for their respective fundamental data structures.

The eight public DLT networks chosen for inclusion here range in ambition and scope, from those that offer only a cryptocurrency capability, through those providing an API for building distributed applications, to those seeking to create a global virtual computer. All though include a cryptocurrency as a fundamental feature. The eight are:

Bitcoin: Blockchain - cryptocurrency with some scripting capability around payments, Proof of Work based

Ethereum: Blockchain - distributed virtual computer, Proof of Work based

Nem: Blockchain – distributed application platform using Proof of Importance

EOS: Blockchain – distributed operating system using BFT-Delegated Proof of Stake hybrid consensus

IOTA: DAG – feeless cryptocurrency for Internet of Things

Hedera Hashgraph: DAG – distributed application platform using unique consensus mechanism offering deterministic asynchronous BFT (aBFT)

Byteball: DAG – decentralised database

Nano (RaiBlocks): DAG – feeless cryptocurrency

Technologies: Blockchains

Blockchains grow a block at a time as the chosen miner groups waiting transactions into a block, executes any actions associated with them and broadcasts the new block to the network. Subject to the consensus model chosen, the chain may fork as more than one miner produces blocks, though eventually a mainchain will be chosen and the fork becomes an orphaned branch. Transactions on an orphaned branch are not considered confirmed and they effectively remain in the pool awaiting inclusion in a future block.

The block production rate is a key design choice in blockchains: longer times, particularly where they are substantially larger than the global network block propagation time, reduce the amount of short-term forking that occurs due to miners concurrently producing valid but different blocks. They also enhance security by making it much harder for fraudulent miners to rewrite history several blocks back from the head of the chain. The main negative is that transaction inclusion and confirmation (latency) can take substantially longer – often more than an hour with Bitcoin.

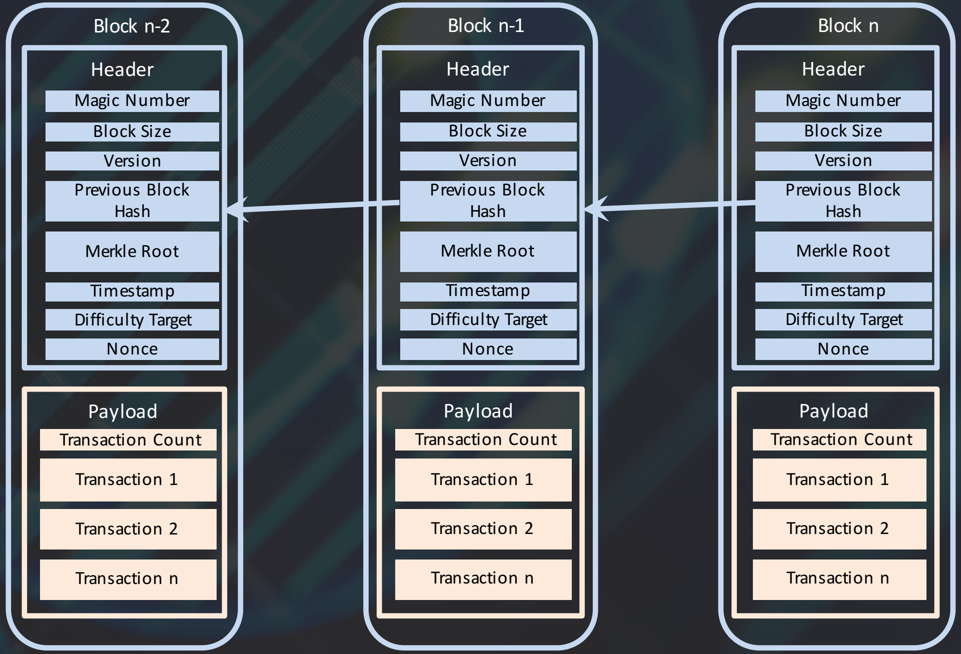

Block n-2

Header
Magic Number
Block Size
Version
Previous Block Hash
Merkle Root
Timestamp
Difficulty Target
Nonce

Payload
Transaction Count
Transaction 1
Transaction 2
Transaction n

Block n-1

Header
Magic Number
Block Size
Version
Previous Block Hash
Merkle Root
Timestamp
Difficulty Target
Nonce

Payload
Transaction Count
Transaction 1
Transaction 2
Transaction n

Block n

Header
Magic Number
Block Size
Version
Previous Block Hash
Merkle Root
Timestamp
Difficulty Target
Nonce

Payload
Transaction Count
Transaction 1
Transaction 2
Transaction n

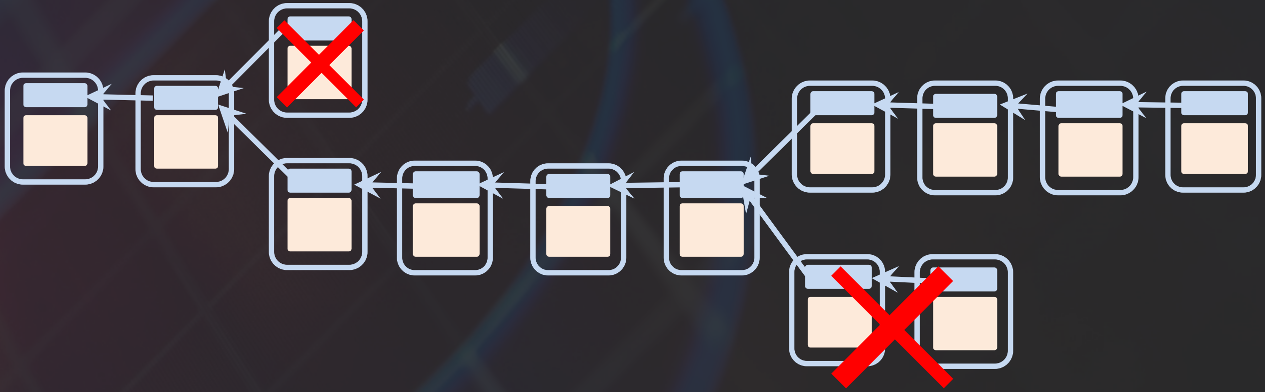

Technologies: DAGs

In contrast to blockchains, DAGs progress in a more granular and parallel fashion an event at a time and although an event may contain a number of transactions, in some cases this may be just a single transaction or a metadata update. With DAGs there can be multiple concurrent events emitted by active nodes, not effectively limited to a group of miners with massive compute (hashrate) power, and the network as a whole eventually determines order. Various DAG algorithms make different choices regarding how to reach this distributed consensus, how fast, how efficiently, and how deterministically that happens. They also make varying choices of the information about the events recorded into the DAG. Some broadcast events to the whole network, others are more point-to-point or use a group of trusted witnesses to validate the growing DAG.

DAGs hold the promise to have higher throughput, lower latency, be more scalable, more robust, and fairer than blockchains.

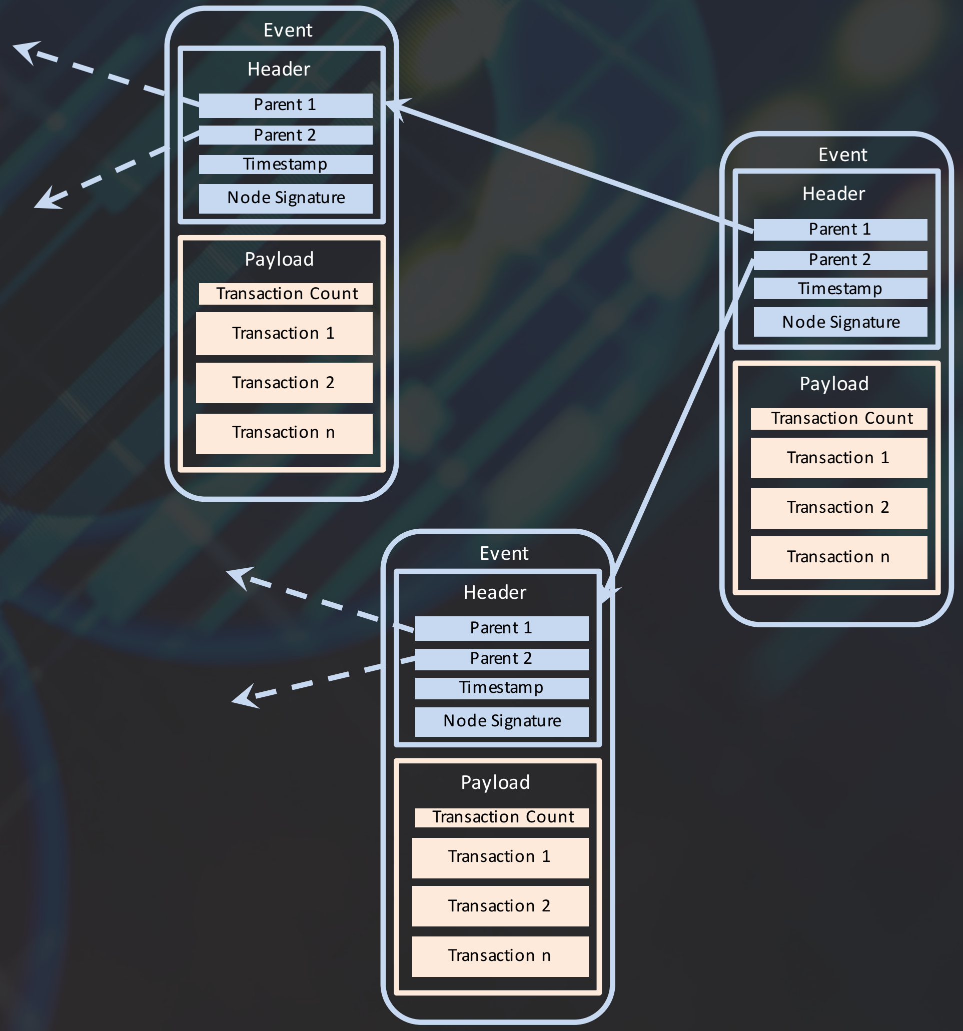

Event

Header

Parent 1

Parent 2

Timestamp

Node Signature

Payload

Transaction Count

Transaction 1

Transaction 2

Transaction n

Event

Header

Parent 1

Parent 2

Timestamp

Node Signature

Payload

Transaction Count

Transaction 1

Transaction 2

Transaction n

Event

Header

Parent 1

Parent 2

Timestamp

Node Signature

Payload

Transaction Count

Transaction 1

Transaction 2

Transaction n

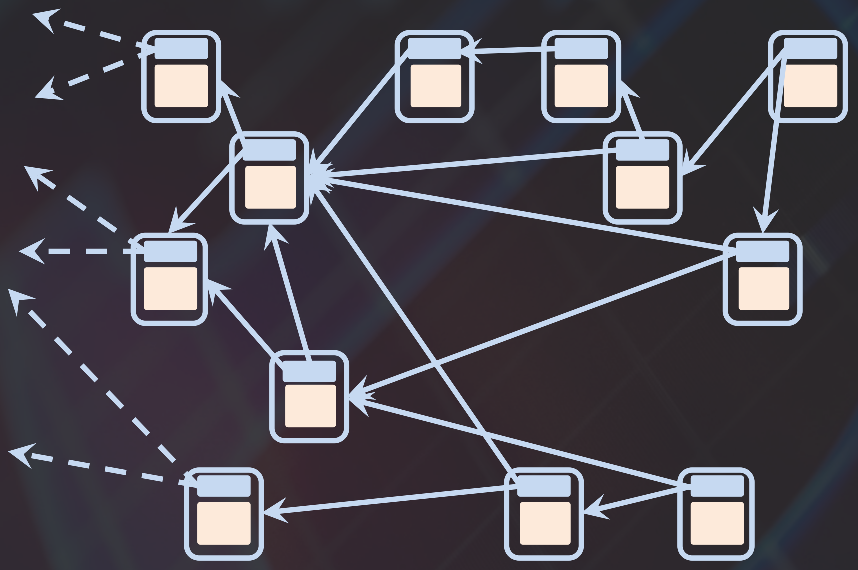

Table of Technology Comparison Categories

Consensus	Performance	Fairness	Security	Programmability	Governance
The method by which the network agrees on the order of events and implied state (e.g. of account balances).	Generally measured in terms of throughput in transactions per second (TPS) with consideration also given to latency and scalability.	Ability of a decentralized technology in treating all transactions in the same way. Our Focus is on timeliness, ordering, and costs.	Consideration of vulnerabilities to: Consensus Hijack, Centralisation, DDoS, Sybil Attack, Smart Contracts, Crypto Strength, Open Source, Other.	The degree of sophistication and usability of any end user programming features in the network.	The Governance of a public distributed ledger is the system by which it is controlled and operated and the mechanisms by which it, and its people, are held to account. Ethics, risk management, compliance and administration are all key elements being considered.
Comparative score out of 10	Comparative scoring poor, moderate, good, excellent	Comparative scoring of poor, good, great	Comparative scoring 0 to 5	Range is from basic scripting through to a virtual global computer	Comparative score out of 10.

It should be noted that the scoring used by PwC is completely subjective from the PwC technical authors' perspectives, taking into consideration all quantitative and qualitative data examined.

Distributed Consensus

In any distributed system consisting of multiple independent actors that seeks to maintain a common view of some shared state, for example an account balance, it is necessary for the system as a whole to reach a consensus on what that state is and therefore the order of any transactions that contribute to that state.

Thus, for Distributed Ledger Technologies (Blockchains, DAGs etc), consensus is the primary functional element and what form it takes and how it is implemented has fundamental impact on the operation of the different products, from performance through security, fairness, and economics. The choice of consensus algorithm defines the art of the possible with any DLT.

In a globally distributed public system, where it can be expected that many bad actors will be present trying to exploit the system and network failures are a daily reality, the consensus mechanism must be robustly secure and reliably converge to a common result in a timely and efficient manner.

Research on distributed consensus extends back decades, and relates to work on State Machine Replication, Distributed / Replicated Databases, ACID compliance (Atomic, Consistent, Isolated, Durable transactions), CAP (Consistency, Availability, Partition tolerance - pick 2), and other core concerns of computer science.

Various mechanisms have been proposed over time and continue to have applicability in appropriate use cases: two phase commit protocol (2PC), using a designated leader; three phase commit (3PC), to address blocking failure modes of 2PC; Paxos; RAFT; …

Paxos for example relies on a leader node (can be any node) to propose itself, gain acceptance from the majority of other nodes, propose a consensus, gain agreement by the majority, and if all this succeeds inform everyone of the outcome. This works well with a limited number of nodes (high communication overhead) and where they are broadly cooperating, but not with Byzantine failure or at scale (exponential message growth).

These mechanisms and issues seem arcane and abstruse, but are fundamental to the safe and performant operation of public ledgers; the choices made amongst the different ledger technologies in this area have deep and far-reaching impact.

Byzantine Fault Tolerance (BFT)

Traditional consensus protocols, such as Paxos, are designed to handle failures amongst a group of cooperating nodes, but fail when presented with nodes behaving unreliably or maliciously such that the rest of the network struggles to reach a consensus as to whether to include those nodes in the more general consensus mechanism. This is termed Byzantine failure from the **Byzantine Generals Problem**.

In 1982 Lamport & Pease showed consensus fails with the possibility of bad actors (Byzantine failure) when more than a third of actors are failed / bad. The solution where less than a third are bad is Byzantine Fault Tolerance, but this requires as many coordination rounds as there are failures and is impractical in most systems.

One of the key advantages of all the BFT protocols is that they reach consensus deterministically rather than probabilistically – i.e., you absolutely know when your transaction is committed.

Practical BFT (pBFT)

Introduced in 1999 by Castro and Liskov, Practical BFT attempts a realistic (practical) version proven for actual deployment as opposed to a theoretical model.

A leader is chosen to coordinate the consensus from a group of known members. If the leader is unresponsive a new leader is chosen. Correct so long as $f < (n-1)/3$, where n is the total number of nodes and f is the number of faulty nodes.

Asynchronous regarding correctness, synchronous for liveness (weakly synchronous). Makes optimizations over previous approaches including cryptographic signatures and digests, tentative early execution, and elimination of redundant full responses.

Potentially vulnerable to successive Denial of Service (DoS) attacks against the leader, requires a known set of members, and still involves significant (point-to-point) communications overhead to reach consensus such that as n becomes large the protocol becomes increasingly impractical - O(n2). Applicable to private enterprise deployments with a limited number of nodes.

Distributed Consensus

Asynchronous BFT (aBFT)

Whereas pBFT assumes that the supporting network is weakly synchronous – that messages will eventually arrive after some bounded time – aBFT loosens this constraint to only require that some message eventually arrives from a node.

Operating in the unreliable Internet with potentially malicious intermediaries this is a significant step forward and represents leading edge consensus technology, currently implemented by HoneyBadgerBFT and the Hashgraph algorithm.

Proof of Work (PoW)

Proof of Work requires that a node prove it has invested significant computational resources (work/cost) in forming a new block for the head of the shared blockchain (mining). It shares this new block with other nodes, who if they agree with its validity will begin trying to produce the following block based on this one. In this way, the chain grows as a series of blocks .

Should another node produce an alternative valid block at the same time (effectively within the propagation delay window of the network), then a fork occurs: some of the network will be building a chain based on one new head and some on the other, subject to which version they received first. Eventually, one branch will grow faster than the alternatives and become the consensual state.

The proof of work is typically finding a cryptographic hash of a block of data subject to certain difficulty criteria and soluble only through brute force (number of different hashes tried), hence the **hashrate** of a miner directly influences its probability of successful mining (and the attached reward) and getting its block accepted as the shared truth.

Problems with PoW include a significant waste of resources with the energy and equipment required for mining, high latency, low throughput, and questionable fairness as the miner chooses which transactions and in which order to include in the block.

Whilst this approach has shown itself to be robust in reaching probable eventual consensus, it remains vulnerable to attacks if some party or cooperating parties gain control of 51% of the total network hashrate.

It only reaches consensus probabilistically for any given transaction as that transaction's block is appended by further blocks until it becomes very unlikely that it is on a fork that will eventually be discarded; this is widely considered to be at 6 blocks by convention from the initial Bitcoin client.

Proof of Stake (PoS)

Proof of Stake avoids the resource and power inefficiencies of PoW. The miner of the next block is chosen at random, but weighted by their proportion of the stake (operating currency) amongst online nodes. This has the proposed added advantage that it defends against the 51% attack of PoW in that those having the larger stakes in the system also have the most to lose through malfeasance.

A concern is the nothing-at-stake problem: in a fork, a miner is not disincentivized from working both branches - introducing penalties can mitigate this.

Delegated Proof of Stake (DPoS)

As for PoS, but allows nodes to delegate their stake to other nodes, introducing a form of trust-based voting and potential for increased efficiency by reducing the number of nodes directly involved in consensus.

Distributed Consensus

Proof of Importance (PoI)

Proof of Importance is related to Proof of Stake but seeks to address some concerns with it by taking into account how well-regarded a node is.

The exact measure of this may vary with implementation, but for instance could be based on the number of valid transactions a node has originated and the value of those transactions.

Proof of Elapsed Time (PoET)

Miner nodes participate in a lottery for the right to mine the next block. Basedon a random delay time they must wait before mining.

Relies on all nodes running trusted hardware (Intel SGX or similar) to guarantee the randomness (of elapsed time) and proof of having waited that period.

This is much more energy efficient than Proof of Work and has applicability to IoT scenarios, implemented in Intel Hyperledger Sawtooth.

Bitcoin	Ethereum	Nem	EOS	IOTA	Hedera Hashgraph	ByteBall	NANO
Probabilistic Proof of Work	Probabilistic Proof of Work	Proof of Importance (PoS + History)	Delegated Proof of Stake	Central Coordinator Markov Chain Monte Carlo Algorithm	Event based virtual voting using the aBFT Hashgraph Algorithm	12 Witness designated nodes	Delegated Proof of Stake
7/10	7/10	6/10	6/10	4/10	9/10	7/10	6/10

Performance

The performance of any public network is important, particularly one that carries payments. For example, the VISA network, widely considered to be the largest payments network, is claimed to be capable of a **throughput** of 56k Transactions Per Second (TPS). This is currently managing primarily human initiated transactions, but with the coming age of the Internet of Things (IoT), and the provision of infrastructures able to handle micropayments with realistic fees, there is an expectation that the volume of payments will grow very significantly, perhaps to 10M TPS. Any public network which intends to compete in this space must demonstrate realistic **scalability** and suitable throughput to be able to serve the next generation of digital commerce. Given the realities of global communication network latencies and throughput which underpin these DLT technologies, it is important that a sharding approach is available to support ultimate scalability.

Another key performance measure of such a system is the **latency**, or time that it takes to complete a payment. A typical use case that highlights this is in a retail point-of-sale situation, the payment needs to complete within a few seconds to a point that the customer is free to leave the store. When we are considering public DLT networks, such as blockchains and DAGs, there are finer degrees of latency to consider:

- the time it takes for a submitted transaction to be initially accepted into the system and broadcast to relevant nodes
- the time it takes for a transaction to be included in the ledger
- and the time it takes for the transaction to become practically irrevocable i.e., system consensus to be reached (generally accepted to be after 6 blocks in blockchain, but this is probabilistic rather than absolute).

NB: For Blockchains; anything that is only in the mempool is ephemeral and there is no guarantee it will ever be included into a block; even after it makes it into a block, there is a chance it is on a branch of the main chain that will be pruned in future, hence the 6 block length stipulation in many cases.

The scoring below is a snap shot in time and will change both positively and negatively as new features are rolled out.

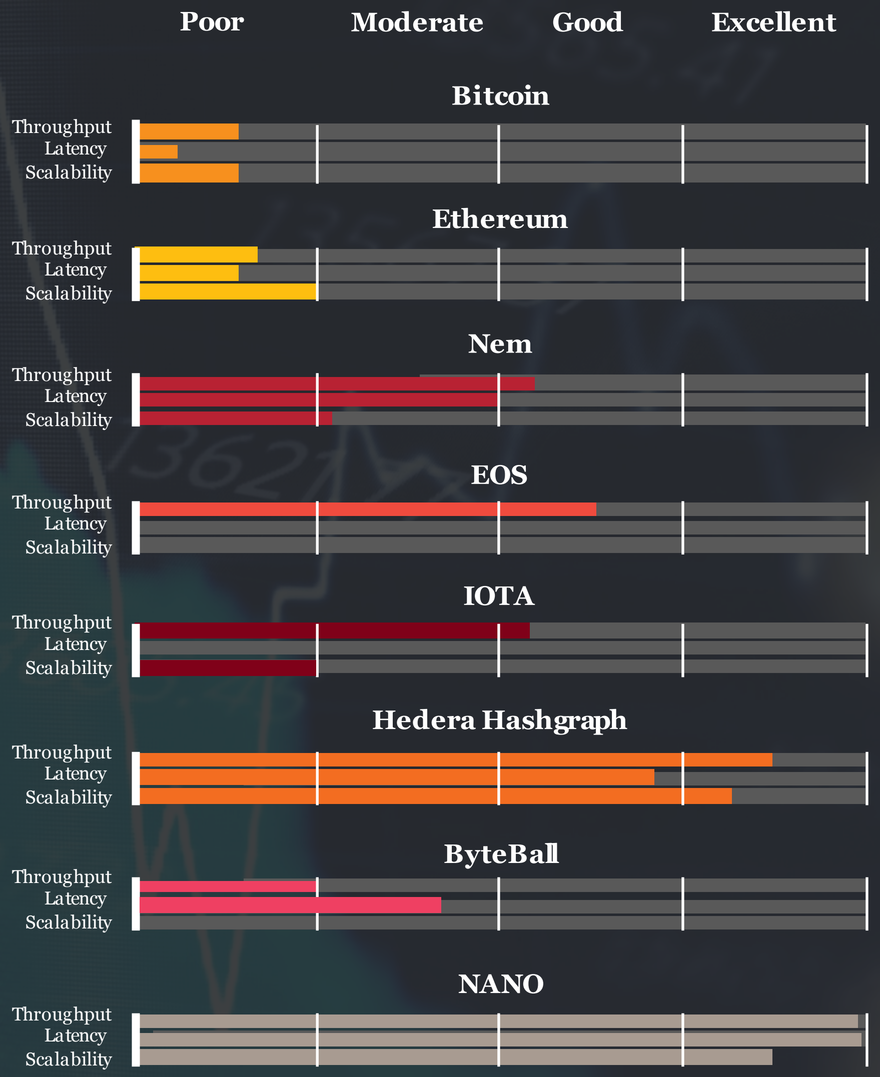

Performance information sourced from public reports including developers own. In some cases wildly conflicting little or no information available.

Fairness

Order of transactions

The order of transactions agreed upon by the network consensus is a fundamental property that can be extremely important in some systems, less so in others. For example, in a pure cryptocurrency, one spend appearing prior to or after another is not terribly important unless they are conflicting transactions (double spend) or constrained by a balance (payment in before payment out). However, in other systems where order is important, such as an auction, the ability of a party (e.g.miner) to alter the order of transactions as recorded in the ledger compared to reality is a real concern.

Censorship

The potential for transactions (from some users) to be delayed, perhaps indefinitely, is a problem in systems where a leader / miner creates the ledger, or where a limited subgroup of the system are responsible for consensus.

Stability (currency)

Users of a currency (other than speculators) generally seek a stable value in order for it to be useful in everyday transactions. The economic model of a digital currency may or may not encourage this.

Fees

Fees are generally necessary for transactions in a public system to pay for the cost of operating the system, but they should be as low as is commensurate with such operation so as not to impede trade. Fees should be designed to account for both short term and long term costs. The Internet of Things / micropayments are good examples where systems based on Proof of Work mining will likely necessitate impractical transaction fees.

Fees may be based on the transaction, its size, its storage requirements, and its compute requirements.

In miner-based systems, the actual fee may be determined by market conditions - the number of transactions awaiting processing and the fees they are offering to the miner for processing. Therefore, low value transactions may incur a higher percentage charge or wait a long time to be processed.

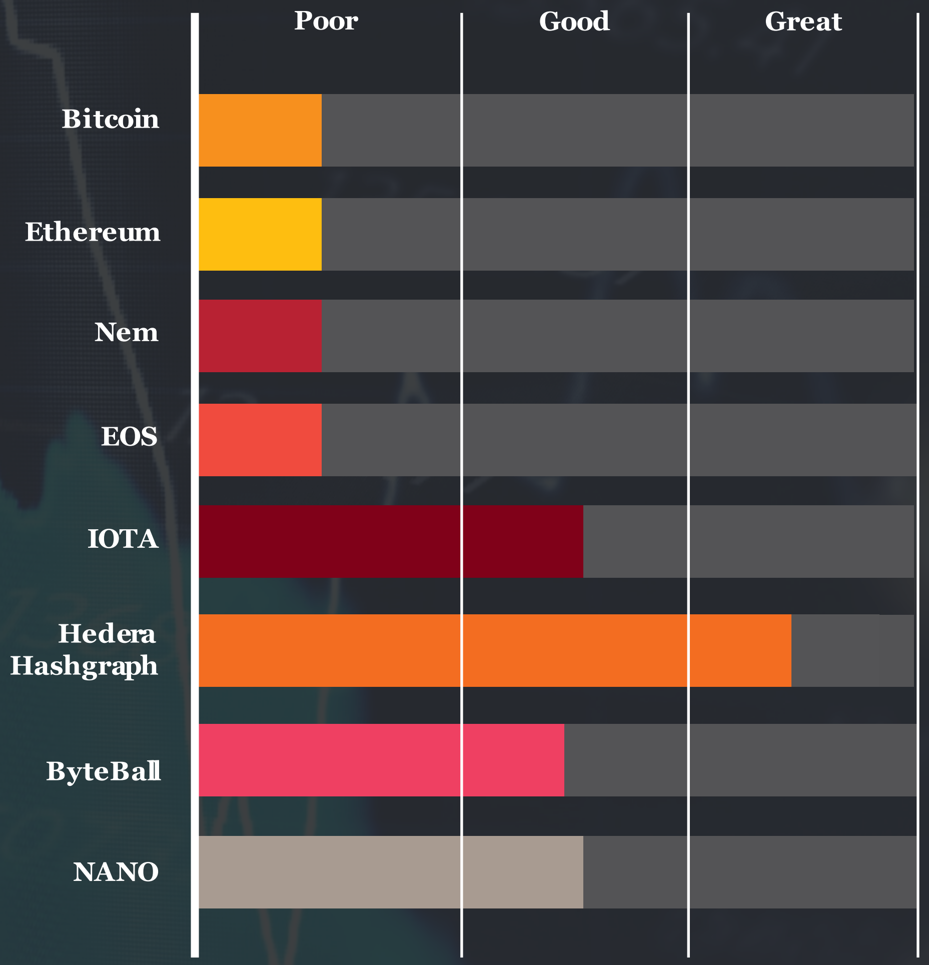

Security

Public Distributed Ledger Technology brings additional security issues but does not supplant many traditional digital security concerns:

Access Control - It is necessary to **Identify** actors in the system, be they software agents or humans (although these actors may be [pseudo-]anonymous), to **Authenticate** the actor and to apply appropriate **Authorisation** to their actions. This is primarily done using **Public Key Infrastructure** (PKI) and brings with it the usual issues of **Key Management** such as key loss and secure storage of **Private Keys**.

The ledger should have data **Integrity**, such that data stored and transmitted is not corrupted, provided by **Cryptographic Checksums** (hashes) and transactions would ideally have Atomic, Consistent, Isolated, Durable (ACID) compliance, which is challenging in highly distributed systems and a core responsibility of the chosen consensus algorithm.

Transactions should be subject to **Non-Repudiation** - "my word is my bond" - which can be achieved through PKI **Digital Signatures**.

Given that much of the transactional payload in public ledgers is financial in nature, **Confidentiality** is important to many users. Some ledgers take the approach that pseudo-anonymity of users through anonymous addresses is sufficient, but recent work has shown that historic analysis of the public ledger correlated with external information can uncover identity. There is also the concern that current PKI approaches, which rely on the computational infeasibility of factoring a large number into primes, may be broken using quantum computers at some point in the future; realistically, most in that field believe such a possibility is at least 10 years away, but some users would be concerned about historic cracking (and public ledger data exists indefinitely), and one cannot rule out a breakthrough sooner. Consequently, at least one of the ledgers we have looked at uses a cryptographic hash function that is claimed to be "quantum-immune", though that in itself is controversial as new cryptographic functions require extensive analysis and testing to verify their safety.

As implied by the above paragraphs, public ledgers make extensive use of **Cryptography** in solving the issues of identity, integrity, confidentiality, and non-repudiation: the careful design and implementation of the ledgers' use of and selection of cryptography is key to their security.

The 24x7 **Availability** of a public ledger network to process transactions and its overall fair and honest behaviour are crucial to widespread adoption. Availability can be compromised through design errors, software bugs, unmitigated hardware failures, or the deliberately disruptive actions of bad actors in the public network. The former elements may all be mitigated by common industry best practices, but bad actors are a special concern that warrants further consideration:

The potential vulnerabilities, or Attack Surface, of public DLT networks are quite extensive exactly because they are publicly networked distributed systems and also in part due to the subtleties of the various consensus algorithms. We consider four main areas of risk:

Consensus Hijack – dominance of the network by gaining more than 50% control and therefore being able to force own version of truth on other members. For example, in blockchain proof of work networks, having >50% hashrate (for a period of time); some proposed attack schemes, such as selfish miner collusion, suggest a vulnerability with as little as 25% hashrate.

Denial of Service – disabling progress on consensus by preventing communication by or with a subset of the nodes in the network, or enforcing own view of consensus by blocking a sufficient number of fair nodes. Usually enacted by a botnet to gain sufficient bandwidth and therefore being considered a Distributed DoS (DDoS). The enacted method may be one of:

- **Spamming** – sending large numbers of invalid messages
- **Transaction Flooding** – large numbers of valid, but unnecessary messages
- **Penny Spend** – large numbers of transactions with trivial value

Sybil / Sock Puppets – large number of "fake" members of the network acting in collusion to overwhelm genuine ones. Only possible where the cost of being an active member is low enough to support the number of fakes required (e.g. typically not with Proof of Work, Proof of Stake).

Smart Contract Vulnerability – subject to the design and coding of the Smart Contract / Distributed Applications and their interaction with the underlying language, libraries, virtual machine, and distributed facilities provided by the system, including the effects of concurrency and timing of consensus, there may be vulnerabilities to be exploited.

Security

Product	Consensus Hijack Resistance	Decentralisation	DDoS Resistance	Sybil Resistance	Smart Contracts	Crypto	Open Source	Other
Bitcoin	> 50% required maybe > 25% for [colluding] selfish miner	>75% hashrate in 6 miners	Broadly immune	PoW and transaction fees mitigate	Basic script only	Robust & proven	Yes	Blocktime >> network latency minimises forking March 2016 slowdown due to wallet spamming
Ethereum	> 50% required maybe > 25% for [colluding] selfish miner	>90% hashrate in 5 miners top pool approaching 50%	Broadly immune	PoW and transaction fees mitigate	Extensive. Solidity EVM	Robust & proven	Yes	10,000's vulnerabilities identified in smart contract code June 2016 $60M DAO split theft / hard fork 2017 $150M DevOp199 contract kill 2017 $32M Parity wallet hack
Nem	PoI mitigates	Unknown, PoI would mitigate	PoI mitigates	PoI mitigates	Off-chain Dapps through API	Mostly mainstream, some unusual choices	Yes	EigenTrust++ for PoI SHA3-512, Twisted Edwards Curve, Ed25519 sig
EOS	DPoS mitigates	DPoS 22 leaders	DPoS and 15/22 leader votes creates vulnerability	DPoS mitigates	Yes Choice of VM	*	Yes	* Insufficient information for analysis at time of writing
IOTA	Mitigated by coordinator currently. Eventually requires active network.	Currently requires central coordinator	Vulnerable with current central coordinator	PoW mitigates, but perhaps not sufficiently	Not yet	Own hash curl p	No	Hash had vulnerabilities, patched, claimed to be quantum resistant
HH	> 33% Proxy Stake required	Multiple Independent organisations to seed network	> 33% denial required	Proxy stake mitigates	Yes Solidity EVM	Proven choices	Open review	Deterministic aBFT algorithm looks like it brings some genuine advances
Byteball	Requires subverting witnesses	Currently witnesses largely under developer control	12 witnesses vulnerable	Trusted witnesses mitigate	Simple declarative	Proven choices	Yes	
Nano	DPoS mitigates	N/A	Vulnerabilities	PoS mitigates	No	*	Yes	Limited info in whitepaper * Insufficient information for analysis at time of writing

0 1 2 3 4 5

Programmability

The programmability of the public ledgers we are examining is a key feature in their utility as enablers of distributed commerce. This takes three broad forms, though opinions and implementations vary in their approach:

Basic Scripting

The technologies which are purely a cryptocurrency play generally include scripting functionality to satisfy more complex payment scenarios, including such things as multi-signatory payment and escrow accounts. However, they are not intended for general purpose computing and the languages are not normally Turing complete (capable of arbitrary computation).

Smart Contracts

Smart Contracts are an attempt to extend the core distributed ledger consensus capability to encompass general commerce without (significant) human intervention. The computer executable contracts are written in either a specialized Domain Specific Language (DSL) for contract processing, which may not be Turing complete by design so as to avoid DAO-like vulnerabilities, or may be a general purpose computer language sitting on a specialized API for contract and blockchain processing. The smart contracts for an application encompass all the actions that may take place and are intended to directly reflect legal agreements and obligations between organizations; they take inputs from other events on the ledger or from approved external sources of information (Oracles) and execute their logic to yield results back onto the ledger. The contract contents are tied to evidence on the ledger so that their state and code is all locked down and cannot be altered unilaterally. Different ledger technologies afford different degrees of support for smart contracts directly in the ledger itself versus essentially providing a toolkit that requires extension to provide full smart contract capability.

Distributed Apps (DApps)

DApps are a more generalised distributed computing capability than smart contracts. Distributed applications have existed for decades and the Internet, Web, and smartphones have made them ubiquitous. Historically they have used private links, mobile telcos, or the Internet for communication, but in the context of distributed public ledgers we gain an immutable and trusted record of transactions between these applications. DApps may be built directly onto a ledger product API, or may interact with the ledger via smart contracts. They may provide a human useable interface, perhaps via a website or mobile app. DApps require less complex supporting facilities from the ledger technology but with correspondingly less trust and control.

In our analysis of the technologies under consideration, we will broadly classify them according to whether they are principally cryptocurrencies, whether they provide an API providing for distributed applications, or whether they provide full smart contract capability.

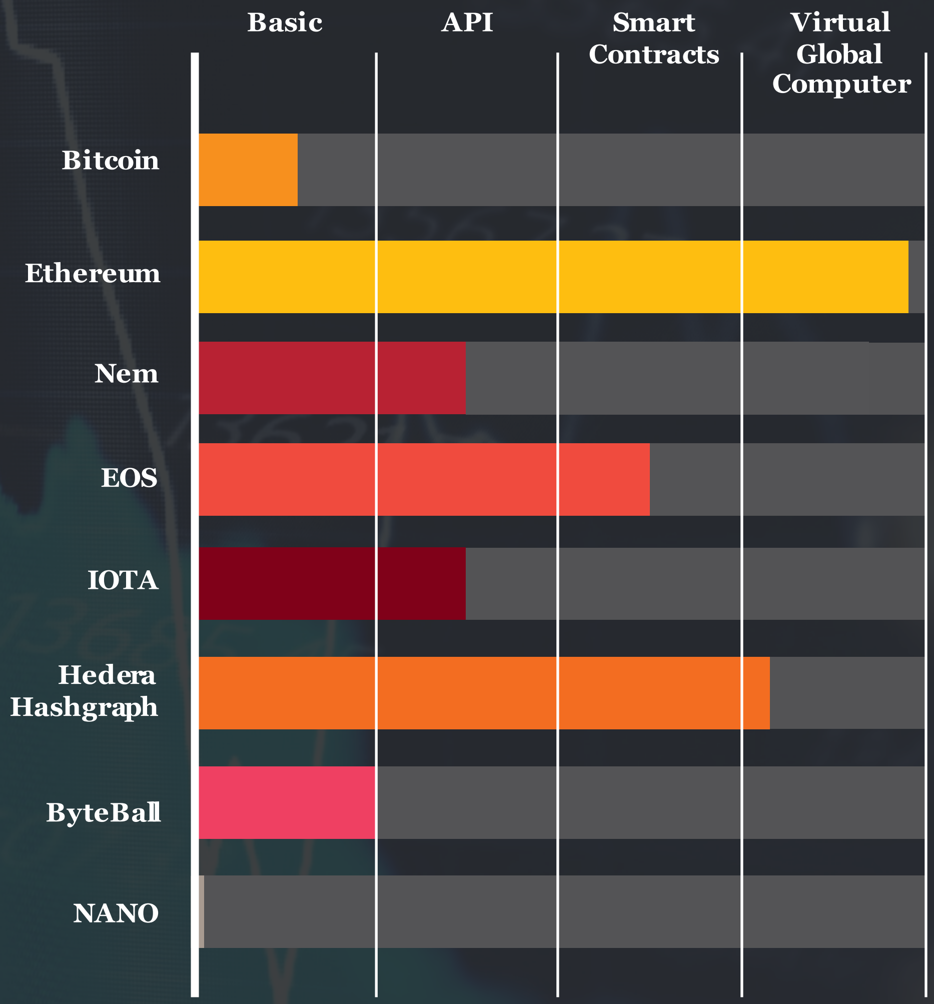

Governance

Drawing comparisons on Governance (or absence of) as a characteristic of any public network or cryptocurrency has a danger of being largely subjective. We need to consider the purpose or role of governance for each network and the "authority" or process that determines other elements such the token supply and the subsequent effect on the perceived underlying value, the technology roadmap and for instance service upgrades and network security.

Governance of the network is a fascinating arena to observe the tradeoffs from totally decentralized to centralized governance models and what falls out and how that affects usage and value of the cyrpto currency. Using Ethereum and Bitcoin as well documented cases, the former is regarded as a benevolent dictatorship which still managed to cause a fork (dilution in value) post the DAO episode. Bitcoin on the other hand original thesis was a social state for decentralized money.

Indeed most network governance has experienced disruption and in some cases mutiny by the human element with stand offs and warring factions over the technology roadmap within the developer community. Problems also exist between the developer community and the mining communities.

The DAG generation have seemingly learned from our two trailblazers and established Social Democracies or Governed public networks. The goals have been to establish foundations with a view for governance however with different levels of success to date.

Hedera Hashgraph are differentiating themselves from the pack by investing such a high degree of effort in signing up a global spread of International businesses to be independent council members with equal voting rights on the governance of the network.

Comparative Score

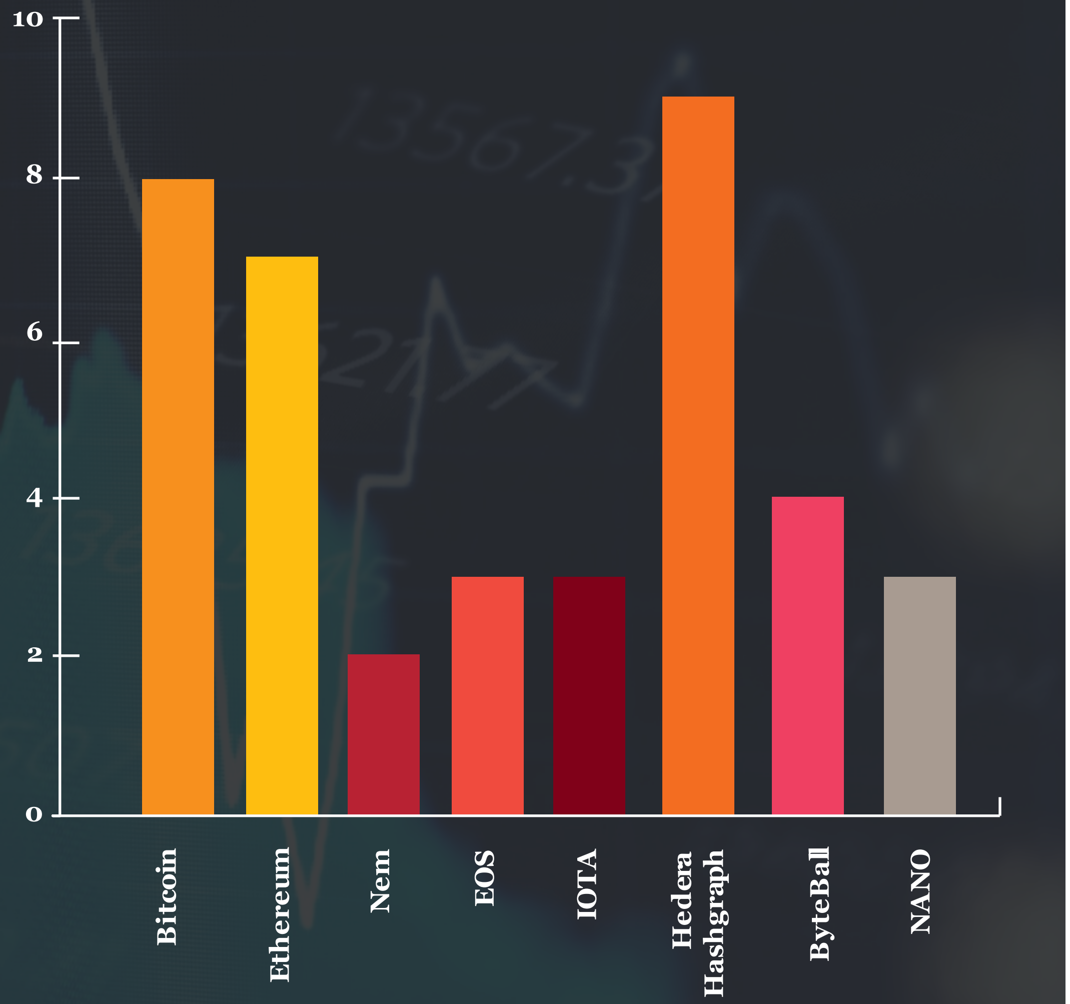

Closing Market Considerations

Summary of comparisons - we found throughout this analysis that the technology is evolving as one would expect, however some of the newer networks have put launching a cryptocurrency as a priority over perhaps the more considered approach of the likes of Bitcoin, Ethereum, and Hashgraph to focus in parallel on establishing the governance. It remains to be seen, but well governed enterprises are less likely to incur significant disruptions to their operations and trust models which is significant to becoming future leaders in this space.

Distributed ledgers are a new architectural combination of proven existing technologies making progress into well established commercial and government sectors. Their success is not dependent on creating brand new market places or consumer choices. They are however technology platforms and to replace or redesign current systems away from centralized databases is a work in progress which will take time and significant resources to properly commercialize. However, because using public DLT networks we can actually now build much more efficient distributed databases that share information more “naturally’ we believe it is not a question of if but when Distributed ledgers will where required become the “current” way of sharing common data.

The winners, and we think there will be more than one, will provide trusted, fair, low-cost, reliable, programmable, and scalable public environments that main street and start ups alike can embrace with their core offerings.

Technology Comparison Summary: Blockchains

	Consensus	Performance	Fairness	Security	Programmability	Ecosystem	Governance
Bitcoin	Proof of Work	Throughput: Low ~5TPS Latency: Very high Scalability: Very limited	Poor Block miner chooses transactions and order	Robust & Proven Concerns of miner centralisation	Limited scripting. Primarily digital currency.	Very extensive	Community: developers and miners. No central authority. Satoshi Roundtable
Ethereum	Proof of Work Plans to switch to Proof of Stake (CASPER / GHOST)	Throughput: Low, ~20 TPS Latency: High Scalability: Very limited	Poor Block miner chooses transactions and order	Robust & Proven Concerns of miner centralisation Smart Contract quality concerns	Full smart contracts	Common ICO platform	Community: developers and miners.
NEM	Proof of Importance	Throughput: Medium, 100s TPS Catapult release promises 1,000s TPS Latency: Average Scalability: limited	Poor Block miner chooses transactions and order	Appears broadly good Some unusual choices but not known to be a problem	Basic Basic operations via blockchain API for custom Dapps, outside blockchain.	Regional focus	NEM Council of few make decisions via committee. Mechanics of the council governance are not public.
EOS	Delegated Proof of Stake	Throughput: High Latency: Scalability:	Poor Block miner chooses transactions and order	Potential DDoS vulnerability in leaders	Full smart contracts	Extensive	Protocol based No accountability Founders Block.one not committed to providing governance.

Technology Comparison Summary: DAGs

	Consensus	Performance	Fairness	Security	Programmability	Ecosystem	Governance
IOTA	Tangle - DAG Minor Proof of Work for anti- spamming. “Coordinator” until network grows	Throughput: Med ~500TPS Latency: Scalability: Claimed very high	Transaction ordering not guaranteed Not time stamped No fee but minor PoW required	Concerns with centralisation of technology and control with developer currently DDoS vulnerability	Smart Contracts on the roadmap.	Growing	Foundation consisting of multinational conglomerates, research institutions and others from industry.
Hedera Hashgraph	DAG Asynchronous Byzantine Fault Tolerance through virtual voting on gossip of gossip	Throughput: V High, 100,000s TPS Latency: Seconds Scalability: Claimed very high through sharding	Transaction timestamp ordering Fees	Strong choices in design Unproven as public network – prior to launch	Smart Contracts (Solidity) Java API API for Dapp Interaction	Pre-launch Growing developer community	Well documented governance model for decentralization and stabilization Aims to prevent consolidation of power.
Byteball	DAG 12 Witness designated nodes under developer control	Throughput: Low ~20TPS Latency: Scalability:	Based on trusted Witness choices	Witnesses are currently primarily under control of the developer	Limited declarative smart contracts		No governance committee established. Still a centralized control model: Only founding development team can change the network protocol.
Nano	Block Lattice Delegated Proof of Stake	Throughput: High 100s demonstrated 10k+ claimed Latency: Very low Scalability: Notionally unlimited	N/A: ordering only within individual accounts No fees	Appears to have DDoS vulnerabilities	None: Cryptocurrency only		Little information on Nano governance. Research implies that only the Nano development team can manage the network.

Appendix – Building on DLT

DLT Implementation of Utility Services

With DLT (be it blockchain or DAG) being improved by many private and public initiatives, there is building promise that it could underpin the consolidation of data utilized by nation states.

Stand-alone and/or bespoke data implementation and management remain the norm across current government services and general utilities that serve the public. However, as and when the capabilities of DLT provides a viable platform upon which to consolidate such pools of data, then considerable operational benefits would start to be realized.

DLT could become the technology of choice that refresh programs migrate to when legacy systems reach end of life, or when operational efficiencies gained on DLT compel an earlier decommissioning.

Example utility services where DLT is being tested

National Identity – Estonia is a leading nation in the implementation of 'e-identity'. Instead of a citizen having their information stored in stand-alone databases, DLT could consolidate under one account relevant information such as a person's birth, marriage(s), and death. It may even be linked to any tax number and/or social security number. It could capture residency status, and issued passport(s) against an individual. Certainly, along with appropriate access control to such information, an individual could allow verification of their identity as and when appropriate. In the case of Estonia they allow e-identity to be utilized in application for bank accounts.

Government Services – Even if a national identity DLT were not set up, the provision of government services such as local, state and national taxation, social security, land titles registry, police/criminal records, voting and public healthcare provision could be consolidated under DLTs. If each were in the end linked to an e-identity, then any nation would have a much clearer view of exactly what services a citizen or foreign national would have consumed, as well as their contributions to the nation.

Energy and Water Utilities – The management of energy and water resources remains a topic of much discussion, especially in the westernized world. DLT as the database that underpins the collection of individual residential or business consumption, could provide much insight into better management of resources be it electricity, gas or water. Paired with internet connected monitoring devices, DLT would serve to inform consumers directly of their consumption, as well as allow utility companies and government regulators to better manage supplies.

Notary and Conveyancing Services – The officiating of any documentation by a Notary could be served by DLT. Though at this stage the face-to-face verification of one's identity by the Notary would need to continue, there is nothing to stop the notarizing be done as part of a digitized process underpinned by a DLT. Equally, the rigor executed by solicitors on conveying property titles could be readily underpinned by DLT. In both cases, as expanded upon in the following section, the governing process that instructs on how participants transact with each other (and hence record such transactions on the DLT), can be done via 'Smart Contracts'.

Building Service Layers

DLT is a foundation technology upon which services could be realised. As per prior use case examples, because of its capability to distribute data amongst participants and keep that data in synch and kept safe from those not participating, it brings with it an architecture that many current services could use to realise operational efficiencies and security, not readily gained via more traditional approaches.

The validation of its value as a new technology is still ongoing. Even now, at almost 10 years on from Bitcoin's launch, the investment in legacy systems, the imbedded processes that rely on them, and the ability of the technology to disintermediate many services (encroaching on existing business models), has meant that adoption is still very much at the exploratory 'proof-of-concept' phase. Albeit starting to gain momentum towards production. As mentioned earlier, nations such as Estonia have embraced blockchain as a fundamental building block for their e-identity system.

The shape of efforts is naturally starting to fall into domain specific applications that may be supported by the likes of Swirlds Hashgraph, IBM Fabric on Hyperledger and R3 Corda for permissioned enterprise services (R3 specifically targeting Finance and Commerce). And those for example like Consensys on Ethereum and Hedera Hashgraph, aim to provide decentralised services that could apply to any given service (be it permissioned or not). In either case, implementations are generally built on a new set of languages called 'Smart Contracts'.

Smart Contracts

The processes that govern most (if not all) transactions between parties can be formally digitized using 'Smart Contracts' as a service provisioning layer. That is, any given agreement between parties could be programmed into a smart contract that in an exact way describes all the transactions between each party, and the governing pre and post conditions to each transaction. A smart contract would prescribe the step-by-step execution of the process and what would be expected of each party with each execution step. For instance, if a smart contract were written to govern the conveyancing process, it would prescribe steps and obligations of each party (buyer, seller, estate agent, mortgage provider) during pre-contract (before striking an agreement on the sale price), at exchange of contracts, then right through to settlement.

The underlying DLT would hence record the transactions between parties as would be governed by the smart contract.

Smart Contracts as a service enabler could in future prove so versatile, that particular process provisioning could in themselves become mainstream offerings. So in the example above, a company could specialize in conveyancing processing on Smart Contracts and emerge as a market leader.

Speed to Market and Crowd Co-operatives

The advent of DLT and Smart Contracts is further evolution in the provisioning of technology that could serve any business or service. They inherently should increase speed to market. Moreover, there is current

momentum in the establishment of crowd co-operatives that challenge traditional businesses such as personal lending and investment. DLT and Smart Contracts, perhaps not unsurprisingly, are a perfect vehicle upon which these new entities could quickly establish operations and speed their journey to market.

The Road Ahead

The Harvard Business Review [18] in an article about blockchain as a DLT highlighted that: "Blockchain is not a 'disruptive' technology, which can attack a traditional business model with a lower-cost solution and overtake incumbent firms quickly. Blockchain is a foundational technology. It has the potential to create new foundations for our economic and social systems."

There is still truth in DLT being able to disrupt industry, in as much as incumbent firms could be usurped if they do not readily utilize its benefits before competitors or indeed crowd-centric cooperatives do.

That said, traditional businesses that have invested in proof of concepts should be well poised to leverage their own efficiencies using the new technology.

DLT is foundational. There is no doubt that unless further advancements in technology is made that would surpass DLT, the very fabric of how data can be utilized and secured, and thus how services could be better formed, DLT is the most likely candidate onto which ageing and less efficient legacy systems would give ground.

The Marketplace for Distributed Ledger Technology

Because of Bitcoin, blockchain as a distributed ledger technology is nearing 10 years of service. Most who have come to understand DLT, have begun to see that its application in the world may be as numerous as the use of databases that serve many users [9]. So, for any service that performs the sharing of information or enacting of transactions between parties that do not know each other but need to trust each other's actions [9], DLT could act as an effective cryptographically safe mediating system. Thus, not only could DLT largely remove the need for a central party (such as a bank), but it could also introduce considerable operational savings by allowing all participants to leverage off very difficult to compromise security measures protecting their data, as well as derive their various positions/stakes/transaction history without laborious and often complicated back-office reconciliation practices.

As such, DLT casts a wide net of application, across industries, businesses, and government organisations. Beyond cryptocurrency like Bitcoin, Ether, Litecoin etc., here are some specific examples of DLT usage:

Financial Services Clearing and Settlement

More efficient settlements (ASX CHESS Replacement)[10] – considered to be the first major market infrastructure project that aims to use blockchain based technology to improve the efficiency of equity clearing and settlement operations. The new system should also greatly reduce the overall reconciliation works by market participants, as well as allow the ASX to possibly introduce new value added services that give participants faster settlement options (better than the normal 'T+2' days of settlement) or options to use the cash or equity during the settlement cycle to generate further income.

Capital Markets

Initial Coin Offerings (ICOs) by April 2018 have raised $US6.3Bn, which has surpassed the total raised during 2017 [11]. Though some controversy surrounds ICOs (Regulators are still ramping up efforts to ensure clearer governance), no doubt that as cryptocurrency becomes more prevalent in the public domain, the more ICOs would move to be an accepted form of capital raise.

Insurance

Insurance firms have started to look into DLT as a means to operationally streamline key processes such as Insurance Contract Lifecycle Management, Claims Management and KYC/ AML [12].

Government

Electoral voting systems are a natural fit for DLT in as much as any polling centre (online or physical) could be a node in a blockchain that directly inserts votes in near real-time. Reconciliation and any recount would be a far less arduous task than paper based methods currently employed [13]

Urban Development

Smart Cities – The provisioning of public and private services and mining of valuable data collected from their interconnections could be greatly facilitated by DLT [14]. One key example being enhanced urban planning that is better in line with neighborhood needs (public transport, local business license approval as examples).

Healthcare

Patient Records – a service based on DLT could digitize, encrypt, and make available upon permission of the patient, medical records to a new hospital or healthcare provider.[15]

References

[1] Ethereum Foundation, 2018 (https://www.ethereum.org/foundation)

[2] NEM Foundation, 2018 (https://nem.io/about/foundation/)

[3] EOS.IO Technical White Paper v2, 16 March 2018 (https://github.com/EOSIO/Documentation/blob/master/TechnicalWhitePaper.md#governance)

[4] IOTA Foundation, 2018 (https://www.iota.org/the-foundation/the-iota-foundation)

[5] Hedera Hashgraph Council (https://www.hederahashgraph.com/council)

[6] Hedera: A Governing Council & Public Hashgraph Network Whitepaper v1.1, 18 May 2018 (https://s3.amazonaws.com/hedera-hashgraph/hh-whitepaper-v1.1-180518.pdf)

[7] Bitcoin Wednesday "Byteball founder Tony Churyumoff talks about his DAG-based ledger 'Byteball", February 2018 (https://www.youtube.com/watch?v=ZpCfucpwwQc)

[8] Nano Organization (https://nano.org/en)

[9] Chartered Accountants Australia New Zealand, "The Future of Blockchain: Applications and Implications of Distributed Ledger Technology", January 2017

[10] Australian Financial Review, "ASX blockchain to go live at end of 2020", 27 April 2018 (https://www.afr.com/technology/asx-blockchain-to-go-live-at-end-of-2020-20180427-h0zcgx)

[11] CoinDesk, "$6.3 Billion: 2018 ICO Funding Has Passed 2017's Total", 19 April 2018 (https://www.coindesk.com/6-3-billion-2018-ico-funding-already-outpaced-2017/)

[12] A Long Finance report prepared by Z/Yen Group, "Chain Reaction: How Blockchain Technology Might Transform Wholesale Insurance", July 2016

[13] PwC Australia, "Advantage Blockchain: Four ways blockchain will shake up business as usual", 22 June 2017

[14] FICCI and PwC India, "Blockchain: The next innovation to make our cities smarter", 2018

[15] PwC US article, "A look at Blockchain Technology", 2017

[16] Forbes, "Explaining Directed Acylic Graph (DAG), The Real Blockchain 3.0", 22 January 2018 (https://www.forbes.com/sites/shermanlee/2018/01/22/explaining-directed-acylic-graph-dag-the-real-blockchain-3-0/#3f016f88180b)

[17] VISA Factsheet, September 2017 (https://usa.visa.com/dam/VCOM/download/corporate/media/visanet-technology/aboutvisafactsheet.pdf)

[18] Harvard Business Review, "The Truth about Blockchain", January 2017 (https://hbr.org/2017/01/the-truth-about-blockchain)