# Distributed Ledger Technology

Assembled by Stephen Downes

| Key Concepts | Underlying Concepts<br>Core Technologies<br>Defining Blockchain<br>Key Concepts<br>Blockchain for Business<br>Applications<br>Public Sector Applications<br>Identity<br>Adoption Strategy<br>Related Technologies<br>Hardware<br>Advanced Concepts<br>Coins<br>Exchanges<br>Platforms<br>Decentralized Storage<br>Distributed Applications<br>Distributed Organizations<br>Activity<br>Issues<br>Future? |
|---|---|

## Distributed Ledger Technology

https://hackernoon.com/an-overview-of-hashgraph-b0900a1fd7bf
"A distributed ledger technology (DLT) is a system where we share information and we don't trust each other individually, but we trust the group as a whole. DLTs allow us to come up with a consensus on the order of transactions and timestamps."

Distributed ledger technology (DLT) - How Blockchain will transform the asset management industry - https://igniteoutsourcing.com/publications/blockchain-asset-management/

- improved data security.
- near-real-time performance in asset tracking
- efficiency primarily comes from the blockchain structure.
- Onboarding - working together, stakeholders can build infrastructures that facilitate the fast sharing of information needed to reduce the onboarding process
- accelerate communications between investors, asset managers, and third-party entities.

- DLT can eliminate the difference between a trade and the clearing of the transaction
- Compliance - Since each blockchain transaction is verified by all vested parties, meeting compliance demands becomes easier.

Distributed ledger diagram, p.7 -
http://www.ey.com/Publication/vwLUAssets/Blockchain_in_wealth_and_asset_management/$FILE/ey-blockchain-innovation-wealth-asset-management.pdf

"R3 Director: 2017 is the Year of the DLT Pilot" -
https://www.coindesk.com/r3-director-2017-year-dlt-pilot/

https://www.oreilly.com/ideas/the-blockchain-beyond-bitcoin
"Shared distributed ledger

This is really as simple as it sounds. The key detail of the ledger is that you cannot go back and change a single item without having to rewrite the entire ledger. This proves useful when dealing with regulatory bodies, as the amount of work required to falsify information is immense—and with proper controls in place, it becomes essentially impossible. The distributed capability is critical to business processes in that it ensures high availability and redundancy for cases such as disaster recovery. The final piece is the shared aspect of the ledger. Consider a two-party agreement and a notary service. If each party has a copy of the contract, either could tamper with it and claim their copy is the correct copy, but with a notary having a third-party copy which cannot be changed, suddenly we have created a case of irrevocable proof. This is a fundamental feature of a blockchain."

Diagrams: https://www.bis.org/publ/arpdf/ar2018e5.pdf
- centralized, distributed (permissioned) and distributed (permissionless) ledgers, p. 96 -
- Valid transactions in different systems p. 97

Distributed Ledger Foundation
https://www.distributedledgerfoundation.org/
- White Paper - https://static1.squarespace.com/static/5aa0b9eaec4eb7958e9f8f4b/t/5aa8301ee2c4839970f4ff40/1520971812762/dlf-whitepaper-2.pdf

# Consensus

SoK: Consensus in the age of blockchains
https://blog.acolyer.org/2018/02/12/sok-consensus-in-the-age-of-blockchains/
Diagram -
https://adriancolyer.files.wordpress.com/2018/02/blockchain-consensus-overview.jpeg?w=566&zoom=2

“Because of the distributed nature of the blockchain database, data about all new transactions must be propagated to all nodes on the network so that the blockchain stays in sync as one “world wide ledger,” and not as many conflicting ledgers. That means that in order to update the blockchain, these multiple, distributed copies of it must be reconciled so that they all contain the same version. This happens in the blockchain via a consensus process: the majority of the nodes in the system must concur.

“his consensus process is one of the key innovations of the blockchain: it is “emergent,” rather than happening at a scheduled time or interval as each new transaction and block is verified computationally.”
https://hackeducation.com/2016/04/07/blockchain-education-guide

“Whenever a new transaction gets broadcasted to the network, nodes have the option to include that transaction to their copy of their ledger or to ignore it. When the majority of the actors which comprise the network decide on a single state, **consensus** is achieved.”
https://medium.com/loom-network/understanding-blockchain-fundamentals-part-1-byzantine-fault-tolerance-245f46fe8419

https://www.oreilly.com/ideas/the-blockchain-beyond-bitcoin
“The concept of consensus
Blockchain implementations rely heavily on the concept of consensus, for this is the determining factor for who can write to the blockchain. For bitcoin, this must be done in a distributed manner so no single person can own the entire blockchain. Within the enterprise, consensus may look a lot like voting or a request for approval and a sign-off approving said request. It could also be a group vote where a quorum is required.”

For Bitcoin:
https://www.federalreserve.gov/econresdata/feds/2014/files/2014104pap.pdf

- The main rule for determining the legitimate ledger is that it is the one that took the most cumulative work to generate.
- consider the special case where two miners encode two different blocks and broadcast them nearly simultaneously, and assume that both blocks are with the same difficulty:
    - to resolve this issue, since a common ledger is the goal, the process waits for the following block to be added to either block C.1 or C.2 with two different parts of the network working on each.
    - If C2 is in the longer chain, accept C2 mark block C.1 as an orphan block.

## Digression: Information Theory

Must know Information Theory concepts in Deep Learning (AI)
https://towardsdatascience.com/must-know-information-theory-concepts-in-deep-learning-ai-e54a5da9769d
“Another way to look at entropy is the average information gained when we observe outcomes of an random experiment. The information gained for a outcome of an experiment is defined as a function

of probability of occurrence of that outcome. More the rarer is the outcome, more is the information gained from observing it.”

## Synchronizing Distributed Databases

---

https://stackoverflow.com/questions/1473014/keeping-distributed-databases-synchronized-in-a-unstable-network

Yahoo's PNUTS system: http://research.yahoo.com/node/2304 and Amazon's Dynamo: http://www.allthingsdistributed.com/2007/10/amazons_dynamo.html

these research papers are relevant (as an example of this research field):

1. Distributed disconnected databases,
    - The dangers of replication and a solution,
    - Improving Data Consistency in Mobile Computing Using Isolation-Only Transactions,
    - Dealing with Server Corruption in Weakly Consistent, Replicated Data Systems,
    - Rumor: Mobile Data Access Through Optimistic Peer-to-Peer Replication,
    - The Case for Non-transparent Replication: Examples from Bayou,
    - Bayou: replicated database services for world-wide applications,
    - Managing update conflicts in Bayou, a weakly connected replicated storage system,
    - Two-level client caching and disconnected operation of notebook computers in distributed systems,
    - Replicated document management in a group communication system,

SymmetricDS. SymmetricDS is web-enabled, database independent, data synchronization/replication software.

---

## The Double Spending Problem

The technical challenge: the Double Spending problem
“any digital form of money is easily replicable and can thus be fraudulently spent more than once. Digital information can be reproduced more easily than physical banknotes. For digital money, solving the double-spending problem requires, at a minimum, that someone keep a record of all transactions. Prior to cryptocurrencies, the only solution was to have a centralised agent do this and verify all transactions.” https://www.bis.org/publ/arpdf/ar2018e5.pdf

“in the words of the original Bitcoin white paper, a cryptocurrency can overcome the double-spending problem in a decentralised way only if “honest nodes control a majority of [computing] power”.” https://www.bis.org/publ/arpdf/ar2018e5.pdf

## The Distributed Ledger Protocol

the key feature of these cryptocurrencies is the implementation of a set of rules (the protocol) that aim to align the incentives of all participants so as to create a reliable payment technology without a central trusted Agent. https://www.bis.org/publ/arpdf/ar2018e5.pdf

- First, the rules entail a cost to updating the ledger. In most cases, this cost comes about because updating requires a “proof-of-work”.
- Second, all miners and users of a cryptocurrency verify all ledger updates, which induces miners to include only valid transactions
- Third, the protocol specifies rules to achieve a consensus on the order of updates to the ledger.

## Byzantine Fault Tolerance

- Synchronous vs asynchronous BFT
-

Nice account of fault tolerance - Two Generals Problem and Byzantine Generals Problem

- What happens when an actor decides to not follow the rules and to tamper with the state of his ledger?
- What happens when these actors are a large part of the network, but not the majority? https://medium.com/loom-network/understanding-blockchain-fundamentals-part-1-byzantine-fault-tolerance-245f46fe8419

“In the absence of BFT, a peer is able to transmit and post false transactions effectively nullifying the blockchain’s reliability. “
https://medium.com/loom-network/understanding-blockchain-fundamentals-part-1-byzantine-fault-tolerance-245f46fe8419

Nakamoto email - 2008 - email -
https://www.mail-archive.com/cryptography@metzdowd.com/msg09997.html

“how can individual computers (*nodes*) in a system reliably communicate truths (in other words, events that have taken place on the network) to each other where a proportion of the nodes are malicious (*Byzantine*) and looking to disrupt the system. Or to put it another way: how can a group of computers agree on which transactions have correctly taken place and in which order?”https://medium.com/safenetwork/parsec-a-paradigm-shift-for-asynchronous-and-permissionless-consensus-e312d721f9d8

### Honey Badger

Github - https://github.com/amiller/HoneyBadgerBFT

- “Most fault tolerant protocols (including RAFT, PBFT, Zyzzyva, Q/U) don't guarantee good performance when there are Byzantine faults. Even the so-called "robust" BFT protocols (like UpRight, RBFT, Prime, Spinning, and Stellar) have various hard-coded timeout parameters, and can only guarantee performance when the network behaves approximately as expected - hence they are best suited to well-controlled settings like corporate data centers.
- “HoneyBadgerBFT is fault tolerance for the wild wild wide-area-network. HoneyBadger nodes can even stay hidden behind anonymizing relays like Tor, and the purely-asynchronous protocol will make progress at whatever rate the network supports.”

The Honey Badger of BFT Protocols
https://eprint.iacr.org/2016/199.pdf
“We present an alternative, HoneyBadgerBFT, the first practical asynchronous BFT protocol, which guarantees liveness without making any timing assumptions.”

### BEAT

BEAT: Asynchronous BFT Made Practical
https://www.csee.umbc.edu/~hbzhang/files/beat.pdf
Through a 92-instance, five-continent deployment of BEAT on Amazon EC2, we show that BEAT is efficient: roughly, all our BEAT instances significantly outperform, in terms of both latency and throughput, HoneyBadgerBFT, the most efficient asynchronous BFT known

BEAT: asynchronous BFT made practical
https://blog.acolyer.org/2018/11/26/beat-asynchronous-bft-made-practical/

## Proof of Work

https://hackeducation.com/2016/04/07/blockchain-education-guide

“New blocks are created by a process called “mining,” which validates new transactions and adds them to the chain. In Bitcoin, a new block is mined every 10 minutes (that rate is different for different cryptocurrencies’ blockchains). The miner (the machine) that mines the new block is rewarded financially – in the case of Bitcoin, the miner receives Bitcoin (currently 25 per block, but that figure will halve later this year), as well as a cut of the transaction fees for all transactions on the block.

“To mine new blocks, miners on the network compete to solve a unique, difficult math puzzle. As noted above, the “proof of work” of that solution is included in the block header which allows the block to be verified. Solving this math problem is nontrivial. ”

https://medium.com/loom-network/understanding-blockchain-fundamentals-part-2-proof-of-work-proof-of-stake-b6ae907c7edb

- Blockchains use consensus algorithms to elect a leader who will decide the contents of the next block.
- That leader is also responsible for broadcasting the block to the network, so that the other peers can verify the validity of its contents.
- in order for an actor to be elected as a leader and choose the next block to be added to the blockchain they have to find a solution to a particular mathematical problem.
- probabilistically speaking, the actor who will solve the aforementioned problem first the majority of the time is the one who has access to the most computing power. These actors are also called **miners**.
- Whenever a new block is *mined,* that *miner* gets rewarded with some currency (block reward, transaction fees) and thus are incentivized to keep mining. Due to the limited supply of computational power, miners are also incentivized not to cheat. -> see diagram: https://cdn-images-1.medium.com/max/1000/0*SxksLCUp1eh_ZuJt.png

From 2016: Blockchain.info estimates that Bitcoin miners are now trying 450 thousand trillion solutions per second to solve these puzzles. As such, in 2015, O’Reilly Media estimated that it takes about $600 million a year to maintain the mining infrastructure of the Bitcoin system.
https://hackeducation.com/2016/04/07/blockchain-education-guide

## Hash Function

https://medium.com/loom-network/understanding-blockchain-fundamentals-part-2-proof-of-work-proof-of-stake-b6ae907c7edb

- A **hash function** is any function that can be used to map data of arbitrary size to data of fixed size[1].
- If a hash function is secure, its output is indistinguishable from random.
- It has been widely successful primarily due to its following properties:
    a. It is hard to find a solution for that given problem
    b. When given a solution to that problem it is easy to verify that it is correct

# Proof of Stake (PoS)

https://medium.com/loom-network/understanding-blockchain-fundamentals-part-2-proof-of-work-proof-of-stake-b6ae907c7edb

- Proof of Stake takes away the energy and computational power requirement of PoW and replaces it with stake. Stake is referred to as an amount of currency that an actor is willing to lock up for a certain amount of time.
- In Proof of Stake, if Bob has more stake than Alice, he is more likely to win ("mine" the next block).
- There are various existing coins which use pure PoS, such as Nxt and Blackcoin.

Can we achieve the same level of security as a Proof-of-Work (PoW) system like Bitcoin while not depleting physical scarce resources to do it?

The Long Road to Proof-of-Stake -
https://blog.cosmos.network/consensus-compare-casper-vs-tendermint-6df154ad56ae

- Practical Byzantine Fault Tolerance (PBFT) - The breakthrough that Satoshi Nakamoto made in 2008 was in bringing internet-scale distributed Byzantine Fault Tolerant (BFT)consensus into a blockchain scheme.
- The breakthrough that Satoshi Nakamoto made in 2008 was in bringing internet-scale distributed Byzantine Fault Tolerant (BFT)consensus into a blockchain scheme.
- huge number of different consensus algorithms (Honeybadger, Ouroboros, Tezos, Casper) popped up that all incorporate elements of BFT research along with other modular observations on the blockchain.

https://arxiv.org/pdf/1710.09437.pdf
There are two major schools of thought in PoS design.

- The first, mimics proof of work mechanics and features a chain of blocks and and simulates mining by pseudorandomly assigning the right to create new blocks to stakeholders. This includes Peercoin[3], Blackcoin[4], and Iddo Bentov's work[5].
- The other school, Byzantine fault tolerant (BFT) based proof of stake, is based on a thirty-year-old body of research into BFT consensus algorithms such as PBFT[]. BFT algorithms typically have proven mathematical properties; for example, one can usually mathematically prove that as long as of protocol participants are following the protocol honestly, then, regardless of network latency, the algorithm cannot finalize conflicting blocks. Repurposing BFT algorithms for proof of stake was first introduced by Tendermint and has modern inspirations such as [8]. Casper follows this BFT tradition, though with some modifications

Detailed article comparoing PoS mopdels:
https://blog.cosmos.network/consensus-compare-casper-vs-tendermint-6df154ad56ae
*recommended*

## Tendermint

https://blog.cosmos.network/consensus-compare-casper-vs-tendermint-6df154ad56ae

Tendermint consists of two chief technical components: a blockchain consensus engine and a generic application interface.

- The application interface, called the Application Blockchain Interface (ABCI), enables the transactions to be processed in any programming language.
- Diagram - https://cdn-images-1.medium.com/max/1000/1*GE_R8CGb3p60iMKzJzp9OQ.png

- Tendermint expmained - https://blog.cosmos.network/tendermint-explained-bringing-bft-based-pos-to-the-public-blockchain-domain-f22e274a0fdb
- Tendermint bug bounty - https://blog.cosmos.network/launched-tendermint-bug-bounty-program-183479658ea1
- See also: Cosmos - https://blog.cosmos.network/understanding-the-value-proposition-of-cosmos-ecaef63350d

## Casper the Friendly Finality Gadget

What is Ethereum Casper Protocol? Crash Course
https://blockgeeks.com/guides/ethereum-casper/

- "They needed a protocol which could implement POS and mitigate the "Nothing at Stake" problem."

This is how POS under Casper would work:

- The validators stake a portion of their Ethers as stake.
- After that, they will start validating the blocks. Meaning, when they discover a block which they think can be added to the chain, they will validate it by placing a bet on it.
- If the block gets appended, then the validators will get a reward proportionate to their bets.
- However, if a validator acts in a malicious manner and tries to do a "nothing at stake", they will immediately be reprimanded, and all of their stake is going to get slashed.

- "As Hudson James and Joris Bontje note in their answers in "StackExchange," Casper designs harsher incentives to guarantee network security, including punishing miners who go offline, unintentionally or not."

Casper the Friendly Finality Gadget (CFFG) is a PoS overlay on top of the existing Ethereum PoW proposal mechanism—a hybrid PoW/PoS implementation led by Vitalik Buterin.
https://blog.cosmos.network/consensus-compare-casper-vs-tendermint-6df154ad56ae (see more below)

What is Holding Back The Blockchain?
https://medium.com/@abhishekkothari/what-worlds-may-come-8f00e631ca31
Two developers Vlad Zamfir and Vitalik Buterin are working on a version of Byzantine Fault Tolerant (BFT) Proof of Stake protocol called Casper. There are two teams working on developing a different versions of Casper

- Casper the Friendly Finality Gadget (FFG)

- Casper the Friendly GHOST: Correct-by-Construction (CBC)

The success of one or both of these versions holds the key to the problem of a scalable Blockchain.

### Casper the Friendly Ghost

- "Casper the Friendly Ghost (CTFG) is Vlad Zamfir's correct-by-construction (CBC) consensus protocol tailored toward combatting an oligopolistic real world environment. CTFG is a PoS adaptation of the Greedy Heaviest-Observed Subtree, or GHOST protocol in PoW, for its fork choice rule. The guiding design principle behind Casper the Friendly Ghost is based on cryptoeconomics using formal methods aimed to achieve estimate safety. CTFG is a pure Proof-of-Stake concept, unlike Casper the Friendly Finality Gadget's hybrid protocol detailed in the earlier section."
  https://blog.cosmos.network/consensus-compare-casper-vs-tendermint-6df154ad56ae

## Delegated Proof of Stake (DPoS)

https://medium.com/loom-network/understanding-blockchain-fundamentals-part-3-delegated-proof-of-stake-b385a6b92ef

"DPoS is a system in which a fixed number of elected entities (called *block producers* or *witnesses*) are selected to create blocks in a round-robin order. Block producers are voted into power by the users of the network, who each get a number of votes proportional to the number of tokens they own on the network (their *stake*)."

- ***Block producers*** are those responsible for creating and signing new blocks. They are limited in number, and are elected by the voters.
- ***Block validators*** in DPoS refer to full nodes who verify that the blocks created by block producers follow the consensus rules. Any user is able run a block validator and verify the network. (This can be confusing, since in Casper's PoS, the word "validators" refers to those who *create* blocks).

- A block is finalized (i.e. cannot be reversed) when it is voted on by (2/3+1) of the block producers. Otherwise, the longest chain rule is followed.
- Voters can also "fire" a block producer if they are found to be malicious (i.e. try to censor transactions or double spend) by not voting for them in the next round.

number of block producers in a DPoS network is up to the consensus rules of that chain.. Numbers for some of the more well-known DPoS chains:

- **EOS:** 21
- **BitShares:** 101
- **Steemit:** 21
- **Lisk:** 101
- **Ark:** 51

https://medium.com/loom-network/understanding-blockchain-fundamentals-part-3-delegated-proof-of-stake-b385a6b92ef

“we (and most other serious blockchain developers) believe that DPoS is not decentralized enough to be a base layer that acts as a store of value and ledger of ownership for Web 3 applications.”

Option:

- Even better—if we use a PoW network like Ethereum as a secure base layer and build our DPoS chain on Layer 2, we can run the majority of an application on the highly-scalable DPoS chain, while still using the secure base layer for parts of the application that require high security, like in-game currency or ownership of assets.
  https://medium.com/loom-network/understanding-blockchain-fundamentals-part-3-delegated-proof-of-stake-b385a6b92ef

### Plasma

- Plasma is a technique that enables the secure transfer of assets through blockchains, in which the security of the system is guaranteed by the “Root Chain”, even if the “Plasma Chain” becomes byzantine.
  https://medium.com/loom-network/understanding-blockchain-fundamentals-part-3-delegated-proof-of-stake-b385a6b92ef
- Loom Network, we are moving along with our Plasma Cash implementation for ERC721 tokens, set to release later this month, which will allow game developers to deposit their ERC721 tokens on Loom DAppChains and use them in games.

## Crypto Asset Classes

https://hackernoon.com/crypto-asset-classes-6dd6ddece456

“Tom Lee break the collections into cohorts based on how they trade—more like sectors. … commodities, platforms, privacy, exchanges and stablecoins.”

“The general partners at Multicoin Capital, break the collections into 3 major cohorts: currencies (stores of value); security tokens (tokens backed by real-world assets); and utility tokens (work tokens).”

“I see eight distinct crypto asset classes—core/reserve, currencies, platforms, utility tokens, security tokens, commodities, appcoins and stablecoins.”
https://hackernoon.com/crypto-asset-classes-6dd6ddece456

## Permissionless

Like money - anyone can join

Before bitcoin, there was no permissionless digital transaction system

Compare Internet - Minitel

- Also - X.25 standard - https://en.wikipedia.org/wiki/X.25 - “The betamax of internet working standards” (Jerry Brito - https://www.youtube.com/watch?v=R0iArSIU0Z8&feature=youtu.be&t=47m16s 1:32:00 )

Internet worked because it was permissionless

## Temporal

Third, Allan; Tiddi, Ilaria; Bastianelli, Emanuele; Valentine, Chris and Domingue, John (2017). Towards the Temporal Streaming of Graph Data on Distributed Ledgers. Lecture Notes in Computer Science, 10577 pp. 327–332. http://oro.open.ac.uk/52930/ “We describe a system in which temporal annotations, and information suitable to validate a given dataset, are stored on a distributed ledger, alongside the results of fixed SPARQL queries executed at the time of data storage. The model adopted implements a graph-based form of temporal RDF, in which time intervals are represented by named graphs corresponding to ledger entries. “