CHAPTER 2

Blockchain

Blockchain technology presents an innovative take on the traditional distributed database. The utilization of existing technology in unconventional settings is the source of the innovation. When we think of a traditional database like Postgres, there is a master node that accepts the write requests and is responsible for synchronizing the state changes to other Postgres nodes. In a decentralized setup like blockchain, there is no master node, but still it serves the same purpose of updating the ledger and making sure that other nodes in the network reflect the updated state correctly.

There are a great number of distinct varieties of blockchains and applications for blockchain technology. The blockchain is an all-encompassing technology that is currently being integrated across a variety of different platforms and pieces of hardware around the world.

A blockchain is a data structure that enables the creation of a digital ledger of data and the sharing of that data across a network of parties that are not affiliated with one another. There are a great number of distinct varieties of blockchains.

2.1 Types of Blockchains

Blockchains can be classified under three broader categories.

S. M. Jain, *A Brief Introduction to Web3*, https://doi.org/10.1007/978-1-4842-8975-4_2

2.1.1 Public Blockchain

Blockchains that are open to the public, like Bitcoin, are large distributed networks that are each managed by their own native token. They have open-source code that is maintained by their community and welcome participation from anyone, regardless of level of involvement.

2.1.2 Private Blockchain

Private blockchains are typically smaller than public ones, and they do not make use of tokens. Their membership is strictly monitored and regulated. Consortiums that comprise reliable members and engage in the confidential exchange of information favor the use of blockchains of this kind. An example of this would be a corporation using a blockchain in its own data center. IBM has an implementation called Hyperledger that is implemented by several corporations.

2.1.3 Permissioned Blockchain

A permissioned blockchain is one in which a user needs permission to access it. This is very different from public and private blockchains. A corporate setup might look for a permissioned blockchain where it uses the blockchain more like a decentralized database within the confines of the corporate boundaries. Examples of permissioned blockchains are Ripple and IBM Food Trust.

The reasons blockchains rose to prominence are as follows:

1. Their decentralized nature – The blockchain setup advocates decentralization in its architecture. This means that in a strong blockchain network like Bitcoin or Ethereum, there are no centralized authorities. Anyone can participate in these

networks by spinning their own node and becoming part of the blockchain network. This is really an empowering feature.

2. Immutable data structure for storage – All data that is stored on the blockchain is permanently recorded across a fleet of nodes/computers across the globe. One can think of this as a permanent ledger of records. Removal of a record from this ledger needs a majority of the network to agree, which is a practical impossibility if the network is big like Bitcoin or Ethereum.

Though the rules vary from blockchain to blockchain, generally any mutation or state chain on the blockchain is done by one entity, which then publishes the changes on the network for other nodes to verify. Only when a majority agrees to that state is the transaction considered valid. The mechanism for who gets the right to change the state can vary from blockchain to blockchain based on the underlying algorithm of the chain. As an example, for a Bitcoin-based blockchain there is the proof-of-work algorithm, which lets the nodes participate in solving a mathematical puzzle, and the node that solves it first gets the right to write on the ledger. The writer then broadcasts the changes to the network for the validator to validate the changes.

2.2 What Is a Blockchain?

A blockchain is a peer-to-peer network that does not rely on a centralized authority to manage the state of the system or ledger. The computers that form the network can be distributed physically. “Full nodes” is another common name for these kinds of computers.

After data has been entered into a blockchain database, it is extremely difficult, if not impossible, to delete or alter that data. This means that once the data is recorded and replicated across tons of nodes across the world, the consumer of the data can have a higher level of trust in the data as compared to data recorded on a centralized ledger. This means that records like property rights, marriage certificates, invoices, and so on can be recorded on the chain for permanence without the worry of tampering. Processes that rely on having a central system in business and banking, such as fund settlements and money wires, can now be completed without having a central authority in place. The implications of having secure digital records are extremely significant for the economy of the entire world.

2.3 Blockchain Building Blocks

There are three main components of a blockchain ecosystem.

2.3.1 Block

A block is a listing of the transactions that have been entered into a ledger over the course of a specific time period. Each blockchain has its own unique dimensions, time intervals, and events that cause blocks to be created.

There are some blockchains whose primary focus is not on maintaining and securing a record of the transactions involving their respective cryptocurrencies. However, every blockchain will record any transactions involving its associated cryptocurrency or token. Consider the transaction to be nothing more than the simple recording of data. The interpretation of what that data means is accomplished by first giving it a value, as is done in the course of a financial transaction.

2.3.2 Chain

A chain can be thought of as a linked list of blocks. The block that comes after the first block (also known as genesis block) will maintain a reference to the previous block and so on and so forth. This reference is maintained by generating a hash of the contents of the previous block. This has some great implications for the immutability of the chain. Anyone who has to alter a previous block will then need to modify all blocks ahead of that particular block. Generally this is a practical impossibility as it might require either huge computing resources in the case of algorithms like proof of work, or a huge stake to be put in the network for proof of stake–based algorithms like Ethereum.

2.3.3 Network

"Full nodes" are the building blocks of the network. Imagine them as a computer that is actively implementing an algorithm to ensure the safety of the network. All nodes in the network have the same copy of records (called the ledger). The ledger comprises all transactions that have ever been recorded on the blockchain.

The network nodes are decentralized and can be run by anyone in any part of the world. Since running these nodes will incur costs, people need some incentive to do so. They are motivated to run a node because they want to earn cryptocurrency, which serves as their incentive. The underlying algorithm of the blockchain provides them with compensation for their service. Typically, a token or cryptocurrency, such as Bitcoin or ether, is awarded as the reward.

2.4 Where Is Blockchain Used?

Applications based on the blockchain are designed around the principle that the network should act as the arbitrator. This kind of system creates an unforgiving and oblivious atmosphere for its users. The computer code effectively becomes the law, and the rules are carried out in the manner in which they were written and are interpreted by the network. The social behaviors and prejudices that are characteristic of humans are absent in computers.

The network is incapable of interpreting the user's intent (at least not yet). As a potential use case for this concept, insurance contracts that are arbitrated on a blockchain have received a lot of attention and research. The ability to keep immutable records is yet another fascinating application of blockchain technology.

You can use them to make a comprehensible timeline of who did what and when. Countless hours are spent by numerous industries and regulatory bodies trying to get a handle on the magnitude of this issue. An immutable ledger will allow us to keep a consistent timeline of transactions that have happened right from the inception of the chain.

Although people frequently use the terms "Bitcoin" and "blockchain" interchangeably, these two concepts are not the same. A blockchain is used by Bitcoin. The underlying protocol that allows for the safe transfer of Bitcoin is known as the blockchain. Bitcoin is the name of the digital currency that serves as the driving force behind the Bitcoin network. So Bitcoin is one of the applications built on top of the blockchain infrastructure.

2.5 Evolution

Blockchain technology was initially developed in conjunction with the cryptocurrency Bitcoin. The whole idea of Bitcoin was to develop a system for decentralized money and banking. The initial purpose of the Bitcoin network's construction was to ensure the cryptocurrency's safety. It has somewhere in the neighborhood of 5,000 full nodes and is spread out across the world. Its primary purpose is to buy and sell Bitcoin and other cryptocurrencies, but the community quickly realized that it could be used for a great deal more than that. Because of its size and the fact that its security has been proven over time, it is also being used to secure other blockchains and blockchain applications that are smaller in scale.

The blockchain technology has undergone a second iteration with the creation of the Ethereum network. It modifies the conventional structure of a blockchain by incorporating a programming language that is embedded within the structure itself. Similar to Bitcoin, it has more than 5,000 full nodes and is distributed across the world. Ether trading, the development of smart contracts, and the formation of decentralized autonomous organizations (DAOs) are the primary uses of Ethereum. Additionally, it is utilized in the process of securing blockchain applications as well as more compact blockchains.

The Ethereum ecosystem is what turned the blockchain ecosystem into a programmable blockchain.

2.6 Consensus

Consensus is one of the most difficult problems to solve in a distributed systems setup.

Blockchains are useful tools because they are capable of self-correcting and do not require the intervention of a third party in order to do so. Through the use of their consensus algorithm, they are able to successfully enforce the rules.

In the world of blockchain technology, “reaching consensus” refers to the procedure of reaching an agreement among a group of nodes that typically do not trust one another. These are the nodes on the network that have full functionality. Transactions that are entered into the network in order for them to be recorded on the ledger are being checked for validity by the full nodes. Each blockchain has its own set of algorithms that it uses to get its network to reach a consensus on the new entries that are being added. Because different kinds of entries are being generated by each blockchain, there is a diverse range of models available for achieving consensus.

Some of the mechanisms used by different blockchains for consensus are as follows:

- Proof of Work – Used by Bitcoin blockchain
- Proof of Stake – Used by Ethereum
- Proof of History – Used by Solana

2.6.1 Proof of Work

The first decentralized cryptocurrency to use a consensus mechanism was Bitcoin, and that mechanism was proof of work. Mining and proof of work are concepts that are closely related to one another. The term “proof of work” comes from the fact that the network needs an extremely high level of processing power to function properly. Proof-of-work blockchains are those in which virtual miners from all over the world compete with one another to see who can solve a mathematical puzzle first. These blockchains are then secured and verified. The victor receives a predetermined amount of cryptocurrency from the network as well as the opportunity to update the blockchain with the most recent transactions that have been verified.

Proof of work offers a number of significant benefits, particularly for the case of Bitcoin, which is a relatively straightforward but extremely valuable cryptocurrency. It is a tried-and-trusted method that can reliably keep a decentralized blockchain in a secure state. As the price of a cryptocurrency continues to rise, more miners will be encouraged to participate in the network, which will result in an increase in both the network's power and its level of security. The Bitcoin network also in its algorithm has a difficulty adjustment mechanism, which is like a feedback loop to alter the difficulty of mining based on demand. If there are huge number of miners and they are able to mine quickly because they have good resources at hand, the algorithm increases the difficulty level of the mining to keep the mining time of one block to 10 minutes.

The mathematical puzzle used by a proof-of-work algorithm is not truly a puzzle in the sense of a physical puzzle. It uses the one-way function of cryptographic hashes as the underlying mechanism. Cryptographic hashes are designed to be deterministic one-way functions. What this means is that any content, be it text of few lines or document or images or videos can be reduced to a specific size. For instance, in a SHA 256–based hash function, the size of the output is 256 bits, no matter how big or small the input is. The other property of hash functions is that the same content will generate the same hash no matter how many times you run the same hash function over it. This is where the determinism comes from.

As I said, it's a one-way function, so going from input to output is easy. Going back from output to input is impossible, or one could say computationally infeasible. Bitcoin uses this property for its proof-of-work algorithm.

So basically the Bitcoin network provides a hash value as the input and asks the miners to generate a hash by taking the following two inputs:

- The block header
- A random number called nonce

The miner needs to add these two together and generate a hash that is less than the value provided by the Bitcoin network for a specific block. As we discussed, since hashing is a one-way function, the miner now starts to add different values of nonce to the header and has to effectively generate the hashes to make it less than the target hash provided by the system. There is no other way to do this except by crunching these hashes.

2.6.2 Proof of Stake

Ethereum blockchain started with proof of work as the algorithm but now is moving to a proof-of-stake algorithm. In proof of stake, to mine the ether (that's the currency on the Ethereum blockchain) the network participant has to stake the currency. Say, for example, I want to be a validator on the network. I then have to put some money at stake to validate the transactions. This puts my skin in the game, and I am disincentivized to cheat. So, instead of putting in energy, as in the case of Bitcoin, Ethereum advocates putting money itself at stake for being a participant on the network.

From the pool of stakers, the one who puts the maximum stake will be chosen as the validator of transactions and will be rewarded for the task.

Other validators will be able to do a proof of validity of the block of transactions once the winner has validated the most recent block of transactions. The blockchain is updated whenever the network reaches a certain predetermined number of confirmations or attestations.

One of the most significant distinctions between the two consensus mechanisms is the amount of energy required. Proof-of-stake blockchains enable networks to function with significantly lower resource consumption than blockchains based on proof of work. The debate still rages as to which is a better mechanism as far as representation of money is concerned.

Both of these consensus mechanisms have economic repercussions that punish malicious actors for disrupting the network and discourage others from doing so. The sunk cost of computing power, energy, and time is the punishment for miners who submit invalid information or blocks in proof-of-work systems. This is the case for Bitcoin and other cryptocurrencies.

The main idea is to create a heavy penalty for cheating the system. In the case of proof of work, it's the energy spent that is at stake, whereas in proof of stake it is the money that is put at stake. If, say, a network participant is found to have accepted a corrupt block, a portion of the funds that they have staked will be "slashed" as a form of punishment. The network will determine the maximum amount by which a validator's reward can be reduced.

2.7 Blockchain Architecture

A client–server network is used in the conventional design of the World Wide Web's architecture. Due to the fact that it is a centralized database that is controlled by a number of administrators who each have permission to make changes, the server in this scenario stores all of the necessary information in a single location so that it can be easily updated.

When it comes to the distributed network architecture of blockchain technology, every participant in the network is responsible for the maintenance, approval, and updating of new entries. Not only are there a number of different people who have control over the system, but everyone who is part of the blockchain network does. Every member is responsible for ensuring that all of the records and procedures are in order, which ultimately leads to the validity and safety of the data. Therefore, it is possible for parties that do not necessarily trust one another to come to an agreement with one another.

In a nutshell, the blockchain can be described as a decentralized distributed ledger of various transactions that is organized into a P2P network. This ledger can be made public or kept private. This network is made up of a large number of computers, but it is constructed in such a way that the data cannot be changed unless there is agreement from all of the computers in the network (each separate computer).

A list of blocks containing transactions arranged in a specific sequence is what the blockchain technology uses to represent its underlying structure. These lists may be kept in the form of a straightforward database or as a simple text file (using the txt format). One can think of the blockchain as a persisted linked list of transaction blocks.

We have already covered the concepts of nodes, blocks, and chains in the context of blockchains. There is one more important entity that needs a mention. This entity is the miner. As an example of the Bitcoin blockchain, the miner is the entity that is responsible for mutating the state of the blockchain. This mutation is achieved by adding a new block to the chain. To get permission to mutate the state, the miner has to solve a mathematical puzzle. The puzzle is trying to find the number for a system-provided hash. So the miner has to run through all different combinations of numbers to see if the hash for one of them matches. This computation is run by many miners across the globe and whoever wins the race gets the permission to write the new block. Once the miner writes the block they publish this on the network for the verifiers to verify the block. The verification process is simple as the miner will provide the hash as well as the number for which this hash was generated. So verification is to just calculate the hash for the number and compare it with the hash the system has asked for. Once a majority of miners has verified the block, this block gets replicated across the network and gets added to all computers on the network gradually.

At a high level, the transaction flow in a Bitcoin blockchain works like the one shown in Figure 2-1.

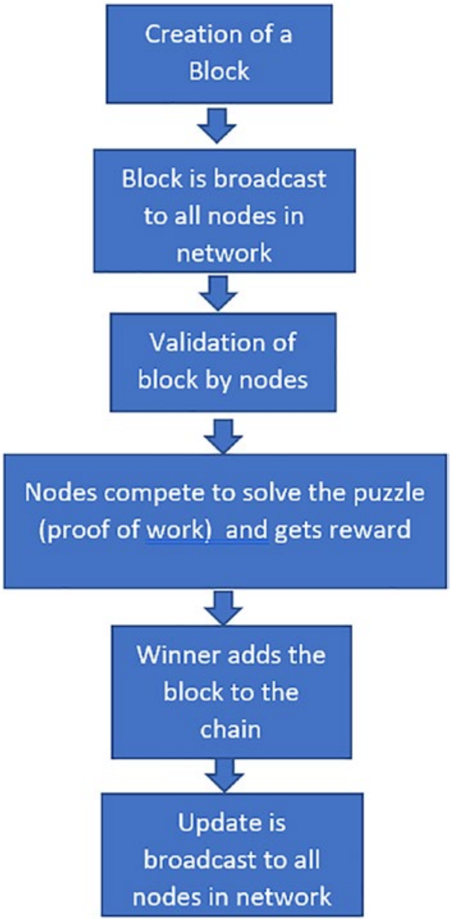
Creation of a Block
Block is broadcast to all nodes in network
Validation of block by nodes
Nodes compete to solve the puzzle (proof of work) and gets reward
Winner adds the block to the chain
Update is broadcast to all nodes in network

Figure 2-1. *How a transaction works on Bitcoin blockchain*

2.8 Cryptographic Keys

Since public blockchain networks are permissionless and trustless, we need a mechanism to authenticate the users in absence of a centralized authority. This task is achieved by providing each participant on the network with a private and public key pair. The private key remains with the user and should be kept highly protected. The private key is used to sign the transactions of the user, whereas the public key (which is visible to everyone) can be used to validate the signer. In this way, without a central authority and using the private and public key infrastructure, security is achieved on the decentralized blockchain network.

2.9 Blockchain Compared to a Singly Linked List

There is a certain degree of similarity between the data structure of a blockchain and that of a singly linked list, as shown in Figure 2-2.

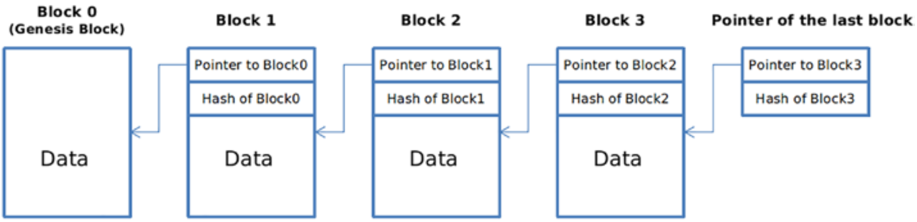

Figure 2-2. *Linked list–like structure depicted for the blockchain*

We can see the first block is called the genesis block, and then from there each new block created holds a reference to the previous block. As an example, block 1 will hold a reference to block 0 and block 2 to block 1 and so on. Apart from this pointer, the hash of the previous block is stored. This appears like a singly linked list data structure. To add to the list we need to use either proof of work or proof of stake, but deletion is extremely difficult on the blockchain, unlike a linked list, where it's easy to remove nodes.

In essence, a blockchain is immutable. No record can be deleted. If we try to think about it, any alteration to any block means all the blocks ahead of that block have to be updated with the hash of previous blocks.

In the case of networks like Bitcoin, this would mean needing immense computing power to use proof of work to alter the blocks. That's one of the reasons it's almost impossible to attack the Bitcoin blockchain network.

Next, we are going to introduce the Ethereum blockchain, which is the programmable blockchain we will use for all our future work in this book.

2.10 Ethereum

Ethereum is a blockchain network that includes a Turing-complete programming language that can be used to build a variety of decentralized applications (also called DApps). The Ethereum network is powered by its own cryptocurrency, ether. The Ethereum network is currently well known for enabling the use of smart contracts. Smart contracts can be compared to cryptographic bank lockers that contain specific values. Certain conditions must be met before these cryptographic lockers can be unlocked. Solidity, a programming language, is primarily used to create smart contracts. Solidity is a relatively simple-to-learn object-oriented programming language. We will see more of Solidity in Chapter 3.

Ethereum works on two types of accounts:

1. **Externally owned accounts (EOA)**

   Private keys are used to control externally owned accounts. Each EOA is protected by a public–private key pair. Users can communicate by creating and signing transactions.

2. **Contract accounts**

   Contract codes are used to manage contract accounts. These codes are saved alongside the account. Each contract account has an associated ether balance. These accounts' contract codes are activated whenever they receive a transaction from an EOA or a message from another contract. When the contract code is enabled, it is possible to read/write messages to local storage, send messages, and create contracts.

2.11 Summary

In this chapter, we covered the basics of blockchain architecture and the different consensus algorithms like proof of work and proof of stake that are commonly used in blockchain networks like Bitcoin and Ethereum. In the next chapter, we will cover Solidity, which is the language used on the Ethereum blockchain, and how Solidity is used for creating smart contracts.