Chapter 8
Distributed Ledger Technology

Xing Liu, Bahar Farahani, and Farshad Firouzi

The chains of habit are too weak to be felt until they are too strong to be broken.

Samuel Johnson

Contents

X. Liu (✉)
Kwantlen Polytechnic University, Surrey, BC, Canada
e-mail: xing.liu@kpu.ca

B. Farahani
Shahid Beheshti University, Tehran, Iran

F. Firouzi
Department of ECE, Duke University, Durham, NC, USA

F. Firouzi et al. (eds.), *Intelligent Internet of Things*,
https://doi.org/10.1007/978-3-030-30367-9_8

8.1 Introduction to Distributed Ledger Technology and IoT

8.1.1 What Is a Distributed Ledger?

Distributed ledger technology is a general term that is used to describe technologies for the storage, distribution, and exchange of data between users over private or public distributed computer networks. Essentially, a distributed ledger is a database that is spread and stored over multiple computers located at physically different locations. Each of such computers is frequently referred to as a node. A distributed ledger can also be considered as a common datasheet stored on multiple distributed nodes.

Figure 8.1 shows a centralized, a decentralized, and a distributed system. Distributed ledger technology is based on distributed systems.

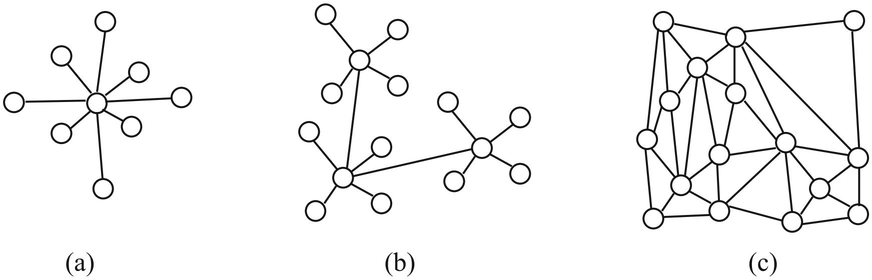

Fig. 8.1 (**a**) Centralized, (**b**) decentralized, and (**c**) distributed system

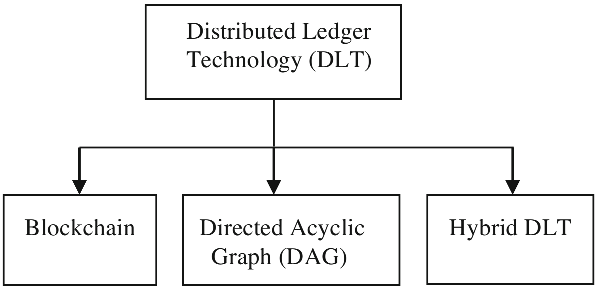
Distributed Ledger Technology (DLT)
Blockchain
Directed Acyclic Graph (DAG)
Hybrid DLT

Fig. 8.2 Three different implementations of distributed ledger technology

Distributed ledger technologies can also be classified into three categories based on the way the technology is implemented. These include blockchain, DAG, and hybrid DLT, as shown in Fig. 8.2.

8.1.2 Blockchain

Blockchain is the underlying technology of Bitcoin. It has become widely known since 2008 because of the publication of a paper titled *Bitcoin: A Peer to Peer Electronic Cash System* [1] which was authored by perhaps a group of people pseudonymously named Satoshi Nakamoto.

Although it was not until 2008 that blockchain became well known to the world, research about digital documents can be dated back to 1991 when Haber and Stornetta published their research paper titled *How to Time-Stamp a Digital Document* [2]. The paper discussed how to use a computer server to timestamp and link digital documents as a chain with pointers attached to the data in each document. Any change in the data would render the pointers invalid. This guaranteed that data stored could not be tampered after the server had signed the documents. The paper also coined key concepts and terms such as timestamping, hashing, signature, data linkage, and distributed trust. These concepts are the cornerstone concepts and terms of modern blockchains.

The idea of digital currency can be traced even further back than blockchain. In 1982, David Chaum published his paper titled *Blind Signatures* [3]. This paper led to the creation of perhaps the first digital currency in the world called *eCash* which was released in 1993. In 1996, Douglas Jackson created a gold-backed digital currency called *e-gold* [4]. In 1998, a digital currency called *b-money* was proposed by a computer scientist named Wei Dai [5]. After that, Adam Back came up with his “hashcash” in 2002 [6] which made the first implementation of Proof-of-Work. About 5 years before 2008 when the Bitcoin paper was published, a decentralized currency system powered by Proof-of-Work (PoW) was created by the team led by Emir Sin Gün who published their paper in 2003 [7].

Since the publication of Satoshi Nakamoto's Bitcoin paper, an overwhelming amount of attention has been paid to Bitcoin and other similar digital currencies, together with the technology used in the currency transfer systems. The enthusiasm was due to the claimed advantages provided by Bitcoin: there is no need for a middle man in the transaction process so trustless parties can do business with each other; there is no central point of failure so system reliability is greatly enhanced; there is no double-spending so fraud can be avoided; transaction history is traceable so transactions can be verified; financial benefits are provided to participants and they are rewarded for their contribution to the operation of the Bitcoin system, for example, by mining blocks. In recent years, numerous Bitcoin-mining data centers have been built around the world to make profits. Due to its enhanced security, privacy, speed, and reduced cost, some organizations have started accepting Bitcoin as a means of payment, such as tuition fees. Up to date, people have been primarily focusing on the financial aspects of Bitcoin.

However, since 2014, greater attention has been paid to the technology which Bitcoin is based on: blockchain. People have realized that blockchain can be separated from Bitcoin and the technology can be beneficial to other areas where data security and privacy are of prime importance. Quick surveys of the literature have revealed that blockchain applications can be found in virtually every aspect of our lives nowadays. For this reason, the business value of blockchain has been predicted to be multibillion dollars by 2030.

8.1.3 Types of Blockchain

Blockchains are of different types. First of all, they can be *permissionless* or *permissioned.*

8.1.3.1 Permissionless Blockchains

Permissionless blockchains are *public*, open source, and based on the Proof-of-Work consensus algorithm. Anybody can participate in a permissionless blockchain without obtaining approval beforehand. The person can simply download the required software program and start running it on his or her own computer. The person can send transactions to the blockchain and these transactions will be included in the blockchain only if they are valid. Transactions are transparent so that everybody on the blockchain can view them, although the transactions are anonymous. In addition, anybody can participate in validating transactions before they are added to the blockchain. Well-known examples of public blockchains are Bitcoin and Ethereum.

8.1.3.2 Permissioned Blockchains

Blockchains can be permissioned as well. Permissioned blockchains can be further divided into two types: *federated* and *private*.

- *Federated* – Federated blockchains are usually operated by a special group of participants. They do not allow the arbitrary participant to validate transactions. Only preselected participants are involved in the process and are responsible for the validity of transactions. In federated blockchains, the public can be granted the right to read the transactions, but only the selected participants can write the transactions. Examples of federated blockchains are R3 which is for banking and EWF which is designed to be used in the energy sector.
- *Private*: A private blockchain is typically centralized to one organization and only this organization can validate transactions. The public can be allowed to read the transactions, or only some selected parties can read the transactions. Multichain is an example of private blockchains.

The difference between a federated blockchain and a private blockchain is the number of organizations that operate the blockchain. A private blockchain is operated by one organization. However, a federated blockchain is operated by multiple organizations, although federated blockchains can still be considered to be private blockchains.

Private and public blockchains differ in the execution of consensus algorithms, maintenance of the common ledger, and the authorization to join to the blockchain network. Table 8.1 shows the differences between the various types of blockchains.

Blockchains use different *consensus algorithms* which make blockchains different from traditional distributed database technologies. Consensus algorithms are essentially about decision-making in a group and how the decisions can be made for the benefit of the majority of group members. Well-known consensus algorithms are Proof-of-Work (PoW) and Proof-of-Stake (PoS).

8.1.4 Directed Acyclic Graph (DAG)

Similar to a blockchain, a DAG can also store data transactions. In a DAG, a transaction is represented by a node and is linked to one or several other transactions.

Table 8.1 Comparison of different types of blockchains

	Public	Private	Federated
Permission type	Permissionless	Permissioned	Permissioned
Reading	Anybody	Restricted	Restricted
Writing	Anybody	Restricted	Restricted
Validation	Anybody	Limited to one	Limited to several

However, the links are *directed* because they point from earlier transactions to newer transactions, in a way called *topological ordering*. A DAG does not allow loops. That means, a node is not allowed to traverse back to itself by following the directed links. In this sense, a DAG is acyclic. Essentially, from the standing point of computer science, a DAG is a graph with transactions being the nodes of the graph and the edges which have directions.

Unlike blockchains, DAGs do not have blocks. There is no mining in DAGs. Transactions provide validation for each other but a transaction cannot validate itself. A new transaction is required to validate one or more previous transactions when it joins the DAG. Every new transaction refers to its parent transactions, signs their hashes, and includes the hashes in the new transaction.

8.1.5 Hybrid DLTs Based on Blockchains and DAGs

Blockchains and DAGs can be combined to create hybrid DLTs. An example of hybrid DLTs is Bexam [8]. Bexam is a platform that leverages blockchain and DAG technologies with greatly improved speed and scalability.

Bexam uses a flexible chain structure together with a node hierarchy so that it has the security of a blockchain and the speed of a DAG with approximately 0.2 seconds per block and about 40 million transactions per second. It is highly scalable because it combines the concepts of DAG. Bexam uses a new consensus algorithm called Proof-of-Rounds (PoR) and a KYC (Know Your Customer) verification process to identify and prevent malicious actions. The electric power and computing resource requirements of Bexam are also very low.

In addition, a token technology is used in Bexam for its transactions. It is convenient to integrate Bexam into existing enterprise infrastructures.

8.1.6 Internet of Things (IoT)

Internet of Things is part of the Industry 4.0 revolution. IoT is a research and industrial focus in recent years.

Essentially, IoT is all about having *smart things* which are equipped with sensors and actuators collaborating over the Internet. IoT systems currently in use have a centralized architecture where data is stored in the cloud. Cloud is a central place with databases and services.

The IoT ecosystem is very complex. The complexity is due to the vast number of devices connected, the types of wired and wireless communication networks involved, and the varieties of software programs used. This complexity makes the IoT ecosystem vulnerable and susceptible to attacks.

The current IoT ecosystem was developed largely based on available Internet technologies in the past and did not have a systematically designed secure structure in the first place. Therefore, people are concerned with the security of current IoT systems.

There is a strong belief that blockchain is the solution for IoT security due to its intrinsic advantages such as distributed data storage and immutability. Blockchain may be able to improve the overall security of the IoT ecosystem.

8.2 Benefits of DLTs

8.2.1 Blockchain Benefits

A blockchain can be considered as a system that stores an identical copy of a spreadsheet called ledger on multiple distributed computers. The system frequently has no central authority. Transactions submitted by participants are validated and recorded in the ledger which is accessible to the community of participants. The transactions are cryptographically signed and are then assembled into blocks. The blocks are linked one after another and are added to the blockchain by consensus mechanisms. Transactions are immune to changes after they are published. Blocks are replicated across all computers of the distributed system so that they all hold the same ledger.

The way a blockchain is created and maintained leads to numerous benefits in comparison to traditional databases. Prominent features of blockchains such as decentralization and consensus are the intrinsic attributes of blockchain that give rise to benefits which are essentially out of the box:

1. First, blockchains *allow trustless participants to interact with each other*. No trusted third parties are required to serve as intermediaries and validate transactions. It is the consensus algorithms that validate the transactions. A blockchain can maintain itself by handling conflicts automatically and creating forks if necessary, so that the ledger is always in good standing.
2. Blockchains also make data storage more *reliable*. This is again due to the elimination of third-party agents which reduces the risk of unauthorized access and unwanted modification of stored data. The ledger is not stored in a single location or managed by any single company. The cryptographic linking between blocks ensures data unchangeability. Blockchains are robust. Transactions are processed by multiple participating nodes; therefore, no single node is critical to the entire database. No central point can be exploited so the system is much better against hacking and fraud. This equips blockchains with high fault tolerance capabilities. Blockchains are immune to malicious modifications. It is impossible to change it back once data has been written into the blockchain. No change of history is allowed either. Third parties cannot make any changes to the system as well.
3. Transactions stored in blockchains are *permanent*. Therefore, they are *verifiable* and *auditable*. Such verifications can be applied not only to data but also to

interactions and message exchanges among participants. Every transaction is tagged with a signature and a timestamp so that data ownership is auditable and traceable. Transactions can be traced back to its origin. Transaction history can be used to verify data authenticity and prevent fraud. This eliminates backdoor transactions and possible disputes and prevents data tampering.

4. *Transparency* is another benefit of blockchain technology. In a public blockchain, transactions, once they are made, are accurate and consistent among participants. All changes to a public blockchain are accessible to all participants. Users have full visibility of transaction information in the system. Anyone can verify the correctness of the system. Using a single public ledger avoids the complications of multiple ledgers. Transparency increases trust among participants too. This is particularly important in scenarios such as fair disbursement of funds or benefits. Everyone can maintain a copy of the ledger and verify its correctness. This provides resiliency and trust among participants.
5. Decentralization in blockchains leads to high *availability*. A blockchain is based on a large number of nodes working in a peer-to-peer manner. Data is replicated and updated on all nodes. Being inaccessible to a single node will not cause the system to stop functioning. Therefore, a system based on blockchains is highly available.
6. Blockchains have enhanced *security* and *integrity* over traditional database systems. Transactions will not be recorded before they are agreed upon by participants. Approved transactions are encrypted and linked as a chain. Together with the distributed copies, a blockchain is very difficult for hackers to break.
7. Blockchains can lead to *reduced transaction costs*. Transactions can be completed in a peer-to-peer or business-to-business manner. No third-party intermediaries are required. This avoids the cost induced by using a third party such as a bank. Cost that can be saved includes overhead, governance, auditing, and other fees.
8. Blockchains make transactions *faster*. In a blockchain, transaction time can be even reduced to just a few minutes. However, current interbank transfers and final settlement could take days. Transaction time is extremely important for industries such as transportation and energy. Time reduction could potentially save billions of dollars. The situation is similar in the financial industry. Blockchains can save time because they eliminate verification, reconciliation, and clearance which are usually lengthy processes. The reason is that a single version of data already agreed upon by participating financial institutions is available on the shared ledger of the blockchain.
9. The other benefit of blockchains is *smart contracts*. A blockchain such as Ethereum not only stores data, but also provides a programming logic called smart contracts. Smart contracts can execute business logic. They are programs that execute agreements and manage the transfers of digital assets between participants under specified conditions in a blockchain. They can be considered a digital version of traditional contracts written in a programming language. Because smart contracts are deployed and executed on blockchains, they are

therefore secured as well. The execution is also transparent, immutable, and decentralized. Essentially, smart contracts make program execution secured.

8.2.2 DAG Benefits

From computer science point of view, a DAG is a graph with directed edges and no cycles. It is a treelike data structure that is suitable for storing, organizing, and finding transactions.

As a distributed ledger technology, DAG has advantages over other technologies. For example, DAG does not have blocks and miners. Validation is done by the transactions themselves. New transactions validate old transactions in a distributed manner when they are added to the DAG. This greatly increases the speed of DAG – hundreds of thousands of transactions can be processed in a DAG in a second.

Distributed validation between transactions leads to much-improved scalability as well. The newer transactions are added to DAG, the more transactions that are available for validation, the faster the validation is done. In theory, DAG has infinite scalability. Because DAG does not have blocks and miners, there is no mining fee associated with DAG. This makes DAG an appropriate technology for the Internet of Things which has a large number of transactions between sensors and devices, and it is not logical and realistic to charge fees.

DAG is also easily made quantum-proof. That is, DAG is safe to use even when quantum computers become available in the future, because DAG does not rely on cryptography which could be potentially broken by quantum computers. Algorithms have been implemented in DAG to make DAG quantum-resistant. An example of DAG is IOTA technology which is already believed to be resistant against quantum computing attacks. However, DAG should have a substantial amount of traffic before it can start working. Greatly reduced traffic will make DAG vulnerable to attacks. Solutions based on coordinators have been suggested to get a DAG system up and running. The effectiveness of the suggestion of using coordinators is still under debate.

It should be noted that there are a number of differences between DAG and blockchains. First, DAG is blockless. In DAG, transactions validate each other and transactions are not assembled into blocks. On the other hand, blockchains assemble transactions into blocks. Secondly, DAG is more scalable. In fact, DAG is infinitely scalable in theory. This means that the performance of DAG will not deteriorate as new transactions are added to the graph. On the contrary, blockchains will experience slowdown when the blockchain gets longer. DAG does not require mining as well. Therefore, DAG uses much less electric power. However, blockchains based on PoW use a lot of electric power.

Another difference is that DAG does not charge fees whereas blockchains do. Furthermore, DAG is much faster because it does not require mining and validation is done in parallel and not in a chained manner. Finally, DAG is quantum-proof. Blockchains are susceptible to quantum attacks because they are based on cryptography and consensus algorithms which are breakable by quantum computers. Table 8.2 lists the differences between DAG and blockchains.

8.3 How Blockchain Works

The functionality of a blockchain in Bitcoin is to facilitate money transactions and the recording of the transactions. A slightly simplified microscopic operation of a blockchain can be illustrated by looking at the process of how a transaction is started and settled, as shown in the flowchart in Fig. 8.3.

In general, a successful transaction has to go through a sequence of steps, as shown in the following:

1. Step 1: Form a transaction. In this step, sender information, receiver information, sender's public key, amount of fund to transfer, receiver's public key, and timing information are required.
2. Step 2: Form a block. In this step, the previous block hash, the current block containing the transaction in Step 1, and other transactions are included.
3. Step 3: The block is broadcasted to the entire network.
4. Step 4: Nodes on the network validate the block.
5. Step 5: The block is added to the blockchain.
6. Step 6: Fund transfer is completed.

Table 8.2 Comparing DAG and blockchain

	DAG	Blockchain
Using block	No	Yes
Scalability	Good	Poor
Mining	No	Yes (for PoW)
Fees	No	Yes
Speed	Fast	Slow
Quantum-proof	Yes	No

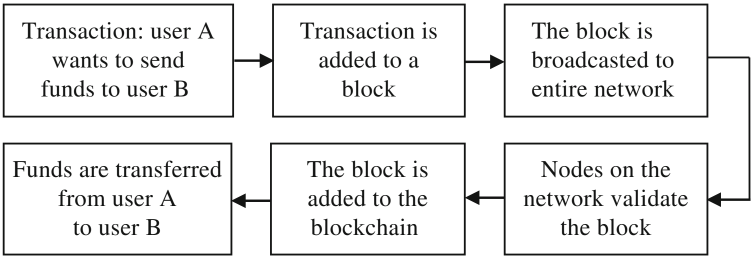
Transaction: user A wants to send funds to user B
Transaction is added to a block
The block is broadcasted to entire network
Nodes on the network validate the block
The block is added to the blockchain
Funds are transferred from user A to user B

Fig. 8.3 Operation of a blockchain

8.3.1 Transaction, Block, Ledger, and Blockchain

There has been some confusion about the concept of the blockchain and the ledger that stores the data for a blockchain. To be precise, a *blockchain* refers to the system of nodes that make a blockchain operational, whereas the term *ledger* is the database stored in the nodes. However, the term *ledger* is used for accurate descriptions. A ledger is a data structure that is replicated and shared among distributed nodes of the blockchain network. A ledger can be considered as *a chain of blocks*. Each block in the chain carries a list of transactions and other data. A transaction has a transaction ID, an input which contains the type of the asset to be transferred and the amount and signed with the sender's public key, and an output which includes the type of asset to be received and the amount and signed with the receiver's public key. Figure 8.4 shows an example of a transaction which is simplified for the convenience of description.

After being validated, transactions are assembled into blocks. A block consists of three parts: the block header, the hash of the block header, and the transactions inside the block. The block header is made in a special way. It contains the hash of the header of the previous block, a timestamp when this block is created, and a Merkle root hash which is derived from the hashes of the transactions of this block. A Merkle root hash is the hash of all the hashes of all the transactions that are part of a block. The block also contains two other important parameters, namely, nonce (which stands for *number used only once)* and difficulty target. These two parameters are what make mining (in a Power-of-Work blockchain) tick. The details of the mining process will be discussed in the next section. Figure 8.5 shows the details of a block (the shaded area is the block header).

	Input	Output
Transaction ID	Type, Amount, Sender Key	Type, Amount, Receiver Key

Fig. 8.4 A transaction

Fig. 8.5 A block

Version
Hash of Header of Previous Block
Timestamp
Difficulty Target
Nonce
Merkle Root Hash
Hash of Header of This Block
Transactions

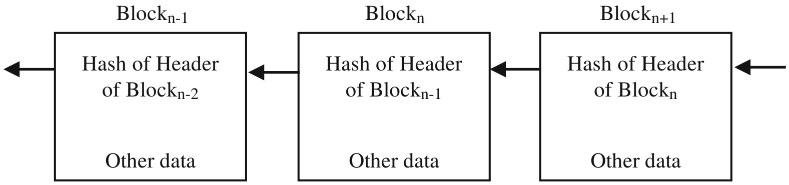
Block$_{n-1}$
Hash of Header of Block$_{n-2}$
Other data
Block$_{n}$
Hash of Header of Block$_{n-1}$
Other data
Block$_{n+1}$
Hash of Header of Block$_{n}$
Other data

Fig. 8.6 A ledger

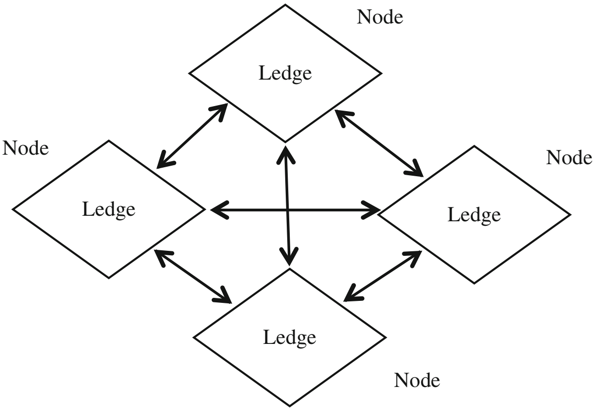

Fig. 8.7 A blockchain as a distributed system of nodes

The very first block of a blockchain is called *genesis block* which is common to all nodes in the blockchain and has no parent. The hash of each block is obtained cryptographically and is the block's identity. Each block contains the hash of the previous block; this way a chain of blocks is established. This chain of blocks is the ledger which is frequently referred to as blockchain. Figure 8.6 depicts a ledger.

It is the ledger shown in Fig. 8.6 that is stored in multiple networked distributed computers. These computers are called *nodes* and they form the blockchain. So, it is the ledger that is stored in the nodes of a blockchain. The nodes communicate via wired or wireless communication networks in a peer-to-peer manner. A blockchain is shown in Fig. 8.7.

8.3.2 Transaction Validation and Block Mining

A user who wants to interact with a blockchain must do it through a node, although several users can share the same node. Through this node, a user can sign and initiate transactions. Every transaction is signed with the user's private key and can be accessed through the user's public key, which essentially serves as the "address" of the transaction. Transactions are broadcasted by a user's node to its immediate neighboring nodes.

The neighboring nodes validate each transaction and propagate it further along possible pathways. All nodes of the blockchain will have this valid transaction after some time. The neighboring nodes will block and discard transactions that are not invalid.

After a given time period, a node will have received a number of valid transactions. The node will then have the transactions organized in order, have them validated and packed into a timestamped candidate block, find a nonce value to create a hash which satisfies the difficulty level set by the blockchain, and have the candidate block broadcasted to all other nodes in the blockchain for verification.

The nodes in the blockchain all participate in verifying the validity of the candidate block. They make sure that the format of the block is correct. They make sure that each transaction in the block is valid and is signed by the suitable parties. They make sure that all hashes in the new block were computed correctly. They also make sure that the candidate block references to the hash of an appropriate previous block in the ledger. If the result of the verification process turns out to be positive, every node will add the block to its own copy of the ledger. If the candidate block is not valid, then it will be discarded. This process will repeat indefinitely as long as the computer network is not down for any reason.

A critically important question is how a node should decide if a transaction is valid. First of all, a node needs to ensure that the signatures (hashes) of the sender and the receiver are valid. That is, the sender and receiver are both legitimate registered participants of the blockchain and they do have valid "accounts" in the blockchain. The amount to be sent should also be valid in terms of the type of assets and minimum allowed value based on the kind of applications. A node also must validate if the sender has sufficient unspent funds. Figure 8.8 shows how a transaction is validated.

However, the above validation process assumes that every node can be trusted, which is usually not the case for a public blockchain. A public blockchain usually consists of a group of non-trusting participants. Therefore, a set of rules are required for the nodes to agree on the validity of the transactions. Because the transactions are assembled into blocks, blocks need to be validated after the transaction validation is carried out. In blockchains, consensus algorithms are employed to validate blocks. The opinion of the majority of the nodes on the blockchain will decide the validity of the blocks.

The problem is that a bad user can create multiple participant identities via one specific node and can therefore potentially control the entire blockchain. In order to avoid such a problem, what Bitcoin does is making the finding of a new valid block very *computationally expensive* so that a bad node is not able to beat other nodes collectively on the blockchain because of a single node's limited computing power. This is the *consensus mechanism* called Power-of-Work. Based on this mechanism, malicious blocks from a bad node are unlikely to be accepted because it is up to the majority of the nodes on the blockchain to approve the validity of a candidate block.

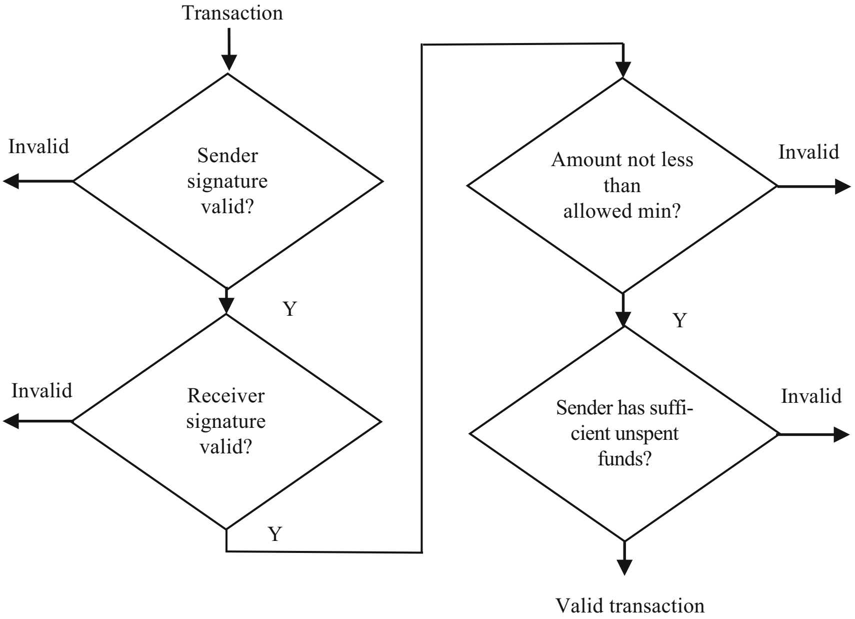
Transaction
Invalid
Sender signature valid?
Y
Invalid
Receiver signature valid?
Y
Amount not less than allowed min?
Invalid
Y
Sender has sufficient unspent funds?
Invalid
Valid transaction

Fig. 8.8 Validate a transaction

In a Power-of-Work blockchain, any node can make efforts by conducting mining to find and recommend a new block as the next valid block for the blockchain.

During mining, a node strives to find a suitable random number called nonce ("number only used once") which is embedded in the block's header (see Fig. 8.9). A valid nonce is a value that makes the hash (e.g., SHA-256 for Bitcoin mining) of a block header have the required number of leading zeros, as set by the difficulty parameter of the blockchain. The number of leading zeroes is called the *difficulty* which is set by the blockchain and can be adjusted over time. In other words, difficulty is a measure of how difficult it is to find a suitable hash based on the given difficulty target. Note that in blockchain, usually the network automatically adjusts the difficulty level for mining over time. The validity of a new block can be easily verified by other nodes because they only need to validate the hash using the nonce value already found by the node that is recommending the new block. This takes only a very short amount of time because it involves only the calculation of one hashing algorithm. Other nodes will adapt and add the validated recommended block to its own copy of the ledger. In general, nodes will validate and adopt broadcasted recommended blocks rather than trying to mine its own. It is better to validate and adopt an existing recommended block and start mining the next recommended block because this at least gives the node the chance of winning a reward for successfully mining the next valid block. The rationale of this strategy is because of the way a blockchain resolves conflicts: only the block which gives the longest chain will be adopted.

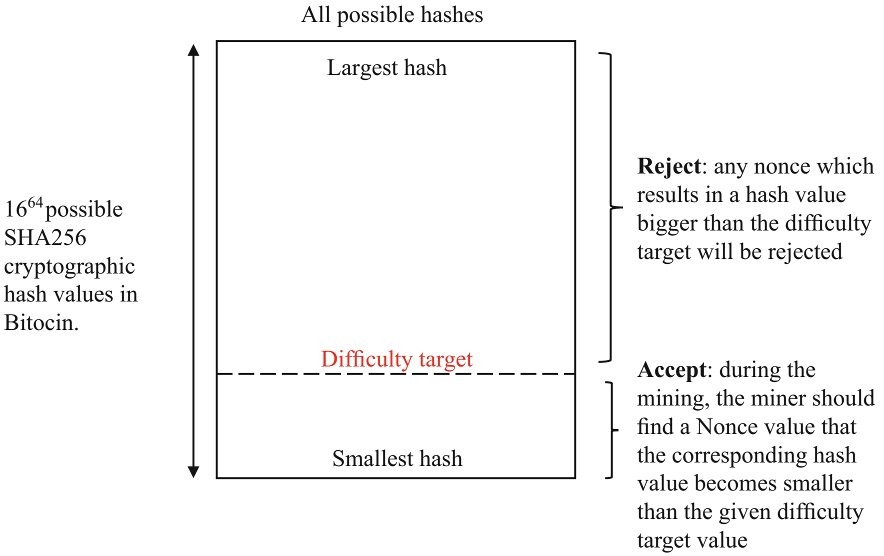
All possible hashes
Largest hash
16^{64}possible SHA256 cryptographic hash values in Bitocin.
Difficulty target
Smallest hash
Reject: any nonce which results in a hash value bigger than the difficulty target will be rejected
Accept: during the mining, the miner should find a Nonce value that the corresponding hash value becomes smaller than the given difficulty target value

Fig. 8.9 How mining works. Miners search for a valid hash that satisfy the given difficulty target. In Bitcoin nonce range contains 4 billion possible values

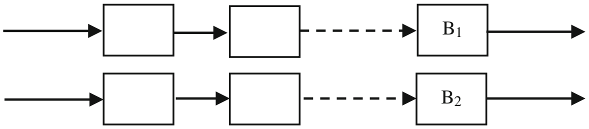

Fig. 8.10 Forking in blockchain

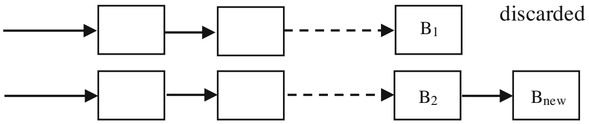

Fig. 8.11 Longer chain is adopted

Therefore, in general, the nodes will validate and adopt the first recommended block broadcasted over the blockchain, and then they will all start mining the next recommended block. If there are two nodes that publish a valid candidate block at the same time, conflict occurs because both candidate blocks will be added to the ledgers on different nodes. It is very likely that these blocks contain different transactions; therefore, the last blocks in the ledgers are not the same. If this happens, a fork is created. The strategy to resolve this conflict is to wait for a new block to be added. Then all nodes will adopt the ledger that has the longest chain because it carries the greatest amount of work in a Proof-of-Work-based blockchain. This way consensus is reached if a block should be in the ledger. The conflict (fork) scenario is shown in Fig. 8.10 where blocks B_1 and B_2 are both valid and have been added to the ledger, but they may contain different transactions inside. Figure 8.11 shows how a conflict is resolved – the longer chain is adopted as a valid ledger.

8.3.3 Smart Contracts

The second generation of blockchains such as Ethereum [9] uses smart contracts. Smart contracts are executable computer code stored on a blockchain. Similar to transactions, smart contract code is transparent and can be examined by all participants of the blockchain. Smart contracts are accessed via their addresses and users can activate them by sending them transactions. In an Ethereum blockchain, smart contracts run on every node of the blockchain which has a virtual machine running. The smart contracts execute under the virtual machine.

To a large extent, smart contracts are very similar to stored procedures in database management systems. Functions in smart contracts are defined based on business rules. Participants can initiate transactions to call functions in smart contracts along with the required data. Smart contracts are deterministic and the same input to a smart contract always generates the same output. Smart contracts can call each other as well. The code of a smart contract can be examined by participants so that they can predict the outcome before they commit the contract. The outcomes of executing smart contracts can be verified by participants too. Smart contracts help prevent possible contract disputes.

8.3.4 Consensus Algorithms

When designing or adopting blockchains, several factors need to be considered. The first factor to consider is who can access the blockchain. This depends on whether the blockchain will be openly accessible by the public or not. If it is, public blockchains must be used. The decision also affects the selection of consensus mechanisms. A private blockchain needs to be adopted if only specific participants can access the blockchain. Private blockchains can employ special consensus algorithms so that time-consuming minings are not required.

As stated previously, a blockchain consists of multiple identical ledgers stored in distributed computers. A blockchain may not be controlled by any central authority. It is therefore obvious that malicious users will be very much attempted to take advantage of the blockchain. This situation of the blockchain is very similar to the famous Byzantine Generals' Problem. From this sense, blockchains require *Byzantine Fault Tolerance* (BFT). Without BFT, bad users will be able to break the blockchains by sending malicious transactions. If damages do occur, they will not be repaired because no central authority is available to carry out corrective actions. For this reason, consensus algorithms are used to provide BFT.

So far, large numbers of consensus algorithms have been developed for blockchains. With consensus algorithms, mutually distrustful participants can work together. Each of the consensus algorithms has its own strengths and weaknesses and is suited for specific applications. Although the number of proposed consensus algorithms is large, only several of them are widely known. Selected consensus algorithms are discussed below.

8.3.4.1 Proof-of-Work (PoW)

Proof-of-Work, or PoW [10], is the most well-known consensus algorithm which was used by Bitcoin. The purpose of this algorithm is to validate transactions and add validated new blocks to the blockchain. Due to its public and distributed nature, a blockchain needs a mechanism to prevent malicious transactions and attacks. This is the responsibility of participating nodes called miners through a process called *mining*. Essentially, PoW presents a complex mathematical puzzle for the nodes to solve. Very strong computing power is needed to solve this puzzle in a timely manner. However, proving the correctness of a solution for the puzzle is easy. In the meantime, miners receive rewards for solving the complex puzzle. In summary, in a PoW-based blockchain, miners strive to validate transactions, solve the puzzle, propose candidate blocks, and receive rewards. The amount of work done by a miner determines its chance of successfully mining a single block and receiving a reward.

PoW not only provides a solution for the Byzantine Generals' Problem, but also provides defense against denial-of-service (DoS) attacks because it takes a tremendous amount of computational power for an attack to be successful. In a PoW-based blockchain, the available fund/credit in the attacker's wallet does not increase its ability to publish new blocks. What really matters is a node's computational power to solve a puzzle and generate new blocks. PoW strongly discourages DoS attacks on a blockchain because it is highly unlikely that an attacker has the ability to acquire enough hardware and energy resources to overpower the rest of the nodes on a blockchain as a whole.

However, PoW has weaknesses and users should be aware of them. First, a PoW-based blockchain is vulnerable to the so-called 51% attack, in which case the attacker has the majority of the mining power for whatever reason. With 51% or more of the mining power, the attacker is able to control the operations of the blockchain and prevent other mining nodes from creating new blocks. By doing this, only the attacker will get the rewards. With 51% of the computing power, the attacker can even reverse transactions.

The second weakness of PoW is huge power consumption due to the need for solving complex puzzles. It has been observed that the Bitcoin blockchain is currently using more power than the whole country of Ireland and will use more power than the whole country of Denmark by 2020.

8.3.4.2 Proof-of-Stake (PoS)

Proof-of-Stake, or PoS [11], was designed to overcome the weaknesses of PoW. The basic rationale of PoS is that a node who owns more stakes in the blockchain will more likely want it to succeed. To be able to be admitted to the blockchain, a node needs to have a specific amount of assets stored in its wallet. Furthermore, a node needs to deposit some assets as stake in order to qualify as a miner. Although every node is entitled to validate and mine new blocks based on their asset possession, actual miners are randomly chosen by the blockchain based on the assets stored in their wallets. The blockchain will examine all nodes with their stakes and choose some of them as miners based on the ratio of their stakes with respect to the overall system stakes. That is, if a node owns 10% of the total stakes, then it has 10% of the chance to be selected as a miner. A node with only 1% of the total stakes will only be selected 1% of the time. The next new block will be voted for by all users with stakes. However, in PoS-based blockchains, although the nodes with more initial stakes can potentially accumulate more and more digital assets, the blockchain is designed in such a way that it is extremely difficult for several nodes to acquire the majority of assets within the blockchain. This way, no nodes will be able to dominantly manipulate the blockchain as they wish.

Comparing to PoW-based blockchains, PoS-based blockchains do not need powerful computing hardware. A functional computer with a stable Internet connection is all that is needed to work as a node. PoS-based blockchains are much more energy efficient than PoW-based blockchains because they do not use much electric power in their operations. Not having mining operations also enables PoS-based blockchains to run much faster than PoW-based blockchain. A PoS-based blockchain has very little chance of having a 51% attack because of its design.

The main disadvantage of PoS-based blockchains is that it is impossible to achieve full decentralization. The reason is that in a PoS-based blockchain, only limited numbers of nodes are participating in creating new blocks.

8.3.4.3 Delegated Proof-of-Stake (DPoS)

Another well-known consensus algorithm is the Delegated Proof-of-Stake (DPoS) [12] invented by Daniel Larimer. In DPoS, there are three groups of entities: stakeholders, witnesses, and delegates. The responsibilities of stakeholders are the election of witnesses. The responsibilities of witnesses are the creation and addition of blocks to the blockchain. The responsibilities of delegates are maintaining the blockchain and suggesting changes to the blockchain.

Witnesses are elected by the stakeholders. Each stakeholder has one vote for one witness. Witnesses with the highest number of votes are elected. Stakeholders vote to increase the number of witnesses until at least 50% of the stakeholders consider the blockchain has achieved sufficient decentralization.

Elected witnesses take turns to produce new blocks in given timeframes. However, the quality of their work is monitored by stakeholders via a reputation scoring system. Poorly performing witnesses will lose scores or their titles. Stakeholders will continuously vote for the witnesses. Part of the witnesses is replaced at regular intervals as well.

Delegates are also elected by stakeholders. However, their responsibility is to maintain the blockchain. For example, delegates can suggest block size changes, paid incentive, and transaction fee changes. The stakeholders will decide if the proposed changes should be implemented. Delegates may receive rewards as well.

Energy saving and decentralization promotion are the two main advantages of DPoS. DPoS needs less energy than PoW because witnesses generate blocks based on specific time schedules, rather than competing with each other to add blocks. The computing hardware requirement is no longer as demanding as PoW as well. In addition, greater decentralization is achieved in DPoS because its consensus mechanism allows stakeholders to choose suitable witnesses to validate transactions.

The main disadvantage of the DPoS consensus mechanism is that it can never achieve full decentralization, although decentralization can be increased by having more witnesses validate blocks, due to scalability constraints.

8.3.4.4 Practical Byzantine Fault Tolerance (PBFT)

The practical Byzantine Fault Tolerance (PBFT) consensus algorithm [13] is another popular consensus algorithm used in blockchains. PBFT enables a blockchain to tolerate Byzantine faults, i.e., defend against attacks from malicious nodes. The algorithm is designed to work in asynchronous systems. PBFT has low overhead time and low latency.

In PBFT, all nodes of a blockchain are organized into a sequence. A specific node is designated as the leader node. Other nodes are designated as backup nodes. When a node sends out a message, the rest of the nodes will exchange information with each other to validate the message in case it is tampered during transmission. It is expected that the good nodes will reach an agreement on the state of the blockchain through majority.

Each round (called *view*) of the PBFT works as follows:

1. A client sends a request to the leader node.
2. The leader node broadcasts the request to backup nodes.
3. The backup nodes execute the request and send a response to the client.
4. The client waits to receive $f + 1$ node responses with the same result which will be used as the result of the operation, where f represents the maximum number of potentially faulty nodes.

To secure its role, the leader node may be changed in a round-robin fashion during every view. The leader node can even be replaced if it does not broadcast a request after a given time interval. The majority of good nodes also have the power to identify a faulty leader node and replace it with the next leader.

To give more details, here is how PBFT works in Fabric:

1. One of the nodes is elected as a leader.
2. Transaction requests are submitted to the leader.
3. The leader organizes the transactions into an ordered list and broadcasts this list to all other nodes in the blockchain for validation.
4. Every validating node executes the ordered transactions one by one. Then it calculates the hash code for the new block which is based on the received transactions. Then this validating node broadcasts the hash code to other validating nodes and starts counting the responses from them.
5. If a validating node realizes that two-thirds of all validation peers have the same hash code, it will add the new block to its own copy of the ledger.

The PBFT model works only if the number of malicious nodes in a blockchain does not exceed one-third of the total nodes in the system in a given time window. The more nodes are there in the blockchain, the more unlikely for the malicious nodes to reach one-third of the total nodes.

The PBFT algorithm has two main advantages compared to other consensus algorithms. The first advantage is that it can finalize transactions and blocks without needing confirmations as what is done in PoW. The second advantage of the PBFT model is that it uses significantly reduced energy, again as compared to PoW.

There are two limitations to the PBFT consensus algorithm. First, it works well only for blockchains of small sizes due to its communication model among nodes. Second, it is susceptible to Sybil attacks. Due to the first limitation, the size of the blockchain cannot be increased significantly just to mitigate Sybil attacks. Luckily, possible solutions have been identified to solve this problem. For example, PBFT can be interlaced with PoW to overcome both limitations.

8.3.4.5 IOTA

A totally different technology in the cryptocurrency family is *IOTA* [14]. IOTA is an open-source distributed ledger with great potential for applications in the Internet of Things.

IOTA works on the platform called Tangle. Tangle hashes use Winternitz signatures [15] which is a hash-based cryptography, unlike blockchains that use elliptic curve cryptography or ECC. Winternitz signatures are much faster than ECC. The actual hash function used by Tangle is Kerl [16] which is a version of SHA-3. Kerl works based on ternary operations, which is more secure than other crypto technologies used in blockchains. Currently, many crypto algorithms can be broken by superfast quantum computers. However, it is very difficult for a quantum computer to break ternary operations used by Kerl. The chances of Tangle suffering from a quantum attack are roughly 1 million times less than the blockchain.

8.4 Directed Acyclic Graph (DAG)

8.4.1 What Is a DAG

As discussed previously, DLT, or "distributed ledger technology," has its set of records (the ledger) held by multiple distributed nodes. For instance, the cryptocurrency Bitcoin has a blockchain which is a DLT with its ledger (transactions) stored in multiple computers. Each new transaction added to the ledger is copied to other computers. This ensures that multiple copies of the ledger are available.

DAG is a type of ledger. A DAG is a graph with directed edges and no cycles. A DAG has its nodes sorted in a special order, which is called *topological sorting*. In a DAG, each transaction is linked to at least one other transaction. The edges are directed from earlier transactions to recent transactions. Loops are not allowed in DAGs, which means that a transaction cannot travel back to itself if it follows along the directed edges. Figure 8.12 shows a DAG.

8.4.2 How IOTA Tangle Works

Tangle is IOTA's DAG that operates in a special way. In Tangle, each new transaction must validate at least two previous transactions before it can be added to the DAG. With Tangle, all nodes on the IOTA network can issue and validate transactions at the same time. In Tangle, data are attached to transactions. However, Tangle does not assemble transactions into blocks. Therefore, Tangle is blockless.

Tangle does not require mining to reach consensus. This avoids powerful mining computers and extensive use of electric energy. No mining also means no fees are needed to reward miners. Users do not need to pay transaction fees as well.

Tangle is highly scalable because of its use of DAG as its ledger and simultaneous transaction processing. Increased transactions in a DAG do not slow down the IOTA network. In fact, performance will improve as the number of transactions increases due to the characteristic of simultaneous validation. IOTA with Tangle has a higher speed than blockchains.

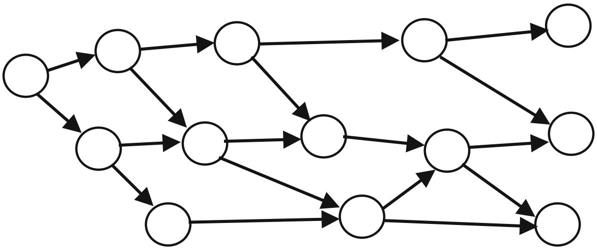

Fig. 8.12 A directed acyclic graph (DAG)

When Tangle gets started, it uses a coordinator to prevent malicious activities because it does not have enough transaction nodes to conduct validation. This coordinator will be obsolete after the IOTA network becomes more established. The use of an initial coordinator creates the possibility of a central point of failure.

8.5 DAG Versus Blockchain

Similar to blockchains, DAGs store transactions on a distributed ledger. However, the ledger is quite different in a blockchain than it is in a DAG. In a blockchain, the distributed ledger is a chain of blocks which are built using transactions. Blocks are validated and chained up in chronological order. Chained blocks are not modifiable. A blockchain is very similar to a linked list concept in computer science. On the contrary, a DAG is a collection of transactions linked in special ways. There are no blocks in a DAG. A DAG can be compared to a tree in computer science. Figure 8.13 compares the structures of blockchain and DAG.

Consensus is achieved differently in blockchains and DAGs. In blockchains, consensus is achieved by validating transactions block by block via mining. On the other hand, DAGs have transactions validate their immediate predecessors.

Based on the above discussions, blockchains have better immutability than DAGs, whereas DAGs are better at handling a large number of transactions. Blockchains do not scale well, but DAGs do. DAGs are vulnerable to attacks if the volume of transactions is too low.

8.6 Blockchain and Internet of Things

8.6.1 Internet of Things

Internet of Things (IoT) [17] is a natural extension of the human being's efforts of connecting the world through computer networks. So far computers around the world have been connected for sharing information. The World Wide Web (or the

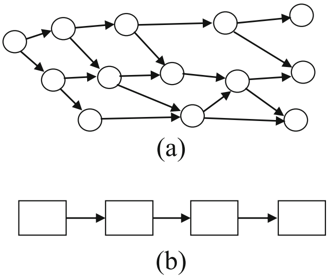

Fig. 8.13 (**a, b**) Comparison of blockchain and DAG structures

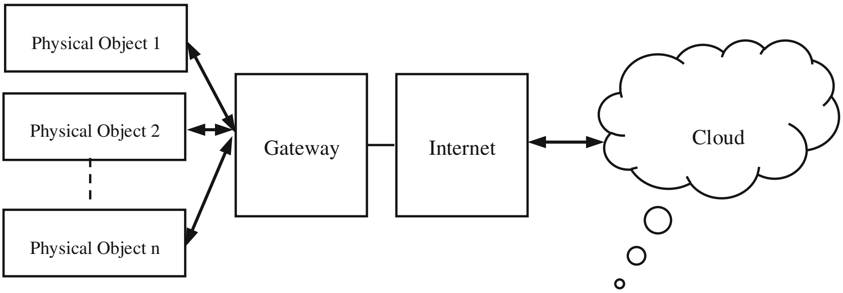

Fig. 8.14 A typical centralized architecture of the Internet of Things

Web) is an indicator of this usage. The Web has been used to share and exchange digital textual or visual information in the form of electronic documents. As the digital revolution continues, the natural next step is to connect all the physical objects in the world, have them not only exchange data, but also interact with each other, in order to make our lives more convenient, more efficient, and safer. These efforts led to the birth of the Internet of Things, or IoT for short.

Currently, most IoT solutions are designed based on a centralized architecture (see Fig. 8.14).

A current IoT system consists of the physical objects/devices, the gateway, the Internet, and the cloud. Physical objects can send data through the gateway and the Internet to the cloud and the data can get stored and analyzed there. Physical objects can also receive commands from other physical objects through the cloud to perform specified actions. Commands can be issued from a central manager from the cloud directly as well. The IoT stack and standard protocols create the layers of an architecture that provides services to IoT physical objects.

In the history of people's efforts in trying to connect physical objects in the real world, large peer-to-peer (P2P) wireless sensor networks (WSNs) were conceived and were once the focus of research. The missing pieces in these researches in terms of fundamental architecture design are privacy and security, which should really have been considered at the beginning of system design. WSNs originally were not designed to operate at a global scope as well.

8.6.2 Weaknesses of Internet of Things

Core recent developments on IoT moved towards the above cloud-based centralized architecture. However, the centralized IoT architecture has numerous *weaknesses*. In this centralized architecture, all information is sent from the physical objects to the cloud where data is processed using analytics tools. Responses are sent back from the cloud to the IoT physical objects if necessary. This type of centralized structure has *poor scalability*. The problem will become even worse when billions of new physical objects are to be added to IoT networks in the near future.

The second weakness of the centralized IoT architecture is a *single point of failure*, because every physical object is potentially a vulnerable point and can compromise the security of the entire IoT network. Failure of a single physical object can potentially bring down the entire IoT network as well.

The third weakness of the centralized IoT architecture is to do with *maintenance*. Updating software in the current IoT network is extremely difficult due to the fact that software updates need to be distributed to a huge number of physical objects which can be physically located anywhere.

The fourth weakness is related to *security and privacy*. Data spoofing and corruption can occur anywhere on the IoT network, ranging from the physical objects, the communication networks over which IoT data travel through, and the cloud storage where IoT data are gathered, stored, and processed. Unauthorized access to personal data in the cloud can happen which has always been the concern of the general public.

The fifth weakness is that IoT systems frequently use *resource-constrained computing devices* such as microcontrollers. These microcontrollers lack the computing power and storage capacity to support advanced and computation-intensive algorithms which can assist in protecting data security and privacy.

The sixth weakness is that current IoT systems have *no immutable records* of the history of interactions among physical objects. Because of this weakness, it is very difficult to track down the causes if problems do occur.

Another weakness of IoT is that the current centralized structure has only one copy of the data stored in the cloud. If this copy of data is tampered, there is no way to know what has been changed. There is no way to prevent the tampering from happening as well.

Because of these weaknesses, IoT faces the challenge of people *lacking trust* in technology, primarily due to their concerns on privacy and security. Their perception of the scale and complexity of IoT systems makes the situation worse because it is beyond their comfort zone. Granting device access and control to technological service providers is frequently a difficult decision and is a sensitive matter for IoT system owners as well.

IoT devices such as connected actuators are often required to perform actions according to the commands they receive from the cloud or other IoT devices. If such commands are hijacked, the consequence could be disastrous. A small example would be that the door of a house is wrongly opened for a burglar. Improper actions of devices could also lead to fires and flood in buildings and offices.

Overall, current IoT systems are subject to physical object identity-based attacks, manipulation-based attacks, cryptanalytic attacks, and service-based attacks.

8.6.3 Blockchains and IoT

Blockchain technologies have exactly what is needed to fix the weaknesses of centralized IoT. It is easy to perceive that the decentralized structure, the way that data is created and stored, and the consensus mechanism used will help overcome most of the weaknesses of the current IoT systems.

Depending on the use cases, blockchain technologies can be applied to each level of the IoT systems. Blockchains can be used to store and manage device IDs, encode and verify data packets on the communication networks, and secure data in the cloud and data stored in the distributed devices.

Blockchain technologies can be applied at a small and local scale such as smart homes and smart buildings, or to larger scales such as in smart cities, or even at a global scale for cross-continent IoT systems.

Blockchain technologies will help reduce IoT operational costs and prevent threats and attacks. Blockchains are unique and attractive because they have the following features: transactional privacy, security, data immutability, auditability, integrity, system transparency, and fault tolerance.

Wired and wireless communication technologies have reached new high levels. The technologies are still evolving, witnessed by the growing interest of adopting 5G technologies in IoT. It can be predicted that the requirements of data transmission speed by IoT will be up to users' expectations.

For this reason, privacy, security, and transparency and trust should be at the center of future IoT system designs. They should be considered right at the beginning when an IoT system is conceived.

In summary, as an emerging technology, IoT is promising and has a great future. Current IoT systems use resource-constrained devices which are ideal targets for cyberattacks. They have poor scalability and have the problem of a single point of failure. Maintenance is difficult. IoT data are not immutable. Privacy and security are critical concerns of IoT.

Blockchains can mitigate IoT risks and issues by using a large number of individual nodes that exchange data on a peer-to-peer (p2p) basis. Data records are immune to tampering and corruption. The consensus mechanism of blockchain can prevent malicious nodes from joining the IoT network, rejecting the data they send, and ensuring data integrity.

Among the various blockchains, practical Byzantine Fault Tolerance (PBFT)-based blockchains appear to be especially suitable for IoT, due to their abilities to defend against attacks from malicious nodes, work in asynchronous systems, and have low overhead time and low latency.

The other promising blockchain IoT platform is IOTA. It was designed specifically for the Internet of Things. IOTA is blockless and does not use computation-intensive mining algorithms. Instead, users verify the transactions of other users. The main advantage of IOTA is greater scalability.

With the support of blockchain technology, IoT systems will have the characteristics of decentralized resource management, robustness against threats and attacks, fault tolerance, and improved trust.

8.6.4 How to Combine Blockchains and IoT

According to the use cases and goals, blockchains can be combined with IoT in different ways. Figure 8.15 shows a block diagram for the current cloud-based centralized IoT system.

Blockchains can be applied to IoT systems in two ways, depending on the purposes of the application. The most comprehensive implementation uses a blockchain to record all data and interactions between physical objects [18], as shown in Fig. 8.16.

In Fig. 8.16, all data and interactions go through the blockchain. In this architecture, data and interactions are validated and their records are immutable. This is useful if both data and device interactions are important for the application. The drawback of this architecture is increased latency, increased bandwidth requirement for the communication network, and increased data flow on the network.

The other choice of combining blockchain and IoT is storing only IoT data in the blockchain, as shown in Fig. 8.17.

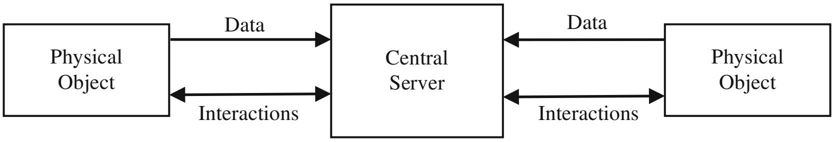

Fig. 8.15 Data and interactions of IoT physical objects are stored in a central server in centralized IoT

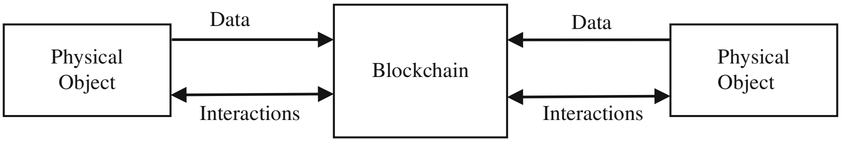

Fig. 8.16 Data and interactions of IoT physical objects are stored in the blockchain

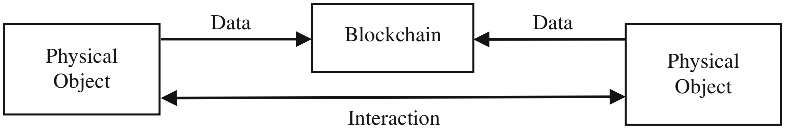

Fig. 8.17 Only data of IoT physical objects are stored in the blockchain

In Fig. 8.17, interactions between physical objects are not stored in the blockchain. In this architecture, only IoT data are validated and only data records are immutable. This is useful if IoT data are important for the application, but interactions between physical objects are not critical. This architecture has less latency, reduced bandwidth requirement for the communication network, and reduced data flow on the network.

In addition, smart contracts can be employed in the blockchain to set up specific requirements and agreements that govern data flow and usage, as well as monitor, allow, or disallow interactions to occur.

8.7 Prominent Enterprise DLT Platforms

Although there are tens of DLT platforms introduced in the literature, so far at the enterprise level, three of the platforms are most prominent. They are Hyperledger Fabric, Ethereum, and IOTA (See Table 8.3).

Enterprises have special requirements for DLTs. Ideally, enterprises require decentralized data storage. Data stored should be immutable and permanent. Security and privacy of data are of prime importance. Enterprise DLTs should support smart contracts and tokens.

8.7.1 Hyperledger Fabric

Hyperledger refers to several technologies. Hyperledger Fabric, developed by IBM, is one of them. Hyperledger Fabric was designed for B2B applications. IBM's client server-based architectures are used in Hyperledger Fabric to provide decentralized

Table 8.3 Comparing Ethereum, Hyperledger, and IOTA

	Data storage	Security and privacy of data	Support for tokens	Support for smart contracts	Immutability and persistency
Ethereum	Truly decentralized	Offered	Offered	Offered	Supported
Hyperledger Fabric	In members of a private consortium	Offered	No Support	Supported	Supported
IOTA	Not decentralized, but maintained by several central components	Offered	Not Offered	Not Supported	Not Supported

data storage. Private transactions are used to provide data security and privacy. However, Hyperledger Fabric does not support tokens.

Hyperledger Fabric supports smart contracts through chaincode. The operation of Hyperledger Fabric depends on a number of central participants.

Overall, Hyperledger Fabric is suitable for use cases where data is to be exchanged between a closed group of companies. It is not suitable for fully distributed applications.

8.7.2 Ethereum

Ethereum is a truly decentralized DLT which can run as a public blockchain, private blockchain, or consortium blockchain. Ethereum provides truly decentralized data storage through its architectural design. Data in Ethereum is less secure and private because it originally focused on the public chain. Tokenization is supported so that real assets can be digitally represented. Companies can build digital business models using Ethereum.

Smart contracts are seamlessly integrated into Ethereum and can be programmed using its built-in programming language called Solidity. Smart contracts are executed in the Ethereum Virtual Machine (EVM). Distributed apps (DApps) can be developed and run under the EVM. DApps can be deployed on Ethereum without additional infrastructure. Data immutability is guaranteed by Ethereum's architecture.

8.7.3 IOTA

IOTA can be used as a data layer on top of IoT to facilitate transactions between machines. The nodes in an IOTA network can both generate and confirm transactions. IOTA is a "feeless" DLT. Currently, IOTA is not truly decentralized because it depends on central maintaining elements. Data security and privacy are achieved via transaction validation, data encryption, and subscriber authorization. IOTA does not support tokenization. IOTA does not support smart contracts as well. High speed and high scalability are two main advantages of IOTA.

The transaction referencing structure in Tangle provides data immunity. In theory, data can be traced back to the very first transaction(s) in Tangle, although the snapshotting mechanism makes this impossible.

8.8 Applications of Blockchain

Blockchain has the potential to be applied to virtually every aspect of our life. Figure 8.18 shows several of the application domains.

8.8.1 Financial Services

Financial service is no doubt the most prominent application domain of blockchain technology due to its relationship with Bitcoin. Unlike traditional financial services, blockchain enables transactions to occur in a peer-to-peer manner without involving third parties. This eliminates intermediary financial services such as banks and saves costly service fees. Blockchain also records transaction history and such records cannot be tampered. This will help with verification and in avoiding disputes.

Blockchain will also greatly speed up transaction processing and can reduce the time needed for processing to seconds, even if the transactions are cross-border, in which case processing delays can be up to several days. Blockchain-based financial services are also available to customers around the clock.

Stock trading platforms based on blockchains allow investors to purchase and sell stocks almost instantly in a secure manner. Funds created from selling stocks can be made available right after the transactions so that investors can reinvest the funds without wait times.

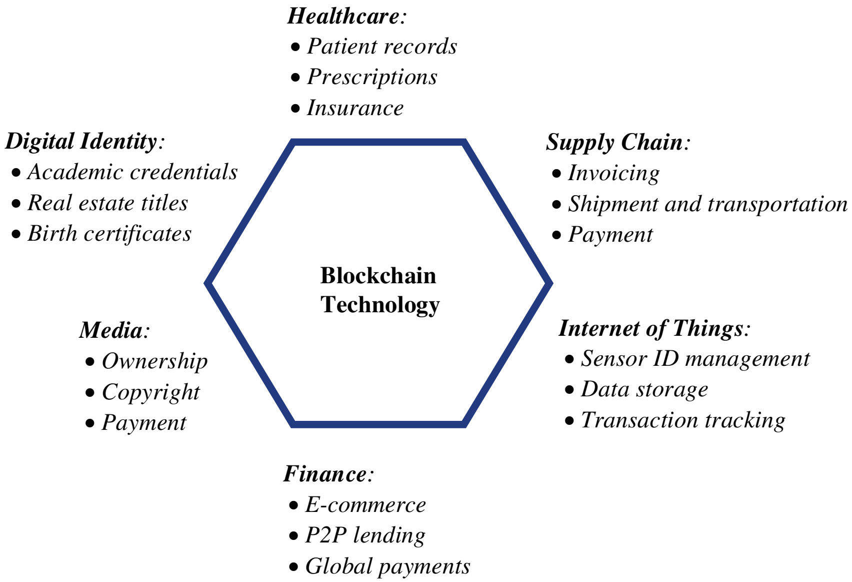

Fig. 8.18 Blockchain applications

The way business is conducted will change because of blockchain. For example, in legal practice, blockchains can be used to store wills and other inheritance records and wills will not be able to be tampered. Smart contracts can be used to set up inheritance criteria as well.

A more effective insurance industry will be in place because of blockchains. It will be very easy to verify asset owners and therefore avoid fraudulent claims. The whole insurance industry will be more effective and reliable.

Copyrighted digital contents can be better protected when blockchains are used. Ownership rights can be made transparent. Content creators will be able to receive royalties speedily.

8.8.2 Healthcare

Healthcare practices will change greatly because of blockchains. Healthcare is a complex business because it involves doctors, nurses, staff, medical service providers, insurance companies, testing labs, and pharmacies. Parties involved in healthcare are distributed in geo-location. However, they can all initiate transactions and the results of these transactions are supposed to be stored in one system and are used in an integrated manner. Currently, such transactions are stored in different systems which are inconvenient, time-consuming, and error-prone to use by stakeholders. Blockchains will change this completely. All transactions will be stored in the same ledger. This means patient medical history, test results, benefits and eligibility, insurance coverage, medication, and allergies are all available on the same blockchain. Management of healthcare systems will be more efficient because demands of medication, equipment, and other consumables can be managed by using blockchains as well.

8.8.3 Energy

The *energy industry* is another major blockchain application domain [19]. Applications can be made in electric grid, or in the oil and gas industry. When applied to the electric grid, blockchains can be used in wholesale or in peer-to-peer electricity distribution systems. When used in wholesale electricity distribution systems, blockchains can directly connect end users to the electric grid. Users can trade energy via the electric grid. No retailers will be necessary any longer so that electricity costs will be lowered. In fact, it has been envisioned that interconnected electric grids will be made available and will allow participants to buy and sell renewable energy at greatly reduced prices. Table 8.4 shows the applications of blockchain in the energy industry.

Blockchains can be specifically implemented in electricity data systems which manage fuel prices, market prices, and marginal costs. Such data can be recorded,

Table 8.4 Blockchain applications in the energy industry

In oil and gas industry	In electric grid
Gas and commodity trading	Wholesale electricity distribution
Supply and data tracking	Peer-to-peer electricity distribution
Consortiums	Electricity data management

stored, and tracked by blockchains. With blockchains, clerical errors can be avoided. Data will not be misreported or unreported. Blockchains will also allow the public to view transactions and their prices and monitor money movements.

Because privacy and trade secrets are particularly important for oil and gas companies, they are more interested in private permissioned blockchains and consortium blockchains. With these blockchains, companies can limit data access to selected parties. Oil and gas companies are looking into using blockchains for commodity trading and supply and data tracking.

The potential benefits of blockchains for oil and gas companies are increased data security, reduced time delays, and reduced data management costs.

8.8.4 Identity Management

Another application area where blockchains are well suited is *digital identity management*. This applies to both the identities of properties and human beings. When blockchains are applied to properties, the use case is called *smart property*. Tangible properties such as cars, bikes, houses, appliances, and jewelry can have their digital identities embedded when they are manufactured. These identities, together with their owners, will then be stored in blockchains. Ownership transfers can be recorded and traced through blockchains. The authenticity of the properties can be verified. Blockchains can manage intangible properties such as patents and company stock shares too.

It has been envisioned that blockchains will be used in managing the identities of human beings. Human beings can receive their identities when they are born. The same identities can be used for governmental registration, health records, motor vehicle licensing, life insurance, schooling, and employment. This will greatly simplify the currently used information system structures which are independent from each other and difficult to synchronize. Current information systems are susceptible to frauds and mishandling, whereas blockchain-based systems are secure.

One of the serious problems with current information systems is that, after users submit their personal data, they do not know exactly where the data is located, for what purpose the data will be used, who will see the data, and who will use the data. These are not only problems for the people who submitted the data, but also problems for organizations who own the databases. Online companies are able to abuse user's personal data or sell the data to advertisers. Blockchains create a so-called single point of trust and protect our privacy. Data will be encrypted and people will have control of their own data.

8.8.5 Supply Chain Management

Supply chain management is another best-suited application for blockchains. Blockchains can be used to record the entire process of transferring physical goods from the producers to the consumers. Details of farms or greenhouses, delivery trucks, warehouses, supermarkets, and retail stores, as well as the movements of physical goods between them, can all be recorded in a blockchain, along with the temperature, humidity, time, etc.

Blockchains make supply chain management more efficient. Managers can use blockchains to help with planning and avoid overstocking. Goods can be located in real time. Causes of problems can be traced back to its origin.

8.8.6 Other Applications

Blockchains will bring revolutions to *voting and elections*. Digital voting will be more secure than ever before. Votes will become transparent and immutable. Voters will be able to find out if their votes have been counted. Election data will not be able to be tampered. Millions of dollars will be saved in running elections.

Another industry that will benefit from blockchains is *real estate*. Property titles and their transfers can be stored in blockchains. The records are transparent to the public and are permanent.

Student records and certificates can be stored in blockchains. Verification of credentials can be made instant and will be reliable.

Driver licenses, violations, and accident records can be stored in blockchains as well. The whole process of vehicle ownership and policy verification will be efficient without errors.

Blockchains can be used just for *data backup.* Data backed up are immune to tampering. However, current cloud-based data storage systems are not immune to hackers.

8.9 Other Aspects of DLTs

8.9.1 Scalability and Other Practical Considerations

When it comes to adopting DLTs in an organization, practical considerations become important. These considerations are to do with common-sense parameters: memory size and speed. For DLTs and blockchains, these parameters are transaction size, block size, and transactions per second (TPS). Before each DLT technology is examined, some information about VisaNet (the credit card processing system) is useful. It is reported that VisaNet can handle an average of 150 million transactions per day [20]. This is equivalent to about 1736 transactions per second on average.

8.9.1.1 Bitcoin

Bitcoin generates 1 block every 10 minutes. The size of the block is 1 megabyte (1,048,576 bytes). This block is broadcasted to the Bitcoin network which had 10,198 nodes on January 17, 2019. The Karlsruhe Institute of Technology reported that, on January 17, 2019, it took 13,989.42 milliseconds or approximately 14 seconds to propagate the block to 99% of the nodes on the Bitcoin network [21]. This means that the block propagation time of Bitcoin is about 14 seconds.

The average Bitcoin transaction size is 380.04 bytes on January 17, 2019 [21], although on May 12, 2019, it was 352.23 bytes per transaction on average [22]. Among the 380.04 bytes, 346 bytes are overheads for the transaction, and only 34 bytes are real data for the transaction.

Therefore, the average number of transactions per block in Bitcoin is 1,048,576/380.04 = 2759.12. This gives 2759.12/(60 × 10) = 4.548 transactions per second. This is far less than the 1736 transactions per second of VisaNet. In summary, Bitcoin has the following parameters as shown in Table 8.5.

8.9.1.2 Hyperledger Fabric

The performance of Hyperledger Fabric [23] is a function of the number of endorsing peers, number of channels, endorsement policy, ordering service configuration

Table 8.5 Bitcoin performance data

Block generation time	Block size	Transactions per block	Block propagation time	Transaction size	Actual data in transaction	Overhead of transaction	Transactions per second (TPS)
10 minutes	1 MB	2759	14 seconds (99% of nodes)	380.04 bytes	34 bytes	346 bytes	4.548

(i.e., block size and frequency), number of organizations, and ledger database used. It is also to do with execution complexity of chaincode or smart contracts, transaction sizes, use of mutual TLS security in network traffic, number of vCPUs, memory allocation, disk type and speed, and network speed. Furthermore, it is to do with data centers, CPU speed, and crypto acceleration. An experiment conducted by [24] tested Hyperledger Fabric 1.3.0 in a single Kubernetes cluster running on the IBM Container Service. The worker nodes were configured as 4vCPU and 16Gb memory with SSDs. A two-organization cluster executed on a single channel and 2, 4, and 8 endorsers were used respectively. The corresponding throughput and average latency are shown in Table 8.6. In this table, TPS stands for "transactions per second." The table indicates that, with 2 endorsers, the tested system carried out 785.58 transactions per second, and it took 715 milliseconds for 95% of the nodes to commit a transaction. Whereas when 8 endorsers were employed, the tested system was able to finish 1265.5 transactions per second, and it took only 686 milliseconds for a transaction to be committed by 95% of the nodes in the system.

It should be noted that Samsung SDS revealed that it had developed an accelerator software to speed up Hyperledger Fabric transactions to 3500 TPS, with experiments succeeded in achieving 20,000 TPS [25].

8.9.1.3 Ethereum

Unlike Bitcoin, Ethereum does not have a fixed block size. Instead, Ethereum has a gas limit for each block which determines how many transactions can fit in a block. The block generation time of Ethereum has achieved about 13 seconds [26]. The TPS of Ethereum is about 15 transactions per second [27]. The history of the Ethereum block size can be found in [28]. Ethereum performance data is shown in Table 8.7.

8.9.1.4 IOTA

IOTA does not have blocks. A transaction in IOTA consists of 2673 trytes [29]. Using the IOTA converter [30], 2673 trytes can be converted to 1589 bytes. IOTA can execute 500–800 transactions per second on average and it will be even faster

Table 8.6 Hyperledger Fabric test results

Number of endorsers		2		4		8	
TPS	95% (ms)	785.58	715	948.2	667	1265.5	686

Table 8.7 Ethereum performance data

Block generation time	Block size	Transactions per second (TPS)
Around 13 seconds	Variable	15

Table 8.8 Ethereum performance data

Transaction size	Transactions per second (TPS)
1.598 bytes	500–800 and above

when more users have participated [31]. IOTA performance data can be found in Table 8.8.

8.9.1.5 Scalability of DLTs

Gartner defines scalability as *the measure of a system's ability to increase or decrease in performance and cost in response to changes in application and system processing demands*. In other words, if a system is scalable, it should be able to grow in size and performance if user demand increases.

Bitcoin does not scale well. In its original design, the block generation time is fixed at 10 minutes, the block size is fixed at 1 megabyte, and the TPS is fixed at 4.5. If more and more users participate, the wait time for a transaction to go through is not acceptable. Technologies such as Segwit (Segregated Witness) have been developed to mitigate the scalability problem of Bitcoin.

IOTA is scalable. This is due to the fact that IOTA does not store transactions in blocks which have limited size. In IOTA transactions approve other transactions. Therefore, the more transactions IOTA has, the more transactions it can approve simultaneously. IOTA performance will increase with the increase of users.

8.9.2 Token and Token Economics

Technically speaking, a token in a blockchain represents a programmable currency unit embedded in a blockchain and is part of smart contract logic. In simple non-technical terms, a token is a kind of private digital currency. A more comprehensive definition is given by Mougayar [32], where a token is defined as *a unit of value that an organization creates to self-govern its business model, and empower its users to interact with its products, while facilitating the distribution and sharing of rewards and benefits to all of its stakeholders*.

Tokens can be used to grant rights to use a product, or the rights to vote. Tokens can also be used as a unit for exchanging values in a blockchain ecosystem. Tokens can be incentives earned by doing useful work and can be spent when using a service or product. Tokens can serve as a payment method. Tokens can be distributed in ICOs (Initial Coin Offerings).

Essentially, tokens help build self-sustainable mini-economies in distributed autonomous organizations (DAOs) based on blockchains. This is interestingly termed as "tokenomics" or "cryptoeconomics."

8.10 Vulnerabilities of Blockchain

Although blockchain technology provides numerous advantages and application potentials, it is not perfect. It is important to be aware of its weaknesses. Potential attacks can occur on several aspects of blockchain technology. Systems based on blockchain technology can even be used to commit crimes.

The first vulnerability of blockchain technology is originated from its *consensus mechanism*, which is susceptible to a 51% attack [33]. Specifically, in a Power-of-Work-based blockchain network, if the computational power of a single miner node exceeds 50% of the total power of the entire blockchain network, then the entire blockchain could potentially be controlled by that attacker. In a Power-of-Stake-based blockchain network, the 51% attack can also occur if the number of stakes owned by a single node is more than 50% of that of the total blockchain network. The attackers of a 51% attack are able to reverse transactions, conduct double-spending, exclude transactions, reorder transactions, cause problems for operations for normal transaction confirmation, and stop the mining operations of other mining nodes.

Sybil attack is also a vulnerability of blockchain which takes advantage of the fact that public blockchain networks have no centrally trusted nodes and every transaction is sent to a number of other nodes for processing. A Sybil attack is initiated by assigning a number of identifiers to the same node. During a Sybil attack, the attacker is able to outvote honest nodes and takes control of the network. Therefore, the consequence of a Sybil attack is equivalent to a 51% attack.

Private keys are another source of vulnerabilities in a blockchain network. A private key is the identity of a user. It is used to sign transactions and verify asset owners. Private keys are also used in transaction validation and candidate block verification. However, a legitimate user's private key can get lost. If this happens, there is no way to recover the private key. The legitimate user will not be able to access his/her account on the blockchain network anymore and will therefore lose the assets he or she owns. If a private key is stolen by a criminal, the legitimate user's blockchain account can get tampered. Whatever damage the criminal does is difficult to track, repair, and recover because there are no centralized third-party trusted institutions to seek assistance from.

Although it is commonly known that, by introducing consensus algorithms, a blockchain network can prevent the *double-spending attack*, as claimed by the Bitcoin paper [1], it is still possible for double-spending to occur in a blockchain network. It is misleading to believe that double-spending is fully eliminated by the consensus mechanism during validation. Among all blockchains, the Power-of-Work-based blockchain network is especially vulnerable, as the attacker can exploit the time interval between the initiation and confirmation of two transactions to quickly launch a double-spending attack. Double-spending refers to the fact that a malicious user spends the same cryptocurrency for multiple transactions. Knowing it takes time to mine a block and reach consensus, the attacker could launch a race attack involving two consecutive transactions. Before the second transaction is invalidated, it is possible that the attacker has already received the output of the first transaction. This results in a double-spending.

As a long-term security problem for the Internet, the *distributed denial-of-service* (DDoS) *attack* is still a threat to blockchain networks. DDoS attacks create a huge amount of traffic on blockchain networks so that valid transactions cannot be processed, giving opportunities for invalid transactions to become successful.

On the other hand, blockchain networks can be used by criminals to *commit crimes*. One such example is ransomware. A typical ransomware is sent out as an email attachment. If the email receiver clicks the attachment, the ransomware starts running as a background process on the receiver's computer system. What it does is that it encrypts the files in the receiver's system so that the victim loses access to the contents of the files. The ransomware demands the receiver to pay funds to a blockchain account of the attacker within a given time frame. Otherwise, there will be no way to restore the encrypted files forever.

Blockchains can also be used by criminals to run underground markets. Bitcoins are used as the currency and hidden services for such markets. Criminals use underground markets to sell drugs, weapons, and other controlled items. Due to blockchain's anonymous nature, it is difficult to track down the sellers and deals.

8.11 Summary

This chapter discusses distributed ledger technologies (DLTs) which include blockchain and directed acyclic graphs. The chapter discusses the benefits of DLTs when they are adopted by current information systems. Detailed descriptions are given on how blockchain and DAG works and what the differences between blockchain and DAG are. The chapter also discusses the Internet of Things (IoT), the weaknesses of current IoT system implementations, why blockchain can help overcome the weaknesses of IoT, and how to integrate blockchain and IoT. Applications of DLTs and practical considerations of DLTs in enterprise environments are also discussed. Vulnerabilities of blockchain are described as well.

References

1. S. Nakamoto, Bitcoin: A peer-to-peer electronic cash system (2008), [Online]. Available: https://bitcoin.org/bitcoin.pdf. Accessed 23 Mar 2019
2. S. Haber, W.S. Stornetta, How to time-stamp a digital document. J. Cryptol. **3**(2), 99–111 (1991)
3. D. Chaum, Blind signatures for untraceable payments, in *Advances in Cryptology Proceedings of Crypto 82*, ed. by D. Chaum, R. L. Rivest, A. T. Sherman, (Plenum (Springer-Verlag), New York, 1983), pp. 199–203

4. T.A. Mahler, Oncologist + gold = revolution? (2018), [Online]. Available: https://medium.com/blockwhat/96-oncologist-gold-revolution-c08a8dc26880. Accessed 23 Mar 2019
5. W. Dai., b-money (1998), [Online]. Available: http://www.weidai.com/bmoney.txt. Accessed 20 May 2019
6. A. Back, Hashcash – A Denial of service counter-measure (2002), [Online]. Available: http://www.hashcash.org/hashcash.pdf. Accessed 23 Mar 2019
7. V. Vishnumurthy, S. Chandrakumar, E.G. Sirer, KARMA : A secure economic framework for peer-to-peer resource sharing (2003), [Online]. Available: https://www.cs.cornell.edu/people/egs/papers/karma.pdf. Accessed 23 Mar 2019
8. Bexam: The next generation blockchain/DAG hybrid platform, (2019), [Online]. Available: https://bexam.io/. Accessed 6 May 2019
9. Ethreum, The blockchain app platform (2019), [Online]. Available:https://www.ethereum.org/. Accessed 23 Mar 2019
10. cointelegraph.com, Proof-of-work explained (2019), [Online]. Available: https://cointelegraph.com/explained/proof-of-work-explained. Accessed 23 Mar 2019
11. S. King, S. Nadal, PPCoin: Peer-to-peer crypto-currency with proof-of-stake (2019), [Online]. Available: https://decred.org/research/king2012.pdf. Accessed 23 Mar 2019
12. B. Asolo, Delegated proof-of-stake (DPoS) Explained (2019), https://www.mycryptopedia.com/delegated-proof-stake-dpos-explained/. Accessed 23 Mar 2019
13. M. Castro, B. Liskov, Practical Byzantine Fault Tolerance, Proceedings of the Third Symposium on Operating Systems Design and Implementation, New Orleans, USA, February 1999 (2019), http://pmg.csail.mit.edu/papers/osdi99.pdf, Accessed 23 Mar 2019
14. IOTA Foundation, IOTA basics overview (2019), https://docs.iota.org/docs/iota-basics/0.1/introduction/overview. Accessed 30 Mar, 2019
15. M. Green, Hash-based signatures: An illustrated primer (2019), https://blog.cryptographyengineering.com/2018/04/07/hash-based-signatures-an-illustrated-primer/. Accessed 10 May 2019
16. E. Hop, Exploring the IOTA signing process (2019), https://medium.com/iota-demystified/exploring-the-iota-signing-process-eb142c839d7f, Accessed 10 May 2019
17. R. Minerva, A. Biru, D. Rotondi, Towards a definition of the internet of things (IoT), 27 May 2015 (2015), https://iot.ieee.org/images/files/pdf/IEEE_IoT_Towards_Definition_Internet_of_Things_Revision1_27MAY15.pdf. Accessed 30 Mar 2019
18. A. Reyna, C. Martín, J. Chen, E. Soler, M. Díaz, On Blockchain and its integration with IoT. Chall. Oppor. Futur. Gener. Comput. Syst. **88**, 173–190 (2018)
19. ConsenSys, Blockchain and the energy industry, (May 25, 2018). https://media.consensys.net/the-state-of-energy-blockchain-37268e053bbd. Accessed 30 Mar 2019
20. Visa Inc, Visa acceptance for retailers (2019), https://usa.visa.com/run-your-business/small-business-tools/retail.html. Accessed 12 May 2019
21. K. Li, The blockchain scalability problem & the race for visa-like transaction speed. (Jan 30, 2019). https://hackernoon.com/the-blockchain-scalability-problem-the-race-for-visa-like-transaction-speed-5cce48f9d44. Accessed 12 May 2019
22. Tradeblock.com (2019), https://tradeblock.com/bitcoin/historical/1h-f-tsize_per_avg-01101. Accessed 12 May 2019
23. M. Mamun, How does hyperledger fabric work? (2019), https://medium.com/coinmonks/how-does-hyperledger-fabric-works-cdb68e6066f5. April 16, 2018. Accessed 12 May 2019
24. C. Ferris, Answering your questions on hyperledger fabric performance and scale, (January 29, 2019), https://www.ibm.com/blogs/blockchain/2019/01/answering-your-questions-on-hyperledger-fabric-performance-and-scale/
25. Ledger Insights Ltd, Samsung Tech to Speed up Hyperledger Fabric (2019), https://www.ledgerinsights.com/samsung-hyperledger-fabric-speed-blockchain/. Feb. 2019. Accessed on 15 May 2019

26. J. Heal, C. Rivet, March 5, 2019, Ethereum block generation time falls following Constantinople upgrade (2019), https://finance.yahoo.com/news/ethereum-block-generation-time-falls-080011370.html. Accessed on 15 May 2019
27. How Will Ethereum Scale?, (2019), https://www.coindesk.com/information/will-ethereum-scale. Accessed on 15 May 2019
28. Ethereum Block Size historical chart, (2019), https://bitinfocharts.com/comparison/ethereum-size.html. Accessed on 15 May 2019
29. The Anatomy of a Transaction, (2019), https://domschiener.gitbooks.io/iota-guide/content/chapter1/transactions-and-bundles.html. Accessed on 15 May 2019
30. IOTA Converters, (2019), https://laurencetennant.com/iota-tools/. Accessed on 15 May 2019
31. Transaction Speed – Bitcoin, Visa, Iota, Paypal, (2019), https://steemit.com/cryptocurrency/@steemhoops99/transaction-speed-bitcoin-visa-iota-paypal. Accessed on 15 May 2019
32. W. Mougayar, Tokenomics — A business guide to token usage, utility and value (Jun 10, 2017), https://medium.com/@wmougayar/tokenomics-a-business-guide-to-token-usage-utility-and-value-b19242053416. Accessed on 15 May 2019
33. S. Sayeed, H. Marco-Gisbert, On the Effectiveness of Blockchain against Cryptocurrency Attacks, UBICOMM 2018 : The Twelfth International Conference on Mobile Ubiquitous Computing, Systems, Services and Technologies, pp.9–14 (2018)