Chapter 3
Cryptocurrencies

A cryptocurrency is a digital asset designed to serve as a medium of exchange that should be an alternative to the classic fiat currency. The idea of bringing money from the physical to the digital realm has been investigated since the 1980s, with many attempts to create digital cash systems. Over the years, several researchers have tried to implement an electronic currency disconnected from the banking system, but none of these projects have been successful until 2008. In that year, an anonymous researcher (or a group of people) known under the pseudonym of Satoshi Nakamoto published a white paper [167] that describes a peer-to-peer electronic cash system—called Bitcoin—completely independent from the traditional banking system.

Making use of existing cryptographic technologies, such as asymmetric cryptography and hash functions, Nakamoto introduces a new technology—called blockchain—which will be crucial in the development not only of cryptocurrencies but of many other applications.

Definitions

Blockchain The term blockchain refers to a technology that implements an immutable and distributed database, consisting of a chain of blocks linked together. Each block, in addition to the stored data, contains the hash of its predecessor in the chain. This implies that if block n is modified, it will be necessary to modify all the other blocks of the chain accordingly, starting from block $n + 1$ until the end of the chain.

Cryptocurrency A cryptocurrency is a decentralized electronic cash system designed to work much like a standard currency, allowing users to make virtual payments, free of a trusted central authority. Cryptocurrencies leverage

(continued)

R. Di Pietro et al., *New Dimensions of Information Warfare*, Advances in Information Security 84, https://doi.org/10.1007/978-3-030-60618-3_3

cryptographic functions to ensure legitimate, unique, and publicly verifiable transactions.

Bitcoin uses the blockchain to store transactions among users, which are then permanently saved in a decentralized and immutable way. New transactions are collected and verified by specific members of the network, called miners. New blocks containing only approved transactions are validated by all participants in the protocol and added to the blockchain through a consensus mechanism, called Proof of Work (PoW), that ensures the security of the network. In this way, Bitcoin implements the first decentralized and open-source payment system that does not need a trusted third party. In fact, nobody controls or owns the Bitcoin network. Anyone can join it and download the transaction ledger—which is public—to verify and validate old and new transactions.

With the introduction of the blockchain, Nakamoto not only created what is still the most famous and used cryptocurrency but also allowed the implementation of other ones. After Bitcoin indeed, numerous other cryptocurrencies were born, some of which are very similar to it while others feature different implementations of the consensus mechanism, but all characterized by the use of the blockchain. This new technology has created novel scenarios for the state economy, enabling use cases that could potentially revolutionize the world economic landscape. Although it is currently difficult to indirectly attack a state's economy by directly hitting a cryptocurrency, it may be different in the near future. The establishment of a state cryptocurrency, as well as other strategic blockchain-based state-sponsored applications, appears to be a very likely scenario [168].

In this chapter, we examine the reasons that could push the development of a state cryptocurrency, citing the nations that have already started planning it. We then describe three scenarios in which a cryptocurrency can be attacked, based on the characteristics of the currently existing cryptocurrencies and the real attacks they have experienced. Finally, we investigate the technical features that a state cryptocurrency should have to be accepted by the population and to restrict the possibilities of attacks, both internal and external.

3.1 State-Sponsored Cryptocurrency

Nowadays, more and more nations are thinking about establishing a state cryptocurrency that will support or replace the standard fiat money. This kind of scenario, on the one hand, introduces several advantages of practical nature, such as no longer having to print physical banknotes, no longer need banking institutions that keep track of balances and transactions, the faster and (supposedly) more secure

transactions, and more. On the other hand, it could expose the nation's economy to a new series of cybersecurity threats. Indeed, the classical physical currency is vulnerable to several indirect attacks that mainly aim to its devaluation, such as speculative attacks. However, other kinds of attacks, such as denial of service, are very difficult if not outright impossible, due to the physical nature of the standard currency. In fact, an attacker could target the electronic systems that allow virtual transaction, causing a temporary block of this service, but there are no ways to stop transactions with cash payments. A cryptocurrency instead, as a virtual asset, is exposed to direct attacks with consequences ranging from blocking the system for a short time to its total destruction. In the first case, malicious entities could prevent legitimate users to join the network or isolate the peers that validate transactions, causing a total blockage of the network. In this eventuality, no transactions are possible in the network because users are not able to create them or the system is not able to receive them. If the attacked cryptocurrency is the only currency available in the state, citizens will no longer be able to make transactions of any kind. Consequently, the sale of goods and services among citizens would fall into anarchy, being possible only through the adoption of antiquated forms of exchange, such as barter. Table 3.1 reports the countries that—at the time of writing—have started to design, develop, or use a national cryptocurrency or have decided to support one.

3.2 Scenario 1: Trust in Maths

This scenario takes into account a cryptocurrency that relies on the intractability of certain mathematical problems to ensure the security of its protocol. To guarantee some properties like confidentiality, integrity, and availability—fundamental for the security of every communication—the cryptocurrency protocol uses different cryptographic techniques based on mathematical problems. Users trust the system because of the difficulty of the crypto-challenges derived from the aforementioned math problems, recognized as computationally hard to solve by the worldwide scientific community.

As an example, Bitcoin relies on the Elliptic Curve Digital Signature Algorithm (ECDSA) to ensure that funds can only be spent by their rightful owners. In fact, each user account is composed of a pair of addresses—i.e., the cryptographic hashes of an ECDSA private key and the derived ECDSA public key. The owner of the private key is the owner of the account and is, therefore, the only person who can spend the money contained in it. The public key is shared with the community and used to receive payments. To transfer Bitcoins from one account to another, the sender creates a new transaction addressed to the recipient's public address and signs it with the sender's private key. For this reason, the core of the security and consistency of the Bitcoin network is the security of ECDSA private keys. ECDSA relies on the math properties of the cyclic groups of elliptic curves over finite fields and on the well-known difficulty of the Elliptic Curve Discrete Logarithm Problem (ECDLP) that points the following: given an elliptic curve over a finite field and

Table 3.1 Countries that already have or are issuing national or regional cryptocurrencies

Region	State	Type	Name	Platform	Status
Caribbean	Anguilla	State sponsored			Legislation to regulate initial offerings of cryptocurrency
	Antigua and Barbuda	State supported		Etherium	No legislation, proof-of-concept
	Dominica	State sponsored			Participates in a pilot program to develop a digital Eastern Caribbean Dollar
	Grenada	State sponsored			Participates in a pilot program to develop a digital Eastern Caribbean Dollar
	Montserrat	State sponsored			Participates in a pilot program to develop a digital Eastern Caribbean Dollar
America	Venezuela	State sponsored	Petro	NEM	Implemented
Europe	Ireland	State supported	Irishcoin	MapleCoin	Implemented, currently in use
	Lithuania	State sponsored	LBChain		Final testing phase
Asia	China	State sponsored			Developing phase

two points G and H on the curve, find the scalar k such that $H = kG$. These properties ensure that deriving a private key from the corresponding public one is computationally infeasible. In the case of Bitcoin, therefore, users trust the protocol because it is based on these known properties and accept that calculating their private key in a reasonable time is an unsolvable problem for a possible attacker. However, there is no guarantee that anyone, at any time, can solve or simplify the mathematical problem from which Bitcoin's transaction security derives. In this case, an adversary could derive any user's private key from its public one, managing to spend the victim's funds without authorization. In such an event, the victim would have no way to get his funds back since the transactions are irreversible in the Bitcoin network. Furthermore, if a cryptographic protocol is not implemented correctly, it could be vulnerable even if the math behind it is still correct and the related problem difficult to solve.

In this scenario, we investigate the possibilities that an attacker can simplify the math on which a cryptocurrency bases its security and consistency. Also, we analyze possible implementation errors that could be exploited by an adversary to attack the protocol.

Definitions

Elliptic Curve Discrete Logarithm Problem (ECDLP) It is a mathematical problem based on the fact that it is infeasible to find the discrete logarithm of a random elliptic curve element with respect to a publicly known base point. The ECDLP is a special case of the discrete logarithm problem. Its apparent intractability is the basis for the security of elliptic curve cryptography. Unlike the finite field discrete logarithm problem, there are no general-purpose subexponential algorithms for solving ECDLP. This implies that it is possible to choose smaller fields than those needed for cryptographic systems based on the finite field DLP, which results in keys of smaller size.

Elliptic Curve Digital Signature Algorithm (ECDSA) The Elliptic Curve Digital Signature Algorithm is used to create the digital signature of a file, or any other digital data, to ensure its authenticity. Essentially, it is a version of the Digital Signature Algorithm (DSA) that uses the elliptic curve. It was accepted as an ANSI standard in 1999 and as an IEEE and NIST standard in 2000. Moreover, it became an ISO standard in 1998.

3.2.1 Threat: Collapse of the Cryptocurrency Foundation

Commonly used encryption methods are believed to be secure because they have been studied for many years by the most important experts in the field. Unfortunately, this doesn't mean that they are provably secure. Someone might find, at any time, a way to reduce their complexity. The two cryptographic schemes, on which the security of almost every communication protocol depends, are based on the accepted difficulty of certain arithmetic operations. In the case of RSA, it is finding the two numbers that have been multiplied together to get the modulo. In the case of Elliptic Curve Cryptography (ECC), given two points P and Q, find out the integer x that satisfies the equation $Q = xP$. Here, "difficult" means that the time spent to perform these operations exceeds the useful time of the secret to be violated. This does not apply if a shortcut is found. In this case, an attacker could be able to get around the difficulties of solving the problem, thus reaching the solution more quickly. For example, different algorithms have been developed for computing the discrete logarithm problem on elliptic curve (on which ECDSA is based) trying to optimize the resolution phase. The two most efficient ones are the "baby-step, giant-step" algorithm and Pollard's rho method. Both mechanisms, however, as well as their various subsequent optimizations, do not currently allow attacking the ECDS protocol in a reasonable time. However, there is no way to know that better shortcuts are not going to be discovered in the future.

As an example of a similar situation, we can cite the story concerning the famous Fermat's Last Theorem (sometimes called Fermat's conjecture). Formulated in 1637, it states that the equation $x^n + y^n = z^n$ has no solution in integer for $n \geq 3$. Although this conjecture is intuitively true, nobody has been able to provide mathematical proof until a British professor, Andrew Wiles, released the first successful demonstration in 1994. Just as Fermat's conjecture, which remained unsolved for three centuries before being proved mathematically, the same could happen at any time for other mathematical problems, such as ECC.

In this scenario, the major threat is represented by an adversary that reduces the mathematical complexity of the problem on which the cryptocurrency relies on, becoming able to solve it in an optimized way. This knowledge makes the adversary able to control the cryptocurrency network by exploiting its capabilities that other peers do not have. The same result could happen if an adversary discovered a 0-day vulnerability in the implementation of one cryptographic function used by the cryptocurrency's protocol and developed a methodology to exploit it.

3.2.2 Attacks and Countermeasures

The complexity of the mathematical properties that support the correctness of the ECDSA protocol has never been reduced. However, several examples of flaws in the protocol implementation that have led to serious security incidents can be found in the literature.

One of the most critical phases of the ECDSA protocol implementation is the choice of the elliptic curve and its domain parameters. These parameters, chosen by the developer at each implementation, are fundamental for the robustness of the protocol. The standard ones (widely accepted as safe) are defined by the scientific community, but every developer is free to use other customized ones. In addition, some parameters must be chosen randomly; otherwise the security of the protocol may be compromised. In 2010, a group of hackers called *fail0verflow* discovered a serious flaw in the implementation of the ECDSA algorithm used by Sony to calculate its own set of keys used to digitally sign original software for the game console PlayStation 3 [169]. The attackers were able to recover the entire private key, then using it to distribute counterfeit software. The attack works only against the algorithm used by Sony, as it used a static parameter (the variable k) instead of one randomly selected at every execution, making the private key solvable by analyzing few digitally signed files.

In some cases, the protocol could be vulnerable even if some parameters are correctly chosen randomly, as recommended by the scientific community. In fact, often the problem lies in how random numbers are generated. In [170], the authors examined the quality of the random number generated by some common Java libraries such as Android Pseudo-random Number Generator (PRNG), Apache Harmony, GNU Classpath, OpenJDK, and Bouncy Castle. They found multiple weaknesses on entropy collector components, with different degrees of severity and

probability of occurrence. In particular, they showed that the overall entropy of the Android PRNG can be reduced to only 64 bits. This weakness was exploited in 2013 to steal balances from Bitcoin users' wallets generated by any Android app. The faulty component seems to be the Java class SecureRandom (used by the vulnerable wallets) that can generate collisions on the produced random numbers. The problem is that the ECDSA algorithm requires that the random number used to sign a private key is only ever used once. If the random generated number is used twice, the private key is recoverable.

In [171], the authors identified a timing attack vulnerability in OpenSSL's implementation of Montgomery's ladder for scalar multiplication of points on elliptic curves over binary fields. This vulnerability allows a full key recovery attack against a TLS server that authenticates with ECDSA signatures.

Definitions

Timing Attack When the execution time of a cryptographic device or function is variable, it may leak information on the secret parameters applied. A careful analysis that includes precise time measurements could allow the reconstruction of the system key involved. The timing attack is an attack based on the leakage of information of secret parameters through variations in the running times of a cryptographic device [172].

Hash functions are another pillar of the most important cryptocurrencies available in the global financial landscape. Several attacks against hash function implementations are discussed in the literature as well as against the compression function they used. The most important are the Chabaud et al. attack against the SHA-0 algorithm [173] and the hash function attack techniques introduced by H. Dobbertin against the MD5 algorithm [174–177]. In [173], the authors presented a methodology to find collisions in SHA-0 by looking for some kind of characteristic masks that can be added to input words with non-trivial probability of unchanging the output of the compression function. They obtained a theoretical attack on the compression function used by SHA-0 with complexity 2^{61}. These techniques are not applicable against SHA-256 and SHA-512—used by Bitcoin and by other major cryptocurrencies—as investigated by several researchers in [178–180].

Definitions

Hash Function A hash function is a computational method that, taking data of an indeterminate size as input (the key), returns a fixed-size string as output (the hash value). A cryptographic hash function uses one-way math functions

(continued)

to generate a hash value from a given input. The fundamental properties of these functions are (i) *pre-image resistance* and (ii) *collision resistance*. The *pre-image resistance* ensures that the function is not invertible. This means that it must be difficult to find the key starting from the hash value. The *collision resistance*, instead, ensures that is difficult to find the same hash value for two different keys. Cryptographic hash functions can be used for many different problems, ranging from integrity and authenticity to pseudo-random number generation and key derivation.

Compression Function A compression function is a one-way function that, taking two fixed-length inputs, produces a fixed-length output. Since it respects the pre-image resistance property, a compression function differs from conventional compression algorithms, which can instead be inverted. Usually, a hash function is defined by repeated applications of a compression function, until the whole message has been processed.

3.2.3 Open Issues

As described in the previous sections, the known attacks against cryptographic functions that are used by cryptocurrencies only concern weaknesses in their implementation and choice of their security parameters. For this reason, the risks associated with the use of these technologies cannot be solved or limited by individual security methodologies (such as the use of stack canary to prevent buffer overflow), but rather by a collection of best practices.

Above all, the use of open-source implementations of security protocol ensures that the code has been reviewed by several independent experts, decreasing the risk of bugs or other errors escaped by developers, which could potentially create security problems. For example, the use of a proprietary implementation of the digital signature algorithm used by Sony caused, as described above, the hacking of the system due to a parameter defined static rather than chosen randomly at each execution. In addition to using open-source software modules, the need to verify the correctness of the whole source code before releasing a software application is crucial. There are several ways to validate the correctness of the software. The most common way to check if a program works as expected is to test it. In this regard, developers submit their program to a wide range of inputs, to check if it behaves as designed. This testing methodology ensures that the software behaves correctly in most cases, which is sufficient for most applications. However, it is impossible to guarantee that the software tested in this way always works correctly, because it is impossible to test it with every conceivable input. Even when testing an algorithm with a very large set of inputs, there is always a small probability that it will fail under some unusual conditions, leading to a security vulnerability. These

possible malfunctions can also be very small and difficult to find, such as memory management errors or input validation errors that can cause a buffer overflow and, even if apparently harmless, can become the weak point that could jeopardize the security of the whole application. One of the most promising approaches to solve, or at least mitigate, this problem is called formal verification. Unlike regular software, written informally and validated mainly based on its behavior, formally verified software looks more like a mathematical proof, where each statement follows logically from the preceding one. This methodology, therefore, allows validating a software with the same certainty with which mathematicians prove a theorem.[1]

The formal verification of the software dates back to the 1970s, from an idea by Edsger Dijkstra. It includes two different domains, formal specification, and formal verification. The formal specification is a methodology to describe precisely and unambiguously what a software should do. The formal verification, instead, proves beyond any doubt that a software meets its specifications and works correctly. The problem with this technology is the difficulty of application, especially regarding formal specifications. In fact, describing in plain terms what a program should do is quite easy. However, expressing the same concept using formal language that a machine can understand and execute is much more complex. As an example, in the context of cryptographic protocol implementation, it is easy to say, at a high level, that the program will never leak the private key used to digitally sign documents, but converting this idea into a mathematical definition is not trivial. This problem particularly affects cryptographic protocols, given the complex nature of the cryptographic algorithms on which they are based. In fact, it is often very difficult to design cryptographic protocols without any weakness. Even relatively simple protocols have turned out to be vulnerable. In many cases, flaws were discovered after some time of publication or even implementation. In general, it is not trivial to see whether a cryptographic protocol is secure simply by looking at it. Even in a simple protocol, the flaws can be very subtle. This has also been demonstrated in many examples in the literature on protocols that have been published, considered sound, and then discovered to be faulty. For this reason, progress has been made in recent years in developing formal methods for the design and validation of cryptographic protocols. Specifically, two types of techniques have been applied to solve this problem, also used in the analysis of conventional communication protocols. The first uses the logic of knowledge and beliefs to shape the behaviors that evolve throughout a protocol. The second one, instead, models the protocol as an interaction between a set of state machines and attempts to demonstrate a safe protocol by specifying unsafe states and attempting to prove them unreachable, mainly by using exhaustive backward state research.

One of the most used tools in this area, called the NRL protocol analyzer, was developed by the US Naval Research Laboratory. It can be used to demonstrate the security properties of cryptographic protocols and to identify security flaws.

[1] https://www.wired.com/2016/09/computer-scientists-close-perfect-hack-proof-code (Last checked August 2020).

The NRL protocol analyzer managed to do both. In particular, it has been used to find previously unknown flaws in the Simmons selective transmission protocol and the Burns-Mitchell resource sharing protocol. Moreover, it has been used also to discover some hidden assumptions in the Neuman-Stubblebine reauthentication protocol and the Aziz-Diffie wireless communication protocol [181].

Resources

NRL Protocol Analyzer A complete presentation of how the analyzer works can be found in [182], together with a description of the basic functions used for the analysis of cryptographic protocols and a working example.

However, both the NRL protocol analyzer and other similar tools have been designed to analyze cryptographic protocols most commonly used for cryptographic key authentication. This was the usual application for cryptographic protocols in the past and remains one of the most common use cases. Subsequently, cryptographic protocols were applied to new problems, including financial transactions and key negotiations. These new applications bring new security problems that put new limits on existing tools, making them less and less effective and reliable. New research efforts are needed to identify new challenges and adapt existing tools accordingly, so that they can also be used successfully in new use case scenarios. The authors investigated in [183] how far current tools could be pushed to analyze complex protocols that must meet new types of requirements and also to find out where they need to be improved. They describe six different emerging areas in the application of formal method to cryptographic protocols, highlighting challenges where more research efforts are needed: open-ended protocols, denial of service, anonymous communication, high fidelity, composability, and negotiation of complex data structures.

3.3 Scenario 2: Trust in the Computational Power

This scenario takes into account a cryptocurrency that relies on the computational power for its security. Some cryptographic protocols, computationally hard to compromise, are used to manage different security features, such as the validity and legitimacy of transactions. Any security problem related to these protocols inevitably reflects on the security of the cryptocurrency and its users, undermining the stability of the system. The most famous example of such a cryptocurrency is Bitcoin. Indeed, its protocol heavily relies on cryptographic functions to secure its network. Public-key cryptography, for example, is used to guarantee that funds can only be spent by their rightful owners. Also, a PoW schema based on a computational puzzle is used to form new blocks and reach consensus among participants.

In general, a PoW is a piece of data that satisfies some requirements, with the peculiarity of being difficult to produce but easy to verify. Bitcoin adopts the so-called HashCash algorithm to manage the block generation. This mechanism requires each miner to bring together well-formed transactions issued by users, creating a new block. The PoW then consists of finding a value, called nonce, to be inserted in the new block such that its hash is lower than a certain target value. The features of the hash function make this task particularly difficult, and, consequently, it is not possible to know who will be the creator of the next block. The difficulty of the PoW can be varied as desired simply by changing the target value. In this way, the creation rate of the new blocks can be changed.

The PoW mechanism plays a fundamental role in system stability, ensuring the correct functioning of the network and preventing various theoretical attacks through a decentralized protocol.

In such a cryptocurrency, users trust the system because of the following two considerations:

1. The cryptographic protocols used to protect the network are well known and formally proven secure. In particular, it has been proved that, at least with the technologies currently available, it is not possible to break their security or circumvent the computational work needed to solve their crypto-challenges.
2. It is difficult for a single entity to have 51% of the whole computational power available in the entire network. This assumption is highly controversial and depends strictly on the characteristics of the system in question. Generally, it is considered valid for mature systems, with a solid and stable community behind it. In young systems, instead, with a community still in the process of settling, this assumption must be considered to be not always valid.

In this scenario, we analyze the vulnerabilities of this kind of cryptocurrency, focusing on two threats in particular: (i) the emergence of new technologies and (ii) the collusion among users. In our analysis, we discussed several real-world examples of the most important cryptocurrencies available in the market. Bitcoin, as the older and most mature one, is theoretically less exposed to such problems but cannot be considered immune.

3.3.1 Threat: New Technologies

In this scenario, one of the main concerns is an attacker equipped with an unexpected computational power derived from a new technology that did not exist at the time of the design and implementation of the cryptocurrency in question. Possible threats include a new generation of hardware that could be used by an adversary to violate the cryptographic protocols used by the network, gaining an illicit advantage over other users. An example taken from the past is the ASIC hardware. The release on the market of this new technology has favored the miners who have immediately used it, compared to those who kept using general-purpose devices.

In this scenario, if the new hardware is immediately made available on the global market, it is unlikely that a single entity will be favored. Being publicly available, the new technology would be used by many distinct and presumably geographically distributed miners, balancing the benefits within the network. This immediate distribution of the new performing resource ensures that the cryptocurrency does not suffer from security issues. In the future, this scenario could repeat itself, triggered, for example, by the development of quantum computers. In this case, however, distribution on the market could be much more complex, due to the high design and development costs of this innovative technology. The company that first obtains a functioning platform may, therefore, be the only one to benefit from it, obtaining hegemony over all the systems vulnerable to its tremendous computational capacity.

Even if the research in this field is still in its infancy, quantum computing promises to efficiently solve problems which are not practically feasible on classical computers. With its huge computational power, a quantum computer could be used to attack cryptocurrency networks whose security is based on cryptographic challenges that require a certain amount of computational power to be solved. By leveraging the advantage over traditional CPUs, an attacker equipped with a quantum computer could easily solve cryptographic schemes in a relatively short time, posing serious security problems to any system based on these mechanisms.

🪄 Definitions

Quantum Computing Classical computers perform logical operations and store data relying on the definite position of individual bits, represented as binary states 0 and 1. Quantum computing, instead, makes use of quantum mechanical phenomena to manipulate and store data, acquiring the potential to process exponentially more data compared to classical computers.

Quantum Supremacy Proposed by John Preskill in 2012, quantum supremacy describes the point where quantum computers can perform a task that no classical computer can feasibly solve, regardless of whether those tasks are useful.

ASIC Hardware Application-specific integrated circuit refers to a device specifically designed for the sole purpose of mining cryptocurrencies. Each ASIC device is customized to solve a particular PoW. As a result, each device is capable of mining only a specific cryptocurrency. This specialization, in terms of both hardware and software, offers ASIC hardware a huge advantage in mining activities compared to general-purpose hardware.

State of the Art

Given that research on quantum computers, although reaching increasingly important milestones, is still in its early stages, the threat it will bring to modern cryptography is still perceived as a remote possibility. However, the scientific community has long started to wonder about the possible impact, preparing the current systems for migration to the post-quantum era.

The potential danger posed to IT security by quantum computing was first established in 1994. That year saw the publication of a quantum computer algorithm [184] by the US mathematician and computer scientist Peter W. Shor. He demonstrated that encryption techniques that were previously assumed secure could be broken in a matter of seconds by factorization or reducing a number into its constituent factors. To do so, the Shor algorithm used the computing power of quantum computers [185].

Several components of a cryptocurrency architecture could be affected by this threat, while some others are immune. Several works demonstrated that quantum computers are capable of solving complex problems unfeasible for classic computers only by using algorithms that exploit the power of quantum parallelism. For example, a quantum computer cannot be faster than a standard one in multiplications [186]. As an example, quantum computers could not be efficiently used to compute the pre-image of a hash function or to generate a collision. For this reason, the hash-based puzzles used by several cryptocurrencies to implement their PoW [187] can be safely used in the post-quantum era as long as they leverage a hash function that provides an output with an adequate length, such as SHA-2 or SHA-3 [186]. On the contrary, quantum computers could be used to efficiently solve some problems underlying the asymmetric cryptography, such as the large prime integer factorization and the discrete logarithm problem, used by several cryptocurrencies to secure wallets—that is, to ensure that funds can only be spent by their rightful owners. However, these eventualities, at the moment, remain purely theoretical. At the time of writing, there is still no real implementation of quantum devices able to run such algorithms. In October 2019, Google claimed to have reached quantum supremacy. A quantum processor with 53 qubits performed a task in just over 3 minutes that, according to Google's calculations, would have taken the world's largest supercomputer 10,000 years, or 1.5 billion times more [188]. This claim has been questioned by IBM, according to which an ideal simulation of the same task can be performed on a classical computer in 2.5 days, and with much greater fidelity.[2] However, the Google's experiment was a first proof of concept. The next step is to build quantum computers with enough qubits to solve useful problems. According to estimates by Google researchers, the methodology they are following requires a number between 100 and 1000 qubits to achieve this goal, but nobody is sure. Once this milestone is reached, an even greater evolution is needed to pose

[2] https://www.ibm.com/blogs/research/2019/10/on-quantum-supremacy (Last checked August 2020).

a danger to modern cryptography. In fact, it is believed that millions of qubits are needed to break the current cryptographic schemes; it may take decades to reach that point.

3.3.2 Threat: Collusion Among Miners

Another serious threat to this scenario is collusion among multiple clients. When multiple nodes organize together, sharing their processing power, entities with high computational capabilities can be created. These entities, known as "pools," are very common in today's cryptocurrency environment. In the Bitcoin network, for example, the computational power required for mining activities has become very high, making "solo mining" almost impossible. For this reason, nowadays, joining a mining pool is the only valid option to mine Bitcoin. Large mining pools could play a key role in a cryptocurrency environment, and if they exceed 50% of the total computational power available in the entire network, they could even control it, by performing the so-called 51% attack. The 51% attack is a typical threat of a permissionless blockchain-based system, due to its open nature. However, also permissioned environments are prone to this attack, as the administrators could give themselves as many participants and nodes on the blockchain as desired [189].

Definitions

51% Attack In the blockchain environment, the 51% attack, also known as majority attack, refers to a situation where an adversary, which controls more than 50% of the total computational power of the network (in case of Bitcoin, the 50% of the entire network hashing power), acts maliciously to disturb the network's operation.

In the particular case of a cryptocurrency, an attacker performing 51% Attack is able to:

- Spend the same coin twice, i.e., double-spend attack
- Prevent any other transaction from being confirmed, i.e., denial of service against typical users
- Prevent any other miner to mine new blocks, i.e., denial of service against legitimate miners

An attacker performing 51% Attack is NOT able to:

- Steal funds from other wallets
- Create new coins
- Change the default reward (coins generated per block)

Attacks

Unlike quantum computing, which is still perceived as a remote threat, collusion between miners is an ongoing problem for almost all blockchain-based systems. In fact, several real-world examples can prove the danger of this threat.

The 51% attack is often considered to be a very remote threat, with little chance of occurring. This belief is based on the extreme difficulty of performing it in the Bitcoin network. On a permissionless blockchain with the PoW as consensus mechanism, the 51% attack requires the attacker to gain 51% of the total hashing power of the network to be successful. Given the maturity reached by the Bitcoin network, in terms of the distribution of miners and the hashing power required to mine a new block, this type of attack is highly unlikely, not only for the difficulty of obtaining the necessary computing power but also for the high costs of this operation which, compared to any profits, make the attack also not very profitable. A demonstration is given by the fact that, historically, only one entity in the whole history of Bitcoin has managed to have the computing power necessary to perform the attack. In January 2014, a mining pool called Gash.io grew so much that it came close to handling 51% of the total hashing power of the network. The event caused tension among the Bitcoin community, but it was resolved very quickly. Not being an intentional growth, the mining pool solved the problem by reducing the number of participants, and therefore the total hash rate available to them, also committing to never exceed 40% in the future. Although this incident has shown the real possibility that a single entity will get more than half of the total computational power, the current conditions of the Bitcoin network make the repetition of this event unlikely. In fact, starting in 2017, the Bitcoin network has registered a tremendous increase in the hashing power available to the network, as shown in Fig. 3.1.[3]

The total hash rate of a permissionless blockchain-based cryptocurrency represents a good security metric. The more hashing power is available in the network, the greater its overall security and its resistance to attacks. The hashing power of a cryptocurrency network is, in general, unknown and difficult to compute exactly. However, it is possible to calculate an estimation given the number of blocks mined in a particular time interval and the current block difficulty. In the particular case of Bitcoin, the estimated hash rate per second (TH/s) is computed by the following formula:

$$TH/s = 2^{32} * \frac{D}{T} \tag{3.1}$$

where T is the average time between the mined blocks and D is the difficulty. Cryptocurrencies with a low total hash rate value certainly have a less negative impact on the environment, in terms of electricity consumption, and miners have

[3] https://blockchain.com/charts/hash-rate.

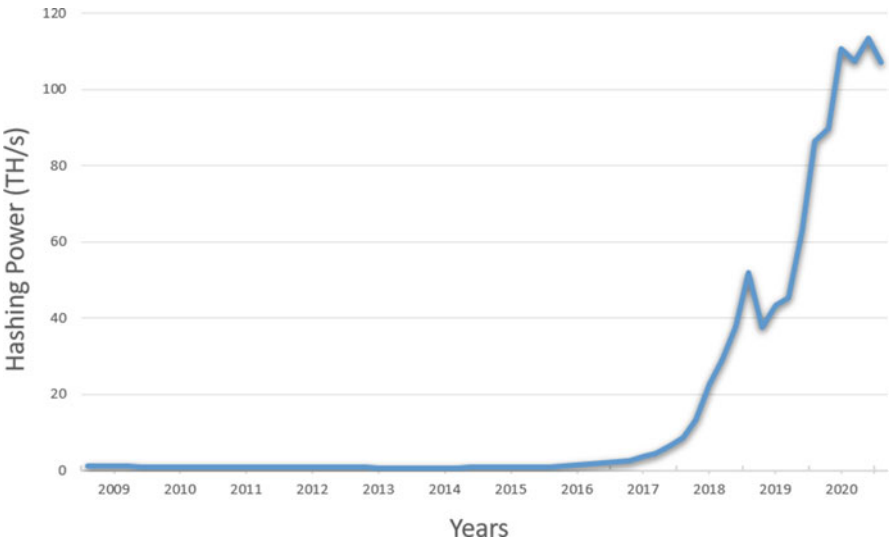

Fig. 3.1 The estimated computational power of the Bitcoin network, expressed in terahashes per second, from its origin to May 2020

a better chance of making profits. However, in this case, also the overall security level is low, exposing the network to different types of attacks.

In May 2018, an unknown malicious entity, with access to a very large amount of hashing power, performed a 51% attack on the Bitcoin Gold network. Once the attacker gained control of the network, he/she performed several double-spend attacks against different exchanges. Immediately after the attack was detected by the Bitcoin Gold community, the targeted exchanges tried to mitigate their loss by waiting for a longer amount of confirmation before approving transactions. This countermeasure does not seem to have helped, and the attacker eventually stole around $18 million. During these attacks, normal users are usually not exposed. The only parties at risk are users who accept payments (usually of large amounts) directly from the attacker. Since the cost of mounting such an attack is high, the attacker can only profit if they can quickly get something of high value from a fake payment. An actor like a cryptocurrency exchange usually accepts large deposits automatically, allowing users to trade into a different coin (including fiat currencies) quickly and then withdraw the desired amount automatically. For this reason, exchanges are more exposed during 51% attacks.

Definitions

Cryptocurrency Exchange Also known as digital currency exchange (DCE), it is an online service that allows customers to trade cryptocurrencies for other assets, such as other virtual currencies or conventional fiat money.

In January 2019, the Ethereum Classic (ETC) network also fell victim to a 51% attack. From a first analysis, it emerged that the attacker has mined only empty blocks, suggesting that he/she aimed more at winning the rewards of the new created blocks rather than executing a 51% attack to double-spend coins. Subsequently, some exchanges said they blocked all ETC transactions as soon as a major reorganization of the blockchain, which included an attempt to double-spend attack, was detected by their systems. This timely reaction has certainly mitigated the damage of the attack, preventing huge losses for the exchanges.

As shown by these examples, early detection of a 51% attack is very important for limiting the damage that can result. For this reason, several research efforts have been made in this direction, also considering that this is a general problem that affects almost every blockchain-based system. As such, it is not only limited to cryptocurrencies. The systems most subject to this attack are certainly the permissionless ones, since anyone is allowed to join the network (without authentication) at any time and become a miner. However, permissioned blockchain-based systems are also vulnerable to 51% attacks, especially in the case of consortium blockchain networks. In these systems, the access is controlled, and permissions are regulated according to the role played by each user. Several institutions, such as public and private companies and governments—and possibly other actors—collaborate with each other to use the blockchain and maintain its security. However, if collusion occurs between a subset of these institutions, network security is no longer guaranteed.

One of the most promising methodologies involves the use of intelligent software agents to monitor the activity of stakeholders in the blockchain networks to detect anomalies, such as collusion. In [190], the author proposed a solution that, by leveraging supervised machine learning techniques and algorithmic game theory, reduces the chances of collusion in decentralized systems. For each new block, the proposed solution estimates the utility function based on the value of the service or goods sold in the transaction examined. Based on this function, the intelligent agent decides how fair each new transaction is, compared to the system under analysis, and therefore the probability of being generated in the context of a majority attack.

3.3.3 *Open Issues*

A possible solution for the threats posed by quantum computing and other advances in technologies certainly revolves around the development of PoW algorithms (or other control protocols) with information-theoretic security. This means that the security of the protocol derives exclusively from information theory. For this reason, it is impossible for an adversary to break the system, even with unlimited computational power, simply because he does not have enough information to compute.

As already discussed, youngest systems are most vulnerable to collusion and, more in general, to 51% attacks. In fact, a relatively young system is most likely characterized by an equally young community, not yet stabilized. This could create conditions, even temporary, in which obtaining 51% of the total computing power of the network is easier. This would jeopardize the stability and security of the entire network. However, even in more mature systems such as Bitcoin, the possibility of collusion between mining pools remains possible. In fact, as depicted in Fig. 3.2, very few large mining pools control a large slice of the total hashing power available in the network. Despite this, collusion between these pools seems unlikely, as the long-term profit of honest mining is higher than what they could have in a short period in which they perform a 51% attack. However, it is necessary to verify whether this balance will remain unchanged even after the next scheduled halving,

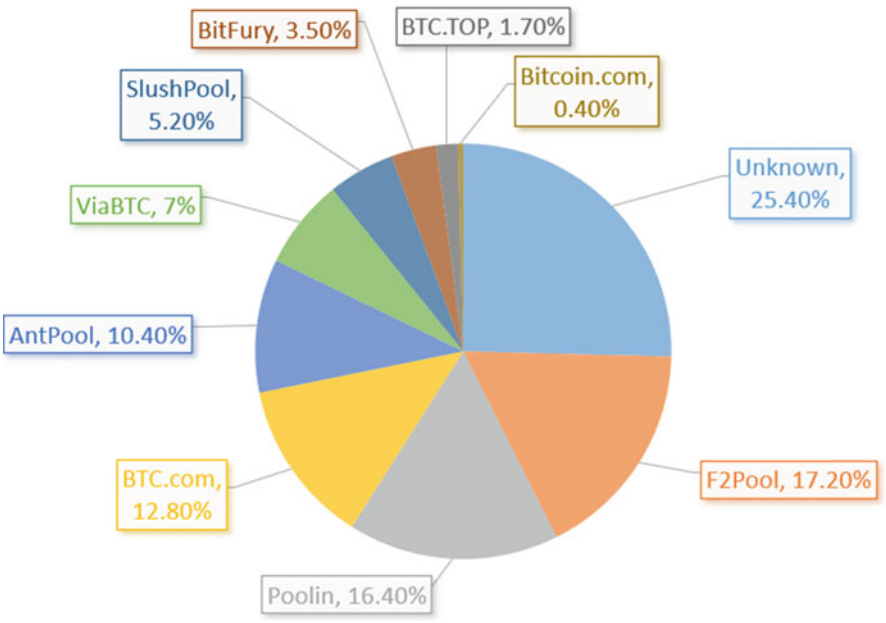

Fig. 3.2 Market share of the most popular Bitcoin mining pools (June 2020)

when the reward for new blocks is cut in half. After each halving phase, in fact, the consequences for the network and its security are unpredictable. The community of Bitcoin users often disagrees on the possible forecasts that precede a halving, with consequences ranging from the change in the price (positive or negative) to the destruction of the system. Furthermore, when the Bitcoin network reaches the available currency limit, and the miners' only profit will be the fees of the transactions included in new blocks, the equilibrium may change forever.

🪄 Definitions

Bitcoin Halving Scheduled event of the Bitcoin network, when the reward for mining Bitcoin blocks is cut in half. At the start of the network, the reward was set to 50 Bitcoins. Then, the Bitcoin protocol imposes a halving event for every 210,000 blocks added to the ledger. The criterion behind this choice was not explained in the foundational paper of the Bitcoin network. However, many believe that this mechanism serves to distribute more coins early on, attracting users, and to keep inflation in check later on.

3.4 Scenario 3: Infrastructure

As already mentioned in this chapter, the security of a cryptocurrency, as well as for any other distributed system, is closely related to the nodes that compose its network. In Sect. 3.3, we discussed the security risks related to a low total computational power available on a cryptocurrency network. When a cryptocurrency is not sufficiently mature in terms of resources to secure its ledger, the major risk is the presence of an adversary with greater resources who takes control of the system. One methodology for mitigating this risk is to increase the resources available to check network operations. In this way, taking possession of the majority becomes more difficult and expensive for an adversary. As a result, such an attack becomes unlikely and possibly less profitable. However, the increase in the number of full-nodes and their resources does not protect the network from any kind of attack. In fact, even without having enough resources to perform a 51% attack, a cryptocurrency could be attacked in other ways. For example, Denial of Service (DoS) attacks could be performed to undermine the availability of the system by targeting the whole network, individual users/entities, or specific services. These attacks can be favored by the structure of the system, exploiting both architectural weaknesses and infrastructural ones, even temporarily present in the network. As an example, in a permissionless blockchain-based system, the geographic distribution of full-nodes is really important to increase the network's stability and resistance to different types of attacks. In fact, these nodes are responsible not only for the validation of transactions and the creation of new blocks but also for maintaining

and distributing the updated copy of the ledger. All these activities, in theory, can be compromised within a specific geographical area by attacking all (or most of) the full-nodes within it. As a result, the version of the ledger in use within the region under attack may differ from the one used in the rest of the network. In addition, some network services in the same area may become unreliable or completely unavailable. For these reasons, it is necessary to consider the geographical position of the full-nodes, distributing them uniformly to increase the overall security level of the network. This approach is not always easy to follow, with difficulties that vary according to the properties of the considered system. In permissioned blockchain-based systems, the institutions managing the network have full control over the privileged nodes responsible for the security of the network. Consequently, these nodes can be deployed and geographically distributed according to the level of security to be achieved. On the contrary, in permissionless blockchain systems, no details about full-nodes are known at the time of design. In fact, during the whole lifetime of the network, the number of full-nodes, their geographical distribution, and their total computational power change continuously. As a result, the network security level varies over time.

Threats that arise from physical infrastructures, such as communication networks, are often underestimated and neglected in the security analysis of a cryptocurrency. In this scenario, we consider a cryptocurrency that relies on the public communication infrastructure to manage communications between nodes provided by its protocol. We highlight the security problems and vulnerabilities that may result from its physical infrastructure. In our analysis, we consider the most important cryptocurrencies currently available, with Bitcoin, once again, as the main example.

3.4.1 Threat: Hijacking Network Infrastructure

In this scenario, the major threat is represented by an adversary that interferes with the cryptocurrency protocol by manipulating the network traffic. The malicious actor could be either an insider or an outsider. In the former case, the malicious actor actively participates in the cryptocurrency's protocol during the attack—for instance, by forging fake network packets. In the latter case, instead, the attacker does not join the protocol, performing a stealthy attack—for instance, by dropping network packets. These types of attacks are directly influenced by several factors, derived both from the technologies used by the protocol and the physical network infrastructure of the cryptocurrency. As an example, the Internet routing infrastructure could be an important attack vector, especially for permissionless blockchain-based systems. In a permissionless cryptocurrency network, anyone from anywhere in the world can join the protocol by running a full-node. Despite this, the nodes that compose the network are unlikely to be uniformly distributed across the globe. This means that, with high probability, most of the full-nodes are hosted in a few Internet Service Providers (ISP)s. Consequently, these few ISPs will

route most of the network packets of the entire system. In such a scenario, several attacks could be performed to target a cryptocurrency, either directly or indirectly, by attacking the ISP's network infrastructure. In this context, we can mention few different malicious activities that could be performed by either an external adversary or a malicious ISP:

- Network packet redirection—e.g., Border Gateway Protocol (BGP) hijacking
- Network packet manipulation—e.g., eclipse attack
- Network packet filtering—e.g., blackhole attack

These malicious activities, classified as Internet routing attacks, are used to perform the following attacks:

- *Partition attack:* aims to partition the network of the cryptocurrency under attack into different disjointed portions, so that the different parts cannot communicate with each other.
- *Delay attack:* aims to delay the new block propagation so as to allow several other attacks, such as double-spending.

These attacks could be performed against any blockchain-based system that makes use of the Internet for its communications. The objectives and possible consequences are manifold and range from the double-spending attack to the Distributed Denial of Service (DDoS). The impact varies by victims: if the attack is performed against a merchant, it is susceptible to double-spending attacks; if the victim is a miner, the attack wastes its computational power; finally, if it is a regular node, it is unable to contribute to the network by propagating the last version of the blockchain.

Definitions

Border Gateway Protocol (BGP) BGP is the standard (de facto) routing protocol that manages how IP packets are forwarded on the Internet. Neighboring networks or Autonomous System (AS) exchanges between each other the routes that lead to different IP prefixes. For each given IP prefix, its AS is responsible for advertising the original route, which is then propagated among ASes until all of them are aware of it.

Autonomous System (AS) An AS is a collection of connected Internet Protocol (IP) routing prefixes under the control of one or more network operators on behalf of a single administrative entity or domain that presents a common, clearly defined routing policy to the Internet.

Internet Service Provider (ISP) An ISP is an organization that provides services for accessing, using, or participating in the Internet.

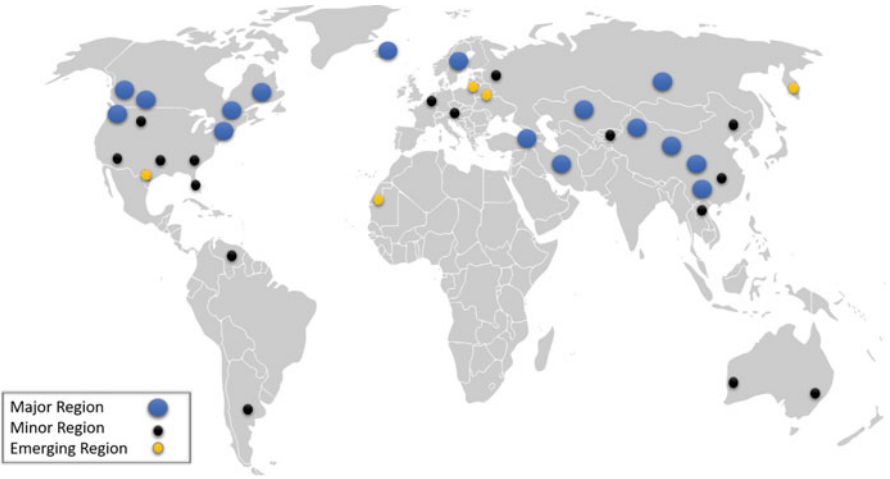

Fig. 3.3 Global overview of the Bitcoin mining regions. Data sourced from [191]

The feasibility and the success rate of both these attacks could be directly influenced by the geographical distribution of the full-nodes of the targeted cryptocurrency. In the particular case of Bitcoin, the hashing power is distributed in very few regions—as shown in Fig. 3.3—with the highest concentration in China (Sichuan region) and North America. This means that, by attacking very few ISPs, it is possible to cut a huge part of the total hashing power off the Bitcoin network, with serious consequences for the security of the system and its users.

3.4.2 Attacks and Countermeasures

In March 2014, several users of the Bitcoin network noticed suspicious activities on mining systems connected to some mining pools. For a few days, the crypto-miners' community has been in turmoil, struggling to understand what was going on. Several users reported a mysterious redirection of their systems to an unknown IP address, which responded with the same protocol—Stratum—used to coordinate the mining pools. After the redirection, everything worked as before. The miners, indeed, continued to receive work from the new server, limiting the blocking of the systems to a few seconds, caused by the server change, keeping the anomaly hidden. Only a small detail had changed: the miners no longer received any reward for their work. Following repeated reports, researchers from the Dell Secureworks Counter Threat Unit discovered several network traffic hijacking activities between February and May 2014. After subsequent investigations, 51 compromised networks, belonging to 19 different ISPs, were discovered. The hijackers exploited the hashing power of legitimate miners by redirecting their mining traffic to a malicious server disguised

as the legitimate pool. The attacker used a technique called BGP hijacking, which takes advantage of the lack of authenticity in routing messages. Indeed, the BGP protocol does not verify the legitimacy of the advertisements. Consequently, any AS can advertise any IP prefix. From the miners' point of view, the attack was completely transparent. Usually, the mining hardware is continuously connected to its pool server to receive tasks. After the attack began, the miners who tried to connect to the legitimate pool server received a new BGP route, pointing to a malicious server maintained by the attackers. Once the hijacker stopped the attack, the miners who were redirected to the malicious pool continued to see activities and tasks but were not rewarded. Miners who have not been redirected remain unaffected. The attackers repeated the hijack in short rounds, allowing the activity to continue undiscovered. The attack lasted for about 4 months, involving also other cryptocurrencies different from Bitcoin. The same researchers who discovered the attack linked BGP malicious announcements to a single router owned by a Canadian ISP and estimated the attackers' profit in $83,000.[4]

The security of blockchain-based systems from a network perspective was at first underestimated compared to other attack scenarios. However, the example provided by the 2014 attack has concretely demonstrated the possibility of indirectly attacking a cryptocurrency by compromising the Internet infrastructure using routing attacks. The hijackers who carried out this attack had the sole purpose of illicitly using the computational power of honest miners to obtain rewards for creating new blocks. For this reason, technically, the Bitcoin network, its protocol, and its regular users have not been compromised. What has been compromised is the BGP routing protocol, used by the Internet infrastructure, to directly attack the Stratum protocol, used to coordinate the pools of miners. Once computational power was obtained, the attacker could have used it to perform other types of attacks against the Bitcoin network and its users, such as the 51% attack and the double-spend attack. By using this technique properly, it may be relatively easy to partition a cryptocurrency network. As shown in Fig. 3.4, the geographical distribution of the mining pools could facilitate this attack. We can observe how few main regions have more than 40% of the computational power of different cryptocurrencies. China, for example, holds 45% of the entire hashing power of the Bitcoin network, while Europe over 40% of the computational power of Ethereum, Zcash, and Monero. This observation implies that a very large chunk of the computational power of a cryptocurrency could be controlled by attacking a relatively low number of ISPs. As noted by Apostolaki et al. in [192], only 13 ASes hosted 30% of the entire Bitcoin network, while 50 ASes host about 50%. The distribution of Bitcoin nodes per IP prefixes is also surprising; only 63 prefixes contain more than 20% of the network. These observations have direct implications on both the delay attack and the partition attack. The former, indeed, becomes much more disruptive with this high concentration of nodes in a few ISPs. At the same time, the latter is much

[4] https://www.secureworks.com/research/bgp-hijacking-for-cryptocurrency-profit (Last checked August 2020).

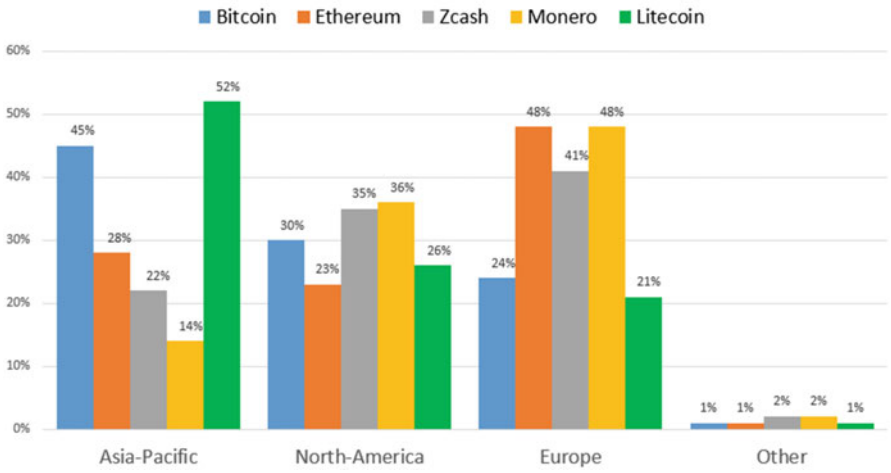

Fig. 3.4 Worldwide distribution of cryptocurrency mining in 2018

easier to execute, but the number of achievable partitions is lower. In the same paper, they also observed that only three ASes intercept more than 60% of the Bitcoin network traffic, making the delay attack much more powerful.As a known open issue, BGP's security has attracted a lot of attention. Several works have been proposed to mitigate routing attacks, mainly following two different approaches: origin validation and path validation. The former aims to filter advertisements from unauthorized ASes, while the latter digitally signs BGP packets to validate the list ofASes through which the announcement passed. To overcome this problem, a countermeasure to the partition attack in the Bitcoin network was presented in [193]. The proposed system, called SABRE, is completely transparent for both BGP and Bitcoin protocols. SABRE implements a Secure Relay Network for Bitcoin by introducing a new type of node, composed by a software module and a hardware component. The software module is a modified Bitcoin full-node (the control component), while the hardware module is a programmable switch (data plane component). The SABRE network is composed of several SABRE nodes and a number of regular Bitcoin full-nodes connected to them via UDP. When an external node advertised a new block to a node connected to the SABRE network, the new block is immediately forwarded to the SABRE data plane. If the switch does not contain the hash of the new block in its memory, the new block is forwarded to the control component. If the control plane validates the new block, the switch memory will be updated, and the new block propagated in the whole network. The proposed countermeasure is focused on Bitcoin. However, its design is general and could be applied to other blockchain systems to mitigate partition attack. Other countermeasures have been proposed to secure routing protocols preventing the attack described above [194–197]. Heilman et al. in [198] examine the eclipse attack on a single node in the context of Bitcoin's Peer to Peer (P2P) network.

Gervais et al. [199] consider other aspects of the centralization of Bitcoin and their consequences to the security of the protocol. Measuring and detecting routing attacks has seen extensive research on BGP hijack [192, 200, 201], as well as interception attacks [202].

3.4.3 Open Issues

Although the problem of attacking a cryptocurrency through the architecture of the Internet has been widely studied, we are still far from a solution. Several works have been proposed to mitigate BGP hijacking. However, none of them has been widely distributed, given the difficulty of choosing and standardizing a single solution and consequently updating all network hardware. Therefore, the Internet is still vulnerable to this attack, cryptocurrency networks included. Some countermeasures have mitigated the partition attack problem in the Bitcoin network, but they should also be extended and tested to the generic case of a permissionless blockchain-based system. Furthermore, the delay attack is still a serious threat, not only in the Bitcoin network but in all blockchain-based systems. The main challenge in solving these security problems in existing cryptocurrencies lies in the difficulty of updating an existing protocol. Usually, an important update often requires a fork of the entire project, with several practical problems. This is not always possible, especially for mature cryptocurrencies with high capitalization, such as Bitcoin. The countermeasures discussed in the previous chapter, however, should be considered in the design of a new cryptocurrency. A state cryptocurrency should consider the discussed threats, taking advantage of the available literature and lessons learned from past attacks.

3.5 Toward a State-Sponsored Cryptocurrency

When you think of a future state-sponsored cryptocurrency, it is natural to imagine it very similar to the existing ones, in particular to the most famous and used in the current crypto-exchange panorama. Bitcoin is the most popular virtual currency in the market today, with a solid technology behind it. Introduced to be the pillar of a new electronic payment technology designed to revolutionize the world of finance, Bitcoin is still suffering from some disadvantages that risk compromising its diffusion on a large scale. While several big players in the financial sector are ready to bet on its potential, the Bitcoin protocol needs several improvements to aim for the creation of a secure, reliable, and robust electronic currency. In this section, we discuss the requirements that a state-sponsored cryptocurrency should have, highlighting the features of Bitcoin (and other existing cryptocurrencies) that do not fit the identified properties.

3.5.1 Bitcoin Limitations

The Bitcoin network suffers from several architectural shortcomings that have limited its diffusion as a means of payment to be used as an alternative to traditional currencies. Consequently, the vast majority of users use Bitcoins as a speculative medium. The most important disadvantages of the Bitcoin protocol, as well as of any other Bitcoin-like cryptocurrency, that cause skepticism among users and restrain the spread of this technology are the following:

- **Fear of losing the wallet:** If a hard drive breaks or a malware corrupts/deletes the data and the wallet file is damaged/deleted, all Bitcoins contained are irreparably lost. Nothing can be done to recover them. These coins will remain in the system forever without anyone being able to spend them. Due to this limitation, a private individual could lose all his savings in seconds, just as a company could go bankrupt because of a corrupt file.
- **Irreversible transactions:** When a good or service is purchased using Bitcoin and the seller does not respect the contract (by not sending the goods or not providing the service agreed upon within the terms), the transaction cannot be canceled. This problem can be solved by using a third-party escrow service, but this implies the existence of a trusted third party, in contrast to decentralized distributed systems.
- **Volatility of the value:** The value of Bitcoins changes continuously based on demand. This constant fluctuation, in addition to lowering customer confidence, causes several problems. For example, if a purchased product is returned a week later, should the merchant return the same amount (in BTC) received at the time of purchase (even if the value has changed at the time of the refund)? What currency should BTC be tied to in this cases? These are still important open questions that the Bitcoin community does not yet have a consensus on.
- **Risk of unknown weaknesses:** The Bitcoin system may contain unknown defects, being a fairly new system based on recent technologies. If Bitcoins were widely adopted and a critical flaw was discovered, this could lead to the destruction of the Bitcoin economy.
- **Deflation:** The Bitcoin system is designed to reward early users. In fact, since the total number of Bitcoins is limited to 21 million, each Bitcoin will be worth more and more as the total number of Bitcoins available in the system reaches its maximum.
- **Poor scaling and weak architecture:** Bitcoin, like many other cryptocurrencies such as Litecoin, Monero, and Bitcoin Cash, suffers from several architectural problems that affect the scalability of the system and cause relatively high transaction fees and transaction times.

To be used on a large scale, Bitcoin must overcome the skepticism of citizens concerned by problems such as volatility of value, the usability of the system, and lack of approval from "trusted" organizations caused by the aforementioned drawbacks. Moreover, Bitcoin must also meet strict government requirements

regarding money laundering and illicit trades to be accepted by a state as an official currency.

3.5.2 Develop a State-Sponsored Cryptocurrency

The best approach to the development of a state-sponsored cryptocurrency is to combine the best technological features of current cryptocurrencies with the properties of a standard fiat currency, under the supervision of a central bank. The result could be a revolutionary payment management framework which, by exploiting the advantages of a distributed system, lowers the management costs typical of a centralized structure. The objective of a state-sponsored cryptocurrency should be to support the national economic system by providing the normal functions of a banking institution, such as the payment circuit, the management of savings accounts, and the provision of loans. All these services would be provided by a secure platform, capable of reducing errors, speeding up the transfer of money, and preserving the anonymity of its customers. Unlike standard cryptocurrencies, a state-sponsored system should have the full endorsement of a government and its central bank. However, to be supported by a government, the cryptocurrency's protocol must ensure the application of specific guidelines for financial services that banks and other financial institutions must now respect all over the world, such as Anti-Money Laundering (AML) and Know Your Customer (KYC).

🪄 Definitions

Anti-Money Laundering (AML) This term refers to a collection of laws, regulations, and procedures designed to prevent criminals from concealing illegally obtained funds as legitimate revenue.

Know Your Customer (KYC) KYC refers to an identification process used mostly by financial companies to verify the identity of their customers and assess the potential risks of illegal activities in the relationship with their client. The term often refers to banking regulations and anti-money laundering laws that govern these activities. The customer recognition processes also concern companies of other types and sizes, to ensure anti-corruption compliance for their agents, consultants, and distributors. Banks, insurance companies, and companies operating at an international level increasingly require their customers to provide the necessary detailed anti-corruption information.

The main technologies behind a state cryptocurrency could be inherited from Bitcoin and other similar protocols. Asymmetric cryptography, for example, could

provide valid support for the management of payments, ensuring pseudo-anonymity to end users. Just like Bitcoin, users would use an asymmetric key pair to manage their deposit account: the public key to receive payments (such as the IBAN in the current SWIFT banking system) and the private key to authorize outgoing transactions. This then translates into the development of easy to use virtual banking applications, such as electronic wallets, already widely used to manage Bitcoin addresses and currently available for any electronic platform. One of the biggest differences between a state-sponsored cryptocurrency and the Bitcoin protocol will certainly be the owner of the ledger. In fact, almost all existing cryptocurrencies use a model called public blockchain, where anyone can join the network without authorization and the ledger does not have an owner. This model is certainly not applicable to a state-sponsored cryptocurrency, which needs to be regulated by a central bank. For this reason, the most suitable model could be a private blockchain, or a consortium blockchain, where participants need consent to join the networks. In this model, the ownership of the decentralized and distributed blockchain would be shared among all the national banks that join the network, previously authorized by the central bank. These financial institutions would be entrusted with the security of the platform. National banks would have the responsibility to provide the pair of asymmetric keys to users, and to verify and validate transactions, such as miners in the Bitcoin protocol. From a privacy point of view, since users utilize their public keys to receive money, transactions would be anonymous. National banks would be the only actors able to link a public key to an identity, having to fulfill AML and KYC requirements. However, these links would be kept confidential in accordance with the law and made public only to the judicial authority if requested. In this way, each user, once obtained his keys from the bank, is able to carry out transactions (both incoming and outgoing) safely, without any third party involved. Just like Bitcoin, users would be able to transfer money without the need for a bank to participate in the transaction. Furthermore, the transaction will be completely anonymous since the only part exposed to the public will be the participants' public addresses, which work as pseudonyms.

In summary, the actors who would participate in this hypothetical scenario are listed in the following, together with their respective duties:

- **Government and its central bank.** The central bank of the issuing state represents the entity that creates and manages the system. It has the authority to authorize national banks to participate in the network, making them part of the system and possibly miners. The central bank also has the task of regulating the amount of money available in the system, which is not created by miners as in the most famous cryptocurrencies like Bitcoin. By applying the government's monetary policies, the central bank is, therefore, able to increase the amount of money available in the system through a transfer of funds from the central bank private key to the public key of the financial institutions of the circuit. On the contrary, to decrease the money available, national banks would transfer funds to the central bank.

- **Banks and other financial institutions.** Financial institutions approved by the central bank are responsible for providing asymmetric keys to end users and possibly participating in the security of the protocol. In fact, a small group of national banks would be authorized to act as miners, collecting the transactions issued by citizens, validating the correctness of the keys involved, and checking for the presence of sufficient funds in the sender's account.
- **Citizens.** Citizens are the end users of the system. Once they have received the pair of asymmetric keys from a qualified financial institution, they can issue incoming and outgoing transactions in complete autonomy using their electronic wallet.
- **International users.** International money transfers would be much simpler. Any international user, regardless of whether he is a bank or a foreign citizen, could obtain an asymmetric key pair and start operating in the desired currency immediately.

This kind of platform would take advantage of all the benefits of a consortium blockchain, providing valuable support to the economy of the country. Transactions issued within the circuit automatically become irreversible, transparent, and auditable, making the system robust without sacrificing citizens' privacy. Also, like all decentralized systems, a state-sponsored cryptocurrency would have different properties such as fault tolerance, availability, and resilience, which help maintain its security even in the event of an attack. Considering the distrust that many users have toward cryptocurrencies, the new electronic currency could take the same name as the national fiat one. The name "virtual EURO," for example, could facilitate the acceptance of a cryptocurrency sponsored by the European monetary union, compared to a new one with an original and unfamiliar name. Most likely, a state-sponsored cryptocurrency would also be adopted to support the national fiat currency instead of replacing it [203].