

Article

Beyond Bitcoin: Recent Trends and Perspectives in Distributed Ledger Technology

Diego Romano ¹ and Giovanni Schmid ^{1,*}

¹ Istituto di Calcolo e Reti ad Alte Prestazioni, Naples, I 80131, Italy

* Correspondence: giovanni.schmid@cnr.it; Tel.: +39-0816139529

Abstract: In the last four years, the evolution and adoption of blockchain and, more generally, distributed ledger systems have shown the affirmation of many concepts and models with significant differences in system governance and suitable applications. This work aims to update the critical analysis of blockchain technologies carried out by our previous contribution to this journal, extending the focus to distributed ledger components and systems. Starting from the topical concept of decentralization, we introduce concepts and building blocks currently adopted in the available systems centering on their functional aspects and impact on possible applications. We present some conceptual framing tools helpful in the application context, and we will propose the concept of process authenticity, which we will discuss through two use cases: blockchain document dematerialization and e-voting.

Keywords: Blockchain technology, Process authenticity, Tokens, Anchors, Oracles

1. Introduction

Bitcoin and related technologies are developing and spreading at unexpected rates and unforeseen ways, impacting almost every industry. Since its conception, Bitcoin has been a discussion subject within restricted communities such as IT and cryptographic experts for nearly a decade. Nevertheless, over the years, the mass media have strongly influenced public opinion by emphasizing the possibility of profit by investing in cryptocurrencies. As a consequence, Bitcoin and few alternative cryptocurrencies have given rise to a constantly expanding market in which people invest fiat money.

In parallel to the soaring exchange rates of Bitcoin, many researchers, companies, and entrepreneurs became interested in the underlying technology and its potentials. Speculations began to circulate about what could potentially be done with this new spectrum of technologies, starting from the vision of an “Internet of value” where assets, in the broadest sense of logical or physical resources with a value, are exchanged as quickly as nowadays information moves around the world just thanks to the Internet.

Currently, distributed ledger technology (DLT) – especially that based on blockchain ledgers – is a subject of interest from many companies rushing to take advantage of the perceived benefits of using a time-oriented, tamper-proof ledger of records as a public or intercompany backbone for transactions, data keeping, and process monitoring.

According to the leading provider of market and consumer data [Statista](#), worldwide spending on blockchain solutions will grow from 1.5 billion in 2018 to an estimated 15.9 billion by 2023. In the first quarter of 2021 alone, blockchain startup companies around the world amassed 2.6 billion U.S. dollars in venture-capital funding, more than the whole year of 2020.

Expanding on the idea of a decentralized currency, as implemented in Bitcoin, many of these startups are explicitly building decentralized business models as an alternative to the existing centralized ones. They promote this idea as “cutting out the middleman” thanks to a peer-to-peer network where customers and suppliers make transactions directly, with the goal of decentralized data storage and saving money on transaction fees.




Figure 1. Global organizations' use cases for blockchain technology as of 2020.

In this regard, DLT offers new paradigms on how companies and people can interact and negotiate and new ways in which assets and data are represented and monetized. The benefits and prospects deriving from this got attention – sometimes perhaps a little too optimistically and superficially – for many application domains and related use cases. From smart grids to supply chains, e-government to e-health, jurisprudence to finance, virtually no sector or application domain did not prospect cases and potential advantages in using some DLT. At the same time, however, non-trivial challenges arise because this technology is still under development, its related business models are largely unexplored, and a comprehensive legal framework is missing. Some surveys conducted by Statista, which involved hundreds of companies worldwide during the three years 2018-2020, reflect this state of affairs. For example, Figure 1 presents use cases for blockchain technologies as of 2020. This statistic shows that companies were working on use cases concerning the following four main broad topics: Digital currencies/payments, Data sharing/reconciliation, Identity Protection, and Asset transfer/traceability, each of which can be in itself declined and adapted to many different applications scenarios.

Figure 2 instead shows the factors that over this span of years have been perceived as the main obstacles that discourage more significant investments in blockchain technology and platforms. As outcomes show, only a strict minority of respondents believed there were no barriers. At the same time, significant uncertainties concerned the actual usefulness and viability of this new technology, the risks associated with its adoption, and the lack of adequate regulations and standards.

1.1. Paper contribution and organization

Many of the contributions in the literature of recent years have tried to mitigate the above problems by orienting operators and people potentially interested in developing or using DLT thanks to an examination of existing solutions and their main characteristics in terms of both functionality and performance, referring to (a set of) specific application domains or use cases. In some cases (see Section 1.2), these contributions take the form of real *vademecum*, which is ready but reasoned reference for a conscious choice of the most appropriate platform (often the best compromise among both DLT platforms and traditional systems and architectures) able to satisfy one's needs.

The purpose of the present work is different: it aims to discuss the evolution of some basic DLT concepts in terms of their scope and implementation, the emerging of new




Figure 2. Perceived barriers for investing in blockchain technology during the three years 2018-2020.

concepts and tools, and which prospects these concepts and tools can have in terms of the effectiveness of the decentralization solution obtained and its cyber-physical security. In particular, we will discuss the challenges for the successful adoption of blockchain when the use case considered provides that the system must interact somehow with the physical world because of the physical nature of the assets to be managed or the input sources for the smart contracts. As we will try to corroborate with examples taken from actual implementations, there is a gap between the cryptographic techniques to protect the digital information at the ledger layer and their interface to the physical realm, which can give rise to the so-called “garbage in-garbage out” issue and nullify purpose and reliability of the entire system. It should be clear that such a problem is very critical for any application scenarios where blockchain is supposed to be the winning solution for implementing an effective and trustful form of decentralization. It indeed affects the success of DLT solutions in at least two of the four broad use cases in which industry and academy are currently investing (see Figure 1), that is, Identity Protection and Asset transfer/traceability. More generally, it represents a threatening barrier to the adoption of DLT outside the finance industry, where the actual use of this technology will depend on how and to what extent this problem will be solved.

Overall, this work intends to provide a broad overview of the main concepts, approaches, and technical solutions that have characterized the development of DLT so far. Section 2 discusses the two opposite models of achieving decentralization through DLT, with their significant differences in system governance, ledger management, and business models. Section 3.4.1 illustrates the core notions and their diverse implementations in major DLT platforms comparing to a reference architecture. In particular, this section compares the different types of technical solutions proposed for the ledger structure and the consensus protocol to allow a better performance and a more extensive set of possible applications. Section 4 discusses some emerging concepts and techniques whose scope is to reliably connect DLT networks to the physical context they must operate. The idea is to obtain cyber-physical systems capable of implementing an extended notion of process authentication, contemplating all the different agents that contribute to realizing the process, whether digital processes, human operators, or physical devices. Finally, Section 5 illustrates the difficulties and opportunities deriving from implementing the authentication concept mentioned above thanks to the analysis of two actual use cases in which the authors of this work are involved as software designers.

1.2. Related works

The scientific literature relating to DLT is vast, even if only considering contributions in reviews, surveys, and tutorials on the subject. Therefore we will limit ourselves to pointing out those articles that most inspired the contents and organization of the present work.

The authors of [1] pursued two primary objectives: (i) providing an entry point for non-security experts to the security and privacy features enabled by blockchain technology and; (ii) helping specialists and researchers to explore the cutting edge cryptographic techniques in this context. Their analysis includes representative consensus protocols, anonymous signatures, secure multiparty computation, non-interactive zero-knowledge proofs, and secure verification of smart contracts.

The survey [2] explores blockchain architectures in terms of their network and data model, execution workflow, consensus protocol, and application domain, focusing on those crucial for deciding whether to adopt the technology or not and which of the available solutions come closer to a given use case. This work's ultimate scope is to clarify the critical properties of blockchain systems, highlighting when and which blockchain technologies to choose and how they can be used and deployed.

Xiao et al. [3] provide a comprehensive and detailed description of consensus protocols for distributed ledger (DL) systems mapping to a five-component framework. Their framework allows complete and in-depth analysis of how the blocks are proposed, transmitted, and validated within the network and the mechanisms and incentives to agree on which valid blocks add to the ledger.

A stimulating review of DLT technologies [4] presents a taxonomy structured in three tiers, focusing on application development. Moreover, the authors present a guide for developers to identify which DLT best fits the business model and the functionality requirements of their project. It is helpful to discriminate among options when the audience is already familiar with the concepts presented and aware of the related implementation difficulties.

About the blockchain e-voting case study, several publications are helpful to understand the general setting of e-voting. The authors of [5] introduce a schematic of key concepts with a review of electronic voting systems at date. A practical example with reasoning on election organization, cryptographic tools, and voting schemes is available in [6]. Focusing on remote e-voting (also called online e-voting), an introductory critic is available in [7]. When going specifically in the setting of blockchain e-voting, the authors of [8] give a schematic overview of several initiatives present in the literature highlighting pros and cons. Finally, in [9] the authors give a severe framework to identify risks in the adoption of blockchain e-voting systems, taking the voters' side and providing a set of critical questions to evaluate any new voting system proposal.

2. The quest for decentralization and its realization through DLT

In addition to the financial speculations possible thanks to the cryptocurrency market, and regardless of the specific use cases, the success of DLT mainly lies in demand for interoperable and reliable decentralization in Internet services and architectures. *Decentralization* is a property that characterizes some distributed systems. A *distributed system* is a system whose components are on different networked computers, possibly in different geographic areas, which communicate and coordinate their actions in order to achieve a common goal. If some of these computers belong to different and independent administrative domains, so that they cannot be collectively managed through a unique organization but require the cooperation of multiple parties, then the distributed system is more appropriately called *decentralized*.

The adoption of a distributed system for implementing Internet services has been a significant trend since the introduction of *microservice architectures* around 2011, a variant of service-oriented architectures where applications are collections of loosely-coupled, fine-grained processes which communicate over a network using technology-agnostic

protocols such as HTTP [10,11]. A growing number of big companies, smaller businesses, and startups have swapped monolithic applications for microservices architectures which are deployed as distributed systems within a data center or at multiple remote sites, thanks also to virtualization and containerization techniques [12]. In this context, Cloud computing is a diffuse approach to offering and consuming IT services since it achieves economies of scale by sharing computer system resources and their on-demand availability without direct active management by users. However, companies usually deploy microservices and cloud architectures in very centric way when administering their components and managing their trust. Amazon, Google, Facebook, Netflix, and many others eventually prompted to shift from their giant and monolithic application stacks to cloud-based microservices, but just in order to improve their performance and scalability and to reduce difficulties in terms of software development, deployment, and management. From a governance point of view, their approaches remain purely centric, if not over-centric, where different service providers are administered by the same company (as in the case of YouTube and Google or Facebook and Whatsapp). In general, this envisioning has its roots in the economy of scale and the currently dominant business model on the Internet, with services substantially financed by online advertising, which in turn is all the more effective, the greater the ability to user profiling thanks to the analysis of their behavior on the net. Thus, the more online services a company owns (e.g., search engines and social media), the more behavioral data it can gather concerning its users, and the greater the chances it can provide advertisers and publishers with analytics for targeted advertising.

The demand for greater decentralization and balance of power among Internet players by a slice of public opinion stems primarily from this state of affairs. There is indeed a widespread concern that the oligarchy currently characterizing the provision of the most common digital information and entertainment services may undermine the rights and freedoms of individuals. Some fear that this excessive concentration of power and wealth in the hands of a few IT companies, joint with the pivotal role that these companies have in managing and distributing virtually any kind of information to billions of individuals, could also degenerate into a dystopian world where people are continually monitored and conditioned to make choices even against their natural, instinctive will. Shoshana Zuboff coined the term *surveillance capitalism* to denote “an emergent logic of accumulation in the networked sphere” based on “unexpected and often illegible mechanisms of [data] extraction, commodification, and control that effectively exile persons from their own behavior while producing new markets of behavioral prediction and modification” [13]. And some other influential thinkers (e.g., technoutopian Georges Gilder [14]) believe that DLT –precisely as enabling technology for the creation of decentralized networks and business models– can subvert organizational, social, and economic imprinting stemmed from the business practices of companies such as Google and Facebook, allowing for greater pluralism, the balance of power and respect for human nature, rights and freedom.

Whether or not one agrees with these analyses and conclusions, other practical circumstances point to decentralization as a critical factor for many application scenarios:

- risks deriving from assuming a single point of failure/trust, as in case of the “trusted third party” assumption;
- costs reduction, both in the financial field and more generally in the commercial one, thanks to the elimination of intermediaries and interoperability “by design” ;
- need to coordinate safely and reliably complex distributed systems which by their nature are better suited, in whole or in part, to peer-to-peer models.

In particular, decentralized-oriented technologies could support a killer feature in the context of a current primary trend in information processing, achieving a particular validated end after a pipeline of tasks performed by multiple and independent agents.

2.1. *Permissionless versus permissioned systems*

A significant divide exists between the cryptocurrency realm and the world of regulated business. It stems from two alternative business models, gives rise to different and somewhat antithetical DLT requirements, and turns out in two opposite approaches to decentralization: permissionless versus permissioned systems. In a *permissionless* system, virtually anyone can participate in managing the ledger: an individual, usually anonymously, only needs to get the right software and possibly be enrolled in the community. On the contrary, only a set of pre-authorized and well-identified parties can participate in building a *permissioned* system; this set usually composes of representatives from companies and organizations which are in charge of managing the ledger and, conversely than in the permissionless case, it is somewhat limited in size and stable over time.

In the vast and heterogeneous amount of blockchain-related literature, the above dichotomy is often confused with another view to classifying DL systems, which instead refers to the way users can interact with the ledger. It is indeed also helpful to distinguish between *public* and *private* systems, depending on whether reading access to the contents of the ledger (or submitting transactions) is allowed to anyone or only to groups of authorized users. This circumstance is unrelated and complementary to the previous one, and in fact, there are examples of blockchains for each of the four possible combinations of the two classifications. For example, regardless of the public, permissionless [Ethereum](#) main network, everyone can deploy an Ethereum-based private chain (for a tutorial on this topic, see [15]), whereas [Hyperledger Fabric](#) is an example of permissioned framework for developing applications both in the private and public models.

From a technical perspective, the private/public classification only affects the access control policies for end-users, while the permissionless versus permissioned distinction has profound implications on system management.

In permissionless DLT-based systems, the system's existence is on a voluntary and not pre-determined basis: peers freely decide if and when to be involved in the ledger management. Because of this, participants receive adequate incentives; otherwise, the network would not survive. On the other hand, the prospect of easy incentives would lead users to misbehavior in order to try to multiply their benefits. In particular, in a similar way to *Sybil attacks* [16] that can plague reputation systems, the same user could hide behind a multiplicity of anonymous identities to be able to grab a more significant number of benefits or to increase the probability of accessing benefits proportionally. Therefore, participation in the management of the ledger must occur at the price of not easily reproducible consumption of a resource and a specific cost, unlike standard digital data.

Bitcoin represents the first practical solution to the above "catch-22" dilemma: in this system, a cryptographic problem makes a user unlikely to be able to find a data capable of allowing the addition of a new block, and thanks to the execution of a computationally burdensome process, called *proof-of-work*, that probability can proportionally increase. To compensate participants for their investments in resources and computation time, those who solve the current proof-of-work first receive a reward in the form of a certain number of bitcoins. The proof-of-work, with its related consensus protocol and the underlying incentive mechanism, has proven to be very effective in preventing both Sybil attacks and another category of very relevant and specific cryptocurrency attacks, namely *double spending* [17].

However, this comes at the price of an underperforming and highly inefficient system, whose energy demand by design grows exponentially with the increase in the number of *miners*, i.e., those who participate in the proof-of-work. In Bitcoin, a miner generates a new block in about 10 minutes; a user can safely consider a transaction permanently registered on the ledger only after an hour and, according to CBECI [18], the current estimate for the annual consumption of the global Bitcoin network is around 68 TWh; that is, more than Austria in the same period. On the other hand, digital transactions relying

on a centralized authority for verification take the order of seconds to be confirmed, and they have a negligible cost when compared to Bitcoin. As of July 22, 2021, the estimated average Bitcoin electric power consumption per transaction exceeds one million Visa transactions [19].

Regardless of the opposition of many to pure cryptocurrencies, for their speculative risks and illicit uses (e.g., money laundering, ransomware, covert cryptomining), it is likely that the growing emergency due to climate change will convince many countries and companies to give up the use of systems such as the current Bitcoin, generally opting for more eco-friendly DLT. Episodes in the spirit of this trend have already occurred (e.g., [20,21]). After all, researchers in business and academia recognized the limits of Bitcoin's proof-of-work since the early years of its launch and made efforts to obtain more efficient and performing permissionless systems. This attitude turned out also in revamping research in consensus protocols based on quorum mechanisms, a field in the context of studies on distributed systems whose first results have occurred since the eighties and which in recent years have undergone substantial developments precisely concerning permissioned or "hybrid" type blockchains. In addition to the staunch supporters of permissionless systems, according to which these would be the only ones able to implement a true decentralization in Internet services, some people believe that decentralization is possible at various degrees and that a consortium-type approach for the provision of digital services can offer various advantages in different contexts of use and application domains.

As shown in Section 3.3, following the studies of the last few years, we have several consensus protocols that offer various trade-offs between scalability, energy efficiency, performance, and the threat assumptions under which they offer an adequate level of reliability. In particular, thanks to the choice of a consensus protocol, it is possible to implement a certain level and type of decentralization according to the usage scenario, both in the permissioned and permissionless model. However, we must emphasize that a consensus algorithm could not be sufficient to implement the requirements of decentralization in a complete sense. First, there are blockchain networks whose functioning depends on data from off-chain sources: the way these systems receive and manage data can significantly affect decentralization, and it is helpful to discuss how in these systems we can preserve decentralization also thanks to cryptography, which is one of the subjects of Section 4. Moreover, at least according to techno-utopians who advocate for permissionless systems, decentralization should be inspired by the principles of freedom, pluralism, and balance of power. Therefore, to be fully decentralized, a network must adopt an adequate governance model: to what extent can a network be considered truly representative of all its participants if a few only decide its constitution and evolution? These aspects transcend the fields of cryptography and, more generally, those technological; however, they are crucial and often underestimated in the current literature. Therefore we will mention them in the following subsection.

2.2. Governance models

Since the inception of Bitcoin by Satoshi Nakamoto, the concept of *governance* has interested actual and potential stakeholders. Who governs Bitcoin, who is accountable in case of problems? These are a few of the many questions users have tried to answer before pouring money into a new digital currency system.

Actually, Nakamoto designed Bitcoin with the idea of leaving it working without any additional governance apart from the one embedded in the protocol. We do not even know who Nakamoto is because no one was intended to be the governor of the currency, depicting Bitcoin as a gift of a talented person to the human community.

The first attempts to introduce some form of governance were linked to adopting new features and overcoming protocol issues. Currently, a group of developers who volunteered to deal with the limits imposed by the original protocol maintains a reference implementation of Bitcoin. They declare to follow the consensus changes within the

community rather than impose them. However, while leading the code distribution, they have a prominent position to control opinions and deployments in the Bitcoin core system, as other influential players do. We can say that governance still resides in the protocol and the consensus mechanism, but with a sort of external *lobbying*. Meetings of relevant people (e.g., miners, developers, investors) can lead to agreements on protocol changes and eventually to fork.

Nevertheless, subsequent blockchain and DLT initiatives took distance from the original idea of a self-consistent currency framework, and faced governance issues, often with fragmentary, individual, and sometimes disordered approaches. Using blockchain technology in applications other than a cryptocurrency implied some redesign, introducing new concepts and tools. The transition from a gift of an anonymous donor to an open-source project automatically introduced some off-chain governance concepts, at least considering the authority to propose protocol changes.

From these few historical reasoning, we understand that talking about blockchain governance means two different concepts, which authors in the literature refer to as on-chain or governance *by* the blockchain, and off-chain or governance *of* the blockchain.

2.2.1. On-chain governance

One aspect that blockchain enthusiasts embrace the most is the trustlessness of the systems. Before Bitcoin, any system was administered by an authority that users trusted for its correct functioning. Trustless systems overcome this limit by providing a protocol involving users through incentives to maintain integrity and functionalities. We can say that trustless systems do not eliminate trust. Instead, they require trust in the protocols that inherently incentives correct behaviors. This aspect strictly links to the consensus mechanism. In Bitcoin, it relies on the economic advantage to keep the system working as expected. In other experiences, more democratic approaches emerged.

The developers of EOS [22] tried to introduce a voting system through Delegated proof-of-Stake (DPoS). Token holders vote to choose block producers, with the underlying idea of preventing actors from controlling multiple nodes. In a world of sincere voters, this would be a good approach. However, criticisms pointed out that corruption is possible, as is cartel forming. If a consistent number of tokens are controllable by few actors, the voting role within the system integrity mechanisms is compromised, so does the trust in the system.

Tezos [23] presents similar criticisms. By using a peculiar proof-of-stake, this platform adapts to the needs of stakeholders by issuing a voting process over an election cycle lasting a certain number of blocks and subdivided into four quarters. In each quarter, proposed amendments of the protocol are subject to a voting stage and, if passed, to promotion in the next quarter, until final adoption:

- initially, an approval voting is issued to accept proposed tarballs containing new protocols. Each proposal collects some preferences;
- stakeholders vote the most preferred protocol again counting votes for, against, and abstained;
- if a certain quorum of votes is reached, including those explicitly abstained, with a minimum approval rate of 80%, the new protocol goes in the test chain. In each cycle, the system automatically updates the quorum to avoid lost coins;
- stakeholders vote again to adopt the tested protocol in the main chain. Also, in this case, quorum and 80% of preferences are the minimums to pass.

The most pivotal point of this mechanism is the ability to completely change the underlying protocol on stakeholders' preferences, exposing all the criticalities of direct democracy. As an extreme example, a misinformed majority could progressively bring the system to a catastrophic failure, even if democratically chosen.

An even more ambitious vision constitutes the building block of Internet Computer governance [24]. Based on DFINITY [25] mechanism, a Network Nervous System (NNS) governs the system through a liquid democracy [26] paradigm. The fundamental element

is called *neuron* and consists of a stake deposited and time-locked that enables the user to vote on several proposal types, like economics, policy, protocol, client upgrades. As an incentive to vote responsibly, the stake gets locked for months following a request to dissolve the neuron. An innovative feature allows stakeholders to delegate votes on specific matters to proxies: by following influencers on social networks, a stakeholder can get acquainted with opinions about complicated matters and explicitly delegate votes to the ID of the trusted influencer. Unlike what happens in representative democracy, where a prefixed number of representatives confront and decide on the majority about discussed proposals, in this liquid democracy system, a charismatic figure can collect enough followers to freely propose and grant critical changes on essential thematic like protocol or economics.

The last tentative of on-chain governance that we present is Decred [27]. This system uses a time-locked voting ticket purchasing to enable stakeholders to vote on governance issues and within the consensus mechanism. Through *Politeia* web platform, community members browse, discuss, and submit proposals about every aspect of project governance, including development funding. Every proposal is subject to an off-chain voting procedure which involves ticket purchasing by time-locking a number of coins. Even if the procedure does not get recorded on-chain, a time-stamping mechanism ensures the immutability of the records.

Stakeholders can also vote on consensus mechanism changes on-chain, adopting a time-locking funds procedure to obtain voting tickets. A significant majority of actors involved in the consensus mechanism must first adopt a proposed change before submitting it to voting. If a proposal reaches a 75% threshold of approval during the voting period, it replaces the previous mechanism after a fixed number of blocks by issuing a hard fork. A last similar voting mechanism is part of the consensus protocol, where five randomly chosen ticket holders vote to approve the proof-of-work block creation by the current mining winner. If a block gets three approvals, the block definitely goes in the chain. The involved ticket holders get a 30% of the transaction fees as a reward.

In general, Decred, like other direct democracy systems, is subject to the issues mentioned above, which can lead to disastrous outcomes.

2.2.2. Off-chain governance

When endowing a blockchain project with time or money, an investor is concerned about governance. The underlying software protocol is just as important as the laws and regulations of a traditional business. Organizations investigating the adoption of blockchain in their asset management usually take governance factors into account and considering the high number of variables involved, the individuation of critical aspects is often challenging. In [28], the authors provide an interesting framework to help identify the peculiarities of specific blockchain governance features. The analyzer can subdivide her interest range into three layers: off-chain community, off-chain development, on-chain protocol. For each of these layers, the authors isolate five factors, called *dimensions*, to help the analyzer make the right questions: roles, incentives, membership, communication, decision-making. Interpreting as many aspects as possible for each factor can help the investor understand how a specific blockchain initiative fits her intentions.

As a case study, the authors present the application of their framework to the Ethereum blockchain. For example, the analysis of the off-chain development layer include:

- **Roles:** Contributors, Maintainers, Ethereum Improvement Proposal (EIP) editors.
- **Incentives:** Contributors expect potential value increase of Ether from working voluntarily, as well as fun and social recognition. Ethereum Foundation (EF) pays some maintainers.
- **Membership:** Everybody is free to contribute. No formal selection procedure for maintainers or EIP editors. Usually, the most recognized contributors are called to relevant positions by EF.

- **Communication:** Contributors communicate via GitHub comments, meet-ups, events, and scheduled calls. Core developers' calls notes are published.
- **Decision Making:** Decisions happen during calls or through the EIP process.

We can quickly notice the analyzed layer weaknesses and strengths, even if some aspects like accountability still miss. For example, how does the EIP process work? Who is involved in the calls? The EIP process consists of filling a proper request by the EIP author, followed by a formal verification by the editor who can send back the proposal for completion, and a reviewing by client developers who decide to accept after an apposite EIP presentation during an *AllCoreDevs* call. Usually, the developers discuss the technical merits of the EIP, evaluate the impact on other clients, and coordinate the eventual implementation of the network update. Even if this seems a pretty technical process, it is rather political: in practice, the developers refuse an EIP if it is divisive and could lead to a network split. Such attitude could inhibit the adoption of effective but unpopular measures and have a profound political value. This kind of impact on governance is difficult to highlight and requires a deep investigation.

Besides the above orienting frameworks, we must point out that off-chain governance uses a looser set of rules, procedures, and social conventions than a code-based system. Because of this less formal structure, it is more challenging to manage and monitor. As a result, users can more easily avoid those policies. On the other hand, off-chain governance has a high degree of malleability, allowing the system to respond swiftly and seamlessly to unanticipated events and rapidly adjust to changes in the environment. The vague definition of off-chain governance rules allows for the required flexibility to limit or broaden the reach of these rules on a case-by-case basis, while the rigidity of on-chain governance rules means that malevolent actors may use them to undermine the system or shape it to their benefit if there is a design fault.

In general, on-chain and off-chain governance may show complementary limits. Let us take chain rewrites as an example: in the case of The DAO, where a set of programming weaknesses led to a catastrophic transfer of 3.6 Million Ethers to a single account, the EF, with an off-chain ethical decision, created a hard fork moving that amount to a recovery address, and implicitly creating the Ethereum Classic blockchain where the on-chain code is law, and code vulnerabilities might be unethical but leading to valid transactions. To deal with such cases, Tezos and DFINITY introduced on-chain voting on rewrites to let the stakeholders revert malevolent transfers. It seems a promising approach to deal with situations like The DAO but suffers in any case of the above-mentioned direct democracy risks. This case represents well the limits of on-chain and off-chain governances, which both tend to overcome national and international legislation by allowing individuals or communities to set impacting rules in grey areas.

3. DLT reference architecture and building blocks

Although DL systems vary significantly in their implementations, they share a typical reference architecture and some key concepts and technologies that give rise to their building blocks. We will use this reference architecture in the present section to overview the main characteristics of techniques and the cryptographic concepts and mechanisms that enforce such features. Moreover, we will use the scheme as a map to place concepts discussed in the following sections in the right place for the system. Figure 3 illustrates the four-layer architecture that virtually characterizes any DL system. The *application layer* defines the types of users involved in the systems and how they can interact with them, offering a suite of concepts and tools that can vary depending on the specific platform considered. The *ledger layer* encompasses all the data structures and mechanisms which give rise to the ledger. Some of them, such as linked lists and Merkle trees, make the ledger data structure, while others enforce consensus or keep track of objects and workflows managed at the application layer. The *consensus layer* concerns the components of the multiparty agreement protocol used to select a unique ledger among possible diverse instances, thus assuring its consistent state. Finally, the *network*

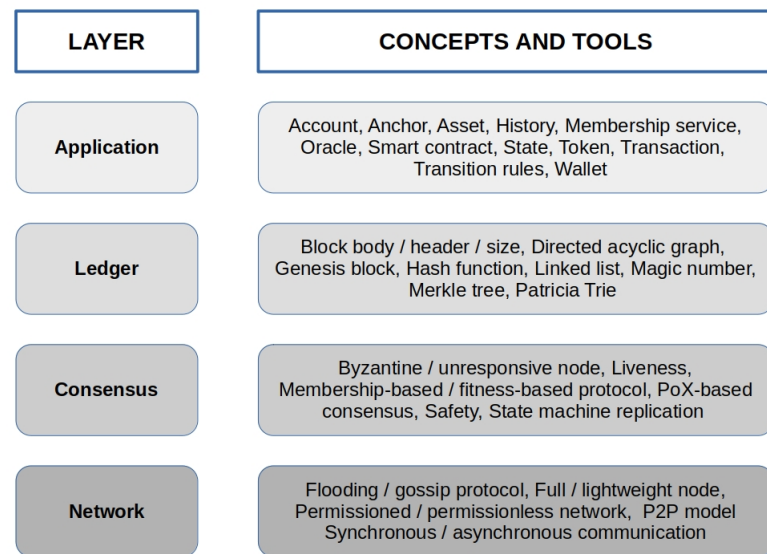


Figure 3. A DLT reference architecture.

layer has to do with how the peer-to-peer (P2P) network of nodes is built and shares information so that the ledger can be queried and managed, and a node can always get its updated state.

3.1. Application

This layer is responsible for two primary tasks: user management and specification of the services that can be delivered by the system, possibly according to users' roles and attributes.

3.1.1. User management

User management is very different depending on whether it is related to a permissioned or permissionless system: in these last, the difference between nodes composing the system network and client nodes lies in a grey area, whereas in permissioned systems, there is a clear distinction between clients and servers, and the latter can have different roles. In general, nodes that submit requests for transactions to be registered on the ledger, or access the ledger in reading mode, are considered users or clients. In contrast, system nodes actively manage the network, representing the set of distributed computational resources that replicate the ledger and offer its related services to the clients. In permissionless systems, clients can opt to take an active role, and vice versa, so that a given node can act as a client, system node, or both, and this happens at the node owner's discretion and not in a predetermined manner, as encoded in some system account policy. Accordingly, there is no distinction between clients and blockchain nodes: all of them are considered blockchain nodes that can transact with the other nodes and participate in building the blockchain. In order to perform the above tasks, a blockchain node needs a unique identifying address, and this address must be such that another node can verify that the emitted and received data are genuinely coming from or destined to that address. A blockchain address, along with its balance in digital coins or some other kind of asset, is named *account* or *wallet*: it represents how users are identified and managed in the system. Accounts can be implemented thanks to a digital signature scheme [29] by generating pairs of private-public keys: if a source uses a private key to sign a transaction, the corresponding public key is the one identifying the source address; analogously, if a transaction has a public key as destination address, only the owner of the corresponding private key will be able to redeem it. In a permissionless

system, private-public keys pairs for a given node are generated independently by the node owner, without any support by a Certification Authority (CA) [30]. Therefore no connection is attested between the identity of the node owner and such keys. Some systems, like Bitcoin, add further protection to users' anonymity thanks to a transaction model that resembles cash transfers, which will be touched upon shortly. Things are very different in the case of permissioned systems, for which multiple CAs can create a Public Key Infrastructure (PKI) [30] architecture capable of reflecting the different organizations that belong to the network and the roles of their units, both in terms of departments, blockchain nodes, and personnel. An example in this regard is Hyperledger Fabric, where a component called Membership Service Provider (MSP) [31] is designed to abstract away all cryptographic mechanisms and protocols behind issuing certificates, validating certificates, and user authentication.

3.1.2. Service specification

At the core of service specification for DL systems, there is the notion of *transaction*. It indeed represents the atomic operation to be tracked in the ledger and the only way to trigger a change in the system; in some platforms (e.g., Hyperledger Fabric), transactions represent the only way to interact with the ledger since they are required not only to write data on it but also to read data from it. In a system backed by a cryptocurrency, transactions are first and foremost static data records tracking digital coin transfers between senders and receivers. In Bitcoin, for example, a transaction specifies a list of inputs and a list of outputs where each input encodes a value in bitcoins (BTC) representing an unspent transaction output (UTXO) for the sender, and each output encodes a value in BTC along with the receiver address, so that the sum of the inputs values is not less than that of the output values. In order to get a valid transaction, the sender has to sign each transaction input properly to redeem its value and be able to use it in an output, possibly summed to other input values. If output values result in a smaller sum than the inputs, the system adds a UTXO for the sender to the transaction to equalize the difference. For example, Alice could redeem 0.05 BTC and 1.181 BTC from two of her UTXOs to transfer 0.67 BTC and 0.63 BTC to Bob and Eve, respectively, getting back a new UTXO of 0.381 BTC by the system.

All decentralized applications rely on state, transition rules, and history, but the implementation of these notions can significantly vary depending on the particular system. The *state* is data currently stored in the ledger that, along with the *transition rules*, determines the following (finite) set of possible valid states. When a transaction is a request for a change in the ledger's state, it will be considered valid or invalid by the system according to its current state and its transition rules. Finally, the *history* is the set of all previous transactions and the order in which they occurred in the ledger.

In the Bitcoin system, the state is the set of all UTXOs with their addresses at a specific time. For each transaction, its encoded program instructions check if the transaction inputs correspond to previous UTXOs for the sender, which she cryptographically signed for redemption. If the answer is yes, the transaction is registered into the ledger, indicating that its output values are UTXOs for the many related output addresses. More recent systems do not adopt the UTXO model, but instead, they have an explicit notion of *account balance* used to keep track of all the cryptocurrency or digital tokens owned by a user. Some systems (e.g., Hyperledger Fabric) further extend this functionality by allowing users to own generic digital assets not necessarily linked to a cryptocurrency.

Already in Bitcoin, the program instructions written in the scripting language Script, which for each transaction redeem its inputs, making them available for its outputs, represent a first and straightforward form of transition rules (see [32] for a detailed description of their implementation). Subsequent systems have expanded the programmable logic related to transactions up to the point of allowing the creation of complex programs for the exchange of digital assets of different nature or for triggering

actions in the physical world. Digital artworks, information, goods or services, funds or rules involving events or other transactions can be managed according to specific contractual requirements, thanks to programs that execute predefined actions based on the state of the system and the rights and obligations of those who participate to bargaining. For this reason, such programs are called *smart contracts*: they indeed represent rules resulting from agreements between the participants in a decentralized network, published and managed through the network. For efficiency reasons, smart contracts are not stored on the ledger since they may result in programs of considerable length. However, systems supporting them provide mechanisms also relying on the ledger to guarantee their authenticity, integrity, and public verifiability, as it happens, for example, for Ethereum smart contract Application Binary Interface (ABI) [33] or Hyperledger Fabric chaincode lifecycle [34]. In these systems, smart contracts work like sets of transaction rules, and they represent the backend code of decentralized applications. Parallel to the extended notion of transaction, blockchain systems of this type support different kinds of assets (e.g., Ethereum tokens), plus tools designed to uniquely and authentically associate a digital asset with a physical good (anchors) or to ascertain the authenticity of external data or processes (oracles). We will explore these emerging topics in Section 4.

3.2. Ledger

Introduced in Bitcoin [35], the blockchain is a tamper-proof, append-only data structure consisting of a linked list of transaction blocks. The key concept at the basis of the blockchain data structure is a *collision resistant hash function* [36]. Haber and Stornetta in 1990 [37] introduced lists of messages linked by including in the current message the hash digest of the previous one. The authors consider the linkage of digital certificates issued by a time-stamping service as a countermeasure to back-date or forward-date a digital document by the owner, even with a colluding service. In Bitcoin, however, the hash function is used in the context of *Merkle trees*, a kind of binary tree for efficiently checking the integrity or authenticity of a set of data [38]. Transactions get grouped in ordered blocks, and each transaction in a block is associated with its hash digest, thus resulting in the tree's leaves. Left and right hash strings are then recursively joined and hashed again to get the digest value for their parent node in the tree until the tree root (Fig. 4, sketch (a)). This way, the hash value associated with the Merkle root uniquely represents the ordered block of transactions with overwhelming probability. Such string is then inserted, along with the hash digest of the header of the previous block, in the header of the current block for reference, to realize the blockchain data structure (Fig. 4, sketch (b)).

Like a physical chain, a blockchain can be resistant to tampering but only provided that its links are sufficiently robust, i.e., difficult to replace with new ones. Assuming that an attacker cannot alter a given block¹, the collision resistance property peculiar of the hash function makes it hard to alter any other previous block in the chain. However, since hash functions are efficiently computable, a given blockchain could be built from scratch starting from its initial (*genesis*) block, with any modification in its subsequent blocks as required by the attacker (*long-range* attack). The critical idea in Bitcoin was to make the construction of a new block relatively tricky, thanks to a kind of cryptographic challenge introduced for anti-spam purposes in [39]. The resulting proof-of-work is a component of the Nakamoto consensus, and as such, we will briefly discuss it in Section 3.3.

3.2.1. Block structure

Although sketch (b) in Figure 4 is a reference model for all blockchain-based systems, the exact structure of a block can vary considerably from system to system, according

¹ A digital signature on the block fulfills this assumption, giving rise to the linked list of digital certificates introduced in [37].

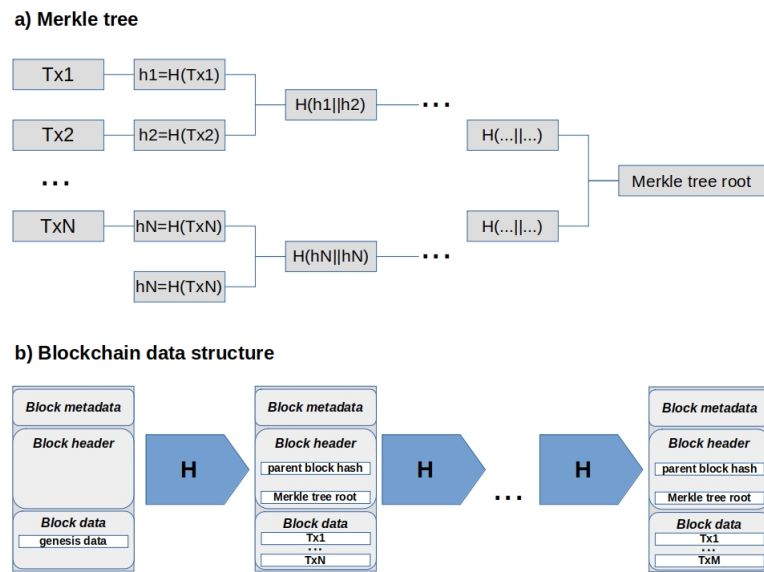


Figure 4. (a) Merkle tree construction for an odd number N of transactions, in which case the last hash digest gets duplicated to get a balanced tree, and; (b) The linked list of blocks starting from the genesis block, which results in the blockchain data structure. In both cases, H denotes a collision-resistant hash function.

to the functionalities and workflows of the application layer and the (one or more) consensus protocol supported. For example, the block header in Hyperledger Fabric contains just three records storing the block number, the parent block hash, and the current block hash, which is the Merkle tree root of the transactions recorded in the block data. In turn, each transaction composes of its header, signature, proposal, response, and endorsements fields. Things are very different in Ethereum, where a block's header is a data structure consisting of 14 fields, and a transaction consists of 7 fields. We illustrate the differences between these two systems in Table 1 for block headers and in Table 2 for transactions. These differences ultimately reflect the different nature of permissionless and permissioned systems concerning the management of decentralized computations.

Hyperledger Fabric is a consortium-oriented system whose primary goal is to support the reliable execution of business processes involving a small number of collaborating parties, among which there are no implicit relationships of trust. In this platform, there is a clear distinction among who is in charge of controlling the correct execution of the business process and who has to build the blockchain accordingly. An *endorsement policy*, a boolean expression defined to fulfill the use case-specific trust requirements, encodes the identities of a set of nodes called *endorsers* and their respective approvals or disapproval contribution to the final decision when checking a transaction proposal. Endorsement is also at the basis of the collaborative development and management of *chaincode*, the program in the form of transactions to be registered on the ledger, logically organized in smart contracts. Instead, a separate set of nodes called *orderers* is in charge of ordering in blocks the endorsed transactions and registering such blocks in the ledger. Unlike endorsers, orderers use an appropriate consensus protocol to agree with a majority on transactions ordering within a block and blocks ordering within the blockchain. In both cases, however, the decision of a node is expressed thanks to its digital signature, each final decision involves a limited number of nodes, and final decisions are reached through the definitive collection of their respective preferences and cannot change at later times. This design turns out in a simple block header without all the fields concerning the consensus protocol, whereas a transaction provides the fields to enforce the endorsement policy.

Table 1: Fields composing the block header in Hyperledger Fabric (HLF) and Ethereum (ETH)

HLF Field	Description
Block number	An integer starting at 0 (the genesis block), and increased by 1 for every new block appended to the blockchain.
Current block hash	The Merkle tree root of all the transactions contained in the current block.
Previous block hash	The hash digest from the previous block header.
ETH Field	Description
Number	Number of ancestor blocks (genesis block = 0).
Hash	Hash digest which combined with Nonce proves that the PoW has been carried out for this block.
Parent hash	The hash digest from the previous block header.
Nonce	The solution to the PoW for this block.
Uncles hash	The hash digest of the list of <i>uncles</i> for this block, that is childs of the parent block which are not parents of this block.
Logs bloom	Bloom filter of the indexable info contained in the log entry from the receipt of each transaction in the transaction list.
Transactions root	Hash digest of the Patricia trie built from the transaction list.
State root	Hash digest of the Patricia trie built from the <i>world state</i> , which maps addresses to account states.
Receipts root	Hash digest of the Patricia trie built from the receipts of the transactions in the transaction list.
Miner	Address for the fees collected because of successful mining.
Extra data	Arbitrary byte array (< 32 byte) containing data relevant to this block.
Gas limit	Current limit of gas expenditure per block.
Gas used	Total gas used for the transactions in this block.
Timestamp	Unix time at the inception of this block.

Ethereum follows an entirely different approach since its permissionless nature does not distinguish nodes based on specific roles: all nodes are potentially in charge of validating transaction proposals and ordering valid transactions in blocks to be registered on the blockchain. Due to the vast number of nodes involved in the above task, the system needs implicit mechanisms to prevent denials of service by constraining the duration of each transaction and its repeated submission. The *gas* concept and its related fields in both the block header and transaction data structures account for the first constrain, while the Nonce field in each transaction serves to ward off reply attacks. Implicit agreement mechanisms are also required to get consensus: the present running version of Ethereum relies on the proof-of-work *Ethash* algorithm, so its block header comprises the fields required to implement it. In addition to the specific fields of the Ethash algorithm (e.g., Hash, Nonce and Miner), the block header provides also the Uncles hash field to take into account the rule for consensus finalization, given that running Ethash can result in several simultaneously valid blocks.²

Although consensus finalization concerns more appropriately one of the characteristics of the protocols presented in Section 3.3, we will briefly discuss the existing alternatives to select the “official” chain here, since this can affect the structure of a block, as shown by the previous example. Furthermore, it may be helpful to spend a few more

² The two more fields, Difficulty and Total difficulty, are provided in the Ethereum block data to account for Ethash and its finalization rule.

Table 2: Fields composing a transaction in Hyperledger Fabric (HLF) and Ethereum (ETH)

HLF Field	Description
Header	Metadata about the transaction (e.g., name of the relevant <i>chaincode</i> , and its version).
Signature	A cryptographic signature by the client application used to check if the transaction details have not been tampered with.
Proposal	Encodes the input parameters supplied by an application to the smart contract which creates the proposed ledger update.
Response	The output of a smart contract that captures the world state before and after this transaction.
Endorsements	A list of signed transaction responses from each required <i>endorser</i> sufficient to satisfy the <i>endorsement policy</i> .
ETH Field	Description
Nonce	A sequence number of transactions from a given address which is increased by one for each new transaction in order to prevent replay attacks.
Gas price	Price of Gas in Gwei (1 Gwei= 10^{-9} ether).
Gas limit	Limit of the amount of ETH the sender is willing to pay for the transaction.
Recipient	An external owned address (EOA) or a contract address which is the destination of this transaction.
Value	The amount of ether/wei from the sender to the recipient. Value is used for both transfer money and contract execution.
Data	Data for activities such as deployment or execution of a contract.
(v, r, s)	Components of the transaction signature by the sender with the Elliptic Curve Digital Signature Algorithm (ECDSA).

words on this point as several people often misunderstood it, and in some systems like Bitcoin and Ethereum, it concerns the significant notion of “blockchain state” introduced in the previous section.

In Nakamoto’s paper [35] there is some ambiguity about the rule for finalizing a block, which is called “longest chain rule” but it is explicitly referred to as the chain of blocks starting from the genesis block that “has the greatest proof-of-work effort invested in it”. Following in a strict sense, the longest chain, as detailed, e.g., in [3], is not safe when the consensus procedure exploits a difficulty target, as in proof-of-work and proof-of-stake protocols (see next section), since longer chains may exist that contain far less mining work. For example, [40] describes two intriguing improvements to the *selfish mining* attack introduced in [41] which can significantly advantage the attacker in catching up the main official chain. In these systems, the correct method for selecting the current valid chain amongst the possible forks deriving from multiple block proposals is to compute the cumulative amount of work required to produce the chain. This fix was implemented early on in Bitcoin, which consists of a *heaviest chain rule*.

The *Greedy Heaviest Observed Subtree* (GHOST) finalization protocol was proposed in 2013 by Vitalik Buterin for Ethereum [42] to achieve a faster block proposal rate without incurring in too many *orphan* blocks, i.e., valid blocks which eventually got cast off as a heavier chain achieved dominance. By allowing new blocks to reference multiple predecessors, a parent and zero or more *uncles*, GHOST includes orphan blocks in the calculation of which chain has the greatest cumulative difficulty; moreover, it establishes

that the miners of its uncles get a fraction of the base reward for a new block. This last rule helps to get fairer management of the efforts by honest nodes, protecting them against nodes (or pools thereof) that exploit their hashing power to create small blocks with only a fraction of the transactions heard over the network. The protocol currently implemented in Ethereum is not GHOST, as erroneously reported by many sources, but a variant that only implements the second rule: uncles do indeed receive a reward, but they do not count towards the total difficulty of a chain [43,44].

However, the Ethereum blockchain has all the fields required for the implementation of GHOST, and the choice not to implement it seems to be due only to the fact that the current protocol, although simpler, still allows for a three times transaction throughput (10-15 tx/sec versus 3-5 tx/sec) and a much shorter transaction confirmation mean time (about 1 minute versus 10 minutes) when compared to Bitcoin, although the number of confirmations after which one can consider a transaction safely recorded into the ledger is 6 for Bitcoin but much more for Ethereum (according to [42], one should wait for at least seven confirmations, but major exchanges wait for around 50 confirmations to consider an Ethereum transaction complete). Nevertheless, after all, the designers of Ethereum are confident in the next *Serenity* release for boosting efficiency and performance, thanks to the switch from the Ethash proof-of-work to the *Casper FFT* proof-of-stake, as we are going to explain in Section 3.3.

3.2.2. Blockchain alternatives

The systems after Bitcoin usually do not deviate from the linked list approach, at least logically (i.e., concerning the way they record and retrieve transaction history); instead, relying on extended or complementary data structures for the blockchain to support more advanced functions. For example, the Ethereum blockchain makes use of (modified) Patricia tries [45] instead of Merkle trees, while Hyperledger Fabric associates a key-value database, storing the current state of the ledger, to a linked list of transaction blocks very similar to Bitcoin.

Permissionless systems such as *IOTA* and *Dagcoin* represent exceptions in which a directed acyclic graph (DAG) [46], first introduced as a data structure suitable for keeping track of transactions in a draft manuscript by the independent cryptocurrency researcher Sergio Demian Lerner [47], replaces the linked list of blocks. DAG ledgers can be block-based or transaction-based. In the first case, every vertex in the graph contains a block of transactions assembled similarly to blockchains, while in the second case, transactions do not get bundled into blocks, but each corresponds to a graph node. In both cases, however, a node references one or more parent nodes identified by their hashes. The links among nodes form a DAG, which confirms earlier nodes and establishes their partial order. As new nodes get added, earlier nodes can receive more and more confirmations, and the users who include them as parents in the node they create get their related fees (if any). If an attacker tries to edit or delete a confirmed node in the graph, the links to all its children and any further descendants will break since the hash and signature of any node depend on its parents' hashes [48].

The rules governing how a node gets added to the graph can vary from system to system and result in its consensus protocol, which usually does not diverge far from Nakamoto's consensus except for the "longest chain rule" (e.g., [49,50]). However, these approaches are far from being consolidated, and they often are based on arguments for which there are insufficient experimental simulations and proofs, lacking ascertained validation by experts in the field. Moreover, it is unclear if these approaches offer real advantages compared to permissionless systems based on blockchain. For example, the *Tangle* design [51] provided in the first *IOTA* implementation has been recently superseded by a new one named *Coordicide* [52]. By the very admission of the *IOTA* principal investigator, one reason was that the "honest transaction majority condition" assumed in the previous approach (i.e., the majority of transactions always come from honest network participants) has to be realized at large scale turning out in a very inefficient

and energy-consuming system. Otherwise, it would be easy for malicious agents to buy enough hashing power and overtake the network.

In the context of permissioned systems, [Corda](#) represents another relevant system that does not adopt the blockchain as its ledger structure. In a Corda network, nodes have long-term stable identities, which they can prove ownership of to others: this way, the network entry process rules out Sybil attacks, and a consensus protocol like the proof-of-work is no more necessary. It is not even necessary to group transactions in blocks, a choice in Bitcoin and other permissionless systems motivated by having to keep the number of new block proposals to reduce conflicting blocks and related forks in the blockchain. In Corda, a transaction is in effect represented by a graph encoding all its transitive dependencies, and each transaction gets presented to a *notary service* that performs a breadth-first search over the transaction graph, downloading any missing transactions into local storage from the counterparty and validating them by checking their signatures. Corda allows for multiple notary services, possibly decentralized over coalitions of different parties using a consensus protocol of their own choice. Since each notary service provides for transaction ordering and timestamping services only for a part of the overall transactions managed by the network, the Corda ledger results in a set of multiple DAGs [53].

3.3. Consensus

Consensus in distributed computing systems refers to the process of achieving agreement about data/states which are managed and stored through a set of network participants so that such nodes share the same data/state as the correct one for the system at a given point in time. In the context of blockchain systems, the consensus has to guarantee that all the ledger replicas deployed through the system are the same (blockchain *consistency*), except possibly for a small number of terminal blocks still awaiting a final decision (*eventual consistency*). These consistent ledger replicas provide a certain degree of tolerance to unresponsive nodes and those that arbitrarily deviate from the application logic, thus representing a more resilient and secure alternative to the trusted third-party approach. Trust gets indeed split over multiple nodes, which must collude or get hacked simultaneously so that the system as a whole deviates from its correct functioning. In order for things to work, however, procedures that can filter out those nodes trying to undermine a correct agreement process are necessary, and this is precisely the goal of every consensus protocol.

A node can become unresponsive because of a network failure resulting in the disconnection of the node from the network, or as a consequence of a *crash fault*, where the node abruptly stops working without quickly resuming because of a power outage, a DoS attack, or an error in its software or hardware. Regardless of the nature of the fault that halted a node, a consensus protocol must terminate correctly as long as the percentage of halt failures does not exceed a certain threshold. As suggested by the intuition, a protocol for which consensus follows from a majority quorum, as it is common practice in the P2P model, can withstand a certain number of unresponsive nodes less than half of the total.

Much more challenging to manage than nodes subject to halt faults are those which try to hinder a correct consensus process thanks to subtle and illegal actions of various types, ranging from the tampering of existing transactions or the creation of false ones (as in double spending) to the modification of the protocol itself, through message manipulation and reordering. These actions may include not responding to messages sent by the other peers, thus encompassing halt failures.

In general, some consensus protocols tolerate unresponsive nodes but not malicious ones (e.g., Zab [54], Raft [55]), whereas a protocol designed to withstand malicious nodes will also tolerate unresponsive nodes to a greater extent. Depending on the blockchain consensus protocol, we can measure fault tolerance as a percentage of the number of malicious nodes or their used resources from the total number of nodes or

resources involved in the consensus process. This measure reflects the first fundamental discrimination for blockchain consensus protocols, which we can divide into the two broad classes of *membership-based* procedures and *fitness-based* procedures, in addition to a further class of hybrid methods obtained by combining the two previous approaches.

In membership-based protocols, a predetermined and reduced set of nodes with explicit identities and roles, often coordinated by a leader (a.k.a. primary) node that acts as an interface with clients, form a committee reaching agreement through multiple rounds of message exchange to achieve a consensus. In this respect, the committee represents a server consortium that has to order transaction proposals as they arrive by clients, arranging them in a valid chain of blocks despite some servers failing. Accordingly, the correctness of consensus for these protocols can be formulated by recasting two concepts introduced in [56]: (i) every server correctly executes the same sequence of (transaction) requests (*safety*), and; (ii) all requests are served (*liveness*). In this context, we can call malicious nodes and faults “Byzantine”, whereas consensus protocols that exhibit some degree of tolerance to them are called *Byzantine fault tolerant* (BFT) protocols. This terminology has come into use after [57], in which the authors introduce the consensus problem making the analogy where a group of generals of the Byzantine army must agree on a common battle plan, communicating only through messengers and despite some traitors hiding among them generate conflicting messages to hamper agreement.

It is shown in [58] that BFT protocols cannot tolerate a critical number of Byzantine nodes greater than or equal to one-third of the total number of nodes involved in the server consortium. Another prominent characteristic of these protocols, and more generally of membership-based protocols, is that they require explicit message passing among the N nodes composing the consortium, typically exhibiting a $O(N^2)$ message complexity which turns out in a considerable networking overhead already for consortia of a few tens of nodes. These protocols are therefore suitable only for permissioned blockchains and can offer a limited degree of decentralization. As for the bright side, membership-based protocols offer *consensus finality*, in that each consensus decision, as returned by the protocol, is conclusive and cannot be subject to changes or cancellation at later times. Thus, a block of transactions chosen to get transcribed on the blockchain can no longer change or get canceled later, allowing block intervals and transaction throughput comparable to those of centralized systems. Fall in the BFT category some protocols (e.g., PBFT [59] and its derivatives) that follow the *state machine replication* (SMR) scheme formalized in [60], together with other more recent protocols having different message workflows (e.g., HoneyBadgerBFT [61] and the related protocol suite BEAT [62]).

Fitness-based consensus protocols have a design suitable for large and open sets of participating nodes whose exact size and members are not known a priori. The crucial element underlying this broad category of protocols is the use of a resource, internal or external to the system, which serves to quantify the fitness of the nodes (i.e., their influence or power) in the consensus process. Such a resource is necessary for a context of permissionless networks, where an arbitrary (and anonymous) node has the opportunity to participate in the management of the blockchain. It indeed serves to constrain node power to the quantity of the resource they wish or can invest in the consensus process, thus counteracting Sybil attacks and their derivatives.

Another central element that virtually distinguishes all protocols suitable for permissionless systems is the ability of the node adding a new block to prove publicly (i.e., to the rest of the network) that it is indeed enabled to do so based on its fitness. Accordingly, authors in the literature denote these protocols as *proof-based* consensus protocols.

Since these protocols operate on P2P networks, allowing the decentralized management of trust and services in a much broader and flexible way than in client-server models, they have been singularly studied in recent years precisely in conjunction with and about the development of DLT, of which they represent a core aspect. Indeed, since the activation of the Bitcoin network, research contributions have multiplied, with

dozens of new protocol proposals that in many cases have been implemented in functioning systems. Due to their importance for DLT, and their primary role for the enforcement of decentralization also besides such technologies, we will allude below to the main types of these protocols; the interested reader should refer to specific literature contributions for a more comprehensive and in-depth treatment of this topic (e.g., [3,63–65]).

Note that these protocols are categorized quite differently by different authors, depending on the feature or approach that, in their opinion, primarily distinguish or unite them. Indeed, these protocols often consist of more than one key element, and the choice of the one considered most relevant influences their classification. In what follows, we have chosen the classification that, in our opinion, best highlights the articulation of proposals in order to overcome security and performance issues concerning previous solutions. This way, we hope to shed some more light on the evolutionary trend of this crucial DLT research field.

Beyond the names assigned by their respective authors, often derived from the type of fitness resource that the protocol allows to verify publicly or “proof” (work, stake, activity), it is significant to note that researchers oriented towards designing more efficient and performing permissionless consensus protocols by recurring to approaches and tools borrowed from the membership-based protocols previously introduced. In this respect, the main distinction concerns the part of the network involved in creating the new blocks. In the Bitcoin protocol and the first alternatives based on crypto coins owned by users (*stakes*), the whole network is involved in this process. Indeed, in these protocols, any node in the network may check by itself if it is fit to add the new current block, giving public evidence of that. As this approach has resulted in poorly performing and computationally costly protocols, many subsequent proposals have adopted a scheme in which a tiny subset of nodes are selected and allowed to add a new block. Then within this “committee”, the nodes that will actually add new blocks, also known as *leaders*, are selected. Accordingly, we will divide fitness-based consensus procedures first and foremost in *network-oriented* and *committee-oriented* protocols, leaving at later levels the distinctions based on the type of fitness resource used and, in the context of committee-oriented protocols, on how committees and leaders are selected.

A list of some significant blockchain consensus protocols, along with their main features, is reported in Table 3.

Table 3: A list of some major network-oriented (NET) and committee-oriented (COM) blockchain consensus protocols.

NET protocol	Proposer selection	Validation check	Finalization rule
PBFT	Client request	Signature	Mutual agreement
Nakamoto	PoW solver	PoW solution	Heaviest chain
Ethash + GHOST	PoW solver	PoW solution	GHOST variant
PoET	TEE waiting time	TEE attestation	Longest chain
Ethash + Casper FFG	PoW solver	PoW solution	Casper FFG
COM protocol	Proposer selection	Validation check	Finalization rule
Tendermint	Round robin	Stake, signature	Mutual agreement
Chain of Activity	FTS output	FTS output	Longest chain
Proof of Authority	Certification	Signature	Mutual agreement
Honey Badger BFT	Client request	Signature	Mutual agreement
Ouroboros	FTS and PVSS	PVSS output	Longest chain
Algorand	VRF + agreement	VRF, signature	Mutual agreement
Ouroboros Praos	VRF output	VRF output	Longest chain
Snow White	PRF (application)	PoS solution	Longest chain
Delegated PoS	Stake delegation	Delegate eligibility	Mutual agreement

3.3.1. Network-oriented consensus protocols

These protocols do not rely on explicit message exchanges among a committee to establish consensus, but rather on an algorithm that allows nodes by themselves to determine if they may add the new current block of transactions, giving corroborated evidence of this to the other nodes thanks to a “proof” that can be easily and publicly verified. Typically, the algorithm involves solving a computational problem by the node, which can publish the solution found as proof that it is fit to insert the new block. In this regard, the solution to the problem must depend on the current transaction block, although not in a predictable way, and it must not get more manageable by the knowledge of additional data such as the blocks previously recorded on the blockchain. For these reasons, the algorithm makes use of a cryptographic primitive which returns pseudorandom outputs (precisely, a cryptographic hash function). The choice of the computational problem to solve is critical since a too complex problem affects the liveness of the system, while a problem that is too easy to solve weakens its consistency. In fact, in the first case, it could be challenging to find a solver in a reasonable time, while in the second, there could be many solvers and many potentially valid blocks to add to the ledger. Designers usually identify an adequate trade-off thanks to using a system resource for which the node must invest a specific effort and whose availability increases the probability of solving the problem in a relatively short period. This way, the amount of the resource mentioned above represents the *fitness* a node has in solving the problem, while its solution is a sort of ticket granting access to the final tender, which establishes who will be able to add the current block to the blockchain. Indeed, there can be multiple problem solvers that correspond as many current blocks, all potentially valid, but which must be selected according to a system rule to obtain a single shared list of consecutive blocks. It follows that network-oriented consensus protocols consist of the three following phases, according to a timeline composed of time frames in each of which potentially all the nodes in the network compete for the addition of a new block in the blockchain:

- *Proposal*: where a node assembles a block of transactions and attach proof to it demonstrating its investment of the fitness resource in such a task;
- *Validation*: in which nodes check the validity of the blocks proposed in the previous phase, alongside with their alleged proofs;
- *Finalization*: in which the network, following a rule usually based on the majority of participants, selects a unique block among those which passed the validation phase.

Because of their implicit agreement model, network-oriented consensus protocols scale much well than membership-based protocols to the number of nodes involved: their message complexity can be indeed limited to $O(1)$ for each node by exploiting the communication models described in Section 3.4. However, this comes at the price of both higher costs and minor performance. Consensus does not conclude with the validation phase, but it requires an extra finalization process which by its nature is not deterministic and could take time since it depends upon a large number of loosely coupled nodes. Therefore, it impacts negatively on the transaction throughput. Cost, on the other hand, does not depend so much on sending and receiving messages and on the local processing required about this, as in the membership-based protocols, but rather on the fact that it is necessary to counteract Sybil attacks and those related to them (e.g., double spending) thanks to considerable consumption of the fitness resource by the nodes. After more than a decade of research in the field, the actual cost of a proof-based protocol can be far lower than for the proof of work implemented in Bitcoin since it strongly depends on the choice and management of the fitness function. However, these costs are decidedly higher for network-oriented consensus protocols than those of the membership-based protocols, if only because many more nodes are involved in the consensus process.

Fall in this category the protocols at the core of many well established cryptocurrencies (e.g., [Bitcoin](#) [35], [Ethereum](#) [66], [Peercoin](#) [67], [Nxt](#) [68]), altcoin proposals never

came to light as Permacoin [69], as well as systems not intrinsically tied to a cryptocurrency (e.g., [Hyperledger Sawtooth](#) [70]). The fitness resource used by the nodes in these systems typically is hashing power or stake value; however, other choices are also possible, as the amount of stored data in Permacoin or the disposal of a trusted execution environment (specifically, Intel SGX [71]) in Sawtooth.

Nakamoto consensus was the pioneering network-oriented protocol; it makes use of a proof-of-work (PoW) algorithm to accomplish the proposal and validation phases, and the so-called *longest chain rule* for finalization, already discussed in Section 3.2. PoW algorithms employ a cryptographic hash function in order to enforce a search problem called *mining*, where nodes have to find a string called *nonce* so to get a hash digest whose value in hexadecimal has not less than a given number of leading zeros, known as the *difficulty target*. The hash function inputs the nonce and a digest resulting from the Merckle tree of the transactions in the current block and the hash digest of the previous block header so that the nodes can instantiate the search problem for the current block. In this way, they cannot circumvent or make it easier by prior computations, but they can only solve it by an iterated error-retry approach.

A *difficulty adjustment algorithm* tunes the difficulty target by computing at each mining node the time required by the network to generate a given number of previous blocks so that a solution is moderately hard to find by the totality of miners in a statistical sense. The goal is to balance the block delivery rate and the frequency in the occurrence of multiple valid blocks, which undermine blockchain consistency and burden the finalization phase.

In some PoW implementations, the hashing algorithm can have some additional properties in order to limit substantial increases in hashing rate thanks to the use of specific hardware, as the *ASIC resistance* provided by the *Ethash* algorithm in Ethereum, which also has other features according to the design rationale described in [72].

Bitcoin developer community conceived *Proof-of-Stake* (PoS) algorithms well before bitcoins became the best known and most used cryptocurrency in the world, as alternatives with much lower and sustainable energy consumption. The term “stake” refers to digital coins or tokens owned by participants and managed through the blockchain network, used as a counterweight to hash power. The difficulty of mining activity is inversely proportional to the number of crypto coins owned by the miner at the time; this way, a virtual resource replaces physical ones (specifically computational devices and electric power).

A pure PoS-based system follows the same principles and construction as Nakamoto consensus, with the sole exception of the implementation of the hashing challenge: the calculation of the hash function can be performed only once for each unit of time (e.g., one second), according to a virtual shared timer realized by keeping synchronized within a given margin of tolerance the local clocks at each node, while the difficulty to solve the search problem is inversely proportional to the stake value invested by the node in the challenge. In this way, PoS avoids the brute-force hashing race characteristic of a PoW, but the significant savings in energy and hardware come at the price of worsening some issues (e.g., selfish mining and centralization) and a series of possible attacks to which PoWs are immune. Indeed, the public availability of staking history in the blockchain weakens the PoS pseudo-randomness while simultaneously incentivizing corruption: two factors exploited by the *stake-grinding* [73] and *bribing* [74] attacks, respectively. On the other hand, the virtual nature of mining paves the way for low-cost simulations: a malicious node may create an alternative branch to the main chain starting at any point of its choice without any substantial cost. That, in turn, makes it easier for users to overtake the main chain by collusion as in the *long-range* attack or to bet in multiple competing chains simultaneously by fractionating their stakes as in the *nothing-at-stake* attack, which in both cases hinder the resiliency of the system to double-spending [75]. This resiliency is theoretically equal to 50% of the fitness resource for both PoW and PoS systems: (pool of) attackers must have more than 50% hashing power or stake value to

mount successful double-spending attacks. However, malicious mining strategies as the selfish mining for Bitcoin [41] and those at the core of the two previous attacks can substantially reduce that threshold: in [41], the authors prove that a colluding group of any size could eventually compromise the Bitcoin network under some circumstances, and a similar result presumably also applies to many other PoS and PoW systems.

By allowing a limited set of users to take control of the network, selfish mining and other illegal mining strategies ultimately threaten the decentralization aimed by permissionless blockchains. However, a centralization risk for proof-based consensus systems stems from the “wealthy get wealthier” paradigm: nodes having more fitness (hashing power or stake value) than others can lawfully exploit this advantage to accumulate more and more wealth over time, potentially reducing the system to an oligopoly. This attitude is particularly true for PoS systems, where miners can reinvest their profits into staking over and over again.

Nxt [68] tries to mitigate some of the above issues thanks to the following rules: (i) designers determined token supply at the origin, and the reward for inserting a new block only comes from transaction fees; (ii) users cannot split stakes over multiple simultaneous mining activities, and; (iii) the stake value is appreciated due to the mining activity during a block cycle, but it is reset to the base value once the cycle ends.

These and other similar rules can encourage users to participate and honestly validate transactions; however, they do not overcome the two primary limits of current network-oriented systems. The first limit concerns performance: the fact that the consensus process potentially involves all the network nodes entails a considerable time for the finalization phase. Even adopting particular strategies for determining the main chain, such as GHOST or the variant currently implemented in Ethereum (see section 3.2), the times for confirming a transaction are orders of magnitude greater than for membership-based systems. The second limit concerns security: the arguments supporting the security of these systems are indeed only of a heuristic nature due to the lack of formal reference models and precise definitions. Therefore, these systems could prove vulnerable by design at a later time, with potential irremediable damage to their investors, which the use of a cryptocurrency can exacerbate. The case of selfish mining for Bitcoin is exemplary in this respect: this network has not yet been the subject of such an attack, presumably solely because its users believe that its longevity and reputation can safeguard their earnings.

Therefore, two somewhat complementary principles guided research efforts in recent years: (i) draw inspiration from skills and methodologies acquired in the thirty-year development of membership-based consensus systems, and; (ii) exploit formal mathematical proofs based on adequate cryptographic primitives to obtain provably secure consensus protocols.

These led to the protocols briefly discussed in the following section.

3.3.2. Committee-oriented consensus protocols

These protocols combine the proof-based approach with the notion of “committee”, already used in membership-based protocols, by adhering to the following high-level procedure, an alternative to the one presented in the previous section.

- *Committee selection*: to designate a (relatively small) subset of nodes in the network as eligible to participate in the consensus process for adding one or more upcoming new blocks;
- *Leader selection*: to designate a node in the committee as the one which is actually in charge of adding the next block or of leading the consensus process for adding the next block;
- *Proposal*: where the leader node returns the new current block along with a proof to demonstrate the block validity;
- *Validation*: in which nodes check the validity of the blocks proposed in the previous phase, alongside with their alleged proofs;

- *Finalization*: in which the network, following a rule usually based on the majority of participants, selects a unique block among those which passed the validation phase.

It is essential to notice here that committee-oriented consensus protocols cannot correctly work without the finalization phase. Unlike in membership-based consensus, a committee might not always get uniquely defined for the entire network, and this could also happen, within a given committee, for the leaders appointed to add the new block. Indeed, in these protocols, the committee is selected starting from a piece of information shared between the nodes thanks to the blockchain to take into account the distribution of the fitness function between them. However, the “blockchain view” required by a given protocol may not be the same for all the nodes in the network because of blockchain forks due to communication delays or illegal behaviors. In other words, these protocols can guarantee that each election gives rise to a single committee only with a specific (albeit high) probability, not absolutely. Similarly, depending on the selection method used for choosing a leader within a committee, this could result in more than one leader and as many respective potentially valid blocks.

In this category fall many protocols for which the committee selection works by taking into account the stake values of the participating nodes. Since the committee has to be chosen so that an attacker can hardly predict the participating nodes, to avoid being subject to corruption or boycott, committee members get usually selected by a pseudo-random process whose output can be publicly verified. Thus, most of these protocols rightfully belong to the PoS class (e.g., [76–79]), whereas few exceptions like *Tendermint* [80] cannot be adequately addressed as PoS systems, although they use the stake values to select the committee. Other committee-oriented systems consider fitness measures alternative to the distribution of stake values among participating nodes. An example is the *proof-of-authority* [81] by the co-founder of Ethereum Gavin Wood, where committee members have to go through a mandatory certification process to build up their authority.

We can trace back the idea of selecting a committee to add a new block to *Tendermint* and the *proof-of-activity* protocol [74] in 2014, although, in these systems, the committee is only implicitly defined and dynamically configured for each new block. In *Tendermint*, the consensus process gets divided into time intervals called rounds, for each of which a chosen leader (also known as proposer) proposes the new block, starting a voting process involving the other nodes in the network through a BFT algorithm. The leader selection function is weighted round-robin, where nodes rotate proportionally to their voting power, and the voting process can only return reliable outputs if less than one-third of all voters are dishonest, as follows from the tolerance limits for BFT protocols. This last condition is challenging to verify, given that the voting nodes are not defined a priori but depend on the expected times for carrying out the various phases of the BFT protocol. Furthermore, the deterministic nature of selecting leaders makes it necessary to protect their true identities by using proxy nodes that are not supposed to divulge those identities. Therefore *Tendermint*’s consensus protocol appears unconvincing, despite the alleged proof of its security given in the more recent draft [82].

Other proposals inject randomness in selecting leaders (or committees) to mitigate corruption or boycott attacks.

The *chain-of-activity* (CoA) protocol [83] uses the pseudo-randomness tied to a sequence of previous blocks in the blockchain (e.g., by computing the hash digest for each block and taking the leading bit of each digest) to generate a shared sequence of N pseudo-random seeds. To protect against possible blockchain forks, each node following the protocol must consider a sequence of blocks such that the last one precedes a predetermined number of final blocks in the received chain. The N seeds are then iteratively passed in input to the *follow-the-satoshi* (FTS) algorithm [74] to get the sequence of leaders responsible for building the following N blocks. The FTS algorithm treats the string received in input as an index of a previously coined satoshi, inspecting the

blockchain to determine which user this satoshi belongs to at the current moment. In this way, the system selects the sequence of leaders, taking into account their respective stake values, but in a way that it is difficult for an attacker to determine their identities until just a short time interval before they contribute to the construction of the blockchain.

Ouroboros [76], the consensus protocol exploited by the *Cardano* cryptocurrency, follows a similar approach to that of CoA. In this case, a *publicly verifiable secret sharing* (PVSS)[84], where the participants in the consensus process and everybody in the network can verify the distribution of the shares, determines the sequence of shared seeds passed in input to the FTS algorithm.

PVSS is a particular case and a building block of the *multiparty computation* (MPC) [85] framework, concerned with enabling some different and untrusted parties to carry out a joint computation of a given function. MPC assumes a threat model very similar to that of BFT protocols, where some malicious entity, even in the form of a subset of participating nodes, tries to hamper the correct computation of the joint function to get a prescribed output or acquire some private information. Accordingly, two essential requirements of secure MPC protocols are *privacy* and *correctness*: the privacy requirement states that parties should learn the function output and nothing else, while the correctness requirement states that each party should receive its correct output.

The public use of the FTS algorithm allows *Ouroboros* to implement the privacy requirement only at a time before the actual work of the new committee of leader nodes, leaving this system as CoA unprotected to the so-called *adaptive corruption attacks*. Moreover, by requiring strong network synchrony for leaders to use precisely their time slots, *Ouroboros* is vulnerable to *desincronization attacks*.

On the contrary, *Snow White* [79] works in *partially synchronous* networks (see Section 3.4) and, thanks to the *sleepy* [86] consensus model, it can also accommodate sporadic participation in which nodes can arbitrarily switch from online to offline status. However, *Snow White* is not able to completely counter adaptive corruption attacks. *Snow white* first determines the composition of the committee of nodes allowed to add new blocks during a time-lapse known as *epoch*, thanks to a fitness function defined at the application layer and to the fitness status of the nodes as deduced by inspecting a blockchain prefix. Then, within each epoch ϵ , it selects the nodes in charge of adding the new block at time $\tau \in \epsilon$ as the solvers of the hashing challenge $H_\epsilon(\tau, pk) < D$, where H_ϵ is a hash function which depends upon the epoch, pk is the node public key, and D determines the challenge difficulty. This way, just before a new epoch begins, an attacker can determine which nodes in the committee will have to insert the new blocks relating to the time slices for that epoch.

Ouroboros Praos [78] is a protocol that, thanks to an analysis of *Ouroboros*, overcomes its two above issues and requires far fewer message exchanges between nodes to accomplish the MPC procedure. First, each elector executes *averifiable random function* (VRF) [87] locally to define its proposal for the following sequence of leaders, so that each other node can verify if got chosen as a leader with respective slots, without being able to determine the other leaders in the sequence. This mechanism turns out in a secure MPC procedure with the requisites of privacy and correctness, assuming that the majority of the stake value belongs to honest participants. Second, some empty time slots up to a maximum value are allowed for each epoch to help electors resynchronize when necessary, thus resulting in a protocol designed for partially synchronous networks.

Like *Ouroboros Praos*, *Algorand* [77] is another provably secure PoS protocol designed for networks that are not firmly synchronous and where attackers can also perform adaptive corruption. *Algorand* too uses a VRF scheme to select the committee randomly, although taking into account the stake values of participating nodes, called cryptographic *sortition*. However, in this case, it is the committee as a whole to decide which block to add thanks to an appropriate BFT algorithm, not a single leader, which turns out in a protocol reaching consensus finality in a slightly stronger assumption than partial synchrony known as *weak synchrony*.

Two other relevant PoS protocols falling in the committee-oriented class are the *delegated proof-of-stake* (DPoS) used in different systems (e.g., [bitshares](#) [88] and [EOSIO](#) [89]) and *Casper FFG* [90], which will appear in the upcoming version 2.0 of Ethereum, also known as *Serenity*. DPoS protocols were already discussed in Section 2.2, as they are generally used to define the system's governance. Casper FFG is an addition to the Ethash PoW which incorporates PoS into block finalization. Ethereum developers intended to change the variant of GHOST currently implemented with a protocol where a set of validators use a BFT algorithm to select the new block, guaranteeing backward compatibility.

3.4. Network

The network layer is the foundational infrastructure of blockchain systems, which establishes the role and functions of its nodes, the way they communicate, and how application workloads get partitioned among them. This layer also serves to specify the modality of participation of nodes (continuous, sporadic) and their ability to synchronize each other through the exchange of messages. In particular, this layer concerns the synchronization assumption underlying the design and usage of a specific consensus protocol by indicating the way and amount of delays affecting message transmissions. We can also consider other characteristics for qualifying or quantifying the reliability of communications, and therefore also concerning the threat model considered, implicitly defining the capabilities of attackers to hinder the consistency and finality of consensus processes.

3.4.1. Network architecture

Permissionless blockchains adopt a pure peer-to-peer (P2P) architecture [91]. Every node can be both supplier and consumer of resources, and it is an equally privileged and equipotent participant in the application that can freely join and leave the network without any authorization. The only relevant difference is between *full nodes* and *lightweight nodes*: the first ones maintain a full copy of the ledger, while the second ones download just a subset of it so to verify if one or more transactions have been included in a block, possibly thanks to the support of a full node.

Permissioned blockchains adopt instead a hybrid model, where nodes must be authorized first and then participate with revealed identity, and often with a specific role and functions. Conversely than in permissionless systems, in this second type of blockchain network, a set of client nodes can generally be distinguished from a server component made up of a multiplicity of nodes divided according to their specific service function.

In both cases, every node (or peer) in the network operates autonomously for a given set of rules encompassing the communication protocol and other functionalities such as transaction processing, consensus participation, and ledger management. These last possibly depending on its role.

3.4.2. Communication models

From the communication viewpoint, peers implement some form of virtual overlay network on top of the Internet, which allows them to exchange information directly through indexing and neighbor node discovery, regardless of the physical network topology. In order to optimize performance and scalability, peers connect only to a limited number of their neighbors through a *flooding* or *gossip* protocol [2]. In flooding protocols, messages get transmitted to all the nodes recognized as neighbors, while gossip protocols relay messages only to a subset of randomly selected neighbor nodes. Fast and pervasive spreading of transactions over the network is a crucial factor for effective management of the ledger status through the consensus protocol, and both approaches are appropriate, although with different transaction bandwidth and delay performance. As a rule of thumb, flooding protocols are more bandwidth-intensive

but less prone to network splitting attacks than gossip protocols because of the more connections involved in transaction propagation. Typically, however, the relative importance of bandwidth and delay of messages depends on the type of traffic workloads; moreover, bandwidth gains are at the expense of delayed transmissions, as in the case of the announce-and-request signaling implemented in the Bitcoin flooding protocol, which requires two more communication steps. Therefore, different systems often implement different protocols, and the choice should propend to obtain the best performance compromise according to the type of network and the functions offered at the application level. In general, transaction flooding best fits in permissionless systems because of its major pervasiveness and resistance to network splitting attacks. In contrast, message dissemination through gossiping goes well for permissioned networks thanks to minor bandwidth requirements and a faster reaction in case of faulty nodes or slow network links. Ethereum tries to take the best of both approaches by adopting a hybrid protocol, in which some messages are sent to all neighboring peers while others go to a limited number of them.

Transaction propagation in blockchain networks has to deal also with privacy and security issues. Transaction payloads can be encrypted depending on whether the blockchain is private or public and the nature of communications among subsets of peers in the network. Nevertheless, in any case, peers must exchange messages with confidence about the authenticity of senders and receivers and the integrity of data payloads. Moreover, depending on the implemented P2P protocol, permissionless networks can be more or less susceptible to techniques that bind IP addresses to blockchain accounts, thus being detrimental to the alleged anonymity of their users. In this respect, flooding usually offers more protection than gossiping, presumably the main reason Bitcoin and the first generation of cryptocurrencies adopted it. However, if a significant amount of transaction traffic is collected and analyzed, the association between blockchain accounts and their IP addresses can be determined regardless of the P2P protocol, as shown in [92,93].

3.4.3. Communication uncertainty

The mere fact of considering a distributed system makes it necessary to consider failures that may concern the nodes of which it is composed or the devices through which they communicate. Some nodes participating in the consensus process or an external attacker may cause malfunctions in communications, manifesting in delays or non-delivery of messages. Therefore, a proper design or deployment must consider assumptions concerning the way and amount of delays affecting message transmissions that the consensus protocol can tolerate without undoing its consistency and finality.

The theory of distributed systems considers the three following main communication uncertainty models, plus some minor variations obtained by strengthening or weakening in some way the assumptions related to these models (e.g., the weak synchronous model considered in [77]):

- *Synchronous*, which assumes the existence of some a priori known finite-time bound Δ , such that each message gets transmitted with a delay at most Δ ;
- *Asynchronous*, where the delivery messages might require any finite amount of time so that each message must eventually be delivered, without prediction of time needed to get to the destination;
- *Partial synchronous*, whose assumption is that there exists some known finite-time bound Δ and a special event called *Global Stabilization Time* (GST) such that: (i) the GST eventually happens after some unknown finite time, and; (ii) the delivery of any message sent at time τ must happen by time $\Delta + \max(\tau, GST)$.

The synchronous assumption is, in practice, too optimistic unless the value of Δ is overestimated, with the result of significantly worse performance. On the other hand, underestimating the value of Δ can negatively impact the consistency of a consensus protocol since nodes receiving messages after the Δ bound will miss some protocol steps.

Notably, all protocols derived from the Nakamoto consensus rely on synchrony for their security (a relatively simple proof of this is in [94]).

Conversely, the asynchronous assumption corresponds to the worst-case scenario where we cannot predict the number of network delays. Consensus protocols designed under this assumption (e.g., HoneyBadgerBFT [61]) are usually very robust in terms of consistency and finality since they depend on neither time-bound for message delivery nor fixed value for timeouts. However, that comes at the price of higher complexity and lower performance.

Two direct results mark the gap between consensus protocols in the synchronous and asynchronous settings. Fisher, Lynch, and Paterson in 1985 [95] showed that any protocol that solves consensus withstanding the hang failure of just one node could not terminate in the asynchronous model. By contrast, several membership-based consensus protocols (see Section 3.3) introduced over the years can withstand up to half minus one dishonest member in the synchronous model (e.g., [96–98]).

Partial synchrony was introduced for consensus protocols in 1988 by [99] as a middle ground between the synchronous and asynchronous models. It indeed assumes that the system behaves asynchronously up to the occurrence of the GST, which turns out to be equivalent to say that there is some finite but unknown upper bound on message delivery. Over time, that has proved to be an excellent assumption for designing practical distributed systems, which usually are built on networks behaving synchronously except in case of unexpected events such as a DoS attack or the temporary unavailability of one or more communication links. Indeed this model allows designing consensus protocols that always guarantee consistency, providing liveness and finality only after the GST occurrence.

4. Emerging concepts and their relevance in applications

The simple Bitcoin transaction workflow sketched at the end of Section 3.2 was questioned in late 2013 by Vitalik Buterin, a programmer and co-founder of Bitcoin Magazine. Buterin argued to the Bitcoin core developers that many other applications and business sectors besides money transfer and fintech could have benefited from blockchain technology, provided, however, that they had a more powerful and flexible language than the Script language [100] provided for Bitcoin. After some work spent trying to bring out this point of view within the Bitcoin community, with his short involvement in the Colored Coins project (see Section 4.1.2), Buterin eventually conceived Ethereum as an alternative general platform for decentralized applications [42].

At the time of writing –shortly before the entry into service of the 2.0 version, which should modify the consensus protocol transitioning from a PoW to a hybrid PoW-PoS system– Ethereum looks like Bitcoin in terms of data structures and protocols for managing the ledger. However, it differs substantially in the application stack, especially for account management and transaction life cycle. Rather than enforcing one specific set of rules targeted toward the creation and usage of a cryptocurrency, Ethereum makes use of a cryptocurrency to allow users to write programs specifying whichever rules they want, upload the programs to the blockchain, and let the blockchain interpret and enforce the rules for them in a decentralized fashion. In order to achieve that, Ethereum design is far more complicated than Bitcoin. Transactions are no more designed to transfer coins by managing the addressing of UXTOs and their corresponding values but rather are cryptographically signed instructions by real accounts in the system to update the state of the network. Furthermore, in addition to having accounts associated with human users called *externally owned accounts* (EOAs), Ethereum also provides a *contract-type* account: while an EOA state keeps the account's balance in *ether*, Ethereum's internal crypto-token, the state of a contract-type account saves the contract's code and storage. The contract's storage is encoded as a key-value database, whereas the contract's code gets encoded in Solidity [101], an object-oriented, high-level language for implementing smart contracts. These arbitrary transition rules in Ethereum govern the behavior of

accounts within the system state. All the accounts and smart contracts which compose the Ethereum state live in the Ethereum Virtual Machine (EVM) [102], a particular distributed state machine deployed through the Ethereum clients and whose interpreter is Solidity. Unlike Script, Solidity is a Turing-complete programming language, for which it is not possible to predict whether a program will terminate or not (the Halting problem undecidability proven by Alan Turing in 1936 [103]). Thus, an inattentive or malicious programmer could create a contract containing not-terminating code, send a transaction to the contract, and crash the system. More generally, users could execute very demanding programs in terms of computational resources, causing system overloads and slowdowns. The solution in Ethereum is the concept of *gas*: the transaction sender specifies a maximum amount of computational resources (computing time and memory space) that they authorize the code execution to take and pay a transaction fee in ethers that is proportional to the number of resources spent. Ethereum must therefore be considered a *computing-as-a-service* platform where ethers serve as an incentive for peers who manage the system and as a currency for processing. This radical change from Bitcoin's approach has sparked new perspectives in using blockchain technologies while entailing the evolution of some previous notions and tools plus the emergence of new ones.

4.1. Assets and Tokens

In financial accounting, an *asset* is anything tangible or intangible owned or controlled since it has positive economic value, usually quantifiable as a given amount of money, and cash itself is also considered an asset. It follows that, since their beginning, blockchains have been systems conceived and aimed at managing certain types of assets. Bitcoin had a native design to manage just the bitcoin cryptocurrency. However, a few years after its launch, several researchers – including Vitalik Buterin mentioned above – were trying to expand the system to allow it to manage other types of digital assets, comprising those suitable for representing physical assets in a unique and unalterable way. In 2012, Meni Rosenfeld [104] introduced a mechanism called *Colored Coins* to take advantage of the decentralized, trustless trading of bitcoins by allowing users to issue transaction outputs tagged for particular purposes, corresponding metaphorically to the emission of (fractions of) bitcoins having a unique “color”, and to state their obligations to owners of coins of these colors. Today, the term “Colored Coins” loosely describes a class of methods for representing and managing real-world assets on top of Bitcoin, thanks to the fact that Script allows storing small amounts of metadata on the ledger, which can represent asset manipulation instructions. There are different labeling schemes (e.g., OBC, EPOBC, Open Assets [105]), but in any case, their purpose is to make it impossible to have coins of a specific color in a different way than receiving coins already of that color, following a chain of transactions traceable back to the color's genesis. This way, for example, it is possible to encode in a transaction that a certain amount in units of a new asset was issued, attaching a real-world value to those units by the asset issuer's promise to redeem them for some goods or services.

Since the emergence of more general-purpose and feature-rich blockchain systems, the idea of Colored Coins and their related tools have lost much of their interest; however, they represent the first attempts to find and build applications based on a cryptocurrency paving the way for an instrumental DLT concept. Today, this concept has reached its maturity and development in implementing Ethereum tokens and Hyperledger Fabric assets, respectively, in the context of permissionless and permissioned systems.

4.1.1. Ethereum tokens

Essentially, tokens are programmable digital assets built on top of a digital ledger like a blockchain, whose life-cycle depends on smart contracts called *token contracts*. Therefore, other than representing a medium of exchange such as a cryptocurrency, tokens are used to trigger certain functions in the smart contract(s) for a decentralized

application, guaranteeing their authenticity and non-clonability. In the Ethereum system, tokens can represent a diverse range of digital assets and real-world, tangible objects. Depending on their scope and function, they are created and managed through a standard application programming interface (API) within smart contracts. The standardization process helps ensure tokens remain *composable* and *interoperable*. Their related smart contracts can be composed of and interoperate with any other standard-compliant smart contract on the Ethereum network. These standards are provided by and for the Ethereum community in EIPs; specifically, contract standards are Ethereum improvement proposals of category ERC, which encompasses application-level standards and conventions. Hence, they are often referenced with ERC instead of EIP in literature but with the same numbering. The most relevant Ethereum token-related standards that at the time of writing have reached the “final” status are:

- *EIP-20: Token Standard* [106] is an API for *fungible* tokens, i.e., tokens all identical in specifications that can get expressed into smaller units, and both tokens and units can be implicitly interchangeable provided that they sum up to the same value. The first and foremost example of a fungible token is cash, and all other examples (e.g., game chips, cryptocurrencies, commodities, common shares) are functionally equivalent to it. Accordingly, this API provides functionalities like the tokens transfer from one account to another, the current token balance of an account, and the total supply of the token available on the network; other functionalities include creating and checking allowances. EIP-20 composes of six mandatory functions, three optional functions, and two events.
- *EIP-721: Non Fungible Token Standard* [107] is an API for non-fungible tokens (NFTs). NFTs are used to identify something or someone uniquely since this kind of token comes up in unique samples, even if generated from the same smart contract. NFTs can be used to encode and manage the ownership of digital files such as collectible items (e.g., photos, audio, and other digital artworks), access keys, nominal tickets. However, access to any copy of the original file does not restrict to the buyer of the NFT. Thus NFTs provide the owner with proof of ownership that is separate from copyright. EIP-721 provides core methods that allow tracking the owner of a unique identifier and a permissioned way for the owner to transfer the asset to others. This standard requires compliant tokens to implement ten mandatory functions and three events; moreover, they must implement the *EIP-165: Standard Interface Detection* [108].
- *EIP-777: Token Standard* [109] defines advanced features to interact with tokens, while remaining backward compatible with EIP-20. This API introduces operators for sending tokens on behalf of another account and sending/receiving hooks offering holders more control over their tokens. It composes of thirteen mandatory functions and five events.
- *EIP-1155: Multi-Token Standard* [110] outlines instead an interface that can represent any combination of fungible and non-fungible tokens in a single contract and transfer multiple token types at once. This API serves as an alternative to EIP-20 and EIP-721, which require the deployment of a separate contract for each token type or collection, placing much redundant bytecode on the blockchain and limiting certain functionalities in some usage scenarios (e.g., blockchain games). It composes six mandatory functions and four events.

Apart from the accepted token standards, the Ethereum community is actively working on more advanced token functionalities. Particularly relevant in this context are securities-related tokens, whose goal is to code assets such as a debt or equity claim on the Issuer and must be therefore more heavily regulated than other tokens. The followings are two interesting token standards for securities currently at the “draft” stage:

- *EIP-1450* [111] extends EIP-20 in order to facilitate the issuance and trading of securities in compliance with the U.S. Securities Exchange Commission (SEC) regulations. According to SEC requirements, this proposed standard mandates that the Issuer is the only role that may create a token and assign the Registered Trade Agent (RTA), while the RTA is the only role that is allowed to execute the mint (creation), burn (destruction) and transfer of tokens. Implementers must maintain off-chain services and databases that record and track the Investor's private information (name, physical address, Ethereum address, and security ownership amount) so that: (i) implementers and the SEC can access it on an as-needed basis; (ii) issuers and the RTA can produce a current list of all investors with names, addresses, and security ownership levels for every security at any given moment, and; (iii) issuers and the RTA can re-issue securities to investors for a variety of regulated reasons.
- *EIP-1462* [112] is another extension to EIP-20 that provides compliance with securities regulations and legal enforceability. However, in contrast to EIP-1450, this API defines a minimal set of additions to comply with the U.S. and international legal requirements. Such requirements include Know Your Customer (KYC) and Anti Money Laundering (AML) regulations, besides the ability to lock tokens for an account and restrict them from transfer due to a legal dispute. Another important requirement provided by this standard is the ability to attach additional legal documentation to set up a dual-binding relationship between the token and off-chain legal entities.

Ethereum's tokens idea is to have different digital assets to effectively represent the state and life cycle of diverse kinds of corresponding assets in the physical world, with their specific characteristics and behaviors. Just as ethers and fungible tokens intend to offer a more efficient and secure alternative to money transfers, mainly if related to fiat-backed stablecoins, NFTs and tokens related to securities aim to optimize the transfer of tangible or intangible assets, improving the assessment of their status and ownership.

In particular, NFTs can unlock the value of many kinds of assets by facilitating their sale and creating new markets. For example, digital artworks: a criticism often raised concerns the ineffectiveness of an NFT in guaranteeing the actual possession of such assets, given that digital information can be replicated and shared at will. However, a closer look shows that the real problem does not derive from the NFT approach but the source of the data, which is the artist who produced the work. Assuming the originality and uniqueness of the artwork, as guaranteed by its creator, an NFT could, in principle, associate a single owner with this work and allow only the latter to have (full) access to it. In this case, of course, the artwork should be stored in a data repository with confidentiality protection, where the owner defines access rules and, for example, could decide to make copies at a reduced resolution of the original for public download or allow just viewing the original artwork through special devices at specific locations. Moreover, the concept of ownership should not be confused with usability: the fact of being able to see a piece of art freely certainly does not give the right to sell it or show it for a fee.

These and similar arguments should convince skeptics that issues and measures concerning a digital artwork (authenticity, rating, fruition, protection) are entirely analogous to those for a physical piece of art and that NFTs can help in implementing their management. The market for digital artifacts supported by NFTs is, in fact, already a very profitable reality, as shown for example by the [CryptoKitties](#) case and by [Christie's](#) online sale for \$69 million of the file "Everydays: The First 5000 Days" by the graphic designer Mike Winkelmann (known as Beeple).

As we will see in Section 4.3, reliable and secure management of digital artworks through NFTs and, more generally, the registration of authentic off-chain data on the blockchain can be guaranteed thanks to the use of appropriate trusted third parties called *oracles*. A different approach is instead necessary to manage physical assets, as in this case, it will be necessary to guarantee a unique and unalterable link between

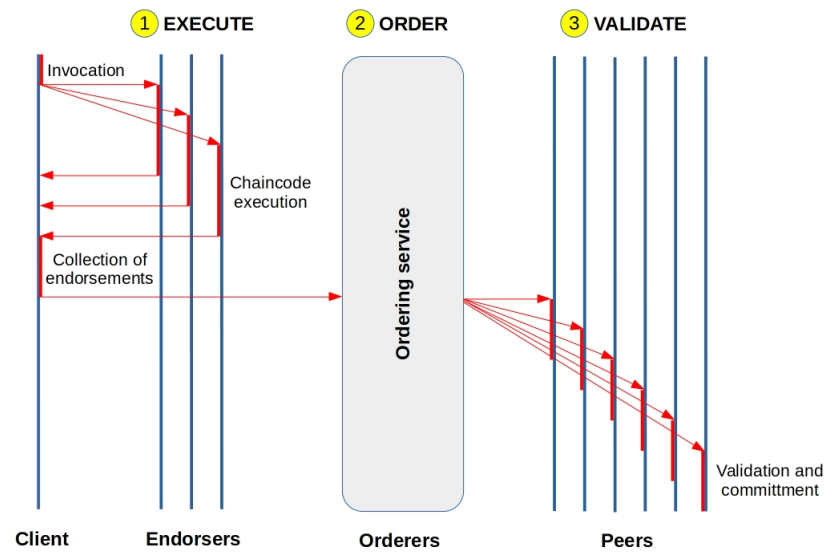


Figure 5. The transaction processing in Fabric entails three types of nodes and as many steps as follows: (1) peers acting as endorsers check for the correctness of a transaction proposal by a client through the verification of its fields and the execution of the operation(s) provided in its payload; (2) orderers assemble all the endorsed transactions provided by clients in ordered blocks, where each block contains a reference to the previous block, and; (3) peers check the ordered blocks for valid transactions and store the resulting blocks in their local copy of the blockchain.

the object and the token representing it. Of course, such a link must connect to the object's specific and unique physical characteristics, easily encodable as digital data. If, for example, a distributed application concerns the buying and selling of cars, then the NFT representing a car could store the license plate number and the chassis number of the car, plus some descriptive elements aimed at identifying the type of car (e.g., make, model, color), in a similar way to what is recorded today on the documents certifying the ownership of a vehicle. However, there are cases where more friendly and secure mechanisms can be put in place to build the tamper-proof binding between the physical object and its token, and this gives rise to the notion of *anchor* presented in Section 4.2.

Resorting to anchors and oracles may be necessary not only in the context of the NFT technology pursued by Ethereum and similar systems. They also concern different approaches for using a blockchain, such as the one pursued in Hyperledger Fabric which we are going to discuss in the next section.

4.1.2. Assets and tokens in Hyperledger Fabric

Hyperledger Fabric was started in 2015 as an open-source project under the aegis of the Linux Foundation to design a blockchain system suitable for supporting interoperability between companies belonging to a consortium organization.

Therefore, its design goals were from the beginning different from those of the blockchain platforms prevailing at the time, offering some key differentiation capabilities in the management of both network nodes and assets. This system adopts a permissioned network model in which nodes with different functions cooperate to commit transactions on the blockchain. The resulting *execute-order-validate* workflow, illustrated in Figure 5, allows overcoming the following major drawbacks of the order-execute paradigm commonly adopted in blockchain systems [113]:

1. the trust model for transaction validation is determined at the consensus layer rather than at the application layer, thus forcing transactions to follow a validation workflow which can diverge from the actual application requirements;
2. consensus is hard-coded within the platform, resulting in monolithic applications which are difficult to upgrade and adapt to different network environments and threat scenarios;
3. the fact that all peers should execute all transactions degrades performance increases consumption, and makes it more challenging to enforce privacy when necessary;
4. smart contracts must be written in a fixed, non-standard, or domain-specific language to avoid the occurrence of non-deterministic or non-terminating executions.

The first of the above issues can find a resolution through the *execute phase*: depending on the trust model at the application layer, an *endorsement policy* in the smart contract specifies which are the nodes, equipped with a local copy of the blockchain and running the smart contract, that have to check for the correctness of a transaction proposal issued by a client. These nodes, also known as *endorsers*, verify the correct encoding of the transaction proposal and its proper execution, returning the client a proposal response. The client then collects endorsements until they satisfy the endorsement policy, proceeds to create the transaction, and passes it to the *ordering service* to give rise to the *order phase*. In this phase, the ordering service establishes a total order on all submitted transactions and outputs transaction blocks suitable to be stored in the blockchain, using a consensus protocol to tolerate faulty or byzantine orderers. Since orderers assemble transactions without verifying them, the consensus protocol can be unplugged from the application layer, thus overcoming the second issue.

The ordering service then sends transaction blocks to the network for the *validate phase*, in which the nodes record the blocks in their local copy of the blockchain after tagging their transactions as valid or invalid. This second validity check compares the endorsement responses in the transactions with the current status of the ledger so that transactions are executed only by their endorsers, and the third issue is solved too. Moreover, the double-check performed through the *execute* and *validate* phases prevents the occurrence of non-deterministic or non-terminating executions. So, we can write smart contracts using a general-purpose language like Go, Java, or Javascript.

By exploiting the above design, Fabric can manage a complex semantic for assets and transactions directly at the ledger layer with its native API, without resorting to additional APIs such as those defined for Ethereum's tokens. Accordingly, an asset in Fabric is represented by a generic but finite collection of key-value pairs, with state changes recorded as transactions on the ledger. Specifying suitable collections of key-value pairs makes it possible to consider virtually any kind of asset, from the tangible to the intangible. However, the management of assets does not inherently encompass their monetary value through the system since a cryptocurrency does not back Fabric; in other words, Fabric requires explicit programming to turn its assets into tokens. A research team at IBM has recently designed and implemented the two components Fabric Smart Client and Fabric Token SDK, which should facilitate token exchanges in enterprise contexts [114]. The goal was to design a modular token management system adaptable to various privacy and security regulations. This system adopts the UTXO model pioneered by Bitcoin with some tweaks to support issuance, redemption, transfer, and atomic swap of assets.

4.2. Anchors

Assuring the authenticity of products or their components is fundamental for many industries, from aerospace to electronics, pharmaceuticals, and the food industry. Unnoticed, faked products like food, drugs, diagnostic tests, electronic components, hardware parts, and raw materials pose a high risk of causing significant harm. Therefore, a related demand is to track and trace the logical and physical route, condition, and chain

of custody or ownership of these goods throughout the supply chain and their lifecycle. Blockchain systems have built-in tamper-proof track and trace functionalities. However, a close link between physical objects and their digital representation is essential to avoid the “garbage-in, garbage-out” issue, where a blockchain-managed digital asset is tied to and used to represent an inauthentic asset in the physical world.

Typically, a product is linked to a digital record through an alphanumeric string (e.g., the Global Trade Item Number (GTIN) [115]) uniquely associated with the individual product or a class of products by model, batch, production site, manufacturer, or similar. This unique identifier (UID) is printed, embossed, or attached like a tag to the product or its packaging, and in many cases, it can be easily copied or transferred to a clone of the original product. UIDs by themselves cannot authenticate physical objects, also if they are cryptographically protected (e.g., authenticated), since they are not intrinsically tied to them in unalterable way. What characterizes the authentication of a physical resource through alphanumeric data is indeed the following:

1. there must be a one-to-one unforgeable association between data items and resources, so that each data item corresponds to only one resource and vice versa, without any possible swapping or substitution in such correspondence;
2. the data item associated to the object must be authentic; that is, only the party who creates or holds the resource must be able to generate a valid data item representing it;
3. if a resource is subject to changes in its status (e.g., ownership, physical properties), then the one-to-one unforgeable association considered in (1) must guarantee a one-to-one correspondence with all the states that characterize the resource life cycle, thanks to an appropriate field in the data item.

We will refer to the above scheme as *authentic data encoding for physical resources* (AD-ExPR). Keep in mind that property (1) and its optional extension (3) in turn requires the following two properties: (i) *non-clonability* of the data items, and; (ii) *non-dissociability* between a resource and its related data item. These should clarify the gap between ADExPR and the classic notion of data authentication, along with the cryptographic tools to guarantee this last (i.e., message authentication codes and digital signatures).

An *anchor* ties a UID to the physical object with a property of the object that is hard to clone, forge and transfer to another object. Such a property may be inherent to the object, or it may be entangled (anchored) so that, if removed, it destroys the object, an object functionality, or the property itself. This way, the property enforces the requisites of non-clonability and non-dissociability and, provided that the creators or owners of the objects authentically generated UIDs, anchors result in an ADExPR scheme. If some part of the above encoding has to rely on a cryptographic tool (e.g., a digital signature to guarantee the authenticity of the source of origin for the UIDs), then the more specific term *crypto-anchor* is often used.

In [116], the classification for anchors depicted in Figure 6 represents as a starting point for the design of a blockchain-based *crypto-anchor platform*. Deployed by IBM in collaboration with some leading vendors of crypto-anchor technologies, this platform supports application development ranging from sustainable sourcing to supply-chain management and consumer engagement.

Configured secrets are useful for electronic devices. A manufacturer, distributor, or owner can configure these to store a secret such as a password or a private key. Public-key cryptography allows for a challenge–response protocol through which the device can prove its identity and thwart reply attacks. Both in symmetric and asymmetric schemes, however, the secret might be retrieved by an adversary able to physically interact with the device or retained by the person who configured the device for being reused multiple times. Tamper-proof hardware and cryptographic protocols resilient to side-channel attacks can mitigate the first issue at the price of more expensive devices.

In some cases, storing a secret in a device can be avoided by exploiting physical patterns or behaviors uniquely attributable to that device. Uncontrollable and unpre-

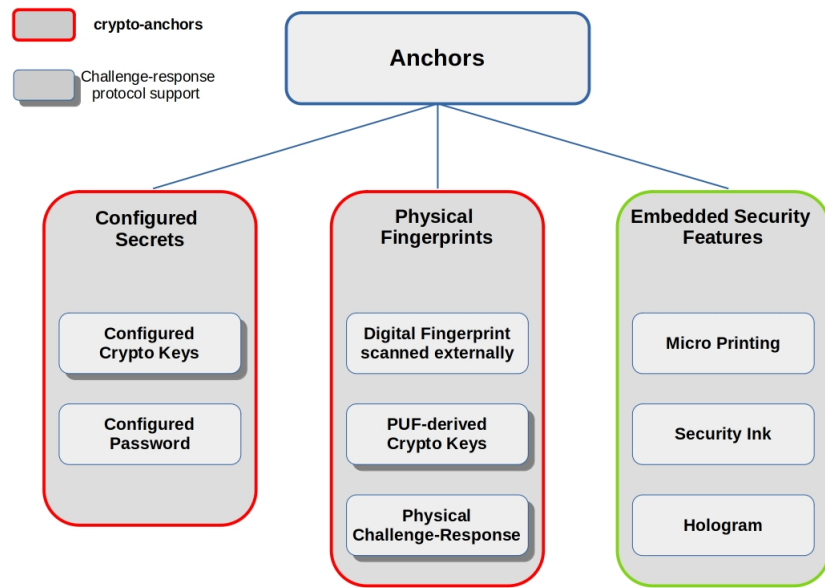


Figure 6. A classification of anchors based on the mechanism used to implement the entanglement of the UID to the physical object (adapted from [116]).

dictable slight variations in the structure of specific materials or the outcome of some production processes can indeed give rise to *physical fingerprints*, like in the production of semiconductors (e.g., transistors) or the texture of a surface. The above variability can be an intrinsic side effect of manufacturing or introduced on purpose and externally, as in doping material with extraneous particles. Since the variations are uncontrollable, they cannot be duplicated, not even by the original manufacturer; moreover, their unpredictable nature can work as a source of randomness for the setup of cryptographic material like secret keys and nonces. However, in order to be eligible as fingerprints of a given set S of objects, the physical variations affecting the objects in S must be characterized in some analytical or statistical sense to be attributable only to the objects in S , in a way that any two different objects in S give rise to two diverse variations. For example, the doping of semiconductors should be attributable to a unique manufacturer, semiconductor type, and semiconductor item. This circumstance means that the physical fingerprints for the objects in S represent the values of a function defined on S and injective therein, called *physical unclonable function* (PUF) [117].

Embedded security features do not rely on cryptographic mechanisms, requiring instead an expensive application process such as micro-printing, hologram generation, or printing with a security ink (e.g., photosensitive ink that is visible to the naked eye but changes color or disappears when placed under a UV light). Anyone can read but not copy a tag produced in this way without special equipment. The difficulty and cost to reproduce an embedded security feature should deter most attackers; however, users often mistake false features for true ones unless they have special equipment for detecting fakes, as in the case of counterfeit banknotes.

4.3. Oracles

Distributed ledger *oracles* are third-party services that provide smart contracts with external information. They intend to prevent another possible cause of the “garbage in–garbage out” problem, i.e., when off-chain data from unverified sources or data providers become input to smart contracts, which could give rise to incorrect processing and related results, with potentially dangerous consequences for both the system and its users. An oracle for a given decentralized application intercepts data originating from one or more sources, possibly processing such data and then delivering them to the smart

contract application along with an *authenticity proof*, a cryptographic guarantee proving the reliability of such data production and that no one tampered it. Data collected by oracles can originate from databases or other online software sources of information and physical devices like electronic sensors, barcode readers, and robots; in some cases, humans can pass data as input to oracles directly.

An essential aspect of oracles is the trust model they implement concerning the multiplicity and heterogeneity of the data they feed a given decentralized system. The oracles needed by smart contracts and their relative trust models can indeed strongly influence the type and level of decentralization implemented at the core system level. Consider, for example, a blockchain system aimed at tracing and guaranteeing the quality of products for a specific production chain. Suppose each party that executes a part of the process independently certifies the quality of their sub-process or by-product through a dedicated and autonomously managed oracle. In that case, the resulting system maintains the degree of decentralization initially provided by the blockchain architecture. Things are pretty different if, on the contrary, a single company acting as certifying body manages all the oracles used by the system. In the latter case, the system preserves decentralization only if the oracles themselves implement this model of trust.

The blockchain ecosystem currently already provides for both centralized and decentralized oracles. For the first ones, a single entity acting as a trusted third-party controls them, while the latter implement some form of DLT to increase the reliability of the information provided to smart contracts by not relying on a single source of truth. In order to preserve decentralization, it is clear that the former kind of oracles require a smart contract that compares multiple sources and selects the data items to be taken into consideration by the majority.

[Provable](#) is an example of an “attestation-as-a-service” company that delivers the data a customer needs along with proof of their authenticity. Provable provides platform-agnostic oracles possibly integrated with centralized or decentralized systems at the customer side. Integration is currently achievable with the smart contracts of some major blockchain platforms like Ethereum, EOS, Hyperledger Fabric, and R3 Corda.

Some of the authenticity proofs provided by Provable are [118]:

- *TLSNotary Proof* relies on a feature of TLS 1.x protocols to enable the splitting of the TLS master key among three parties: the server, an auditee, and an auditor. Provable is the auditee in this scheme, while a locked-down instance of a specially designed, open-source virtual machine on Amazon [Elastic Computing Cloud](#) acts as the auditor.
- *Android Proof* makes use of a software remote attestation technology developed by Google, called SafetyNet [119], to validate that a given Android application is running on a safe, non-rooted physical device connected to Provable’s infrastructure. It also remotely validates the application code hash, enabling authentication of the application running on the device.
- *Ledger Proof* leverages a trusted execution environment comprising an STMicroelectronics secure element, a controller, and BOLOS [120], an operating system developed for hardware wallets by [Ledger](#). BOLOS exposes a set of kernel-level cryptographic APIs, including attestation: any application can ask the kernel to measure its binary and produce a signed hash.

Provable’s authenticity proofs can be either delivered directly to a smart contract or uploaded and stored on some alternate storage medium like the *interplanetary file system* (IPFS) [121], a P2P file system that combines a distributed hashtable, an incentivized block exchange, and a self-certifying namespace so that its nodes do not need to trust each other.

[ChainLink](#), on the other hand, is an example of a company that is revamping in a joint effort with academia the approach of decentralized oracles. In the recently updated whitepaper [122] concerning ChainLink goal and architecture, the authors envision an increasingly expansive role for decentralized oracles thanks to a broader oracle notion and

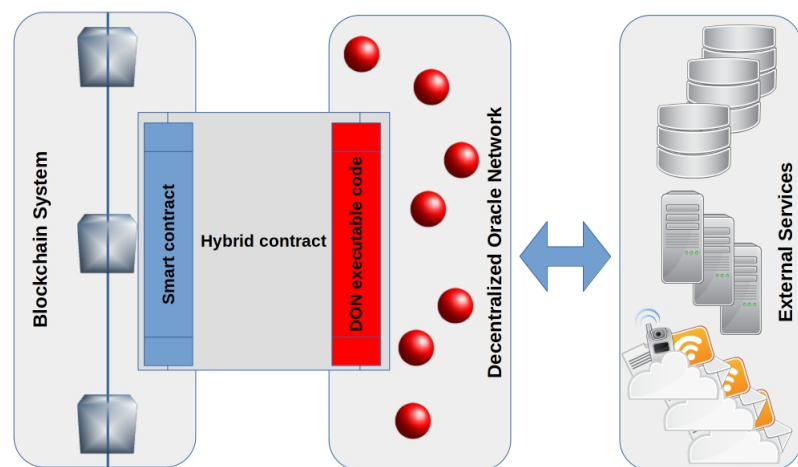


Figure 7. A hybrid smart contract consists of a smart contract component and an off-chain component that executes on a DON. The DON connects the two components and the hybrid contract with off-chain resources such as web services, databases, other blockchains.

a new definition for the oracle network concept previously introduced in [123]. In this up-to-date design, an oracle is a general-purpose, bidirectional, and “compute-enabled” interface between on-chain and off-chain systems, while a *decentralized oracle network* (DON) is a set of oracles provided by third parties and maintained by a committee of Chainlink nodes through a consensus protocol. A DON supports a range of oracle functions chosen for deployment by the committee, acting as a blockchain abstraction layer and providing interfaces to off-chain resources for both smart contracts and other systems. Its goal is to enable a flexible and secure combination of on-chain and off-chain computation connected to external resources, resulting in the *hybrid smart contract* oriented architecture sketched in Figure 7.

Ultimately, the ambition of the above design is to push the degree of abstraction achieved by Chainlink to the point of implementing a layer of *meta contracts* that abstract away the on-chain/off-chain distinction for all classes of developers and users, allowing the seamless creation and use of decentralized services.

5. Toward multiparty process authenticity

Cryptography, and more generally computer security, provides a series of concepts and tools to corroborate evidence of the authenticity of data stored or in transit and the identity of users and devices during their interactions. These are, for example, the concepts of message and entity authentication as well as the related tools of message authentication codes, digital signatures, and challenge-response protocols. However, no existing single tool can encompass the notion of *process authenticity*, intuitively corresponding to the validity of all its steps and their proper binding. Now that many processes are realized thanks to the cooperation of various and independent parties, relying on a single trusted party to guarantee the authenticity of these processes is no more possible. Conversely, DL systems –with their tracking, linking, and tamper-proof features– represent a natural choice to enforce the above notion. However, reaching this goal is a task whose feasibility and difficulty strongly depend on the application scenario and the specific process considered.

The following two sections illustrate the possible strategies and techniques to try to achieve process authentication and the relative limitations in two very different use cases: document dematerialization and e-voting.

5.1. Use case: blockchain and document dematerialization

For about a year, we participated in a project concerning the dematerialization of paper documents in compliance with the updated provisions on the subject issued by [Agenzia per l'Italia Digitale](#) [124]. The project goal is to achieve corroborating evidence of the authenticity of the dematerialization process thanks to an information system that monitors each phase of the process and uses techniques to ascertain the conformity of digital documents with their paper originals. Some of these techniques fall under DLT, and in this section, we will illustrate the benefits and limitations of their use due to careful analysis of design requirements for this use case.

In order to guarantee the authenticity of a document dematerialization process, all the entities (human beings and devices) executing or controlling (a phase of) the process, as well as the resources processed (paper documents and related digital copies), must be suitably identified and traced along the entire period of the process itself. On the other hand, business models and deployment methods for this type of service turn out in a modular and hierarchical framework to enforce and manage trust between the parties involved. The trusted root is associated with a consortium of companies, while subsequent levels correspond to roles and functions of people and devices in a company within the consortium.

Therefore, this use case requires a permissioned blockchain network; moreover, the analysis of the functional characteristics and requirements of the dematerialization process resulted in the blockchain conceptual model shown below.

The notion of *asset* as a set of key-value attributes and a unique identifier (UID) referencing them is well suited to encode each archival item or unit (AU) [125] with its specific features and metadata for cataloging and search purposes. Assuming we implemented the UID through an anchor (see Section 4.2), we accomplished the above by assuring a one-to-one unforgeable association between physical AUs and their digital counterparts. Alternatively, a less severe threat model allows implementing this feature for packaging, as described in the following.

The notion of *transaction* as a change of state for an asset natively offers the possibility of keeping track of the actions performed by the different agents in the different phases of the dematerialization process for a given AU. Thus, we can encode the cause-and-effect correlation among different actions on an AU as a transaction chain in the ledger.

Finally, we can use the notion of *block* to associate multiple AUs with the package used to transport them from the customer site to the processing center. Block creation and verification should correspond to assembling a package at the customer site and opening it at the processing center. This workflow is substantially different from that implemented in blockchain platforms currently on the market, where the transactions submitted by different agents get assembled in blocks only for efficiency reasons. We can use the hash digest obtained as the Merkle tree root of the transactions composing a block as the UID of the package. If the customer monitors AUs packaging at its site, then an anchor implementation for the package UIDs and a tamper-proof packaging could ensure a one-to-one unforgeable association between physical AUs and their digital counterparts, even if the AU UIDs got not anchored.

As for any cyber-physical system, the main challenge in this context is to correctly bind the digital entities managed through the blockchain network to their counterparts in the physical world. Standard authentication methods (e.g., digital signatures and access control) can accomplish the above binding for people, but processing devices and the physical resources subject to processing require advanced mechanisms. The most

compelling case concerns AUs and their packages since they require a technology easy to deploy and at a low cost.

The most straightforward solutions consist of tapes and labels printed on special supports to provide visual evidence of their possible tampering. These unequivocally indicate the possible first opening thanks to unique graphics that can be checked both with the naked eye and through UV rays. The labels contain a QR-code able to identify them uniquely, and the appropriate application of tapes and labels to the AUs or their packages can protect them against tampering or theft.

If replacement or cloning are plausible threats, then more sophisticated tags resulting in anchors can be obtained by generating a special QR-code whose properties are checked by querying a remote server. In this way, besides the anti-counterfeiting mechanism, it is possible to have movement tracking of the AUs, thanks to the geolocation of the reading devices.

An example of an anchor based on *embedded security features* (see Section 4.2) is the *ValiGate* tag produced by tesa scribos, a company of the [tesa group](#). *ValiGate* looks like a normal QR-code but contains a data field protected with a dual encryption algorithm and printed thanks to a proprietary micro-printing technique, requiring a dedicated smartphone scanning app. In this case, encryption can provide an authenticated source of origin, while micro-printing counteracts cloning attacks by preventing a simple “scan and print” approach. A recent change of ownership of *ValiGate* to [SCRIBOS](#), a subsidiary of the german KURZ Group specialized in protecting brands with product markings and digital tools, introduced a new version of the tag removing special printing features, making it compatible with the world’s most widely available printing machines and not requiring dedicated scanning apps.

An alternative in the class of the *configured secrets* is a special kind of signature patented by one of the authors [126], which is only valid if the associated physical asset has a given status and gets univocally tied to a characteristic data set for the asset. In this case, the resilience to cloning is because a valid tag can be generated only by a source of authority (e.g., a product brand), and one or more data uniquely tied to the asset get only disclosed when the asset reaches a given status (e.g., when a product on sale gets sold).

Both of the two previous alternatives can be successfully used not only for the protection of UAs but more generally in different types of supply chains to guarantee the originality of the products they produce and distribute.

5.2. Use case: blockchain and e-voting for public consultations

In many discussions about e-voting, the tacit subject is remote e-voting, which is a relevant and enabling idea when mobility is difficult, or gatherings are discouraged. Several systems exist [7] to express preferences via web or mobile app, and they are effectively adopted to consult opinions in limited assemblies or specific communities. Consequently, authors usually pay little attention to corruption or coercive pressures when narrowing the idea of e-voting in such limited contexts. If this is the case, a proposal of blockchain-based e-voting systems [8] overstates actual risks while introducing an immature layer in the structure.

On the other hand, when thinking about public consultations involving a multitude of people on sensible themes (Referenda) or to assign democratically chosen governing seats, political forces come into play, and we need to consider both specific legal cases and technological issues. In [127], an EU Cooperation Group created a taxonomy to classify the origins and consequences of large-scale cyber-attacks on technologies related to elections. A list of possible treats regarding the solely digital voting procedure includes:

- tampering or DoS of voting and/or vote confidentiality during or after the elections;
- software bug altering election results;
- tampering with logs/journals;
- breach of voter privacy during the casting of votes;

- tampering, DoS, or overload of the systems used for counting or aggregating results;
- tampering or DoS of communication links used to transfer (interim) results;
- tampering with the supply chain involved in the movement or transfer of data.

The above risks directly depend on the inherent properties of a large-scale networked e-voting system, regardless of the assumed technological layers. May the blockchain mitigate those risks? What are the legal points to consider in the e-voting process?

In order to investigate these questions, in 2018, we decided to answer a call of the municipality of Naples (Italy) addressed to associations, universities, research centers, and students to form a volunteer working group on the usage of blockchain technology in municipal referenda. The underlying idea aimed to directly involve the citizens in the democratic decisional process on town-related issues, limiting the costs of a traditional voting consultation.

Since the first day meeting, thanks to suggestions from officers embroiled in past voting sessions, the group focused on possible strategies to limit corruption, vote-buying, and malicious manipulation of ballots. Those generic issues in the first instance convinced us to abandon the idea of remote unattended e-voting. Whatever technology, secured protocol, or cryptographic algorithm we would introduce, we could not receive assurance about private voting: if a voter uses a privately owned unattended device to express a preference, unwanted third parties can record or assist that expression, opening to any corrupting, vote-selling, or menacing activity. Remote e-voting goes against the principle of anonymity requested by public consultations.

When discussing the protection of process authenticity in the context of e-voting, we need to consider many competing factors, and blockchain does not always work toward a reliable outcome. Starting from the involved actors, we considered all the components present in the traditional paper ballot process –resulting from century-long debates on modern vote expression– and their possible equivalent role in a hypothetical blockchain electronic system. The legal framework currently enacted for voting must be respected, and since each nation employs a different system, our reasonings were tailored to the Italian case for local referenda. But the approach herein described can be reshaped to different situations.

As well described in [9], elections frequently have a diverse range of stakeholders and opponents. System designers and legislators conceived of balancing the different forces playing during the voting procedure by identifying each actor and assigning them roles, duties, and powers:

- **voters**, attentive in the correctness of the voting procedure, but also in affirming their interest that could lead to malicious behavior;
- **candidates or referendum promoters**, who want their position chosen by the electors;
- **election officials and poll workers**, who are directly involved in the correctness of the procedure and could misbehave for interest;
- **scrutineers and list representatives**, requiring access to all possible evidence to check for correctness;
- **organizers and suppliers**, who are responsible for the actuation of the process and, therefore, the most exposed figures.

The current voting procedure explicitly or implicitly limits many possibly misconducted activities, and the designers of an e-voting system should keep existing good ideas while correcting, where possible, weaknesses and limiting the introduction of new fragilities.

Principal weaknesses in our local context regard corruption/vote-buying coupled with a tolerated sloppiness in the process execution, leading to errors or ease of misconduct. An unmodifiable ledger like blockchain could positively impact problems like missing ballot papers, tallying errors, inconsistencies, transmission issues while giving more audibility options to scrutineers, list representatives, and the general public.

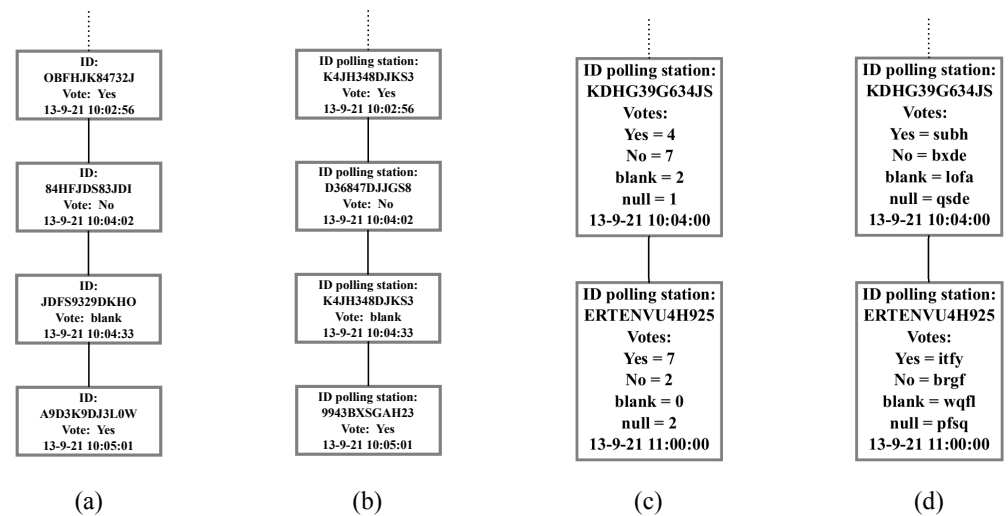


Figure 8. Schematic of possible text for e-voting transactions on a blockchain. (a) Remote voting with anonymization (privacy) by voter ID; vote and timestamp are exposed. (b) Polling station voting with anonymization by collective ID; vote and timestamp are exposed. (c) Polling station voting with anonymization by vote aggregation; votes and timestamp are exposed. (d) Polling station voting with anonymization by vote aggregation and ciphering; timestamp is exposed.

However, better audibility is an enemy of anonymity. One of the strengths of paper ballots is their tangible anonymity: we can count the pieces of paper, which should correspond to the voters, but we cannot understand who put that single piece in the ballot box or at what time. In a digital context, a voter who can use a record in the ledger to track her single vote will be happy to receive proof of counting her intentions, but at the same time, that record will expose her choices (Figure 8/a). In general, registering single votes on a blockchain, even if expressed through a presided voting booth through a shared wallet, can expose a timestamp that malicious parties can track (Figure 8/b). As a general rule, voters should never have a receipt identifying their vote, and a physical and presided ballot box is a reasonable compromise between verifiability and receipt-freeness in most cases.

Can a blockchain help reduce the weaknesses described above without impacting the required anonymity in public consultations? We believe this is the focal question for any case study involving blockchain and e-voting, and no immediate answer is available.

Before proceeding with a tentative, we need to step back to the original questions on possible technological treats on e-voting systems. The highest risks in any system reside in the central controller, and for this reason, any successful attack on that controller would invalidate the outcome of the process at the whole scale. On the contrary, one of the strengths of the traditional paper ballot in Italy is the distributed topology with local responsibilities. Each polling station has officers, workers, and scrutineers that run the local process, from preparing the premises to tallying and packaging of reports and records. If we want to mimic the actors of the traditional paper ballot in an e-voting system with polling stations, we should keep that topology in our networked system: any attack to a station would not interfere with the global system, and an eventual revoting would be local and not complete. A blockchain ledger could fit in this idea.

At this point, we can sketch an initial draft of a blockchain e-voting system with polling stations. Let us consider a certain number of polling stations distributed on the territory basing on voters' population. Each station has a virtual ballot box, not yet well defined, collecting the electronic votes expressed in voting booths through proper voting interfaces (e.g., tablet, laptop, PC). Votes in the ballot box can be aggregated and anonymized before sending the local tallying to a public blockchain address (Figure 8/c). At the end of the polling period, anybody holding access to the adopted blockchain can make the final tallying and declare the winner of the consultation.

Summarizing, during two years-long activity of the working group, we identified some fundamental principles to ensure that the blockchain e-voting process could mitigate traditional risks and improve some desired values for a correct democratic representation:

1. A distributed system can circumscribe possible attacks to local polling stations as long as the chosen blockchain infrastructure is solid and reliable.
2. Voters should go to a polling station and present a valid voting certificate to an officer who will consent to access a polling booth.
3. Votes must be aggregated and anonymized before recording on the blockchain.
4. Each polling station should have a blockchain wallet to make transactions with a central wallet and transmit aggregated votes. Those transactions should be periodic during the polling hours but with a minimum consistency in the number of votes to prevent time-related traceability.
5. Italian regulations prohibit releasing exit polls during voting hours, and therefore voting information within transactions should be encrypted (Figure 8/d), and decryption keys publicly released at the end of the voting period.
6. The distributed ledger adopted should be public to allow free scrutinizing.
7. Anybody provided with suitable software, free access to the adopted blockchain, and the decryption keys must have the opportunity to do the final overall tallying.
8. All software must be open source to let the public check functionalities.

An additional opportunity, not specific to blockchain but electronic voting in general, is the ability to implement revoting functionalities: if a voter can change her vote repeatedly, any visual record of her choice (e.g., a picture on a mobile phone) does not represent a final decision, and a malevolent third party will not have a proof for coercive finalities.

About point 3 above, much discussion is still open. How can a voter trust the implemented system for the local virtual ballot box? Since receipt-freeness [6] is compulsory to prevent coercion and vote-buying, are end-to-end voting systems feasible to implement this task? Are zero-knowledge-proof or shuffle techniques feasible in this context? Using complex cryptographic protocols could deter skeptical voters from e-voting systems, obtaining a contrary effect from the desired increased democratic participation.

6. Conclusions

Following the critical approach used in [32] during the booming of DLT, in this work, we presented the most relevant and impacting concepts within the context of technologies sprouted from the affirmation of Bitcoin. Such concepts, sometimes already present in past literature, now receive an injection of concrete applicability. Opportunities to make a profit or affirm new governance approaches have nurtured a plethora of initiatives that we tried to organize by concepts and impact in implementing the actual applications.

This paper represents our tentative to exhibit sometimes underexposed but still impacting new or consolidated aspects of DLT. It was born from our need to have a clearer view about tangled and often just sketched proposals available in literature or on the market during our involvement in applicative projects. We do not propose a taxonomy or a classification but a digest of valuable concepts for beginners and experts looking for unexpected links among ideas in the DLT scene. In particular, in Section 3.3 about consensus, we followed a historical path marking the progression of ideas emerging with the evolution of DL systems. We also introduced the ADExPR scheme to characterize the authentication of physical resources through alphanumeric data, which can facilitate the reader to discriminate from the qualities of digital data authentication. Finally, we propose the concept of process authenticity, which we believe will be more and more prominent in the following years, and we took advantage of two actual use cases to show how blockchain and DLT, in general, can help in its validation.

7. Patents

This section is not mandatory, but may be added if there are patents resulting from the work reported in this manuscript.

Funding: Please add: “This research received no external funding” or “This research was funded by NAME OF FUNDER grant number XXX.” and “The APC was funded by XXX”. Check carefully that the details given are accurate and use the standard spelling of funding agency names at <https://search.crossref.org/funding>, any errors may affect your future funding.

Institutional Review Board Statement: In this section, please add the Institutional Review Board Statement and approval number for studies involving humans or animals. Please note that the Editorial Office might ask you for further information. Please add “The study was conducted according to the guidelines of the Declaration of Helsinki, and approved by the Institutional Review Board (or Ethics Committee) of NAME OF INSTITUTE (protocol code XXX and date of approval).” OR “Ethical review and approval were waived for this study, due to REASON (please provide a detailed justification).” OR “Not applicable” for studies not involving humans or animals. You might also choose to exclude this statement if the study did not involve humans or animals.

Informed Consent Statement: Any research article describing a study involving humans should contain this statement. Please add “Informed consent was obtained from all subjects involved in the study.” OR “Patient consent was waived due to REASON (please provide a detailed justification).” OR “Not applicable” for studies not involving humans. You might also choose to exclude this statement if the study did not involve humans.

Written informed consent for publication must be obtained from participating patients who can be identified (including by the patients themselves). Please state “Written informed consent has been obtained from the patient(s) to publish this paper” if applicable.

Data Availability Statement: In this section, please provide details regarding where data supporting reported results can be found, including links to publicly archived datasets analyzed or generated during the study. Please refer to suggested Data Availability Statements in section “MDPI Research Data Policies” at <https://www.mdpi.com/ethics>. You might choose to exclude this statement if the study did not report any data.

Acknowledgments: In this section you can acknowledge any support given which is not covered by the author contribution or funding sections. This may include administrative and technical support, or donations in kind (e.g., materials used for experiments).

Conflicts of Interest: Declare conflicts of interest or state “The authors declare no conflict of interest.” Authors must identify and declare any personal circumstances or interest that may be perceived as inappropriately influencing the representation or interpretation of reported research results. Any role of the funders in the design of the study; in the collection, analyses or interpretation of data; in the writing of the manuscript, or in the decision to publish the results must be declared in this section. If there is no role, please state “The funders had no role in the design of the study; in the collection, analyses, or interpretation of data; in the writing of the manuscript, or in the decision to publish the results”.

Sample Availability: Samples of the compounds ... are available from the authors.

1. Zhang, R.; Xue, R.; Liu, L. Security and privacy on blockchain. *ACM Computing Surveys (CSUR)* **2019**, *52*, 1–34.
2. Belotti, M.; Božić, N.; Pujolle, G.; Secci, S. A vademecum on blockchain technologies: When, which, and how. *IEEE Communications Surveys & Tutorials* **2019**, *21*, 3796–3838.
3. Xiao, Y.; Zhang, N.; Lou, W.; Hou, Y.T. A survey of distributed consensus protocols for blockchain networks. *IEEE Communications Surveys & Tutorials* **2020**, *22*, 1432–1465.
4. Antal, C.; Cioara, T.; Anghel, I.; Antal, M.; Salomie, I. Distributed Ledger Technology Review and Decentralized Applications Development Guidelines. *Future Internet* **2021**, *13*, 62. doi:10.3390/fi13030062.
5. Wang, K.H.K.; Mondal, S.K.; Chan, K.C.; Xie, X. A Review of Contemporary E-voting: Requirements, Technology, Systems and Usability. *Data Science and Pattern Recognition* **2017**, *1*, 31.

6. Baudron, O.; Fouque, P.A.; Pointcheval, D.; Stern, J.; Poupard, G. Practical Multi-Candidate Election System. *Proceedings of the Twentieth Annual ACM Symposium on Principles of Distributed Computing*; Association for Computing Machinery: New York, NY, USA, 2001; PODC '01, p. 274–283. doi:10.1145/383962.384044.
7. Gibson, J.P.; Krimmer, R.; Teague, V.; Pomares, J. A review of E-voting: the past, present and future. *Annals of Telecommunications* **2016**, *71*, 279–286. doi:10.1007/s12243-016-0525-8.
8. Abuidris, Y.; Kumar, R.; Wenyong, W. A Survey of Blockchain Based on E-voting Systems. *Proceedings of the 2019 2nd International Conference on Blockchain Technology and Applications*. Association for Computing Machinery, 2019, ICBTA 2019, p. 99–104. doi:10.1145/3376044.3376060.
9. Park, S.; Specter, M.; Narula, N.; Rivest, R.L. Going from bad to worse: from Internet voting to blockchain voting. *Journal of Cybersecurity* **2021**, *7*. doi:10.1093/cybsec/tyaa025.
10. Fowler, M.; Lewis, J. *Microservices*, 2014. Accessed: 2021-06-14.
11. Newman, S. *Building microservices: designing fine-grained systems*; "O'Reilly Media, Inc.", 2015.
12. Guo, D.; Wang, W.; Zeng, G.; Wei, Z. Microservices architecture based cloudware deployment platform for service computing. *2016 IEEE Symposium on Service-Oriented System Engineering (SOSE)*. IEEE, 2016, pp. 358–363.
13. Zuboff, S. Big other: surveillance capitalism and the prospects of an information civilization. *Journal of Information Technology* **2015**, *30*, 75–89.
14. Gilder, G. *Life after Google: The fall of big data and the rise of the blockchain economy*; Simon and Schuster, 2018.
15. Pathirana, N. *How to set up a Private Ethereum Blockchain*, 2019. Accessed: 2021-07-23.
16. Douceur, J.R. The sybil attack. *International workshop on peer-to-peer systems*. Springer, 2002, pp. 251–260.
17. Osipkov, I.; Vasserman, E.Y.; Hopper, N.; Kim, Y. Combating double-spending using cooperative P2P systems. *27th International Conference on Distributed Computing Systems (ICDCS'07)*. IEEE, 2007, pp. 41–41.
18. Bevand, M. *Cambridge Bitcoin Electricity Consumption Index*, 2017. Accessed: 2021-07-27.
19. Bevand, M. *Energy consumption of a Bitcoin (BTC, BTH) and VISA transaction*, 2021. Accessed: 2021-07-27.
20. Cox, J. *Yellen sounds warning about "extremely inefficient" bitcoin*, 2021. Accessed: 2021-07-27.
21. Kolodny, L. *Elon Musk says Tesla will stop accepting bitcoin for car purchases, citing environmental concerns*, 2021. Accessed: 2021-07-27.
22. Xu, B.; Luthra, D.; Cole, Z.; Blakely, N. *EOS: An Architectural, Performance, and Economic Analysis* **2018**. p. 25.
23. Goodman, L.M. *Tezos – a self-amending crypto-ledger White paper* **2014**. p. 17.
24. *Internet Computer Governance*. Accessed: 2021-08-1.
25. *DFINITY Technical Library*. Accessed: 2021-08-1.
26. Kahng, A.; Mackenzie, S.; Procaccia, A. Liquid Democracy: An Algorithmic Perspective. *Journal of Artificial Intelligence Research* **2021**, *70*, 1223–1252. doi:10.1613/jair.1.12261.
27. *Decred Documentation*. Accessed: 2021-08-1.
28. Pelt, R.v.; Jansen, S.; Baars, D.; Overbeek, S. Defining Blockchain Governance: A Framework for Analysis and Comparison. *Information Systems Management* **2021**, *38*, 21–41. doi:10.1080/10580530.2020.1720046.
29. Johnson, D.; Menezes, A.; Vanstone, S. The elliptic curve digital signature algorithm (ECDSA). *International Journal of Information Security* **2001**, *1*, 36–63.
30. Cooper, D.; Santesson, S.; Farrell, S.; Boeyen, S.; Housley, R.; Polk, W. Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile. <http://tools.ietf.org/html/rfc5280>, 2008.
31. *Membership Service Providers*, 2020. Accessed: 2021-08-14.
32. Romano, D.; Schmid, G. Beyond bitcoin: A critical look at blockchain-based systems. *Cryptography* **2017**, *1*, 15.
33. *Contract ABI Specification*, 2020. Accessed: 2021-08-14.
34. *Fabric chaincode lifecycle*, 2020. Accessed: 2021-08-14.
35. Nakamoto, S. *Bitcoin: A Peer to Peer Electronic Cash System*, 2008.
36. Rogaway, P.; Shrimpton, T. Cryptographic hash-function basics: Definitions, implications, and separations for preimage resistance, second-preimage resistance, and collision resistance. *Fast Software Encryption*. Springer, 2004, pp. 371–388.
37. Haber, S.; Stornetta, W.S. How to time-stamp a digital document. *Conference on the Theory and Application of Cryptography*. Springer, 1990, pp. 437–455.
38. Merkle, R.C. A digital signature based on a conventional encryption function. *Advances in Cryptology—CRYPTO'87*. Springer, 1987, pp. 369–378.
39. Dwork, C.; Naor, M. Pricing via processing or combatting junk mail. *Advances in Cryptology - CRYPTO'92*. Springer, 1993, pp. 139–147.
40. Howell, A. *The Longest Blockchain is not the Strongest Blockchain*, 2019. Accessed: 2021-07-27.
41. Eyal, I.; Sirer, E.G. Majority is not enough: Bitcoin mining is vulnerable. In *Financial Cryptography and Data Security*; Springer, 2014; pp. 436–454.
42. Buterin, V. *Ethereum White Paper*, 2013. Accessed: 2021-07-03.
43. Gervais, A.; Karame, G.O.; Wüst, K.; Glykantzis, V.; Ritzdorf, H.; Capkun, S. On the security and performance of proof of work blockchains. *Proceedings of the 2016 ACM SIGSAC conference on computer and communications security*, 2016, pp. 3–16.
44. Szilagy, P.; Wilcke, J.; Lange, F.; Zsolt, F. [go-ethereum/core/blockchain.go:difficulty/calculation](https://go-ethereum.org/core/blockchain.go:difficulty/calculation), 2017. Accessed: 2021-07-27.

45. Morrison, D.R. PATRICIA—practical algorithm to retrieve information coded in alphanumeric. *Journal of the ACM (JACM)* **1968**, 15, 514–534.
46. Thulasiraman, K.; Swamy, M.N. *Graphs: theory and algorithms*; John Wiley & Sons, 2011.
47. Lerner, S.D. [DagCoin draft](#), 2015. Accessed: 2021-07-27.
48. Ribero, Y. [DagCoin whitepaper](#), 2018. Accessed: 2021-07-27.
49. Sompolinsky, Y.; Lewenberg, Y.; Zohar, A. SPECTRE: a fast and scalable cryptocurrency protocol. *IACR Cryptol. ePrint Arch.* **2016**, 2016.
50. Li, C.; Li, P.; Zhou, D.; Xu, W.; Long, F.; Yao, A. Scaling nakamoto consensus to thousands of transactions per second. *arXiv preprint arXiv:1805.03870* **2018**.
51. Popov, S. [The tangle](#), 2018. Accessed: 2021-07-27.
52. Popov, S.; Moog, H.; Camargo, D.; Caposelle, A.; Dimitrov, V.; Gal, A.; Greve, A.; Kusmierz, B.; Mueller, S.; Penzkofer, A.; others. [The coordicide](#), 2020. Accessed: 2021-07-27.
53. Hearn, M.; Brown, R.G. [Corda: A distributed ledger](#), 2019. Accessed: 2021-07-27.
54. Junqueira, F.P.; Reed, B.C.; Serafini, M. Zab: High-performance broadcast for primary-backup systems. *IEEE/IFIP 41st International Conference on Dependable Systems & Networks (DSN)*. IEEE, 2011, pp. 245–256.
55. Ongaro, D.; Ousterhout, J. In search of an understandable consensus algorithm. *USENIX Annual Technical Conference ATC 14*, 2014, pp. 305–319.
56. Lamport, L. Proving the correctness of multiprocess programs. *IEEE transactions on software engineering* **1977**, pp. 125–143.
57. Lamport, L.; Shostak, R.; Pease, M. The Byzantine Generals Problem. *ACM Transactions on Programming Languages and Systems* **1982**, 4, 382–401.
58. Pease, M.; Shostak, R.; Lamport, L. Reaching agreement in the presence of faults. *Journal of the ACM (JACM)* **1980**, 27, 228–234.
59. Castro, M.; Liskov, B.; others. Practical byzantine fault tolerance. *OSDI*, 1999, Vol. 99, pp. 173–186.
60. Schneider, F.B. Implementing fault-tolerant services using the state machine approach: A tutorial. *ACM Computing Surveys (CSUR)* **1990**, 22, 299–319.
61. Miller, A.; Xia, Y.; Croman, K.; Shi, E.; Song, D. The honey badger of BFT protocols. *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, 2016, pp. 31–42.
62. Duan, S.; Reiter, M.K.; Zhang, H. BEAT: Asynchronous BFT made practical. *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security*, 2018, pp. 2028–2041.
63. Cachin, C.; Vukolić, M. Blockchain consensus protocols in the wild. *arXiv preprint arXiv:1707.01873* **2017**.
64. Bano, S.; Sonnino, A.; Al-Bassam, M.; Azouvi, S.; McCorry, P.; Meiklejohn, S.; Danezis, G. Consensus in the age of blockchains. *arXiv preprint arXiv:1711.03936* **2017**.
65. Wang, W.; Hoang, D.T.; Hu, P.; Xiong, Z.; Niyato, D.; Wang, P.; Wen, Y.; Kim, D.I. A survey on consensus mechanisms and mining strategy management in blockchain networks. *IEEE Access* **2019**, 7, 22328–22370.
66. Wood, G. [Ethereum: a Secure Decentralized Generalized Transaction Ledger](#), 2016.
67. King, S.; Nadal, S. [Ppcoin: Peer-to-peer crypto-currency with proof-of-stake](#), 2012.
68. Community, N. [Nxt White Paper](#), 2014. Accessed: 2021-07-03.
69. Miller, A.; Juels, A.; Shi, E.; Parno, B.; Katz, J. Permacoin: Repurposing bitcoin work for data preservation. *2014 IEEE Symposium on Security and Privacy*. IEEE, 2014, pp. 475–490.
70. Chen, L.; Xu, L.; Shah, N.; Gao, Z.; Lu, Y.; Shi, W. On security analysis of proof-of-elapsed-time (poet). *International Symposium on Stabilization, Safety, and Security of Distributed Systems*. Springer, 2017, pp. 282–297.
71. Costan, V.; Devadas, S. Intel sgx explained. *IACR Cryptol. ePrint Arch.* **2016**, 2016, 1–118.
72. [Ethash Design Rationale](#), 2020. Accessed: 2021-07-03.
73. [Proof of Stake FAQs](#), 2020. Accessed: 2021-07-03.
74. Bentov, I.; Lee, C.; Mizrahi, A.; Rosenfeld, M. Proof of Activity: Extending Bitcoin’s Proof of Work via Proof of Stake. *SIGMETRICS 2014 Workshop on Economics of Networked Systems, NetEcon*. ACM, 2014.
75. Li, W.; Andreina, S.; Bohli, J.M.; Karamé, G. Securing proof-of-stake blockchain protocols. In *Data Privacy Management, Cryptocurrencies and Blockchain Technology*; Springer, 2017; pp. 297–315.
76. Kiayias, A.; Russell, A.; David, B.; Oliynykov, R. Ouroboros: A provably secure proof-of-stake blockchain protocol. *Annual International Cryptology Conference*. Springer, 2017, pp. 357–388.
77. Gilad, Y.; Hemo, R.; Micali, S.; Vlachos, G.; Zeldovich, N. Algorand: Scaling byzantine agreements for cryptocurrencies. *Proceedings of the 26th symposium on operating systems principles*, 2017, pp. 51–68.
78. David, B.; Gaži, P.; Kiayias, A.; Russell, A. Ouroboros praos: An adaptively-secure, semi-synchronous proof-of-stake blockchain. *Annual International Conference on the Theory and Applications of Cryptographic Techniques*. Springer, 2018, pp. 66–98.
79. Daian, P.; Pass, R.; Shi, E. Snow white: Robustly reconfigurable consensus and applications to provably secure proof of stake. *International Conference on Financial Cryptography and Data Security*. Springer, 2019, pp. 23–41.
80. Buchman, E. Tendermint: Byzantine fault tolerance in the age of blockchains. PhD thesis, 2016.
81. Wood, G. [PoA Private Chains](#), 2015. Accessed: 2021-07-27.
82. Buchman, E.; Kwon, J.; Milosevic, Z. The latest gossip on BFT consensus, 2019, [[arXiv:cs.DC/1807.04938](#)].

83. Bentov, I.; Gabizon, A.; Mizrahi, A. Cryptocurrencies without proof of work. *International conference on financial cryptography and data security*. Springer, 2016, pp. 142–157.
84. Stadler, M. Publicly verifiable secret sharing. *International Conference on the Theory and Applications of Cryptographic Techniques*. Springer, 1996, pp. 190–199.
85. Cramer, R.; Damgård, I. Multiparty computation, an introduction. In *Contemporary cryptology*; Springer, 2005; pp. 41–87.
86. Pass, R.; Shi, E. The sleepy model of consensus. *International Conference on the Theory and Application of Cryptology and Information Security*. Springer, 2017, pp. 380–409.
87. Micali, S.; Rabin, M.; Vadhan, S. Verifiable random functions. *40th annual symposium on foundations of computer science (cat. No. 99CB37039)*. IEEE, 1999, pp. 120–130.
88. [Bitshares documentation: Delegated Proof of Stake \(DPOS\)](#). Accessed: 2021-09-11.
89. Lavin, J.; Larimer, D.; Hourt, N.; Ma, Q.; Prioriello, W. [EOS.IO Technical White Paper v2](#), 2018. Accessed: 2021-07-27.
90. Buterin, V.; Griffith, V. Casper the friendly finality gadget. *arXiv preprint arXiv:1710.09437* **2017**.
91. Schollmeier, R. A definition of peer-to-peer networking for the classification of peer-to-peer architectures and applications. *Proceedings First International Conference on Peer-to-Peer Computing*. IEEE, 2001, pp. 101–102.
92. Fanti, G.; Viswanath, P. Deanonymization in the bitcoin P2P network. *Proceedings of the 31st International Conference on Neural Information Processing Systems*, 2017, pp. 1364–1373.
93. Biryukov, A.; Tikhomirov, S. Deanonymization and linkability of cryptocurrency transactions based on network analysis. *2019 IEEE European Symposium on Security and Privacy (EuroS&P)*. IEEE, 2019, pp. 172–184.
94. Ren, L. Analysis of Nakamoto Consensus. *IACR Cryptol. ePrint Arch.* **2019**, 2019, 943.
95. Fischer, M.J.; Lynch, N.A.; Paterson, M.S. Impossibility of distributed consensus with one faulty process. *Journal of the ACM (JACM)* **1985**, 32, 374–382.
96. Liu, S.; Viotti, P.; Cachin, C.; Quéma, V.; Vukolić, M. {XFT}: Practical fault tolerance beyond crashes. *12th {USENIX} Symposium on Operating Systems Design and Implementation ({OSDI} 16)*, 2016, pp. 485–500.
97. Dolev, D.; Strong, H.R. Authenticated algorithms for Byzantine agreement. *SIAM Journal on Computing* **1983**, 12, 656–666.
98. Abraham, I.; Malkhi, D.; Nayak, K.; Ren, L.; Yin, M. Sync hotstuff: Simple and practical synchronous state machine replication. *2020 IEEE Symposium on Security and Privacy (SP)*. IEEE, 2020, pp. 106–118.
99. Dwork, C.; Lynch, N.; Stockmeyer, L. Consensus in the presence of partial synchrony. *Journal of the ACM (JACM)* **1988**, 35, 288–323.
100. Bitcoin Wiki: [Script](#). Accessed: 2021-07-03.
101. [Solidity Documentation](#). Accessed: 2021-07-03.
102. [Ethereum Virtual Machine \(EVM\)](#), 2020. Accessed: 2021-08-14.
103. Turing, A.M. On computable numbers, with an application to the Entscheidungsproblem. *Proceedings of the London mathematical society* **1937**, 2, 230–265.
104. Rosenfeld, M. [Overview of colored coins](#), 2012. Accessed: 2021-08-14.
105. Bitcoin Wiki: [Colored Coins](#). Accessed: 2021-07-03.
106. Vogelsteller, F.; Buterin, V. [EIP-20: Token Standard](#), 2015. Accessed: 2021-08-14.
107. Entriiken, W.; Shirley, D.; Evans, J.; Sachs, N. [EIP-721: Non-Fungible Token Standard](#), 2018. Accessed: 2021-08-14.
108. Reitwießner, C.; Johnson, N.; Vogelsteller, F.; Baylina, J.; Feldmeier, K.; Entriiken, W. [EIP-165: Standard Interface Detection](#), 2018. Accessed: 2021-08-14.
109. Dafflon, J.; Baylina, J.; Shababi, T. [EIP-777: Token Standard](#), 2017. Accessed: 2021-08-14.
110. Radomski, W.; Cooke, A.; Castonguay, P.; Therien, J.; Binet, E.; Sandford, R. [EIP-1155: Multi-Token Standard](#), 2018. Accessed: 2021-08-14.
111. Shiple, J.; Marks, H.; Zhang, D. [EIP-1450: A compatible security token for issuing and trading SEC-compliant securities](#), 2018. Accessed: 2021-08-14.
112. Kupriianov, M.; Svirsky, J. [EIP-1462: Base security token](#), 2018. Accessed: 2021-08-14.
113. Androulaki, E.; Barger, A.; Bortnikov, V.; Cachin, C.; Christidis, K.; De Caro, A.; Enyeart, D.; Ferris, C.; Laventman, G.; Manevich, Y.; others. Hyperledger fabric: a distributed operating system for permissioned blockchains. *Proceedings of the thirteenth EuroSys conference*, 2018, pp. 1–15.
114. Androulaki, E.; De Caro, A.; El Khyaoui, K. [Making fungible tokens and NFTs safer to use for enterprises](#), 2021. Accessed: 2021-07-27.
115. [Global Trade Item Number \(GTIN\)](#), 2020. Accessed: 2021-08-14.
116. Balagurusamy, V.S.K.; Cabral, C.; Coomaraswamy, S.; Delamarche, E.; Dillenberger, D.N.; Dittmann, G.; Friedman, D.; Gökçe, O.; Hinds, N.; Jelitto, J.; Kind, A.; Kumar, A.D.; Libsch, F.; Ligman, J.W.; Munetoh, S.; Narayanaswami, C.; Narendra, A.; Paidimarri, A.; Delgado, M.A.P.; Rayfield, J.; Subramanian, C.; Vaculin, R. Crypto anchors. *IBM Journal of Research and Development* **2019**, 63, 4:1–4:12. doi:10.1147/JRD.2019.2900651.
117. Prada-Delgado, M.Á.; Baturone, I.; Dittmann, G.; Jelitto, J.; Kind, A. PUF-derived IoT identities in a zero-knowledge protocol for blockchain. *Internet of Things* **2020**, 9, 100057.
118. [Authenticity Proofs](#), 2019. Accessed: 2021-07-27.
119. [SafetyNet Attestation API](#), 2021. Accessed: 2021-07-27.
120. [Introducing BOLOS: Blockchain Open Ledger Operating System](#), 2016. Accessed: 2021-07-27.

-
121. Benet, J. IPFS—content addressed, versioned, p2p file system. *arXiv preprint arXiv:1407.3561* **2014**.
 122. Breidenbach, L.; Cachin, C.; Chan, B.; Coventry, A.; Ellis, S.; Juels, A.; Koushanfar, F.; Miller, A.; Magauran, B.; Moroz, D.; Nazarov, S.; Topliceanu, A.; Tramer, F.; Zhang, F. [Chainlink 2.0: Next Steps in the Evolution of Decentralized Oracle Networks](#) **2021**. p. 136. Accessed: 2021-08-1.
 123. Ellis, S.; Juels, A.; Nazarov, S. [ChainLink – A Decentralized Oracle Network](#) **2017**. p. 38. Accessed: 2021-08-1.
 124. [Codice dell'amministrazione digitale](#), 2020. Accessed: 2021-08-14.
 125. [Levels of Archival Description](#), 2016. Accessed: 2021-09-23.
 126. Schmid, G. [Un procedimento di anti-contraffazione su base collaborativa](#), 2016. Accessed: 2021-09-23.
 127. [Compendium on Cyber Security of Election Technology](#) **2018**. Accessed: 2021-09-18.