Blockchain for Enterprise: Overview, Opportunities and Challenges

Elyes Ben Hamida, Kei Leo Brousmiche, Hugo Levard, Eric Thea

Blockchain for Enterprise: Overview, Opportunities and Challenges

Elyes Ben Hamida*, Kei Leo Brousmiche *, Hugo Levard *† and Eric Thea *
*Institute for Technological Research SystemX, France
Email: elyes.ben-hamida,kei-leo.brousmiche,eric.thea@irt-systemx.fr
†SQLI, France
Email: hlevard@sqli.com

***Abstract*—The Blockchain technology has received increased interests in recent years, from both the scientific community and the industry. This technology represents a major paradigm shift in the way smart cities solutions will be built, operated, consumed and marketed in the near future. Even though Blockchains will have a tremendous potential impact on businesses and societies, there are many open challenges that need to be carefully tackled. This article focuses on enterprise Blockchains and provides a detailed analysis on its core components, technologies and applications. Finally, various research challenges and opportunities are discussed.**

Keywords–Blockchain; Distributed Ledger Technology; Smart Contract; Consensus Algorithm; Data Privacy and Security; Scalability.

I. Introduction

Blockchains have recently attracted increased interests within the governments, businesses and research community, with applications in key industries, such as finance, insurance, logistics, energy and transportation. Indeed, the blockchain technology is foreseen as the core backbone of future smart cities and Internet of Things by enhancing its security, data management and process automation.

A blockchain [1] is essentially a trustless, peer-to-peer and continuously growing database (or ledger) of records, including distributed applications (or smart contracts), that have been executed and shared among the participating entities. It enables applications and systems to operate in a fully decentralized fashion without the need for any third party or trust authority.

This technology per se is not novel, but is rather a combination of well-known building blocks, including peer-to-peer protocols, cryptographic primitives, distributed consensus algorithms and economic incentives mechanisms. A blockchain is more a paradigm shift in the way applications and solutions will be built, deployed, operated, consumed and marketed in the near future, than just a technology. Blockchain is secure by design and relies on well-known cryptographic tools and distributed consensus mechanisms to provide key characteristics, such as persistence, anonymity, fault-tolerance, auditability and resilience. Indeed, each record in the chain is verified by consensus of a majority of the blockchain's participants, and once committed on the ledger, cannot be easily tampered with.

More recently, smart contracts [2] have emerged as a new usage for blockchains to digitize and automate the execution of business workflows (*i.e.,* self-executing contracts or agreements), and whose proper execution is enforced by the consensus mechanism. This makes the blockchain technology particularly suitable for the management of medical records [3], notary services [4], users' identities [5] and reputations [6], data traceability [7], *etc.*

However, several challenges will need to be addressed to unlock the tremendous potential of blockchains, especially before this paradigm shift becomes technically, economically and legally viable in business environments. The first category of these challenges concerns the technical aspects of blockchains, including in terms of governance (*i.e.,* open, private or consortium), scalability, data privacy, and validity of smart contracts. The second set of challenges is related to the development of viable underlying business models and incentives mechanisms. Last but not least, the legal aspects of blockchains represent a challenge, especially in France and Europe, where this technology should be analyzed in the light of upcoming new regulations, such as the General Data Protection Regulation (GDPR) (Regulation (EU) 2016/679 [8]), and whose objective is to strengthen users' data privacy and protection within the European Union.

This article focuses on the technical aspects of blockchains and their potential benefits to enterprises and industrial use cases. The remainder of this article is organized as follows. Section II discusses the technical aspects of blockchains in terms of taxonomy, system architecture, consensus algorithms and technologies. Section III provides a classification of blockchains applications and highlights typical use cases in the finance, energy, mobility and logistics sectors. Section IV draws and discusses various research challenges and opportunities. Finally, Section V concludes the article.

II. Public and Private Blockchain Technology

While public blockchains enable parties to make transactions in a secured manner in trust-less environments, they show certain limitations when applied to industrial use cases. Indeed, we believe that aspects, such as controlled data reversibility (1), data privacy (2), transactions volume scalability (3), system responsiveness (4) and ease of protocol updatability (5) that are crucial for the majority of corporate applications are not covered by public blockchain implementations. These shortcomings led industrials to develop alternative blockchain technologies tackling the aforementioned aspects and intended for restricted audience. These technologies can generally be classified into two categories: private and consortium blockchains [9]. The distinction between them comes down to the governance scheme. In private chains, one participant rules the whole system whereas members of consortium blockchains share the authority among them. Nowadays, new terms and concepts are flourishing for categorizing approaches between public and private blockchains such as semi-private or enterprise technologies. However, their differences concerns the application level and not architectural aspects. For the

sake of simplicity and clarity, we will assume that the term of private blockchain encompasses these different non-public concepts in the following. In next sections, we overview the architecture of both public and private blockchain and highlight their inherent differences.

A. Blockchain Architecture

a) Data structure: The data structure of a blockchain, whether public or private, corresponds to a linked list of blocks containing transactions. Each element of the list, has a pointer to the previous block. Moreover, each pointer of a block contains the hash of the previous block. This hash is the key element of the blockchain security. Indeed, if an adversary tries to modify the content of a block, anyone can detect it by computing its hash and comparing it to the hash stored in the next block to see the inconsistence. In order to avoid this detection, the adversary could try to change all the hashes from the tampered block to the latest block. However, this is not feasible without the consenting of more than the half of the participants (see Section II-A0c). Therefore, modifying the content of a block is impossible over public chains. On the other hand, private chains members can easily come to an agreement off-line and modify data content (1). Private chains can be seen as append-only databases where the main goal consists in sharing and syncing data within a consortium.

b) Network and privacy: Along with its data structure, a blockchain is based on a peer-to-peer network that ties its participants. Depending on the implementation of the blockchain, the network can be public (*i.e.,* anyone can access it) or private (*i.e.,* only accounts that are allowed can participate). This restricted access to the network assures data privacy (2). Moreover, some private blockchains allows to control data visibility at a more finer grain by enabling data encryption at transaction level (*e.g.*, [10]). Nodes can read data and ask the network to add new data, these pending data are then picked by some special nodes called the miners (also known as block generators or validators).

c) Security and scalability: Miners, or validators, are nodes that are willing to share their computational power to add blocks to the blockchain. The process for selecting the actual node that will add the next block among all the validators is referred as a consensus protocol. In a trust-less public configuration, this consensus is crucial for the integrity and security of the data. Thus, to prove their commitment and prevent malicious activity, miners usually have solve a computationally demanding cryptographic puzzle (*i.e.,* Proof of Work [11]).

On the other hand, in private chains, since miners or rather validators are preliminary known and trusted to some degree, this process of selection can be lowered in terms of computational power. This reduction of complexity in the consensus protocol leads directly to an increased scalability in terms of transactions throughput (3). An overview of the major consensus algorithm is proposed in section II-B.

d) Forks and responsiveness: Once a miner's block has been selected, it is added to the blockchain and the information is broadcast. Due to network effects, there are cases where multiple miners blocks are selected, so there are different versions of the blockchain in different regions of the network. This is called a fork: the blockchain splits into branches. In this case, nodes should somehow converge towards acknowledging a unique and same version of the blockchain. In practice, the Proof of Work consensus achieves this result by requiring miners to work on the longest branch that they see. However, this means that even if a transaction has been validated, we cannot be sure that it will remain on the main chain. In Bitcoin, users usually wait 6 blocks of confirmation before considering a transaction as valid. Thus, there is a correlation between the probability of fork occurrence and the *responsiveness* of the blockchain. On private chains, the use of adapted consensus algorithm lowers the risk of forks and increases the system responsiveness by shrinking waiting time for confirmations (4).

e) Forks and updates: In addition, miners software is sometimes updated to fix bugs or add functionalities. This also can create forks, as different nodes might handle transactions differently depending on their software versions. We usually distinguish:
- soft forks where the transactions considered valid by the new version are also valid for the old version.
- hard forks where the transactions considered invalid by the old version might be valid for the new version.

While it is complicated to synchronize the software over public blockchains due to the huge amount of anonymous participants and potential disagreements among them, it is easily feasible on private chains where members know each other and can quickly come to a mutual agreement (5).

This overview of blockchain architecture components highlighted the main differences between public and private implementations. In the next section, we present major consensus algorithms used in private blockchains that enables high transaction throughput compared to the regular algorithm such as the Proof of Work or the Proof of Stake.

B. Consensus Algorithms

As we saw in the previous section, generating and adding blocks to the linked list is a crucial step in terms of security and scalability. In this section we describe the major consensus algorithms that are used today in private blockchains by which participants choose the block generator node (*i.e.,* miner).

1) Proof of Elapsed Time: At each cycle, miners wait for a given random time. The first miner for which the waiting time has elapsed is selected to validate a block before repeating this process. In other words, the miner with the shortest wait time is elected the leader. For instance, this mechanism is used in Quorum, a permissionned fork of Ethereum [10]. This is one of the least secured consensus protocol aimed for private blockchain with high trust among the block makers since miners can cheat on the random generation.

However, Intel proposed to use Trusted Execution Environment (TEE) such as Software Guard Extensions (Intel SGX) to ensure safety an randomness. The idea is to protect code and data from disclosure or modification through the use of enclaves, which are protected areas of execution [12]. Another way to prevent miners to cheat and monopolize leadership on the network is to add a voting consensus on top of the elapsed time protocol. While this protocol is less secured that its competitors (without using safety protocols), it remains very fast and scalable.

2) Leader based consensus: This category gathers algorithms that attempt to solve the problem of *agreement* in which distributed/asynchronous processes have to agree on a

TABLE I. BENCHMARKING OF ENTERPRISE BLOCKCHAIN PLATFORMS AND TECHNOLOGIES.

Company	Platform	Popularity	Activity (GitHub)	Consensus	Performance	Data Encryption	Smart Contract	Virtual Machine	Oracle	Additional Features
Coin Sciences	Multichain	Medium	Medium	Round robin (diversity)	100-1000/s	No	No	No	No	Assets, streams
J.P. Morgan	Quorum	High	Medium	Time and vote based	12-100/s	Yes	Yes	EVM	-	-
IBM	Hyperledger Fabric	High	High	PBFT	10k-100k/s	Yes	Yes	Chaincode	No	-
Coinprism	OpenChain	Medium	Low	Partitioned	Thousands/s	Yes	No	Yes	-	Assets, side-chains
Chain	Chain Core	High	High	Federated consensus	N/A	Yes	No	Yes	-	Assets
R3	Corda	Medium	High	BFT, etc.	N/A	Yes	Yes	JVM	Yes	Assets, market
Monax	Monax	High	Medium	Tendermint	10k/s	No	Yes	EVM	-	-

leader process believed to be valid. While it is mathematically proven that this problem is impossible to solve [13] (*i.e.,* if one process fails, the termination/validity of election cannot be guaranteed), in practice some algorithms manage to achieve the agreement with a probability close to 1 [14]. This is achieved by using either an oracle that provides random numbers or failure detection (or a combination of these) [15][14].

3) Practical Byzantine Fault Tolerance (PBFT): PBFT is a replication algorithm that is able to tolerate Byzantine faults [16]. To put it simply, this algorithm ensures the consistency of consensus as long as $2/3$ of the network's nodes are safe (*i.e.,* not malicious or faulty). This is enabled by replicating behaviors (*i.e.,* state machines) of generating nodes and applying protocols for choosing a leader among them. However, this method requires that all the generating nodes know each other since they need to communicate. In other words, all the parties have to agree on the exact list of participants.

4) Federated Byzantine Agreement (FBA): The FBA consensus protocol breaks the prerequisite of the unanimously accepted membership list of PBFT by letting any new participants to grow the network [17]. Each participant knows some nodes that are considered as important and waits that the majority of them and the majority of the rest of the network agree on a new transaction before considering it as valid. The main limitation of this protocol remains its performance: it costs a lot of messages (*i.e.,* communication over the network) that come with latency.

5) Tendermint: Tendermint is another byzantine fault tolerant algorithm based on a state machine that enables nodes to propose and vote for the next validator [18]. It makes the assumption that the network is partially synchronized since the time factor is central to this protocol. For each new block, a validator node is selected in a round-robbin manner which has to propose a block. This block is then spread into the network and has to gather more than two third of votes of *members* within a given time period before being added to the blockchain. However, these members are selected based on their stake and thus ties trust to resource ownership.

6) Diversity Mining consensus: The mining diversity consensus approach was proposed by MultiChain [19] to resolve the case where one participant of a private bockchain could monopolize the mining process. The solution consists in limiting the number of blocks that might be created by one specific miner within a given time period. This implicitly enforce a round robin schema where each permitted miner must create blocks in rotation. A *mining diversity* parameter defines the strictness of the rotation, where a value of 1 means that every permitted miner should be included in the rotation, whereas a value of 0 means no restriction at all.

C. Benchmarking Existing Technologies

Blockchain is currently under extensive research and development, leading to a high market fragmentation, with more than 20 different technologies and frameworks, which have been released by companies, open-source communities and universities. Table I compares the key characteristics of some popular blockchain technologies, especially for the context of enterprise and consortium based case studies.

III. Applications

A. Classification

Many criteria can be used to classify blockchain applications. We will start here with a technology approach: we will first describe use cases where the blockchain is a self-sufficient technology, and then move on to explore new scenarios, where the combination of blockchain and other technologies/competencies can enable new perspectives.

1) Assets and Data Management: The blockchain can be used as an immutable distributed ledger where transactions are timestamped by block, therefore directly enabling asset tracking, ownership transfert certification and history record. The appearance of Bitcoin and its crytotoken has opened an incredible potentiality: it is now possible to create a digital asset that is unique. Indeed, as opposed to a MP3 file that can be infinitely duplicate without alteration, it is not possible to give away a bitcoin without losing it. So from this perspective, a bitcoin resembles a physical object, except it lives in the

digital world, and tying digital asset and difital identity leads to proof of ownership.

2) Market Places: The wish to exchange or sell these digital assets on a peer-to-peer network, *i.e.,* without relying on intermediary, leads to blockchain-based market places. One potential benefit from a decentralized market place could be reduced costs. By removing a trusted third party, and its associated fees, the created value should be better shared between the buyer and the seller. Another benefit is the system resilience, as it is not relying on a central actor that could be a single point of failure. But more importantly, we could foresee that this model generates more end-user empowerment. From an operation perspective, we could rely on peers, for example for conflicts arbitration. From a content perspective, we could imagine that end-users would have better control about which of their data is shared, and with whom. For this market place use case, optimization algorithms and multi-agents simulation could be used to enrich the trading mechanisms.

3) Data Exchanges and Processes Automation: Smart contracts (a.k.a. chaincodes) are programs that can be used to automate company internal processes, or even B2B/B2C services. But in order to be efficient, smart contracts should be combined with Artificial Narrow Intelligence (ANI) so that the workflow is smooth and fast. This requires domain specific data to build up this machine expertise. Then, ANI-driven calls to smart contracts will be possible, making the most of automation and data knowledge. Obviously, this functional layer can sit on top of the marketplaces defined above, therefore creating new opportunities for data monetization in an automatic manner (say between connected objects) or in a permissioned manner (say a marketplace where the end user keeps control of its data and decides who can access them, for how long and to do what).

4) Decentralized Autonomous Organizations (DAOs): A decentralized autonomous organization (DAO) is an organization that relies on rules implemented in smart contracts. This requires yet another level of sophistication, namely artificial general intelligence, to make it fully efficient. One can imagine that this new type of ventures could decide how to invest its money to crowdfund projects for profit, how the eventual benefits from its proceedings should be distributed, how the governance should evolve in case of disagreement and so on. So in this example, the DAO fully replicates in the blockchain world the behavior of a company board. And one can imagine many other life-similar examples.

B. Case Studies

While we described above the blockchain use cases from a technology perspective, we can also use a sectorial approach to map them.

1) Finance and Insurance: The first blockchain application was the cryptocurrency Bitcoin. But many use cases have followed since. As an example, Chaincore implements the distributed ledger technology for clearing and settlement, as a way to lower costs and improve efficiency. It can also be used to issue and trade assets, such as bonds, in a decentralized market place (see the proof-of-concept from Caisse Des Depots in France). The blockchain can also help with processes such as KYC (Know Your Customer), by sharing the proof of identity and not the data itself between banks (see KYC-chain as an implementation example). Finally, crowdsharing an insurance deductible can be a good DAO application in the insurance sector.

2) Energy: With the rise of solar panels and other green sources of energy, the energy production is becoming more decentralized and offers a promising field for blockchain applications. As an example, the distributed ledger technology can be used to certify the source of energy production, therefore guaranteeing that it is green. It can also be used to trade energy at the local grid level, between individual producers and consumers (see the proof-of-concept from LO3 Energy in Brooklyn). We can imagine further benefits in the home where devices can schedule their energy charging to optimize costs and exchange data autonomously between them.

3) Mobility: In this sector, the distributed ledger technology can be used to safely store the car data (for example, its mileage). Another example is arcade city, which is a blockchain-based ridesharing platform that matches passengers and drivers. So this is basically an uber-like service, in a decentralized architecture. One more example would be a decentralized transportation ecosystem, where people can use a same token to ride on a bus, rent a bike or carpool, without any central authority to organize its operation.

4) Logistics: In this sector, the distributed ledger technology can be used to track an asset. For example, Everledger tracks diamonds to ensure their authenticity, Provenance can track food origin to guarantee its sanitary safety. Another example would be using a blockchain to create a collaborative IT system, which matches transporters and customers timetable for efficient delivery.

IV. Research Directions and Opportunities

Blockchain is currently under extensive research and development from both the academia and the industry, however, there are still major challenges to be overcome before mass market penetration and adoption. In this section, we highlight major research directions and opportunities that we believe are important to investigate.

A. Data Analysis and Visualization

A blockchain being no more than a ledger of transactions between accounts, data from a blockchain can be seen as no more than nodes connected by occasionally existing multi-property edges. Under which structural form should they be tackled depends on the aimed the analysis. From a blockchain *network supervision* point of view, crucial in a private companies consortium, the relevant data aggregation level is the block, with a time-series scheme. From the point of view of auditing the *quality of the user activity*, transactions should be considered the atomic level to investigate, under a graph scheme, and more specifically under a time-varying graph (TVG) scheme [20].

The aim of efficiently auditing a blockchain brings several challenges:

1) Real-time analysis: Because of the possibility of forks, there is no such thing as absolute reliability of the data retrieved from the blockchain. It is decreasingly high toward the most recent blocks data, as one only get the version of the ledger stored on a node at a given time, so that a blockchain-specific time-dependent reliability weight has to be determined. This procedure must be highly dependent on the chosen consortium governance scheme.

2) Exploitable visual representation of TVG: From a graph point of view, each edge (transaction) represents a unique and directed communication bridge between nodes, having an infinitesimally narrow timewidth. To be able to graphically analyse a blockchain networks, or to compute common graph indicators such as centrality or community borders, systematic smart ways to define edges weight based on non-Dirac delta function in time have to be conceived.

3) Smart contract internal transactions unravelling: Unless explicitely coded as so, the transactions from and to smart contracts, or from smart contracts to users, are no written down in the ledger, and this can be used for transaction obfuscation allowing token laundering [21], Ponzi scheme [22] or other uses where the blockchain only serves itself. In order to determine whether or not blockchain transactions are related to real-world event, or more generaly what it is used for, studies on specific key quality indicators related to smart contract have to be conducted.

B. Blockchain Audit

Data immutability is generally put forward when referring to Blockchain technologies. However, as already discussed in Sub-section II.A.a, the written data could still be tampered and the blockchain rebuilt as long as the majority of the participants (or miners) have reached a consensus. This is especially true in consortium and private blockchains where the number of miners is generally limited in comparison with public Blockchains.

In this context, it becomes extremely difficult for a regulation authority to audit consortium based Blockchains and to check whether the data and transactions have been tampered with or not. A commonly adopted solution, consists in piggybacking data hashes from the consortium Blockchain into the Bitcoin network, by embedding those hashes inside the OP_RETURN field of Bitcoin transactions. However, this contribute in polluting and increasing the size of the Bitcoin network with nonsense and non-financial data.

More recently, alternatives solutions have been proposed to reduce the impact of piggybacking on public blockchains, including the concepts of side-chains and notary chains whose main objective is to make it extremely hard for malicious users and/or the network participants to alter the blockchain data.

C. Governance

The governance in a private blockchain assigns authority and responsibility among the consortium members. It determines nodes that will be able to create blocks (*i.e.,* miners), to read/write data, to contribute in the consensus mechanism (*e.g.,* voting for a miner) and/or to participate in decisions for the system evolution (*e.g.,* software updates, allow new nodes to join the system etc.). This power distribution has an impact not only within the system but also on the business model of the use case.

Costs linked with the system activity such as the system set-up, its execution or maintenance are shared within the consortium according to the governance scheme. It also affects future incomes or losses at a business level since the governing nodes decide the rules of the system. For example, the majority of governing nodes can decide to allow the membership a of new company into the consortium that is concurrent with a member who has no power over this decision that could jeopardize the viability of the system.

The viability of the system can also be affected by the governance definition. In many cases, to be durable, the consortium has to be able to grow by allowing new members to integrate the system. It is the case for example of new services over blockchain like dematerialized car service books. The more companies join the consortium such as car manufacturers, car repair shops or insurance companies, the more durable and available is the system. On the other hand, the power is dissolved with the growth of consortium.

One should also take into account the impacts on the business model when building the governance scheme as it will be discussed in the next section.

D. Incentives and Business Models

Blockchain solves the issues of trust between actors in situations of exchange where the temptation of cheating is high by *removing* this need of trust. Any business model based on a solution that would not claim to solve a trust issue would inevitably fail, as its solution could be replaced by a less constraining and probably already existing centralized system.

In a blockchain whose users are exclusively individuals, the pecuniary incentives must ensure that, because members either receive additional incomes or just lessen their expenses, they find a financial interest in participating to the process. In a consortium of commercial entities however, it should be pointed out that the simple fact *not to* be part of the consortium might represent a handicap that could lead to loss of turnover or customers attrition, because of the latter attraction to blockchain promises and interest in financial incentives.

E. Data Privacy

Data privacy is an imperative for enterprise blockchains. But lets first distinguish anonymity and privacy. A transaction is considered anonymous if we cannot identify its owner, whereas a transaction is called private if the object and the amount of transaction are unknown.

We have seen many schemes on public blockchains to improve privacy: Stealth Addresses, Pedersen Commitments, Ring signature, Homomorphic encryption, Zero-knowledge-proof. No scheme can hide the sender, the receiver and the amount at the same time, so we see actual implementations mixing these techniques in order to achieve the desired level of privacy. In addition, there are some known drawbacks such as computational time, so further research is needed. But we can expect that these initiatives on public blockchains will drive improvement on enterprise blockchains privacy as well.

F. Security

Guaranteeing End to End security means identifying vulnerabilities and mitigating risks at each element level and at the system level. This goes beyond looking at the blockchain building blocks (consensus, distributed network, cryptographic tools) and includes evaluating the virtual machine, the Smart Contracts, the Oracle, the user client, the hardware component, the keys management and PKI, etc. Some areas of research are the following: Formal verification of smart contracts, Usage of trusted platform modules for key storage, Identification of the

different types of attack vectors and their counter strategies (sybil attacks, double spending attacks, distributed denial of service attacks, botnet attacks, storage specific attacks, censorship), Audit (detect issues a priori or a posteriori), Supervision (detect issues during run time).

G. Scalability

As usual, there is always a trade-off between costs, security and performance. Because participants are known in enterprise blockchains, the scalability issue is therefore easier to solve, as compared to public blockchains. Yet, in order to achieve scalability, we first need to keep in mind the usage context and the performance metrics we want to optimize: transactions throughput, validation latency, number of participant nodes, number of validating nodes, energy costs, computation costs, storage costs or other criteria? As always, remember the trade-off principle: A round robin consensus algorithm will scale well, but the participants need to be honest. A PBFT algorithm can recover from malicious behaviors (up to 1/3) but the validating nodes should not be too many (tens of nodes at most) if the system is to work [23]. All in all, scalability is an active area of research and we can mention some initiatives such as: fragmenting the global ledger into smaller sub-ledgers run by sub-groups of nodes, removing old transactions in order to optimize the storage, using a hierarchy of blockchains (transactions are done at a higher level and settled optionally afterwards in the blockchain), and so on.

V. Conclusions

The Blockchain technology represents a major paradigm shift in the way business applications will be designed, operated, consumed and marketed in the near future. In this paper, we analyzed the technical component of this technology and we provided a taxonomy of applications and use cases. Finally, we highlighted the major research challenges that need to be addressed before achieving mass market penetration, including the issues related to governance, audit, scalability, incentives, data privacy, security and data analytics.

Acknowledgment

This research work has been carried out under the leadership of the Institute for Technological Research SystemX, and therefore granted with public funds within the scope of the French Program Investissements d'Avenir.

References

[1] Z. Zheng, S. Xie, H.-N. Dai, and H. Wang, "Blockchain challenges and opportunities: A survey," International Journal of Web and Grid Services, 2017, pp. 1–23.

[2] K. Christidis and M. Devetsikiotis, "Blockchains and smart contracts for the internet of things," IEEE Access, vol. 4, 2016, pp. 2292–2303.

[3] M. Mettler, "Blockchain technology in healthcare: The revolution starts here," in 2016 IEEE 18th International Conference on e-Health Networking, Applications and Services (Healthcom), Sept 2016, pp. 1–3.

[4] "stampd.io: A document blockchain stamping notary app," accessed: 2017-07-01. [Online]. Available: https://stampd.io/

[5] A. Yasin and L. Liu, "An online identity and smart contract management system," in 2016 IEEE 40th Annual Computer Software and Applications Conference (COMPSAC), vol. 2, June 2016, pp. 192–198.

[6] R. Dennis and G. Owen, "Rep on the block: A next generation reputation system based on the blockchain," in 2015 10th International Conference for Internet Technology and Secured Transactions (ICITST), Dec 2015, pp. 131–138.

[7] F. Tian, "An agri-food supply chain traceability system for china based on rfid blockchain technology," in 2016 13th International Conference on Service Systems and Service Management (ICSSSM), June 2016, pp. 1–6.

[8] "Regulation (eu) 2016/679 of the european parliament and of the council of 27 april 2016," accessed: 2017-07-01. [Online]. Available: http://eur-lex.europa.eu/eli/reg/2016/679/oj

[9] V. Buterin, "On Public and Private Blockchains," 2015, accessed: 2017-07-01. [Online]. Available: https://blog.ethereum.org/2015/08/07/on-public-and-private-blockchains/

[10] "Quorum | J.P. Morgan," accessed: 2017-07-01. [Online]. Available: https://www.jpmorgan.com/country/US/EN/Quorum

[11] S. Nakamoto, "Bitcoin: A peer-to-peer electronic cash system," 2008, accessed: 2017-07-01. [Online]. Available: http://www.cryptovest.co.uk/resources/Bitcoin%20paper%20Original.pdf

[12] V. Costan and S. Devadas, "Intel sgx explained." IACR Cryptology ePrint Archive, 2016, p. 86, accessed: 2017-07-01. [Online]. Available: https://pdfs.semanticscholar.org/2d7f/3f4ca3fbb15ae04533456e5031e0d0dc845a.pdf

[13] M. J. Fischer, N. A. Lynch, and M. S. Paterson, "Impossibility of distributed consensus with one faulty process," Journal of the ACM (JACM), vol. 32, no. 2, 1985, pp. 374–382.

[14] A. Mostefaoui and M. Raynal, "Leader-based consensus," Parallel Processing Letters, vol. 11, no. 01, 2001, pp. 95–107.

[15] L. Lamport, "Paxos made simple," ACM Sigact News, vol. 32, no. 4, 2001, pp. 18–25.

[16] M. Castro and B. a. Liskov, "Practical Byzantine fault tolerance," in Third Symposium on Operating Systems Design and Implementation (OSDI), vol. 99, 1999, pp. 173–186.

[17] D. Mazieres, "The stellar consensus protocol: A federated model for internet-level consensus," 2015, accessed: 2017-02-10. [Online]. Available: https://www.stellar.org/papers/stellar-consensus-protocol.pdf

[18] J. Kwon, "Tendermint: Consensus without mining," 2014, accessed: 2017-07-01. [Online]. Available: https://cdn.relayto.com/media/files/LPgoWO18TCeMIggJVakt_tendermint.pdf

[19] "MultiChain | Open source private blockchain platform," accessed: 2017-07-01. [Online]. Available: http://www.multichain.com/

[20] A. Casteigts, P. Flocchini, N. Santoro, and W. Quattrociocchi, "Time-varying graphs ans dynamic networks," International Journal of Parallel, Emergent and Distributed Systems, vol. 27, no. 5, 2012, pp. 387–408.

[21] M. Moser, R. Bohme, and D. Breuker, "An inquiry into money laundering tools in the bitcoin ecosystem," in eCrime Researchers Summit (eCRS), 2013. IEEE, 2013, pp. 1–14.

[22] M. Bartoletti, S. Carta, T. Cimoli, and R. Saia, "Dissecting ponzi schemes on ethereum: identification, analysis, and impact," CoRR, vol. abs/1703.03779, 2017.

[23] M. Vukolic, "The Quest for Scalable Blockchain Fabric: Proof-of-Work vs. BFT Replication," Open Problems in Network Security (iNetSec), 2015, pp. 112–125.