NIST Author Manuscript

NIST Author Manuscript

NIST Author Manuscript

Author Manuscript

Accepted for publication in a peer-reviewed journal

NIST National Institute of Standards and Technology • U.S. Department of Commerce

Published in final edited form as:

Computer (Long Beach Calif). 2019 February ; 52(2): . doi:10.1109/mc.2019.2898162.

Rethinking Distributed Ledger Technology

Rick Kuhn,

Jeff Voas,

Dylan Yaga

National Institute of Standards and Technology, Gaithersburg, MD 20899, USA

Abstract

Blockchains were designed to solve the problem of double-spending in cryptocurrencies, and the success of the Bitcoin design has generated vastly more interest than previous proposals for digital currencies. Blockchains are being used in other areas as well, but the design choices that made blockchains effective for cryptocurrencies often do not fit well with other applications. In this paper we review the properties of distributed ledger technology (DLT) for use in typical data management applications and show how two recently developed distributed ledger ideas can be used to retain valuable aspects of blockchain while simplifying design and adding new but often necessary capabilities to permissioned distributed ledger applications. In particular, we are interested in the ability to delete or modify blocks, and the ability to provide a timestamping mechanism to provide a highly accurate time for applications that are time order dependent.

Introduction

While most of the excitement around blockchain stems from its use in cryptocurrencies, designers are beginning to find interesting ways to solve system problems using blockchain and other forms of distributed ledger technology (DLT). The most commonly used data structure for distributed ledgers is the blockchain. A key feature of a blockchain based system is the decentralized, replicated data synchronized among separate network nodes, which may be geographically dispersed. There is substantial discussion around some terms in DLT. We believe it is better to distinguish blockchain systems based on their permission model: permissioned or permissionless; since it is directly tied to the technology; where private or public may apply to the visibility of the network or ledger itself.

With its features providing distributed, trusted data using no central server, DLT seems a natural tool for many complex distributed systems, and a number of implementations have been proposed. However, some environments and applications are not well suited to using an append-only ledger. For example, an analysis of DLT for the international banking consortium SWIFT found that the permissionless model used by Bitcoin and other cryptocurrencies “does not provide the level of trust, transparency, and accountability required by the financial industry” [4]. The SWIFT analysis noted that permissioned ledgers are helpful, but “existing implementations of permissioned ledgers remain basic”.

kuhn@nist.gov .

NIST Author Manuscript NIST Author Manuscript NIST Author Manuscript

Of particular concern is the “immutable” aspect of transactions recorded in blockchains [6] [8]. As noted by the European Banking Institute, “once an error is embedded in the blockchain, this may be highly problematic, legally, in that often law requires the ability to rectify errors as a matter of law in a way foreign to DLT” [5]. Some proponents of DLT argue that errors can be corrected by issuing a new transaction which supersedes the older erroneous transaction. The argument put forth is that the error should remain recorded in the ledger (to provide a full history of events as they happened) and not be overwritten or removed. While this may be possible for certain transactions, privacy laws lead to additional complications, as discussed below.

Although the decentralization, data replication, and data integrity assurance of the original blockchain design have some valuable properties for building information systems, other properties of conventional blockchains are not a good match for applications beyond cryptocurrency [9][10][11]. Consequently, modifications to distributed ledger designs are being developed to meet new needs, such as faster consensus times [9][11]. Blockchains are a valuable distributed ledger technology for providing trust, but there are many ways to construct distributed ledgers. We describe an alternative that provides the trust features of blockchains with a more flexible data structure and ordering protocol.

DLT and Data Management

A distributed ledger, as the name suggests, is a distributed record of transactions, maintained by consensus among a network of peer to peer nodes (possibly geographically dispersed). The most widely recognized form of DLT is the blockchain structure, which provides the basis for cryptocurrencies and a variety of other applications. Most currently available distributed ledger designs using blockchain provide certain properties:

- *Pseudo-anonymity* – especially for cryptocurrency, blockchains enable participation using only identifiers. Permissioned blockchains may not include this property.
- *Public access, transparency* – every participant can see all transactions on the blockchain, although they may be anonymized. This property may also not be provided in permissioned systems.
- *Small transaction size* – Blockchains were originally designed for monetary transactions, so messages are assumed to be relatively small.
- *Immutability* – As a consequence of the linked chain of cryptographic hashes of records, a change to one record would cause the hash of subsequent records to be invalid, so changes require recomputing the entire chain. As a result, it is generally intractable to change any record in a blockchain.
- *Proof of work or other expensive consensus models* – A consequence of the need to prevent double spending. Permissioned blockchains do not generally need this feature and can use simpler consensus.
- *Block ordering guarantee* – the consensus mechanism ensures ordering of the blocks and therefore transactions, preventing the possibility of double spending.

NIST Author Manuscript

NIST Author Manuscript

NIST Author Manuscript

- *Decentralization* – there is no central authority for records. With each update, records are dispersed to peer nodes simultaneously, who ensure the updates are correct.
- *Replication and Synchronization guarantee* – transactions are duplicated across all nodes of the network, so that every node has an identical copy of all transaction records, current to the most recent update cycle. Consensus protocols are designed such that when the consensus is complete, all nodes have an identical copy of the distributed ledger records.
- *Integrity protection* – Cryptographic hashes are used to guarantee that records have not been changed.

We compare these properties with the needs of more typical applications of distributed data storage and retrieval in Table 1. Note that six of the nine blockchain properties designed for cryptocurrency are at odds with the requirements of many other applications.

Need for an Alternative Solution

The mismatch between blockchain properties and many application needs has led to a number of problems in applying blockchain designs to data management problems. For example, Bitcoin is designed to provide some degree of anonymity in transactions (i.e., only public identifiers, not real-world identities are used), but the law may prohibit anonymity for many types of transactions and require participants to be identified for tax or other purposes. Laws such as the European Union General Data Protection Regulation (GDPR), that require the ability to delete privacy relevant information, may limit the type of information that can be stored in a blockchain [6][7][8].

For system engineers, the price of distributed trust is often added complexity. The design choices that were made to incorporate anonymity and prevent double spending in blockchains often lead to seemingly unnecessary complications when applied to areas beyond cryptocurrency. For example, immutability has resulted in designs where alterable records must be kept off of the blockchain, with only pointers to them stored in the blockchain itself. Alternatively, some designs involve encrypting data on the blockchain, then destroying the encryption key to “delete” the data. Neither of these options may be desirable for many applications, as the first option leads to unnecessary complications, and the second risks the data being decrypted in the future, when data must be protected for decades. These are serious design issues for supporting privacy requirements such as those of GDPR, resulting in proposals such as an “editable blockchain” [12] using new forms of hashing. A consensus algorithm is needed to guarantee record ordering in the absence of a central time authority (i.e., transactions are ordered based on group consensus, rather than time of entry into a system), and this ordering is used to prevent double spending. Designs for access control using blockchain may involve tokenizing permissions, then passing these to users, and spending down the value to remove a permission from a user. All of these strategies are needed to take advantage of blockchain’s trust properties, but blockchains would probably not be used if a more conventional database could provide the desired distributed trust.

NIST Author Manuscript NIST Author Manuscript NIST Author Manuscript

At first glance, blockchain solutions for applications such as supply chain, financial settlement, and others may appear to offer nothing more than added complexity in comparison with a conventional database. However, when more than one organization is involved, the decentralized trust of blockchains and other distributed ledgers can be a tremendous advantage. For example, consider regulated industries where audit is a part of doing business. Every node on the system can have a full set of records detailing the movement of assets. Any shared database can keep track of asset movement, but DLT adds trust by maintaining current, integrity-protected records at every organization, making it easy to audit the process. Thus, the financial industry views *full traceability* and *simplified reconciliation* of transactions among the key advantages of DLT [4]. We can view DLT as adding a layer of distributed trust to the problem of data storage and retrieval, clearly a desirable property, but industry is still struggling with how to use DLT in practical ways.

A Permissioned Distributed Ledger Model for Decentralized Trust

Much of the current DLT research seems to center on how to get around properties that were baked in to blockchain. Adaptations such as faster consensus algorithms are gradually moving DLT from its origin in cryptocurrency towards a more general-purpose database technology. But instead of tweaking blockchain designs, we can rethink the idea of a distributed ledger to reflect the needs of data management applications as discussed earlier.

Can we provide a simpler model, that gives the decentralized trust of a blockchain, but otherwise behaves as a conventional database? In this section, we describe an approach to achieving this goal, using two recent developments: a *data block matrix* [13], and *verified time* [14]. The data block matrix retains hash-based data integrity guarantees, while allowing controlled modification or deletion of specified records, with integrity guarantees for all other records. A data block matrix can be implemented in a decentralized system to provide data replication among peers. The verified time protocol allows guaranteed timestamps to be used in place of consensus algorithms to ensure record ordering.

The data block matrix uses an array of blocks, with hash values for each row and column. This structure makes it possible to delete or modify a particular block with hash values assuring that other blocks have not been affected. An example is shown in Fig 1. Suppose that it is desired to delete block 12, by writing all zeroes to that block, or otherwise modifying it. This change disrupts the hash values of $H_{3,-}$ and $H_{-,2}$ for row 3 and column 2. However, the integrity of all blocks except the one containing “X” is still ensured by the other hash values. That is, other blocks of row 3 are included in the hashes for columns 0, 1, 3, and 4. Similarly, other blocks of column 2 are included in the hashes for rows 0, 1, 2, and 4. Thus the integrity of blocks that have not been deleted is assured. Blocks can be deleted by overwriting with zeroes or other values, with one row and one column hash recalculated; specifically, after deleting block *i, j,* row *i* and column *j* hash values are recalculated.

As shown in [13], the data structure ensures the following properties:

- *Balance*: upper half (above diagonal) contains at most one additional cell more than the lower half.

NIST Author Manuscript NIST Author Manuscript NIST Author Manuscript

- *Hash sequence length*: number of blocks in a row or column hash proportional to $\sqrt{N}$ for a matrix with N blocks, by the balance property.
- *Number of blocks*: The total number of data blocks in the matrix is $N^2 - N$ since the diagonal is null.
- *Block dispersal*: No consecutive blocks appear in the same row or column

Clearly, this data structure is not suited to all DLT applications, but it offers features that are difficult to provide with a conventional blockchain. Our goal is not to replace blockchains, but to offer a new form of data storage structure that provides the integrity guarantees of blockchain with the addition of reversibility, which can be used in a wide range of applications. A comparison is shown in Table 2.

While looking at distributed ledger designs, the role of time is often an afterthought. Some DLT systems transactions have no inherit timestamp for when they were submitted to the system and instead adopt the time in which they were included into the blockchain ledger. This approach has worked for applications where just having a transaction accepted is good enough (e.g., we do not need to know that a cryptocurrency transaction was submitted down to the millisecond, just that it was indeed submitted and eventually recorded in the ledger).

However, when time dependent situations arise, a timestamp becomes more important, and knowing when a transaction was submitted to a system may be more important than when it was incorporated into the ledger. Often, systems will rely on local system time, or a network time – both may differ from one system to the next. Distributed ledgers must be able to operate in environments that include rules mandated by government or contract. For some applications, the ordering of transactions into blocks within the blockchain may not be enough, and there is a need for a global timestamp service, providing verified time [14]. Time is a key component of this because things happen outside of the blockchain that matter for applications – orders must be fulfilled by a specified date and time, legal papers filed on schedule, and so on, requiring timestamps of events that take place outside of the blockchain.

A global time-stamping approach would include an agreed upon and accepted service. This approach could incorporate a high resolution blockchain time mechanism, such as the open-source Chainpoint protocol [16], into the distributed ledger to produce a final and agreed upon timestamp. Chainpoint uses the Network Time Protocol with the NIST Randomness Beacon to provide provable timestamps, and might be adapted to the needs described here.

Conclusions

The blockchain data structure and proof-of-work protocol were designed to solve the problem of double spending in cryptocurrencies. Although blockchain has found many applications outside of cryptocurrency, many of its features are not well suited to common data management applications, leading many to argue that distributed ledgers are only databases with more complex features. As we have described, the added trust of distributed ledgers is a valuable feature, providing greatly simplified auditability and verification of actions among multiple parties in applications such as supply chain and others.

NIST Author Manuscript NIST Author Manuscript NIST Author Manuscript

The blockchain design for hash-based integrity verification provides trust, at the cost of an inability to delete or update records, leading to design complications that would not arise with conventional database management systems. Similarly, the sequencing guarantees of blockchain consensus protocols are needed for cryptocurrency in the absence of a universal timestamp. Moreover, actions within the distributed ledger must be connected with other actions in the real world, through accurate timestamps. We have presented a new architecture that provides the trust features of blockchains, with characteristics that allow for simpler designs and greater practicality in conventional data management problems. We believe this alternative can lead to new approaches to incorporating trust into distributed systems applications.

References

[1]. Locher T, Obermeier S, & Pignolet YA (2018). When Can a Distributed Ledger Replace a Trusted Third Party?. arXiv preprint arXiv:1806.10929.

[2]. Maesa DDF, Mori P, & Ricci L (2017, June). Blockchain based access control. In IFIP International Conference on Distributed Applications and Interoperable Systems (pp. 206–220). Springer, Cham.

[3]. Ouaddah A, Abou Elkalam A, & Ait Ouahman A (2016). FairAccess: a new Blockchain-based access control framework for the Internet of Things. Security and Communication Networks, 9(18), 5943–5964.

[4]. Le Borne F, Treat D, Dimidschstein F, & Brodersen C (2017). SWIFT on distributed ledger technologies-Delivering an industry-standard platform through community collaboration.

[5]. Zetzsche DA, Buckley RP, & Arner DW (2017). The Distributed Liability of Distributed Ledgers: Legal Risks of Blockchain.

[6]. Berberich M, Steiner M,“Blockchain Technology and the GDPR,” 2 Eur. Data Protection Law Review. 422, 2016.

[7]. Chang H, “Blockchain: Disrupting Data Protection?”, Privacy Law and Business Intl. Rpt, Nov.2017.

[8]. Kharif O, “Is Your Blockchain Doomed?”, Bloomberg Business Week, Mar. 22, 2018.

[9]. Simonite T, “Banks Embrace Blockchain’s Heart but not it’s Soul”, MIT Technology Review, Sept. 24, 2015.

[10]. Lehdonvirta V, “The Blockchain Paradox: Why Distributed Ledger Technologies may do Little to Transform the Economy”, Oxford Internet Institute, Nov. 21, 2016.

[11]. UK Govt. Chief Scientific Advoser, “Distributed Ledger Technology: beyond Block Chain”, https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/492972/gs-16-1-distributed-ledger-technology.pdf

[12]. Condliff J, “Is an Editable Blockchain the Future of Finance?”, MIT Technology Review, Sept. 20, 2016.

[13]. Kuhn DR, A Data Structure for Integrity Protection with Erasure Capability, NIST Cybersecurity Whitepaper, May 31, 2018.

[14]. Stavrou A, & Voas J (2017). Verified Time. Computer, 50(3), 78–82.

[15]. Lombardi MA, 2002. NIST time and frequency services. US Department of Commerce, Technology Administration, National Institute of Standards and Technology.

[16]. Chainpoint – Innovations in Blockchain Timestamp Proofs. https://blog.tierion.com/chainpoint-innovations-in-blockchain-timestamp-proofs

NIST Author Manuscript

NIST Author Manuscript

NIST Author Manuscript

	0	1	2	3	4	
0	•	1	3	7	13	$H_{0,-}$
1	2	•	5	9	15	$H_{1,-}$
2	4	6	•	11	17	$H_{2,-}$
3	8	10	12	•	19	$H_{3,-}$
4	14	16	18	20	•	$H_{4,-}$
	$H_{-,0}$	$H_{-,1}$	$H_{-,2}$	$H_{-,3}$	$H_{-,4}$	etc.

Figure 1.
Data block matrix with numbered cells

NIST Author Manuscript

NIST Author Manuscript

NIST Author Manuscript

Table 1.

Comparing characteristics of DLT applications

Cryptocurrency	**Finance, supply chain, e-commerce, etc.**
1. Anonymity	ID required for contracts or government regulation
2. Public access/transparency	Controlled access
3. Small transaction size	Range of message sizes up to large documents, images
4. Immutable records	Changes and deletions, often required by law
5. Proof of work	Flexible consensus models
6. Block ordering guarantees	Timestamps often required
7. Decentralization	Same in many applications
8. Replication	Same in many applications
9. Data integrity guarantees	Same in many applications

NIST Author Manuscript

NIST Author Manuscript

NIST Author Manuscript

Table 2.

Blockchain and data block matrix features

Blockchain – provides integrity, sequencing	**Data block matrix** – provides integrity, erasure
Integrity protection, no erasure possible	Integrity protection for all blocks not erased
Double-spend problem solved by distributed transaction ordering guarantees	Ability to erase values obviates need for ordering guarantees through consensus algorithms
Ordering guarantees require consensus algorithms	Ordering guarantees granted by time authority