Preprint: Roberts, J. D., Defranco, J. F., & Kuhn, D. R. (2023). Data Block Matrix and Hyperledger Implementation: Extending Distributed Ledger Technology for Privacy Requirements. ACM Distributed Ledger Technologies: Research and Practice, 2(2), 1-11.

Data Block Matrix and Hyperledger Implementation: Extending Distributed Ledger Technology for Privacy Requirements

Joshua D. Roberts

National Institute of Standards and Technology, joshua.roberts@nist.gov

Joanna F. DeFranco

The Pennsylvania State University, jfd104@psu.edu

D. Richard Kuhn

National Institute of Standards and Technology, d.kuhn@nist.gov

Distributed ledger technology (DLT), including blockchain, has a number of properties that make it useful for distributed systems. However, the immutability of blockchain and most forms of DLT make it impossible to delete data, as is required for compliance with many privacy rules regarding personally identifiable information. Thus, there is a need for DLT that can provide the integrity-preserving property of DLT while also allowing support for privacy rules. The data block matrix (DBM) is a variant of distributed ledger technology. It provides the integrity assurance of blockchain but allows for controlled revision or deletion of data. This property is essential for using DLT in applications that must guarantee privacy requirements by the deleting of a user's private data at their request. The DBM design solves the blockchain privacy conflict thus expanding the range of blockchain applications by also allowing exception management. It has been implemented and is available (https://csrc.nist.gov/projects/redactable-distributed-ledger) as a configurable option for Hyperledger Fabric (HF), with a proof-of-concept application for data sharing in a health care environment. Other potential applications include logistics management and digital currency. This paper will cover the DBM properties and data structure, the DBM implementation in HF, and a use case and application design of the DBM implementation using the pharmaceutical industry supply chain.

CCS CONCEPTS • Data structures design and analysis • Security and privacy

Additional Keywords and Phrases: blockchain, data structures, distributed ledger, security and privacy

1 INTRODUCTION

Blockchain technology provides integrity protection through immutability. This is fundamental to solving the problem of cryptocurrency double spending [1]. But this integrity protection is also a valuable property for other applications in addition to digital currency. For example, logistics and e-commerce are candidate applications, because "distributed trust" is valuable as any node in the network can have guaranteed accuracy of the data without relying on a single, centrally stored record. The mechanism to ensure that data blocks are not altered is a chain of hashed values, that are impossible to delete or change without changing the other blocks. As such, a blockchain is immutable because it is not possible to change any bit without requiring the entire chain to be rebuilt, which is generally infeasible for a large blockchain.

The added trust of distributed ledger technology (DLT) is a valuable feature as it provides greatly simplified auditability and verification of actions among multiple parties in applications. However, the immutability property of conventional blockchains can be difficult to use in many distributed system applications. Many privacy requirements (e.g., EU General Data Privacy Regulation (GDPR)), allows users to request that their private data be deleted thus making the immutability property of a blockchain solution impractical when privacy rules are required [2] [3]. In other words, redactable distributed ledgers are useful and needed for some applications [4]. Proposals for redactable DLT include a blockchain based on chameleon hash [5] and data block matrix (DBM) from NIST [6] [7].

The DBM is a new form of DLT, which provides the integrity assurance of blockchain along with controlled revision or deletion of data. This is an essential property when using DLT in applications that require the support of privacy requirements (e.g., deletion of private data). We introduce a new data structure and associated operations, via the DBM, to extend the range of applications for blockchain solutions by solving the conflict between privacy regulations and blockchain. This is accomplished by allowing block edits or data deletion when there are privacy requirements or needs for exception management. Applications for DBMs are currently being investigated and developed within international standards bodies [8][9][10].

The DBM solution was implemented in Hyperledger Fabric (HF) to make it usable and practical in real-world applications. With this implementation, users can configure individual Fabric channels to use either a DBM or a blockchain. Once a channel is configured to use a DBM, the integration is designed to be transparent to the HF user. Essentially, this design focuses on adapting the block storage mechanisms to be compatible with the DBM functionality. The remainder of this paper will cover the DBM properties and data structure, the DBM implementation in HF, and a use case and application design of the DBM implementation using the pharmaceutical industry supply chain.

2 RELATED WORK

A number of proposals have been presented to address the conflict between blockchain properties and the 'right to erasure' regulations of GDPR and related laws. Under GDPR, personal data are defined as "any information concerning an identified or identifiable natural person" [11]. Some argue that personal data should not be included in blockchains [12]. A limitation of this approach is that most current uses of blockchains involve personal data such as financial transactions and purchase history.

A more fundamental problem with avoiding the inclusion of personal data on blockchains is that EU rulings have indicated that "Personal data which have undergone pseudonymization, which could be attributed to a natural person by the use of additional information should be considered to be information on an identifiable natural person." [13]. For example, an IP address can be considered personal if it can be linked to an individual at a particular time [14; 15]. The rules regarding pseudo-anonymized data conflict with attempts to address the erasure problem by storing only indirect links on the blockchain and erasing the data pointed to by the links. Even if a user's name is not stored on the chain, the indirect link may itself be considered personal data if it can be traced to an individual [13]. In this case, the question is what is included in the blockchain hash. If it is a person's name, then that is private data and needs to be deleted, which would of course disrupt the blockchain. If it is a pseudonym or other indirect linking to a person, then under existing rulings that indirect link to an individual is also considered privacy sensitive and must be deleted. This approach thus would require managing potentially multiple levels of indirection, increasing the risk of errors while still possibly not solving the problem.

Other proposals for supporting 'right to erasure' in blockchains involve encrypting the personal data, so that only possessors of the key can access it [16, 12]. To "delete" the data, the key can be destroyed, although the data is clearly still present. In addition to failing to truly support the letter of the law for data deletion, a serious limitation of this approach is that the encrypted data must be made permanently inaccessible, or at least for many decades. The history of cryptography shows that cryptosystems will eventually become vulnerable. This is why the Data Encryption Standard (DES) was withdrawn as a standard and replaced by the Advanced Encryption Standard (AES) only 26 years after DES was initially released.

The most commonly proposed technical approach to resolving the erasure problem is the use of chameleon hash functions [5; 17]. In this approach, a specialized hash function is used, allowing the computation of collisions if a particular trapdoor (or key) is known. Thus, to remove personal data, a new block of data can be computed that erases the original, but computes to the same hash value, leaving the chain of hashes in the blockchain undisturbed. The tradeoffs with this approach are that it must use a non-standard hash function and requires key management operations outside the conventional blockchain functions. A mechanism must also be provided for controlling which parties have access to the block erasure functions. In contrast, the DBM Hyperledger implementation allows use of the standard SHA 256 hash function, rather than a chameleon hash.

One alternative approach [39] to allow block modification while using the standard hash function was proposed in [40]. This method would replace the block to be modified or deleted with a "special record" that includes the edited block data, a hash of the previous block, and a flag indicating that the block was modified. Nodes of the blockchain then use the hash of the non-modified blocks to do the blockchain hash, with a separate hash signature for the new value of the blocks, and the protocol specifies that the nodes use the updated value.

A feature in common among the approaches is the requirement for controlling access to block erasure, in a permissioned distributed ledger. Some form of conventional access control would be required for this task. It should be noted that the determination of whether data deletion requests are legitimate is a legal and policy question, which would be addressed by policies established by laws and managing organizations. A choice of DBM Hyperledger or chameleon hash functions would depend on the requirements of the data sharing being designed, and the engineering tradeoffs involved.

3 DBM DATA STRUCTURE

The data block matrix structure uses two hash values to provide the key feature of blockchain, distributed integrity protection, while also allowing for controlled edits. Figure 1 shows a basic DBM with numbered data blocks. Where each block may contain multiple or single record transactions. Every row or column is terminated with a block that contains a hash of that row or column (e.g., H0,− is the hash of row 0). In addition, various forms of the hash structure are also possible: hash value can be stored in the last block of the row or column instead of a separate hash block or concatenate hashes from each block in a row or column, similar to the blockchain process. It is important to note that a hash concatenation would then serve as the hash value for that row or column.

	0	1	2	3	4	
0						$H_{0,-}$
1						$H_{1,-}$
2						$H_{2,-}$
3			X			$H_{3,-}$
4						$H_{4,-}$
	$H_{-,0}$	$H_{-,1}$	$H_{-,2}$	$H_{-,3}$	$H_{-,4}$	

Figure 1: Basic Data Block Matrix

	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 2. Data Block Matrix with Numbered Cells

To illustrate an example, consider that a block labeled ”X” may be deleted by either writing all zeroes to that block (overwriting the data) or can be revised with different values. Either of these changes will affect the hash values of H3,− and H−,2 for row 3 and column 2 as they will be recalculated due to that change. However, the integrity of all blocks, except the one containing ”X”, remain ensured by the other hash values (i.e., the 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 are assured. An algorithm to maintain this structure is shown below along with a description of its properties.

Within the data structure, blocks are numbered 1..*k*, and are added to the data structure starting with cell 0,1. Please note: It is desirable to keep cells on the diagonal null, for reasons explained later. Variables *i*,*j* are column indices, and swap(*i*,*j*) ex-changes the values of *i* and *j*, i.e., $i' = j$ and $j' = i$. With this algorithm, cells are filled as shown in Algorithm 1.

ALGORITHM 1: DBM Construction

Invariant: $i < j \wedge odd(\mathrm{B}i,j) \vee i > j \wedge even(\mathrm{B}i,j)$
$i \leftarrow 0$
$j \leftarrow 1$
$\mathrm{B} \leftarrow 1$
while new blocks, do
    if $i = j$ then
        add null block B*i,j* // diagonal
        $i \leftarrow 0$
        $j \leftarrow \mathrm{j} + 1$
    else if $i < j$ then
        add block B*i,j* // upper half
        $\mathrm{B} \leftarrow \mathrm{B} + 1$
        swap(*i,j*)
    else if $i > j$ then
        add block B*i,j* // lower half
        $\mathrm{B} \leftarrow \mathrm{B} + 1$
        $\mathrm{j} \leftarrow \mathrm{j} + 1$
        swap(i,j)
    end if

end while

4 PROPERTIES

It is shown in [7] that certain desirable properties are maintained with this data structure. These features allow for efficient storage and retrieval of data blocks.

Balance property. Cells are filled in a balanced manner, so that the upper half (above diagonal) contains at most one additional cell more than the lower half. This property ensures balance among hashed block sequences as the matrix size increases.

Hash chain length. The number of blocks in a row or column hash chain is proportional to sqrt(N) for a matrix with N blocks. This property provides efficient computation of hashes across block sequences.

Block numbering. All even numbered blocks are placed below the diagonal and all odd numbered blocks are placed above the diagonal. This property, and the block dispersal property (below) make it possible to delete or edit two consecutive blocks i and i+1 without both disrupting the same row or column hashes, to simplify hash updating and reduce performance impact of hash computation when blocks are edited.

Block dispersal. No consecutive blocks appear in the same row or column, i.e., for any two blocks labeled $a,b,$ where $b = a + 1,$ in rows i_a and $i_b,$ and columns j_a and j_b respectively, $i_a \neq i_b$ and $j_a \neq j_b$.

5 HYPERLEDGER IMPLEMENTATION AND DESIGN CONSIDERATIONS

Hyperledger Fabric was identified as the best open source blockchain solution to implement the features of the DBM, because it is modular, robust, and widely used in domains with strong privacy requirements. It is also a permissioned blockchain which fits this use case well. The implementation is available on https://csrc.nist.gov/projects/redactable-distributed-ledger.

Prior to implementing the DBM using HF, we had to understand how the block storage mechanism worked in HF. The anatomy of a blocks consists of a header and data. The block header stores the hash of the block data and other relevant metadata. The block data contains a set of transactions which store read-write sets. These read-write sets reference the key-value pairs that are or will be stored on the ledger [18]. Blocks are stored as byte arrays in files on the peer file system. These files are append-only as that is the primary operation needed in a traditional blockchain implementation. To facilitate query operations, blocks and transactions are indexed based on their location in the file system. These indexes are stored in a key-value database also on the peer's system [19].

We completed two iterations to implement the DBM using HF. In the first iteration, we focused on changing how blocks were created and stored by reserving each block for a specific key, where new writes to the key would be appended to the end of the block's data byte array. Upon key deletion, the whole block was deleted, and the row/column hashes updated. Essentially, this approach fundamentally changed the way transactions were packaged and how blocks were built.

Unfortunately, such a fundamental change to the Fabric system required too many changes to the code base outside of the block storage mechanism and was deemed unfeasible.

In the second iteration we decided to restrict our modifications to only the block storage mechanism (i.e., how blocks are stored, retrieved and validated). In order to support the write operations of the DBM we took advantage of the already in place key-value database used for indexing blocks to store the actual blocks. Then, we modified the internal process of adding a block to the ledger. When a new block is added to the matrix, the following steps occur:

(a) The block is added, and the row and column hashes are updated based on the insertion process described in section (3).
(b) The block's data is then inspected to identify keys that have been written and deleted in valid transactions only (transactions marked as invalid are ignored).
(c) Each key referenced in the block data is added to an index that tracks the blocks a key is referenced in.
(d) For each key that is marked as deleted in the block:
    1. Get the blocks the key is referenced in from the index described in step (c).
    2. Sequentially, but in no specific order, remove the deleted key from the blocks. The row and column hashes should be updated after each block is updated. This is to prevent collisions if blocks share a row or column. Then, update the block's metadata to record the transaction that caused the block to be updated. This provides validation that a block was correctly updated.

The DBM functionality in Hyperledger Fabric is illustrated below. Figures 3a and 3b show the modified block storage mechanism.

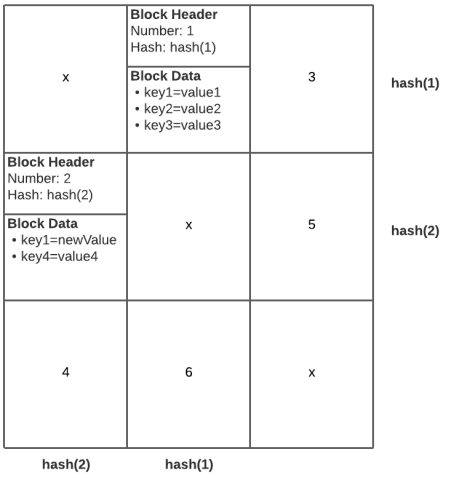

Figure 3a. Two blocks are added

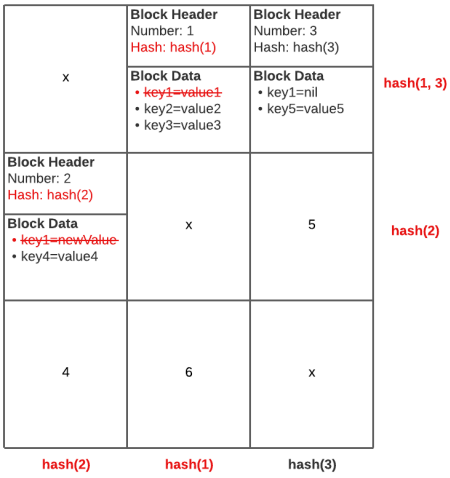

Figure 3b. A block deletes key "key1"

Blocks are stored according the DBM structure where odd numbered blocks are above the diagonal and even numbered blocks are below the diagonal. Hashes of the blocks in each row and column are also stored as denoted by *hash(b…b_n)* in the diagrams. If a new block is added to the matrix and deletes a key, any block that had previously written a value for that key is updated and the key-value pair is removed. The block is then rehashed, and the corresponding row and column hashes are also updated. Affected blocks are processed one at a time to ensure each update to a block only affects the appropriate row and column hashes. For example, in Figure 3b, block 3 deletes the value of "key1". As a result, block 1 will be updated and then block 2.

6 USE CASE: PHARMACEUTICAL SUPPLY CHAIN

Counterfeit medication is an economic burden on society and major problem to the viability of the pharmaceutical industry. More importantly, counterfeit medications can be dangerous, even deadly. It is estimated that 1 in 10 medical products in low- and middle-income countries are substandard or falsified [20]. This is largely due to lack of "tools and technical capacity to enforce quality standards" and "inadequate regulation and governance compounded by unethical practice" within the pharmaceutical supply chain [20]. Clearly it is essential that a clear and consistent definition of drug supply-chain visibility is carefully communicated across industries. In other words, being able to verify the source of raw materials as well as maintaining visibility as the product moves through the chain to avoid ethical issues (where low- and middle-income countries are not regulating and identifying fake medication), lack of the capability to distinguish fake from valid medications, and lack of testing facilities and drug inspectors [21]. This is also a safety issue, in that when there is a problem with drugs and drug ingredients, there needs to be an effective and accurate recall tracking method.

6.1 Blockchain Solutions and Related Work

A multidisciplinary review of current and emerging digital solutions to combat the counterfeit drug incidences cited blockchain as one of technologies with the potential to better establish drug supply chain root [22]. Several researchers have designed or proposed applications based on blockchain technology that intends to trace the drug pathway starting from the manufacturer to the consumer [23]. One application design includes five starting nodes for the prototype to represent the participants in a pharmaceutical supply chain: manufacturer, wholesaler, retailer, Food and Drug Administration (FDA), and the consumer. With such an application, a consumer would receive a code (to be scanned with a mobile phone camera) along with drug product purchased. Upon scanning the code, the consumer would be directed then to the portal, where they can view a display of drug distribution history. Other ideas proposed for using blockchain as a solution incorporate IoT to track product information, location, time, and dates as well as sensors to monitor temperature for temperature-sensitive drug storage [21] [24] [25] [26] [27] [28]. One product, modum.io, uses smart contracts to ensure temperature compliance [24]. Another is proposed to be a mobile application linked to the NEM blockchain nano wallet to purchase medication [21].

6.2 DBM Solution

A blockchain solution offers the needed transparency to mitigate the drug counterfeit problem. However, patient data must be protected according to privacy rules that conflict with the blockchain immutability property [2]. Privacy regulations that require a "right to erasure" of personal data, where a consumer can order their private data held by an organization to be deleted, include the European Union General Data Protection Regulation (GDPR) regulations [27], and US laws in California, Virginia, and Colorado [29, 30, 31, 32, 33]. GDPR also influences practices internationally because companies exchanging data with EU organizations must abide by GDPR requirements [32, 34].

In a pharmaceutical supply chain design, where each node represented a part of the supply chain (e.g., supplier, production, distribution, consumption), the consumer PII information (e.g., ID) could be deleted to comply with GDPR. Additionally, to comply with HIPAA privacy and 42 part 2 regulations, the consumer may want to delete a certain medication (data stored on the blockchain) to conceal a sensitive diagnosis, and patients can refuse to have their identity and data stored permanently in the blockchain [27]. The DBM makes it possible to comply with these regulations.

6.3 Access Control

Next Generation Access Control (NGAC), an ANSI/INCITS ABAC standard, is an access control framework to support attribute-based access control policy [35] [36]. NGACs framework is a set of configurable data elements and relations to express access control policy such as for the read and write operations on data elements. This is made possible by NGAC's use of graphs to store policies. Changing the policy to grant or revoke access to a resource can be as simple as adding or removing an edge in the graph. The graph is stored in memory which allows for quick decisions and real time updates. An NGAC Policy Decision Point (PDP) can be deployed as a shim between the Hyperledger Fabric smart contracts and end user applications to control access to sensitive patient data.

.

6.4 Pharmaceutical Supply Chain Design Using the HF DBM

An implementation of the DBM using HF would address the compliance of the patient privacy and confidentiality regulations discussed earlier. In this section we describe an approach that works alongside of a traditional pharmaceutical blockchain – where the goal is drug transparency and safety.

The overall system architecture is shown in Figure 4. The four members of the HF network are: manufacturer (M), pharmacy (Ph), patients' (P), and doctor's office (D). The members have their own channel (i.e., a DBM ledger). In this design, the M, Ph, and D members join the "Drug Channel" to monitor the supply chain history of a particular drug. The drug channel accessibility is as follows: M has *write* access to update drug safety information (i.e., they may get this information from a traditional pharmaceutical blockchain supply chain) and Ph and D have *read* access.

The "Patients' Channel" will contain each patient's record. Ph and D are able to join the patient channel in order to gain *read* and *write* access. The doctor will update the patient information such as prescriptions and diagnoses for each patient. In addition, the D's can query data from the drug channel then write prescriptions to the patient channel. Ph can check safety information form drug channel and read a prescription from the patient channel to fill, and write notes (e.g., prescription filled, issue with supply chain etc.).

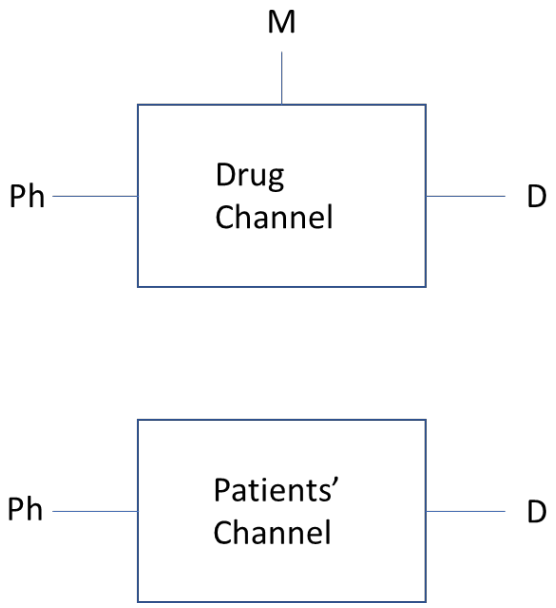

Figure 4. Overall Architecture

Specific Design features (Figure 5):

- Doctor member creates and joins the patients' channel and adds a record for each new patient.
- Patients have control over their record. That control is provided by the doctor's NGAC policy – (i.e., read only with deletion requests, adding caregiver access, and adding other specialists).
- After the doctor adds a patient record to the patients' channel, each patient needs to register as an end user at the doctor's office (note: they can only see their own record).
- The API services that interact with the DBM channels can be located on a central server or other architecture.
- Channels can have both HF identities (i.e., systems or applications) and human end user identities (note: the human identities use HF identities to interact with the DBM). Both types of identities able to query each DBM channel.

The scenario shown in Figure 5 is as follows: The end users use the application to call an API specific to a Doctor's organization. An NGAC layer will authorize users to access permitted information. The server will use a HF identity to call the appropriate smart contract function(s) where the interaction with the patient's record is facilitated.

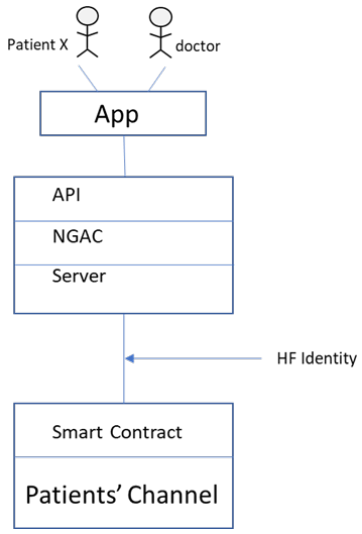

Figure 5: Organization Structure

This implementation adds the Pharmaceutical DBM supply chain system to a traditional blockchain enabled supply chain to address privacy compliance along with transparency and safety.

7 DISCUSSION AND CONCLUSIONS

Distributed ledgers in general are useful to secure and validate data in a distributed system. However, if there are privacy laws that require certain private data to be deleted, a traditional distributed ledger such as blockchain may be unacceptable as a solution due to its immutability feature. To gain the data security and integrity of a blockchain as well as meet privacy regulations for data deletion, the DBM technology is needed, because DBM provides the distributed trust and data integrity guarantees of blockchain, without the immutability property.

We chose Hyperledger Fabric to implement the DBM because it provided the base of a production-ready blockchain solution that we were able to modify slightly to achieve the DBM functionality. During the development process we determined that using a key value datastore, which Fabric uses for indexing blocks, was the better solution to storing the blocks themselves instead of the file system Fabric uses to store blocks. The key value datastore better facilitates the read and write operations of the blocks instead of the file system which prioritizes read operations.

HF is designed for high transaction rates, scalability, and other features essential in large-scale distributed systems. We believe that the DBM component should not have a significant impact on HF's current performance in production databases, as only minimal changes are needed to integrate the DBM component (i.e., HF's standard SHA will continue to be used). The possible slight performance impact is only from writing operations since any update to the DBM triggers a rehash of the corresponding row and column. Thus, the latency for write operations is comparable to regular fabric. The only difference is writing the bytes to a key value database instead of a file.

The performance of delete operations depends on the how many blocks are updated as a result of the delete. The more blocks that need to be updated to delete a key, the more processing that is required to reflect the delete in the DBM, but the increase is linear in the number of blocks to update. Thus, any latency for delete operations is linear to the number of blocks updated by the delete operation. Read operations are not affected and clearly comparable to regular fabric in that the only difference is reading the bytes from a key value database instead of a file. Future work will measure and evaluate performance characteristics of the Hyperledger Fabric DBM. The process of sharing this code in a Hyperledger Foundation lab repository is also underway.

The pharmaceutical supply chain was an effective use case to demonstrate the need for a distributed ledger with privacy built in, but there are many other possible applications and in multiple domains such as clinical trials [37], law enforcement, open banking and healthcare data [38]. For example, a secure federated data sharing system could use a DBM two ways. First to store, share, and manage user-to-attribute assignments. A DBM is ideal for this requirement because of the need to delete and alter attribute assignments. An additional DBM could be used to store a data summary and a pointer to the detailed data located at different organizations.

The DBM technology is not disruptive to database operations but is an effective solution to improve security and integrity of data when privacy is a requirement. That is, it is not designed to be a replacement for blockchain, but as an alternative component for distributed systems designers where compliance with privacy regulations is essential. As such, it has potential for extending the range of applications for distributed ledger technology to problems where blockchain use would be impractical due to regulatory requirements.

DISCLAIMER:

Commercial products may be identified in this document, but such identification does not imply recommendation or endorsement by NIST, nor that the products identified are necessarily the best available for the purpose.

REFERENCES

[1] Nakamoto, S., "Bitcoin: A Peer-to-Peer Electronic Cash System," 2008, https://bitcoin.org/bitcoin.pdf. [Accessed 5-16-22]

[2] Dirk A Zetzsche, Ross P Buckley, and Douglas W Arner. The distributed liability of distributed ledgers: Legal risks of blockchain. *U. Ill. L. Rev.*, page 1361, 2018.

[3] Henry Chang. Blockchain: Disrupting data protection? *Privacy Law and Business International Report, November*, 2017.

[4] Kondapally Ashritha, M Sindhu, and KV Lakshmy. Redactable blockchain using enhanced chameleon hash function. In *2019 5th International Conference on Advanced Computing & Communication Systems (ICACCS)*, pages 323–328. IEEE, 2019.

[5] Ateniese, G., Magri, B., Venturi, D., & Andrade, E. (2017, April). Redactable blockchain–or–rewriting history in bitcoin and friends. In *2017 IEEE European symposium on security and privacy (EuroS&P)* (pp. 111-126). IEEE.

[6] Rick Kuhn, Dylan Yaga, and Jeffrey Voas. Rethinking distributed ledger technology. *Computer*, 52(2):68–72, 2019.

[7] D Richard Kuhn. A data structure for integrity protection with erasure capability. *NIST Cybersecurity Whitepaper*, 2018. https://admin.cms.csrc.nist.gov/csrc/media/Projects/enhanced-distributed-ledger-technology/documents/NIST.CSWP.25.pdf

[8] European Telecommunications Standards Institute, Introduction to Permissioned Distributed Ledger (PDL), Jan 24, 2022 [accessed 5/11/22]

[9] European Telecommunications Standards Institute, IPv6 Security, Cybersecurity, Blockchain ETSI GR IP6 031 V1.1.1 (Tech. Rpt. 2020). [accessed 5/11/22]

[10] European Telecommunications Standards Institute, standards work item: IPv6 and Cloud using DataBlockMatrix for Food Supply Chain Tracking and Tracing IPv6-based DataBlockMatrix [accessed 5/11/22]

[11] GDPR, Recital 26. Regulation 2016/679 of the European Parliament and of the Council on the Protection of Natural Persons with Regard to the Processing of Personal Data and on the Free Advancement of Such Data, and repealing Directive 95/46/EC, 2016 O.J. L 119/1 [

[12] Mearian, L., "Will Blockchain Run Afoul of GDPR? Yes and No", Computerworld, May 7, 2018. https://www.computerworld.com/article/3269750/will-blockchain-run-afoul-of-gdpr-yes-and-no.html
[13] EPRS | European Parliamentary Research Service Scientific Foresight Unit (STOA) PE 634.445 – July 2019 https://www.europarl.europa.eu/RegData/etudes/STUD/2019/634445/EPRS_STU(2019)634445_EN.pdf
[14] Kelleher, D., "In Breyer decision today, Europe's highest court rules on definition of personal data", Oct. 19, 2016 https://iapp.org/news/a/in-breyer-decision-today-europes-highest-court-rules-on-definition-of-personal-data/
[15] Breyer v. Germany, https://curia.europa.eu/juris/document/document.jsf?text=&docid=184668&pageIndex=0&doclang=en&mode=lst&dir=&occ=first&part=1&cid=1116945
[16] Baker-Hostetler, Five Things Blockchain Companies Need to Know About the GDPR https://www.bakerlaw.com/webfiles/Privacy/2018/Brief/Five-Things-Blockchain-Cos-Need-to-Know-About-GDPR.pdf
[17] Deuber, D., Magri, B., & Thyagarajan, S. A. K. (2019, May). Redactable blockchain in the permissionless setting. In *2019 IEEE Symposium on Security and Privacy (SP)* (pp. 124-138). IEEE.
[18] Christian Cachin, et al. "Hyperledger Fabric: A Distributed Operating System for Permissioned Blockchains", January 30 2018, https://arxiv.org/abs/1801.10228v2 [accessed 10/17/2022].
[19] Hyperledger Fabric [Source code] https://github.com/hyperledger/fabric/tree/release-2.3.
[20] World Health Organization, "1 in 10 Medical Products in Developing Countries is Substandard or Falsified," 2019. November 28, 2017, https://www.who.int/news-room/detail/28-11-2017-1-in-10-medical-products-in-developing-countries-is-substandard-orfalsified. [Accessed 2-16-22]
[21] Ganesan Subramanian, Anand Sreekantan Thampy, Nnamdi Valbosco Ugwuoke, and Baghwan Ramnani. Crypto pharmacy–digital medicine: A mobile application integrated with hybrid blockchain to tackle the issues in pharma supply chain. *IEEE Open Journal of the Computer Society*, 2:26–37, 2021.
[22] Tim K Mackey and Gaurvika Nayyar. A review of existing and emerging digital technologies to combat the global trade in fake medicines. *Expert opinion on drug safety*, 16(5):587–602, 2017.
[23] Patrick Sylim, Fang Liu, Alvin Marcelo, Paul Fontelo, et al. Blockchain technology for detecting falsified and substandard drugs in distribution: pharmaceutical supply chain intervention. *JMIR research protocols*, 7(9):e10163, 2018.
[24] Shankar D Nawale and Rahul R Konapure. Blockchain & iot based drugs traceability for pharma industry. In *2021 IEEE International Conference on Engineering, Technology and Innovation (ICE/ITMC)*, pages 1–4.IEEE, 2021.
[25] Thomas Bocek, Bruno B Rodrigues, Tim Strasser, and Burkhard Stiller. Blockchains everywhere-a use-case of blockchains in the pharma supplychain. In *2017 IFIP/IEEE symposium on integrated network and service management (IM)*, pages 772–777. IEEE, 2017.
[26] Krishna Mohan Botcha, Vedula VSSS Chakravarthy, et al. Enhancing traceability in pharmaceutical supply chain using internet of things (iot) and blockchain. In *2019 IEEE International Conference on Intelligent Systems and Green Technology (ICISGT)*, pages 45–453. IEEE, 2019.
[27] Ahmad Musamih, Khaled Salah, Raja Jayaraman, Junaid Arshad, Mazin Debe, Yousof Al-Hammadi, and Samer Ellahham. A blockchain-based approach for drug traceability in healthcare supply chain. *IEEE Access*, 9:9728–9743, 2021.
[28] Sandip Jangir, Ajit Muzumdar, Alok Jaiswal, Chirag N Modi, Sheetal Chandel, and C Vyjayanthi. A novel framework for pharmaceutical supply chain management using distributed ledger and smart contracts. In *2019 10th International Conference on Computing, Communication and Networking Technologies (ICCCNT)*, pages 1–7. IEEE, 2019.
[29] Wiley, S., Barwig, J., "California Law Requires Legal Compliance Scruitiny to Maintain Pharma Data Sharing," PharmExec.com, January 6, 2020, https://www.pharmexec.com/view/california-law-requires-legal-compliance-scrutiny-maintain-pharma-data-sharing [accessed 5/10/22].
[30] BUCKLEY, "Virginia Passes Amendments on CDPA for Data Deletion," BUCKLEY, March 3, 2022, https://buckleyfirm.com/blog/2022-03-03/virginia-passes-amendments-cdpa-data-deletion [accessed 5/10/22].
[31] Houser, K. A., & Voss, W. G. (2018). GDPR: The end of Google and Facebook or a new paradigm in data privacy. *Rich. JL & Tech.*, *25*, 1.
[32] Klar, M. (2020). Binding Effects of the European General Data Protection Regulation (GDPR) on US Companies. *Hastings Sci. & Tech. LJ*, *11*, 101.
[33] Forcier, M. B., Gallois, H., Mullan, S., & Joly, Y. (2019). Integrating artificial intelligence into health care through data access: can the GDPR act as a beacon for policymakers?. *Journal of Law and the Biosciences*, *6*(1), 317.
[34] Jones, M. L., & Kaminski, M. E. (2020). An American's Guide to the GDPR. *Denv. L. Rev.*, *98*, 93.
[35] INCITS, INCITS 565-2020 -Information Technology – ext Generation Access Control, 4/24/2020 ,https://standards.incits.org/apps/group_public/project/details.php?project_id=2328 [accessed 5/10/22].
[36] Ferraiolo, D., Chandramouli, R., Hu, V., Kuhn, R., "A Comparison of Attribute Based Access Control (ABAC) Standards for Data Service Applications," NIST Special Publication 800-178, October 2016, https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-178.pdf [accessed 5/10/22].
[37] Ferraiolo, D., DeFranco, J.F., Kuhn, D.R., Roberts, J., "A New Approach to Data Sharing and Distributed Ledger Technology: A Clinical Trial Use Case." IEEE Netw. 35 (1), 4-5
[38] DeFranco, J., Ferraiolo, D., Kuhn, R., Roberts, J., "A Trusted Federated System to Share Granular Data Among Disparate Database Resources," IEEE Computer, 54 (3)

[39] Bernabe, J. B., Canovas, J. L., Hernandez-Ramos, J. L., Moreno, R. T., & Skarmeta, A. (2019). Privacy-preserving solutions for blockchain: Review and challenges. *IEEE Access*, *7*, 164908-164940.
[40] V. Demianets and A. Kanakakis, *Distributed Ledger With Secure Data Deletion—Revision 1.4*. Stockholm, Sweden, 2016. https://github.com/TarantulaTechnology/Documents-Blockchain-vol-002/blob/master/Distributed%20Ledger%20With%20Secure%20Data%20Deletion.pdf