Can Blockchains and Data Privacy Laws be Reconciled?

A Fundamental Study of How Privacy-Aware Blockchains are Feasible

Christoph Stach
christoph.stach@ipvs.uni-stuttgart.de
University of Stuttgart, IPVS / AS
Stuttgart, Germany

Clémentine Gritti
clementine.gritti@canterbury.ac.nz
University of Canterbury
Christchurch, New Zealand

Dennis Przytarski
dennis.przytarski@ipvs.uni-stuttgart.de
University of Stuttgart, IPVS / AS
Stuttgart, Germany

Bernhard Mitschang
bernhard.mitschang@ipvs.uni-stuttgart.de
University of Stuttgart, IPVS / AS
Stuttgart, Germany

ABSTRACT

Due to the advancing digitalization, the importance of data is constantly increasing. Application domains such as smart cars, smart cities, or smart healthcare rely on the permanent availability of large amounts of data to all parties involved. As a result, the value of data increases, making it a lucrative target for cyber-attacks. Particularly when human lives depend on the data, additional protection measures are therefore important for data management and provision. *Blockchains*, i. e., *decentralized*, *immutable*, and *tamper-proof* data stores, are becoming increasingly popular for this purpose. Yet, from a data protection perspective, the immutable and tamper-proof properties of blockchains pose a privacy concern. In this paper, we therefore investigate whether blockchains are in compliance with the *General Data Protection Regulation* (*GDPR*) if personal data are involved. To this end, we elaborate which articles of the GDPR are relevant in this regard and present technical solutions for those legal requirements with which blockchains are in conflict. We further identify open research questions that need to be addressed in order to achieve a *privacy-by-design blockchain system*.

CCS CONCEPTS

• **Security and privacy → Distributed systems security**; **Privacy protections**; *Usability in security and privacy*.

KEYWORDS

blockchains, immutable, tamper-proof, GDPR, privacy assessment

ACM Reference Format:
Christoph Stach, Clémentine Gritti, Dennis Przytarski, and Bernhard Mitschang. 2022. Can Blockchains and Data Privacy Laws be Reconciled?: A Fundamental Study of How Privacy-Aware Blockchains are Feasible. In *The 37th ACM/SIGAPP Symposium on Applied Computing (SAC '22), April 25–29, 2022, Virtual Event*. ACM, New York, NY, USA, 10 pages. https://doi.org/10.1145/3477314.3506986

SAC '22, April 25–29, 2022, Virtual Event

ACM ISBN 978-1-4503-8713-2/22/04…$15.00
https://doi.org/10.1145/3477314.3506986

1 INTRODUCTION

"Data is the new oil." is a commonly cited quote by Clive Humby used to emphasize the importance of data in modern times. Unlike oil, however, which was a key driver of the *Technological Revolution* only, data are revolutionizing society as a whole. Smart cars are able to drive autonomously, smart cities enable more environment-friendly traffic control, and smart healthcare facilitates the lives of both patients and physicians, just to name a few examples. However, all of this is only possible if the data of each participant is reliably made available to all other parties involved [30].

Due to the high value which data as a commodity have in our society, they become an attractive target for cyber-criminals. However, cyber-attacks can not only cause immense economic damage, but they also pose a threat to life and limb. For instance, cyber-criminals could tamper with location data of cars or traffic management data, causing accidents in the process [18], or they could render medical data unreadable, impeding the proper treatment of patients [41]. Therefore, modern data management systems require specialized security mechanisms, especially if human lives depend on the data they are dealing with. First and foremost, they must ensure that the data are *immutable* and *tamper-proof*. Since *blockchains* possess these two key properties, it is hardly surprising that they are commonly used as decentralized data stores in such instances [39].

Despite all of these undeniable benefits of blockchains regarding the protection of sensitive data, their usage is not uncontroversial if personal data are involved. Several legal requirements imposed by the *General Data Protection Regulation* (*GDPR*) [9] cannot be satisfied when using blockchains. For example, immutability is an inherent violation of the right to be forgotten, while tamper-proofness renders practicable anonymization of data subjects impossible [40].

That is why we investigate how a *privacy-by-design blockchain system* can be achieved without losing the immutability and tamper-proofness required from a security point of view. To this end, we provide the following three contributions in this paper: (1) We elaborate on which articles of the GDPR blockchains are in conflict with when handling personal data. (2) We assess which research approaches can be used to resolve these conflicts and how they can be applied to blockchains. (3) We identify open research questions that need to be addressed in this context in order to provide an efficient privacy control in blockchains. Additionally, we outline how these research gaps can be overcome.

Figure 1: Internal Structure of a Blockchain.

The remainder of this paper is structured as follows: In Section 2, we present the fundamental principles of blockchains that are responsible for the conflicts with the GDPR. Then, in Section 3, we elaborate on the articles of the GDPR with which blockchains are intrinsically in conflict. Section 4 presents related work and addresses how our work differs from these studies. We discuss technical approaches towards a GDPR-compliant blockchain in Section 5, before identifying open research questions in this regard in Section 6. Finally, Section 7 concludes this paper.

2 BLOCKCHAIN FOUNDATIONS

Whenever multiple parties operate on a common database and share their data with each other, centralized databases often pose a problem. On the one hand, the availability of the data depends entirely on this database — i. e., it represents for all participants a *single point of failure* for their operability. On the other hand, the central authority that operates the database has full control over the data and is capable of establishing the *single point of truth*, e. g., by manipulating the data or by withholding the data. To address such issues, *distributed ledger* technology has come to forth recently. A distributed ledger represents a decentralized data storage, where each participant maintains the entire data stock. A consensus is reached among all participants as to which data or which transactions are authorized and added to the ledger [4].

The blockchain is a subtype of a distributed ledger. Its main focus is on the immutability and tamper-proofness of data. For this purpose, the blockchain bundles data into blocks that are protected against manipulation by means of digital signatures. Figure 1 illustrates this approach. From a pool of data that should be added to the blockchain, a subset is selected and assembled into an initial block, the so-called *genesis block*. When the block is ready, a so-called *cryptographic hash* is used to protect it and all the data contained in it against tampering. This hash serves not only as a signature, but also as a unique fingerprint for the block. If further data have to be added to the blockchain, a new block is generated in the same way. In addition to the actual payload data, this new block has a header which contains the hash of its predecessor. Thereby, the two blocks are inherently linked together since the cryptographic hash also protects the header against manipulation [19].

Figure 2 outlines how such a chain of blocks is managed in a blockchain architecture and how data is added to it. All participants gather their data that have to be added to the blockchain in a common data pool. From this pool, a subset of the data is bundled into a block. A consensus protocol is used to agree on which data subset is selected. To this end, there are different approaches. These approaches can be divided into two main classes: *absolute-finality consensus protocols* and *probabilistic-finality consensus protocols*. While absolute-finality consensus protocols (e. g., *Practical Byzantine Fault Tolerance* [6]) render a data record immediately valid and make it available to all parties as soon as it has been inserted into a block, probabilistic-finality consensus protocols support only eventual consistency. That is, a data record can be removed from the blockchain retrospectively under certain circumstances. As illustrated in Figure 1, the last block of a blockchain can be removed without causing any problems, since it is not yet validated by other subsequent blocks. In probabilistic-finality consensus protocols, a data record in a blockchain is therefore not considered valid until it is in a block with a certain depth, i. e., a block which has a certain number of subsequent blocks. Despite this limitation, however, probabilistic-finality consensus protocols are generally preferred in blockchains because absolute-finality consensus protocols require a single central leader that dictates which data records are valid for all parties [45].

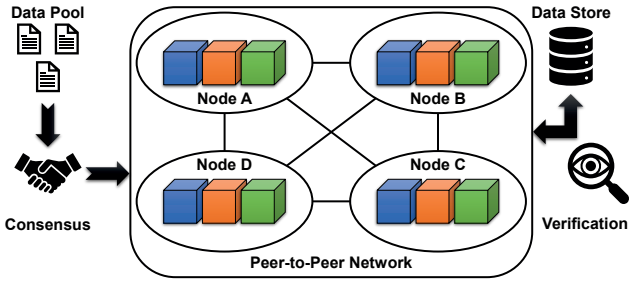

Figure 2: Simplified Architecture of a Blockchain System.

Common examples of probabilistic-finality consensus protocols are *Proof-of-Work* and *Proof-of-Stake*. In the Proof-of-Work approach, for instance, so-called *miners* have to solve a hard cryptographic challenge. The block of the miner who solves the challenge first is added to the blockchain and linked to its latest block. However, solving the cryptographic challenge requires a tremendous amount of computational power. The Proof-of-Stake approach therefore simplifies this process by randomly selecting a participant who is entitled to generate the next block of the blockchain. Although the selection is basically random, it depends on the *stake* of a participant. In other words, the more coins of the respective cryptocurrency a participant owns, the more likely s/he will be selected in the next round [3]. Despite the significantly higher energy consumption, however, it is evident in practice that the large long-established blockchain systems such as *Bitcoin* [24] and *Etherium* [44] rely on the Proof-of-Work approach.

The valid state of the blockchain is held in a peer-to-peer network. On the one hand, this ensures availability of the data at all times, and on the other hand, it guarantees immutability. Although the integrity of each block is ensured by the cryptographic hash and all inner blocks are linked to each other via their headers, the last block can be deleted unnoticed because it is not linked yet [23]. Due to the redundant distribution of the blockchain on several nodes, however, an attacker would have to control the majority of the computational power or s/he would have to be able to manipulate the majority of the nodes in order to delete the last block [31].

Since the computational complexity of managing the data in the blockchain increases with the size of the data, it is often decided not to store the actual data in the blockchain. Instead, the payload

data are stored in an external data store and only unambiguous and tamper-proof references to the data are kept in the blockchain. Verification tools can be used to check both, the integrity of the blockchain itself and the integrity of the data. To this end, there are also many approaches that extend the blockchain by a database layer on which all transactions are performed. The blockchain serves to ensure the integrity of the data. One example of this approach is *FalconDB* [26], which uses a relational database schema to manage the payload data.

Accordingly, the method of operating the blockchain also differs. Using *on-chain data management* — i. e., the complete storage of payload data in the blockchain — it is possible to perform the entire data processing in the blockchain as well, creating full transparency as all operations are publicly auditable. Using *off-chain data management* — i. e., the outsourcing of the actual payload data to an external data source — this comprehensive audit capability is omitted, but resource-efficient data management and processing is achieved instead. Since the data processing capabilities of a blockchain are inherently rather limited, off-chain data management can be used to enable big data analytics [43].

If analyses should be performed on-chain nevertheless, so-called *smart contracts* can be defined to process the data in the blockchain. A smart contract describes which transactions are to be executed on certain data when a specified condition applies. The results of these transactions are automatically added to the data pool and are thus eventually added to the blockchain in a new block after successful validation by a miner. Since smart contracts are executed automatically and their code is fully transparent as part of the blockchain, all parties can rely on the correct execution of the transactions specified therein [13].

Basically, there are two different types of blockchains: *public blockchains* and *private blockchains*. In a public blockchain, anyone can join the peer-to-peer network and thus contribute to the blockchain. Due to this high level of replication and distribution, public blockchains offer maximum protection against data manipulation. However, each participant also has full access to the data, whereas a private blockchain has a central authority that controls the blockchain and determines who can contribute to the blockchain. This reduces the number of parties that have access to the data, but partially enables the central authority to manipulate the blocks and thus the data. Although only verified participants can join the network, and they are only authorized to perform specific actions, there is not a single authority controlling all nodes. *Hybrid and consortium blockchains* are somewhere in between. Here, a group of participants has joint control over the blockchain instead of a single central authority [19]. In the context of our work, however, they can be considered as subtypes of private blockchains. For more information on the design and operation of blockchains, please refer to the given literature.

3 BLOCKCHAINS AND THE GDPR

The GDPR is intended to give data subjects full control over their personal data in an increasingly digitalized world — they must be empowered to control who has access to their data. So, it is not surprising that blockchains, which are primarily designed to make data permanently and immutably accessible to all interested parties, are in conflict with such regulations if personal data are involved.

In her study, Finck [10] therefore examines whether blockchains can be squared with the GDPR. Here, a fundamental problem becomes apparent, namely that the GDPR presumes that there is a **data controller** who is responsible for compliance with the data protection rights of data subjects (Article 24), e. g., the duty to **inform the data subjects** about the collection and processing of their personal data (Article 12 – 15). However, due to the decentralized nature of blockchains, such a central control authority does not exist. As a result, the study concludes that it is difficult to achieve GDPR-compliance, especially for public blockchains. Therefore, we focus on *permissioned blockchains* (e. g., private blockchains), where there are organizational and technological regulatory means.

Going over the articles of the GDPR in consecutive order, the first articles that seem to be relevant for blockchains are Articles 5 and 7. They specify the legal framework within which processing of personal data is allowed. If the data are processed directly in the blockchain, smart contracts can specify exactly for which purpose the data are processed as well as in which way they are processed. Thereby, a kind of **purpose limitation** (Article 5(1)(b)) is achieved. As all data stored in the blockchain are available to all participants of the peer-to-peer network, this nevertheless raises a problem with regard to **data minimization** (Article 5(1)(c)). Furthermore, since the data in the blockchain are immutable, neither the **accuracy** of the data can be improved retroactively (Article 5(1)(d)) nor any **storage limitation** (Article 5(1)(e)) can be enforced as blockchain are an *append-only* data structure. Moreover, the **consent** of the data subject (Article 7) is only reliably respected within the scope of a smart contract. If the data are processed outside of the blockchain, the agreements reached in the smart contracts no longer apply.

If the data stored in the blockchain are incorrect, a data subject also has no means to have them corrected as required by the **right to rectification** (Article 16) due to immutability and tamper-proofness. It is also not possible for a single data subject to exercise its **right to erasure** (Article 17) — due to the linkage between the blocks only the last block can be erased without destroying the structure of the blockchain. Furthermore, even the last block can only be deleted completely or not at all. Since a block contains an arbitrary subset of the data from the data pool, the deletion of a block therefore always affects the data of several data subjects.

As smart contracts execute transactions automatically and without human intervention, data subjects also have issues exercising their **right to restriction of processing** (Article 18). Only when a smart contract has been modified on the majority of the nodes of the peer-to-peer network according to the requested restrictions, the modifications will take effect. In any case, the **automated individual decision-making** (Article 22) bears another conflict potential, since smart contracts can be used for such decision-making. Therefore, the usage of smart contracts in the context of personal data has to be regarded as problematic in general.

Other regulations such as the **territorial scope** (Article 3) and the **lawfulness of processing** (Article 6) are primarily organizational issues. In our work, however, we focus on technical aspects of the blockchain that inherently conflict with the GDPR.

In summary, it can be observed that the immutability and tamper-proofness of blockchains in particular cause problems with regard to

the correction and deletion of data. Furthermore, the decentralized management of the data poses a challenge in terms of restricting access to the available data. This also results in an issue regarding a central controller that ensures compliance with data protection regulations. These issues must be overcome in order to support **data protection by design** (Article 25) for blockchains.

Table 1 outlines the key conflicts that we identify in blockchains with regard to the GDPR. Here, however, we focus only on the first and foremost technical aspects.

Table 1: Summary of the GDPR Articles with which Blockchains Inherently Conflict due to Technical Reasons.

GDPR Article	Conflicting Blockchain Property
Article 5(1)(b)	By default, blockchains do not impose a *purpose limitation*. However, a kind of purpose limitation can be achieved via well-defined smart contracts.
Article 5(1)(c)	All data on the blockchain are accessible to all nodes. Therefore, there is no *data minimization*.
Article 5(1)(d)	Data on the blockchain are immutable, i. e., their *accuracy* cannot be improved retroactively.
Article 5(1)(e)	Since blockchains are append-only data structures, *storage limitation* cannot be achieved.
Article 7	In general, blockchains do not require the *consent* of a data subject to process its data. However, this can be realized via smart contracts.
Article 12 – 15 & Article 24	In order to fulfill the duty to *inform the data subjects*, an all-embracing *data controller* is required. Yet, such a central authority fundamentally contradicts the decentralized nature of a blockchain.
Article 16	The *right to rectification* cannot be enforced in blockchains as the data are stored immutable and tamper-proof.
Article 17	The *right to erasure* cannot be enforced in blockchains as this would destroy the internal blockchain structure.
Article 18	When using smart contracts, the *right to restriction of processing* can only be enforced if a majority of blockchain nodes agree to the requested changes.
Article 22	If data processing is handled by autonomously acting smart contracts, the *automated individual decision-making* is violated.
Article 25	Only if all the ten technical issues listed above are addressed, a blockchain can support *data protection by design*.

4 RELATED WORK

Due to an increasing number of novel use cases for blockchains, there is a large body of research regarding blockchains and privacy in addition to the aforementioned study by Finck [10]. Haque et al. [12] conduct a comprehensive literature review on different aspects of how to improve the GDPR-compliance of blockchains. They conclude that there is basically a lot of prior works on the topic of GDPR-compliant blockchains. However, besides some well researched application areas, such as the healthcare sector, there are many unexplored areas where there are still open research questions regarding GDPR-compliance issues with blockchains.

This is due to the fact that studies such as those by Campanile et al. [5] or Miyachi and Mackey [21] deal with a very specific use case for blockchains in the area of smart cars or smart healthcare, respectively. They are developing a privacy-aware blockchain solution for exactly these use cases. However, these solutions require a dedicated infrastructure and cannot be transferred to other application areas and use cases due to their high degree of specialization.

While these studies focus on technical solutions to make blockchains GDPR-compliant for specific use cases, studies like the one by Shuaib et al. [33] provide administrative guidelines on how blockchains can be used to store sensitive data, such as electronic health data. In a similar direction, the work by Molina et al. [22] presents high-level design guidelines for administrators to set up a GDPR-compliant infrastructure with blockchains.

Furthermore, blockchains are also assessed from a purely legal perspective. Poelman and Iqbal [27] come to the disillusioning conclusion that GDPR-compliant blockchains are basically impossible. However, they water this statement down by adding that it might be possible in a permissioned private blockchain with appropriate extensions. However, this requires that certain limitations have to be accepted regarding the key characteristics of the blockchain, namely decentralization, immutability, and tamper-proofness. In contrast, Manteghi [20] concludes that the current data privacy laws also need to be adjusted to enable a "peaceful" coexistence with blockchains. One way or the other, there is a need for action.

Related work thus can be divided into four categories: literature reviews, privacy-aware blockchain solutions tailored to specific use cases, administrative guidelines, and legal assessments. Our work differs significantly as we investigate technical measures that can be added to any blockchain to achieve compliance with the GDPR. To this end, we discuss techniques that are well-known from other application areas and describe how they can be applied to blockchains in order to comply with data protection requirements.

5 APPLICABLE TECHNICAL SOLUTIONS

As discussed in Section 3, a major problem of blockchains with regard to the GDPR is their inability to rectify or erase personal data. To this end, we explain in Section 5.1 how hierarchical data encryption can be used to achieve fine-grained data purging in blockchains. However, better than correcting data retrospectively is to ensure that data quality is as high as possible beforehand. Therefore, we explain in Section 5.2 how attribute-based authentication can be used to prevent data from dubious sources from being included in the blockchain in the first place. Moreover, these techniques also enable purpose-based permission control, as we

show in Section 5.3. These permissions are able to minimize the disclosed information about the data subject by applying privacy filters to the data. In Section 5.4, we discuss how these filters can be used to realize the principles relating to processing of personal data in blockchains. Finally, the distributed nature of blockchains poses a problem with respect to the GDPR, as there is no distinct data controller. For this reason, we conclude in Section 5.5 with a reflection on how all these techniques can be incorporated into a central privacy control architecture for blockchains.

5.1 Data Purging by Encryption

In order to permanently delete data, there are two methods that are considered to be reliable: Either the data carrier on which the data is stored is physically destroyed or the sectors containing the data in question are overwritten several times. Both approaches guarantee that the data cannot be restored. However, there are use cases in which neither of the two methods can be applied, as organizational reasons speak against the destruction of the data carrier — e. g., because there are also data on it that must not be deleted or because the costs for frequent deletions would skyrocket — or because it is not possible for technical reasons to access an explicit data sector via the available interfaces — e. g., when dealing with databases.

In such cases, another method has proven to be extremely reliable: *data purging by encryption*. Here, all data in the data store are encrypted. This is done completely automatically in the background and is entirely transparent to the user. The keys are kept outside of the data store containing the payload data. To delete data, it is sufficient to destroy the associated keys, since subsequently the data cannot be decrypted and are therefore rendered unreadable. Since the keys are much smaller than the payload data, they can be held in special data stores that ensure secure erasure, e. g., by providing interfaces via which read and write operations can be performed at sector level. Such an encryption-based erasure procedure also fulfills the purging requirements of privacy laws [32].

Although it is not possible to delete data in blockchains due to the cryptographic hashes and the links between the blocks, a data purging by encryption approach can grant the right to rectification as well as the right to erasure, without thereby rendering the immutability and tamper-proofness as such obsolete. If personal data are stored in the blockchain in encrypted form, the blockchain can still guarantee immutability via the hashes. However, data can only be processed as long as the key required for decryption exists. If this key is stored externally in a trusted environment and the data subjects are given full control over their keys, they can make their data unreadable at any time. Also, not all data on the blockchain have to be encrypted, but only those that are considered personal data according to Article 4 of the GDPR.

For instance, the blockchain system *Hyperledger Fabric*[1] uses *CouchDB*[2] to represent the world state, i. e., the consolidated view of all nodes. Stach and Mitschang [37] have shown that secure deletion is feasible efficiently on such document-oriented databases via an encryption-based approach. Opposed to a regular database, blockchains are append-only, i. e., such an approach also results in less overhead since update operations involving multiple decryption and encryption operations are omitted. Thus, data purging by encryption can be considered an applicable technical solution for a blockchain, e. g., to implement the right to erasure.

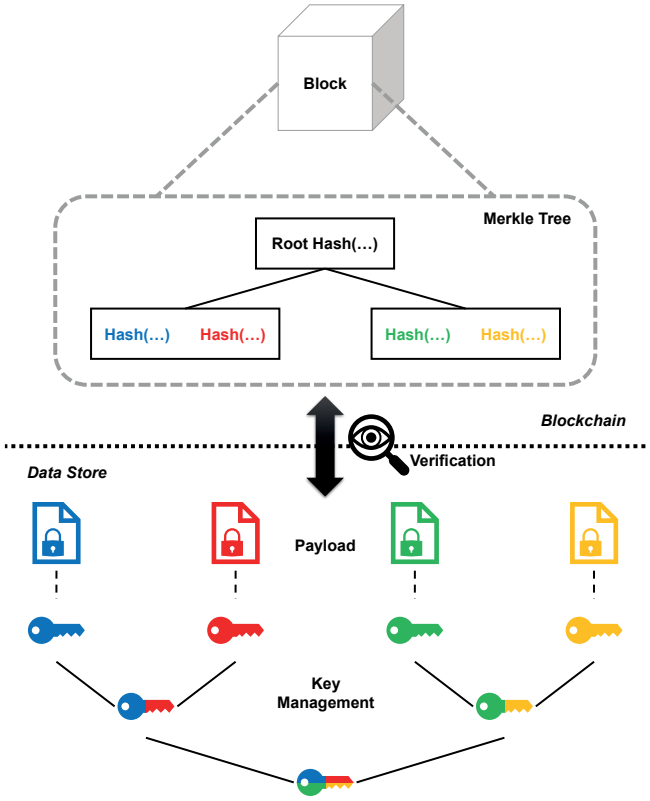

Figure 3: Data Purging by Encryption in a Blockchain.

Yet, with regard to the management of the keys, an effective strategy has to be adopted due to the large amount of data that can accumulate in a blockchain and thus the potentially large number of keys. Here, a structural property of many blockchains can be exploited. In blockchains, so-called *Merkle trees* are often used for data verification. This is a hash tree in which the leaves contain the hashes of the payload data, and the inner nodes contain a hash of its child nodes (see Figure 3 upper part). That is, a hierarchical structure is established, where each node is responsible for the consistency of all data contained in the subtree rooted at that node [19].

Waizenegger et al. [42] have introduced a tree-like data structure for managing keys. The keys with which payload data are encrypted are located at the leaf level. Each key is based on its parent node. In this way, all keys in a subtree become invalid if the key in its root node is deleted (see Figure 3 lower part). These two tree-structures can be mapped to each other, so that the required keys can be deleted very easily as soon as the node in whose subtree the data to be purged is located has been identified in the Merkle tree. This interrelation is outlined in Figure 3.

[1] see https://www.hyperledger.org/use/fabric (accessed on 17 December 2021)
[2] see http://couchdb.apache.org/ (accessed on 17 December 2021)

5.2 Attribute-Based Data Authentication

Data purging by encryption can also be used to correct data in a blockchain by deleting the incorrect data and then adding the corrected data to the data pool of the blockchain. However, this is a costly process. It is therefore much better to ensure the highest possible data quality in advance. One way to achieve this is to accept data from reliable sources, only. That is, an authentication of data sources is required.

Attribute-based authentication methods are suitable for this purpose as they can distinguish between different sources at a fine-grained level. These methods are based on a digital signature that contains certain attributes of the signer — i. e., the data source. As a result, the signature can not only be used to verify that the data has been transmitted genuinely, but it can also be used to determine unambiguously which properties the sender has. These properties are then checked against a policy. Only if they meet this policy the data are considered authentic [7].

Thereby, it is possible, e. g., to verify that the device that captured data about a data subject has the necessary software and hardware to capture this kind of data with a sufficient degree of accuracy. The attribute-based approach enables an arbitrary fine-grained distinction of entities, since the number of attributes used is not restricted. Nevertheless, it is an effective way to specify a policy to determine which requirements a source has to meet in order to provide a certain kind of data. It is only necessary to specify a threshold for the relevant attributes that a source must at least satisfy. All other attributes can be ignored when verifying the signature.

However, the virtually unlimited number of attributes that a signature can contain also harbors an inherent threat with regard to the privacy as some of the attributes might reveal too much information about the sender. This would represent a significant drawback if, in order to protect the privacy of one data subject, another data subject (in this case the sender) is exposed. Gritti et al. [11] therefore introduce a privacy-preserving attribute-based authentication. For this purpose, they use *delegated authentication*. That is, a trusted control authority acts as an intermediary between the source and the designated sink, i. e., in our case the blockchain. A source signs the data with its full signature and sends it to the control authority. The control instance filters out all attributes from the signature that are not required by the sink for authentication and applies the resulting *delegated signature* to the payload data. This reduced signature is still sufficient for the blockchain to verify the authenticity and origin of the data, but it does not reveal any privacy-critical information about the source.

5.3 Purpose-Based Permission Control

As discussed in Section 2, public blockchains are not suitable for storing sensitive information because anyone can join the network and thus gain unrestricted access to all data in the blockchain. This is not the case with private blockchains, since the number of parties with access to the data is severely restricted in such blockchains. Furthermore, smart contracts can be used to further regulate the processing of data by making it dependent on certain conditions.

Smart contracts, however, have to be hard-coded in *chaincode*. Therefore, they are comparable to the transformation operators defined in a data warehouse. There, the data are also automatically pre-processed according to predefined rules and optimized for certain use cases that are fully known in advance [15]. Yet, this implies that these use cases have to be identified and specified in advance. In dynamic environments, like today's smart environments, such a concept is too rigid. Data consumers require more flexibility, as new use cases are constantly emerging. The goal should therefore be to keep the data as generic and unprocessed as possible and to leave the processing entirely to the data consumers [14]. Therefore, it must also be possible to process the data outside of smart contracts.

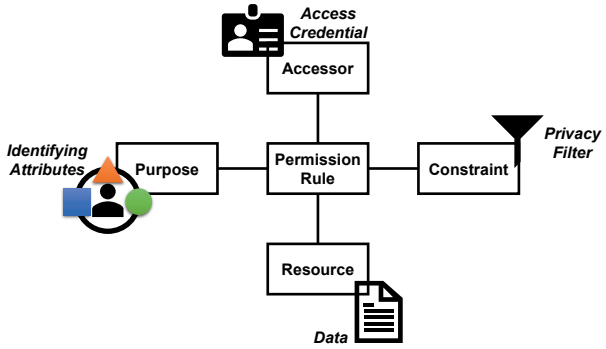

Figure 4: A Permission Model for Blockchain Applications.

However, this also entails that there have to be well-defined permissions as to which parties are allowed to access which data on the blockchain. In permissioned blockchains such as Hyperledger Fabric, this is regulated by means of *access control lists*. With these policies, it is not only possible to define who can participate in the network in general, but also which resources they are allowed to access. In addition, it is possible to restrict who is allowed to make updates — in terms of adding new data — to the blockchain. These access control lists rely on *role-based access control*. Yet, this rather traditional form of access control is often not dynamic and flexible enough, which is why Khan et al. [17] introduce *DistU*. DistU monitors the data of a blockchain permanently and grants or revokes permissions depending on how a data object is used.

Nevertheless, a data consumer still has either full access to a data object or none at all. That is, the permission model itself also needs to be extended in order to enable effective data minimization. Stach et al. [36] present such a *fine-grained permission model* for distributed Internet of Things applications that can be adapted to blockchains. Figure 4 shows this adapted permission model. A *permission rule* describes which *accessor* — i. e., which data consumer — may access which *resource* — i. e., which data. In the case of a permissioned blockchain, the data consumer can be identified using its access credentials. Since personal data may only be processed for a given *purpose*, such a purpose can be attached to a permission rule. Using an attribute-based authentication method as described in Section 5.2, the purpose can be specified by means of identifying attributes of the data consumer as well as the processing environment. Finally, *constraints* can be imposed on the processing. They are described in terms of privacy filters (see Section 5.4) that have to be applied to the data prior to processing.

In addition to the higher flexibility, as these permission rules do not have to be hard-coded as chaincode, and the possibility of

assigning very fine-grained permissions, this approach has another advantage over smart contracts. They are much easier to define as no coding skills are required. That is, they can also be comprehended and specified by IT laymen according to their privacy requirements.

5.4 Privacy Filters

By means of the permission rules, data access can be restricted quite well, but in order to be able to ensure data minimization effectively, it must also be possible to reduce the information contained in the data. Smart contracts could realize this, as they can transform the data and thus, e. g., filter out certain features during processing. However, the implementation of such a function is far too complex, so that data subjects are not capable of specifying such a smart contract to reduce the information content.

A more user-friendly solution is to provide out-of-the-box *privacy filters* that are able to blur certain privacy-relevant aspects in the data before releasing them to an accessor. However, it is important that the data are not rendered invalid in this process. Therefore, a collection of privacy filters adapted to specific data types and use cases is needed [2]. In addition to generic filters that can be applied to any type of data (e. g., withholding some data or adding noise to the data), also specialized filters are required for cloaking of location data [1] or distortion of time series data [8]. By applying the appropriate filter, it is possible to filter out certain aspects that are less relevant for processing but contain a lot of privacy-relevant information. In this way, information minimization can be achieved for any use case.

Besides such filters that operate on the data of a single user or even single data points, it is also feasible to use privacy filters tailored to large multi-user data stores like a blockchain. For instance, a privacy filter based on differential privacy enables statistical analyses without identifying individual users [46]. There are also filter operators that are designed to filter out large amounts of data without impairing the usability of the underlying data too much [25].

Stach et al. [34] present an architecture in which a set of different privacy filters is gathered in a repository and a suitable filter for the respective type of data and use case is selected. A utility metric is used to determine which filter provides the best privacy protection but at the same time has the least impact on the quality of the data for the particular use case. The code of the selected filter algorithm is then loaded into a data processor and applied to the data before releasing them to the data consumer.

Yet, this requires a trusted environment in which the filter operator is executed. A blockchain, however, is a trustless system, i. e., the individual parties cannot trust the other participants in the blockchain. Only by reaching a consensus among all participants for any operation, trust in the overall system is established. That is, for the application of privacy filters, a *Trusted Execution Environment* (*TEE*) is required in the blockchain system. A TEE is an isolated execution environment on which only approved applications can be executed. Cryptographic primitives ensure the integrity of the code executed in this environment and other processes have no influence on the execution as well as the outcomes [16]. In a TEE, it can therefore be ensured that the privacy filters cannot be manipulated and are executed correctly.

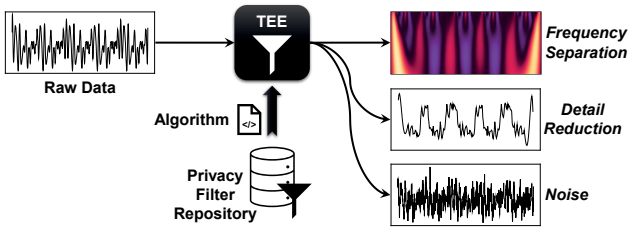

Figure 5: Application of Privacy Filters in a TEE.

Figure 5 shows how the privacy filters are applied. Depending on the requested type of data, applicable algorithms are selected from the privacy filter repository. For instance, for time series data, frequency separation can be used to abstract the data progression so that only changes in frequency are visible, the resolution of the data can be lowered to reduce details, or noise can be added to the data. Depending on the use case, the most suitable privacy filter is selected and applied to the raw data in the TEE. This way, both sides (data subject and data consumer) can trust that the privacy filter is applied correctly. For the data subject this means that the desired privacy level is maintained and for the data processor that the promised data quality is delivered.

5.5 Trusted Privacy Control Environment

The technical solutions shown in this section for ensuring data protection principles in a blockchain system, such as purpose limitation, data minimization, right to rectification, or right to erasure, however, also require an extension of the conceptual infrastructure of a blockchain so that they can be applied reliably. This is necessary whenever personal data are managed and processed by the blockchain, otherwise, as discussed in Section 3, a blockchain cannot be operated in a GDPR-compliant manner. Such sensitive data cannot be kept confidential in public blockchains for obvious reasons. Thus, if such data is involved, a closed set of participants, i. e., a private blockchain, can be assumed. Nevertheless, each node is considered trustless in its own right (as well as the data sources and data consumers) — otherwise, one would not need a blockchain.

Our approach is therefore to embed the trustless distributed components in a *trusted environment* that can be controlled by a central authority. In this way, our approach also meets the demand for a data controller. Stach et al. [35] present such a control environment for distributed Internet of Things applications. Figure 6 shows how we adapted this approach to a blockchain environment.

The blockchain is completely isolated from both, data sources and data consumers. If a source wants to add data to the blockchain (or rather its data pool), this must be done via an interface controlled by the trusted environment. Here, the attribute-based data authentication comes into play (see Section 5.2), which can verify whether the source has the necessary properties to be able to provide trustable data. If the verification is successful, personal data are encrypted before they are forwarded to the blockchain to enable data purging by encryption (see Section 5.1). The blockchain then processes the data autonomously and unaffected by the control environment. That is, its crucial key properties (decentralized, immutable, and tamper-proof) are not impaired by any means.

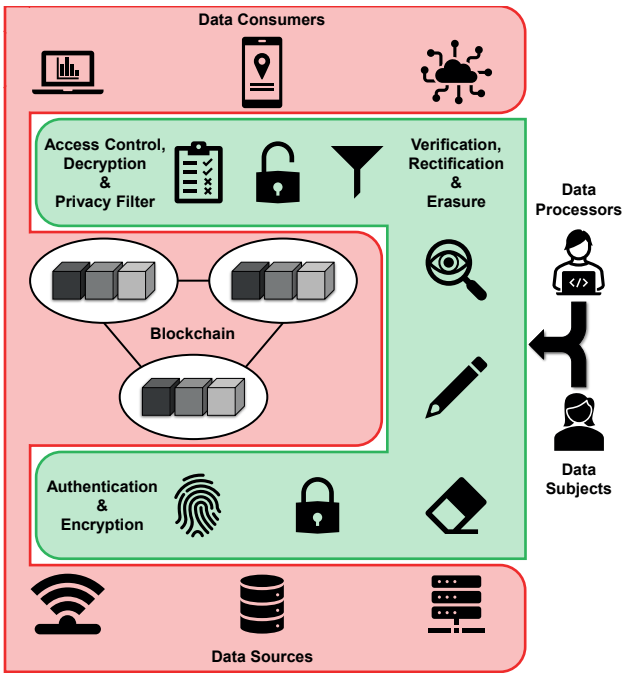

Figure 6: Embedding of a Blockchain in a Trusted Privacy Control Environment (trusted components are depicted in green while trustless components are depicted in red).

If data consumers want to gain access to the data stored on the blockchain, this is also done via a restricted interface. In this interface, the purpose-based permission rules (see Section 5.3) are checked to determine whether the consumer has the required access rights. If this is the case, personal data is decrypted, and privacy filters are applied to them according to the relevant permission rule (see Section 5.4). Yet, these privacy filters tamper with the data. Therefore, an additional verification interface is required for data processors to prove the authenticity of the underlying raw data (guaranteed by the blockchain itself) as well as the correct execution of the applied privacy filters (guaranteed by the TEE). This user interface also needs to enable data subjects to enforce rectification and erasure via the data purging techniques and express new or changed privacy requirements in terms of permission rules.

While all of this is feasible from a technical perspective, from an organizational perspective, however, it has to be resolved who has the responsibility for operating the trusted control environment. This operator has complete control over the data, as s/he controls which data are added to the blockchain and which data from the blockchain are made available to whom. Therefore, both, data subjects and data processors have to trust this controller implicitly. From our point of view, only the data protection officer of the organization which operates the private blockchain is eligible, as s/he is unbiased and trustworthy.

All Technical Solutions discussed in this paper, as well as their role in terms of a Privacy-Aware Blockchain, are outlined in Table 2.

6 FUTURE RESEARCH DIRECTIONS

As shown in the previous section, the five discussed technical solutions represent a major step towards a privacy-aware blockchain in the sense of the GDPR. By means of data purging by encryption, compliance with Article 16 and Article 17 — i. e., the right to rectification and the right to erasure — is achieved. Since rectification is realized by deleting the incorrect data and resubmitting

Table 2: Summary of the Discussed Technical Solutions and their Contribution towards Data Protection.

Technical Solution	Contribution towards Data Protection
Data Purging by Encryption	Due to the full encryption of all data in the blockchain, it is possible to delete data by deleting their respective decryption key. From a technical perspective, the data are still available, but they are no longer readable. This addresses all privacy issues that are related to the *revision or deletion of data*, e. g., Article 5(1)(d), Article 5(1)(e), Article 16, and Article 17.
Attribute-Based Data Authentication	By authenticating data sources and determining certain characteristics, inappropriate data sources can be easily identified. Their data can thus be excluded from the blockchain, as the expected data quality from such sources is low. This addresses all privacy issues that are related to the *quality and correctness of data*, e. g., Article 5(1)(d), Article 16, and Article 17.
Purpose-Based Permission Control	Fine-grained access control enables data subjects to specify who gets access to their data and for what purpose, without requiring a smart contract. This addresses all privacy issues that are related to the *processing of data*, e. g., Article 5(1)(b), Article 7, Article 18, and Article 22.
Privacy Filters	By applying privacy filters, data quality can be adjusted to reduce the amount of disclosed sensitive information. This addresses all privacy issues that are related to the *information value of data*, e. g., Article 5(1)(c), Article 5(1)(d), and Article 18.
Trusted Privacy Control Environment	By embedding these four techniques in a central control environment and isolating the blockchain from data sinks and data sources, a privacy-aware operation of the blockchain can be realized. This addresses all privacy issues that are related to the *management of data*, e. g., Article 12 – 15, Article 22, and Article 24, and enables *data protection by design* (Article 25).

the corrected data (which requires considerable effort due to the consensus protocol), attribute-based data authentication helps to ensure that the data sources are appropriate before they can add data to the blockchain in order to maintain high data accuracy (Article 5(1)(d)). Via the purpose-based permission control, data subjects are empowered to exercise their right to restriction of processing (Article 18), as they can specify in fine-grained manner which data are processed for which purpose (Article 5(1)(b)). The associated privacy filters achieve data minimization (Article 5(1)(c)), since only the information required for processing is passed on to a data consumer. The trusted privacy control environment, in which all of these concepts can be embedded, provides an additional *virtual* storage limitation (Article 5(1)(e)), since on the one hand the incoming data and on the other hand the visibility of the available data can be restricted. Furthermore, with this environment, data protection officers are enabled to take on the role of data controllers and thus represent a central point of contact for data subjects (Article 24). This environment also limits the power of smart contracts, as they can no longer be used for automated individual decision-making (Article 22), as their results initially remain completely isolated in the blockchain until they are approved by the data controller.

These technical solutions make blockchains compliant with the GDPR without having to sacrifice their security-relevant characteristics, such as immutability and tamper-proofness. However, often single data objects do not reveal highly privacy-critical insights about a data subject. Yet, if this data object is combined with other data, compromising information can be derived from it. Therefore, what is more important for data subjects than restricting access to individual data objects is to define what knowledge about them must not be disclosed [38]. However, to find (and thus protect) such complex knowledge patterns efficiently within a dataset, powerful query engines, index structures, and a modified data model for blockchains are required [28]. In the blockchain context, this especially concerns three types of queries:

Queries on Temporal Relationships. It is necessary to find temporal relationships in the data since privacy-critical knowledge is derived from a certain sequence of events. It must therefore be supported by a blockchain system to formulate queries that find data gathered immediately before or after a given date. Also querying data within a certain time window are required to this end.

Queries on Structural Relationships. As outlined in Section 2, new data initially lands in a data pool. When and in what order these data are added to a block cannot be predicted. So, the time at which data are captured can differ greatly from the time at which they are added to the blockchain. Therefore, queries regarding the structure of the blockchain — i. e., queries on data that are in specific blocks — also have to be supported, e. g., to determine at which point in time a certain piece of information was disclosed to all participants of the blockchain or as of when a rectification became visible.

Queries on Chronological Relationships. Due to the immutability of the data stored on a blockchain, it is not possible to execute an update on them. If a data object changes over time, this must be reflected as a new entry in the blockchain. To reduce the data size, such changes are often only described as compressed delta against the previous version. To obtain the data object and its history, all changes over time must be found in addition to the object itself. These data artifacts can be distributed over several blocks.

Efficient support for such complex queries is not available in current blockchain systems [29]. To us, this is the key research gap towards a comprehensive privacy-by-design blockchain (Article 25).

7 CONCLUSION

Whenever data have to be shared securely between several parties, the use of a blockchain is a suitable option. Blockchains ensure that the data are immutable, tamper-proof, and available to all participants in a transparent manner. Yet, it is due to these characteristics that they conflict with data privacy laws such as the GDPR.

To this end, we assessed in this paper, whether privacy-aware blockchains are feasible. (1) First, we identified with which articles of the GDPR there is a conflict. (2) Then, we presented five technical solutions that address these conflicts (namely data purging by encryption, attribute-based data authentication, purpose-based permission control, privacy filters, and a trusted privacy control environment) and described how they can be applied to a blockchain. (3) Finally, we discussed why more powerful query engines for blockchains have to be developed to facilitate a comprehensive privacy-by-design blockchain fully compliant with the GDPR. With queries on temporal relationships, queries on structural relationships, and queries on chronological relationships, we outlined three query types that require special attention in this regard.

This also provides an answer to our opening question whether blockchains and data privacy laws can be reconciled. It is possible, however technical and organizational adjustments are required and there is still a lot of research necessary to make these data protection measures in blockchain systems efficient and effective.

ACKNOWLEDGMENTS

Some work presented in this paper was performed in the project 'NUCLIDE' as part of the Software Campus program, which is funded by the German Federal Ministry of Education and Research (BMBF) under grant number 01IS17051.

REFERENCES

[1] Zahrah A. Almusaylim and NZ Jhanjhi. 2020. Comprehensive Review: Privacy Protection of User in Location-Aware Services of Mobile Cloud Computing. *Wireless Personal Communications* 111 (2020), 541–564. https://doi.org/10.1007/s11277-019-06872-3

[2] Sascha Alpers, Andreas Oberweis, Maria Pieper, Stefanie Betz, Andreas Fritsch, Gunther Schiefer, and Manuela Wagner. 2017. PRIVACY-AVARE: An approach to manage and distribute privacy settings. In *Proceedings of the 2017 3rd IEEE International Conference on Computer and Communications (ICCC '17)*. 1460–1468. https://doi.org/10.1109/CompComm.2017.8322784

[3] Iddo Bentov, Ariel Gabizon, and Alex Mizrahi. 2016. Cryptocurrencies Without Proof of Work. In *Proceedings of the 20th International Conference on Financial Cryptography and Data Security (Workshops) (BITCOIN '16)*. 142–157. https://doi.org/10.1007/978-3-662-53357-4_10

[4] David Berdik, Safa Otoum, Nikolas Schmidt, Dylan Porter, and Yaser Jararweh. 2021. A Survey on Blockchain for Information Systems Management and Security. *Information Processing & Management* 58, 1 (2021), 102397:1–102397:28. https://doi.org/10.1016/j.ipm.2020.102397

[5] Lelio Campanile, Mauro Iacono, Fiammetta Marulli, and Michele Mastroianni. 2021. Designing a GDPR compliant blockchain-based IoV distributed information tracking system. *Information Processing & Management* 58, 3 (2021), 102511:1–102511:23. https://doi.org/10.1016/j.ipm.2021.102511

[6] Miguel Castro and Barbara Liskov. 1999. Practical Byzantine Fault Tolerance. In *Proceedings of the Third Symposium on Operating Systems Design and Implementation (OSDI '99)*. 173–186. https://doi.org/10.5555/296806.296824

[7] Yu Chen, Jiguo Li, Chengdong Liu, Jinguang Han, Yichen Zhang, and Peng Yi. 2021. Efficient Attribute Based Server-Aided Verification Signature. *IEEE*

Transactions on Services Computing (Early Access) (2021), 1–9. https://doi.org/10.1109/TSC.2021.3096420

[8] Mi-Jung Choi, Hea-Suk Kim, and Yang-Sae Moon. 2012. Publishing Sensitive Time-Series Data under Preservation of Privacy and Distance Orders. *International Journal of Innovative Computing, Information and Control* 8, 5(B) (2012), 3619–3638.

[9] European Parliament and Council of the European Union. 2016. *Regulation on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (Data Protection Directive).* Legislative acts L119. Official Journal of the European Union. https://eur-lex.europa.eu/eli/reg/2016/679/oj

[10] Michèle Finck. 2019. *Blockchain and the General Data Protection Regulation: Can distributed ledgers be squared with European data protection law?* Study PE 634.445. European Parliamentary Research Service. https://www.europarl.europa.eu/thinktank/en/document/EPRS_STU(2019)634445

[11] Clémentine Gritti, Melek Önen, and Refik Molva. 2019. Privacy-Preserving Delegable Authentication in the Internet of Things. In *Proceedings of the 34th ACM/SIGAPP Symposium on Applied Computing (SAC '19)*. 861–869. https://doi.org/10.1145/3297280.3297365

[12] Akm Bahalul Haque, A. K. M. Najmul Islam, Sami Hyrynsalmi, Bilal Naqvi, and Kari Smolander. 2021. GDPR Compliant Blockchains – A Systematic Literature Review. *IEEE Access* 9 (2021), 50593–50606. https://doi.org/10.1109/ACCESS.2021.3069877

[13] Tharaka Mawanane Hewa, Yining Hu, Madhusanka Liyanage, Salil S. Kanhare, and Mika Ylianttila. 2021. Survey on Blockchain-Based Smart Contracts: Technical Aspects and Future Research. *IEEE Access* 9 (2021), 87643–87662. https://doi.org/10.1109/ACCESS.2021.3068178

[14] Bill Inmon. 2016. *Data Lake Architecture: Designing the Data Lake and avoiding the garbage dump.* Technics Publications, Basking Ridge, New Jersey, USA.

[15] William H. Inmon, Derek Strauss, and Genia Neushloss. 2008. *DW 2.0: The Architecture for the Next Generation of Data Warehousing.* Morgan Kaufmann Publishers Inc., Burlington, Massachusetts, USA.

[16] Patrick Jauernig, Ahmad-Reza Sadeghi, and Emmanuel Stapf. 2020. Trusted Execution Environments: Properties, Applications, and Challenges. *IEEE Security & Privacy* 18, 2 (2020), 56–60. https://doi.org/10.1109/MSEC.2019.2947124

[17] Muhammad Yasar Khan, Megat F. Zuhairi, Toqeer Ali, Turki Alghamdi, and Jose Antonio Marmolejo-Saucedo. 2020. An extended access control model for permissioned blockchain frameworks. *Wireless Networks* 26, 7 (2020), 4943–4954. https://doi.org/10.1007/s11276-019-01968-x

[18] Priya Kohli, Sachin Sharma, and Priya Matta. 2021. Security Challenges, Applications and Vehicular Authentication Methods in VANET for Smart Traffic Management. In *Proceedings of the 2021 2nd International Conference on Intelligent Engineering and Management (ICIEM '21)*. 327–332. https://doi.org/10.1109/ICIEM51511.2021.9445337

[19] Saravanan Krishnan, Valentina E. Balas, E. Golden Julie, Y. Harold Robinson, S. Balaji, and Raghvendra Kumar (Eds.). 2020. *Handbook of Research on Blockchain Technology.* Academic Press, London, San Diego, Cambridge, and Oxford.

[20] Maryna Manteghi. 2021. *Blockchain and the European Union's General Data Protection Regulation: From Conflict to "Peaceful" Coexistence?* White Paper. SSRN. https://doi.org/10.2139/ssrn.3805647

[21] Ken Miyachi and Tim K. Mackey. 2021. hOCBS: A privacy-preserving blockchain framework for healthcare data leveraging an on-chain and off-chain system design. *Information Processing & Management* 58, 3 (2021), 102535:1–102535:24. https://doi.org/10.1016/j.ipm.2021.102535

[22] Fernanda Molina, Gustavo Betarte, and Carlos Luna. 2021. Design principles for constructing GDPR-compliant blockchain solutions. In *Proceedings of the 2021 IEEE/ACM 4th International Workshop on Emerging Trends in Software Engineering for Blockchain (WETSEB '21)*. 1–8. https://doi.org/10.1109/WETSEB52558.2021.00008

[23] Muhammad Muzammal, Qiang Qu, and Bulat Nasrulin. 2019. Renovating blockchain with distributed databases: An open source system. *Future Generation Computer Systems* 90 (2019), 105–117. https://doi.org/10.1016/j.future.2018.07.042

[24] Satoshi Nakamoto. 2008. *Bitcoin: A Peer-to-Peer Electronic Cash System.* White Paper. Bitcoin Project. 1–9 pages. https://bitcoin.org/bitcoin.pdf

[25] Ripon Patgiri, Sabuzima Nayak, and Naresh Babu Muppalaneni. 2021. Is Bloom Filter a Bad Choice for Security and Privacy?. In *Proceedings of the 2021 International Conference on Information Networking (ICOIN '21)*. 648–653. https://doi.org/10.1109/ICOIN50884.2021.9333950

[26] Yanqing Peng, Min Du, Feifei Li, Raymond Cheng, and Dawn Song. 2020. FalconDB: Blockchain-Based Collaborative Database. In *Proceedings of the 2020 ACM SIGMOD International Conference on Management of Data (SIGMOD '20)*. 637–652. https://doi.org/10.1145/3318464.3380594

[27] Michelle Poelman and Sarfraz Iqbal. 2021. Investigating the Compliance of the GDPR: Processing Personal Data On A Blockchain. In *Proceedings of the 2021 IEEE 5th International Conference on Cryptography, Security and Privacy (CSP '21)*. 38–44. https://doi.org/10.1109/CSP51677.2021.9357590

[28] Dennis Przytarski. 2019. Using Triples as the Data Model for Blockchain Systems. In *Proceedings of the 18th International Semantic Web Conference (Workshops) (BlockSW '19)*. 1–2.

[29] Dennis Przytarski, Christoph Stach, Clémentine Gritti, and Bernhard Mitschang. 2022. Query Processing in Blockchain Systems: Current State and Future Challenges. *Future Internet* 14, 1 (2022), 1:1–1:31. https://doi.org/10.3390/fi14010001

[30] Marcelo Romero, Wided Guédria, Hervé Panetto, and Béatrix Barafort. 2020. Towards a Characterisation of Smart Systems: A Systematic Literature Review. *Computers in Industry* 120 (2020), 103224:1–103224:17. https://doi.org/10.1016/j.compind.2020.103224

[31] Sarwar Sayeed and Hector Marco-Gisbert. 2019. Assessing Blockchain Consensus and Security Mechanisms against the 51% Attack. *Applied Sciences* 9, 9 (2019), 1788:1–1788:17. https://doi.org/10.3390/app9091788

[32] Nick Scope, Alexander Rasin, James Wagner, Ben Lenard, and Karen Heart. 2021. Purging Data from Backups by Encryption. In *Proceedings of the 32nd International Conference on Database and Expert Systems Applications (DEXA '21)*. 245–258. https://doi.org/10.1007/978-3-030-86472-9_23

[33] Mohammed Shuaib, Shadab Alam, Mohammad Shabbir Alam, and Mohammad Shahnawaz Nasir. 2021. Compliance with HIPAA and GDPR in blockchain-based electronic health record. *Materials Today: Proceedings* (2021), 1–6. https://doi.org/10.1016/j.matpr.2021.03.059

[34] Christoph Stach, Julia Bräcker, Rebecca Eichler, Corinna Giebler, and Clémentine Gritti. 2020. How to Provide High-Utility Time Series Data in a Privacy-Aware Manner: A VAULT to Manage Time Series Data. *International Journal on Advances in Security* 13, 3 & 4 (2020), 88–108.

[35] Christoph Stach, Frank Dürr, Kai Mindermann, Saravana Murthy Palanisamy, and Stefan Wagner. 2018. How a Pattern-based Privacy System Contributes to Improve Context Recognition. In *Proceedings of the 2020 IEEE International Conference on Pervasive Computing and Communications (Workshops) (CoMoRea '18)*. 238–243. https://doi.org/10.1109/PERCOMW.2018.8480227

[36] Christoph Stach, Clémentine Gritti, and Bernhard Mitschang. 2020. Bringing Privacy Control Back to Citizens: DISPEL — A Distributed Privacy Management Platform for the Internet of Things. In *Proceedings of the 35th ACM/SIGAPP Symposium on Applied Computing (SAC '20)*. 1272–1279. https://doi.org/10.1145/3341105.3375754

[37] Christoph Stach and Bernhard Mitschang. 2018. CURATOR—A Secure Shared Object Store: Design, Implementation, and Evaluation of a Manageable, Secure, and Performant Data Exchange Mechanism for Smart Devices. In *Proceedings of the 33rd Annual ACM Symposium on Applied Computing (SAC '18)*. 533–540. https://doi.org/10.1145/3167132.3167190

[38] Christoph Stach and Frank Steimle. 2019. Recommender-based Privacy Requirements Elicitation – EPICUREAN: An Approach to Simplify Privacy Settings in IoT Applications with Respect to the GDPR. In *Proceedings of the 34th ACM/SIGAPP Symposium on Applied Computing (SAC '19)*. 1500–1507. https://doi.org/10.1145/3297280.3297432

[39] Noshina Tariq, Ayesha Qamar, Muhammad Asim, and Farrukh Aslam Khan. 2020. Blockchain and Smart Healthcare Security: A Survey. *Procedia Computer Science* 175 (2020), 615–620. https://doi.org/10.1016/j.procs.2020.07.089

[40] Unal Tatar, Yasir Gokce, and Brian Nussbaum. 2020. Law versus technology: Blockchain, GDPR, and tough tradeoffs. *Computer Law & Security Review* 38 (2020), 105454:1–105454:11. https://doi.org/10.1016/j.clsr.2020.105454

[41] Noor Thamer and Raaid Alubady. 2021. A Survey of Ransomware Attacks for Healthcare Systems: Risks, Challenges, Solutions and Opportunity of Research. In *Proceedings of the 2021 1st Babylon International Conference on Information Technology and Science (BICITS '21)*. 210–216. https://doi.org/10.1109/BICITS51482.2021.9509877

[42] Tim Waizenegger, Frank Wagner, and Cataldo Mega. 2017. SDOS: Using Trusted Platform Modules for Secure Cryptographic Deletion in the Swift Object Store. In *Proceedings of the 20th International Conference on Extending Database Technology (EDBT '17)*. 550–553. https://doi.org/10.5441/002/edbt.2017.67

[43] Kete Wang, Yong Yan, Shaoyong Guo, Xin Wei, and Sujie Shao. 2021. On-Chain and Off-Chain Collaborative Management System Based on Consortium Blockchain. In *Proceedings of the 7th International Conference on Artificial Intelligence and Security (ICAIS '21)*. 172–187. https://doi.org/10.1007/978-3-030-78618-2_14

[44] Gavin Wood. 2021. *Ethereum: A Secure Decentralised Generalised Transaction Ledger.* Ethereum Yellow Paper Berlin Version 888949c. Ethereum Project. 1–41 pages. https://ethereum.github.io/yellowpaper/paper.pdf

[45] Shijie Zhang and Jong-Hyouk Lee. 2020. Analysis of the main consensus protocols of blockchain. *ICT Express* 6, 2 (2020), 93–97. https://doi.org/10.1016/j.icte.2019.08.001

[46] Muhammad Talha Zia, Manzoor Ahmed Khan, and Hesham El-Sayed. 2020. Application of Differential Privacy Approach in Healthcare Data – A Case Study. In *Proceedings of the 2020 14th International Conference on Innovations in Information Technology (IIT '20)*. 35–39. https://doi.org/10.1109/IIT50501.2020.9299084