PriFoB: a Privacy-aware Fog-enhanced Blockchain-based system for Global Accreditation and Credential Verification

This paper was downloaded from TechRxiv (<https://www.techrxiv.org>).

LICENSE

CC BY-NC-SA 4.0

SUBMISSION DATE / POSTED DATE

13-01-2022 / 19-01-2022

CITATION

Baniata, Hamza; Kertesz, Attila (2022): PriFoB: a Privacy-aware Fog-enhanced Blockchain-based system for Global Accreditation and Credential Verification. TechRxiv. Preprint.
<https://doi.org/10.36227/techrxiv.18319322.v1>

DOI

[10.36227/techrxiv.18319322.v1](https://doi.org/10.36227/techrxiv.18319322.v1)

PriFoB: a Privacy-aware Fog-enhanced Blockchain-based system for Global Accreditation and Credential Verification

Hamza Baniata*, Attila Kertesz

*^aDepartment of Software Engineering, University of Szeged
Szeged 6720, Hungary*

Abstract

Trusted online credential management solutions are needed for instant and practical verification. Most of the available frameworks targeting this field violate the privacy of end-users or lack sufficient solutions in terms of security and Quality-of-Service (QoS). In this paper, we propose a Privacy-aware Fog-enhanced Blockchain-based online credential solution, namely PriFoB. Our proposed solution adopts a public permissioned Blockchain model with different reliable encryption schemes, standardized Zero-Knowledge-Proofs (ZKPs) and Digital Signatures (DS) within a Fog-Blockchain integrated framework, which is also GDPR compliant. We deploy both the Proof-of-Authority (PoA) and the Signatures-of-Work (SoW) consensus algorithms for efficient and secure handling of Verifiable Credentials (VCs) and global accreditation of VC issuers, respectively. Furthermore, we propose a novel three-dimensional DAG-based model of the Distributed Ledger (3DDL) and provide ready-to-deploy PriFoB implementation. We discuss insights regarding the utilization and the potential of PriFoB, and evaluate it in terms of security, privacy, latency, throughput and power utilization. We analyze its performance in different layers of Fog-enabled cloud architecture with simulation and emulation, and we show that PriFoB outperforms several Blockchain-based solutions utilizing Ethereum, Hyperledger Fabric, Hyperledger Besu and Hyperledger Indy platforms.

Keywords: Blockchain, Fog Computing, Privacy, Global Accreditation, Credential Verification, Self Sovereign Identity (SSI)

1. Introduction

Credential recognition is the process where a (inter)national body, called Verifier, validates the legitimacy of a document that was issued by another body, also called as Issuer. A credential is issued upon an event occurrence to certify that this event has indeed happened, such as educational credentials, vaccination certificates, governmental passports/IDs, etc. Within one country, area, or continent, one can find agreed-on regulations to recognize a named type of credentials for purposes like governmental treatment, hiring, travelling, etc. However, once a person/entity that was issued a legitimate credential needs to approve it abroad, usually a painfully lengthy and costly process needs to be carried out. This is because credential documents generally include different types of stamps, proofs, identification numbers and other data that have to be verified individually and carefully for each credential referring to distinct, centralized, locally maintained databases (DBs). That is, no standard is used by issuers around the world and no constant way to prove different credentials is guaranteed. A credential may even need to be approved by the representing body of the issuer's country at the foreign country, e.g. Embassy, where the representing body needs to communicate with the national recognition, and several governmental, bodies that are related to the management of the credential subject. The higher the

sensitivity of a credential, the more complicated and costly it is to validate it abroad.

Several previous studies looked for solutions to overcome the above mentioned drawbacks. The approach usually taken is to propose a central Trusted Third Party (TTP) that carries out international issuer recognition process [1], which makes it a viable reference for credential validation. The most recent example of such a solution is the EU Digital COVID Certificate ¹, where authorized governmental bodies within Europe updates a central DB that includes personal data on vaccination. Using this service, vaccinated people, or their agents, can prove that they have received a vaccine within a European country, allowing them to travel abroad without, e.g. quarantine, restrictions. However, private data of those agents in such a scheme is not only exposed to national, but also to international governmental personnel/systems. Additionally, locally vaccinated people need to individually register to several different platforms, before they can obtain an EU accredited vaccination certificate. Although such data management schemes are not compliant with the General Data Protection Regulations (GDPR) [2], it apparently was the only available approach to relax the pandemic's restrictions as soon as possible. In other, more sensitive cases, such as foreign educational diploma recognition or voting systems, privacy awareness is a critical factor that needs to be adopted by design in any proposed solution.

*Corresponding Author.

Email address: baniatah@inf.u-szeged.hu (Hamza Baniata)

¹<https://ec.europa.eu/info/live-work-travel-eu/coronavirus-response/safe-covid-19-vaccines-europeans/eu-digital-covid-certificate>

Fog Computing (FC) is an emerging trend for extending the Cloud Computing (CC) technology to address computing and networking bottlenecks in large scale deployment of CC-assisted systems [3]. The basic idea of FC is to create a layer of distributed fog entities between the centralized cloud data centers/processors and end-user devices (or Things in the case of Internet-of-Things (IoT) systems [4]) at the edge of the network. The FC layer should, conceptually, control the handling of users' data in a private and secure manner, while providing cloud services. Utilizing the fog was proven to be more efficient, than classical, centralized cloud-based services, in terms of overall system throughput [5], response latency [6], storage efficiency [7], and privacy [8].

A Blockchain (BC) system is a distributed computer system that is able to store and process a distributed ledger (DL) in a Peer-to-Peer (P2P) network model [9]. A BC can be permissioned, where participants added by authorized entities within the system are capable of appending (or mining) new blocks onto the DL. Permissionless BCs, on the other hand, allow anyone to participate as a miner in the system without the need for an access grant. Data saved on BC DLs can be either public or private, depending on the domain of users able to request or perform tasks within the system. That is, anyone can access a public BC (e.g. read data on the chain and/or request to be a miner), while only users within a predefined domain within the organization, region, etc. can access data on a private BC. Different BC functionalities appeared in the literature, mostly data management [10], payment and trading management [11], reputation management [12], and identity management [13]. In the scenarios of Smart Everything [14], the required automation criterion needs highly trusted and secure processing of end-users data. The GDPR enforced in Europe back in 2018, suggests wide variety of practices that should be used in order to protect the privacy of end-users. These regulations are hard to achieve within a BC-based system, yet doable, due to the immutability of data saved on-chain. Thus, system architects need to take precise and clear measures of what data is saved on/off chain, when BC-based solutions to be utilized.

We have previously surveyed the integration of BC and FC systems in [15], and found that already proposed approaches in the academia for integrating BC and FC majorly targeted IoT applications. The integration of these two technologies revealed the potential lying beyond, if collaboratively interacted. Regarding credential and identity management in such an innovative integration, projects that consider GDPR regulations mostly deployed centralized databases for handling private data, such as the cloud or a TTP server. Others, that save private data on the immutable DL, proposed some access control mechanisms, so that only authorized entities can read or add private data (which is still a non GDPR-compliant practice).

In this paper, we utilize a public-permissioned Blockchain and a Fog layer to propose an efficient system for global institution accreditation and credential verification, namely PriFoB. The BC in PriFoB acts as a Distributed Trusted Third Party (DTTP), in which miners are national accreditation bodies (e.g. national ministry of higher education, ministry of foreign affairs or ministry of health affairs, etc.). On the other hand, fog

nodes are realized by credential issuer bodies (e.g. universities, hospitals, vaccination centers, etc.). We designed PriFoB to guarantee privacy and security of system entities, by deploying the robust Proof-of-Authority (PoA) [16] Consensus Algorithm (CA), a realized application of the secure Signatures-of-Work (SoW) [17] scheme, RSA-based PKI-infrastructure and Digital Signatures (DS), and Zero Knowledge Proof (ZKP) mechanisms. The SoW deployment, specifically, is an additional and optional consensus sub-layer for practical realization of PriFoB as a global accreditation solution. We also propose a relaxed and efficient multi-dimensional BC model, where blocks are partially (not fully) immutable. Each dimension holds a different type of transactions (TXs), which enhances the overall efficiency of the validation process. Furthermore, we use an improved Directed Acyclic Graph (DAG) based block relations within each dimension, which outperforms the classical linear model in terms of total throughput and response latency [18].

The remaining part of the paper is structured as follows: Section 2 describes our proposed solution and highlights the main methods deployed within. Section 3 presents the evaluation results in terms of security, privacy, latency, throughput and power utilization. Section 4 presents the state-of-the-art regarding BC-based solutions proposed for solving similar problems we target. Finally, Section 5 concludes our work.

2. System Modelling

The main objective of our proposed system is the simultaneous provision of two major services: i) Institution Accreditation and ii) Credential Verification. The proposed system must be privacy-preserving by design, meaning that the deployed communications protocols and interaction/processing methods shall allow no window for private data leakage. The efficiency of the proposed solution can then be evaluated in terms of security, privacy, average throughput, storage cost and energy consumption (or processing cost). To address all of these, we propose Privacy-aware Fog-enhanced Blockchain-based Institution Accreditation and Credential Verification (PriFoB), which utilizes different technologies namely Fog Computing (FC), Blockchain (BC), Digital Signatures (DS) and Zero-Knowledge-Proof (ZKP) methods. Specifically, we utilize a public permissioned BC with PoA and a realized SoW CAs. Meanwhile, we use the robust SHA-256, AES and RSA encryption methods. We deploy a BC infrastructure to serve as a DTTP for a decentralized pool of (inter)national accreditation organizations. The general architecture of the proposed PriFoB, and a simplified control flow scheme within, are depicted in Figure 1. We made the source-code of the PriFoB solution, along with a tutorial on setup and deployment, publicly available at Github².

We believe it is more inclusive and comprehensive to describe a BC solution with reference to the BC layers described in Figure 2. Accordingly, we will separately discuss each layer of PriFoB design all the way up to the application layer.

²<https://github.com/sed-inf-u-szeged/PriFoB>

Figure 1: The general architecture and framework of entities in a PriFoB system. The circled numbers indicates the order of steps to remotely accredit an issuer, publish new schemes, issue a new VC by an accredited issuer and verify that VC

Figure 2: Foundation layers of Blockchain systems

2.1. Infrastructure modelling

Mainly, there are three elements of PriFoB: issuers, agents, and DTTP. An issuer is an end-user, with national/international accreditation, who is formally allowed to issue Verifiable Credentials (VCs) according to its business model. An agent is also an end-user, yet it represents a typical client of the issuer. That is, an agent can be issued a VC, but cannot issue one.

The communications between agents and issuers can be conceptually viewed as an end-user to fog communication. That is, an agent, who is granted a VC, can directly request a VC from a fog server (i.e. issuer server) closer to them, than from a web-based cloud application. Lastly, a DTTP is the element required to validate issued VCs with high efficiency. Noticing that PriFoB is mainly targeting global VC verification, it is recommended to have a distributed TTP to handle requests from anywhere. The advantage of such DTTP is demonstrated through task distribution leading to higher throughput compared to centralized TTPs. However, the DTTP should consist of trusted bodies representing national accreditation organizations. Examples of such include a ministry of higher education, or a governmental regulatory authority. End-users may be PCs, Mobile devices or even Microcomputers. A Gateway (GW) is then needed to load-balance the tasks received by the DTTP, and to control communications from/to the miners, which are the main components of the DTTP. Connecting miners into a P2P network, and the expected heavy computational and storage load on them, as the aim is having a global DTTP, suggest that miner nodes and the GW can not be handled by limited-resource devices like end-users'. Each miner node, along with the GW, can be administered by a human factor, or by an automation tool according to a regularly updated DB.

Our implementation of all PriFoB entities are containerized using Docker [19]. This allows for easy enhancement of the system in terms of scalability. That is, a clustering and/or scaling service can be utilized, using e.g. Kubernetes, which allows for container orchestration. For example, multiple GWs can be

utilized instead of one, and multiple containers can be utilized per miners. To clarify, a miner is attributed using its IP-address, location, and the administrative organization that owns it. Each miner can then be logically represented to the GW as one machine, yet several machines, with several local containers and addresses, can process the requests received by the cluster. Similar deployment options can be utilized for PriFoB issuers, i.e. the fog layer components, leading to a multi-level fog set-up [20]. For example, a department within a faculty is allowed to issue a limited number of schemes, or it is allowed to utilize specific schemes' private keys. Thus, no department should issue VCs for degrees that are not obtained within. On a higher level, a VC maybe revoked only by the faculty head office or the university head office, which implies that the private keys of the university should be managed, specifically, by the office allowed to use it.

2.2. Data layer modelling

In this subsection, we discuss how and when different types of data can be encrypted, decrypted, signed, and verified. Furthermore, we explain types of messages and ZKPs exchanged between different entities of PriFoB.

2.2.1. One-way, Symmetric and Asymmetric encryption

In cryptography, there are many types of encryption used to hide shared information. A hashing function $h(\cdot)$, or a one-way encryption function, is a mathematical function that takes a variable-length input string and converts it into a fixed-length binary sequence that is computationally difficult to invert [21]. A hashing function enables the determination of a message's integrity: any change to the message will, with a very high probability, result in a different message digest [22]. In PriFoB, we use the SHA-256 [23] function for one-way encryption.

Symmetric encryption is the process of turning a readable text (plain text P) into a non-understandable text (cipher C) using an encryption function $E(\cdot)$ [24]. The input of a symmetric encryption function is P or C , and a key k , leading to $E(P, k) = C$. The processes performed by $E(\cdot)$ shall be traversed, using k , if P to be derived from C . That is, $E^{-1}(C, k) = P$. In PriFoB, we use the AES methods [25] for the symmetric encryption.

Asymmetric encryption is a secure method $S(M, k)$ used to ensure that only the receiver is able to decrypt $D(C, g)$ a cipher and read the original message M . We deploy this type of encryption in PriFoB, in addition to one-way and symmetric encryption methods, so that the security of exchanged messages that include private data is guaranteed. This type of encryption implies the generation of two keys, g which decrypts C that was originally encrypted by k . A message M that was encrypted using k is computationally hard to be decrypted unless g is known. One of the most secure and famous asymmetric encryption algorithms is the RSA algorithm [26]. The RSA key generation process is described in Appendix A.1, while Appendix A.2 describes the encryption (steps 1 and 2) and decryption (steps 3 and 4) processes.

2.2.2. Digital signature and verification

The RSA keys generated above can be swapped without the loss of generality. That is, g may be used to decrypt a cipher C that was encrypted using k , or to encrypt a message M to verify the credibility of M 's origin. Specifically:

1. The original sender of M computes:
 $S(h(M), g) \rightarrow Sig$, where $h(\cdot)$ is an agreed on hashing function (e.g. SHA-256),
2. The sender sends the resulting Signature (Sig) along with M [M, Sig] to the receiver,
3. The receiver computes $h(M)$,
4. The receiver computes $D(Sig, k) \rightarrow h'(M)$,
5. if $h(M) = h'(M)$, the receiver shall be confident that M was sent by the original sender who is the only one that can read g .

Because this method is typically used to prove the sender credibility, while anyone can decrypt Sig using the publicly available k , it is called Signing and Verification rather than Encryption and Decryption. Following this remark, both M and Sig shall be encrypted at the sender side using symmetric encryption with a shared key, or asymmetric encryption with the public key k of the receiver. As described in the previous subsection, only the receiver then shall be able to *read* and *verify* the contents of M and Sig , respectively.

2.2.3. Zero-Knowledge-Proofs (ZKPs)

A ZKP is a verification technique which, using cryptography, allows one substance to prove to another component that it knows a specific data or fulfills a particular requirement without disclosing any actual data that supports that evidence [27]. We deploy the ZKP technique in PriFoB with the goal of end-users being able to verify VCs without disclosing any private data within. PriFoB implies that each VC is coupled with a Sig which is the encrypted hash of the issued VC $h(VC)$. Referring to the definitions presented in [28], BC miners and end-users in PriFoB can be considered verifiers and provers, respectively. The ZKPs in PriFoB fulfill all the properties of a successful ZKP deployment as follows:

1. **Completeness:** an honest prover can always convince an honest verifier. In PriFoB, No system entity is able to generate a *correct* Sig other than the original issuer of the VC because only the issuer knows the private key used for signing VCs.
2. **Proof of knowledge:** a malicious prover not knowing the secret cannot convince the verifier, except with negligible small probability. A prover in PriFoB sends the Sig along with $h'(VC)$. The verifier then decrypts Sig and obtains $h(VC)$. If $h(VC)$ is equal to $h'(VC)$, the verifier can be confident that the prover obtained this Sig from the original issuer and thus $h'(VC)$ is correct. Note that a prover needs to know both $h(VC)$ and the private key g in order to generate a correct Sig .
3. **Zero-knowledge:** an honest verifier that follows the protocol cannot learn additional information about the secret.

According to the previously presented definition of hashing functions, any alteration within the input of $h(\cdot)$ results a different output that can not be expected. Additionally, the input can not be known from the hash string (hence, the name one-way-encryption). By allowing the verifier to read $h(VC)$, $h'(VC)$, Sig and k the verifier shall not be able to read/deduce any private data that belongs to the prover nor to the issuer.

2.2.4. Exchanged messages

Next, we discuss different types of messages exchanged between PriFoB entities. Accordingly, we arrive at a conclusion about when to use the above described asymmetric encryption schemes. Knowing that using asymmetric encryption is computationally expensive and bounded, we only use it in PriFoB when necessary. As depicted in Figure 1, the first step for an issuer to be able to issue new VCs is to be accredited. To do that, the issuer generates a public key and sends it to the DTTP along with non-private issuer information (e.g. official name, IP-address, etc.). The corresponding private key is then saved and managed locally by the issuer. No encryption is needed in this step, other than ordinary symmetric encryption performed within the frame of the TCP/IP protocol. Similarly, publishing a new schema upon issuer accreditation includes sending non-private information (e.g. Schema title, public key, etc.). Once the issuer is accredited and its schemes have been published, it can generate new VCs to its clients. That is, the issuer admin issues the VCs with its clients' private data (e.g. name, grade, birth-info, etc.) and perhaps with its own shareable private data as well (e.g. courses, professors' names, admin and registrar's signatures, etc.). Thus, this VC is only saved locally, as we assume that, trivially, clients of an issuer do trust their issuers with their private data.

An agent (e.g. student or hospital patient, we also interchangeably use the term client) sends a VC request to the issuer, which also includes the client's public key. The request shall include some private identifying information (e.g. full name, SSN, year of credential issuing, etc.), so that the issuer can share a VC representing the original credential with high confidence that the requester is the client herself. Because of that, the client connects with the BC network to ask for the issuer's public key (which was initially saved on-chain in the first step of the protocol). Using this public key, the client can encrypt her credential request. Consequently, no entity but the issuer can read private data within a VC request, as only its private key can decipher the request. The issuer can then use the client's public key to encrypt its response. Mainly, the issuer response shall include two parts (assuming the requested credential was indeed issued) the digital credential, and Sig (using the schema's private key). The encrypted response, which is in fact the VC, can only be then read by the client, as only her private key can decipher it.

Once the response is decrypted, the client is free to share and verify the obtained VC. Regarding the verification process, the client may send a verification request to the DTTP. The verification request includes only non-private data, such as the issuer

name (to find its DID), the title of the corresponding schema (to find its public key), the hash of the digital credential and the accompanying Sig . Note that none of these data can reveal any private information about the client. Accordingly, asymmetric encryption is not required here. Once the BC receives the verification request, a miner, randomly selected by the GW according to the implemented load-balancing criteria, performs a VC verification as described in subsections 2.2.2 and 2.2.3. The output of this step defines the response from the DTTP to the client. That is, the response is either Valid or Invalid, yet the reason for considering a VC invalid shall be also provided. Reasons for considering a VC invalid include: the DID or Schema has not been published on-chain, the hash of the credential is not equivalent to the decrypted Sig , or the VC has been revoked by the issuer.

Similarly, the client can share her VC(s) with other agents after encrypting it using the agents' public keys. Note that an issuer is an agent as well, which means that issuers can also request VCs from other issuers within similar framework.

2.2.5. The Distributed Ledger

The BC as a PriFoB system element, including all its components, represents a DTTP for different types of agents to hold their public information and/or to securely validate issued VCs. On the other hand, many verifiers are required to perform tasks instead of a single central entity. Thus, it is recommended, in order to fulfil the system globalization feature, to have this DTTP designed as a BC. Additionally, those verifiers need to be granted equal provisioning and maintenance rights of the DL, as their roles in their territories are alike. As general criteria of a system that needs a BC include several equal participants, performing similar tasks, and maintaining a DL using an agreed-on CA, this description perfectly fits the scenario of the DTTP in PriFoB.

The DAG-based DL proposed in PriFoB has three dimensions, each dimension is used for a specific type of blocks. Simplified views of our proposed DL are depicted in Figure 3. In Figure 3a, the order of confirmed blocks, present at a given time slot within the DL, is demonstrated with reference to the depth of each dimension up to the genesis block. It is noticeable here that each child block is pointing to specifically one parent while each parent block is allowed to have several children blocks. Additionally, several blocks at a given dimension can have similar index.

In Figure 3b, mature blocks (will be discussed later) appearing in the DTTP are demonstrated with reference to the time they appear. It is noticeable here that DID blocks can either point to the genesis block or to other DID blocks. Schema blocks can either point to the DID block of, specifically, their issuer, or to other schema blocks generated earlier by their issuer. Revoke blocks can either point to the schema block that was used to issue the revoked VC, or to any other revoke block that was generated earlier by the same issuer using the same schema block. It is also noticeable here that the time at which a mature block appears does not necessarily define the index of the block nor the position at which it is placed within the DL.

Figure 3: A simplified view of our implemented DAG-based 3DDL (dashed links in (a) and arrows in (b) represent the usage of higher depth block of the signature of the linked lower depth block. Green node: Genesis Block, red nodes: DID blocks, blue nodes: schema blocks and orange nodes: revoke blocks)

Although both demonstrations within Figure 3 are captured from one miner's point of view, other honest miners within the DTTP will have exact similar views. The index and the previous signature of a mature block are decided by the miner who generated that block. In comparison with linear BC models, if a miner receives a valid block with a previous signature that is not of the, specifically, previous block, the new block is rejected. Additionally, the index of a received new valid block is determined by the receiver not the generator of the block.

According to the classification of DAG-based DLs, detailed in [29], each dimension in our proposed DL is of Type-II DAG, where TXs need to be organized in blocks for packaging and the topology is a natural graph. However, our proposed DL model allows for parent blocks being pointed to by several child blocks, while each child block is allowed to point to only one parent block. This results in a tree-like DL, with several blocks potentially having similar index but unique identifiers (i.e. official name of issuer, type of schema, and hash of revoked VC, respectively). To efficiently allocate a block, an orthogonal parameter (*identifier*, *index*) needs to be provided. The identifier of every newly advertised block is checked and repetitions are not allowed. In a rare temporary case where an orthogonal parameter is similar for two different blocks within the same dimension, the longest-chain extension method [30] is used leaving a DL with no orphaned blocks and no repetitions. This concludes that our proposed DL provides a deterministic finality [31]. The consensus mechanisms utilized to add new blocks will be detailed in later sections.

We denote by a Three-Dimensional Distributed Ledger (3DDL) the set of all directed acyclic block graphs $G = (V, E)$ with vertices V (blocks) and directed edges E , where each $B \in V$ has in addition a disordered $< B$ over all its outgoing edges. In our setup, an edge goes from a block to its parent, thus childless vertices (i.e. leaves of the tree-like DL) are those with no incoming edges. Graphs in 3DDL are required to have a unique maximal vertex, the genesis block. We further denote by sub (B, G) the subgraph that includes all blocks in G reach-

able from B .

The DTTP maintains a 3DDL, where each dimension is a DAG structured. The first dimension holds confirmed DID blocks, each of those blocks includes:

1. an IMMUTABLE Header consisting the block type and miner signature,
2. an IMMUTABLE Body used in miner signature generation, consisting:
 - the DID TX sent by the issuer,
 - signatures of active miners (each indicates whether this issuer is accredited or not by the signer),
 - and the signature in the Header of the previous block,
3. a MUTABLE independent chain of schemes (i.e. The second dimension), where future schema blocks issued by this specific issuer shall be saved.

Similar to the first dimension design, the second dimension holds a DAG of confirmed schema blocks each includes:

1. an IMMUTABLE Header,
2. an IMMUTABLE Body (with or without the accreditation signatures of all miners according to application specifications),
3. and a MUTABLE independent chain of revoked credentials (i.e. the third dimension), where future confirmed revoke blocks issued by the DID owner shall be saved.

The third dimension of the proposed 3DDL is dedicated to saving the hashes of revoked credentials. This would allow BC miners to check if the VC was revoked by the issuer, without actually reading any private data within the VC.

In all the three dimensions, the previous signature of a given block is not necessarily the signature in the header of the *last* confirmed block, yet the claimed previous signature must exist in one of the previously confirmed blocks for credibility. Our design allows for flexible data confirmation and non-linear

chaining, leading to higher block confirmation rates, higher throughput rates, and lower response latency.

To this end, there is no need to implement a time synchronization method between miners (which typically appears in PoA-based BCs) as there is no restrictions on the order of confirmed blocks. Several blocks can simultaneously appear in the network, instead of a single block per time slot, without a consistency problem as they are all considered valid by all miners who confirmed their claimed previous blocks. Sample DID, Schema and Revoke blocks, newly confirmed into the 3DDL, are presented in Appendix B.

New schemes can be appended later into the "schemes.chain" field of a given DID block, in which hashes of revoked VCs can be appended later into "Hashes_of_revoked_credentials" field of a given schema within the given DID block. Each newly confirmed schema shall consist of its own list of revoked VCs. Note that the immutability property of confirmed blocks is still preserved as changing a block, and all its consequent confirmed blocks for successful attack, requires the attacker to know all private keys of all miners who mined the consequent blocks. The detailed security evaluation will be discussed later.

For obtaining a maximal efficiency reading from the 3DDL, a locally-saved **Sorted.Chain** object is implemented. Once a block is added to the chain, a tuple including its unique identifier and its on-chain index (probably not unique) is added to this object. This way, when a TX is later received, for which a miner needs to search for a block by its unique identifier instead of its index, the miner extracts its index from the **Sorted.Chain** and then reads the block directly on the chain. We implemented a binary search algorithm for this purpose to enhance the search efficiency in PriFoB to reach $O(\log n)$. Note that the blocks are ordered on-chain according to their time of confirmation while they are sorted alphabetically in the Sorted.Chain using their unique identifiers. Similar approach was utilized in [32] and found most efficient.

We also decided to set the block size to exactly one TX per block, as each TX requires its own group of signatures even if multiple TXs per block are utilized. As increasing block sizes lead to proportional decrease in block generation rate, due to higher verification latency [33], the overall throughput shall not improve with changing the block size as proven in [34].

2.3. Network Layer modelling

In PriFoB, the BC network is a typical interconnected P2P network, with a randomly selected group of neighbors per peer. Each peer represents a national accreditation body, which is usually the focal point for accreditation in a given country. Peers can either be actual computers at the facility or cloud-based instances that can be accessed through the internet. In either cases, each peer shall connect to the PriFoB GW to get updated data regarding other active miners, the current state of the DL, and to receive and respond to the submitted TXs by agents and other miners.

The GW in PriFoB is a load-balancing and a monitoring element for the DTTP network in general. The GW keeps track

on active miners, accept new mining requests and broadcasts miners' identification information (including their public keys) into the BC network. Shared data among neighbors are sent directly, while responses from miners to agents are sent through the GW.

To realize the connections between different PriFoB components, we utilize the open Secure Sockets Layer (OpenSSL) framework. Each entity dedicates two threads, namely Server and Client, to asynchronously receive and send messages, respectively. These threads deploy the Transport Layer Security (TLS) Protocol [35] with a static port number (i.e. 5050) for sending and receiving.

Communications between system entities are ruled by a strict protocol for the sake of achieving the Privacy-by-Design method. All communications with the GW of the BC, and between Miners, are performed in plain text because none of these communications include private data of any system entity. However, Miners can only communicate with the GW and with their neighbors, while the GW can communicate with all of them. All exchanged messages are JSON strings for the ease of future deployment of heterogeneous system entities.

Agents can communicate directly with each other or with the GW. DID, schema, and revoke requests sent by issuers to the GW need not to be encrypted as they hold no private data. Similarly, verification requests sent by agents to the GW need not to be encrypted as they only include the hash of the credential and the signature of the issuer. Only communications between agents, described as step no. 4 in Figure 1, must be encrypted. The implemented protocol requires a VC requester agent, who intends to request a VC from another agent, to encrypt its request using the public key of the requested agent (be it another client or an issuer). This way, only the requested agent can read the VC request. The public keys of issuers are publicly available at the DTTP side, thus, it is requested in plain text by the end-user. Once the encrypted VC request (which includes the public key of the requester) arrives at the requested agent side, the requested agent decrypts it using its private key. If the requested agent is the issuer of the VC, it validates the request according to private end-user data provided within the request (e.g. student no., health insurance no., full name, SSN, etc.), and searches for a previously issued VC at its local DB that matches the requester data. This process is what makes an issuer behave as a fog node because fog nodes are, fundamentally, entities that communicate directly with end-users/edge devices, and perform computations instead of a classical centralized server (i.e. mostly a cloud server). Once the VC is found, it is encrypted using the public key of the requester and sent back. Accordingly, only the requester can read the VC, which includes private data. If end-user 'a' would like to view a VC of end-user 'b', 'a' must send a share request including their names/identification, IP-address and their public key to 'b'. 'b' then responds to this request by accepting or rejecting it. If accepted, 'b' selects a VC that is locally saved, encrypts it using the public key of 'a', and directly sends it to 'a'. User 'a' then is the only entity who can read the VC as it holds the private key. Through our CLI-based implementation, end-users have the option to request, validate, delete, rename or share VCs, and they

Table 1: Computational complexities required to generate new blocks referring to different types of requests, and the expected appearance rates of different types of TXs throughout the life-cycle of an accredited institution

Request type	appearance	computational complexity (wrt. no. Miners)
DID TX	Few	$(n+1)$ signing + $(2n-1)$ verification
Schema TX	Moderate	1 signing + n verification
VC validation	Most	1 verification
Revoke TX	Rare	1 signing + $(2n - 1)$ verification

are able to download part or full of the publicly available DL.

2.4. Consensus modelling

In PriFoB, we utilize two types of CAs for two different layers of consensus regarding new blocks. By default, we assume that a successful verification of a new schema TX or a new revoke TX means a successful verification of the TX issuer (among other things to validate). Thus, those two types of TXs are processed by only one miner leading to mining a new block using the signature of this miner (i.e. Proof of Authority). The last type of TXs that could be received by the BC network is the DID TX type. Here, we decided to propagate the TX throughout the BC network asking each miner, which represents a national accreditation authority, to declare whether the DID requester is accredited or not in its country. The decision can be performed both manually by the miner admin, or automatically referring to a local/remote DB (both approaches are implemented).

Once a miner’s declaration is ready, the miner’s signature using its private key is added to the declaration and the signed declaration is added to the TX. The TX is propagated throughout the network, in the state of unready/immature TX, until the following strict conditions are met. When a miner receives a DID TX that meets all these conditions, the TX is considered ready or mature to be mined:

1. All active miners have signed the TX (periodically pinged by the GW, e.g. every 5 seconds)
2. Recipient miner has already signed the TX
3. All signatures are correct
4. TX is not found in any previously confirmed DID block (i.e. the claimed DID identifier is unique)

Unlike classical PoA CA, no restrictions are enforced in PriFoB regarding the number of blocks that can be generated within a single time slot. The miner who finds a mature DID TX mines and broadcasts it immediately.

All miners are authorized to mine new blocks, according to the data layer definitions, in any time slot by signing the block body, and adding the signature to the Header of this new block. Note that the accumulated and complete group of correct miners’ signatures represents a trigger for the last miner to start the mining process. In other words, the trigger to mine a new DID block in PriFoB is pulled by the network (instead of a single leader node in typical PoA-based BCs). Consequently, the miner adds a PoA to the new block and propagates it throughout the network. The PoA is then verified, as well as the block

body, by all recipient miners, and is added upon successful validation and verification.

Table 1 presents the computational complexities expected, for each type of TXs, from the time it is delivered to the GW, until it is confirmed and responded to. Note that the expected relative appearance of different TX types can be deduced logically. That is, each issuer is allowed only one unique DID, while it is expected to issue unlimited number of VCs referring to a limited number of schemes. Additionally, it is very rare that an institution would revoke a VC it issued. This being said, an expiration data of a VC can be injected within by the issuer, and consequently the accompanying signature, so that it can be checked once the VC is found valid. A similar approach can be used for DIDs and schemes as well.

Compared to different consensus complexities presented in [36], PriFoB performs optimally for schema and revoke TXs as both require $O(1)$ computational complexity by the system (i.e. signing). For DID blocks, PriFoB requires $O(n)$ computational complexity by the system, which is better than the PBFT and RBFT CAs with complexities $O(n^2)$ and $O(n^3)$, respectively.

From the values presented in the table, it can be observed that the verification is used much more for all types of TXs than signing. It can also be observed that the most appearing type of TXs is the VC validation request which requires no signing by any BC entity as it is not saved on-chain. This, in fact, was our initial motivation to adopt the RSA encryption methods. For more details, we present our thorough comparison between the famous ECC and our adopted RSA encryption methods in Section 4.

In PriFoB, VC issuers (e.g. institutions in the case of deployment for educational diploma verification) that want to start issuing credentials to their clients (e.g. students) need to first obtain accreditation from the DTTP. The accreditation is represented by the acceptance of the DID of this issuer, and saving it on-chain (the first dimension). Once the DID is published, the institution may generate new schemes, each scheme is correspondent to one type of credentials. A schema should at least include a public key used to generate the signatures of the corresponding type of credentials. This schema will also be saved on-chain (the second dimension). Figure 4 presents a simplified sequence of communications and interactions for publishing a DID and a schema.

Depending on the saved on-chain data, clients shall be able to verify the correctness of their received VCs. In PriFoB, a confirmed public key of a schema, published by an issuer, is the reference the DTTP uses to confirm a VC validity. Once a schema is saved on-chain, no further action from the issuer side

Figure 4: PriFoB Control Flow for publishing DIDs and schemas. Blue and Light Salmon colored areas indicate, where the SoW and the PoA consensus algorithms are utilized, respectively.

is required towards the DTTP. However, the issuer can generate as many VCs as it needs, as long as those VCs are signed using the private key of the schema (saved and managed locally at the issuer side). If a new type of VCs to be issued, a new schema needs to be submitted and, only after saved on-chain, the new type of VCs can be issued. Note that old schemas remain saved on-chain as the BC provides immutable storage, thus old fashioned VCs remain verifiable despite a new schema application. This is both beneficial and critical. It is beneficial for old scholars who are guaranteed they will not lose their credibility even if the issuer's system is changed. However, it is critical if, for some reason, the issuer decided that some of its previously issued VCs should be considered invalid. In our proposed PriFoB system, we solve these issues by utilizing our proposed 3DDL as discussed in subsection 2.2.5.

2.5. Credential Verification

Both end-users and issuers can request the DTTP to validate a VC through our implementation. The validation request sent

to the DTTP mainly includes four pieces of data: the DID block identifier and index (e.g. issuer official name and its index on-chain), the schema block identifier and index (which might be similar for different issuers but unique for each issuer), the hash of the credential to be validated, and the signature originally provided by the issuer within the VC.

Once a miner receives this request (randomly selected by the GW), it searches for the DID with the claimed identifier and index, if found, it searches within the schemas chain within this DID block for the schema identifier and index. Once found, it searches within the revoke chain within this schema block for the hash of the credential provided within the request. If not found, then the signature is verified using the public key of the schema found. if the DID/schema was not found, a response is sent to the requester that the issuer/the schema is not registered. If the signature is valid, and the hash is not in the revoke chain, a response is sent to the requester that the credential is valid. if the signature is valid but the hash is in the revoke chain, a

response is sent to the requester that the credential is revoked.

Note that no private data are saved on-chain, or provided for validation. This is the ZKPs scheme we use as the VC is validated without any knowledge requirement. The requester need not to expose any private data other than its address, to which the response should be send. The VC validation is performed automatically and publicly without any restrictions or conditions. If an agent decides to download the publicly available DL, it can verify any VC without referring to the DTTP in PriFoB. Figure 5 presents a simplified sequence of events for issuing and revoking a VC in PriFoB.

2.6. Application layer modelling

In an abstract description of the ZKP-based verification protocol of PriFoB, a prover can be either a client of an issuer or the issuer entity itself. An issuer entity can then issue a valid VC for itself. However, this should not be a trust problem since the application of PriFoB does not give a meaning of such behaviour. For example, consider a hospital that issues vaccination certificates to patients, or an educational institution that issues degree certificates to graduate students. Both of these entities do not need to be certified as vaccinated or graduate, respectively. On the other hand, an issuer shall be able to validate a VC before/after it is issued, in order to deliver guaranteed services to clients. Thus, issuers are able to perform ZKP-based verification just as their clients.

Note as well that issuers can be clients to each other, which is indeed a realistic scenario. That is, organisations (e.g. hospitals, education institutions, etc.) shall be able to be granted certificates from other, higher-level organizations (e.g. international committee of specialized hospitals, international corporation of education quality, etc.).

The PriFoB framework provides a ready to deploy BC-based VC verification solutions for a wide variety of applications. That is, DID blocks only provide a global identifier of issuers, while Schema blocks are much flexible for any type of VCs. Issuer accreditation is also optional at the time of deployment, and if it does not apply for the business model, then it can be simply deactivated by the miners. Accordingly, any received DID TX will be verified only depending on the correctness of the signatures within. In addition to the educational diploma and the global vaccination certificates examples, PriFoB can be used for realizing other applications such as a BC-based smart parking solution [37]. Depending on the business model, we can envision a scenario where several international companies need to co-operate their parking lots that are distributed in different countries. Each parking lot is maintained by a local management team and needs to digitally verify cars entering their areas. The international companies then can be connected to perform a DTTP, which accepts DID TXs from the parking lots. Each miner in the DTTP can individually accredit (or not) a parking lot according to the agreement between the companies.

Once a parking lot has a confirmed DID, it can issue its own different schemes relating to different sections/floors or different types of vehicles (e.g. small cars, trucks, etc.). Once a schema is confirmed, each vehicle is granted a VC certifying that it is allowed to enter and park at the lot which, specifically,

owns this DID, in which a specific section is allowed to be utilized. Such a VC can be validated by security checkers at the gates of the parking lots and/or the sections to be entered. The security checker might be a human factor with a simple end-user account or an automated verification machine using, e.g. a QR-code scanner.

The mentioned scenario is one of many several applications where PriFoB could be used for global co-operation. PriFoB, in short, provides the highly required flexibility for distributing a global server, while maintaining a unified interface for end-users at different locations, regions and countries. Other examples, including (inter)national e-Voting, e-Health, IoV applications, can be easily realized using our open-source PriFoB solution.

It is also beneficial to highlight the award policy in PriFoB, as all system entities do perform computational and storage tasks throughout the life cycle of a VC. As a start, miners are not awarded for newly mined blocks, but rather granted a fare share of revenues according to their provided computational and storage capacities. Every TX submitted to the BC can be accompanied with a proof of pre-payment, either using cryptocurrency or bank-based accounts. The amount of valid pre-payment can be adjusted during system setup according to the application business model (e.g. type of TX, expected cost of a block confirmation, etc.). A verified TX can then be checked by the recipient miner for valid pre-payment TX on another platform. End-users can use the services of PriFoB (i.e. the PriFoB wallet interface) freely as revenue can be obtained by adds on the application UI. However, issuers can individually enforce a pre-payment scheme for issuing a VC depending on their own business model as well.

For example, a valid DID TX can be defined to cost \$10, a valid schema TX can be defined to cost \$5 and a valid revoke TX can be defined to cost \$50. Accordingly, if 10k issuers are expected to register to the PriFoB-based system, with an average of 20 schemes per issuer, then an initial revenue of \$100k is expected for DID blocks plus \$1m is expected for schema blocks. If each schema is expected to be used for issuing 100 VCs per year, then the number of yearly active end-user wallets can be estimated by $100 \text{ users} \times 20 \text{ schemes} \times 10\text{k issuers} = 20 \text{ million active wallets per year}$. Assuming that only 1% of those users open their wallets per day (i.e. 200k users), and an average revenue rate of \$11.5 per 1,000 views [38] as of 2021, the expected daily revenue can be estimated then by \$2300 (i.e. \$839,500 per year). Assuming the DTTP is composed of 300 miners, each is providing equal computational and storage capacities, each of those miners can be granted $\approx \$2798$ per year. As a basic 24/7 virtual machine at a public cloud provider platform (e.g. E-2 instance at the Google Cloud Platform) may cost $\approx \$350$ per year [39], then each miner is left with a yearly revenue of more than \$2440. Note that in the case where the DTTP miners have agreements/collaboration with issuers, such as described in parking lots case above, the revenue for issuing the 20 million VCs can also be divided amongst the collaboration entities. That is, assuming that a VC is issued for \$10 with an annual expiration date, then a total yearly revenue of \$200 million is expected. Similar revenue calculations, can be conducted

Figure 5: PriFoB Control Flow for issuing a new Verifiable Credential and revoke a previously issued one. Light Salmon colored area indicates where the PoA consensus algorithm is utilized.

with different TX predefined costs for different applications.

3. Evaluation

In this section, we analyze the security in the Data Layer and the Consensus Layer of PriFoB. The security discussion of the remaining layers are beyond the scope of this work. We also discuss the privacy in PriFoB in its generality referring to the GDPR rules and the detailed description of requests and responses among system entities. After that, we compare real measures of PriFoB deployed in the cloud against well documented implementations of Ethereum and different Hyperledger platforms, including Indy, Besu and Fabric. Finally, we parameterize a Discrete Event Simulation environment, using the real data we obtained, to predict the throughput and power utilization of PriFoB in different sittings.

3.1. Security

Distributed systems are usually evaluated in terms of security referring to several concepts. The **Strong Validity** requirement implies that if all honest nodes propose the same value v , then no honest node decides a value different from v . The value v here corresponds to the state of the DL. In PriFoB, this property is guaranteed as all honest nodes that are connected to the

network will eventually receive and accept a new block, leading to its addition to their local BC. Thus, all honest nodes will eventually agree on the contents of the DL.

The **Agreement** requirement implies that no two honest nodes decide differently. This is also referred to as the consistency of the DL [40]. This property is usually the hardest to achieve in BC-based system due to the need of agreeing not only on valid blocks, but also on the order of those blocks. PoW-based systems solve this by hardening the puzzle difficulty so that only one block is generated every 10 minutes, giving sufficient time to new blocks to propagate throughout the network. However, PriFoB does not suffer in this regard as the order maintenance requirement is relaxed because we utilize our previously described DAG-based 3DDL. This being said, PriFoB is not suitable for applications that require strict chronological order of blocks, such as digital currency applications, but is suitable for credential verification frameworks as described earlier.

The **Termination** and the **Integrity** requirements imply that every honest node eventually decides some value and that no honest node decides twice, respectively. Both of these properties are true for PriFoB as we fine-tuned all system elements to either accept or reject new TXs/blocks, according to strictly defined conditions. Furthermore, if an immature TX is signed then a miner does not sign it again.

Originally, a generic SoW scheme for consensus in distributed systems was described in [17]. The generic scheme was proven correct, secure against Tampering and Chosen-Message Attacks, and t-Verifiable (the verifier is able to verify a SoW in t steps, where t is a lot smaller than the time needed to produce a signature). The generic SoW scheme has been shown feasible for deployment in even less trusted environment, namely permissionless DLs. In such settings, miners are allowed to join and leave the network without restrictions.

Specifically, for n miners collaborating to produce a group of signatures that represents a proof of correctness for a given block, h is the hardness level of signature generation, H is a collision resistant hash function, λ is a security parameter and T is the number of failed/adversary nodes ($T < \frac{n}{2}$), the lower bound on the rate of generating uniquely successful blocks by an adversary, denoted γ , for such system can be calculated using Equation 1.

$$\gamma_{DID} = (n - T) \cdot (1 - (\frac{h}{2^\lambda})^{T_H}) \cdot (\frac{h}{2^\lambda})^{(n-1)T_H} \quad (1)$$

For example, a SoW-based system, such as ours, that consists of $n = 6$ nodes, with $h = 50$ milliseconds, H is SHA-256, $\lambda = 1024$ bits and $T = 2$, generates each DID block with its unique SoW with $\gamma = 6.66 \times 10^{-1533}$. Note how small this value is, indicating nearly Zero probability of an adversary miner to be able to change any DID blocks saved on the DL. To clarify, differential cryptanalysis cannot be used by a dishonest miner to deduce other miners' private keys and generate forged blocks.

For other types of TXs, i.e. schema and revoke TXs, that consist of one signature that proves its correctness, the probability of forging is determined by λ , h and T . We can deduce Equation 2 from Equation 1 to assess the probability of a successful schema or revoke block forging. The simplest way to describe it is to assume that the longer the key (in bits) the stronger it is. More details can be found in several references such as [41] and [17].

$$\gamma_{scheme, revoke} = (1 - (\frac{h}{2^\lambda})^{T_H}) \cdot (\frac{h}{2^\lambda})^{T_H} \quad (2)$$

We further assess PriFoB using the Common Vulnerability Scoring System (CVSS v3.1) [42]. The CVSS v1.0 was originally implemented and proposed in 2006 and had been continuously developed to v2.0 in 2007, v3.0 in 2015 and finally v3.1 in 2019 [43]. CVSS is a tool used to quantify the fuzzy security of a given system into a numeric score. The CVSS scores given for a system range between 0.0 to 10.0. The lower the score of a given system, the more secure it is against security attacks. To obtain an accurate security score of a system, the CVSS tool needs to be parameterized according to the system design and the probable attacks characteristics.

To obtain a security score for our proposed PriFoB, we parameterized the CVSS tool as follows:

- Attack Vector: Local (attacker must be connected to the DTTP network)
- Attack Complexity: High (attacker needs high computational power to forge a block)

- Privileges Required: High (attacker must be accepted as a miner by the GW)
- User Interaction: Required (other miners must accept the attacker as a legitimate miner and start using its public key distributed by the GW)
- Scope: Changed (forged blocks with correct signatures leads to changing the state of the DL)
- Confidentiality: None (a forged block does not imply a confidentiality vulnerability)
- Integrity: Low (a forged block does imply an integrity vulnerability but only for one customer not for the whole system)
- Availability: None (a forged block does not imply an availability vulnerability)

The above parameterization resulted in a rather low CVSS **Base Score** of 2.3/10.0. Hardening the assessment, PriFoB received another low **Environmental Score** of 3.9/10.0 as follows:

- Modified Attack Vector: Network (attacker can join the DTTP network from anywhere)
- Modified Privileges Required: low

Regarding the security of our proposed 3DDL, such model implies higher efficiency and throughput compared to linear models as proven in several previous works. However, the drawback of utilizing it is demonstrated by higher probability of launching a selfish mining breach leading to double-spend attacks [44] (e.g. payment and smart contract solutions such as Bitcoin and Ethereum). Such type of attacks indeed makes sense for applications that require a chronological order of confirmed TXs/blocks. However, those attacks do not make sense within the framework of PriFoB [45]. That is, the simultaneous confirmation of multiple blocks/TXs, leading to different order of blocks at different physical locations of the DL (termed forks), does not imply that one of the blocks in the tie/fork is incorrect or misleading. Each of the blocks/TXs in a tie of a fork shall represent a different issuer with a different unique identifier although they may have similar block indices. To solve a probable mistaken response for a given block index that is common for multiple blocks a requested miner searches for all the blocks with the given index and requires an equal value of the unique identifier in the request to the value of the unique identifier in the block.

3.2. Privacy

As PriFoB complies with the W3C standard and further does not allow for any private data to be saved on the DL, PriFoB can be trivially considered a GDPR-compliant system. We have described earlier how VCs are only saved locally at the end-user layer and are never sent to the DTTP. Additionally, the deployed PKI schemes, the ZKPs with collision resistant hash

Table 2: PriFoB comparison with previous related works, that proposed solutions for distributed credential verification systems, in terms of utilized Blockchain platform, granting institution accreditation services, number of Miners (M), assisted request type (T), lower and upper bounds of request per second rates (req/s) and the lower and upper bounds of response Latency measured as second per request (s/req)

Solution	CA	w/Accreditation?	M	T	req/s	Latency (s/req)
Hyperledger Indy [46]	Plenum (PBFT)	NO	4	DID/Schema (write)	1–250	2–6
			8	Any (read)		0.08–1.6
Hyperledger Besu [47]	PoA	NO	4	DID/Schema (write)	10–100	2–10
				Any (read)		0.1–2.5
Hyperledger Fabric [48]	RAFT	NO	2	Any (write)	50–250	0.6–0.8
Hyperledger Fabric [47]		NO	4	Any (write)		0.7–0.95
Ethereum [47]	PoW	NO	2	Any (read)	10–100	5.03–5.58
Ethereum [49]	PoA	NO	2	Any (write)		0.02–0.06
Ethereum [50]	PoW	YES	N/A	Any (read)	25–100	5–34
Ethereum [51]	PoW	YES	N/A	Any (read)		0.2–0.4
	PoW	YES	Main Net	DID (write)	1–100	47–114
	PoW	YES	Main Net	DID (write)	N/A	5–40
PriFoB	SoW+PoA	YES	2–6	DID (write) Schema (write) Revoke (write) Any (read)	1–250	0.013–1.09 0.006–0.6 0.005–0.09 0.003–0.14

functions, along with the utilization of digital signatures, all lead to a Privacy-by-Design implementation.

Nevertheless, we referred to the GDPR checklist for data controllers available at [52]. In our paper, we have conducted an information audit to determine what information PriFoB processes and who has access to it. Thus, we provided clear information about PriFoB data processing and legal justification in its privacy policy (PriFoB-based solutions will have a legal justification for its data processing activities as end-users voluntarily sign up into the system). Encryption, pseudonymization, and anonymization of personal data wherever possible is performed in PriFoB. Additionally, it is easy for PriFoB customers to request and receive all the information it has about them. That is, the whole blockchain is directly downloadable via the application interface. Finally, different types of TXs can be revoked and regenerated, while private data is only saved locally, making the private data controllable only by its owners (there is no need to request private data deletion).

3.3. Latency

Searching recently proposed solutions, we found the following works that are similar to PriFoB in terms of objectives and BC-deployment. Feher [46] analyzed the latency of BC-based solutions utilizing Hyperledger Indy. Urbančok [47] described and compared four open-source BC platforms, namely Ethereum, Hyperledger’s Fabric, Besu, and Iroha, in different terms including latency. Xu et al. [48] analyzed the latency of BC-based solutions utilizing Hyperledger Fabric. Zhang et al. [49] experimentally evaluated the performance of Ethereum

testnets in terms of Account balance query latency, Block generation time and End-to-end TX acceptance latency. Härer and Fill [50] utilized the main net of Ethereum platform and proposed a credential verification solution, on which they performed latency assessment. Bampatsikos et al. [51] proposed a solution for mitigating probable Computational Denial of Service (CDoS) attacks when utilizing Ethereum for credential verification purposes, and provided latency assessment for their solution. We clarify the differences and similarities, along with brief comparison of the results, for those works and PriFoB in Table 2.

Some of these works evaluated their solutions using simulation (all miner nodes run on a single machine), while others evaluated their solutions using real test-beds (each miner is allocated a different machine). We compare the latency of PriFoB, however, with all of them, to prove PriFoB outperformance and provide insights regarding expected latency with different scalability measures. Additionally, we run several test scenarios in order to comprehensively compare PriFoB with all of them. The results of the scenarios we tested, along with each scenario parameters, are detailed in Figure 6. To facilitate reading and comparing the performance results, we color-coded the tables’ cells according to the scale provided within the figure.

We evaluated the performance of PriFoB using a Proof of Concept (PoC) system composed of a single GW, a network of miners with different sizes (2, 4 and 6 miners), and a script we implemented, that emulates several issuers and agents simultaneously communicating with the DTTP. We deploy each of those entities on a separate VM at the Google Cloud Platform, each is of type E2-medium (2 Intel(R) Xeon(R) vCPUs clocked

Solution	no. Miners	request / second								
		1	10	20	25	50	100	150	200	250
Indy	4	0.08				0.3	0.6	0.9	1.1	1.6
	8	0.1				0.4	0.5	0.8	1.1	2.5
Besu	4		0.04	0.1	0.3	0.4	0.56			
Fabric	2		0.6	0.65		0.7	0.8			
Ethereum	2		0.02	0.03		0.06	0.06			
Ethereum	N/A				0.2	0.3	0.4			
PriFoB (read)	2	0.003	0.006	0.009	0.009	0.015	0.02	0.03	0.05	0.05
	4	0.003	0.008	0.01	0.012	0.02	0.03	0.06	0.08	0.14
	6	0.004	0.006	0.016	0.02	0.036	0.057	0.08	0.1	0.13

Solution	no. Miners	request / second								
		1	10	20	25	50	100	150	200	250
Indy	4	2.7				2.5	2.2	4.5	5.4	6.4
	8	2.6				2.6	2.5	4.9	5.6	10.4
Besu	4		3.34	3.37		4.6	4.32			
Fabric	2					0.65	0.7	0.72	0.77	0.8
Fabric	4					0.7	0.8	0.82	0.85	0.9
Ethereum	2		5.03	5.04		5.58	5.07			
Ethereum	N/A	7			12	15	20			
Ethereum	main net	47					114			
Ethereum	main net	5		25		35	40			
PriFoB DID	2	0.013	0.032	0.03	0.08	0.15	0.3	0.5	0.75	0.9
	4	0.02	0.04	0.037	0.08	0.17	0.4	0.5	0.72	0.89
	6	0.046	0.06	0.12	0.14	0.26	0.47	0.71	0.98	1.09
PriFoB schema	2	0.006	0.02	0.03	0.05	0.1	0.2	0.3	0.45	0.6
	4	0.007	0.02	0.04	0.04	0.09	0.18	0.28	0.38	0.46
	6	0.008	0.016	0.04	0.054	0.095	0.19	0.26	0.36	0.45
PriFoB revoke	2	0.005	0.01	0.01	0.01	0.02	0.03	0.04	0.06	0.09
	4	0.005	0.009	0.01	0.012	0.02	0.05	0.06	0.07	0.08
	6	0.006	0.016	0.016	0.013	0.02	0.03	0.05	0.07	0.08

Figure 6: Color-coded average response latency for PriFoB verification requests per second (upper table), and average response latency for PriFoB DID, Schema and Revoke write requests per second (lower table) against average read and write latency measurements, reported in the literature, for major Blockchain solutions utilized for objectives similar to PriFoB

at 2.30GHz with 10GB of RAM), running Linux 20.10 OS. Alternatively, issuer and agent implementations are available and tested within the project repository but they can not be used to stress-test the DTTP in PriFoB.

According to the aforementioned description of PriFoB, the GW component is the bottleneck of a PriFoB-based system. Thus, we tested PriFoB using one GW to obtain the worst results possible. The bottleneck effect can trivially justify the observable proportional relation, in Figure 6, between the number of miners and the average latency. However, it will be shown in the following subsection that the Read latency in PriFoB should not be affected by the number of miners when deploying a computationally powerful GW. As a result, our containerized implementation of PriFoB elements, including the GW, shall show better measurements than those presented here, if more powerful machines and/or more GWs were deployed using e.g. Kubernetes orchestrator. It is worth noting here that despite the apparent bottleneck effect in our experiments, PriFoB still outperforms all the compared related systems.

3.4. Throughput

Next, we mathematically model our system to predict its general behaviour while the rate of requests per second (λ) and the number of miners (n) increase. To do so, we refer to the Queuing Theory [53] and characterize PriFoB using the Kendall's notation as $(\lambda=M/D=M/n):(FCFS/\infty/\infty)$, where M, D and FCFS stands for Poisson distribution, output distribution and the First-Come-First-Served discipline, respectively. Utilizing such approach to model BC-based systems has been predominantly used in the literature [54]. To compute the expected average processing time (W_s) in PriFoB, we use Equation 3:

$$W_s = \frac{\mu(\lambda/\mu)^n}{(n-1)!(n\mu - \lambda)^2} P_0 + \frac{1}{\mu} \quad (3)$$

where μ and P_0 are the mean service rate per busy server (request per second i.e. $\mu = 1/\text{Latency}$) and the probability that there are Zero requests in the system, respectively. For all $n \times \mu > \lambda$, we compute P_0 using Equation 4:

$$P_0 = \frac{1}{[\sum_{i=0}^{n-1} \frac{1}{i!} (\frac{\lambda}{\mu})^i] + \frac{1}{n!} (\frac{\lambda}{\mu})^n \frac{n\mu}{n\mu - \lambda}} \quad (4)$$

All the following experiments were run using the Discrete Event Simulation tool provided at [55], with suitable modifications and tuning referring to our model and real parameterization. The tool provides analysis for a single queuing case per run so we injected several FOR loops as needed to go through all the cases in each scenario we attempted to test. Using our mathematical model of PriFoB, we measured W_s for $\lambda \in [1, 250]$ with a static $n = 10$ and using $\mu = 335$ referring to the real READ results we detailed previously. Accordingly, we could simulate the general effect of increasing λ and compare it with our real data to validate the simulation tool. The results we obtained for the first scenario are depicted in Figure 7. As can be noticed, the simulation results comply with the real measures we obtained as increasing λ results in linear increment of W_s . However, average W_s is less than the real data presented earlier, as expected, due to the deployment of more mining nodes. This indeed shall increase the overall throughput of the system as more servers are able to process an equivalent λ .

To capture the general effect on W_s when increasing n , we tested $\lambda = 10$, $n \in [1-100]$ and using $\mu = 335$. The results we obtained are depicted in Figure 8. Here, the simulation assumes that all incoming requests are immediately distributed among available miners (i.e. no bottleneck effect). Accordingly, the latency is nearly constant between 3–5 millisecond per request. This proves that the proportional relation between n and the Read latency is mainly attributed to the limited computational power of the GW. Uncontrollable communication delays and/or unpredicted hit ratios when searching the BC can also cause such effect. For all of these reasons, we used a static μ in both scenarios we simulated so far, as μ is not related to n for all types of TXs, except for DID TXs.

For DID TXs, however, μ is affected by changing n as more miners in the system implies more signature generations. Consequently, increasing n should, theoretically, increase W_s for DID TXs. For this reason, we could state that the results depicted in Figure 8 do not represent the expected behaviour of PriFoB for DID TXs.

To capture the general effect of increasing n on W_s for DID TXs, we estimate the values of μ for increased n using the latency observations we obtained in the previous subsection. Let $X_{i,l}$ be the real DID TX latency for $i \in Y = [2, 4, 6]$, $l \in Z = [1-250]$. We can then calculate the increase in latency per each added miner, denoted by $\gamma_{i,j,l}$, using Equation 5:

$$\gamma_{i,j,l} = \frac{X_{i,l} - X_{j,l}}{j - i} \quad \forall i, j \in Y, l \in Z \text{ and } i < j \quad (5)$$

Using data in the set Γ , which consists of all γ values we obtained, we calculated $\gamma_{avg} = 0.03$ second per request per added miner and we calculated the Maximum Likelihood Estimation (MLE) $\gamma_{MLE} = 0.01$ second per request per added miner. The MLE was obtained from a Poisson Distribution obtained from the set Γ . Note that MLE is generally more accurate than the average [56]. We injected both γ_{avg} and γ_{MLE} into Equation 3 resulting Equation 6:

$$W_{sDID} = \frac{\mu(\lambda/\mu)^n}{(n-1)!(n\mu-\lambda)^2} P_0 + \frac{1}{\mu} + (n \times \gamma) \quad (6)$$

Using Equation 6 to predict W_{sDID} values for different n values is now realistic using both of γ values. We tested $\lambda = 10$ DID TXs, $n \in [1, 100]$. The results we obtained are depicted in Figure 9.

3.5. Power Utilization

Finally, deploying more miners may indeed increase the latency for DIDs, a type of TXs that is used less frequently than Schema TXs, and further appears even more rarely compared to VC validation requests where PriFoB system flourish. However, we need to further evaluate PriFoB in terms of power utilization. To achieve this, we injected several variables into our mathematical model implementation. We could then assess the percentage of system power utilization for known λ and n . We tested for $\lambda = 250$, $n \in [1, 100]$, a dynamic μ for DID TXs with $\gamma_{MLE} = 0.01$, and a static $\mu = 335$ for all other types of TXs. Figure 10 describes the percentage of computational power consumption out of the total available computational power, to process all received requests. For all types of TXs other than DID TXs, it is clear from the results we obtained that adding more miners to the system shall enhance the overall efficiency in terms of computational power consumption per request. However, increasing n implies higher power consumption rates due to the required SoW computational complexity. Note that if the optional accreditation service in PriFoB is deactivated, DID TXs shall consume as much energy as any other type of TXs. Also note that the power utilization here is directly proportional with μ which is rather low in our experiments, due to the low computational capacity of our test-bed infrastructure.

4. Related work

Several previous works approached BC-based solutions for realizing digital credential verification. Fauteux et al. [57] stated how BC deployment can solve the long list of challenges when an academic credential needs to be validated. The BC solution was discussed along its pros and challenges, and some technical details were further presented. Only in the years of 2019 and 2020, some few testable implementations were proposed to address credential verification problem using BC technology. Recently launched BC-based credential verification projects, namely BlockCerts [58], OpenCerts [59], and trustED [60] were surveyed in [61], where the most mature and suitable consensus methods, BC architectures, and BC platforms were presented and discussed. Consequently, the authors proposed AcaChain, which is a private, permissioned BC system that allows issuers to track the achievements of their agents, and then issue the graduation proof once all conditions of the credential are fulfilled. As the three BC-based credential systems, namely AcaChain, BlockCerts and OpenCerts, save all credentials information, along with relevant personal identifiers of students on the immutable chain, they are non GDPR-compliant. Furthermore, saving all data on the chain is considered an inefficient approach of using a BC system, as this can rapidly drain

Figure 7: Processing time per READ request in PriFoB for $\lambda \in [1, 250]$, $n = 10$ and $\mu = 335$

Figure 8: Processing time per READ request in PriFoB for $\lambda = 10$, $n \in [1, 100]$ and $\mu = 335$

Figure 9: Processing time per DID request in PriFoB for $\lambda = 10$ DID TXs, $n \in [1, 100]$, $\gamma_{avg} = 0.03$ and $\gamma_{MLE} = 0.01$

Figure 10: Power utilization with $\lambda = 250$, $n \in [1, 100]$ for DID TXs ($\gamma_{MLE} = 0.01$) and all other types of TXs ($\mu = 335$)

storage and computation resources. Thus, these projects may not be considered practical in industrial deployment. Similar approach was utilized by other solutions as well, e.g. [62].

Anant et al. [63] proposed deploying BC for SRM Institute of Science and Technology digital credential validation, using Ethereum smart contracts. This approach also saves all students credentials and data on the public chain. Similarly, Ahammad MS Tomal MH [64] proposed a BC-based system that saves only the signed hash of each certificate instead of the certificate itself, while deployed the inefficient Proof-of-Work (PoW) CA to maintain the consistency of the DL. In [65], a BC-based privacy preserving protocol is proposed so that users can access on-chain services. The proposed protocol suggests that a verifier is needed to personally verify the correctness of a claimed credential. Furthermore, in the generality of this protocol, the assumption of the user having to be a member of the blockchain raises some questions regarding deployment feasibility.

Hyperledger Indy [66] is an open source project, administered by the Linux foundation, which aims at providing a BC-based Verifiable Credential (VC) system. The project implementation provides a platform for Decentralized Identifiers (DID) rooted on BCs or other DLs so that they are interoperable across administrative domains, applications, etc. The project deploys privacy preserving mechanisms such as ZKPs, and takes good care on what data is saved on-chain. Indy uses the Plenum CA which is a special purpose RBFT CA [67]. The Sovrin BC [68] is built on top of Indy project for providing a general purpose, global DIDs. After years of development, this BC was officially launched in September 2019. Similarly, The GraphChain [69] was proposed for exploiting the advantages of Indy and facilitate the issuance of Legal Entity Identifiers.

Tariq et al. [70] proposed Cerberus, where Ethereum-based, private-permissioned BC architecture was utilized. Here, VCs and (if applicable) revoke TXs are saved on-chain. A QR code-scanning approach is deployed then to verify a VC is indeed on-chain and that it is not revoked. The network in Cerberus is monitored and maintained by a single accreditation body that refers to its local DB for a full list of issuers. Thus, an issuer is automatically accredited by the system only if it was accredited through a legacy accreditation channel. Although the architecture seemed very promising, the authors have only discussed their proposal theoretically while no implementation or experimental comparisons were provided.

In [45], a strategy for modelling, designing and developing BC-based healthcare accreditation and verification solutions was discussed. Although the paper aimed at e-Health applications, it provided rich insights on design principles, trade-offs, and critical terminology definitions. Similarly to most proposals for BC-based accreditation and verification solutions, this proposal was discussed only in theory.

In November 2019, the World Wide Web Consortium³ (W3C) published a standard recommendation for solutions targeting VC solutions, including data models, system concepts, approved methods, privacy and security challenges with pro-

posed solutions, and validation. EBSI⁴ and BCDiploma⁵ are examples of services built according to this standard with the BC as the TTP where system entities save their data. However, the standard recommendation of W3C⁶ and the solutions following it did not consider the issuer accreditation challenge assuming any entity should be able to own a DID and issue any type of VCs.

Comparing EBSI with our proposed solution, EBSI utilizes the general purpose Hyperledger Fabric⁷ which uses the Raft CA. VC issuers in EBSI apply to an accreditation body for accreditation outside EBSI scope. Once an issuer is accredited in a legacy approach, the accreditation body issues a VC on the EBSI network certifying that this issuer is eligible to issue credentials. Miners in EBSI refer to bodies who attain accreditation confirmation. The BCDiploma platform directly saves all issued VC hashes on the Blockchain which means that they can never be deleted, and it uses a PoW-based BC with linear DL model. Additionally, issuer accreditation service is not provided.

It is worth noting that different asymmetric encryption schemes are available in the literature. Many previous works deployed the most famous alternative scheme of the RSA namely Elliptic Curve Cryptography (ECC) [71]. For this we performed a brief comparison between RSA and ECC to decide on which is more suitable to deploy in PriFoB. Table 3 presents a brief comparison between the two schemes referring to different benchmarks. To clarify, ECC keys are shorter than RSA keys for an equivalent level of security. Thus, ECC scheme is considered more secure than RSA comparing their keys of the same length. ECC is relatively fast in key generation and cipher decryption, while RSA is relatively fast in plain-text encryption and signature verification. As will be discussed later, we found that PriFoB entities mostly utilize the asymmetric encryption framework for signature verification, while signing is occasionally utilized. We also found that an RSA key with length of 1024 bits is currently considered post-quantum and secure, while such length of a key performs fast enough to outperform state-of-the-art solutions. For these reasons, we decided to use the RSA encryption framework with modifiable 1024-bit keys.

Table 3: Comparison between RSA and ECC encryption schemes

Benchmark	ECC	RSA	ref.
Post-Quantum?	yes	temporarily yes	[72]
Security level	better	worse	[72]
Key Generation Time	better	worse	[73]
Encryption Time	worse	better	[41]
Decryption Time	better	worse	[41]
Signing Time	better	worse	[73]
Verification Time	worse	better	[73]

⁴<https://ec.europa.eu/cefdigital/wiki/display/CEFDIGITAL/EBSI>

⁵<https://www.bcdiploma.com/en-GB>

⁶<https://www.w3.org/TR/vc-data-model/#introduction>

⁷<https://www.hyperledger.org/use/fabric>

³<https://www.w3.org/>

in the late 2020, Wang et al. [29] have comprehensively investigated the state-of-the-art regarding DAG-based DL systems. They provided a general mathematical model of such systems, and categorized existing structures into six types. They could then systematically define potential applications and drawbacks of those structures as they were found either rough in summaries, superficial in analysis, or incomplete in evaluations. The main performance bottleneck of linear DL structure, which is not the case in DAG-based ones, was identified to be the utilized consensus mechanisms. Specifically, the competition among a group of miners for the right of block packaging does not appear in DAG-based DLs as each newly added block is allowed to refer to more than one parent (many-to-many cardinality model of the blockchain). To solve this issue in linear DLs, block confirmation must be artificially suppressed (e.g. adjust the puzzle difficulty in the consensus method) so that each block is fully attached before the next one's arrival (resulting in one-to-one cardinality model of the blockchain). DAG-based DLs, on the other hand, support concurrent operations as multiple nodes can simultaneously add TXs/blocks to the DL, thereby significantly improving the throughput.

Type-II DAGs, specifically, have been utilized in several previous works including Spectre [74], Phantom [75] and Meshcash [76]. However, Spectre uses it temporarily for swift processing but the final DL is linearized by a majority voting. Phantom utilizes a PoW consensus and enforces a strict linear ordering over blocks and transactions in the network, resulting in a DL with probabilistic finality. Meshcash utilizes a PoW consensus as well, with blocks pointing for necessity to *every* block confirmed in the previous round. None of those solutions offer a Tree-like multi-dimensional DL as PriFoB does. Specifically, all previous works adopt a many-to-many cardinality model; each newly added block may refer to more than one parent, and each parent may have several children. Our 3DDL, on the other hand, adopts a one-to-many cardinality model; each newly added block refers to specifically one parent, while each parent may have several children. To the best of our knowledge, our 3DDL is the first to propose a permanent, multi-dimensional and PoA- and SoW-based DL with deterministic finality.

Although our implementation is similar to few several recently proposed credential verification solutions (e.g. the Indy and Sovrin BCs) in terms of purpose and the BC access control modelling, we argue that our implementation is more appropriate. In PriFoB, accreditation is performed remotely and directly operated by the PriFoB network/platform, which reduces the whole process complexity. This is because miners in PriFoB are administered by the accreditation bodies themselves. Additionally, PriFoB has the following properties:

1. The BC model in PriFoB is Multi-Dimensional, which implies enhanced structure of the classical linear BC model, which makes it harder to breach. Additionally, it is more efficient to search the ledger for a schema during e.g. credential validation phase.
2. PriFoB provides both services, namely global accredita-

tion and credential validation, with GDPR- and W3C-compliance.

3. We deploy two consensus layers, including SoW and PoA, to realize both of these services. Accordingly, all nodes, rather than a majority of nodes in Plenum, should validate and sign, and later confirm a DID block.
4. The maximum number of miner nodes in systems deploying versions of the RAFT CA is recommended to be 25 nodes at maximum. This observation is inline with the findings presented in [34], where Hyperledger projects were reported to be limited by a maximum of 16 nodes. On the other hand, PriFoB can dynamically scale up and down according to the necessary number of miner nodes for an application, with more efficient latency and throughput as will be shown later.

5. Conclusion

In this paper, we have proposed PriFoB: a Privacy-aware Fog-enhanced Blockchain-based solution for global accreditation and digital credential verification. We used Zero-Knowledge-Proofs, Digital Signatures, SHA-256, AES and RSA encryption schemes, and two different consensus algorithms, namely PoA and SoW. Furthermore, we have proposed a novel Three-Dimensional DAG-based Distributed Ledger (3DDL) with efficient deterministic finality. We evaluated PriFoB in terms of security, privacy, latency, throughput and power utilization. Additionally, we compared our realized cloud-based deployment of PriFoB with similar blockchain-based solutions, for various system parameterization. PriFoB outperformed all of the recently proposed solutions utilizing Ethereum and Hyperledger projects (Besu, Fabric and Indy). We provided a ready-to-deploy implementation of PriFoB, and we made it available at a public, open-source repository.

Our future directions include the investigation of deploying new techniques into PriFoB, such as Sharding and Merkle Trees with light nodes. Specifically, we will investigate available opportunities for increasing the block size in PriFoB, as it is currently set to 1 TX/B. In the future, we also plan to provide a web-based implementation of PriFoB entities and a wallet application for mobile devices. To optimize the block finality time and message propagation, we plan to enhance future versions of PriFoB by deploying the DONS protocol [77] into the DTTP.

Appendices

A. Encryption Modelling

A.1. RSA keys generation

To generate RSA Public and Private keys ($[n, e]$ and d):

1. Select two primes $P, Q \in \mathbb{Z}^+$, and compute $\langle n \leftarrow P \times Q \rangle$
2. Find $\langle \lambda(n) \leftarrow \text{lcm}^8(P-1, Q-1) \rangle$
3. Select a third prime number $e \in \mathbb{Z}^+$, which is not a factor of n nor $\lambda(n)$, such that $1 < e < \lambda(n)$.
4. compute the modular multiplicative inverse of $e \bmod \lambda(n)$. That is $\langle e \times d \bmod \lambda(n) \rightarrow 1 \rangle$. Note that d shall then satisfy: $\text{gcd}^9(d, \lambda(n)) \rightarrow 1$.

A.2. RSA encryption and decryption

To utilize the RSA algorithm for encrypting a plain text M into a cipher C and vice versa:

1. Represent M as an integer using one of the standard representation methods (say $F(\cdot)$).
2. The cipher C can then be computed using the function $C \leftarrow F(M)^e \bmod n$.
3. To decipher C , one can compute $C^d \bmod n \rightarrow F'(M)$, which must equal $F(M)$.
4. Consequently, M can be deduced from $F'(M)$ by using the same representation method in step 1.

B. Sample PriFoB blocks

Listing 1: A sample DID block in PriFoB, newly confirmed into the Blockchain

```
{
  "Header": {
    "Type": "DID_block",
    "miner_id": "10.128.0.32",
    "Signature": "\u000b... (\"Body\" signed by authorized Miner with its private key)",
    "Index": 10
  },
  "Body": {
    "Transaction": {
      "Identifier": "University of Szeged",
      "institution_address": "10.128.0.29",
      "institution_public_key": "MEgC...",
      "Accredited By": {
        "USA": {
          "Address": "10.128.0.28",
          "Person Who Signed This": "General Registrar",
          "Signature": "0\u0009f... (of Accreditation Body \u2192 using its private key)"
        },
        "UAE": {
          "Address": "10.128.0.27",
          "Person Who Signed This": "General Registrar",
          "Signature": "UPXd\u0000..."
        },
        "Jordan": {
          "Address": "10.128.0.33",
          "Person Who Signed This": "General Registrar",
          "Signature": "d<\u0000..."
        }
      },
      "Not Accredited by": {
        "Hungary": {
          "Address": "10.128.0.32",
          "Person Who Signed This": "General Registrar",
          "Signature": "D^\u000f..."
        }
      },
      "Previous_Signature": "\u0000f...(in the \"Header\" of the previous DID block )"
    },
    "schemes_chain": [... new Schema blocks issued by this issuer are appended here]
  }
}
```

⁸Least Common Multiple

⁹Greatest Common Divisor

Listing 2: A sample Schema block in PriFoB, newly confirmed into the Blockchain

```
{
  "Header": {
    "Type": "Schema_block",
    "miner_id": "10.128.0.28",
    "Signature": "\u0000...("Body" signed by authorized Miner using its private key)",
    "Index": 1},
  "Body": {
    "Transaction": {
      "institution_name": "University of Szeged",
      "institution_address": "10.128.0.29",
      "Identifier": "PhD_schema",
      "schema_public_key": "mDc...",
      "schema_attributes": [
        ["name", "Mandatory"],
        ["Grade", "Not Mandatory"],
        ... (
          ↪ more can be added here by issuer),
      ],
      "DID_index": 10},
      "Previous_Signature": "\u0006j...(in the "Header" of the previous Schema block)",
      "Issuer_Signature": "-jjK?\u...( "Transaction" signed by issuer using DID private key)"}
    "Hashes_of_revoked_credentials": [
      ... new Revoke blocks, which uses this schema of this specific
      ↪ issuer are appended here]
  }
```

Listing 3: A sample Revoke block in PriFoB, newly confirmed into the Blockchain

```
{
  "Header": {
    "Type": "Revoke_block",
    "miner_id": "10.128.0.27",
    "Signature": "\u00b6\u...( "Body" signed by authorized Miner using its private key)",
    "Index": 8},
  "Body": {
    "Transaction": {
      "DID_Identifier": "University of Szeged",
      "DID_index": 10,
      "Schema_index": 1,
      "institution_address": "10.128.0.29",
      "Schema_Identifier": "PhD_schema",
      "Identifier": "943fc4a7...(hash of VC to be revoked)"},
      "Previous_Signature": "00e4\u...(in the "Header" of the previous Revoke block)",
      "Issuer_Signature": "3\u00e4\u...(signed by the issuer using Schema private key)"}
  }
```

Acknowledgment

The research leading to these results has received funding from the TruBlo project of the European Union's Horizon 2020 research and innovation programme under grant agreement No. 957228.

References

- [1] B. Chakroun, J. Keevy, Digital credentialing: implications for the recognition of learning across borders, United Nations Educational, Scientific and Cultural Organization 7 (2018).
- [2] E. Commission, 2018 reform of eu data protection rules, ec.europa.eu/commission/sites/beta-political/files/data-protection-factsheet-changes_en.pdf, 2020.
- [3] P. Habibi, M. Farhoudi, S. Kazemian, S. Khorsandi, A. Leon-Garcia, Fog computing: a comprehensive architectural survey, *IEEE Access* 8 (2020) 69105–69133.
- [4] L. Loffi, C. M. Westphall, L. D. Grüttner, C. B. Westphall, Mutual authentication with multi-factor in iot-fog-cloud environment, *Journal of Network and Computer Applications* 176 (2021) 102932.
- [5] A. A. Mutlag, M. Khanapi Abd Ghani, M. A. Mohammed, M. S. Maashi, O. Mohd, S. A. Mostafa, K. H. Abdulkareem, G. Marques, I. de la Torre Díez, Mafc: Multi-agent fog computing model for healthcare critical tasks management, *Sensors* 20 (2020) 1853.
- [6] F. Khosroabadi, F. Fotouhi-Ghazvini, H. Fotouhi, Scatter: Service placement in real-time fog-assisted iot networks, *Journal of Sensor and Actuator Networks* 10 (2021) 26.
- [7] T. Wang, J. Zhou, A. Liu, M. Z. A. Bhuiyan, G. Wang, W. Jia, Fog-based computing and storage offloading for data synchronization in iot, *IEEE Internet of Things Journal* 6 (2018) 4272–4282.
- [8] M. Arif, G. Wang, V. E. Balas, O. Geman, A. Castiglione, J. Chen, Sdn based communications privacy-preserving architecture for vanets using fog computing, *Vehicular Communications* 26 (2020) 100265.
- [9] Z. Zheng, S. Xie, H.-N. Dai, X. Chen, H. Wang, Blockchain challenges and opportunities: A survey, *International Journal of Web and Grid Services* 14 (2018) 352–375.
- [10] H. T. Vo, A. Kundu, M. K. Mohania, Research directions in blockchain data management and analytics., in: *EDBT*, 2018, pp. 445–448.
- [11] B. Notheisen, J. B. Cholewa, A. P. Shanmugam, Trading real-world assets on blockchain, *Business & Information Systems Engineering* 59 (2017) 425–440.
- [12] R. Dennis, G. Owen, Rep on the block: A next generation reputation system based on the blockchain, in: *2015 10th International Conference for Internet Technology and Secured Transactions (ICITST)*, IEEE, 2015, pp. 131–138.
- [13] P. Dunphy, F. A. Petitcolas, A first look at identity management schemes on the blockchain, *IEEE security & privacy* 16 (2018) 20–29.
- [14] D. J. Langley, J. van Doorn, I. C. Ng, S. Stieglitz, A. Lazovik, A. Boonstra, The Internet of everything: Smart things and their impact on business models, *Journal of Business Research* 122 (2021) 853–863.
- [15] H. Baniata, A. Kertész, A survey on blockchain-fog integration approaches, *IEEE Access* 8 (2020) 102657–102668.
- [16] P. K. Singh, R. Singh, S. K. Nandi, S. Nandi, Managing smart home appliances with proof of authority and blockchain, in: *International Conference on Innovations for Community Services*, Springer, 2019, pp. 221–232.
- [17] J. A. Garay, A. Kiayias, G. Panagiotakos, Consensus from signatures of work, in: *Cryptographers' Track at the RSA Conference*, Springer, 2020, pp. 319–344.
- [18] H. Pervez, M. Muneeb, M. U. Irfan, I. U. Haq, A comparative analysis of dag-based blockchain architectures, in: *2018 12th International Conference on Open Source Systems and Technologies (ICOSST)*, IEEE, 2018, pp. 27–34.
- [19] J. Shah, D. Dubaria, Building modern clouds: using docker, kubernetes & google cloud platform, in: *2019 IEEE 9th Annual Computing and Communication Workshop and Conference (CCWC)*, IEEE, 2019, pp. 0184–0189.
- [20] S. Shaik, S. Baskiyar, Hierarchical and autonomous fog architecture, in: *Proceedings of the 47th International Conference on Parallel Processing Companion*, 2018, pp. 1–8.
- [21] C. Thomas Porter, C. CCNP, M. Gough, et al., How to cheat at VoIP security, Syngress, 2011.
- [22] L. Johnson, Security controls evaluation, testing, and assessment handbook, Academic Press, 2019.
- [23] H. Yoshida, A. Biryukov, Analysis of a sha-256 variant, in: *International Workshop on Selected Areas in Cryptography*, Springer, 2005, pp. 245–260.
- [24] C. De Cannière, Analysis and design of symmetric encryption algorithms, Doctoral Dissertaion, KULeuven (2007).
- [25] M.-L. Akkar, C. Giraud, An implementation of des and aes, secure against some attacks, in: *International Workshop on Cryptographic Hardware and Embedded Systems*, Springer, 2001, pp. 309–318.
- [26] R. L. Rivest, A. Shamir, L. Adleman, A method for obtaining digital signatures and public-key cryptosystems, *Communications of the ACM* 21 (1978) 120–126.
- [27] R. S. Malyan, A. K. Madan, Blockchain technology as a tool to manage digital identity: A conceptual study, in: *Advances in Manufacturing and Industrial Engineering*, Springer, 2021, pp. 635–647.
- [28] J. Bacerlar Almeida, M. Barbosa, E. Bangerter, G. Barthe, S. Krenn, S. Zanella Béguelin, Full proof cryptography: verifiable compilation of efficient zero-knowledge protocols, in: *Proceedings of the 2012 ACM conference on Computer and communications security*, 2012, pp. 488–500.
- [29] Q. Wang, J. Yu, S. Chen, Y. Xiang, Sok: Diving into dag-based blockchain systems, *arXiv preprint arXiv:2012.06128* (2020).
- [30] E. Shi, Analysis of deterministic longest-chain protocols, in: *2019 IEEE 32nd Computer Security Foundations Symposium (CSF)*, IEEE, 2019, pp. 122–12213.
- [31] E. Anceaume, A. Pozzo, T. Rieutord, S. Tucci-Piergiovanni, On finality in blockchains, *arXiv preprint arXiv:2012.10172* (2020).
- [32] E. Bandara, W. K. Ng, K. De Zoysa, N. Fernando, S. Tharaka, P. Maurakirathan, N. Jayasuriya, Mystiko—blockchain meets big data, in: *2018 IEEE international conference on big data (big data)*, IEEE, 2018, pp. 3024–3032.
- [33] H. Baniata, A. Kertész, Pf-bvm: A privacy-aware fog-enhanced blockchain validation mechanism., in: *CLOSER*, 2020, pp. 430–439.
- [34] T. T. A. Dinh, J. Wang, G. Chen, R. Liu, B. C. Ooi, K.-L. Tan, Blockbench: A framework for analyzing private blockchains, in: *Proceedings of the 2017 ACM International Conference on Management of Data*, 2017, pp. 1085–1100.
- [35] T. Dierks, E. Rescorla, The transport layer security (tls) protocol version 1.2, *RFC 5246* (2008) 1–104.
- [36] L. Eichhorn, T. Shreedhar, A. Zavadovski, N. Mohan, Distributed ledgers for distributed edge: Are we there yet?, in: *Proceedings of the Interdisciplinary Workshop on (de) Centralization in the Internet*, 2021, pp. 26–33.
- [37] W. Al Amiri, M. Baza, K. Banawan, M. Mahmoud, W. Alasmay, K. Akkaya, Towards secure smart parking system using blockchain technology, in: *2020 IEEE 17th Annual Consumer Communications & Networking Conference (CCNC)*, IEEE, 2020, pp. 1–2.
- [38] A. DOGTIEV, Mobile advertising rates (2021), businessofapps.com/ads/research/mobile-app-advertising-cpm-rates/, 2021. Accessed: 2021-12-01.
- [39] G. corp., Vm instances pricing, cloud.google.com/compute/vm-instance-pricing, 2021. Accessed: 2021-12-01.
- [40] A. Kertész, H. Baniata, Consistency analysis of distributed ledgers in fog-enhanced blockchains, in: *European Conference on Parallel Processing*, Springer, 2021.
- [41] D. Mahto, D. K. Yadav, Rsa and ecc: a comparative analysis, *International journal of applied engineering research* 12 (2017) 9053–9061.
- [42] P. Mell, K. Scarfone, S. Romanosky, Common vulnerability scoring system, *IEEE Security & Privacy* 4 (2006) 85–89.
- [43] M. Nowak, M. Walkowski, S. Sujecki, Conversion of cvss base score from 2.0 to 3.1, in: *2021 International Conference on Software, Telecommunications and Computer Networks (SoftCOM)*, IEEE, 2021, pp. 1–3.
- [44] A. Begum, A. Tareq, M. Sultana, M. Sohel, T. Rahman, A. Sarwar, Blockchain attacks, analysis and a model to solve double spending attack, *International Journal of Machine Learning and Computing* 10 (2020) 352–357.

- [45] J. De Souza Jr, D. De Araújo, G. Barbosa, P. Letouze, An international accreditation system for healthcare professionals based on blockchain, *Int. J. Info. Educ. Tech* 9 (2019) 462–469.
- [46] Z. Feher, Development of fog-based blockchain framework, Master's thesis, University of Szeged, Faculty of Informatics, Szeged, Hungary, 2021.
- [47] D. Urbančok, Blockchain open-source software comparison, Master's thesis, Masaryk University, Faculty of Informatics, Brno, The Czech Republic, 2019.
- [48] X. Xu, G. Sun, L. Luo, H. Cao, H. Yu, A. V. Vasilakos, Latency performance modeling and analysis for hyperledger fabric blockchain network, *Information Processing & Management* 58 (2021) 102436.
- [49] L. Zhang, B. Lee, Y. Ye, Y. Qiao, Ethereum transaction performance evaluation using test-nets, in: *European Conference on Parallel Processing*, Springer, 2019, pp. 179–190.
- [50] F. Härer, H.-G. Fill, Decentralized attestation of conceptual models using the ethereum blockchain, in: *2019 IEEE 21st Conference on Business Informatics (CBI)*, volume 1, IEEE, 2019, pp. 104–113.
- [51] M. Bampatsikos, C. Ntantogian, C. Xenakis, S. C. Thomopoulos, Barrett blockchain regulated remote attestation, in: *IEEE/WIC/ACM International Conference on Web Intelligence-Companion Volume*, 2019, pp. 256–262.
- [52] P. T. AG, Gdpr checklist for data controllers, <https://gdpr.eu/checklist/>, 2021. Accessed: 2021-12-12.
- [53] Q.-L. Li, J.-Y. Ma, Y.-X. Chang, Blockchain queue theory, in: *International Conference on Computational Social Networks*, Springer, 2018, pp. 25–40.
- [54] S. Smetnin, A. Ometov, N. Kannengießer, B. Sturm, M. Komarov, A. Sunyaev, Modeling of distributed ledgers: Challenges and future perspectives, in: *2020 IEEE 22nd Conference on Business Informatics (CBI)*, volume 1, IEEE, 2020, pp. 162–171.
- [55] M. A. R. Gonzalez, Discrete event simulation of a queue using python, github.com/miguelrizzog96/QueueSimulationPython, 2020.
- [56] T. A. Severini, *Likelihood methods in statistics*, Oxford University Press, 2000.
- [57] L. Fauteux, P. Alexandre, L. Maxime, M. Guillaume, BLOCKCHAIN-TACKLING THE ACADEMIC CREDENTIAL VERIFICATION PROBLEM, Technical Report, Tampere University, Faculty of Management and Business, 2020.
- [58] Blockcerts home page, <https://www.blockcerts.org/>, 2021. Accessed: 2021-12-03.
- [59] Opencerts home page, <https://opencerts.io/>, 2021. Accessed: 2021-12-04.
- [60] K. Batzavalis, R. Bala, A. Norta, O. Norta Partners, A platform for leveraging blockchain technology for the storage, issuance and authentication of academic credentials., <https://www.trustededucation.io/>, 2021. Accessed: 2021-12-16.
- [61] K. Bhumichitr, S. Channarukul, Acachain: Academic credential attestation system using blockchain, in: *Proceedings of the 11th International Conference on Advances in Information Technology*, 2020, pp. 1–8.
- [62] Z. Li, H. Wu, L. H. Lao, S. Guo, Y. Yang, B. Xiao, Pistis: Issuing trusted and authorized certificates with distributed ledger and tee, *IEEE Transactions on Parallel and Distributed Systems* 33 (2021) 1636–1649.
- [63] S. Anant, I. Sneha, T. Tanisha J., S. Ujjwala, Academic credential verification technique using blockchain, *International Journal of Advanced Science and Technology* 29 (2020) 4244 – 4254.
- [64] I. M. Ahammad MS Tomal MH, Rahman A, Distb-cvs: A distributed secure blockchain based online certificate verification system from bangladesh perspective, in: *2nd International Conference on Advanced Information and Communication Technology 2020 (ICAICT 2020)*, IEEE, 2020, pp. 1–6.
- [65] K. Singh, O. Dib, C. Huyart, K. Toumi, A novel credential protocol for protecting personal attributes in blockchain, *Computers & Electrical Engineering* 83 (2020) 106586.
- [66] L. Foundation, Hyperledger indy, hyperledger.org/use/hyperledger-indy, 2020.
- [67] P.-L. Aublin, S. B. Mokhtar, V. Quéma, Rbft: Redundant byzantine fault tolerance, in: *2013 IEEE 33rd International Conference on Distributed Computing Systems*, IEEE, 2013, pp. 297–306.
- [68] P. J. Windley, How sovryn works, Windely. com (2016).
- [69] M. Sopek, P. Gradzki, D. Kuzinski, R. Trojczak, R. Trypuz, Legal entity identifier blockchained by a hyperledger indy implementation of graphchain, in: *Research Conference on Metadata and Semantics Research*, Springer, 2018, pp. 26–36.
- [70] A. Tariq, H. B. Haq, S. T. Ali, Cerberus: A blockchain-based accreditation and degree verification system, *arXiv preprint arXiv:1912.06812* (2019).
- [71] D. Hankerson, A. J. Menezes, S. Vanstone, *Guide to elliptic curve cryptography*, Springer Science & Business Media, 2006.
- [72] M. Bafandehkar, S. M. Yasin, R. Mahmood, Z. M. Hanapi, Comparison of ecc and rsa algorithm in resource constrained devices, in: *2013 International Conference on IT Convergence and Security (ICITCS)*, IEEE, 2013, pp. 1–3.
- [73] K. Maletsky, *Rsa vs ecc comparison for embedded systems*, White Paper, Atmel 5 (2015).
- [74] Y. Sompolinsky, Y. Lewenberg, A. Zohar, Spectre: a fast and scalable cryptocurrency protocol., *IACR Cryptol. ePrint Arch.* 2016 (2016).
- [75] Y. Sompolinsky, A. Zohar, Phantom, *IACR Cryptology ePrint Archive*, Report 2018/104 (2018).
- [76] I. Bentov, P. Hubáček, T. Moran, A. Nadler, Tortoise and hares consensus: the meshcash framework for incentive-compatible, scalable cryptocurrencies., *IACR Cryptol. ePrint Arch.* 2017 (2017) 300.
- [77] H. Baniata, A. Anaqreh, A. Kertesz, Dons: Dynamic optimized neighbor selection for smart blockchain networks, *Future Generation Computer Systems* 130 (2022) 75–90.