Energy-Aware Privacy Controls for Clouds

Jianzhou Mao, Ting Cao, Xiaopu Peng, Tathagata Bhattacharya
Wei-Shinn Ku, and Xiao Qin
Department of Computer Science and Software Engineering
Auburn University, Auburn, AL.
{jzm0122, tzc0028, xzp0007, tzb0063, weishinn, xqin}@auburn.edu

***Abstract*—Cloud computing has radically changed the landscape of computing, storage, and communication infrastructures and services. Cloud computing's benefits encompass on-demand capacity, low cost of ownership, and flexible pricing. While moving towards the concept of on-demand services and resource pooling in distributed computing environments, privacy protection becomes a major concern due to the sharing and consolidation features of clouds. At the same time, blockchain is an ideal privacy protection technology characterized by decentralization, transparency, data security, and system autonomy. In this paper, we navigate leading-edge energy saving and privacy protection techniques for clouds. Next, we investigate privacy controls in blockchain systems. Inspired by modern blockchain and cloud computing techniques, we articulate a research roadmap towards future energy-aware privacy protection mechanisms in clouds. As a case study, we propose a blockchain-based VM consolidation framework accompanied by the DVFS (Dynamic Voltage and Frequency Scaling) technique to offer energy savings and privacy controls in clouds. We expect that the roadmap will open up the potential to develop energy-efficient blockchain-based cloud computing platforms.**

***Index Terms*—Privacy preserving, energy saving, cloud computing, blockchain.**

I. Introduction

We navigate energy-aware privacy preserving techniques in realms of centralized and decentralized computing systems. We start our investigation by focusing on cloud data centers, the backbone of cloud infrastructure platforms supporting large-scale data processing and storage, followed by blockchain techniques that safeguard energy transactions in a distributed network. This study is inspired by the following two motivations - (1) privacy-aware energy-efficient data storage (see Section I-A) and (2) blockchain-based privacy preserving techniques for clouds (see Section I-B).

A. Motivation 1: Privacy-aware Green Data Storage

BP or British Petroleum forecasts that global energy demand continues to grow in the predictable future, driven by increasing prosperity and living standards [1]. Moreover, ExxonMobil predicts that global energy demand will rises by 20 percent to 2040 and; during the same time period, the electricity consumption will rise by 60 percent [2]. The trend to electrify buildings, factories, cars, and buses, along with smart appliances, spurs the pressing need for more electricity everywhere. Constructing energy-efficient data centers catering to cloud computing aims to address the concerns of increasing electricity demands. Cloud data center is energy friendly by the virtue of the on-demand deployment of resources through cutting-edge virtualization technology. Cloud users are enabled to swiftly allocate computing power and resources according to dynamically changing needs at any time without maintaining the underlying physical structure of computing platforms. Growing evidence demonstrates that cloud computing platforms are slated to conserve energy by providing information resources in a pay-as-you-go model [3].

The virtualization technology in cloud data centers brings forth versatility and reliability to cloud services. To facilitate virtual machine (VM) management in cloud data centers, system administrators make use of shared data storage to handle VMs' data in uniform storage space. The concept of shared data storage is implemented by adopting a centralized structure, where all physical nodes are connected to a centralized storage unit such as network-attached storage (NAS) [4]. Even it is convenient to build high-end centralized storage systems, a centralized structure is prone to data leakage of VMs running in cloud data centers when access privileges of some nodes are comprised [5].

B. Motivation 2: Blockchains and Privacy Protections

The preceding discussions emphasize the importance of decentralization, one trait commonly associated with blockchain [6]. Blockchain is characterized by decentralization, transparency, data security, and system autonomy. It has been applied widely in areas such as finance, education and employment, culture and entertainment, public service, information security, healthcare, supply chain, and internet of things. Moreover, blockchain gains its popularity in the energy sectors [7] thanks to blockchain's underpinning characteristics such as anonymity, decentralization, transparency, and reliability.

Despite the observable benefits of using blockchain in a diversity of areas including the energy sectors, privacy concerns are restricting blockchain's applications. For instance, users may need to disclose private energy demand data to a third-party in order to schedule the use of shared energy resources. This process may reveal sensitive personal data such as working patterns, number of occupants, and vacation periods. On the other hand, the privacy-related attacks should not be overlooked. Representative privacy concerns include linking attacks, which utilize open information recorded in blocks and obtain privacy from linking the information with other datasets. Moreover, in the arena of privacy preserving methods, we confront the following challenges. First, more

times than not, attackers are capable of obtaining privacy with inaccurate data. Thus, any features pertaining to privacy control ought to be hidden to prevent privacy leakage. Second, most existing differential privacy schemes are inadequate for accurately recording energy trading operations because a noised record stored in blocks results in a malfunction of a transaction ledger. Hence, there is an urgent demand for designing adoptable privacy-preserving mechanisms catering for blockchain-enabled energy applications.

C. Organization

The remainder of this paper is organized as follows. Section II outlines the energy saving techniques of cloud computing systems from the perspectives of infrastructure, hardware techniques, and software solutions. In Section III, we survey various privacy preserving methods in clouds. In Section IV, we introduce representative blockchain-based privacy protection mechanisms. We present a research roadmap, where new approaches and directions are discussed in Section V. Finally, Section VI presents concluding remarks.

II. Green Cloud Data Centers

Cloud computing has radically changed the landscape of computing, storage, and communication infrastructures. With strong interest and investment from the industry and government, cloud computing infrastructures are being increasingly patronized by both organizations and individuals. With increasing energy prices and data center scaling, high energy consumption has become an impediment to the development of cloud-computing environments. Reducing operational costs of data centers should be achieved through boosting energy efficiency. As such, optimizing the energy efficiency of data centers supporting cloud computing has captured much attention. A flood of intriguing studies have been recently conducted to facilitate the development of energy-efficient data centers. As shown in Fig. 1, we classify leading-edge energy saving techniques from the following perspectives of infrastructure, hardware techniques, and software solutions.

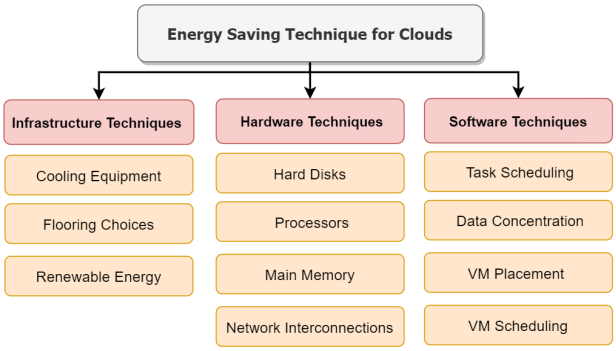

Fig. 1. Commonly adopted energy conservation techniques for clouds.

A. Energy-Efficient Infrastructure Techniques

Infrastructure techniques intend to curb energy consumption by building green data centers. The infrastructural energy conservation techniques entail cooling equipment [8], flooring choices [9], and renewable energy sources [10].

With regard to the cooling principle, existing cooling solutions fall into three camps, namely, air-cooling, liquid-cooling or free-cooling schemes. Thanks to low operational cost accompanied by simple maintenance, the air-cooling technology is the most conventional way of cooling down large-scale data centers [8]. Unlike air-cooling solutions, liquid cooling is one of the most prominent and practical methods to be directly or indirectly implemented in data centers. An indirect liquid-cooling system embraces a heat dissipation process where heat sources and liquid coolants contact indirectly [11]. In contrast, liquid coolant in a direct liquid-cooling method directly contacts electronic devices, where dielectric fluid offers electrical insulation [12]. Furthermore, a raft of data centers leverage the free-cooling technology to conserve cooling cost by the virtue of natural free cooling sources [13]. For instance, Facebook constructed a naturally cooled data center in northern Sweden in 2013 [14].

Apart from the above versatile cooling techniques, a growing number of green data centers adopt flooring techniques with perforated tiles and a raised floor plenum for cool air intake [9]. The application of renewable sources of energy (e.g., solar, geothermal, wind, hydro power) unequivocally and considerably cut back energy consumption in data centers. For example, Zhang *et al.* devised *GreenWare*, a novel middleware system that conducts dynamic request dispatching to enhance the percentage of renewable energy powering a distributed data center [10]. *GreenWare* fully utilizes renewable energy sources while meeting the desired cost budget for cloud service providers.

B. Energy-Aware Hardware Techniques

Acquiring a diversity of energy-efficient hardware components makes it possible and desirable to construct green data centers with high energy efficiency. Energy-efficient hardware techniques deployed in modern data centers include hard disks [15], processors [16], main memory [17], and network interconnections [18].

When it comes to data storage systems, solid-state disks and multi-speed disks are proved to be capable of trimming energy consumed by disks [15]. The dynamic voltage frequency scaling (DVFS) technique [16] is a popular technique, which is a feasible solution to conserve energy consumption of DVFS-enabled CPU and main memory - key underpinnings in computing servers. It is evident that DVFS enables processors and main memory to consume less power by electing the most appropriate frequency and supplied voltage. For example, Garg *et al.* developed the near-optimal energy-efficient scheduling algorithms, where DVFS is employed to decrease carbon emission by scaling down CPU frequency and optimizing cloud providers' profits [19]. Speaking of energy-aware network interconnects for data centers, an array of network architectures have been proposed and customized for data centers. Representative techniques include, but not limited

to, energy proportional networks [20], networks based on the elastic tree topology [21] and the *Proteus* network [22].

C. Energy-aware Software Techniques

A wide range of software techniques were designed to conserve energy consumption in clouds. Such cutting-edge software solutions include task scheduling [23], data concentration [24], virtual machine (VM) placement [25] and scheduling [26].

Dong *et al.* developed a greedy task-scheduling policy (the most efficient-server-first scheduling) to enhance energy efficiency of servers deployed in data centers. This scheduling scheme shortens the average task response time while minimizing the energy expenditure of servers [23]. Pinheiro *et al.* devised the popular data concentration (PDC) technique, a promising technique that migrates frequently accessed data to a small subset of disks [24]. The overarching goal of PDC is to skew the load towards a few of the disks, allowing the other disks to be transitioned to a low-power energy-saving mode [24]. VM placement, consisting of VM migration and consolidation, is one of the most common approaches to achieving high energy efficiency by dynamically scaling down the size of running clusters. With the help of virtualization, energy consumed by cluster computing infrastructures are immensely reduced by applying energy-aware VM migrations and consolidations [25]. In the arena of VM management, VM scheduling is an outstanding energy-saving method in cloud environments. For example, Li *et al.* proposed GRANITE - a holistic virtual machine scheduling algorithm being capable of minimizing total energy consumption in a data center [26]. GRANITE embraces an elaborate thermal model, which is adept at analyzing the temperature distribution of airflow and processors [26].

III. Privacy Protections in Cloud Computing

The key benefits of cloud computing, from the cloud provider's perspective, include resource consolidation, uniform management, and cost-effective operation. When it comes to cloud computing users, cloud computing's benefits encompass on-demand capacity, low cost of ownership, and flexible pricing. The sharing and consolidation features that bring forth such benefits inevitably introduce potential security and privacy concerns. Security and privacy issues arising from illegal and unethical use of data as well as disclosure of confidential information can tremendously hinder users' willingness to participate in cloud-based services. Recognizing such security concerns, a growing research and development efforts in the industry and academia have been devoted to preserving the cloud's data privacy. Domingo-Ferrer *et al.* classified the privacy protection techniques into three categories (see Fig. 2), namely, (1) data splitting mechanisms, (2) data anonymization methods, and (3) cryptographic techniques [27].

A. Data Splitting

Privacy-preserving data splitting protects data privacy by deploying multiple-CSPs-based (cloud service providers) architectures [28]. It is evident that data splitting minimizes information leakage through distributed data among an array of CSPs. This technique is proved to be a practical solution as long as the distributed CSPs have zero communication with one another.

A horde of data splitting mechanisms devised in the prior studies undertake data partitioning at the binary level. For instance, Zhang *et al.* implemented a scheme to split sensitive files into bits, which are reassembled to form numerous part-files before being uploaded to various cloud storage servers. After part-files are downloaded from multiple cloud servers, the part-files are concatenated to build an original file [29]. To strength the security protection for split chunks, Gai *et al.* mixed the byte-level data splitting technique with an encryption module [30]. The secure-efficient data distributions algorithm or *SED2* was proposed to spill data in a way of preventing sensitive information from leaking on clouds. The SED2 algorithm was realized through two underpinning algorithms, which are slated to efficiently encrypt and decrypt data [30]. Dev *et al.* assigned each file a privacy level in accordance with the sensitivity of the file's content [31]. Then, file fragments are created under the governance of a standalone RAID storage system, in which data with high privacy levels are stored in trustworthy locations [31].

Apart from the aforementioned binary-level strategies, innovative split methods built at the attribute level capture much attention. In this technique category, data splitting may be executed in a horizontal or vertical fashion. A horizontal format implies that sets of data records are separately stored, whereas a vertical layout indicates that sets of attributes are separately stored. In the case of horizontal data splitting, data sets are structured in a tabular format according to attributes. To achieve confidentiality at the record level, vertical chunks are comprised of all data records on a single attribute. Aggarwal *et al.* designed an approach to decomposing a dataset into two privacy-preserving vertical fragments [32]. With the deployment of the graph-coloring techniques, the proposed decomposition algorithm cuts back the data querying cost. In case sensitive attribute pairs require more than two chunks to preserve data privacy, an encryption module will be incorporated [32]. Ganapathy *et al.* investigated a solution based on two fragments coupled with a data encryption service [33]. The three additional heuristics were developed to shorten query time by applying the greedy hill-climbing algorithm. In this study, the time complexity of the proposed data splitting solution was articulated [33].

B. Data Anonymization Methods

A key advantage of anonymized data over encrypted data and data splitting is rooted in ease of data processing. In the realm of *data anonymization*, masking is merely performed at the data storage phase. More times than not, data anonymization are implemented by linear or quasi-linear algorithms. In this type of approaches, anonymized data's query, being transparent, does not incur hefty overhead for clouds. Original data of a micro data set are manipulated to originate new data, which are applicable for statistical analysis. In doing so, the

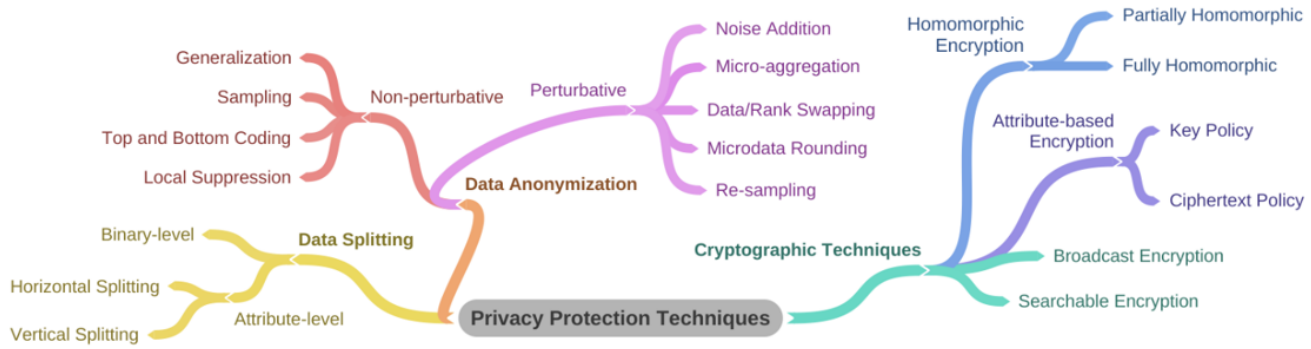

Fig. 2. Commonly adopted privacy protection techniques for clouds.

confidentiality of respondents is enforced. Masking methods in turn are be divided into two camps, depending on the effect on original data (see Fig. 3).

Because Non-perturbative methods have no intent to alter data, these methods yield partial suppressions or reductions of detail in an original dataset. Representative non-perturbative masking methods are sampling, generalization, top and bottom coding, and local suppression. A sampling scheme masks a sample of original files rather than publishing the original files [34]. If an intruder identifies a unique record in released file (sample), the intruder will be unsure about the data uniqueness in the original file. Generalization, which is referred to as global recording, forms a new less specific attribute by combining several more specific categorical attributes [35]. Top and bottom coding approaches can be envisioned as special cases of generalization techniques, where a new category is forged by gleaning values that are above (top coding) or below (bottom coding) a threshold [36]. A local suppression mechanism aims to suppress the values of individual attributes to increase a set of records supporting quasi-identifiers. Chen *et al.* explored a local suppression method to build a customized privacy model for trajectory data anonymization [37].

A perturbative masking technique manipulates a dataset in a way that respondents' privacy is preserved to a certain degree. It is noteworthy that such an approach aims to protect statistical properties of the dataset. Representatives of leading-edge perturbative masking methods are noise addition, micro-aggregation, data/rank swapping, microdata rounding, and re-sampling. A noise addition scheme intends to mask an original dataset by injecting random noise [38]. A micro-aggregation solution groups individual tuples into small aggregates of a fixed dimension k, where an average over each aggregate rather than individual values is published [39]. Data/rank swapping techniques are adroit at transforming databases by switching values of confidential attributes across stored records [40]. Rounding techniques substitute rounded values for the original values of attributes [41]. The key idea behind re-sampling schemes is to exchange the values of a continuous attribute with an average value derived from a set of samples gleaned from original data [42].

C. Cryptographic Techniques

Cryptography is one of the predominant building blocks deployed to address privacy concerns in the clouds. For example, growing evidence indicates that homomorphic encryption [43], attribute-based encryption [44], broadcast encryption [45], searchable encryption [46] are adopted to offer remote secure computation solutions for clouds. These diversity of techniques furnish fine-grained access controls in cloud storage.

In homomorphic cryptosystems, encryption functions are a homomorphism that protects group operations by the virtue of ciphertexts. With homomorphic encryption algorithms in place, one is enabled to perform computations on ciphertexts without decrypting data in advance, thereby preserving data privacy. After Rivest *et al.* introduced homomorphism in 1978 [43], all the homomorphism schemes are classified into two categories, namely, partially homomorphic encryption and fully homomorphic encryption. A partially homomorphic encryption solution supports merely one type of operation repeatedly running for unlimited times [47]. In contrast, a fully homomorphic encryption allows an unlimited number of operations performed on encrypted data, where output data range within a given ciphertext space [48] [49].

Attribute-based encryption approaches are proved to be a practical and promising technique. These solutions cater to facilitate encrypted fine-grained access controls for outsourced data. An attribute-based encryption or *ABE* applies public-key encryption where the secret key of a user as well as ciphertext rely on attributes. The decryption of a ciphertext becomes feasible only if the attribute set of the user key matches the ciphertext's attributes [44]. Given an access policy, we group these solutions into two camps - (1) key policy attribute-based encryption schemes or KP-ABE and (2) ciphertext-policy attribute-based encryption schemes or CP-ABE. A KP-ABE scheme employs an attribute set to model encrypted data and to construct an access policy in private key [50]. On the flip side, ciphertext in CP-ABE is associated with an access policy, whereas secret keys are associated with attributes [51]. In this case, data owners are enabled to elect users who have the privilege to decode. If the policy ought to be frequently managed, the CP-ABE schemes will be flexible because the data owner can readily update the ciphertext's access structure.

Now let us introduce broadcast encryption and searchable

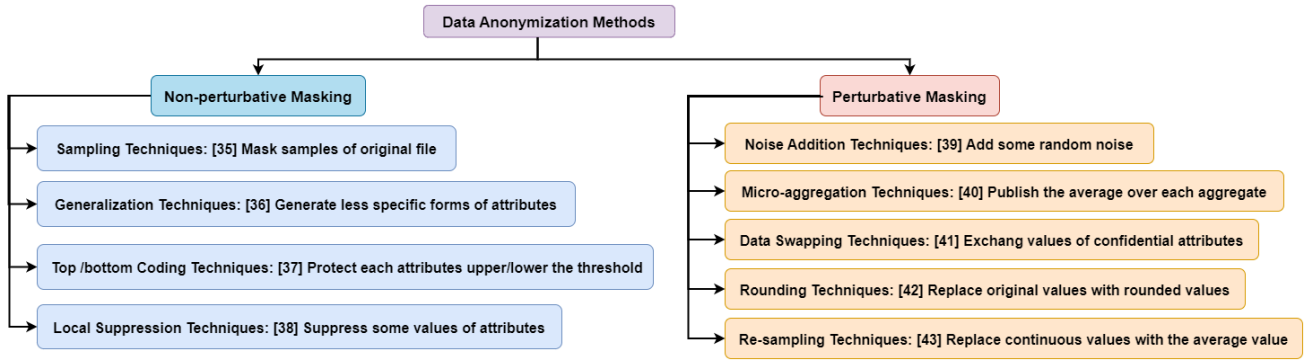

Fig. 3. Nine data anonymization methods for clouds.

encryption schemes - two popular cryptographic techniques. Broadcast encryption solutions allow a broadcaster to encrypt messages and to transmit the messages to any subset of authorized users. Given a broadcast encryption scheme, a broadcasting sender is positioned to encrypt a message by combining receivers' public identities in the subset coupled with system parameters. In doing so, only a dynamically changing privileged subset of users are able to decode encrypted messages [45]. Searchable encryption techniques equip data users with the capability to securely search over encrypted data using keywords without having to decrypt the data. The overall objective of a searchable encryption system is to combine confidentiality with respect to cloud providers with a powerful search functionality. Kamara *et al.* made use of the multicore architecture to realize a searchable encryption scheme [46]. Such a multicore-based implementation makes the searchable encryption module highly scalable. This novel solution offers sublinear search time by the virtue of a tree-based multimap data structure per keyword, which is referred to as red-black trees [46].

IV. Privacy of Blockchain Systems

The blockchain technology has emerged as a creative way of maintaining distributed systems thanks to its high efficiency, high data security, and high credibility at a low cost [6]. Blockchain techniques employ a linked block structure to store and verify data, the changes of which are synchronized by a trusted consensus mechanism. A growing number of novel consensus mechanisms catering to cryptocurrencies have been proposed in the past few years. For example, an innovative consensus method was incorporated in the Kraft system to avert multiple hash-rate scenarios [52], thereby offering stable average block times. Sompolinsky and Zohar designed the *GHOST chain selection* rule, which weights branches to speed up selection tasks for miners [53]. The PeerCensus system is capable of maintaining a strong consistency in Bitcoin-like systems [54]. *Discoin*, built atop PeerCensus, enhances consensus efficiency by decoupling block creations from transaction confirmation operations.

The blockchain techniques make it feasible to forge a tamper-proof storage system catering to data storage. It is noteworthy that data privacy challenges may hinder the wide applications of the blockchain technology. For instance, Kosba *et al.* unveiled that blockchain may not guarantee the privacy of transactions because the values of all transactions and balances for a public key are publicly visible [55].

Similar to the privacy-preserving methods for clouds (see Section III), a raft of privacy protection mechanisms were developed in the arena of the blockchain techniques. Representative mechanisms include, but not limited to, mixing service [56], anonymous signatures [57], and encryption techniques [58].

A. Mixing Service

Chaum *et al.* proposed a mechanism of coin mixing [56]. Fig. 4 (a) depicts that this mechanism allows privacy-seeking coin users to deliver transactions to a mixer service, which blends a pool of coins to delink a transaction trail. Bonneau *et al.* [59] developed *Mixcoin*, where a central server is in charge of mixing transaction addresses to offer external anonymity. Such a centralized mixing mechanism relies on third-party servers, where dishonest mixers may stealthily archive transaction records or provide poor mixing services. To tackle the thread imposed by untrusted servers, decentralized mixing pattern was proposed (see Fig.4 (b)). Because decentralized mixing mechanisms are independent of the credibility of third-party servers, decentralized designs effectively avert third party thefts and leakages of coin mixing information. The decentralized approaches also eliminate mixing fees. *CoinJoin* is the earliest decentralized mixing scheme first proposed by Maxwell *et al.* [60].

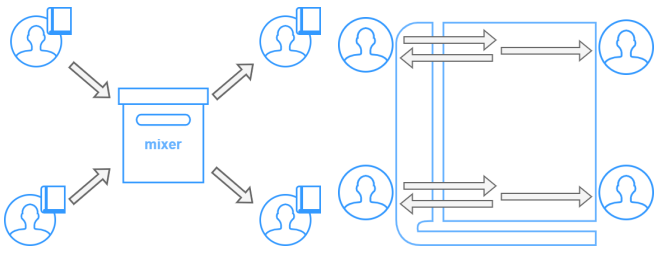

(a) Mixing transaction (b) CoinJoin transaction

Fig. 4. Mixing service mechanisms in blockchain.

B. Anonymous Signatures

Unlike the mixing methods suffer a delay in which participants discover their partners for transactions to be mixed, anonymous signature schemes have strengths in furnishing anonymity for signers. Among anonymous signature schemes, the group signature and ring signature solution are two representative schemes [57] [61]. The group signature method, designed by Chaum and Heyst [57], enables group members to set up anonymous signatures on behalf of a group, the managers of which may open signatures when the signature is disputed. A ring signature is constructed by a group member, whose identify is protected from the other members [61]. *Monero* is a successful implementation of the ring signature approach, where ring signatures coupled with hidden addresses to camouflage the linkage between input and output addresses [62].

C. Encryption Methods

The homomorphic encryption and attribute-based encryption methods are widely applied in blockchains. Recall that (see Section III-C) homomorphic encryption allows computations to be accomplished on encrypted data without accessing a decryption key. Distributed electronic voting and bidding systems deploy the homomorphic encryption technology to protect data privacy, to enhance the anonymity of participants, and to boost data reliability and verifiability [58]. It is evident that attribute-based encryption is powerful. For example, Lewko *et al.* implemented a decentralized attribute-based encryption scheme on a blockchain [63]. Nevertheless, applications of attribute-based encryption schemes ought to be further explored.

Another cryptographic technology that embraces privacy-preserving properties is *zero knowledge proofs (ZKP)* [64]. A prover in a zero knowledge proof system makes a verifier believe that a message is correct without sending valid information to the verifier. Non-interactive zero-knowledge proof or NIZK is an extension of ZKP, where only a single message is transferred from a prover to a verifier. NIZK was deployed in *Zcash* - a privacy-protecting digital currency system that shields transaction information [65].

All the aforementioned privacy-preserving techniques (mixing, anonymous, and encryption) are tabulated in Table 1, which summarizes the strengths and weaknesses of the leading-edge privacy protection techniques.

V. Roadmap

A. Energy-aware Privacy Protections in Clouds

Recall that (see Sections II and III) high energy efficiency and privacy protections are two vital design objectives for cloud computing platforms. In the first phase of our research roadmap, we will focus on bringing forth energy-efficient privacy protection techniques in clouds.

Splitting takes constant work to split an original data set into fragments. Extra energy consumption is expected because operations on the fragments lead to additional input/output overhead. Jaikar *et al.* devised a secure data distribution scheme anchored on secret splitting to preserve data privacy over clouds [66]. Although this technique protects sensitive data, the secret splitting technique inevitably increases energy usage due to extra bandwidth and storage utilization.

Most anonymization methods cost energy to generate anonymized data sets. Once anonymized data sets are generated and uploaded to the cloud, no further intervention is required. Any query like search and retrieval on anonymized data incurs no overhead on clouds. When it comes to homomorphic encryption, the complexity of encryption and decryption are normally higher than that of plain encryption. Searchable encryption's energy efficiency lies between those of plain encryption and homomorphic encryption.

There are four future research directions in developing energy-aware privacy protection systems in clouds.

- *Research Direction 1-1.* We propose to build energy-efficient data splitting mechanisms by storing fragments on energy-aware data storage systems such as *Eco-storage* [67].
- *Research Direction 1-2.* We plan to design energy-efficient anonymization and encryption strategies, which expect to become technological underpinnings of energy-aware privacy preserving mechanisms in clouds.
- *Research Direction 1-3.* We will navigator an approach to making a good trade-off between privacy protection and energy efficiency in clouds.
- *Research Direction 1-4.* We will delve in the development of blockchain-based privacy-preserving techniques, which offer high energy efficiency in clouds. Please refer to V-C for the details.

B. Energy-Efficient Blockchains for Privacy Controls

Unlike traditional centralized solutions, the *blockchain* technology securely manages chain data across a distributed and interlinked network of nodes. Blockchains, serving as a tamper-resistant distributed ledger, naturally offer data privacy protections in clouds.

Blockchain-based data provenance provides tamper-proof records, enables the transparency of data accountability in clouds, and enhances data privacy. Very recently, Ali *et al.* demonstrated that the blockchain techniques embrace immutable, deterministic, and public natures that play a vital role in data provenance [68]. The concept of *smart contract* balances data provenance, functionality, and trusted environment, regardless of on-chain or off-chain data storage. Liang *et al.* devised a decentralized and trusted cloud data provenance architecture - *ProvChain* - powered by the blockchain technology [69]. To glean provenance data on storage clouds, ProvChain is slated to detect user operations on cloud files. User's privacy is protected by ProvChain because users' identities are constructed in a hashed form, where only service providers are authorized to map hashed values to the identities.

Because wasting resources of mining networks becomes a key drawback of blockchain technology, we will explore the following research directions to construct energy-efficient blockchain techniques.

TABLE I. Summary of Privacy Techniques on Blockchian.

Techniques	Applications	Advantages	Disadvantages
Mixing	Mixcoin [59] CoinJoin [60]	It can obfuscate users' addresses from being linked.	Cause a Delay waiting to be mixed. High risky on unprotected transaction content.
Group signature		It is efficient anonymity and revocability.	Need a trusted manager.
Ring signature	Monero [62]	It can hide tradition origin and no need for trusted participant.	The identity of the signer cannot be revealed even in a dispute. The storage overhead is heavy.
Homomorphic encryption		It enable to perform computations on ciphertexts without decrypting data in advance.	Low efficiency for complex functions and no support for auditing.
NIZK	Zcash [65]	It can simultaneously achieve anonymity and transaction privacy.	Heavy computation overhead.

- *Research Direction 2-1.* We will pilot low-energy architecture designs to furnish the development of energy-efficient blockchains to preserve user privacy.
- *Research Direction 2-2.* We intend to explore new ways of enhancing the energy efficiency of consensus mechanisms in blockchains. We will kick off this direction by profiling energy consumption of popular consensus mechanisms on edge computing platforms.
- *Research Direction 2-3.* We plan to extend novel ideas [70] of applying renewable energy to the blockchain techniques. We will evaluate the energy efficiency of blockchain algorithms powered by solar and wind farms.

C. Blockchain-based Energy Management for Clouds

In one of our recent studies [71], we proposed a frequency-aware DVFS (Dynamic Voltage/Frequency Scaling) model aiming to conserve energy consumption of tasks imposing QoS requirements. Our model advocates for specifying QoS requirements with respect to frequency rather than execution time.

A proposed framework depicted in Fig. 5 combines virtual machine (VM) migrations with the DVFS technique to further improve energy efficiency. To protect data during VM migrations and data movement, we propose to make use of blockchain-enabled resource allocation to offer a transparent and trustworthy service on clouds (see also Section V-B). We promote the blockchain technique as an advanced decentralized structure to avoid privacy leakage during VM migrations while guarding data against malicious tampering.

To facilitate an energy-efficient cloud platform, our *DVFS model* intends to derive an optimal frequency ratio that leads to the minimum energy consumption of each active server while shutting down idle servers. The model makes power management decisions by incorporating the servers' hardware information such as static power P_c^{sta} and maximum dynamic power P_c^{dmax}. The *frequency adjusting module* compares the optimal frequency ratio and an overall minimum frequency requirement to appropriately configure frequency levels to cut back the energy consumption of VMs running on clouds.

We will spearhead this research effort along with the two directions below.

- *Research Direction 3-1.* We will develop a blockcahin-based VM consolidation mechanism accompanied by the DVFS technique to offer energy savings and privacy protection in clouds.
- *Research Direction 3-2.* We will incorporate reinforcement-learning algorithms into our privacy-aware energy management system to optimize the performance of resource allocation in the realm of cloud computing.

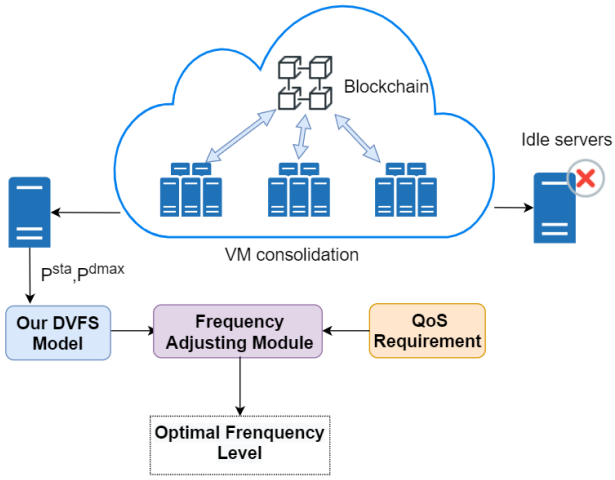

Fig. 5. The blockchain-based privacy protection for the VM consolidation mechanism coupled with the frequency-aware DVFS model.

VI. Concluding Remarks

We introduced in this paper the cutting-edge energy conservation techniques and privacy protection methods of cloud computing systems. An increasing number of energy-efficient datacenters have been built because high energy consumption has become an impediment to the development of cloud-computing environments. We navigated leading-edge energy saving techniques from the perspectives of infrastructure, hardware techniques, and software solutions. Recognizing that the sharing and consolidation features of clouds bring potential security and privacy concerns, we classified the privacy protection techniques into three categories, namely, data splitting mechanisms, data anonymization methods, and cryptographic techniques. Moreover, the blockchain techniques make it feasible to forge a tamper-proof storage system catering to data services on clouds. We surveyed an array of representative mechanisms such as mixing service, anonymous signatures, and encryption techniques. Among all the energy-saving and privacy protection schemes for cloud computing, we shed bright a light on blockchain-based VM consolidation combin-

ing DVFS to offer energy savings and privacy protection in clouds.

As the research roadmap towards the privacy-aware energy management in clouds, we proposed three connected research activities: (1) building energy-aware privacy protection services, (2) developing energy-efficient blockchains, and (3) devising blockchain-enabled energy management modules in clouds. Currently, we are in the process of designing a privacy-aware energy management system for cloud computing environments. Our novel energy management system is expected to achieve high privacy and energy efficiency in clouds by seamlessly integrating the blockchain, VM consolidation and frequency-aware DVFS model.

ACKNOWLEDGMENT

This work is supported in part by the U.S. National Science Foundation under Grants IIS-1618669, OAC-1642133, and CCF-0845257.

REFERENCES

[1] "Energy Outlook: 2020 edition ," BP p.l.c., Tech. Rep., 2020.

[2] "2019 Outlook for Energy: A perspective to 2040," ExxonMobil, Tech. Rep., 2019.

[3] A. Vafamehr and M. E. Khodayar, "Energy-aware cloud computing," *The Electricity Journal*, vol. 31, no. 2, pp. 40–49, 2018.

[4] S. B. Vaghani, "Virtual machine file system," *ACM SIGOPS Operating Systems Review*, vol. 44, no. 4, pp. 57–70, 2010.

[5] W. Itani, A. Kayssi, and A. Chehab, "Privacy as a service: Privacy-aware data storage and processing in cloud computing architectures," in *2009 Eighth IEEE International Conference on Dependable, Autonomic and Secure Computing*. IEEE, 2009, pp. 711–716.

[6] Z. Zheng, S. Xie, H.-N. Dai, X. Chen, and H. Wang, "Blockchain challenges and opportunities: A survey," *International Journal of Web and Grid Services*, vol. 14, no. 4, pp. 352–375, 2018.

[7] V. Brilliantova and T. W. Thurner, "Blockchain and the future of energy," *Technology in Society*, vol. 57, pp. 38–45, 2019.

[8] C. Nadjahi, H. Louahlia, and S. Lemasson, "A review of thermal management and innovative cooling strategies for data center," *Sustainable Computing: Informatics and Systems*, vol. 19, pp. 14–28, 2018.

[9] K. C. Karki and S. V. Patankar, "Airflow distribution through perforated tiles in raised-floor data centers," *Building and environment*, vol. 41, no. 6, pp. 734–744, 2006.

[10] Y. Zhang, Y. Wang, and X. Wang, "Greenware: Greening cloud-scale data centers to maximize the use of renewable energy," in *ACM/IFIP/USENIX International Conference on Distributed Systems Platforms and Open Distributed Processing*. Springer, 2011, pp. 143–164.

[11] Y. J. Lee, P. K. Singh, and P. S. Lee, "Fluid flow and heat transfer investigations on enhanced microchannel heat sink using oblique fins with parametric study," *International Journal of Heat and Mass Transfer*, vol. 81, pp. 325–336, 2015.

[12] E. A. Silk, E. L. Golliher, and R. P. Selvam, "Spray cooling heat transfer: technology overview and assessment of future challenges for microgravity application," *Energy Conversion and Management*, vol. 49, no. 3, pp. 453–468, 2008.

[13] H. Zhang, S. Shao, H. Xu, H. Zou, and C. Tian, "Free cooling of data centers: A review," *Renewable and Sustainable Energy Reviews*, vol. 35, pp. 171–182, 2014.

[14] A. Vonderau, "Scaling the cloud: Making state and infrastructure in sweden," *Ethnos*, vol. 84, no. 4, pp. 698–718, 2019.

[15] M. Song, "Minimizing power consumption in video servers by the combined use of solid-state disks and multi-speed disks," *IEEE Access*, vol. 6, pp. 25 737–25 746, 2018.

[16] W. Kim, M. S. Gupta, G.-Y. Wei, and D. Brooks, "System level analysis of fast, per-core dvfs using on-chip switching regulators," in *2008 IEEE 14th International Symposium on High Performance Computer Architecture*. IEEE, 2008, pp. 123–134.

[17] Q. Deng, D. Meisner, A. Bhattacharjee, T. F. Wenisch, and R. Bianchini, "Coscale: Coordinating cpu and memory system dvfs in server systems," in *2012 45th annual IEEE/ACM international symposium on microarchitecture*. IEEE, 2012, pp. 143–154.

[18] N. Akhter and M. Othman, "Energy aware resource allocation of cloud data center: review and open issues," *Cluster computing*, vol. 19, no. 3, pp. 1163–1182, 2016.

[19] S. K. Garg, C. S. Yeo, A. Anandasivam, and R. Buyya, "Environment-conscious scheduling of hpc applications on distributed cloud-oriented data centers," *Journal of Parallel and Distributed Computing*, vol. 71, no. 6, pp. 732–749, 2011.

[20] D. Abts, M. R. Marty, P. M. Wells, P. Klausler, and H. Liu, "Energy proportional datacenter networks," in *Proceedings of the 37th annual international symposium on Computer architecture*, 2010, pp. 338–347.

[21] B. Heller, S. Seetharaman, P. Mahadevan, Y. Yiakoumis, P. Sharma, S. Banerjee, and N. McKeown, "Elastictree: Saving energy in data center networks." in *Nsdi*, vol. 10, 2010, pp. 249–264.

[22] A. Singla, A. Singh, K. Ramachandran, L. Xu, and Y. Zhang, "Proteus: a topology malleable data center network," in *Proceedings of the 9th ACM SIGCOMM Workshop on Hot Topics in Networks*, 2010, pp. 1–6.

[23] Z. Dong, N. Liu, and R. Rojas-Cessa, "Greedy scheduling of tasks with time constraints for energy-efficient cloud-computing data centers," *Journal of Cloud Computing*, vol. 4, no. 1, pp. 1–14, 2015.

[24] E. Pinheiro and R. Bianchini, "Energy conservation techniques for disk array-based servers," in *Proceedings of the 18th annual international conference on Supercomputing*, 2004, pp. 68–78.

[25] S. K. Mishra, D. Puthal, B. Sahoo, P. P. Jayaraman, S. Jun, A. Y. Zomaya, and R. Ranjan, "Energy-efficient vm-placement in cloud data center," *Sustainable computing: informatics and systems*, vol. 20, pp. 48–55, 2018.

[26] X. Li, P. Garraghan, X. Jiang, Z. Wu, and J. Xu, "Holistic virtual machine scheduling in cloud datacenters towards minimizing total energy," *IEEE Transactions on Parallel and Distributed Systems*, vol. 29, no. 6, pp. 1317–1331, 2017.

[27] J. Domingo-Ferrer, O. Farras, J. Ribes-González, and D. Sánchez, "Privacy-preserving cloud computing on sensitive data: A survey of methods, products and challenges," *Computer Communications*, vol. 140, pp. 38–60, 2019.

[28] V. Balasaraswathi and S. Manikandan, "Enhanced security for multi-cloud storage using cryptographic data splitting with dynamic approach," in *2014 IEEE International Conference on Advanced Communications, Control and Computing Technologies*. IEEE, 2014, pp. 1190–1194.

[29] W. Zhang, X. Sun, and T. Xu, "Data privacy protection using multiple cloud storages," in *Proceedings 2013 International Conference on Mechatronic Sciences, Electric Engineering and Computer (MEC)*. IEEE, 2013, pp. 1768–1772.

[30] K. Gai, M. Qiu, and H. Zhao, "Security-aware efficient mass distributed storage approach for cloud systems in big data," in *2016 IEEE 2Nd international conference on big data security on cloud (bigdatasecurity), IEEE international conference on high performance and smart computing (HPSC), and IEEE international conference on intelligent data and security (IDS)*. IEEE, 2016, pp. 140–145.

[31] H. Dev, T. Sen, M. Basak, and M. E. Ali, "An approach to protect the privacy of cloud data from data mining based attacks," in *2012 SC Companion: High Performance Computing, Networking Storage and Analysis*. IEEE, 2012, pp. 1106–1115.

[32] G. Aggarwal, M. Bawa, P. Ganesan, H. Garcia-Molina, K. Kenthapadi, R. Motwani, U. Srivastava, D. Thomas, and Y. Xu, "Two can keep a secret: A distributed architecture for secure database services," *CIDR 2005*, 2005.

[33] V. Ganapathy, D. Thomas, T. Feder, H. Garcia-Molina, and R. Motwani, "Distributing data for secure database services," in *Proceedings of the 4th International Workshop on Privacy and Anonymity in the Information Society*, 2011, pp. 1–10.

[34] L. Willenborg and T. De Waal, *Elements of statistical disclosure control*. Springer Science & Business Media, 2012, vol. 155.

[35] P. Samarati and L. Sweeney, "Protecting privacy when disclosing information: k-anonymity and its enforcement through generalization and suppression," 1998.

[36] V. Ciriani, S. D. C. Di Vimercati, S. Foresti, and P. Samarati, "Micro-data protection," in *Secure data management in decentralized systems*. Springer, 2007, pp. 291–321.

[37] R. Chen, B. C. Fung, N. Mohammed, B. C. Desai, and K. Wang,

"Privacy-preserving trajectory data publishing by local suppression," *Information Sciences*, vol. 231, pp. 83–97, 2013.

[38] Y. Chen, J.-F. Martínez, P. Castillejo, and L. López, "A privacy-preserving noise addition data aggregation scheme for smart grid," *Energies*, vol. 11, no. 11, p. 2972, 2018.

[39] J. Soria-Comas, J. Domingo-Ferrer, D. Sanchez, and S. Martinez, "t-closeness through microaggregation: Strict privacy with enhanced utility preservation," *IEEE Transactions on Knowledge and Data Engineering*, vol. 27, no. 11, pp. 3098–3110, 2015.

[40] M. Rodriguez-Garcia, M. Batet, and D. Sanchez, "Utility-preserving privacy protection of nominal data sets via semantic rank swapping," *Information Fusion*, vol. 45, pp. 282–295, 2019.

[41] C. M. O'Keefe and D. B. Rubin, "Individual privacy versus public good: protecting confidentiality in health research," *Statistics in medicine*, vol. 34, no. 23, pp. 3081–3103, 2015.

[42] S. Martínez, D. Sánchez, and A. Valls, "Towards k-anonymous non-numerical data via semantic resampling," in *International Conference on Information Processing and Management of Uncertainty in Knowledge-Based Systems*. Springer, 2012, pp. 519–528.

[43] R. L. Rivest, L. Adleman, M. L. Dertouzos *et al.*, "On data banks and privacy homomorphisms," *Foundations of secure computation*, vol. 4, no. 11, pp. 169–180, 1978.

[44] A. Sahai and B. Waters, "Fuzzy identity-based encryption," in *Annual international conference on the theory and applications of cryptographic techniques*. Springer, 2005, pp. 457–473.

[45] K. He, J. Weng, J.-N. Liu, J. K. Liu, W. Liu, and R. H. Deng, "Anonymous identity-based broadcast encryption with chosen-ciphertext security," in *Proceedings of the 11th ACM on Asia Conference on Computer and Communications Security*, 2016, pp. 247–255.

[46] S. Kamara and C. Papamanthou, "Parallel and dynamic searchable symmetric encryption," in *International conference on financial cryptography and data security*. Springer, 2013, pp. 258–274.

[47] A. Kawachi, K. Tanaka, and K. Xagawa, "Multi-bit cryptosystems based on lattice problems," in *International Workshop on Public Key Cryptography*. Springer, 2007, pp. 315–329.

[48] C. Gentry, *A fully homomorphic encryption scheme*. Stanford university, 2009.

[49] F. Armknecht, C. Boyd, C. Carr, K. Gjøsteen, A. Jäschke, C. A. Reuter, and M. Strand, "A guide to fully homomorphic encryption." *IACR Cryptol. ePrint Arch.*, vol. 2015, p. 1192, 2015.

[50] Y. Rahulamathavan, S. Veluru, J. Han, F. Li, M. Rajarajan, and R. Lu, "User collusion avoidance scheme for privacy-preserving decentralized key-policy attribute-based encryption," *IEEE Transactions on Computers*, vol. 65, no. 9, pp. 2939–2946, 2015.

[51] L. Li, T. Gu, L. Chang, Z. Xu, Y. Liu, and J. Qian, "A ciphertext-policy attribute-based encryption based on an ordered binary decision diagram," *IEEE Access*, vol. 5, pp. 1137–1145, 2017.

[52] D. Kraft, "Difficulty control for blockchain-based consensus systems," *Peer-to-peer Networking and Applications*, vol. 9, no. 2, pp. 397–413, 2016.

[53] Y. Sompolinsky and A. Zohar, "Accelerating bitcoin's transaction processing. fast money grows on trees, not chains (2013)," *URL https://eprint. iacr. org/2013/881*, 2018.

[54] C. Decker, J. Seidel, and R. Wattenhofer, "Bitcoin meets strong consistency," in *Proceedings of the 17th International Conference on Distributed Computing and Networking*, 2016, pp. 1–10.

[55] A. Kosba, A. Miller, E. Shi, Z. Wen, and C. Papamanthou, "Hawk: The blockchain model of cryptography and privacy-preserving smart contracts," in *2016 IEEE symposium on security and privacy (SP)*. IEEE, 2016, pp. 839–858.

[56] D. L. Chaum, "Untraceable electronic mail, return addresses, and digital pseudonyms," *Communications of the ACM*, vol. 24, no. 2, pp. 84–90, 1981.

[57] D. Chaum and E. Van Heyst, "Group signatures," in *Workshop on the Theory and Application of of Cryptographic Techniques*. Springer, 1991, pp. 257–265.

[58] R. Tso, Z.-Y. Liu, and J.-H. Hsiao, "Distributed e-voting and e-bidding systems based on smart contract," *Electronics*, vol. 8, no. 4, p. 422, 2019.

[59] J. Bonneau, A. Narayanan, A. Miller, J. Clark, J. A. Kroll, and E. W. Felten, "Mixcoin: Anonymity for bitcoin with accountable mixes," in *International Conference on Financial Cryptography and Data Security*. Springer, 2014, pp. 486–504.

[60] G. Maxwell, "Coinjoin: Bitcoin privacy for the real world," in *Post on Bitcoin forum*, 2013.

[61] R. L. Rivest, A. Shamir, and Y. Tauman, "How to leak a secret," in *International Conference on the Theory and Application of Cryptology and Information Security*. Springer, 2001, pp. 552–565.

[62] "The Monero Project." [Online]. Available: https://www.getmonero.org//index.html

[63] A. Lewko and B. Waters, "Decentralizing attribute-based encryption," in *Annual international conference on the theory and applications of cryptographic techniques*. Springer, 2011, pp. 568–588.

[64] S. Goldwasser, S. Micali, and C. Rackoff, "The knowledge complexity of interactive proof systems," *SIAM Journal on computing*, vol. 18, no. 1, pp. 186–208, 1989.

[65] E. B. Sasson, A. Chiesa, C. Garman, M. Green, I. Miers, E. Tromer, and M. Virza, "Zerocash: Decentralized anonymous payments from bitcoin," in *2014 IEEE Symposium on Security and Privacy*. IEEE, 2014, pp. 459–474.

[66] S. P. Jaikar, R. C. Maheshwar, S. P. Mamadapure, and A. A. Bhosle, "Secure data distribution using secret splitting over cloud," *Global Journal of Computer Science and Technology*, 2017.

[67] M. M. Al Assaf, X. Jiang, M. R. Abid, and X. Qin, "Eco-storage: A hybrid storage system with energy-efficient informed prefetching," *Journal of Signal Processing Systems*, vol. 72, no. 3, pp. 165–180, 2013.

[68] S. Ali, G. Wang, M. Z. A. Bhuiyan, and H. Jiang, "Secure data provenance in cloud-centric internet of things via blockchain smart contracts," in *2018 IEEE SmartWorld, Ubiquitous Intelligence & Computing, Advanced & Trusted Computing, Scalable Computing & Communications, Cloud & Big Data Computing, Internet of People and Smart City Innovation (SmartWorld/SCALCOM/UIC/ATC/CBDCom/IOP/SCI)*. IEEE, 2018, pp. 991–998.

[69] X. Liang, S. Shetty, D. Tosh, C. Kamhoua, K. Kwiat, and L. Njilla, "Provchain: A blockchain-based data provenance architecture in cloud environment with enhanced privacy and availability," in *2017 17th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing (CCGRID)*. IEEE, 2017, pp. 468–477.

[70] N. Gogerty and J. Zitoli, "Deko–currency proposal using a portfolio of electricity linked assets," *Available at SSRN 1802166*, 2011.

[71] J. Mao, T. Bhattacharya, X. Peng, T. Cao, and X. Qin, "Modeling energy consumption of virtual machines in dvfs-enabled cloud data centers," in *2020 39th IEEE International Performance Computing and Communications Conference*. IEEE, 2020.