DS-06-2017: Cybersecurity PPP: Cryptography

PRIViLEDGE
Privacy-Enhancing Cryptography in Distributed Ledgers

D6.4 – Second Scientific & Research Impact Measurement

Due date of deliverable: 30 June 2021
Actual submission date: 29 June 2021

Grant agreement number: 780477
Start date of project: 1 January 2018
Revision 1.0

Lead contractor: Guardtime AS (GT)
Duration: 42 months

Project funded by the European Commission within the EU Framework Programme for Research and Innovation HORIZON 2020	
Dissemination Level	
PU = Public, fully open	X
CO = Confidential, restricted under conditions set out in the Grant Agreement	
CI = Classified, information as referred to in Commission Decision 2001/844/EC	

D6.4

Second Scientific & Research Impact Measurement

Editor
Kaoutar Elkhiyaoui (IBM)

Contributors
Ahto Truu (GT)
Markulf Kohlweiss (UEDIN)
Toon Segers (TUE)
Ivan Visconti (UNISA)
Sven Heiberg (SCCEIV)
Panos Louridas (GRNET)
Nikos Voutsinas (GUNET)
Nikos Karagiannidis (I.O.Research)

Reviewers
Markulf Kohlweiss (UEDIN)
Daniele Friolo (UNISA)

29 June 2021
Revision 1.0

The work described in this document has been conducted within the project PRIViLEDGE, started in January 2018. This project has received funding from the European Union’s Horizon 2020 research and innovation programme under grant agreement No. 780477.

The opinions expressed and arguments employed herein do not necessarily reflect the official views of the European Commission.

Contents

Executive Summary

As a Research and Innovation Action, PRIViLEDGE produced several research results that advance the state of the art in the areas of distributed ledger technology and security, and developed a number of toolkits and prototypes that demonstrate the utility of the conducted research in an industrial context.

The present deliverable analyses the impact of PRIViLEDGE's outputs and discusses their exploitability and commercialization potential. We recall that the project yielded 42 publications at scientific conferences, several of which are considered top-tier. Many of these publications gathered a respectable number of citations, indicating a healthy adoption by the academic community. Beyond their academic impact, the project's publications informed the design and development of 6 toolkits and 4 prototypes that will be exploited by stakeholders in the project and beyond; in fact, 7 of the toolkits and prototypes are fully open-source and ready to be used by external parties.

1 Introduction

The areas of security and privacy are rich with ingenious techniques that ensure anonymity, confidentiality, verifiability and consensus. However, these techniques are rarely adopted in real-world systems, for reasons that vary from being overly complex, to lacking compatibility with existing infrastructure, to failing to address performance requirements. To bridge this gap between theory and practice, PRIViLEDGE has two primary objectives:

1. Advance the state of the art in the areas of:
    - Privacy-enhancing technologies to enable privacy-preserving transaction processing.
    - Privacy-preserving and accountable data publication to allow users to submit transactions anonymously while ensuring some level of attributability (if the need arises).
    - Secure multi-party computation to design more efficient and fair protocols that leverage distributed ledger technology as a building block.
    - Consensus algorithms to accommodate high-throughput applications without sacrificing security or governance guarantees. We focus in particular on proof of stake and federated consensus methods.
    - Quantum-safe security protocols to protect the integrity of distributed ledgers well into the post-quantum era.
    - Decentralized and secure software updates to ensure that all nodes in the blockchain execute the same version of the code (i.e. no forks).
2. Provide tools for effective exploitation of PRIViLEDGE's results in real operational environments, namely:
    - Toolkits that implement PRIViLEDGE's results and help assess their performance and practicality.
    - Prototypes that integrate the toolkits in real-world systems and validate their usability.

Pursuing the first objective resulted in 54 publications that covered a broad range of research areas: ranging from the practical, yet challenging ones – such as *efficient* consensus and secure decentralized software updates– to the more forward-looking, exemplified by post-quantum primitives and efficient protocols for zero-knowledge proofs and secure multi-party computation.

The second objective was a catalyst for the design and implementation of 6 toolkits and 4 prototypes (cf. Table 1 and 2 respectively). These toolkits and prototypes were developed in the aim of bringing the results of the research papers closer to market. In fact, 7 are fully open-source, and as such, can be picked up by the larger blockchain community and help gauge the relevance of the project's findings and inform the research roadmap of the partners.

Toolkit	**Partners**	**Description**	**Open Source**
Post-Quantum Secure Protocols for Distributed Ledgers	GT	It provides access to hash-then-publish time-stamping and hash-based digital signatures.	The hash-then-publish time-stamping is open-sourced at [1]
Anonymous Authentication for Hyperledger Fabric	IBM	It allows users to anonymously authenticate their Hyperledger Fabric transactions while supporting revocation and audit functionalities.	Open-sourced at [2]
Mir-BFT: High-throughput Consensus for Hyperledger Fabric	IBM	It introduces an efficient consensus algorithm to Hyperledger Fabric to accommodate high-throughput applications.	Open-sourced at [3]; a Hyperledger Labs project

Secure Multi-Party Computation on Ledgers	UNISA, TUE	This toolkit consists of two modules. The first generically allows libraries for two/multi-party computation to use a ledger as a communication channel instead of point-to-point connections. The second module adds support for verifiable multi-party computation and secure groups to the MPyC framework [4].	First module is open-sourced at [5]; second module is open-sourced at [6] and [7]
Zero-knowledge Proofs for Ledgers	GRNET, UEDIN	The toolkit, called SNARKY, is a Rust implementation of the Snarky Ceremonies protocol [68] over the BLS12-381 elliptic curve. It is the first experimental version of a production-grade library for running the protocol in real-life applications.	Open sourced at [8]
Privacy-Preserving Data Storage on Ledgers	GRNET	It is is a ledger-oriented solution for bridging data residing on a blockchain with those stored in an off-chain database. It is responsible for coordinating interaction between the backend ledger, storage, and the cryptographic library of an API service.	Open sourced at [9]

Table 1: PRIViLEDGE's Toolkits

Prototype	**Partners**	**Description**	**Open Source**
Tiviledge: Verifiable Online Voting	SCCEIV	It allows universal auditability with everlasting privacy on HyperLedger Fabric blockchain with applications to online voting.	
Prototype Application for Health Insurance	GT	This is an internal research prototype for assessing feasibility of verifiable multi-party computation for privacy-preserving reports on detailed medical records.	
University Diploma Record Ledger	GRNET	This is a research prototype that implements the DIPLOMATA protocol described in Chapter 4 of [73].	Open sourced at [10]
Decentralized Software Updates for Stake-based Ledgers	I.O.Research	This is a research prototype that enables decentralized software updates in public stake-based blockchains in a way that is tolerant to chain splits and security attacks.	Open-sourced at [11]

Table 2: PRIViLEDGE's Prototypes

In light of these results, we measure the project impact with respect to three metrics: *academic performance*, *results exploitability* and *commercialization potential*. *Academic performance* refers to the papers produced within the project and their influence; influence is quantified by the number of citations, quality of targeted publication venues and follow-up work or collaborations that the papers have triggered. *Results exploitability*, on the other hand, refers to how the project findings advance the state of the art and how the developed toolkits address the partners' needs. For example, a toolkit can enhance the performances of an existing system, or help these systems comply with GDPR or become more environment-friendly. Finally, *commercialization potential* refers to the ability to generate external interest in the developed technologies: this describes how the project's results differentiate the partners' offerings. For example, the project's results can help the partners meet existing (or projected) market or regulatory demands, develop new innovative products through which the partners can boost their brand recognition or improve their customer engagement and retention.

The deliverable is organized as follows: Section 2 lists the research publications produced by the partners and discusses their impact. Section 3 discusses the exploitability of the project's results (i.e. publications, toolkits and prototypes) and their commercialization potential. Section 4 concludes the deliverable.

2 Academic Performance

Paper	Partners	Publication Date	Venue	Number of Citations	Follow-up
[57]	UEDIN	August 2018	CRYPTO	85	Used in [62] and in updatable and/or universal SNARK proof systems such as Plonk [12] and Marlin [13]. Ideas from this paper are also used in toolkit [8]
[65]	UEDIN, UT, I.O.Research, SCCEIV	September 2018	SCN	27	
[50]	UEDIN	September 2018	SCN	45	
[28]	IBM	September 2018	ESORICS	55	Informs discussions on cross-channel transactions in Hyperledger Fabric
[33]	UEDIN, I.O.Research	October 2018	CCS	150	Subsequent research by UEDIN, I.O.Research
[58]	SCCEIV, UT	October 2018	e-Vote ID	16	
[34]	UNISA	December 2018	ASIACRYPT	6	
[26]	TUE	March 2019	CT-RSA	3	
[75]	UNISA	April 2019	PKC	8	Subsequent research by UNISA
[59]	UNISA	April 2019	C2SI	–	
[64]	UEDIN, I.O.Research	May 2019	S&P	36	
[56]	UEDIN, I.O.Research	May 2019	S&P	62	
[23]	UT	June 2019	ACNS	6	Used in [24]
[24]	UT	July 2019	AFRICACRYPT	14	
[36]	UT	July 2019	AFRICACRYPT	12	

[30]	UT	September 2019	DPM/CBT	13	
[71]	UEDIN	November 2019	CCS	108	Used in [62] and in updatable and/or universal SNARK proof systems such as Plonk [12] and Marlin [13]. Ideas from this paper are also used in toolkit [8]
[35]	UNISA	December 2019	ASIACRYPT	3	
[53]	GT	January 2020	CPP	4	
[54]	UEDIN	February 2020	CT-RSA	80	
[27]	UNISA	May 2020	EUROCRYPT	4	
[55]	UNISA	September 2020	SCN	–	
[51]	UEDIN	September 2020	SCN	6	
[48]	UEDIN, I.O.Research	September 2020	ESORICS	4	Subsequent research by I.O.Research
[77]	GT	September 2020	BPM	–	
[22]	UEDIN, I.O.Research	October 2020	ACNS	3	
[43]	TUE	October 2020	ACNS	1	
[29]	IBM	October 2020	AFT	8	
[74]	UEDIN	November 2020	TCC	17	
[38]	I.O.Research, UEDIN	December 2020	ASIACRYPT	1	
[31]	UNISA	February 2021	CoronaDef (NDSS)	36	
[41]	UNISA	March 2021	FC	–	Ideas from this paper are also used in [5]
[60]	UEDIN, I.O.Research	March 2021	FC	–	Subsequent research by I.O.Research
[37]	UEDIN, I.O.Research	March 2021	FC	9	
[62]	UEDIN, I.O.Research	March 2021	FC	5	
[76]	UNISA	May 2021	PKC	1	
[52]	UEDIN	May 2021	PKC	2	
[32]	UNISA	June 2021	ACNS	–	
[61]	UEDIN, I.O.Research	June 2021	CSF	7	
[39]	TUE	July 2021	CSCML	5	
[63]	UEDIN, I.O.Research	August 2021	CRYPTO	2	
[47]	UEDIN	October 2021	EUROCRYPT	–	

Table 3: Peer-reviewed Papers: Citations and follow-up work. Numbers are from Google Scholar.

The most objectively measurable impact of a research project is academic publications and their uptake by the academic community. This section analyzes the impact of the research results produced during the course of PRIViLEDGE.

The venues targeted by PRIViLEDGE vary in terms of size and topics. Some venues cover broad areas of security research, such as ACM CCS, IEEE S&P, or ESORICS, whereas others are clearly focused on cryptography such as EUROCRYPT, CRYPTO or ASIACRYPT, or on specific use-cases such as e-voting (E-VoteID) or contact tracing during the Covid-19 pandemic (CoronaDef). Each type of venue has its advantages: the wider the reach of the venue, the easier it is to publicize the project's findings; specialized venues however allow for focused discussions with participants who are more likely to follow up on the presented solutions.

Table 3 depicts the list of peer-reviewed papers produced by the project partners. 42 papers have been published, with 12 being at some of the top 10 conferences in security and cryptography (cf. Table 4 and Figure 1b). The reputation of a venue is defined in relation to its impact factor, as expressed by *h5-index* and *h5-median* (numbers from Google Scholar). Taking *h5-index* as a measure of potential impact in the next 5 years, we see that 15 papers can garner more than 40 citations.

Venue	**Ranking in Top20**	**Number of Publications**	**h5-index**	**h5-median**
CCS	1	2	88	140
S&P	4	2	74	142
EUROCRYPT	6	2	61	89
CRYPTO	9	2	52	87
FC	10	4	46	74
ASIACRYPT	11	3	42	61
TCC	13	1	38	58
ESORICS	18	2	34	43
PKC	–	3	29	39
CSF	–	1	29	38
CT-RSA	–	2	26	45
ACNS	–	4	21	28
DPM/CBT	–	1	NA	NA
AFT	–	1	NA	NA
CPP	–	1	NA	NA
C2SI	–	1	NA	NA
CoronaDef	–	1	NA	NA
AFRICACRYPT	–	2	NA	NA
SCN	–	4	NA	NA
BPM	–	1	NA	NA
CSCML	–	1	NA	NA

Table 4: Impact of Targeted Venues. h5-index and h5-median from Google Scholar.

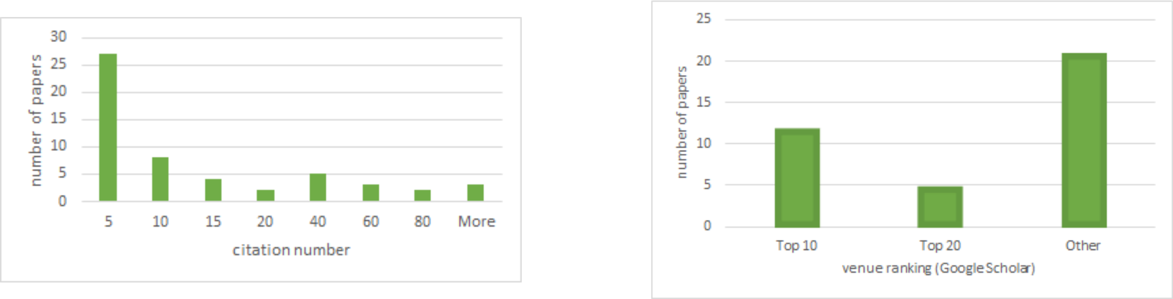

(a) Citation Number Histogram

(b) Venue Ranking Histogram

Figure 1

Looking at the citation numbers of PRIViLEDGE's papers (cf. Figure 1a) – including the technical reports, we see that 13 papers have a citation number of over 20. Three papers have a citation number larger than 80, whereas one paper has the impressive citation number of 150. These numbers show a healthy uptake by the academic community and bode well for the more recently-produced papers.

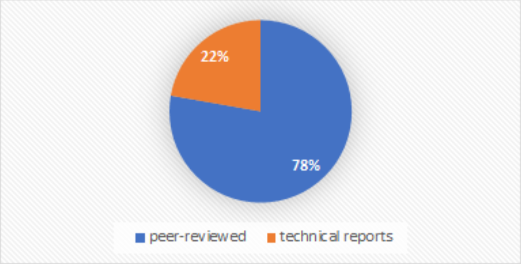

Figure 2: Peer-reviewed Papers

Paper	Publication Date	Number of Citations	Follow-up
[42]	July 2018	5	
[25]	September 2018	15	
[67]	October 2018	41	Extended in [56]
[66]	October 2018	21	
[45]	January 2019	10	Used in [29]
[70]	March 2019	–	
[69]	March 2019	3	
[44]	June2019	–	
[78]	June 2019	32	Implemented as toolkit [3]
[40]	September 2019	4	Implemented as toolkit [2]
[49]	April 2020	–	
[46]	October 2020	–	

Table 5: Technical Reports: Citations and follow-up work. Numbers are from Google Scholar.

Table 5 lists the technical reports, which to be yet peer-reviewed. Out of 54 papers, only 12 are not published yet. Figure 2 depicts the percentage of peer-reviewed papers and technical reports, which corresponds to 78% and 22% respectively. This demonstrates that most of the results produced by PRIViLEDGE are valued externally and believed to advance the state of the art in the areas of security, privacy and consensus.

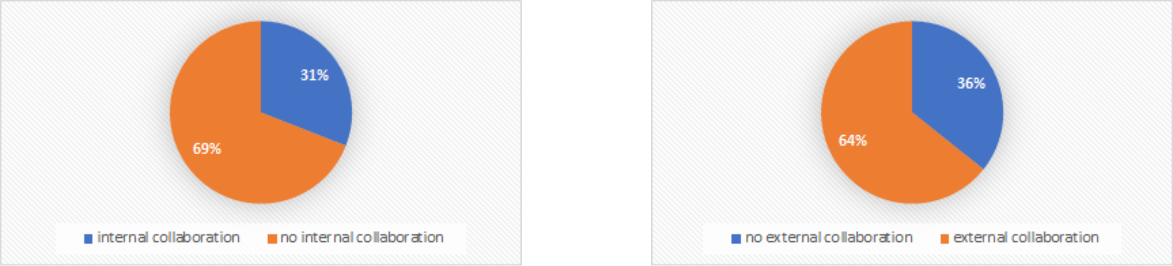

(a) Internal Collaboration

(b) External Collaboration

Figure 3

Figure 3a and Figure 3b reveal that out of the 42 peer-reviewed papers, 31% were the result of internal collaborations, whereas 64% the result of external collaborations. These numbers demonstrate that the project fostered a healthy amount of joint efforts and that its reach is not restricted to PRIViLEDGE's partners.

3 Exploitability and Commercialization

3.1 Guardtime AS

3.1.1 Results Exploitability

A major focus of Guardtime is quantum-safe distributed ledger platforms. To authenticate transactions, ledgers rely on asymmetric cryptography, often using primitives that are vulnerable to quantum attacks (most of them are based on the difficulty of computing discrete logarithms over elliptic curve points; a few are based on the difficulty of factoring large integers). Our toolkit [1] enables authentication and integrity protection based on hash functions which are believed to be resilient against quantum attacks. Hash function based time-stamping can also indemnify transactions authenticated with current technology against future quantum attacks.

The toolkit makes use of Guardtime's KSI time-stamping, which is a mature technology that has been available for bespoke integrations for a long time. The KSI component of the toolkit will make it easier to integrate KSI time-stamping into distributed ledger solutions based on the Hyperledger Fabric platform. The BLT component of the toolkit will enable authentication using hash-based signatures on Hyperledger Fabric. The toolkit can also serve as an example for similar integrations into other distributed ledger platforms.

Another focus of Guardtime is outcome-based contracting (OBC), which despite seeing growth in the last few years, still faces hurdles like expensive technology, outmoded infrastructure and regulations that get in the way. The shift to OBC, while beneficial to patients' health outcomes, has also been hindered by the tension between trust and privacy requirements. Currently, to prove the medical care goals have been achieved, the accountable care organization (ACO) would need to give insurers access to not only the list of treatments performed, but also to detailed test results. The patients, however, may be unwilling to consent to sharing all the details with the insurers. There's the option that ACO could report only the aggregate numbers over sufficiently large groups of patients, which would not violate any one patient's privacy, but then the insurers are reluctant to trust that the reported values are correct. From a data sharing viewpoint, this inherent conflict is the primary hurdle that needs to be overcome in order to pave the way to personalized and more effective health-care. A way to relieve this tension is to post to a shared ledger cryptographic commitments of patient records and provide for each report a proof that it is consistent with the posted commitments.

Guardtime's prototype combines secure multi-party computation (MPC) among the ACO members with zero-knowledge proofs that enable the insurers to verify the correctness of the reports without leaking the details of individual patients. The prototype is based on the verifiable MPC extensions added to the MPyC framework as part of the PRIViLEDGE project. Showing the possibility of such privacy-preserving reporting, as the prototype does, will encourage wider deployment of the OBC model and thus advance the efficiency of medical insurance in particular and the healthcare sector in general.

3.1.2 Commercialization Potential

The cryptographic tools developed in the project support creating new value to Guardtime's product portfolio in several privacy-conscious verticals. Guardtime expects the project's results to strengthen its offerings in the health insurance sector. In particular, the ability to provide verifiable reports on healthcare outcomes without disclosing the underlying patient records would be key in enabling wider use of outcome-based contracting in the sector. What is more, the cryptographic techniques developed for the prototype are universal and thus might become applicable to various other future solutions (e.g. in the financial sector).

Our RWD-E product, based on auditable multi-party computation technology, is already in different stages of deployment in Sweden (in cooperation with Roche and AstraZeneca) and in Spain (in cooperation with AstraZeneca), and there is clear commercial interest in solutions allowing better reporting based on actual patient data. Adding cryptographic verifiability to such reports makes the case much stronger.

There are still both legal and technical barriers to overcome. On the legal side, there's quite a bit of uncertainty how encrypted data, hashed data, or data shares passed around in MPC protocols will be treated under privacy regulations. Preliminary legal opinions we have obtained so far seem to indicate that such data will not be considered personally identifiable information under GDPR. However, additional local regulations need to be considered in each new market we want to enter. On the technical side, the need to operate on the data in encrypted form or represented as cryptographic shares significantly complicates detection and correction of any mistakes stemming from data quality in general or uniformity of representation in particular.

3.2 IBM Research – Zurich

3.2.1 Results Exploitability

The rapid rise of blockchain technology and the ensuing demand for high-performance and cheap decentralized applications brought to the fore the importance of efficient and secure consensus algorithms. Efficiency is measured in terms of *throughput* and *latency*, whereas security is captured by the properties of *fairness* and *liveness*. While existing solutions [72, 14] ensure security, they suffer from low throughput and/or high latency. Throughput is particularly critical for retail payments, IoT or track and trace applications; which are among the most prominent applications for blockchains, especially permissioned ones, such as Hyperledger-Fabric. Taking this into account, IBM plans to leverage the *MIR-BFT* toolkit [3] to enhance the performances of Hyperledger-Fabric. We recall that MIR-BFT achieves a whopping throughput of 60000 tps for a network of 100 nodes for transactions of size 500 bytes. To put these numbers in perspective, VISA, a fully-centralized system, is capable of performing 24000 transactions per second [15].

Another goal of IBM is to design blockchain solutions that minimize personal data exposure, in compliance with *GDPR*. Take for example, payment applications: every time a user performs a payment, they are required to show that they are authorized to make the payment. If implemented naively, payment systems will expose the identity of the user whenever they initiate a payment. Pseudonymous solutions such as Bitcoin can to some extent protect user identity, but they fail to prevent transaction/user linkability. Additionally, permissioned blockchains, such as Hyperledger-Fabric, require users to hold long-term identities as a way to ensure accountability if the need ever arises. To marry accountability and privacy, Hyperledger-Fabric uses the Idemix technology that enables users to submit transactions anonymously using their long-term identities. However, current Idemix is only suitable for single-issuer settings and supports neither revocation nor transaction auditability [40]. Fortunately, the toolkit for *Anonymous Authentication* [2] mitigates these limitations with almost no additional cost. With straightforward integration with Hyperledger-Fabric and potential applications to decentralized identity and transaction monitoring, this toolkit is well-positioned to play an important role in IBM's future blockchain solutions.

3.2.2 Commercialization Potential

The decline of cash usage and the prevalence of private payment systems have given rise to a world-wide push for central-bank digital currency (CBDC) [16, 17]. CBDC can be either centralized or federated, with centralized solutions enjoying high-performances and federated ones being more secure and resilient. Without high-throughput, permissioned blockchains cannot compete with fully-centralized systems. MIR-BFT luckily demonstrates that with the right consensus algorithm, throughput may no longer be an issue for permissioned systems. Privacy is also critical for CBDC systems: linking payment transactions of individuals is a serious privacy threat that can obstruct the adoption of blockchain-based CBDC solutions. The toolkit for anonymous credentials combined with privacy-preserving decentralized payment systems [29] provide a ready-for-use solution that CBDC systems can benefit from.

Another emerging application for blockchain is supply chain track and trace. This application in particular is greedy in terms of throughput (e.g. pharmaceutical supply chains in the US deal with 33 to 55 millions transactions a day [18]). Without high-throughput, blockchain cannot be used effectively. Alternatives such as Layer2 solutions can be leveraged to artificially increase throughput, but this comes at the cost of security. Increasing the throughput Hyperledger Fabric using MIR-BFT is a first step towards accommodating supply-chain applications, for which the demand is predicted to steadily grow [19].

3.3 University of Edinburgh

3.3.1 Results Exploitability

One of the areas of interest of UEDIN is *software updates in blockchain systems*. Traditionally, these updates have been handled in an ad-hoc, centralized manner: somebody, often a trusted authority, or the original author of the software, provides a new version of the software, and users download and install it from that authority's website. Even if the system follows an open source software development model, and therefore an update can potentially be implemented by anyone, the final decision of accepting, or rejecting, a new piece of code is always taken by the main maintainer(s) of the system, who essentially constitutes a central authority.

UEDIN put forth a novel mechanism to securely realize decentralized software updates that focuses on public stake-based blockchain systems. More precisely, we introduce in these systems the capability to take stake-based decisions based on: a) the software update priority, b) the correctness of the new code, c) the maintenance of the code base and d) the authenticity and safety of the downloaded software. To the best of our knowledge, there is no related work that addresses the problem of the decentralization of software updates in the context of blockchain systems in a holistic manner, i.e., one that takes into consideration all phases in the lifecycle of a software update. Bitcoin, Bitcoin Cash, Ethereum and Zcash use a "social governance" scheme, in which decisions on update proposals is reached through discussions on social media. This type of informal guidance is too unstable and prone to chain splits, or prone to becoming de-facto centralized. There exist blockchain systems that adopt a decentralized governance scheme, in which the priorities, as well as the funding of update proposals is voted on-chain as part of a maintenance protocol. However, these proposals do not follow a holistic approach to the decentralization software updates problem. Instead, their focus is merely on the ideation phase in the lifecycle of a software update, where an update proposal is born as an idea, and the community is called to accept if it will be funded or rejected. These solutions do not deal with the residual phases in this lifecycle. In contrast, our work formalizes the problem of *decentralized software updates* and introduces the first solution that activates the changes on the blockchain without risking a chain split. The outcome of this research is the basis of the Decentralized Software Updates for Stake-based Ledgers prototype [11] developed by I.O.Research.

Another area of interest are *Structured Reference String Updates* for blockchain systems that use SNARK-based zero-knowledge proofs to achieve privacy-preserving transaction validation. The proofs generally rely on a trusted setup in which a *common reference string* is generated. The soundness of the proofs hinges upon the trustworthiness of the parties partaking in the setup. MPC-based solutions relax the trust assumptions by ensuring soundness in the presence of one single honest participant. In contrast, we guarantee update knowledge soundness: an adversary cannot forge a SNARK proof even if they participate in the setup. We propose a new security framework in which the common reference string is updated by a large number of parties, e.g. the blockchain miners, in such a way that participating in the updates does not give the adversary any significant advantage as long as there is at least one honest update.

In particular, we prove the security of the Groth16 SNARK with a setup ceremony of zcash in our new security framework. We intentionally try not to change the original MPC ceremony protocol too much so that our security proof would apply to protocols already used in practice. Security is proven with respect to algebraic adversaries in the random oracle model. We require a single party to be honest in each phase of the protocol in order to guarantee update knowledge soundness and subversion zero-knowledge hold unconditionally. Unlike zcash's ceremony, our security proof does not rely on the use of a random beacon. However, our security proof does apply to protocols that have been implemented using a (potentially insecure) random beacon because the beacon can just be treated as an additional malicious party. We see this as an important security validation of real-life protocols that cryptocurrencies depend on.

The research work on these two areas contributed to the Blockchains and Distributed Ledgers course by

Aggelos Kiayias at UEDIN [20] and shaped the PhD of Mikhail Volkov. It also allowed us to invite prominent researchers to our Blockchain Technology Lab Security & Privacy Seminars [21].

3.3.2 Commercialization Potential

In addition to their theoretical interest, UEDIN's contributions have a practical impact. The work on secure decentralized software update is implemented as part of the I.O.Research open-source prototype by the same name. Given the importance of decentralized governance in open blockchain systems, we anticipate that decentralized software updates is a feature that will definitely differentiate I.O.Research's offering. On the other hand, the proposed security framework *Structured Reference String Updates* will help existing blockchain systems support privacy-preserving transaction validation with relaxed security assumptions, fostering hence, more trust and transparency.

3.4 Technical University of Eindhoven

3.4.1 Results Exploitability

The Technical University of Eindhoven sought results that enhance its portfolio in privacy-protecting cryptographic protocols and help extend its MPyC framework. A focus of TUE's research is the interplay between secure multiparty computation (MPC) and blockchain technology (e.g., by committing the inputs and outputs for secure computations using blockchain) and its applications to real-world use-cases.

In particular, TUE's work yielded solutions for secure computation of the Moore-Penrose Pseudo-inverse (ACNS 2020), secure comparison of medium-sized integers (RSA Conference 2019) and secure Ridge Regression (accepted for CSCML 2021). All of these results have direct applications to machine learning and neural networks.

Moreover, TUE's research on verifiable MPC using new Zero Knowledge Proof primitives (IEEE S&P 2018) and (CRYPTO 2020) led to the design of a conceptual verifiable MPC scheme and a practical construction that takes as inputs encryptions on a bulletin board and permits encrypted outputs. The main advantage of the developed solution is that it has a reusable and non-trusted setup based on standard cryptographic assumptions, making it well-suited for security-sensitive distributed applications. This solution is implemented as part of the second module of the toolkit on ledger-oriented two/multi-party computation, and is made available at [6] to third parties wishing to develop MPC-based applications with public verifiability of the outputs. This public verifiability property is very relevant in e-voting systems, for example.

TUE also developed a Secure Group scheme that simplifies the implementation of group-based cryptographic operations in MPC. Secure Groups implement finite groups as oblivious data structures, ensuring that no information can be inferred about the values of the group elements after a sequence of operations. One of TUE's goals is efficient and constant-time secure protocols; an example of which is TUE's extended GCD of two secret-shared integers that adapts constant-time gcd algorithms by Bernstein and Yang (CHES 2019) from the p-adic setting to the finite field setting. This new protocol is of independent interest and particularly useful for threshold cryptography applications. The secure group scheme and extended gcd implementations are available at [7], and can be readily used by distributed applications relying on threshold trust for privacy and fault-tolerance.

These promising research results and the developed prototypes are expected to further strengthen the position of TUE as a knowledge partner in blockchain technology, at the national level within the Netherlands. In fact, TUE's advice on privacy technology is already sought after by public agencies, such as Statistics Netherlands (Centraal Bureau voor de Statistiek). They are also expected to shape advanced courses on Applied Cryptography and Cryptographic Protocols taught at the TUE graduate school (master level), and keep them up-to-date with modern technology.

Finally, PRIViLEDGE's results will guide TUE's future work and collaborations. More specifically, we plan to extend Secure Groups and verifiable MPC to support a wider spectrum of arithmetic operators and problem statements. We also plan to investigate how verifiable MPC can be leveraged in blockchain systems to support private outsourcing and private smart contracts. This focuses on efficient on-chain verification of private and expensive off-chain work that often involves multiple *mistrustful* parties.

3.4.2 Commercialization Potential

The MPyC framework (155 Stars on GitHub) and its enhancement with verifiable MPC and Secure Groups will most likely inform the development of third-party MPC engines. As the number of MPC-based applications increases, the demand for ready-to-use blueprints will also increase. To accommodate such a demand, TUE researchers launched (during PRIViLEDGE) *Roseman Labs*, an MPC company that aims to valorize the research output of the project in a commercial context, bringing it closer to the market.

3.5 University of Salerno

3.5.1 Results Exploitability

A primary goal of UNISA is to leverage the competence and expertise acquired during PRIViLEDGE to produce significant scientific results on blockchain and its applications. As a matter of fact, the research group at University of Salerno was successful to publish several papers that advance the state of the art of blockchain technology on both the theoretical and practical fronts. For example, a paper published in FC 2021 shows how to design more efficient smart contracts that retain security even in case of forks of the underlying blockchains. The results of this paper guided the design of the first module of the toolkit on ledger-oriented two/multi-party computation [5]. This toolkit is aimed for researchers to quickly test prototypes of their blockchain-based applications that rely on secure multi-party computation.

Covid-19 pandemic triggered research on vulnerabilities of contact tracing applications. This yielded two publications, one in ACNS 2021 and another in CoronaDef 2021. Aside from their timeliness, these publications have two major contributions: (1) they show how smart contracts can be exploited to pollute the most used systems (i.e., GAEN based) by generating fake at-risk exposures; and (2) propose mitigation mechanisms that use the blockchains as a decentralized bulletin board to provide transparency and foster trust in such systems.

Additionally, two PKC publications (2019, 2021) show how blockchain systems can potentially serve as a trusted infrastructure to construct cryptographic protocols and primitives like zero-knowledge proofs. This is particularly useful in decentralized payment systems that heavily rely on zero-knowledge proofs to meet privacy requirements.

These published results have contributed to the visibility of UNISA's research group both in academia and industry. In fact, UNISA established collaborations with industrial partners interested in the use of confidential data in blockchain, and joined the *Algorand Foundation's Global University Program*, which includes only few selected universities in the world.

Furthermore, the expertise developed during the course of PRIViLEDGE positively affected the content of cybersecurity classes and shaped the research activities of master and PhD students.

3.5.2 Commercialization Potential

The ledger-oriented two/multi-party computation toolkit can be a starting point to obtain a fully-versatile library with production-level code intended for applications that require both transparency and confidentiality. Notice that by being built on top of blockchain, the toolkit offers transparency by design, whereas the primitive for secure multi-party computation provably guarantees confidentiality.

Additionally, the applicability of the research output can be widened by investigating new research areas such as: (i) federated self-regulated privacy-preserving machine learning; (ii) on-chain secure computation in the presence of smart contracts with limited power (e.g., Algorand); (iii) protecting business secrets in blockchain-enabled supply chains; (iv) compact zero-knowledge proofs with stable/standard assumptions. With the growing demand for privacy, accountability and decentralization, we anticipate that solutions for these open problems will not only have a theoretical interest, but also find direct applications in the real-world.

While UNISA does not plan to commercialize the outcomes of PRIViLEDGE, these outcomes are available to third parties, which can leverage them to build new products or improve existing ones.

3.6 Smartmatic-Cybernetica Centre of Excellence for Internet Voting

3.6.1 Results Exploitability

SCCEIV has developed a prototype for online voting system called Tiviledge. Tiviledge approaches the auditability of online elections differently from the state of the art. Tiviledge focuses on ensuring everlasting privacy of the data made available to the auditors. To that end, Tiviledge uses Hyperledger Fabric technology to distribute data to the auditors in a controlled manner while leaving immutable and privacy-preserving logs in the shared ledger.

The prototype makes it possible for election organizers to directly involve external parties in the storage of election audit material – privacy preserving commitments. This would be difficult with existing technologies, as it would require distributing the encrypted votes to the auditors as well, raising questions of the long-term secrecy of the ballot.

3.6.2 Commercialization Potential

The work on the Tiviledge prototype has shown, what aspects must be considered to securely audit election data. The proposed architecture could be used already in the near future, without vastly changing the assumptions about the organization of the election. However, in order to use the prototype the way it's actually intended, further reconsideration of the role of the external auditors in the elections is required. The way elections are organized today relies heavily on a central organization. Tiviledge will work best in environments where multiple organizations use the platform for their elections, thus providing auditing services for each other and also benefiting from the elections themselves.

3.7 GRNET

3.7.1 Results Exploitability

There is a widespread interest in providing privacy-preserving means to validate certifications with blockchain technology. The DIPLOMATA protocol fills this need, as the protocol guarantees that the intended recipient will be able to verify the credential proofs. Besides the development of the protocol, GRNET also worked on its implementation, which demonstrates how it can be used in practice in actual blockchain platforms, and how it could inter-operate with university record systems. GRNET is actively involved in the European Blockchain Services Infrastructure (EBSI) and is following the design of the Diplomas Use Case under development. In the coming months, as the Diplomas Use Case design progresses, GRNET will seek to exploit the DIPLOMATA implementation for the purposes of inter-operating with diplomas certification at a cross-country level.

Apart from the implementation of the actual protocol, considerable work went into building the user interface and bridging with storage in order to create an integrated service. This technological know-how, i.e., how to extend beyond a cryptographic protocol to a real working application and service leveraging distributed ledgers, can be readily exploited in other projects and services.

While the SNARKY toolkit is certainly valuable as a standalone implementation of the Snarky Ceremonies protocol, it was implemented with real-world applications in mind. In particular, considerable effort was dedicated to endow the toolkit with efficiency and speed way above a research prototype implementation. This paves the road for direct exploitation of the implementation in applications; moreover, as with DIPLOMATA, the engineering know-how gained from developing this toolkit can fertilize future protocol implementations.

Similarly, the storage toolkit shows how an abstraction bridging application logic, cryptography, distributed ledgers, and off-ledger storage, can be engineered in practice to deliver an actual working product, whose outcome can be readily utilized by other projects.

3.7.2 Commercialization Potential

There are many initiatives and efforts around the world for blockchain-based diplomas validation; it is clear that the technologies developed by GRNET could be leveraged by third parties (note that, due to its status as a provider of services to the Greek academic and research community and operating under the auspices of the Ministry of Digital Governance, GRNET itself may not be able to directly commercialize DIPLOMATA). The same applies to SNARKY and the storage toolkit: although GRNET is not a commercial organization, its outcomes are offered as open source, and could be commercialized by third-parties.

3.8 GUNET

3.8.1 Results Exploitability

The topic of verifiable digital credentials in the higher education landscape has been gaining attraction among technology providers and Higher Education Information management. During the course of PRIViLEDGE, GUnet contributed the design of the interface between Student Information Systems and the diplomas prototype. GUnet's efforts accordingly focused on the digital representation of the diploma object and the protocol/API for retrieving that information from authoritative sources. The goal was to facilitate the credentials portability across the chain and enhance the semantic interoperability. The outcome of this effort has already propelled the basis for the eDiplomas initiative in Greece and helped GUnet to achieve wider adoption of core design concepts among Greek Universities.

3.8.2 Commercialization Potential

Considering that GUnet is a non-profit company, there are no commercialization plans. However, the PRIViLEDGE results will be used to enhance GUnet's service portfolio for its members.

3.9 I.O.Research

3.9.1 Results Exploitability

The basic driver for our work in PRIViLEDGE has been that today software updates for public blockchains are neither secure nor decentralized. So we aimed to design a software update mechanism that would: a) decentralize software updates in their *whole lifecycle*, i.e., end-to-end (from ideation to activation) and not only partially, b) enable a secure activation of changes by defining what security means in this context and propose specific secure activation protocols and c) realize our ideas into a prototype implementation, and most importantly, integrate this with the Cardano node. In particular, I.O.Research foresees that the results from PRIViLEDGE will help the Cardano blockchain become a community-governed and self-sustaining blockchain. Decentralized governance is the key factor to achieve this goal, a core component of which is decentralized software updates. The research results on decentralized software updates developed within PRIViLEDGE, as well as the corresponding prototype implementation [11], will greatly influence the future versions of Cardano enabling decentralized governance.

3.9.2 Commercialization Potential

I.O.Research expects that given the above-described research outcomes, their implementation in the Cardano blockchain will provide a competitive advantage to the project, increasing hence, its commercial success. This expectation is based on the current understanding that decentralized governance is a major milestone in the Cardano blockchain product roadmap, which has been announced from the early beginning and is greatly anticipated by the Cardano community. We expect that this will significantly increase the value of the Cardano blockchain, since it will cover all the aspects of decision making in Cardano: from funding and initial ideation to the very last stage where changes are activated on the mainnet. This holistic approach to decentralized governance is often missing in existing projects in the cryptocurrency space and we strongly believe that it will be appreciated greatly by the Cardano community.

4 Conclusion

PRIViLEDGE produced a healthy number of scientific publications (42 to be exact), with 15 papers gathering more than 20 citations. The project sparked a number of external and internal collaborations: 27 papers resulted from external collaborations whereas 13 from internal collaborations.

The toolkits and prototypes developed during the project leveraged the research results to address a broad range of requirements (e.g. high-throughput consensus, post-quantum security, decentralized governance) and use-cases (e.g. healthcare, decentralized software updates, diploma certification). These more practical outcomes illustrate the applicability of PRIViLEDGE's research in real-world settings, and they are a starting point to build new innovative products, or set off further research endeavors that focus on performance, privacy, new applications or all three.

References

[1] https://github.com/guardtime/ksi-hlf.

[2] https://github.com/IBM/dac-lib.

[3] https://github.com/hyperledger-labs/mirbft.

[4] https://github.com/lschoe/mpyc.

[5] https://github.com/danielefriolo/ledgerMPC.

[6] https://github.com/toonsegers/verifiable_mpc.

[7] https://github.com/toonsegers/sec_groups.

[8] https://github.com/grnet/snarky.

[9] https://github.com/grnet/db-chain-bridge.

[10] https://github.com/grnet/e-diplomata.

[11] https://github.com/input-output-hk/decentralized-software-updates.

[12] https://eprint.iacr.org/2019/953.

[13] https://eprint.iacr.org/2019/1047.

[14] https://github.com/ethereum/wiki/wiki/White-Paper.

[15] https://usa.visa.com/run-your-business/small-business-tools/retail.html.

[16] https://www.ecb.europa.eu/paym/digital_euro/html/index.en.html.

[17] https://www.riksbank.se/en-gb/payments--cash/e-krona/.

[18] https://hbr.org/2020/05/building-a-transparent-supply-chain.

[19] https://www2.deloitte.com/us/en/pages/operations/articles/blockchain-supply-chain-innovation.html.

[20] http://www.drps.ed.ac.uk/20-21/dpt/cxinfr11144.htm.

[21] https://www.ed.ac.uk/informatics/blockchain/events2/previous-events.

[22] Aydin Abadi, Michele Ciampi, Aggelos Kiayias, and Vassilis Zikas. Timed signatures and zero-knowledge proofs—timestamping in the blockchain era—. In Mauro Conti, Jianying Zhou, Emiliano Casalicchio, and Angelo Spognardi, editors, *Applied Cryptography and Network Security*, pages 335–354. Springer International Publishing, 2020.

[23] Behzad Abdolmaleki, Karim Baghery, Helger Lipmaa, Janno Siim, and Michal Zajac. DL-extractable UC-commitment schemes. In Robert Deng and Moti Yung, editors, *ACNS*, volume 11464 of *LNCS*, pages 385–405. Springer, 2019.

[24] Behzad Abdolmaleki, Karim Baghery, Helger Lipmaa, Janno Siim, and Michal Zajac. UC-secure CRS generation for SNARKs. In Johannes Buchmann, Abderrahmane Nitaj, and Tajjeeddine Rachidi, editors, *Africacrypt*, LNCS. Springer, 2019.

[25] Behzad Abdolmaleki, Helger Lipmaa, Janno Siim, and Michal Zajac. On QA-NIZK in the BPK model. Cryptology eprint archive, report 2018/877, September 2018.

[26] Mark Abspoel, Niek J. Bouman, Berry Schoenmakers, and Niels de Vreede. Fast secure comparison for medium-sized integers and its application in binarized neural networks. In Mitsuru Matsui, editor, *Topics in Cryptology – CT-RSA 2019*, pages 453–472. Springer International Publishing, 2019.

[27] Divesh Aggarwal, Maciej Obremski, João L. Ribeiro, Luisa Siniscalchi, and Ivan Visconti. How to extract useful randomness from unreliable sources. In Anne Canteaut and Yuval Ishai, editors, *Advances in Cryptology - EUROCRYPT 2020 - 39th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Zagreb, Croatia, May 10-14, 2020, Proceedings, Part I*, volume 12105 of *Lecture Notes in Computer Science*, pages 343–372. Springer, 2020.

[28] Elli Androulaki, Christian Cachin, Angelo De Caro, and Eleftherios Kokoris-Kogias. Channels: Horizontal scaling and confidentiality on permissioned blockchains. In Javier Lopez, Jianying Zhou, and Miguel Soriano, editors, *European Symposium on Research in Computer Security*, volume 11098 of *LNCS*, pages 111–131. Springer, 2018.

[29] Elli Androulaki, Jan Camenisch, Angelo De Caro, Maria Dubovitskaya, Kaoutar Elkhiyaoui, and Björn Tackmann. Privacy-preserving auditable token payments in a permissioned blockchain system. In *Proceedings of the 2nd ACM Conference on Advances in Financial Technologies*, AFT '20, page 255–267, New York, NY, USA, 2020. Association for Computing Machinery.

[30] Shahla Atapoor and Karim Baghery. Simulation extractability in groth's zk-snark. In Cristina Pérez-Solà, Guillermo Navarro-Arribas, Alex Biryukov, and Joaquin Garcia-Alfaro, editors, *Data Privacy Management, Cryptocurrencies and Blockchain Technology*, pages 336–354, Cham, 2019. Springer International Publishing.

[31] Gennaro Avitabile, Vincenzo Botta, Vincenzo Iovino, and Ivan Visconti. Towards defeating mass surveillance and sars-cov-2: The pronto-c2 fully decentralized automatic contact tracing system. Cryptology ePrint Archive, Report 2020/493, 2020. https://eprint.iacr.org/2020/493.

[32] Gennaro Avitabile, Daniele Friolo, and Ivan Visconti. Terrorist attacks for fake exposure notifications in contact tracing systems. Cryptology ePrint Archive, Report 2020/1150, 2020. https://eprint.iacr.org/2020/1150.

[33] Christian Badertscher, Peter Gaži, Aggelos Kiayias, Alexander Russell, and Vassilis Zikas. Ouroboros genesis: Composable proof-of-stake blockchains with dynamic availability. In *CCS*. ACM, 2018.

[34] Saikrishna Badrinarayanan, Abhishek Jain, Rafail Ostrovsky, and Ivan Visconti. Non-interactive secure computation from one-way functions. In Thomas Peyrin and Steven Galbraith, editors, *Advances in Cryptology – ASIACRYPT*, volume 11274 of *LNCS*, pages 118–138. Springer, 2018.

[35] Saikrishna Badrinarayanan, Abhishek Jain, Rafail Ostrovsky, and Ivan Visconti. Uc-secure multiparty computation from one-way functions using stateless tokens. In Steven D. Galbraith and Shiho Moriai, editors, *Advances in Cryptology – ASIACRYPT 2019*, pages 577–605, Cham, 2019. Springer International Publishing.

[36] Karim Baghery. On the efficiency of privacy-preserving smart contract systems. In Johannes Buchmann, Abderrahmane Nitaj, and Tajjeeddine Rachidi, editors, *Africacrypt*, LNCS. Springer, 2019.

[37] Karim Baghery, Markulf Kohlweiss, Janno Siim, and Mikhail Volkhov. Another look at extraction and randomization of groth's zk-snark. Cryptology ePrint Archive, Report 2020/811, 2020. `https://eprint.iacr.org/2020/811`.

[38] Foteini Baldimtsi, Aggelos Kiayias, Thomas Zacharias, and Bingsheng Zhang. Crowd verifiable zero-knowledge and end-to-end verifiable multiparty computation. In Shiho Moriai and Huaxiong Wang, editors, *Advances in Cryptology – ASIACRYPT 2020*, pages 717–748. Springer International Publishing, 2020.

[39] Frank Blom, Niek J. Bouman, Berry Schoenmakers, and Niels de Vreede. Efficient secure ridge regression from randomized gaussian elimination. Cryptology ePrint Archive, Report 2019/773, 2019. `https://eprint.iacr.org/2019/773`.

[40] Dmytro Bogatov, Angelo De Caro, Kaoutar Elkhiyaoui, and Björn Tackmann. Anonymous transactions with revocation and auditing in hyperledger fabric. Cryptology ePrint Archive, Report 2019/1097, 2019. `https://eprint.iacr.org/2019/1097`.

[41] Vincenzo Botta, Daniele Friolo, Daniele Venturi, and Ivan Visconti. Shielded computations in smart contracts overcoming forks. Cryptology ePrint Archive, Report 2019/891, 2019. `https://eprint.iacr.org/2019/891`.

[42] Niek J. Bouman and Niels de Vreede. New protocols for secure linear algebra: Pivoting-free elimination and fast block-recursive matrix decomposition. Cryptology ePrint Archive, Report 2018/703, 2018. `https://eprint.iacr.org/2018/703`.

[43] Niek J. Bouman and Niels de Vreede. A practical approach to the secure computation of the moore–penrose pseudoinverse over the rationals. In Mauro Conti, Jianying Zhou, Emiliano Casalicchio, and Angelo Spognardi, editors, *Applied Cryptography and Network Security*, pages 398–417. Springer International Publishing, 2020.

[44] Christian Cachin and Björn Tackmann. Asymmetric distributed trust. *CoRR*, abs/1906.09314, 2019.

[45] Jan Camenisch, Manu Drijvers, and Björn Tackmann. Multi-protocol UC and its use for building modular and efficient protocols. Cryptology eprint archive, report 2019/065, January 2019.

[46] Michele Ciampi, Alexandru Cojocaru, Elham Kashefi, and Atul Mantri. Secure two-party quantum computation over classical channels. Cryptology ePrint Archive, Report 2020/1286, 2020. `https://eprint.iacr.org/2020/1286`.

[47] Michele Ciampi, Vipul Goyal, and Rafail Ostrovsky. Threshold garbled circuits and ad hoc secure computation. Cryptology ePrint Archive, Report 2021/308, 2021. `https://eprint.iacr.org/2021/308`.

[48] Michele Ciampi, Nikos Karayannidis, Aggelos Kiayias, and Dionysis Zindros. Updatable blockchains. In Liqun Chen, Ninghui Li, Kaitai Liang, and Steve Schneider, editors, *Computer Security – ESORICS 2020*, pages 590–609. Springer International Publishing, 2020.

[49] Michele Ciampi, Yun Lu, and Vassilis Zikas. Collusion-preserving computation without a mediator. Cryptology ePrint Archive, Report 2020/497, 2020. `https://eprint.iacr.org/2020/497`.

[50] Michele Ciampi and Claudio Orlandi. Combining private set-intersection with secure two-party computation. In Dario Catalano and Roberto De Prisco, editors, *International Conference on Security and Cryptography for Networks*, volume 11035 of *LNCS*, pages 464–482. Springer, 2018.

[51] Michele Ciampi, Roberto Parisella, and Daniele Venturi. On adaptive security of delayed-input sigma protocols and fiat-shamir nizks. In Clemente Galdi and Vladimir Kolesnikov, editors, *Security and Cryptography for Networks*, pages 670–690. Springer International Publishing, 2020.

[52] Michele Ciampi, Luisa Siniscalchi, and Hendrik Waldner. Multi-client functional encryption for separable functions. In Juan A. Garay, editor, *Public-Key Cryptography – PKC 2021*, pages 469–498. Springer International Publishing, 2021.

[53] Denis Firsov, Ahto Buldas, Ahto Truu, and Risto Laanoja. Verified security of blt signature scheme. In *Proceedings of the 9th ACM SIGPLAN International Conference on Certified Programs and Proofs*, CPP 2020, page 244–257, New York, NY, USA, 2020. Association for Computing Machinery.

[54] Juan Garay and Aggelos Kiayias. Sok: A consensus taxonomy in the blockchain era. In Stanislaw Jarecki, editor, *Topics in Cryptology – CT-RSA 2020*, pages 284–318. Springer International Publishing, 2020.

[55] Sanjam Garg, Xiao Liang, Omkant Pandey, and Ivan Visconti. Black-box constructions of bounded-concurrent secure computation. In Clemente Galdi and Vladimir Kolesnikov, editors, *Security and Cryptography for Networks*, pages 87–107. Springer International Publishing, 2020.

[56] Peter Gaži, Aggelos Kiayias, and Dionysis Zindros. Proof-of-stake sidechains. In *IEEE Symposium on Security and Privacy*. IEEE, 2019.

[57] Jens Groth, Markulf Kohlweiss, Mary Maller, Sarah Meiklejohn, and Ian Miers. Updatable and universal common reference strings with applications to zk-snarks. Cryptology ePrint Archive, Report 2018/280, 2018. `https://eprint.iacr.org/2018/280`.

[58] Sven Heiberg, Ivo Kubjas, Janno Siim, and Jan Willemson. On trade-offs of applying block chains for electronic voting bulletin boards. In *E-Vote ID*. Tallinna Tehnikaülikooli Raamatukogu Digikogu, 2018.

[59] Vincenzo Iovino and Ivan Visconti. Non-interactive zero knowledge proofs in the random oracle model. In Claude Carlet, Sylvain Guilley, Abderrahmane Nitaj, and El Mamoun Souidi, editors, *Codes, Cryptology and Information Security*, pages 118–141. Springer International Publishing, 2019.

[60] Dimitris Karakostas, Nikos Karayannidis, and Aggelos Kiayias. Efficient state management in distributed ledgers. Cryptology ePrint Archive, Report 2021/183, 2021. `https://eprint.iacr.org/2021/183`.

[61] Thomas Kerber, Aggelos Kiayias, and Markulf Kohlweiss. Kachina - foundations of private smart contracts. Cryptology ePrint Archive, Report 2020/543, 2020. `https://eprint.iacr.org/2020/543`.

[62] Thomas Kerber, Aggelos Kiayias, and Markulf Kohlweiss. Mining for privacy: How to bootstrap a snarky blockchain. Cryptology ePrint Archive, Report 2020/401, 2020. `https://eprint.iacr.org/2020/401`.

[63] Thomas Kerber, Aggelos Kiayias, and Markulf Kohlweiss. Composition with knowledge assumptions. Cryptology ePrint Archive, Report 2021/165, 2021. `https://eprint.iacr.org/2021/165`.

[64] Thomas Kerber, Aggelos Kiayias, Markulf Kohlweiss, and Vassilis Zikas. Ouroboros crypsinous: Privacy-preserving proof-of-stake. In *2019 IEEE Symposium on Security and Privacy (SP)*, pages 157–174, 2019.

[65] Aggelos Kiayias, Annabel Kuldmaa, Helger Lipmaa, and Janno Siim. On the security properties of e-voting bulletin boards. In Dario Catalano and Roberto De Prisco, editors, *International Conference on Security and Cryptography for Networks*, volume 11035 of *LNCS*, pages 505–523. Springer, 2018.

[66] Aggelos Kiayias and Alexander Russell. Ouroboros-BFT: A simple Byzantine fault tolerant consensus protocol. Cryptology eprint archive, report 2018/1049, October 2018.

[67] Aggelos Kiayias and Dionysis Zindros. Proof-of-work sidechains. Cryptology eprint archive, report 2018/1048, October 2018.

[68] Markulf Kohlweiss, Mary Maller, Janno Siim, and Mikhail Volkhov. Snarky ceremonies. Cryptology ePrint Archive, Report 2021/219, 2021. `https://eprint.iacr.org/2021/219`.

[69] Helger Lipmaa. Key-and-argument-updatable QA-NIZKs. Cryptology eprint archive, report 2019/333, March 2019.

[70] Helger Lipmaa. Simple yet efficient knowledge-sound and non-black-box any-simulation-extractable ZK-SNARKs. Cryptology ePrint Archive, Report 2019/612, 2019. `https://eprint.iacr.org/2019/612`.

[71] Mary Maller, Sean Bowe, Markulf Kohlweiss, and Sarah Meiklejohn. Sonic: Zero-knowledge snarks from linear-size universal and updatable structured reference strings. CCS '19, page 2111–2128, New York, NY, USA, 2019. Association for Computing Machinery.

[72] Satoshi Nakamoto. Bitcoin: A peer-to-peer electronic cash system. http://bitcoin.org/bitcoin.pdf, 2008.

[73] Report on Architecture for Privacy-Preserving Applications on Ledgers (D4.2). Technical report, PRIViLEDGE project, 2020.

[74] Arka Rai Choudhuri, Michele Ciampi, Vipul Goyal, Abhishek Jain, and Rafail Ostrovsky. Round optimal secure multiparty computation from minimal assumptions. In Rafael Pass and Krzysztof Pietrzak, editors, *Theory of Cryptography*, pages 291–319. Springer International Publishing, 2020.

[75] Alessandra Scafuro, Luisa Siniscalchi, and Ivan Visconti. Publicly verifiable proofs from blockchains. In Dongdai Lin and Kazue Sako, editors, *IACR International Workshop on Public Key Cryptography*, volume 11442 of *LNCS*, pages 374–401. Springer, 2019.

[76] Alessandra Scafuro, Luisa Siniscalchi, and Ivan Visconti. Publicly verifiable zero knowledge from (collapsing) blockchains. In Juan A. Garay, editor, *Public-Key Cryptography – PKC 2021*, pages 469–498. Springer International Publishing, 2021.

[77] Joosep Simm, Jamie Steiner, and Ahto Truu. Verifiable multi-party business process automation. In Adela Del Río Ortega, Henrik Leopold, and Flávia Maria Santoro, editors, *Business Process Management Workshops*, pages 30–41. Springer International Publishing, 2020.

[78] Chrysoula Stathakopoulou, Tudor David, and Marko Vukolic. Mir-bft: High-throughput BFT for blockchains. *CoRR*, abs/1906.05552, 2019.