Designing a Central Bank Digital Currency with Support for Cash-like Privacy

Jonas Gross[a], Johannes Sedlmeir[a,b], Matthias Babel[a], Alexander Bechtel[c], Benjamin Schellinger[a,d]

[a]University of Bayreuth; [b]Project Group Business & Information Systems Engineering of the Fraunhofer FIT; [c]University of St. Gallen; [d]FIM Research Center, University of Bayreuth

ARTICLE VERSION

January 14, 2022

ABSTRACT

Most central banks in advanced economies consider issuing central bank digital currencies (CBDCs) to address the declining use of cash as a means of payment and to position themselves against increased competition from Big Tech companies, cryptocurrencies, and stablecoins. One crucial design dimension of a CBDC is the degree of transaction privacy. Existing solutions are either prone to security concerns or do not provide full (cash-like) privacy. Moreover, it is often argued that a fully private payment system and, in particular, anonymous transactions cannot comply with anti-money laundering (AML) and countering the financing of terrorism (CFT) regulation. In this paper, we follow a design science research approach (DSR) to develop and evaluate a holistic software-based CBDC system that supports fully private transactions and addresses regulatory constraints. To this end, we employ zero-knowledge proofs (ZKPs) to enforce limits on fully private payments. Thereby, we are able to address regulatory constraints without disclosing any transaction details to third parties. We evaluate our artifact through interviews with leading economic, legal, and technical experts and find that a regulatorily compliant CBDC system based on ZKPs that supports full (cash-like) privacy is feasible.

KEYWORDS

Anonymity, CBDC, Compliance, Design Science, Digital Identity, Digital Wallet, Electronic Cash, Payment System, Privacy by Design, Regulation, Self-Sovereign Identity, Zero-Knowledge Proof

Corresponding author: Jonas Gross, University of Bayreuth, Universitätsstraße 30, 95447 Bayreuth, Germany, jonas.gross@uni-bayreuth.de.

Electronic copy available at: https://ssrn.com/abstract=3891121

Excutive Summary

We propose a two-tiered (retail) central bank digital currency (CBDC) system consisting of transparent and private CBDC accounts. The main focus of this paper lies on the private accounts provided by the *privacy pool*. Our general architecture is illustrated in Figure 1. The central bank issues CBDC into transparent accounts. These accounts are maintained either by the central bank or by payment service providers (PSPs), such as banks, on behalf of the central bank. CBDC account holders can deposit CBDC into the privacy pool either via transparent CBDC accounts or deposit cash via special automated teller machines (ATMs). End-users can conduct three different types of payments differing in the degree of privacy, namely fully private, semi-private, and fully transparent transfers.

- Fully private (cash-like) transfers take place inside the privacy pool. In fully private transfers, the identities of both involved parties, as well as the transaction amount, remain hidden to third parties.
- Semi-private transfers take place between the privacy pool and the transparent CBDC accounts. They reveal the amount of the transfer to the PSP and the central bank, and only the identity of the holder of the transparent CBDC account to the involved PSPs and potentially, depending on the design, also to the central bank.
- Fully transparent transfers take place between transparent CBDC accounts. In fully transparent transfers, the sender's and receiver's PSPs and, depending on the design, potentially also the central bank, know the identities of the involved parties and the transaction amount, similar to current electronic payments via commercial banks.

For semi-private and fully private transfers, the privacy guarantees are provided *by design*, i.e., no third parties and, in particular, neither PSPs nor the central bank nor regulatory authorities, will learn about the transaction amount and the involved parties even if they collude or, in the most extreme case, behave maliciously.

Addressing regulatory constraints related to anti-money laundering (AML) and combating the financing of terrorism (CFT) regulation while allowing for fully private payments, we impose balance, transfer, and turnover limits. We enforce these limits without sharing transfer details beyond the information that we discussed previously with any third parties by using general-purpose zero-knowledge proofs (ZKPs)

Electronic copy available at: https://ssrn.com/abstract=3891121

in the form of zero-knowledge succinct non-interactive arguments of knowledge (zk-SNARKs).

We develop an *unspent account state (UAS)* system to obtain the privacy and, in particular, unlinkability guarantees of Zcash with a similar technical construction based on commitments, nullifiers, and ZKPs. More precisely, we do not use a token-based model (corresponding to *coins* or *digital banknotes*), as it is unclear how turnover and balance limits could be implemented in such an approach. Instead, we use commitments and nullifiers to obfuscate transfers in an account-based model (including an *identity* and a *balance*) and make them unlinkable. This approach facilitates the enforcement of turnover and balance limits for fully private payment, as it funnels all associated transfers conducted by one particular user into one unique privacy pool account that they control. In our design, all end-users privately register and maintain their own private CBDC account and only send cryptographic proofs of the correct local accounting and compliance with the imposed limits to the ledger that is maintained by the central bank. A transfer proposal is only accepted by the central bank if the ZKP rightfully proves the compliance with the limits imposed. To ensure that every user only has access to one account, and hence make turnover and balance limits effective, our construction relies on the availability of a *unique* digital identity for end-users. For this purpose, we suggest using government-issued, certificate-based, digital IDs, as currently explored in various countries. Alternatively, PSPs such as banks could issue these identities, which would, however, require coordination among these intermediaries to prevent multiple registrations.

From the perspective of a user, an anonymous transfer corresponds to invalidating their previous account state and creating a new account state, where the difference between the previous balance and the new balance equals the amount of the transfer. Due to the fully-fledged user-sided accounting, our CBDC system would also allow integrating remuneration in private CBDCs accounts that directly affects the difference between the previous and the new balance. Third parties cannot trace back a transaction history to one specific user, and transfers by the same user are unlinkable. The only entities in the system that have access to their account details, including their ID information, balance, and turnover, are the users themselves. Users keep their CBDC in a digital wallet, e.g., in a dedicated app on a mobile device. The digital wallet stores the users' digital ID, cryptographic keys, and account details for their private CBDC account as well as access credentials for their transparent CBDC account. The wallet also supports users in maintaining their private and transparent accounts. For example, if one of the limits in the privacy pool prevents a user from conducting an

Electronic copy available at: https://ssrn.com/abstract=3891121

anonymous payment, the wallet may automatically suggest conducting the payment on the transparent side.

The implementation of our system can already be regarded as feasible. From a user's perspective, the processing time for a transfer would be comparable to private transfers in Zcash. Hence, payments would be processed within a few seconds or even less. Given the rapid improvements on the performance of generating ZKPs observed over the last few years, it seems likely that processing times will see further significant improvements in the near future. Additionally, the operational burden for the central bank is relatively low because the succinct proofs in zk-SNARKs are small and offer fast verification in few milliseconds. The main challenge of our approach is the prevention of a black market for private CBDC accounts: Getting control over a large number of private CBDC accounts by buying them from users that are not interested in using the privacy pool or by blackmailing would allow illicit actors to circumvent the limits and use the privacy pool for a *money mule* business. However, this threat can be effectively mitigated by binding users' digital identities to secure hardware, e.g., their mobile phones, and by requiring a ZKP of ownership of the private key stored in this secure hardware and associated with a non-expired, non-revoked unique digital identity in each transaction. This approach provides a high degree of assurance that users own the digital ID associated with the privacy pool account for every private transfer without the need to reveal correlatable identity-related information to third parties. Users that seek to pass on their accounts would need to pass on their unlocked mobile phone, on which their digital ID is stored, including all associated rights and permissions. Our discussions with stakeholders indicate that this risk is comparatively low and thus acceptable.

Electronic copy available at: https://ssrn.com/abstract=3891121

Figure 1.: High level architecture of the proposed CBDC design.

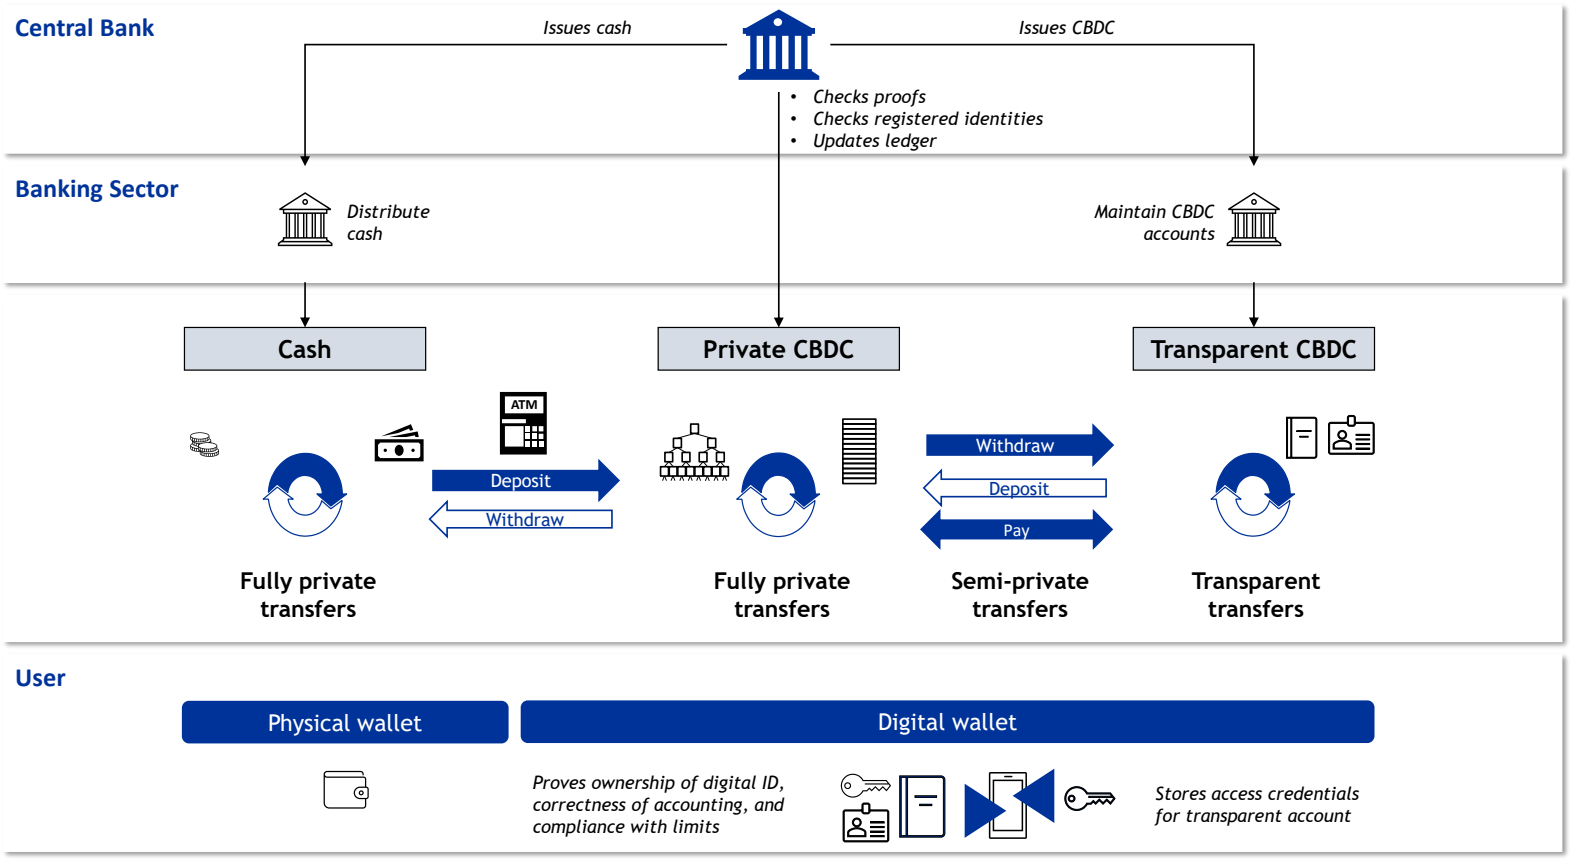

Electronic copy available at: https://ssrn.com/abstract=3891121

1. Introduction

The monetary system is changing. In many advanced economies, the use of cash as a means of payment has declined steadily over the last decade and in an accelerated way during the COVID-19 pandemic (European Central Bank, 2020b). Moreover, public money faces increasing competition from novel, private sector-issued forms of money, such as cryptocurrencies and stablecoins, and from Big Tech payment systems (European Central Bank, 2020a). Consequently, central banks take actions to preserve their monetary sovereignty. Today, 86 % of central banks around the world consider issuing their own digital currencies, known as central bank digital currencies (CBDCs) (Boar & Wehrli, 2021). While the Bahamas has already launched a CBDC and some other countries have introduced CBDC pilots (e.g., China), most jurisdictions are still debating and analyzing design options. In this context, the appropriate degree of transaction privacy receives great attention from central bankers, policy-makers, and academics. For instance, in its announcement to start a project on the digital euro, the European Central Bank (ECB) stressed that the two-year investigation phase aims to identify "the design options to ensure privacy and avoid risks for euro area citizens, intermediaries and the overall economy" (European Central Bank, 2021c).

Keeping transaction data private is crucial, among other reasons, to avoid identity theft, threats to personal security, data exploitation, and harassment based on potentially embarrassing but legal purchases (e.g., Chaum, Grothoff, & Moser, 2021; Choi, Henry, Lehar, Reardon, & Safavi-Naini, 2021; Kahn, 2018; Kahn, McAndrews, & Roberds, 2005). Privacy is also considered essential from an economic perspective, as it can help to avoid price discrimination (Acquisti, Taylor, & Wagman, 2016; Odlyzko, 2004), making privacy a public good (Garratt & Van Oordt, 2021). Moreover, privacy constitutes a fundamental civil right enshrined in Article 12 of the United Nations Declaration of Human Rights (UDHR) (United Nations, 1948), in Article 8 of the Convention for the Protection of Human Rights and Fundamental Freedoms (European Court of Human Rights, 1950) as well as in Article 7 and 8 of the Charter of Fundamental Rights of the European Union (European Convention, 2000). In the context of CBDC, a consultation of European citizens demonstrated that they see privacy as the *most* important requirement for a CBDC (European Central Bank, 2021b).

A CBDC that stores transaction details in a centralized database operated by the central bank (or payment service providers (PSPs) on behalf of the central bank) bears the risk of losing trust and causing security incidents such as data breaches, e.g., due to human misbehavior or cyber attacks. Incidents like the hack of New Zealand's

Electronic copy available at: https://ssrn.com/abstract=3891121

central bank in 2021 demonstrates that cyber risks are indeed a threat that should be taken seriously (The Guardian, 2021). Furthermore, if sensitive data are stored centrally, end-users have to *trust* the operator that the privacy promises will not be compromised in the future. However, in such a setup, operators could potentially change their mind or secretly analyze (historic) transaction data and share it with further parties, thereby potentially undermining privacy and trust. Trust in a payment system that inevitably processes sensitive data can be increased by following a *privacy-by-design* approach in which customers do not need to trust the operator for privacy protection and where large-scale data breaches are naturally excluded. In this case, private data would only be stored with the end-user involved in a transaction and not aggregated in a centralized system, thereby providing *trustless privacy*.

Today, cash is the only regulatorily compliant form of money that provides full privacy by design. If payments are conducted digitally, e.g., through mobile payments, bank transfers, or credit cards, the transaction data is stored with the involved PSP. Contrary to public perception, cryptocurrencies such as Bitcoin and Ether do not ensure a high degree of privacy either, as transaction details are stored on a public ledger, and the pseudonymous addresses that send and receive cryptocurrencies can often be traced back to the users that control them (Biryukov, Khovratovich, & Pustogarov, 2014) through analyzing metadata (such as IP addresses) and information from exchanges that need to conduct know-your-customer (KYC) measures (Silfversten et al., 2020). Against this background, privacy-oriented cryptocurrencies such as Zcash and Monero have been developed. They use cryptographic techniques such as zero-knowledge proofs (ZKPs) to enable fully private payments (Fauzi, Meiklejohn, Mercer, & Orlandi, 2019). However, these cryptocurrencies do not conform with prevailing regulations, as unlimited anonymous payments open the door for illicit activities, such as money laundering and terrorist financing (Silfversten et al., 2020).

To secure access to a fully private, regulatorily compliant form of money in an increasingly digital environment, a CBDC should provide a high degree of transaction privacy and offer (at least) the same privacy-preserving features as cash. Multiple central banks have already indicated their willingness to consider privacy-enhancing features for their CBDCs (Bank of Canada, 2020; European Central Bank, 2021a; Lane, 2020), and first CBDC solutions have been proposed by both central bankers and academic researchers that provide some degree of transaction privacy. Naturally, these suggestions also consider regulatory constraints. However, software-based CBDC designs proposed by central banks, e.g., the ECB's anonymity voucher proposal (European Central Bank, 2019), or by academic researchers (Chaum et al., 2021; Dold, 2019;

Electronic copy available at: https://ssrn.com/abstract=3891121

Tinn & Dubach, 2021), do not support fully private transactions. Besides software-based designs, CBDC solutions can use hardware elements, e.g., built in computers, mobile phones, or smart cards, as gateways to access the CBDC infrastructure (European Central Bank, 2020a) and could, therefore, technically replicate the trustless privacy guarantees of cash. As an example, a hardware-based smart card system is currently being tested by the Central Bank of the Bahamas (Mastercard, 2021). However, such hardware-based solutions still exhibit considerable security challenges (Chaum et al., 2021; European Central Bank, 2021c), and mitigating the risks of sophisticated attacks would likely compromise privacy guarantees (Chaum et al., 2021).[1]

Fortunately, the maturity of privacy-enhancing cryptographic techniques, and particularly of ZKPs, has grown considerably in recent years, offering new opportunities for enhanced privacy. ZKPs have already seen considerable adoption in the context of privacy-oriented cryptocurrencies, where they are used to ensure the integrity of payment systems, e.g., to prevent double-spending, while maintaining a high degree of privacy for the user. However, ZKPs can be more broadly employed, with particular emphasis on enforcing further monetary or regulatory rules in a privacy-oriented payment system. Literature on cryptography already acknowledges the suitability of ZKPs for reconciling privacy and integrity or compliance requirements for electronic payments, e.g., through imposing turnover or per-transaction limits (e.g. Garman, Green, & Miers, 2016). Bontekoe (2020) specifically proposed an extension of Zcash in which third parties escrow users' digital identities and ZKPs allow the enforcement of turnover limits.

Still, regulators and central bankers have repeatedly claimed that reconciling full privacy with regulatory constraints is not possible (e.g. Armelius, Claussen, & Hull, 2021; Auer & Boehme, 2021). This statement clearly indicates a lack of communication between the different research streams. Moreover, neither CBDC nor cryptographic literature has so far provided a rigorous, holistic evaluation of a payment system design that addresses regulatory requirements while supporting fully private payments and that evaluates the design with key stakeholders, such as central bankers and regulators. To address this research gap, we follow a rigorous design science research (DSR) approach to design and evaluate a holistic, software-based CBDC system that is based on ZKPs and supports fully private payments. We first consolidate proposals from the cryptographic and CBDC literature to develop an account-based CBDC payment sys-

[1] We recommend to refer to the comprehensive discussion in Chaum et al. (2021) for a more detailed overview of the challenges associated with hardware-based solutions, particularly if deployed on a larger scale, and to Grothoff and Dold (2021) for additional arguments why a software-based CBDC design may be beneficial.

Electronic copy available at: https://ssrn.com/abstract=3891121

tem that is fully private by design while addressing regulatory requirements by using per-transaction, turnover, and balance limits. We also instantiate our design through an implementation of the core transaction types using ZKPs. We then evaluate and refine our IT artifact in four evaluation cycles consisting of a total of 22 interviews with 44 experts in the area of regulation, cryptography, central banking, identity, and payments. We find using ZKPs for CBDCs can replicate cash-like privacy in the digital realm and ensure adherence to regulatory constraints. Against this background, ZKPs enable strict privacy protection by design, storing personal transaction data only on the end-users' devices (*trustless privacy*).

The theoretical contribution of our paper is twofold: First, our innovative software-based CBDC payment system combines elements from different strands of the literature, including cryptography, privacy-by-design concepts, and CBDCs. Second, the rigorous evaluation of our CBDC system by key stakeholders in the cadre of DSR allows us to assess the practical feasibility of a CBDC design that provides cash-like privacy using ZKPs and also to discuss risks and potential mitigation measures.

Our paper is structured as follows: In Section 2, we introduce essential background knowledge on CBDCs, different notions of transaction privacy, regulatory aspects of CBDCs, and ZKPs. We then present our DSR approach in Section 3. Subsequently, we discuss related work and describe the design of our IT artifact in Section 4, followed by the presentation of our evaluation cycles in Section 5. Section 6 summarizes our main findings, describes limitations, and gives an outlook on potential future applications of our design and related research opportunities.

2. Theoretical Foundations

2.1. Central Bank Digital Currencies

In general, there are two forms of CBDCs, wholesale and retail CBDCs (Bech & Garratt, 2017). A wholesale CBDC is a digital form of central bank money accessible for financial institutions to optimize the settlement of wholesale payments and tokenized financial assets. A retail CBDC, in contrast, constitutes a novel form of central bank money available to the general public. In this paper, we solely refer to a retail CBDC, as we focus on end-user payments. A retail CBDC unites features of today's predominant forms of money: cash and bank deposits (Bech & Garratt, 2017). While cash is issued by central banks in physical form, bank deposits are issued by commercial banks in digital form. As central bank money, CBDCs bear no counterparty risk

Electronic copy available at: https://ssrn.com/abstract=3891121

because the central bank issuing the CBDC cannot – by definition and in contrast to commercial banks – become bankrupt.[2] CBDCs would hence provide a safer and practically riskless form of money for end-users.

CBDCs can be designed and implemented in different ways (Allen et al., 2020; Auer & Böhme, 2020; European Central Bank, 2020b; Kiff et al., 2020). Auer and Böhme (2020) identify *architecture*, *access*, and *technology* as the three main design considerations for a CBDC.[3] The architecture specifies the role of the central bank and other market participants in the CBDC ecosystem. The account management, onboarding processes, and distribution of a CBDC might be conducted directly by the central bank (direct model) or by private sector PSPs (intermediated model). The access model defines how CBDC transaction data is stored and how access is managed. In an account-based model, the CBDC is stored in accounts, and hence the ownership of a CBDC is tied to an identity. In a token-based model, the central bank issues digital bearer instruments and ties the CBDC ownership to the (proof of) ownership of the CBDC units itself, similar to cash today. Regarding technology, a CBDC can be issued either via a centralized or a distributed ledger. If a centralized ledger is used, the central bank manages and controls the CBDC system. In the case of a distributed ledger, data processing, storage, and governance can be distributed across additional private or public sector institutions.

2.2. Privacy and Regulatory Compliance of Payments

In this paper, we distinguish between private, anonymous, and fully (cash-like) private transactions. In a *private* transaction, the transaction amount remains unknown, but the sender and receiver, i.e., the transaction parties, might be known to third parties (e.g., PSPs, the central bank, or regulatory authorities). In an *anonymous* transaction, the identities of the sender and receiver remain hidden, but the transaction amount might be known. *Fully private* transactions are private and anonymous; neither the transaction amount nor the sender or receiver are revealed to third parties. Therefore, our definition of full privacy is similar to the concept of secrecy as the concealment of

[2]Today, deposit insurance schemes are established to address the risk of commercial bank money and avoid that, in the case of bankruptcy of a commercial bank, costumers face substantial financial losses. However, deposit insurance schemes are not available in all countries equally, and commercial bank money is only secured until a specific threshold, e.g., in the euro area up to 100,000 Euro per client per financial institution.

[3]The fourth dimension refers to retail and wholesale interlinkages. Such interlinkages are especially relevant for cross-border CBDC payments. As we abstract from cross-border use in this paper, we do not consider this dimension.

Electronic copy available at: https://ssrn.com/abstract=3891121

information (Bok, 1989; Tefft, 1980). Full privacy or secrecy describes the attempt of consumers to avoid sharing information in order to prevent third parties from creating a digital representation of the real self (Dinev, Xu, Smith, & Hart, 2013; Zwick & Dholakia, 2004).

Today, fully private and regulatorily compliant transactions are only possible with physical cash and, to a certain extent, with e-money. All other payment methods are either not fully private (e.g., credit cards, bank transfers, Bitcoin), or they are not regulatorily compliant (e.g., privacy-oriented cryptocurrencies such as Monero and Zcash). In order to restrict the large-scale financing of illicit activities, regulators usually enforce per-transaction, turnover, and/or balance limits for anonymous payments. For instance, there are *per-transaction limits* for fully private cash payments in many euro area countries such as Greece (500 EUR), France and Portugal (1,000 EUR), Italy (2,000 EUR), Spain (2,500 EUR), Belgium (3,000 EUR), and Slovakia (15,000 EUR) (Pocher & Veneris, 2021). For anonymous e-money transactions, the 5th Anti-Money Laundering Directive specifies a monthly *turnover* and *balance limit* of 150 EUR (European Union, 2018). As, to date, regulatory frameworks do not capture CBDCs, there are no such regulatory limits for CBDCs yet. However, it seems reasonable to expect that similar limits would need to be introduced for anonymous CBDC payments, similar to today's restrictions for anonymous cash and anonymous e-money transactions.

2.3. Zero-Knowledge Proofs

The notion of ZKPs was first introduced in the 1980s, describing "proofs that convey no additional knowledge other than the correctness of the proposition in question" (Goldwasser, Micali, & Rackoff, 1989, p. 186). ZKPs refer to cryptographic protocols in which a *prover* can convince a *verifier* about a mathematical statement, for example, that the prover knows a piece of data that has specific properties. This statement may refer to the knowledge of a pre-image of a publicly known value under a hash function or about properties of the result of a publicly known algorithm that was executed on public or private data. In this setting, with a ZKP, the prover can convince the verifier without disclosing *any* information beyond the statement under consideration (Ben-Sasson, Bentov, Horesh, & Riabzev, 2018; Ben-Sasson, Chiesa, Genkin, Tromer, & Virza, 2013). If the statement refers to the output of an algorithm that was applied to data that is public or that was publicly committed to, a ZKP can enforce *computational integrity* without the need for the verifier to replicate the computation. Besides

Electronic copy available at: https://ssrn.com/abstract=3891121

providing *confidentiality* for data and intermediate steps in a computation, an appealing property of many ZKPs is that they are *succinct*, i.e., the size of proofs and the computational complexity of proof verification is significantly smaller than applying the algorithm. In the case of the zero-knowledge succinct non-interactive arguments of knowledge (zk-SNARKs) that we use in our CBDC system, both proof size and verification complexity are even *independent* from the complexity of the computation that is to be verified (Ben-Sasson et al., 2013). However, general-purpose ZKPs that can cover a large class of statements come at a high computational overhead for the prover.

In the 25 years after their discovery, researchers have leveraged special types of ZKPs in some contexts, such as enforcing *correct behaviour* in multiparty computations or *selective disclosure* in digital identity management schemes with *anonymous credentials*. The latter describes digital certificates that a trusted organization signed digitally and that their owner can use to prove claims about parts of the content of these certificates without revealing all of the contained information. In particular, when verifiably presenting attributes attested in the anonymous credential, strongly correlating contents such as the value of the digital signature itself do not need to be revealed (Ben-Sasson et al., 2013). Lately, these anonymous credentials have seen first adoption in so-called decentralized or self-sovereign identity projects as explored by the public and private sector in Canada and Germany, among others (Kubach & Sellung, 2021). However, practical applications remained rare, as for general-purpose ZKPs beyond these very specific cases, the computational complexity for the prover was prohibitive. Also, seemingly, there was not a considerable need for deploying ZKPs because information systems (IS) were generally designed with a service provider that was trusted with respect to both integrity and confidentiality. However, this paradigm started to shift with the advent of Bitcoin and the decentralization as facilitated by blockchain technology. Building on blockchain technology, a new type of IS, decentralized applications, emerged that do not involve a third party that is trusted with respect to integrity by performing computations redundantly (Rossi, Mueller-Bloch, Thatcher, & Beck, 2019). However, the replicated execution of operations on blockchains immediately leads to considerable challenges from a scalability and confidentiality perspective (Ben-Sasson et al., 2018; Kannengießer, Lins, Dehling, & Sunyaev, 2020).

In this context, general-purpose ZKPs started to find applications in privacy-oriented cryptocurrencies such as Zcash or applications on Ethereum such as Tornado-Cash, building on prior academic work (Ben-Sasson et al., 2014) to provide a Bitcoin-like payment system with fully private transactions. In the last few years, addi-

Electronic copy available at: https://ssrn.com/abstract=3891121

tionally, the succinctness of proofs has been leveraged by various projects on the Ethereum blockchain and novel cryptocurrencies. In zk-rollups, an untrusted third party batches many operations and proves the correctness of the resulting state transition with a ZKP. Through their ability to solve privacy challenges in cryptocurrency and blockchain projects (Partala, Nguyen, & Pirttikangas, 2020), ZKPs have received increased attention in academia and business, and have hence considerably matured in terms of performance and applicability. Consequently, IS building on blockchains and general-purpose ZKPs have already seen first adoption in industry consortia that leverage blockchain technology, e.g., in the context of medical supply chains (Mattke, Hund, Maier, & Weitzel, 2019).

3. Method

This paper follows a rigorous DSR approach (Hevner, March, Park, & Ram, 2004; March & Smith, 1995; Peffers, Tuunanen, Rothenberger, & Chatterjee, 2007) to design, develop, and evaluate a CBDC system that provides full privacy while addressing regulary requirements related to anti-money laundering (AML) and combating the financing of terrorism (CFT). We structure our paper as proposed by Gregor and Hevner (2013). To ensure methodological rigor, we use the widely accepted DSR methodology proposed by Peffers et al. (2007). Thus, we apply the following six steps procedure to derive our IT artifact: (1) Problem identification, (2) objectives definition, (3) design and development, (4) demonstration, (5) evaluation, and (6) communication (see Figure 2).

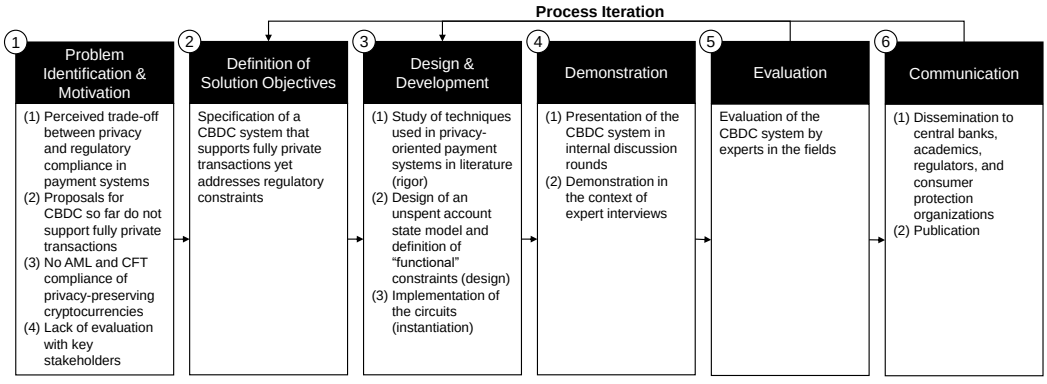

Figure 2.: DSR approach to design our account-based CBDC system according to Peffers et al. (2007).

Electronic copy available at: https://ssrn.com/abstract=3891121

DSR was established to enable IS practitioners to find solutions to previously unsolved problems through a continuous build-and-evaluate process (Hevner et al., 2004; March & Smith, 1995). For an IT artifact to make a valuable contribution to IS research, it must address both a relevant business need (Hevner et al., 2004) and a general problem (Iivari, 2015). First, we identified the underlying problems and derived design requirements for our CBDC system. We screened the most relevant primary literature on CBDCs, namely ECB (2020a) (European Central Bank, 2020a), BoC (2020) (Bank of Canada, 2020), Fed (2021) (Cheng, Lawson, & Won, 2021), BoE (2020) (Bank of England, 2020), and BIS et al. (2020) (Bank for International Settlements et al., 2020), and identified both users' and central banks' requirements for a CBDC (see Table 1).

We derived the following key requirements for end-users: privacy protection, high security and transaction speed, fast settlement, low costs, high usability, and availability. For central banks, the following requirements are important: AML and CFT compliance, market neutrality, resilience, cooperation with market participants, universal access, cost efficiency, and interoperability. In Section 1, we argued that privacy features should be at the core of a CBDC system, but also that regulatory constraints need to be addressed. Thus, in our DSR approach, we focus on these two core requirements. As CBDC implementations are currently still at an early stage, it is not (yet)

Table 1.: Literature review on CBDC requirements based on central bank statements. We included requirements that have been referred to in at least four of the five studies.

	ECB (2020a)	**BoC (2020)**	**Fed (2021)**	**BoE (2020)**	**BIS et al. (2020)**
Liability of the central bank	●	●	●	●	●
Market neutrality	●	●	●	○	●
Security	●	●	●	●	●
Convenience and ease of use	●	●	●	●	●
Transaction speed and fast settlement	●	●	●	●	●
Privacy protection	●	●	○	●	●
Usability	●	●	●	○	●
Low cost	●	●	○	●	●
Regulatory compliance	●	●	●	●	●
Resilience	●	●	●	●	●
Cooperation with market participants	●	●	●	●	○
Universal access	○	●	○	●	●
Cost efficiency	●	●	●	●	○
Interoperability	●	●	○	●	●

Electronic copy available at: https://ssrn.com/abstract=3891121

feasible to address all requirements simultaneously in one artifact. This hypothesis was also confirmed in the expert interviews that we conducted. The already proposed CBDC approaches do not enable fully private transactions, and related work in cryptography lacks a concrete design that can be used for the rigorous evaluation of their regulatory compliance and feasibility from the perspective of stakeholders (see Section 4). In this light, we designed and instantiated a solution that uses cryptographic techniques, i.e., ZKPs, to address these requirements and to enable a discussion with stakeholders. In particular, we aimed to develop a CBDC system that ensures cash-like privacy by design, where the transaction amount and the identities of involved transaction parties are not shared with any third party. Compliance with regulation is enforced by per-transaction, balance, and turnover limits for fully private payments.

Next, we designed and developed our CBDC system in cycles that iterated between conceptualization, instantiation, and internal evaluation. We present the overall CBDC system architecture, onboarding procedure, depositing, withdrawing, and fully private transaction processes in Section 4. We then discussed our CBDC proposal in internal discussion rounds and presented it to leading experts from various fields. An overview of the interviewed experts is depicted in Table 2. As one key result, our evaluation confirmed the feasibility and adequacy of our CBDC system. The adjustments to our CBDC architecture after each evaluation cycle are discussed in Section 5. Sixth, as a final step, we disseminated our key findings to the interviewees and other stakeholders, in particular, to decision-makers in central banks and regulatory authorities and to researchers. In addition, we publish the source code of our prototype for the proposed CBDC system on GitHub.[4]

[4]The repository can be accessed at: https://github.com/applied-crypto/cbdc.

Electronic copy available at: https://ssrn.com/abstract=3891121

Table 2.: Overview of interviewed experts.

Cycle	Field of expertise	Int. No.	Exp. No.	Role	Organization
1	Law	01	01	Assistant Professor	University
	Law	02	02	Specialist Payment Fraud	Europol
	CBDC	03	03	Senior Financial Sector Expert	ex-IMF
	Payments	04	04	Head of Payments	Central Bank
	Payments		05	Head of Digitalization and Payment Systems	
	CBDC	05	06	Chief Economic Advisor	SFB Technologies
	Cryptography	06	07	Global Managing Director Digital Assets	Accenture
2	IT	07	08	Technical Lead Digital Currencies	CBDC-developing company
	Economics		09	Business Lead Digital Currencies	
	Business		10	Product Manager Digital Currencies	
	Computer Science		11	Data Engineer Digital Currencies	
	Information Systems	08	12	Senior Researcher	Research Institute
	Law	09	13	Banking Supervision Expert	Banking Association
	Payments		14	Lead Digitalisation	
	Economics		15	Chief Economist	
	Law		16	Legal Lead	
	CBDC		17	CBDC Expert	
	Payments		18	CBDC Expert	
	Law	10	19	Professor	University
	IT / Business	11	20	Head of DLT Product	Bank
3	Economics	12	21	Alternate Member of the Governing Board	Central Bank
	Computer Science		22	Professor	University
	CBDC	13	23	Senior Economist	International Organization
	Payments		24	Senior Financial Market Analyst	
	Cryptography	14	25	Professor	University
	Economics		26	PhD Candidate	Télécom Paris
	CBDC	15	27	Head of Blockchain	Association
	Payments	16	28	Market Infrastructure Specialist	Central Bank
	Computer Science		29	IT Application Development Specialist	
4	Digital Identities	17	30	Head of SSI Consortium (IDUnion)	Main Incubator
	Law	18	31	Expert on CBDC and AML Regulation	University
	Business	19	32	Senior Manager Marketing & Public Affairs	Federal Printer
	CBDC		33	Senior Project Manager CBDC	
	Digital Identities		34	Senior Consultant Trusted Services	
	Cyptography		35	Senior Principal Security Systems	
	Computer Science		36	Technological Expert CBDC	
	Business		37	Senior Account Manager	
	Business		38	Regional Sales Director	
	Business		39	Senior Business Development Manager	
	Computer Science	20	40	CBDC Technology Expert	Central Bank
	Computer Science	21	41	Research Group Lead	Research Institute
	Economics		42	PhD Candidate	
	Finance	22	43	Research Group Lead	Consumer Protection Org.
	Economics		44	Advisor	

Electronic copy available at: https://ssrn.com/abstract=3891121

4. Our CBDC Design

4.1. Related Work

To conceptualize our design, we investigated cryptography-related research with a focus on privacy-oriented digital payment systems. Chaum (1983) founded the research stream on *e-cash* that aims to develop cryptography-based payment systems that are private by design and make payments untraceable. He proposes a design in which users need to exchange their received *digital banknotes* for new ones in a compulsory interaction with a trusted PSP, e.g., a bank. To make the spending and receiving of a specific banknote unlinkable, *blind signatures* hide the serial numbers of unique and thus distinguishable digital banknotes. However, copying and thus double-spending these digital banknotes cannot be prevented technically. Instead, to hold users accountable, the cryptographic protocol allows retrieving a user's identity that is hidden in the digital banknote from combining the information obtained in two different payments with the same digital banknote. Despite being computationally very efficient, this approach implies that, while the sender remains anonymous, the receiver is identified by its PSP, and the payment amount is transparent. Moreover, the design cannot practically enforce per-transaction limits. Moreover, turnover and balance limits are cumbersome to implement because in a CBDC system that involves multiple PSPs, as currently planned in most jurisdictions, a synchronization among PSPs would be required to detect double-spending attempts and to prevent users from visiting multiple PSPs to circumvent these limits. Furthermore, it is not clear how to implement basic programmability functionalities, such as interest payments, on top of this design.

Sander and Ta-Shma (1999) address some limitations of Chaum's approach, for instance, hiding the transaction amount by dividing it into discrete shares. Nevertheless, the PSP still learns the transaction amount paid and received in this epoch and the identity of the receiver. The first e-cash system that hides the identity of both sender and receiver and also the transaction amount without the need for a trusted third party was proposed by Camenisch, Hohenberger, and Lysyanskaya (2006). In addition to hiding the identity of the receiver and the transaction amount, this payment system guarantees the sender's anonymity as long as they do not double-spend and exceed a pre-defined per-transaction limit. Technically, the proposal is based on ZKPs that are mathematically tailored specifically to this use case. However, as this approach is still designed for digital banknotes and different payments in which the same individual is involved cannot be identified and not even linked, turnover limits cannot

Electronic copy available at: https://ssrn.com/abstract=3891121

be enforced. Additionally, it is challenging to extend this model to incorporate basic programmability features.

While none of the previous academic work seems to have received considerable adoption, the first practical emergence of privacy-oriented digital payment systems happened in the context of cryptocurrencies. Researchers developed privacy-oriented modifications of Bitcoin (Nakamoto, 2008) as the use of a public distributed ledger increases the need for a privacy-enhancing design. The two arguably most relevant academic studies in this context are Zerocoin, which hides transaction parties but not the transaction amount (Miers, Garman, Green, & Rubin, 2013), and Zerocash (Ben-Sasson et al., 2014), which additionally hides the transaction amount and laid the foundation for the well-known cryptocurrency Zcash. Both approaches use general-purpose ZKPs, in particular zk-SNARKs, to demonstrate that transactions are valid and respect agreed-upon monetary policies (e.g., no double-spends) without revealing transaction details. Garman et al. (2016) acknowledge that "from an investigative standpoint, Zerocash is no different than cash" (Garman et al., 2016, p. 81). However, they also point to the conflict between regulation and private payment systems, as both Zerocoin and Zerocash do not take into account legal frameworks that restrict the use of anonymous payments to comply with AML and CFT regulation. Consequently, Garman et al. (2016) sketch potential extensions of the Zerocash model to address regulatory constraints, such as considering specific jurisdictions involved in payments, per-transaction limits, the tracing of specific coins, and the payment of appropriate taxes. Since exceeding the limits enforced by ZKPs would imply that no transactions are possible and thus reduce the utility of the payment system substantially, they propose a tiered approach for a private payment system so that any transaction above the spending limit should be additionally signed by an authority that conducts KYC and AML checks. Similarly, Bontekoe (2020) follows the approach of modifying Zcash by adding an account-based system based on a previous KYC-process. This setup allows limiting the transactions of an account within a specific epoch that can ensure balance and turnover limits and, thus, regulatory compliance. The study also describes the possibility of enabling transactions that exceed a limit by verifiably encrypting the associated transaction data and allowing for subsequent checks through a dedicated authority.

Both Garman et al. (2016) and Bontekoe (2020) also mention that users' digital identity management based on digital certificates can take a valuable role in this context. Literature that considers the role of identity in privacy-oriented payment systems also emphasizes that a mechanism for revoking identities and accounts is required (Choi et

Electronic copy available at: https://ssrn.com/abstract=3891121

al., 2021). From a regulatory perspective, this resonates well with sanctions lists such as the Office of Foreign Assets Control (OFAC)'s Specially Designated Nationals.

In parallel to these developments, cryptographic research in the context of CBDCs continued pursuing Chaum's initial *asymmetric* approach to transaction privacy, offering anonymity for the sender but not for the receiver (Chaum et al., 2021; Tinn & Dubach, 2021). Chaum et al. (2021) base their design on blind signatures and hence employ similar techniques as earlier work by Chaum (1983) from a technical perspective, whereas the approach by Tinn and Dubach (2021) is based on zk-SNARKs. Veneris, Park, Long, and Puri (2021) propose a system for CBDC that makes transaction amounts transparent yet enables anonymity through identity escrow (which is consequently not trustless). They also add a hardware-based solution that provides essentially cash-like privacy but requires regular online settlement. In these designs, the privacy protection of the sender covers many practical requirements, especially when a consumer purchases a product from a business that needs to disclose its accounting transparently. However, as discussed in Section 1, particularly for payments between end-users, trustless, cash-like privacy may be desirable.

To date, several CBDC designs have been proposed that aim for regulatory compliance and, at the same time, preserve users' privacy – at least to some extent (see Table 3). However, these designs differ substantially in the extent to which privacy is guaranteed. In this context, *privacy by design* and *compliance by design* play a central role, representing systems where no trust in the operator is needed. Privacy by design and compliance by design can be achieved through extending a payment system like Zcash that allows fully private transactions with the possibility of person-related monthly turnover or per-transaction limits, as proposed by Bontekoe (2020); Garman et al. (2016). Yet, since the perception that enabling fully private transactions in digital form fundamentally contradicts AML and CFT regulation is still widespread, and regulators and central banks have repeatedly claimed that reconciling full privacy with regulatory constraints is not possible (e.g. Armelius et al., 2021; Auer & Boehme, 2021), we present a holistic approach for a *privacy/compliance-by-design* CBDC in detail. We also show how digital identities can replace an efficient and privacy-preserving KYC process.

A frequently stated caveat when designing a private payment system is that it is necessary to consider not only the processing and storage of transaction data but also the metadata from the corresponding communication. For example, anonymizing the parties involved in a payment is hardly useful if the parties' static IP addresses are attached to transaction requests. In this context, sidechannel attacks have already been

Electronic copy available at: https://ssrn.com/abstract=3891121

Table 3.: Comparison of privacy-oriented CBDC solutions.

		ECB (2019) Anonymity Vouchers	Chaum et al. (2021) Blind signatures	Tinn and Dubach (2021) P-hybrid CBDC	Choi et al. (2021) Banking$^+$ and Cash$^+$	Veneris et al. (2021) CBDL	Our approach UAS model with ZKPs
Privacy	Anonymity of the sender	●	●	●	◐	●	●
	Anonymity of the receiver	●	○	○	○	●	●
	Privacy of transaction amount	○	◐	○	○	●	●
	Trustless privacy (privacy by design)	○	●	●	○	◐	●
Regulation	Regulation by design	○	●	○	○	◐	●
	Per-transaction limits	●	●	○	○	●	●
	Turnover and balance limits	●	●	○	○	●	●
	Anonymous onboarding	○	○	○	◐	○	●
Other	Offline payments	○	○	○	○	●	○
	Account-based system	○	○	●	◐	●	●
	DLT-based system	●	○	●	●	●	○
	Involvement of PSPs	●	●	◐	●	●	◐

used to de-anonymize shielded transactions in Zcash (Tramèr, Boneh, & Paterson, 2020). Consequently, related work on private payment systems (e.g. Tinn & Dubach, 2021) state that privacy on the networking layer must also be provided. However, given the existence of onion routing mechanisms as implemented by the Tor network (Dingledine, Mathewson, & Syverson, 2004), this problem can be considered solved and we will not refer to it further in the remainder of this paper.

4.2. High-Level Overview of our CBDC Architecture

Following Garman et al. (2016), our overall CBDC design is based on a two-tiered approach that supports both fully private and transparent CBDC payments. On the transparent side, the CBDC is distributed to users via PSPs – i.e., we use an intermediated model – and access is linked to an identity – i.e., we use an account-based model. Transaction data is stored in a centralized ledger. However, our system could also accommodate a distributed ledger. The disadvantage of a distributed ledger is that it may introduce challenges regarding performance, and privacy issues may arise when storing transaction- or account-related information on the transparent side. Thus, modifications, such as storing only PSPs' balances in an obfuscated way, may be necessary on

Electronic copy available at: https://ssrn.com/abstract=3891121

the transparent side. In general, the transparent side is flexible to implement central bank-specific designs, e.g., introducing a maximum limit on transactions in general or using a direct instead of an intermediated model.

We propose an account-based model that allows the funnelling of all of a user's transactions into one single account, which has many similarities to the approaches by Garman et al. (2016) and Bontekoe (2020). Recently, the security of an account-based approach in combination with ZKPs has been studied in more depth also in Wüst, Kostiainen, and Capkun (2021), including several improvements in terms of performance. Yet, none of these publications incorporates a connection to digital identity management, which — as we will discuss later — is likely indispensable to address some of the challenges related to sharing access to the privacy pool. In our system, users maintain their accounting privately, and a transaction corresponds to updating their private account and sending a ZKP that proves to the central bank that the transaction's expected policies have been met. The transaction amount then essentially corresponds to the delta, i.e., the difference of the previous and the new account states' balances. In essence, users store and manage their own accounts and prove the correctness of their local accounting to the operator of the ledger, i.e., the central bank.

4.3. The Privacy Pool in Detail

Most cryptocurrencies, such as Bitcoin, record every transaction, including the sender, receiver, transaction amount, and authorization (in terms of the sender's digital signature) on a public ledger (Zhang, Xue, & Liu, 2019). Additionally, transactions either point directly to one or more previously received but so far unspent coins (unspent transaction output (UTXO) model) or to the sender's and receiver's pseudonymous accounts with a public balance (account-based model). This setup generally allows one to track the transaction history of digital banknotes and corresponding metadata as well as additional information retrieved from exchanges to identify the involved parties in most transactions (Meiklejohn et al., 2016). Consequently, even if the ledger is not public, such a construction would allow the operator of the ledger to link transactions and correlate even pseudonymous accounts with real-world identities and, thus, to retrieve personal information.

In contrast, Zcash only stores cryptographic hashes of transactions in a ledger that hides details and ensures unlinkability. The payment details are only known to the sender and receiver. In particular, for every transaction (including its details), Zcash

Electronic copy available at: https://ssrn.com/abstract=3891121

applies two different one-way cryptographic functions to generate two unique but different outputs, i.e., (1) a *commitment* and (2) a *nullifier*. The commitment and nullifier hence essentially hide the same transaction details but are computationally infeasible to correlate without knowledge of the transaction details. To spend a previously received transaction in the form of a commitment, the corresponding nullifier is then published together with a ZKP that the nullifier corresponds to a previously published commitment (proof of knowledge of the joint pre-image) (Ben-Sasson et al., 2014). The central bank can then check that the associated commitment has not been spent before by checking whether this particular nullifier has been published before. The ZKP-based construction hence certifies that the sender has access to digital banknotes that have not been spent previously and that the amount of money spent equals the amount received, without the need to disclose *any* additional information about the transaction, such as sender, receiver, amount, or a direct reference to the previous transaction in which the banknote was received (Ben-Sasson et al., 2014).

The design of our privacy pool is based on the construction of Zcash with an append-only ledger that stores commitments and nullifiers. Adding a per-transaction limit to Zcash would be an easy task (Garman et al., 2016). However, our approach contains a crucial modification, replacing the UTXO-based model with an unspent account state (UAS) model to (privately yet provably) funnel all of a users' transactions into one account and hence enable a user to prove that balance and turnover limits are also satisfied. Each commitment and nullifier thus represents a unique account state. By publishing a commitment and a new nullifier, the previous account state is invalidated and a new one is created. To validate transactions, the central bank receives the associated commitments, nullifiers, and proofs of the correct update of the account state using ZKPs. In particular, the central bank verifies the validity of the ZKPs and that the nullifier has not been revealed before and then adds the transaction to the ledger by including the commitment and nullifier in the existing associated ledger. Due to the succinctness property of the ZKPs used, the workload for the central bank in verifying the transactions is very small (Ben-Sasson et al., 2013). The pre-image resistance of hash functions (with a pre-image of high entropy) and the unlinkability of commitments and nullifiers ensure that information that is revealed to the central bank is not sensitive (Ben-Sasson et al., 2014). Consequently, only the account owners know their respective transaction and account details, such as the amount, current balance, and turnover. Nevertheless, they can still prove compliance with predefined rules by using ZKPs. This approach ensures *privacy by design* and enables the creation of *proofs of compliance with limits by default*. In such a system, entities that maintain

Electronic copy available at: https://ssrn.com/abstract=3891121

the ledger only need to be trusted with respect to *integrity* and not with respect to protecting the users' privacy.

As common in the context of blockchains, for storing the commitments and enabling efficient proofs of inclusion for a commitment, we use Merkle trees. A Merkle tree is a cryptographic data structure that represents many entries by one identifier, i.e., the Merkle root (Merkle, 1987). This setup allows for proofs of inclusion that can be checked only by comparing with the Merkle root. The Merkle root changes with each new entry in the tree. By using these Merkle trees in combination with ZKPs, it is possible to prove that a transaction proposal refers to an existing commitment in the ledger without pointing to (and thus revealing) it. If the ledger is not public, for instance in a setup based on a centralized ledger or a permissioned blockchain, it is necessary to prevent the correlatability of commitments and nullifiers that may occur through querying the Merkle path of a specific commitment, e.g., through querying Merkle subtrees.

To ensure the integrity of the payment system and the compliance with turnover, balance, and per-transaction limits, our UAS model contains the following information that is stored in digital wallets and consequently is only known to their respective holder:

- Identity information: A public key and a digital certificate that includes the account owner's identity information (digital ID card, see, e.g., (Sedlmeir, Smethurst, Rieger, & Fridgen, 2021)).
- Balance: The balance of the account holder.
- Epoch turnover: An accumulation of all amounts of spending transactions in the current epoch (whose length can be specified by the regulator).
- Epoch reset: The last reset of the epoch (the last time the epoch turnover was set to zero).

This structure should be seen as an initial proposal that one can flexibly extend when more account details need to be checked for compliance, e.g., one could include the nationality or the type of the account, i.e., private or corporate.

For our CBDC system, we determine three core transaction types for interacting with the privacy pool, illustrated in Figure 3. In the following, we outline and discuss these three types of transactions that are related to the privacy pool – onboarding, semi-private transfers, and fully private transfers. All three types of transactions involve private inputs, public outputs, and a ZKP that connects them and proves the correctness of the local accounting to the central bank. The account owner provides

Electronic copy available at: https://ssrn.com/abstract=3891121

private inputs, which contain the already mentioned account information. This data is sensitive, therefore, it stays hidden. However, the private inputs are represented by the public outputs using one-way functions. Hence, the account owner shares the public output (commitment, nullifier, etc.) with the central bank along with a ZKP that ensures that the public outputs were computed from the private input according to the rules of the payment system.

4.3.1. Onboarding

Each account is tied to an individual cryptographic key pair and, using a digital certificate, to a government-issued identity. We assume that each user has a single digital ID card that is also bound to a cryptographic key pair. Many countries already integrate keypairs into their physical ID cards and even provide dedicated mobile apps to store the ID card in digital form, e.g., Germany and Estonia. To onboard a new user, the user has to create a cryptographic proof that they possess a valid ID that is not expired or revoked and that the initial commitment is deterministically derived from the ID card via a one-way function and contains the correct initial account entries. Therefore, each onboarding procedure of one individual is based on the same key pair and always results in the same commitment. In detail, a ZKP for onboarding and, thus, opening a new account would be structured as follows:

> **Onboarding**
>
> INPUT: Key pair and digital certificate (government-issued ID card), timestamp
>
> OUTPUT: Commitment to initial account state, timestamp, ID issuer's public key
>
> CHECKING THAT:
>
> - *the account holder controls a valid ID card (signature, binding, non-revocation)*
> - *the identity lodged in the account corresponds to the ID card*
> - *the account's initial balance and epoch turnover are set to zero*
> - *the account's initial epoch reset equals the timestamp*
> - *the commitment equals the hash of the signed initial account state*

The onboarding process works as follows: First, using their legit key pair and digital ID, users create the onboarding transaction, consisting of the initial commitment and the ZKP and send it to the central bank. Second, the central bank verifies the

Electronic copy available at: https://ssrn.com/abstract=3891121

Figure 3.: Transaction types in the privacy pool.

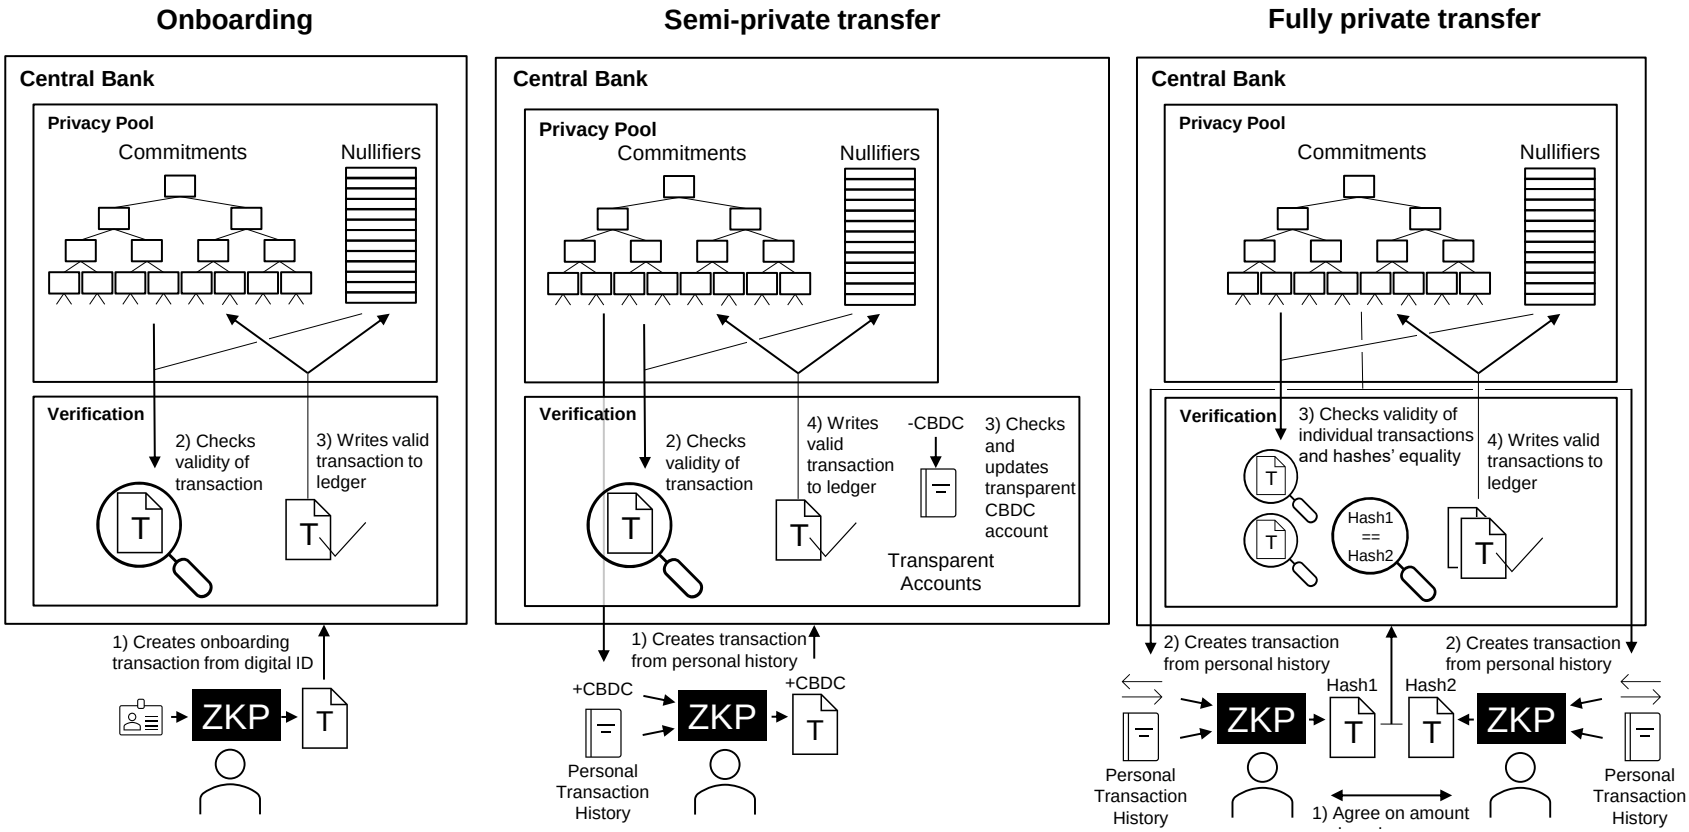

Electronic copy available at: https://ssrn.com/abstract=3891121

validity of the ZKP that refers to the commitment and adds the commitment to the current state of the ledger. Since the private key is only used for the generation of a signature that serves as private inpupt for the generation of the ZKP, and since the commitment represents the encryption of a hash, the central bank does not learn anything about the onboarded user. The central bank can also detect multiple attempts to create an onboarding transaction, as the commitment is deterministically derived from the ID card.

It is important to note that the anonymity of the onboarding process is *not* required to guarantee the anonymity of a user's subsequent transactions in the privacy pool due to the unlinkability of commitments and nullifiers. Consequently, the central bank could even demand that users that register on the privacy pool present some of their identity attributes, e.g., their nationality or the issuing authority of their digital ID.

Notably, it is not even necessary to check whether the commitment has already been used for previous onboarding to ensure that every user only has a single account: the commitment can only be spent once, independent on how often it is included in the Merkle tree because the corresponding nullifier is also unique. Nevertheless, to avoid attacks that spam the ledger with correct yet useless transactions, it may be useful to check for such collisions.

4.3.2. Semi-private transfers

Semi-private transfers describe the exchange of funds between an account in the privacy pool and a transparent account. These transfers include deposits and withdrawals from the same user so that a user can transfer CBDCs from their transparent account to their privacy pool account, or vice versa. In a semi-private transfer, a user combines an update of their account in the privacy pool with an update of their transparent account that is confirmed by the PSP. As it has to be ensured that the total money supply in the system is unchanged when money is deposited or withdrawn, the transaction amount, i.e., the difference in balances between the spent and created account state, must be disclosed to check whether it is equal to the counter-transaction on the transparent side.

In the following, we consider the depositing process as an example of a semi-private transfer. In more detail, the ZKP would be specified as follows:

Electronic copy available at: https://ssrn.com/abstract=3891121

Semi-private transfer

INPUT: Key pair, Merkle path of the previous commitment,

previous account state, amount

OUTPUT: Merkle root, nullifier, new commitment, (deposit/withdrawal) amount

CHECKING THAT:

- *the previous commitment is contained in the tree represented by the Merkle root*
- *the previous account state belongs to the previous commitment*
- *the nullifier equals the hash of the previous account state*
- *the new account state is correct (e.g., new balance = old balance + amount)*
- *the new account state complies with the rules*

  (e.g., positive balance, epoch turnover below turnover limit)
- *the commitment equals the hash of the signed new account state*

First, the user creates a new commitment and nullifier and attaches a ZKP proving that the account update is legitimate and corresponds to the public outputs. Second, the central bank verifies the ZKP and checks if the public outputs match the requirements, i.e., whether the Merkle root specified by the public outputs matches the Merkle tree of commitments in the ledger, whether the nullifier is not already included in the nullifier list maintained by the central bank, and whether the amount equals the transparent counter-transaction's amount. If all these requirements are satisfied, the central bank adds the new commitment and the nullifier to the ledger and notifies the client about the successful transaction.

The Merkle tree of commitments is append-only, and the mechanism that protects against double-spending, i.e., using the same old account state multiple times, is facilitated by the list of nullifiers. Consequently, the Merkle root does not necessarily have to be the most recent one. This allows users to create transactions that are accepted even if the Merkle root of the ledger has changed in the meantime due to a transaction by another user. Specifically, this option decreases the computational burden for the central bank, as it is sufficient to recompute the Merkle tree in larger epochs. Nevertheless, the epochs should not be too long, as the sequential processing of transactions by the same user requires a new Merkle tree that includes their previous commitment.

Processing this transaction, the central bank inevitably learns the amount of CBDC that is transferred from a transparent account to the privacy pool. However, since the

Electronic copy available at: https://ssrn.com/abstract=3891121

commitments are unlinkable, it is not possible for anyone except the holder of the account to further trace these CBDC units. The same applies to the nullifier that invalidates a previous account state from the Merkle tree of commitments without revealing which specific commitment it refers to. Thus, it is impossible to determine whether a specific semi-private transaction (deposit, withdrawal) was already followed by another semi-private or fully private transaction and, particularly, whether a deposit has already been spent.

In addition to using the transparent CBDC accounts for depositing money in the privacy pool, a user could use CBDC-specific automated teller machines (ATMs) to deposit cash into the privacy pool: The user would create a transaction proposal that contains the desired amount and a ZKP similar to the depositing process described above and send it to the ATM, e.g., via Bluetooth, Wi-Fi Direct, or NFC. Then, the user inserts the corresponding amount of cash directly into the ATM. The ATM confirms the receipt of cash and forwards the user's commitment, nullifier, and ZKP to the central bank. The central bank processes the data as described above. However, the user's transparent CBDC account is not involved in this case, making the depositing and withdrawal process anonymous. Instead, the transaction is conducted via the CBDC account linked to the ATM, which could be provided by a PSP or the central bank directly.

4.3.3. Fully private transfers

The transfer of funds in the privacy pool enables the private bilateral exchange of CBDC. A transfer consists of two transaction proposals, i.e., an individual payment instruction provided by both the sender and receiver. First, the sender and receiver need to agree on the amount of CBDC that should be transferred and on a common randomly generated number, i.e., a nonce, to increase entropy and to be able to link the two update proposals. Next, each participant creates their individual transaction proposal, including the corresponding commitment, nullifier, and ZKP, as they would in the semi-private case, with the difference that the transaction amount is not revealed but hidden by a hash (salted with the nonce). In detail, the ZKP of each of the two involved users looks as follows:

Electronic copy available at: https://ssrn.com/abstract=3891121

Fully private transfer

INPUT: Key pair, Merkle path of a previous commitment,

previous account state, (transaction) amount, nonce, role (sender/receiver)

OUTPUT: Merkle root, nullifier, new commitment, hash of amount | nonce, role

CHECKING THAT:

- *the previous commitment is contained in the tree represented by the Merkle root*
- *the previous account state belongs to the previous commitment*
- *the nullifier equals the hash of the previous account state*
- *the new account state is correct*

  (e.g., new turnover = old turnover + amount if the role is "sender")
- *the new account state complies with the rules*

  (e.g., positive balance, epoch turnover below turnover limit)
- *the commitment equals the hash of the signed new account state*
- *the hash of amount | nonce was computed correctly*

Either the sender or the receiver batches both individual transaction proposals (including their public outputs and ZKP) and sends them to the central bank. Afterward, the central bank validates the integrity of this data through verifying the ZKP that the user generated and compares the outputs with the current state of the ledger. In particular, the central bank checks whether both Merkle roots correspond to the Merkle root of the current ledger and that the two nullifiers are not yet part of the nullifiers' list. Then, the central bank verifies whether the hashes of the value concatenated with the nonce are the same in both transactions. This check guarantees that the amount deducted from one account is equal to the amount added to the other account, without the central bank learning the actual amount due to the pre-image resistance of hash functions. Furthermore, the central bank verifies that the roles specified in the two update proposals are different, i.e., there is one sender and one receiver. Finally, the central bank adds the two new commitments to the Merkle tree and the two nullifiers to the list and notifies the client application about the successful transfer.

Overall, the central bank only learns that a valid transaction was conducted and which two new commitments and nullifiers are involved. However, these hashes are neither related to each other in any further way nor could they be associated with the

Electronic copy available at: https://ssrn.com/abstract=3891121

hashes of previous or future transactions due to the unlinkability guarantees achieved through the system's construction based on commitments, nullifiers, and ZKPs. Thus, the transfers facilitated by the UAS do not provide any information about the individuals themselves, their account balances, or the amount of the transaction, and are therefore fully private.

5. Evaluation

As suggested by Hevner et al. (2004), Peffers et al. (2007), and Peffers, Tuunanen, and Niehaves (2018), we asked key stakeholders to evaluate our IT artifact and to assess the practical feasibility of a CBDC design that provides cash-like privacy using ZKPs. To obtain valuable insights from expert interviews, we followed the recommendations for conducting qualitative interviews by Myers and Newman (2007). In this context, we minimized social dissonances between the researcher's team and the interviewees by first introducing all interview participants to each other. To obtain a rich collection of perspectives, we talked to professionals from various backgrounds, including experts from regulation, cryptography, central banking, identity, and payments. We also made use of specific interview models, such as the waterfall technique, by first motivating our research project, providing context and general definition. Next, we provided a high-level architecture overview and a technical deep dive to improve disclosure in our interviews.

We also discussed the potential risks that may occur using our CBDC with experts and demonstrate adequate mitigation measures. In addition, we presented our overall CBDC system, including the onboarding procedure, semi-private, and fully private transfers. In total, we conducted 22 interviews with a total of 44 international experts with profound expertise in the fields of law, economics, technology, and others (see Table 2) to reinforce the rigor of our methodological approach. Additionally, we sought to gain insights from key stakeholders regarding their key requirements for a CBDC and to receive feedback on our proposed CBDC system. Both the diversity and CBDC-specific expertise of the interviewees provided us with valuable feedback to iteratively adjust and improve our CBDC system. In addition, feedback was collected via internal discussions within the research team's organizations and presentations at various events to test the general feasibility of our approach, thereby continuously improving the artifact (Peffers et al., 2007).

During the first evaluation cycle, the experts confirmed that our design is highly innovative, that it addresses the relevant need for privacy, and that the implemented

Electronic copy available at: https://ssrn.com/abstract=3891121

per-transaction and turnover limits on fully private transactions fit smoothly into existing regulatory frameworks from an AML and CFT compliance perspective. Many experts were surprised or even impressed by the technical capabilities of ZKPs and the maturity of the design (experts 1, 2, 3, 5, 6). In addition, the capability of our CBDC system to flexibly accommodate possible future regulatory changes, e.g., dynamically adaptable thresholds, was considered highly valuable (expert 1). Expert 2 confirmed that our approach of enforcing turnover limits on anonymous transactions would be in line with the 5th European Anti-Money Laundering Directive. The expert further stressed that tracing these small-scale transactions is neither required nor desirable from a law enforcement perspective. Expert 3 emphasized that, in some jurisdictions, mistrust towards the government is considerably high, indicating that a privacy-by-design approach may be helpful to improve broad adoption of a digital currency. He also pointed out that, for the usability of the payment system, it is important to reconcile and settle transparent transactions seamlessly when the limits in the privacy pool are exceeded. However, experts 4 and 5 noted that, although our adaption of the Zcash architecture is a "good trick", it may still be challenging to convince skeptical users that a solution based upon this model is, in fact, fully private. Against this background, providing the code open-source to facilitate independent audits by cryptographers and consumer protection organizations and taking large educational efforts may be required to increase trust in such a cryptography-based privacy-by-design solution.

Experts 3 and 5 noted that our design could also interact smoothly with the current account-based banking systems. Moreover, expert 6 considered our approach of "digital cash" to be intuitive, particularly illustrated through the possibility of depositing and withdrawing CBDC via an ATM. Expert 6 also mentioned that our design, in which the central bank performs the highly automatable task of verifying ZKPs and maintaining the ledger but does not need to conduct resource-intensive KYC processes, fits the roles that central banks aim for in CBDC systems (see, e.g., European Central Bank (2020a)).

However, the experts in the first evaluation cycle also raised two major concerns related to identity management and addressing the needs of corporations additional to those of private end-users, as they may require different limits and identity concepts. Specifically, expert 1 noted that the concept needs to be able to handle the recovery of funds in the case of theft or lost access to the mobile phone. Also, it must provide a way to disable a privacy pool account in the case of blacklisting, e.g., if an individual is put on a sanctions list. Moreover, while our approach to provide one wallet for both a digital identity and payments was considered efficient and appealing from a

Electronic copy available at: https://ssrn.com/abstract=3891121

user perspective (experts 6 and 7), experts 1, 6, and 7 consider the identity-based concept complex and difficult to implement in practice. According to the experts, it is unlikely that such a digital identity can be bootstrapped in the short-term, and expert 7 even considered the availability of a standardized, unique digital identity across multiple European member states a task that is as complex as introducing a CBDC. According to the expert, a particular challenge is that many citizens in the EU have multiple nationalities and, thus, various ID cards, which makes ID cards less suitable for guaranteeing that any citizen can only open and control one account in the privacy pool. Consequently, experts 1, 3, and 7 suggest adding intermediary-based onboarding procedures as another venue to our design.

We incorporated this feedback in our design through the following modifications:

- We added the opportunity for a joint (centralized or decentralized) ledger managed by PSPs that contains identifiers that are deterministically derived from citizens' identity, such as Hash(first name | last name | date of birth). A PSP can then sign an onboarding transaction that proves the possession of a digital certificate issued by the PSP instead of an ID card. As we already discussed in Section 4.3, the anonymity guarantees of private payments do not depend on the privacy of the onboarding process in our solution. Consequently, while the PSP learns that a specific user has registered for the privacy pool, this approach does not compromise the opportunity to conduct fully private transactions.
- We incorporate periodic proofs of non-expiration and non-revocation of the ID card (or the PSP-provided digital certificate) into our design whenever the epoch reset is performed. This modification ensures that once per epoch the user needs to prove that their ID card is still valid and hence allows to block accounts connected to an ID card that has been already revoked, e.g., in the event of loss or theft or the inclusion of an individual on a sanctions list.

In the next cycle, the experts in interview 7 pointed out that metadata, such as IP addresses, need to be taken into account for analysing privacy, and hence that pseudonymization is not sufficient to ensure privacy. This diagnosis confirms our path of ensuring the perfect unlinkability of transactions. They further noted that the *verifiability* of transactions is an important feature, i.e., a user should have the opportunity to prove that a payment did indeed happen, e.g., in the case of a lawsuit. In fact, our design already provides this capability, as a user stores their account history on their local device and can consequently reveal two consecutive previous account states in combination with the transaction confirmation signed by the central bank to demon-

Electronic copy available at: https://ssrn.com/abstract=3891121

strate that a payment was indeed conducted with the claimed details. Furthermore, the bilateral communication between the sender and recipient preceding a transaction, where they agree upon, a transaction amount and a transaction ID, can also be used to make the parties accountable bilaterally if desired by the involved parties, e.g., by revealing parts of their ID cards to the counterparty. The experts also appreciated the capability of our approach to account for embargo lists that prevent a sanctioned user from registering and further using their account through periodic checks of expiration and non-revocation of their ID card.

In interview 9, one expert noted that a balance limit might not be necessary, as a transaction cannot be larger than the turnover in a specific epoch. Nevertheless, the feasibility of balance, per-transaction, and turnover limits can account for the particular needs of various regulators. For instance, if any of these limits is not required in a particular jurisdiction, our system can be implemented easily without such a limit. The experts also pointed out that reversing a transfer must be possible if both parties agree to do so. Moreover, they emphasized the importance to publish the source code of a solution to be able to gain broad acceptance by the public and allow for auditing by consumer protection organizations. This in turn, will, ultimately increase trust in centrally-operated payment systems and also address the concerns regarding education on privacy-by-design approaches as raised by experts 4 and 5. Considering the opportunity to prove that a payment happened, the experts in interview 9 and expert 19 added that, while this proof is indeed a desirable feature, it is also crucial to implement the possibility to delete these records to prevent measures that aim to force users to reveal them (e.g., considering coercive detention or even torture).

While confirming that a limit-based approach is suitable to make fully private payments regulatorily compliant, expert 19 argued that the current limits for cash are relatively low, and further reducing these limits may be difficult to justify, particularly in view of privacy-oriented regulatory norms (see Section 1) and "shadow economies" that may appear when limits are too low to be practical. Moreover, the expert acknowledged that the compliance by design may even help justify higher limits and that, compared to the asymmetric privacy approach by Chaum et al. (2021), our design provides true cash-like privacy. Expert 20 confirmed the suitability of ZKPs for a privacy-oriented yet regulatorily compliant form of money from a technical perspective. The expert also pointed to similar approaches based on the Ethereum blockchain that aim to create private, account-based forms of money but yet do not address regulatory constraints. He also emphasized that due to the complexity of the privacy-oriented cryptographic

Electronic copy available at: https://ssrn.com/abstract=3891121

tools, such as ZKPs, only a few stakeholders that work on CBDC are indeed aware of their technical capabilities.

An issue that was already briefly mentioned by experts 2 and 3 in the first cycle and also caused intensive discussions with experts 12, 15, 18, and 19 in the second cycle refers to the integration of businesses in our CBDC system. A business should not be allowed to spend large amounts of money in the privacy pool as it could potentially evade documentation, such as avoiding tax declarations, or support money laundering on large scale. We discussed two different options to incorporate businesses in our CBDC design: The first approach is to allow them to open a private account via a business ID, e.g., provided by tax authorities or ingrained in trade registers. Our design is sufficiently flexible in assigning other limits to businesses and distinguishing between receiving and spending transactions or withdrawals of a business' privacy pool account to their transparent account. The second approach is based on the concept of asymmetric privacy considered in Chaum et al. (2021) and Tinn and Dubach (2021), tailored to business-to-customer interactions. Indeed, by extending the semi-transparent transactions to not only include deposits and withdrawals between the accounts of one single entity, it is possible to hide the identity of the buyer and only disclose the transaction amount and the receiving business. One essential advantage of this approach may be the opportunity to be able to choose whether to keep the sender's or the receiver's identity private. For example, when the purchase of a potentially embarrassing but legal product needs to be refunded, it may be desirable to hide the identity of the end-user that previously paid anonymously.

The third design cycle further confirmed the suitability of a ZKP-based system for enabling fully private transactions and aligning with regulatory constraints through enforcing limits (experts 22, 25, and 27). Besides, expert 25 emphasized that in general, ZKPs are a "perfect tool" to align privacy and compliance requirements, but also acknowledged that, from a cryptographic perspective, it may be useful to use other forms of ZKPs, such as zk-STARKs (Ben-Sasson et al., 2018), in a potential implementation, specifically as they are regarded post-quantum secure. Expert 28 highlighted that a flexible design that allows for remuneration of CBDC deposits is desirable from a central bank perspective, which illustrates that basic programmability features based on our account-based design is advantageous. Expert 23 also acknowledged that, in contrast to solutions such as the ECB's anonymity vouchers (European Central Bank, 2019), our design can provide true, cash-like privacy. On the other hand, expert 22 expressed concerns that the presumably low limits for fully private transactions might

Electronic copy available at: https://ssrn.com/abstract=3891121

imply that, despite a small share of transactions being fully private, most transactions will eventually be transparent, and hence privacy is not considerably improved overall.

One focus of this design cycle was the mitigation of risks that can potentially arise from our design. On the one hand, expert 22 warned that criminals could abuse the fully private payment system by getting access to several user accounts by purchasing accounts from other users on a black market or via blackmailing. In such a case, the effective limits could be circumvented by possessing a considerable number of accounts in the privacy pool. Although these problems can be mitigated through digital IDs that are bound to secure hardware (expert 22) and connecting the ID to other ID systems is a "great idea" (expert 23), expert 22 still pointed out that the use of secure hardware for storing keys conflicts with recovery capabilities. Experts 28 and 29 employed with a central bank also raised security concerns, particularly related to the high level of obfuscation of transaction-related information inside the privacy pool. Specifically, central banks require strong guarantees that, even if the implementation of the ZKP has a security gap, the extent to which this gap allows illicit activities or harms the monetary system remains marginal. In this regard, they also referred to an implementation error in Zcash that was detected only in 2019 and that would have allowed to "create money out of thin air" (Fortune, 2019).

We addressed the concerns raised in the third cycle by

- conceptualizing backup capabilities with the use of secure hardware through the precautionary creation of a transaction that withdraws all funds that a user has deposited from the privacy pool to their transparent account. The user can then store this recovery backup in the cloud, potentially in encrypted form. This transaction does not provide a proof of non-revocation and non-expiration, as this would quickly be outdated, but these proofs can be given through an in-person visit to the PSP that manages the recovery.
- mitigating risks by
    - pointing out the all-or-nothing transferability that the combination with a digital ID bound to secure hardware allows, i.e., if a user wants to sell their private account, this implies that they need to give away their complete digital identity, meaning that they can no longer use their digital ID card, including all credentials associated with it, such as digital credit cards, diplomas, or health insurances that are bound to this ID.
    - periodically closing and clearing the privacy pool so that users need to transfer their private funds to their transparent account. This process is

Electronic copy available at: https://ssrn.com/abstract=3891121

also helpful to improve performance, as the Merkle tree and the list of commitments do not need to grow quasi infinitely.

- suggesting a hybrid approach with moderate limits for fully private transactions that are primarily meant for customer-to-customer interactions and for semi-private transactions that are primarily meant for business-to-customer interactions.

By implementing these changes, the central bank can detect some of the most critical issues imposed by potential flaws in the implementation of ZKPs, e.g., if after closing the privacy pool, the money supply would be higher than expected. Via small limits for fully private payments, the severity of the impact of selling accounts, which cannot be excluded with certainty, can be mitigated.

As our first three design cycles raised the concern that the most fragile component of our construction is its dependence on a universal digital ID, we conducted another design cycle that – besides discussing the appropriateness of our risk mitigation measures – specifically includes experts on digital ID systems. Expert 30 agreed that our proposed risk mitigation strategies and the combination with a digital ID might be a useful approach, although the integration with secure hardware may be considered inflexible from an end-user perspective, as they cannot access their privacy pool from multiple devices. He also suggested checking the validity of the user's ID in every transaction by default to prevent misuse. Moreover, the experts in interview 19 employed at a federal printer in an advanced economy confirmed that the digital ID-based approach is not only more elegant but will also be possible relatively soon as many federal printers are working on digital IDs and general-purpose ZKPs that are used in such systems.

Within this group, expert 35 confirmed that our proposed backup capabilities seem suitable to recover funds when losing the mobile phone. The experts also stressed that the first digital IDs that integrate with secure hardware on users' devices will be available soon and that they are potentially the only way to efficiently impede the theft of digital IDs or their sharing or selling on a black market. The expert also found the combination of fully private transactions between end-users and semi-private transactions between end-users and businesses very suitable. Moreover, expert 35 pointed out that the security of implementing ZKPs has increased substantially due to cryptographic progress in the last years, so ZKP-related security gaps discussed previously are relatively unlikely today if state-of-the-art guidelines are considered. Further, the group noted that detecting security flaws in well-audited ZKP is significantly less promising

Electronic copy available at: https://ssrn.com/abstract=3891121

for criminals than counterfeiting paper-based money. Finally, expert 40 confirmed that our risk mitigation design may be a promising proposal to consider for central banks and could not identify obvious shortcomings. He particularly appreciated the coupling to a hardware-bound digital identity. However, as with every design proposal for a critical infrastructure, he emphasized that a thorough risk analysis would be required that is beyond the scope of this work. Nevertheless, the expert highlighted that our proposal is well presented and visualized also for non-technical experts, presenting a solid discussion basis for more in-depth analyses, and that our ZKP-based flexible design is a promising approach for the future. Experts 41 and 42 confirmed the suitability as well, emphasizing the insufficient communication between institutions that work on CBDC designs and the state of the art in cryptographic research, so the knowledge transfer via our DSR approach is highly valuable. Furthermore, experts 41 and 42 confirmed that our design incorporates privacy by design well and poses an attractive solution from the perspective of privacy-seeking users. Experts 40 and 41 also noted that, although our software-based solution cannot provide offline payments, our proposal avoids the inherent centralization of risks that comes with a hardware-based approach that may be particularly relevant in regions that do not have local businesses that develop secure hardware (expert 41). For example, there is currently no provider for secure hardware on smartphones where the manufacturing takes place in Europe. Expert 41 also pointed out that abuse cannot be excluded completely even when using hardware-bound digital identities. For instance, users could install proxies on their smartphone that transact on behalf of criminals. Yet, abuse is considerably more complex to organize than it is today with cash.

Since no considerable needs for improvements were brought forward by the experts in the fourth design cycle, and the experts interviewed in this cycle represented diverse fields, including central banks, digital identity issuers, and academics with interdisciplinary expertise from cryptography and economics, we concluded that we have reached a high level of saturation in the design and development of our artifact (Peffers et al., 2007). The critical feedback we received from key stakeholders positively influenced the design of our CBDC system, allowing us to continuously improve our artifact and thus to ultimately answer our research question.

6. Conclusion

In this paper, we followed a rigorous DSR approach to develop and evaluate a CBDC system based on an unspent-account-state model. To this end, we make use of re-

Electronic copy available at: https://ssrn.com/abstract=3891121

cent advancements in cryptography, especially related to ZKPs. Contrary to common beliefs (e.g. Armelius et al., 2021; Auer & Boehme, 2021), we demonstrate that a software-based CBDC system can support full privacy while addressing constraints related to AML and CFT regulation by imposing limits on anonymous payments. In our system, both privacy and compliance are provided *by design*, i.e., end-users do not have to trust third parties for preserving privacy and conducting compliance checks, as transaction data are stored only on the end-users' devices (trustless privacy). We assess the feasibility and suitability of our technical artifact in 22 interviews with 44 leading experts from various fields, including regulation, computer science, cryptography, central banking, and payments.

In addition to enabling full privacy and regulatory compliance, our ZKP-based CBDC system provides novel applications for end-users that go beyond existing forms of fully private money and especially cash. For instance, when faced with a legal accusation such as money laundering, our system enables users to reveal their payment history and provide evidence for its integrity and completeness. Thereby, users can address accusations, e.g., by proving that a payment did or did not happen.

Our proposed solution can also be used in more general applications, e.g., for account-based IT artifacts that enable full privacy by design while addressing certain compliance requirements. For example, future systems for documenting and exchanging carbon certificates between organizations and/or individuals will likely require high privacy guarantees, as well. Moreover, our solution can be applied in decentralized settings. Illustrated by the progress in the fields of blockchain technology in general and decentralized finance in particular, one can observe a tendency towards decentralizing the financial system. Specifically, tokenization promises to simplify the exchange of ownership (Sunyaev et al., 2021), but so far, existing blockchain-based solutions for managing and exchanging these tokens have not focused on stakeholders' (or even regulatory, e.g., the general data protection regulation (GDPR)) privacy requirements and regulation has rather focused on service providers that provide access to token-related services than end-users directly. Privacy-by-design approaches are the only way to establish full privacy in such blockchain-based systems because there are, by definition, no trusted third parties, and data is replicated among many different nodes. Our approach of leveraging ZKPs makes them a digital substitute for physical hardware-based approaches without a single point of failure to achieve integrity and compliance while storing data only on end-users' devices, as the correct local accounting is ensured through providing cryptographic proofs instead of hardware-based attestation.

Electronic copy available at: https://ssrn.com/abstract=3891121

Detecting an error in a well-audited cryptographic protocol seems less likely than a successful attack on a single secure hardware device.

Our artifact provides a starting point that balances privacy and compliance and may provide many avenues for future research. So far, we have presented our design to experts and instantiated the core logic in a technical implementation. Expert 22 also stated that it is important to focus on key requirements as most design criteria tend to have trade-offs and that it is almost impossible in one research project to implement interfaces to end-users and businesses. Thus, our focus on full (cash-like) privacy and regulatory compliance is a reasonable first step. However, there are additional important CBDC design dimensions, including security, scalability, and cost, that have to be considered. Thus, there is a need for future research for a rigorous evaluation of the extent to which our IT artifact can potentially also address these other important CBDC design dimensions. Specifically, besides more detailed analyses on performance, future analyses may involve other aspects related to user experience. For example, the implications of the added complexity of a two-tiered approach, the limited recovery options, and the integration of multiple devices require innovative solutions that shield complexity from the user and survey their impact on usability. Finally, future research could study the interplay of our design with potential extensions, such as using secure hardware for facilitating offline payments, that may, however, potentially come with restricted privacy guarantees to account for mitigating the related security challenges.

Electronic copy available at: https://ssrn.com/abstract=3891121

References

Acquisti, A., Taylor, C., & Wagman, L. (2016). The economics of privacy. *Journal of Economic Literature*, *54*(2), 442–92.

Allen, S., Čapkun, S., Eyal, I., Fanti, G., Ford, B. A., Grimmelmann, J., ... others (2020). *Design choices for central bank digital currency: Policy and technical considerations* (Tech. Rep.). National Bureau of Economic Research. Retrieved 2022-01-10, from https://www.nber.org/system/files/working_papers/w27634/w27634.pdf

Armelius, H., Claussen, C. A., & Hull, I. (2021). *On the possibility of a cash-like CBDC.* Retrieved 2022-01-10, from https://www.riksbank.se/globalassets/media/rapporter/staff-memo/engelska/2021/on-the-possibility-of-a-cash-like-cbdc.pdf

Auer, R., & Boehme, R. (2021). *Central bank digital currency: The quest for minimally invasive technology.* Retrieved 2022-01-10, from https://www.bis.org/publ/work948.pdf

Auer, R., & Böhme, R. (2020). *The technology of retail central bank digital currency.* Retrieved 2022-01-10, from https://www.bis.org/publ/qtrpdf/r_qt2003j.pdf

Bank for International Settlements, Bank of Canada, European Central Bank, Bank of Japan, Sveriges Riksbank, Swiss National Bank, ... Federal Reserve Board of Governors (2020). *Central bank digital currencies: Foundational principles and core features.* Retrieved 2022-01-10, from https://www.bis.org/publ/othp33.pdf

Bank of Canada. (2020). *Contingency planning for a central bank digital currency.* Retrieved 2022-01-10, from https://www.bankofcanada.ca/2020/02/contingency-planning-central-bank-digital-currency/

Bank of England. (2020). *Central bank digital currency: Opportunities, challenges and design.* Retrieved 2022-01-10, from https://www.bankofengland.co.uk/paper/2020/central-bank-digital-currency-opportunities-challenges-and-design-discussion-paper

Bech, M. L., & Garratt, R. (2017). Central bank cryptocurrencies. *BIS Quarterly Review September*.

Ben-Sasson, E., Bentov, I., Horesh, Y., & Riabzev, M. (2018). *Scalable, transparent, and post-quantum secure computational integrity.* Retrieved 2022-01-10, from https://eprint.iacr.org/2018/046.pdf

Ben-Sasson, E., Chiesa, A., Garman, C., Green, M., Miers, I., Tromer, E., & Virza, M. (2014). Zerocash: Decentralized anonymous payments from Bitcoin. In *IEEE symposium on security and privacy* (pp. 459–474).

Ben-Sasson, E., Chiesa, A., Genkin, D., Tromer, E., & Virza, M. (2013). SNARKs for C: Verifying program executions succinctly and in zero knowledge. In *Annual cryptology conference* (pp. 90–108).

Biryukov, A., Khovratovich, D., & Pustogarov, I. (2014). Deanonymisation of clients in Bitcoin P2P network. In *Proceedings of the 2014 ACM SIGSAC conference on computer*

Electronic copy available at: https://ssrn.com/abstract=3891121

and communications security (p. 15–29). ACM.

Boar, C., & Wehrli, A. (2021). *Ready, steady, go? – Results of the third BIS survey on central bank digital currency.* Retrieved 2022-01-10, from https://www.bis.org/publ/bppdf/bispap114.pdf

Bok, S. (1989). *Secrets: On the ethics of concealment and revelation.* Random House Digital, Inc.

Bontekoe, T. (2020). *Balancing privacy and accountability in digital payment methods using zk-SNARKs* (Master's thesis, University of Twente). Retrieved 2022-01-10, from http://essay.utwente.nl/83617/1/Bontekoe_MA_EEMCS.pdf

Camenisch, J., Hohenberger, S., & Lysyanskaya, A. (2006). Balancing accountability and privacy using e-cash. In *International conference on security and cryptography for networks* (pp. 141–155).

Chaum, D. (1983). Blind signatures for untraceable payments. In *Advances in cryptology* (pp. 199–203).

Chaum, D., Grothoff, C., & Moser, T. (2021). *How to issue a central bank digital currency.* Retrieved 2022-01-10, from https://arxiv.org/ftp/arxiv/papers/2103/2103.00254.pdf

Cheng, J., Lawson, A. N., & Won, P. (2021). *Preconditions for a general-purpose central bank digital currency.* Retrieved 2022-01-10, from https://www.federalreserve.gov/econres/notes/feds-notes/preconditions-for-a-general-purpose-central-bank-digital-currency-20210224.htm

Choi, K. J., Henry, R., Lehar, A., Reardon, J., & Safavi-Naini, R. (2021). *A proposal for a Canadian CBDC.* Retrieved 2022-01-10, from https://papers.ssrn.com/sol3/Delivery.cfm/SSRN_ID3786426_code244292.pdf?abstractid=3786426&mirid=1

Dinev, T., Xu, H., Smith, J. H., & Hart, P. (2013). Information privacy and correlates: an empirical attempt to bridge and distinguish privacy-related concepts. *European Journal of Information Systems*, *22*(3), 295-316.

Dingledine, R., Mathewson, N., & Syverson, P. (2004). Tor: The second-generation onion router. In *13th USENIX security symposium.* USENIX Association.

Dold, F. (2019). *The GNU Taler system: Practical and provably secure electronic payments* (Doctoral dissertation, Université Rennes 1). Retrieved 2022-01-10, from https://tel.archives-ouvertes.fr/tel-02138082/file/DOLD_Florian.pdf

European Central Bank. (2019). *Exploring anonymity in central bank digital currencies.* Retrieved 2022-01-10, from https://www.ecb.europa.eu/paym/intro/publications/pdf/ecb.mipinfocus191217.en.pdf

European Central Bank. (2020a). *Report on a digital euro.* Retrieved 2022-01-10, from https://www.ecb.europa.eu/euro/html/digitaleuro-report.en.html

European Central Bank. (2020b). *Study on the payment attitudes of consumers in the euro area.* (https://www.ecb.europa.eu/press/pr/date/2020/html/ecb

Electronic copy available at: https://ssrn.com/abstract=3891121

.pr201202~0645677cf6.en.html, accessed 2021-04-06)

European Central Bank. (2021a). *A digital euro to meet the expectations of Europeans.* Retrieved 2022-01-10, from https://www.ecb.europa.eu/press/key/date/2021/html/ecb.sp210414_1~e76b855b5c.en.html

European Central Bank. (2021b). *ECB publishes the results of the public consultation on a digital euro.* Retrieved 2022-01-10, from https://www.ecb.europa.eu/press/pr/date/2021/html/ecb.pr210414~ca3013c852.en.html

European Central Bank. (2021c). *Eurosystem launches digital euro project.* Retrieved 2021-07-17, from https://www.ecb.europa.eu/press/pr/date/2021/html/ecb.pr210714~d99198ea23.en.html

European Convention. (2000). *Charter of fundamental rights of the European Union.* Retrieved 2022-01-10, from https://www.europarl.europa.eu/charter/pdf/text_en.pdf

European Court of Human Rights. (1950). *European convention of human rights.* Retrieved 2022-01-10, from https://www.echr.coe.int/documents/convention_eng.pdf

European Union. (2018). *Directive (EU) 2018/843 of the European parliament and of the council.* Retrieved 2022-01-10, from https://eur-lex.europa.eu/legal-content/EN/TXT/HTML/?uri=CELEX:32018L0843&from=EN

Fauzi, P., Meiklejohn, S., Mercer, R., & Orlandi, C. (2019). Quisquis: A new design for anonymous cryptocurrencies. In *International conference on the theory and application of cryptology and information security* (pp. 649–678).

Fortune. (2019). *Zcash discloses vulnerability that could have allowed 'infinite counterfeit' cryptocurrency.* Retrieved 2022-01-10, from https://fortune.com/2019/02/05/zcash-vulnerability-cryptocurrency/

Garman, C., Green, M., & Miers, I. (2016). Accountable privacy for decentralized anonymous payments. In *International conference on financial cryptography and data security* (pp. 81–98).

Garratt, R. J., & Van Oordt, M. R. (2021). Privacy as a public good: A case for electronic cash. *Journal of Political Economy*, *129*(7).

Goldwasser, S., Micali, S., & Rackoff, C. (1989). The knowledge complexity of interactive proof systems. *SIAM Journal on Computing*, *18*(1), 186–208.

Gregor, S., & Hevner, A. (2013). Positioning and presenting design science research for maximum impact. *MIS Quarterly*, *37*, 337-355.

Grothoff, C., & Dold, F. (2021). *Why a digital euro should be online-first and bearer-based.* Retrieved 2021-07-18, from https://taler.net/papers/euro-bearer-online-2021.pdf

Hevner, A. R., March, S. T., Park, J., & Ram, S. (2004). Design science in information systems research. *MIS Quarterly*(1), 75–105.

Iivari, J. (2015). Distinguishing and contrasting two strategies for design science research. *European Journal of Information Systems*, *24*(1), 107-115.

Kahn, C. M. (2018). *Payment systems and privacy.* Retrieved 2022-01-10, from https://

Electronic copy available at: https://ssrn.com/abstract=3891121

papers.ssrn.com/sol3/papers.cfm?abstract_id=3315945

Kahn, C. M., McAndrews, J., & Roberds, W. (2005). Money is privacy. *International Economic Review*, *46*(2), 377–399.

Kannengießer, N., Lins, S., Dehling, T., & Sunyaev, A. (2020). Trade-offs between distributed ledger technology characteristics. *ACM Computing Surveys*, *53*(2).

Kiff, J., Alwazir, J., Davidovic, S., Farias, A., Khan, A., Khiaonarong, T., ... others (2020). *A survey of research on retail central bank digital currency.* Retrieved 2022-01-10, from https://www.imf.org/-/media/Files/Publications/WP/2020/English/wpiea2020104-print-pdf.ashx

Kubach, M., & Sellung, R. (2021). On the market for self-sovereign identity: Structure and stakeholders. In *Open identity summit 2021.* Gesellschaft für Informatik eV.

Lane, T. (2020). *Money and payments in the digital age.* Retrieved 2022-01-10, from https://www.bis.org/review/r200311d.pdf

March, S. T., & Smith. (1995). Design and natural science research on information technology. *Decision Support Systems*, *15*(4), 251–266.

Mastercard. (2021). *Mastercard and island pay launch world's first central bank digital currency-linked card.* Retrieved 2021-07-18, from https://www.mastercard.com/news/press/2021/february/mastercard-and-island-pay-launch-world-s-first-central-bank-digital-currency-linked-card/

Mattke, J., Hund, A., Maier, C., & Weitzel, T. (2019). How an enterprise blockchain application in the US pharmaceuticals supply chain is saving lives. *MIS Quarterly Executive*, *18*(4).

Meiklejohn, S., Pomarole, M., Jordan, G., Levchenko, K., McCoy, D., Voelker, G. M., & Savage, S. (2016). A fistful of Bitcoins: Characterizing payments among men with no names. *Communications of the ACM*, *59*(4), 86–93.

Merkle, R. C. (1987). A digital signature based on a conventional encryption function. In *Conference on the theory and application of cryptographic techniques* (pp. 369–378).

Miers, I., Garman, C., Green, M., & Rubin, A. D. (2013). Zerocoin: Anonymous distributed e-cash from Bitcoin. In *IEEE symposium on security and privacy* (p. 397-411).

Myers, M., & Newman, M. (2007). The qualitative interview in is research: Examining the craft. *Information and Organization*, *17*, 2-26.

Nakamoto, S. (2008). *Bitcoin: A peer-to-peer electronic cash system.* Retrieved 2022-01-10, from http://bitcoin.org/bitcoin.pdf

Odlyzko, A. (2004). Privacy, economics, and price discrimination on the internet. In *Economics of information security* (pp. 187–211). Springer.

Partala, J., Nguyen, T. H., & Pirttikangas, S. (2020). Non-interactive zero-knowledge for blockchain: A survey. *IEEE Access*, *8*, 227945-227961.

Peffers, K., Tuunanen, T., & Niehaves, B. (2018). Design science research genres: Introduction to the special issue on exemplars and criteria for applicable design science research. *European Journal of Information Systems*, *27*(2), 129-139.

Electronic copy available at: https://ssrn.com/abstract=3891121

Peffers, K., Tuunanen, T., Rothenberger, M., & Chatterjee, S. (2007). A design science research methodology for information systems research. *Journal of Management Information Systems*, *24*(3), 45–77.

Pocher, N., & Veneris, A. (2021). *Privacy and transparency in CBDCs: A regulation-by-design AML/CFT scheme.* Retrieved 2022-01-10, from `https://www.eecg.utoronto.ca/~veneris/21icbc2.pdf`

Rossi, M., Mueller-Bloch, C., Thatcher, J. B., & Beck, R. (2019). Blockchain research in information systems: Current trends and an inclusive future research agenda. *Journal of the Association for Information Systems*, *20*(9), 1388–1403.

Sander, T., & Ta-Shma, A. (1999). Flow control: A new approach for anonymity control in electronic cash systems. In *International conference on financial cryptography* (pp. 46–61).

Sedlmeir, J., Smethurst, R., Rieger, A., & Fridgen, G. (2021). Digital identities and verifiable credentials. *Business & Information Systems Engineering*, *63*(5), 603–613.

Silfversten, E., Favaro, M., Slapakova, L., Ishikawa, S., Liu, J., & Salas, A. (2020). *Exploring the use of Zcash cryptocurrency for illicit or criminal purposes.* RAND.

Sunyaev, A., Kannengießer, N., Beck, R., Treiblmaier, H., Lacity, M., Kranz, J., ... Luckow, A. (2021). Token economy. *Business & Information Systems Engineering*, 1–22.

Tefft, S. K. (1980). *Secrecy a cross-cultural perspective.* Human Sciences Press.

The Guardian. (2021). *New Zealand's central bank says its systems have been hacked.* Retrieved 2022-01-10, from `https://www.theguardian.com/world/2021/jan/11/new-zealands-central-bank-says-its-systems-have-been-hacked`

Tinn, K., & Dubach, C. (2021). *Central bank digital currency with asymmetric privacy.* Retrieved 2022-01-10, from `https://www.mcgill.ca/engineering/files/engineering/central_bank_digital_currency_with_asymmetric_privacy_mcgill_tinn_dubach.pdf`

Tramèr, F., Boneh, D., & Paterson, K. (2020). Remote side-channel attacks on anonymous transactions. In *29th {USENIX} security symposium* (pp. 2739–2756).

United Nations. (1948). *Universal declaration of human rights.* Retrieved 2022-01-10, from `https://www.un.org/en/udhrbook/pdf/udhr_booklet_en_web.pdf`

Veneris, A., Park, A., Long, F., & Puri, P. (2021). *Central bank digital loonie: Canadian cash for a new global economy.* Retrieved 2022-01-10, from `https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3770024`

Wüst, K., Kostiainen, K., & Capkun, S. (2021). *Platypus: A central bank digital currency with unlinkable transactions and privacy preserving regulation.* Retrieved 2022-01-10, from `https://eprint.iacr.org/2021/1443`

Zhang, R., Xue, R., & Liu, L. (2019). Security and privacy on blockchain. *ACM Computing Surveys*, *52*(3), 1–34.

Zwick, D., & Dholakia, N. (2004). Whose identity is it anyway? Consumer representation in the age of database marketing. *Journal of Macromarketing*, *24*(1), 31–43.

Electronic copy available at: https://ssrn.com/abstract=3891121