# ZEXE: Enabling Decentralized Private Computation


Sean Bowe
sean@z.cash
Zcash

Alessandro Chiesa
alexch@berkeley.edu
UC Berkeley

Matthew Green
mgreen@cs.jhu.edu
Johns Hopkins University

Ian Miers
imiers@cs.cornell.edu
Cornell Tech

Pratyush Mishra
pratyush@berkeley.edu
UC Berkeley

Howard Wu
howardwu@berkeley.edu
UC Berkeley



First published: October 14, 2018
Current revision: November 6, 2025



**Abstract**

Ledger-based systems that support rich applications often suffer from two limitations. First, validating a transaction requires re-executing the state transition that it attests to. Second, transactions not only reveal which application had a state transition but also reveal the application's internal state.

We design, implement, and evaluate ZEXE, a ledger-based system where users can execute offline computations and subsequently produce transactions, attesting to the correctness of these computations, that satisfy two main properties. First, transactions *hide all information* about the offline computations. Second, transactions can be *validated in constant time* by anyone, regardless of the offline computation.

The core of ZEXE is a construction for a new cryptographic primitive that we introduce, *decentralized private computation* (DPC) schemes. In order to achieve an efficient implementation of our construction, we leverage tools in the area of cryptographic proofs, including succinct zero knowledge proofs and recursive proof composition. Overall, transactions in ZEXE are 968 bytes *regardless of the offline computation*, and generating them takes less than a minute plus a time that grows with the offline computation.

We demonstrate how to use ZEXE to realize privacy-preserving analogues of popular applications: private decentralized exchanges for user-defined fungible assets and regulation-friendly private stablecoins.



**Keywords**: decentralized computation; zero knowledge proofs; succinct arguments

# Contents

# 1 Introduction

Distributed ledgers are a mechanism that maintains data across a distributed system while ensuring that every party has the same view of the data, even in the presence of corrupted parties. Ledgers can provide an indisputable history of all "events" logged in a system, thereby offering a mechanism for multiple parties to collaborate with minimal trust (any party can ensure the system's integrity by auditing history). Interest in distributed ledgers has soared recently, catalyzed by their use in cryptocurrencies (peer-to-peer payment systems) and by their potential as a foundation for new forms of financial systems, governance, and data sharing. In this work we study two limitations of ledgers, one about *privacy* and the other about *scalability*.

**A privacy problem.** The main strength of distributed ledgers is also their main weakness: *the history of all events is available for anyone to read*. This severely limits a direct application of distributed ledgers.

For example, in ledger-based payment systems such as Bitcoin [Nak09], every payment transaction reveals the payment's sender, receiver, and amount. This not only reveals private financial details of individuals and businesses using the system,[1] but also violates fungibility, a fundamental economic property of money. This lack of privacy becomes more severe in smart contract systems like Ethereum [Woo17], wherein transactions not only contain payment details, but also embed function calls to specific applications. In these systems, every application's internal state is necessarily public, and so is the history of function calls associated to it.

This problem has motivated prior work to find ways to achieve meaningful privacy guarantees on ledgers. For example, the Zerocash protocol [BCG+14] provides privacy-preserving payments, and Hawk [KMS+16] enables general state transitions with data privacy, that is, an application's data is hidden from third parties.

However, all prior work is limited to hiding the inputs and outputs of a state transition but not *which* transition function is being executed. That is, prior work achieves *data privacy* but not *function privacy*. In systems with a single transition function this is not a concern.[2] In systems with multiple transition functions, however, this leakage is problematic. For example, Ethereum currently supports thousands of separate ERC-20 "token" contracts [Eth18], each representing a distinct currency on the Ethereum ledger; even if these contracts each individually adopted a protocol such as Zerocash to hide details about token payments, the corresponding transactions would still reveal *which* token was being exchanged. Moreover, the leakage of this information would substantially reduce the anonymity set of those payments.

**A scalability problem.** Public auditability in the aforementioned systems (and many others) is achieved via direct verification of state transitions that re-executes the associated computation. This creates the following scalability issues. First, note that in a network consisting of devices with heterogeneous computing power, requiring every node to re-execute transactions makes the weakest node a bottleneck, and this effect persists even when the underlying ledger is "perfect", that is, it confirms every valid transaction immediately. To counteract this and to discourage denial-of-service attacks whereby users send transactions that take a long time to validate, current systems introduce mechanisms such as *gas* to make users pay more for longer computations. However, such mechanisms can make it unprofitable to validate legitimate but expensive transactions, a problem known as the "Verifier's Dilemma" [LTKS15]. These problems have resulted in Bitcoin forks [Bit15] and Ethereum attacks [Eth16].

In sum, there is a dire need for techniques that facilitate the use of distributed ledgers for rich applications, without compromising privacy (of data or functions) or relying on unnecessary re-executions. Prior works only partially address this need, as discussed in Section 1.2 below.

[1] Even if payments merely contain *addresses* rather than, say, social security numbers, much information about individuals and businesses can be gleaned by analyzing the flow of money over time between addresses [RH11, RS13, AKR+13, MPJ+13, SMZ14, KGC+17]. There are even companies that offer analytics services on the information stored on ledgers [Ell13, Cha14].

[2] For example, in Zerocash the single transition function is the one governing cash flow of a single currency.

## 1.1 Our contributions

We design, implement, and evaluate Zexe (*Zero knowledge EXEcution*), a ledger-based system that enables users to execute offline computations and subsequently produce publicly-verifiable transactions that attest to the correctness of these offline executions. Zexe simultaneously provides two main security properties.

- **Privacy:** *a transaction reveals no information about the offline computation, except (an upper bound on) the number of consumed inputs and created outputs.*[3] One cannot link together multiple transactions by the same user or involving related computations, nor selectively censor transactions based on such information.
- **Succinctness:** *a transaction can be validated in time that is independent of the cost of the offline computation whose correctness it attests to.* Since all transactions are indistinguishable, and are hence equally cheap to validate, there is no "Verifier's Dilemma", nor a need for mechanisms like Ethereum's gas.

Zexe also offers rich functionality, as offline computations in Zexe can be used to realize state transitions of multiple applications (such as tokens, elections, markets) simultaneously running atop the *same* ledger. The users participating in applications do not have to trust, or even know of, one another. Zexe supports this functionality by exposing a simple, yet powerful, *shared execution environment* with the following properties.

- **Extensibility:** users may execute arbitrary functions of their choice, without seeking anyone's permission.
- **Isolation:** functions of malicious users cannot interfere with the computations and data of honest users.
- **Inter-process communication:** functions may exchange data with one another.

**DPC schemes.** The technical core of Zexe is a protocol for a new cryptographic primitive for performing computations on a ledger called *decentralized private computation* (DPC). Informally, a DPC scheme supports a simple, yet expressive, programming model in which units of data, which we call *records*, are bound to scripts (arbitrary programs) that specify the conditions under which a record can be created and consumed (this model is similar to the UTXO model; see Remark 2.3). The rules that dictate how these programs interact can be viewed as a "nano-kernel" that provides a shared execution environment upon which to build applications. From a technical perspective, DPC can be viewed as extending Zerocash [BCG$^+$14] to the foregoing programming model, while still providing strong privacy guarantees, not only within a single application (which is a straightforward extension) but also across multiple co-existing applications (which requires new ideas that we discuss later on). The security guarantees of DPC are captured via an ideal functionality, which our protocol provably achieves.

**Applications.** To illustrate the expressivity of the RNK, we show how to use DPC schemes to construct privacy-preserving analogues of popular applications: private user-defined assets, private decentralized or non-custodial exchanges (DEXs), and private stablecoins. Our privacy guarantees in particular protect against vulnerabilities of current DEX designs such as front-running [BDJT17, BBD$^+$17, EMC19, DGK$^+$20]. Moreover, we sketch how to use DPC to construct a privacy-preserving smart contract system. See Sections 2.3 and 6 for details.

**Techniques for efficient implementation.** We devise a set of techniques to achieve an efficient implementation of our DPC protocol, by drawing upon recent advances in zero knowledge succinct cryptographic proofs (namely, zkSNARKs) and in recursive proof composition (proofs attesting to the validity of other proofs).

[3] One can fix the number of inputs and outputs (say, fix both to 2), or carefully consider side channels that could arise from revealing bounds on the number of inputs and outputs.

Overall, transactions in ZEXE with two input records and two output records are 968 bytes and can be verified in tens of milliseconds, *regardless of the offline computation*; generating these transactions takes less than a minute plus a time that grows with the offline computation (inevitably so). This implementation is achieved in a modular fashion via a collection of Rust libraries (see Fig. 15), in which the top-level one is `libzexe`. Our implementation also supports transactions with *any* number $m$ of input records and $n$ of output records; transactions size in this case is $32m + 32n + 840$ bytes (the transaction stores the serial number of each input record and the commitment of each output record).

**Delegating transactions.** While verifying succinct cryptographic proofs is cheap, producing them can be expensive. As the offline computation grows, the (time and space) cost of producing a cryptographic proof of its correctness also grows, which could become infeasible for a user.

To address this problem, we further obtain *delegable DPC*. The user communicates to an untrusted worker details about the desired transaction, then the worker produces the transaction, and finally the user authorizes it via a cheap computation (and in a way that does not violate indistinguishability of transactions). This feature is particularly relevant for prospective real-world deployments, because it enables support for weak devices, such as mobile phones or hardware tokens.

In fact, our delegable DPC protocol also extends to support *threshold transactions*, which can be used to improve operational security, and also to support *blind transactions*, which can be used to realize lottery tickets for applications such as micropayments.

All of these extensions are also part of our Rust library `libzexe`.

**A perspective on costs.** ZEXE is not a lightweight construction, but achieves, in our opinion, tolerable efficiency for the ambitious goals it sets out to achieve: *data and function privacy, and succinctness, with rich functionality, in a threat model that requires security against all efficient adversaries*. Relaxing any of these goals (assuming rational adversaries or hardware enclaves, or compromising on privacy) will lead to more efficient approaches.

The primary cost in our system is, unsurprisingly, the cost of generating the cryptographic proofs that are included in transactions. We have managed to keep this cost to roughly a minute plus a cost that grows with the offline computation. For the applications mentioned above, these additional costs are negligible. Our system thus supports applications of real-world interest today (e.g., private DEXs) with reasonable costs.

## 1.2 Related work

**Avoiding naive re-execution.** A number of proposals for improving the scalability of smart contract systems, such as TrueBit [TR17], Plasma [PB17], and Arbitrum [KGC+18], avoid naive re-execution by having users report the results of their computations *without* any cryptographic proofs, and instead putting in place incentive mechanisms wherein others can challenge reported results. The user and challenger engage in a so-called *refereed game* [FK97, CRR11, CRR13, JSST16, Rei16], mediated by a smart contract acting as the referee, that efficiently determines which of the two was "telling the truth". In contrast, in this work correctness of computation is ensured by cryptography, regardless of any economic motives; we thus protect against all efficient adversaries rather than merely all rational and efficient ones. Also, unlike our DPC scheme, the above works do not provide formal guarantees of strong privacy (challengers must be able to re-execute the computation leading to a result and in particular must know its potentially private inputs).

**Private payments.** Zerocash [BCG+14], building on earlier work [MGGR13], showed how to use distributed ledgers to achieve payment systems with strong privacy guarantees. The Zerocash protocol, with some modifications, is now commercially deployed in several cryptocurrencies, including Zcash [ZCa15]. Solidus [CZJ+17] enables customers of financial institutions (such as banks) to transfer funds to one another in

a manner that ensures that only the banks of the sender and receiver learn the details of the transfer; all other parties (all other customers and banks) only learn that a transfer occurred, and nothing else. zkLedger [NVV18] enables anonymous payments between a small number of distinguished parties via the use of homomorphic commitments and Schnorr proofs. None of these protocols support scripts (small programs that dictate how funds can be spent), let alone arbitrary state transitions as in Zexe.

**Privacy beyond payments.** Hawk [KMS+16], combining ideas from Zerocash and the notion of an evaluator-prover for multi-party computation, enables parties to conduct offline computations and then report their results via cryptographic proofs. Hawk's privacy guarantee protects the private inputs used in a computation, but does not hide *which* computation was performed. That said, we view Hawk as complementary to our work: a user in our system could in particular be a semi-trusted manager that administers a multi-party computation and generates a transaction about its output. The privacy guarantees provided in this work would then additionally hide *which* computation was carried out offline.

Zether [BAZB20] is a system that enables *publicly known* smart contracts to reason about homomorphic commitments in zero knowledge, and in particular enables these to transact in a manner that hides transaction amounts; it does not hide the identities of parties involved in the transaction, beyond a small anonymity set. Furthermore, the cost of verifying a transaction scales linearly with the size of the anonymity set, whereas in Zexe this cost scales logarithmically with the size of anonymity set.

**Succinct blockchains.** Coda [MS18] uses arbitrary-depth recursive composition of SNARKs to enable blockchain nodes to verify the current blockchain state quickly. In contrast, Zexe uses depth-2 recursive composition to ensure that all blockchain transactions are equally cheap to verify (and are moreover indistinguishable from each other), regardless of the cost of the offline computation. In this respect, Coda and Zexe address orthogonal scalability concerns.

**MPC with ledgers.** Several works [ADMM14b, ADMM14a, KMB15, KB16, BKM17, RCGJ+17] have applied ledgers to obtain secure multi-party protocols that have security properties that are difficult to achieve otherwise, such as *fairness*. These approaches are complementary to our work, as any set of parties wishing to jointly compute a certain function via one of these protocols could run the protocol "under" our DPC scheme in such a way that third parties would not learn any information that such a multi-party computation is happening.

**Hardware enclaves.** Kaptchuk et al. [KGM19] and Ekiden [CZK+18] combine ledgers with hardware enclaves, such as Intel Software Guard Extensions [MAB+13], to achieve various integrity and privacy goals for smart contracts. Beyond ledgers, several systems explore privacy goals in distributed systems by leveraging hardware enclaves; see for example M2R [DSC+15], VC3 [SCF+15], and Opaque [ZDB+17]. All of these works are able to efficiently support rich and complex computations. In this work, we make no use of hardware enclaves, and instead rely entirely on cryptography. This means that on the one hand our performance overheads are more severe, while on the other hand we protect against a richer class of adversaries (all efficient ones). Moreover, the techniques above depend on a working *remote attestation capability*; we note that our techniques can be used to achieve stronger security guarantees, even in the face of a compromise in the remote attestation capabilities of an enclave system (as recently occurred with Intel SGX [VBMW+19]).

# 2 Techniques

We now summarize the main ideas behind our contributions. Our goal is to design a ledger-based system in which transactions attest to offline computations while simultaneously providing *privacy* and *succinctness*.

We first note that if privacy is not required, there is a straightforward folklore approach that provides succinctness and low verification cost: each user accompanies the result reported in a transaction with a succinct cryptographic proof (i.e., a SNARK) attesting to the result's correctness. Others who validate the transaction can simply verify the cryptographic proof, and do not have to re-execute the computation. Even this limited approach rules out a number of cryptographic directions, such as the use of Bulletproofs [BCC$^+$16, BBB$^+$18] (which have verification time linear in the circuit complexity), but can be accomplished using a number of efficient SNARK techniques [GGPR13, BCTV14, BCS16, BCTV17]. In light of this, we shall first discuss how to achieve privacy, and then how to additionally achieve succinctness.

The rest of this section is organized as follows. In Sections 2.1 and 2.2 we explain why achieving privacy in our setting is challenging. In Section 2.3 we introduce the shared execution environment that we consider, and in Section 2.4 we introduce *decentralized private computation* (DPC), a cryptographic primitive that securely realizes it. In Section 2.5 we describe how we turn our ideas into an efficient implementation.

## 2.1 Achieving privacy for a single arbitrary function

Zerocash [BCG$^+$14] is a protocol that achieves privacy for a specific functionality, namely, *value transfers within a single currency*. Therefore, it is natural to consider what happens if we extend Zerocash from this special case to the general case of a *single arbitrary function* that is known in advance to everybody.

**Sketch of Zerocash.** Money in Zerocash is represented via *coins*. The commitment of a coin is published on the ledger when the coin is created, and its serial number is published when the coin is consumed. Each transaction on the ledger attests that some "old" coins were consumed in order to create some "new" coins: it contains the serial numbers of the consumed coins, commitments of the created coins, and a zero knowledge proof attesting that the serial numbers belong to coins created in the past (without identifying which ones), and that the commitments contain new coins of the same total value. A transaction is private because it only reveals how many coins were consumed and how many were created, but no other information (each coin's value and owner address remain hidden). Also, revealing a coin's serial number ensures that a coin cannot be consumed more than once (the same serial number would appear twice). In sum, data in Zerocash corresponds to coin values, and state transitions are the single invariant that monetary value is preserved.

**Extending to an arbitrary function.** One way to extend Zerocash to a single arbitrary function $\Phi$ (known in advance to everybody) is to think of a coin as a *record* that stores some arbitrary data *payload*, rather than just some integer value. The commitment of a record would then be published on the ledger when the record is created, and its unique serial number would be published when the record is consumed. A transaction would then contain serial numbers of consumed records, commitments of created records, and a proof attesting that invoking the function $\Phi$ on (the payload of) the old records produces (the payload of) the new records.

Data privacy holds because the ledger merely stores each record's commitment (and its serial number once consumed), and transactions only reveal that some number of old records were consumed in order to create some number of new records in a way that is consistent with $\Phi$. Function privacy also holds but for trivial reasons: $\Phi$ is known in advance to everybody, and every transaction is about computations of $\Phi$.

Note that Zerocash is indeed a special case of the above: it corresponds to fixing $\Phi$ to the particular (and publicly known) choice of a function $\Phi_{\$}$ that governs value transfers within a single currency. However the foregoing protocol supports only a single hard-coded function $\Phi$, while instead we want to enable users to select their own functions, as we discuss next.

## 2.2 Difficulties with achieving privacy for user-defined functions

We want to enable users to execute functions of their choice concurrently on the same ledger without seeking permission from anyone. That is, when preparing a transaction, a user should be able to pick *any* function $\Phi$ of their choice for creating new records by consuming some old records. If function privacy is not a concern, then this is easy: just attach to the transaction a zero-knowledge proof that $\Phi$ was correctly evaluated offline. However, because this approach reveals $\Phi$, we cannot use it because function privacy is a goal for us.

An approach that *does* achieve function privacy would be to modify the sketch in Section 2.1 by fixing a single function that is *universal*, and then interpreting data payloads as user-defined functions that are provided as inputs. Indeed, zero knowledge would ensure function privacy in this case. However merely allowing users to define their own functions does *not* by itself yield meaningful functionality, as we explain next.

**The problem: malicious functions.** A key challenge in this setting is that malicious users could devise functions to attack or disrupt other users' functions and data, so that a particular user would not know whether to trust records created by other users; indeed, due to function privacy, a verifier would not know what functions were used to create those records. For a concrete example, suppose that we wanted to realize the special case of value transfers within a single currency (i.e., Zerocash). One may believe that it would suffice to instruct users to pick the function $\Phi_{\$}$ (or similar). But this does *not* work: a user receiving a record claiming to contain, say, 1 unit of currency does not know if this record was created via the function $\Phi_{\$}$ from other such records and so on. A malicious user could have used a different function to create that record, for example, one that illegally "mints" records that appear valid to $\Phi_{\$}$, and thus enables arbitrary inflation of the currency. More generally, the lack of any enforced rules about how user-defined functions can interact precludes productive cooperation between users that are mutually distrustful. We stress that this challenge arises specifically due to the requirement that functions be private: if the function that created (the commitment of) a record was public knowledge, users could decide for themselves if records they receive were generated by "good" functions.

One way to address the foregoing problem is to augment records with a new attribute that identifies the function that "created" the record, and then impose the restriction that in a valid transaction only records created by the same function may participate. This new attribute is contained within a hiding commitment and thus is never revealed publicly on the ledger (just like a record's payload); the zero knowledge proof is tasked with ensuring that records participating in the same transaction are all of the same "type". This approach now *does* suffice to realize value transfers within a single currency, by letting users select the function $\Phi_{\$}$. More generally, this approach generalizes that in Section 2.1, and can be viewed as running multiple segregated "virtual ledgers" each with a fixed function. Function privacy holds because one cannot tell if a transaction belongs to one virtual ledger or another.

**The problem: functions cannot communicate.** The limitation of the above technique is that it forbids any "inter-process communication" between different functions, and so one cannot realize even simple functionalities like transferring value between different currencies on the same ledger. It also rules out more complex smart contract systems, as communication between contracts is a key part of such systems. It is thus clear that this crude "time sharing" of the ledger is too limiting.

## 2.3 The records nano-kernel: a minimalist shared execution environment

The approaches in Section 2.2 lie at opposite extremes: unrestricted inter-process interaction prevents the secure construction of even basic applications such as a single currency, while complete process segregation limits the ability to construct complex applications that interact with with each other.

Balancing these extremes requires a shared execution environment: one can think of this as an *operating system* for a shared ledger. This operating system manages user-defined functions: it provides process isolation, determines data ownership, handles inter-process communication, and so on. Overall, processes must be able to concurrently share a ledger, without violating the integrity or confidentiality of one another.

However, function privacy (one of our goals) dictates that user-defined functions are hidden, which means that an operating system cannot be maintained publicly atop the ledger (as in current smart contract systems) but, instead, must be part of the statement proved in zero knowledge. This is unfortunate because designing an operating system that governs interactions across user-defined functions within a zero knowledge proof is not only a colossal design challenge but also entails many arbitrary design choices that we should not have to take.

In light of the above, we choose to take the following approach: we formulate a *minimalist* shared execution environment that imposes simple, yet expressive, rules on how records may interact, and enables programming applications in the UTXO model (see Remark 2.3 for why we make this choice). This execution environment can be viewed as a "nano-kernel" that enables users to manage records containing data by programming two boolean functions (or predicates) associated with each record. These predicates control the two defining moments in a record's life, namely creation (or "birth") and consumption (or "death"), and are hence called the record's *birth* and *death* predicates. A user can create and consume records in a transaction by satisfying the predicates of those records. In more detail,

> **The records nano-kernel (RNK)** is an execution environment that operates over units of data called records. A record contains a *data payload*, a *birth predicate* $\Phi_{\mathsf{b}}$, and a *death predicate* $\Phi_{\mathsf{d}}$. Records are created and consumed by *valid transactions*. These are transactions where the death predicates of all consumed records and the birth predicates of all created records are simultaneously satisfied when given as input the transaction's *local data* (see Fig. 3), which includes: (a) every record's contents (such as its payload and the identity of its predicates); (b) a piece of shared memory that is publicly revealed, called *transaction memorandum*; (c) a piece of shared memory that is kept hidden, called *auxiliary input*; and (d) other construction specifics.

The foregoing definition enables predicates to see the contents of the entire transaction and hence *to individually decide if the local data is valid according to its own logic*. This in turn enables predicates to communicate with each other in a secure manner without interference from malicious predicates. In more detail, a record $\mathbf{r}$ can protect itself from other records that contain "bad" birth or death predicates because the $\mathbf{r}$'s predicates could refuse to accept when they detect (from reading the local data) that they are in a transaction with records having bad predicates. At the same time, a record can interact with other records in the same transaction when its predicates decide to accept, providing the flexibility that we seek.

We briefly illustrate this via an example, *user-defined assets*, whereby one can use birth predicates to define and transact with their own assets, and also use death predicates to enforce custom access control policies over these assets.

**Example 2.1** (user-defined assets)**.** Consider records whose payloads encode an asset identifier $\mathsf{id}$, the initial asset supply $\mathbb{v}$, and a value $v$. Fix the birth predicate in all such records to be a *mint-or-conserve* function $\mathsf{MoC}$ that is responsible for creating the initial supply of a new asset, and then subsequently conserving the value of the asset across all transactions. In more detail, $\mathsf{MoC}$ can be invoked in one of two modes. In *mint mode*, given as input a desired initial supply $\mathbb{v}$, $\mathsf{MoC}$ deterministically derives (in a way that we discuss later) a fresh unique identifier $\mathsf{id}$ for a new asset and stores $(\mathsf{id}, \mathbb{v}, v = \mathbb{v})$ in a *genesis record*. In *conserve mode*, $\mathsf{MoC}$ inspects all records in a transaction whose birth predicates equal to $\mathsf{MoC}$ and whose asset identifiers equal the identifier of the current record, and ensures that among these records, the asset values are conserved.

Users can program death predicates of records to enforce conditions on how assets can be consumed, e.g., by realizing *conditional exchanges* with other counter-parties. Suppose that Alice wishes to exchange

100 units of an asset $\mathsf{id}_1$ for 50 units of another asset $\mathsf{id}_2$, but does not have a counter-party for the exchange. She creates a record $\mathbf{r}$ with 100 units of $\mathsf{id}_1$ whose death predicate enforces that any transaction consuming $\mathbf{r}$ must also create another record, consumable by Alice, with 50 units of $\mathsf{id}_2$. She then publishes out of band information about $\mathbf{r}$, and anyone can subsequently claim it by creating a transaction doing the exchange.

Since death predicates can be *arbitrary*, many different access policies can also be realized, e.g., to enforce that a transaction redeeming a record (a) must be authorized by two of three public keys, or (b) becomes valid only after a given amount of time, or (c) must reveal the pre-image of a hash function.

One can generalize this basic example to show how the RNK can realize a specific class of smart contract systems, namely those in which the transaction creator knows both the contract code being executed, as well as the (public and secret) state of the contract. At a high level, these contracts can be executed within a single transaction, or across multiple transactions, by storing suitable intermediate state/message data in record payloads, or by publishing that data in transaction memoranda (as plaintext or ciphertext as needed). We discuss in more detail below.

**Example 2.2** (smart contracts with caller-known state)**.** At the highest level, smart contract systems operate over a set of individual contracts, each of which consists of a function (or collection of functions), some state variables, and some form of *address* that serves to uniquely identify the contract. The contract address ensures that the same code/functions can be deployed multiple times by different individuals, without two contracts inadvertently sharing state.[4] A standard feature of smart contract systems is that a contract can communicate with other contracts: that is, a contract can invoke a second smart contract as a subroutine, provided that the second contract provides an interface to allow this behavior. In our setting, we consider contracts in which the caller knows at least part of the state of each contract.

In this setting, one can use the records nano-kernel to realize basic smart contracts as follows. Each contract can be implemented as a function $\Phi_{\mathsf{sc}}$. The contract's state variables can be stored in one or more records such that each record $\mathbf{r}_i$ is labeled with $\Phi_{\mathsf{sc}}$ as the birth and death predicate. Using this labeling, $\Phi_{\mathsf{sc}}$ (via the RNK) can enforce that only it can update its state variables, thus fulfilling one requirement of a secure contract. Of course, while this serves to prevent *other functions* from updating the contract's state, it does not address the situation where multiple users wish to deploy different instances of the same function $\Phi_{\mathsf{sc}}$, each with isolated state. Fortunately (and validating our argument that the RNK realizes the *minimal* requirements needed for such a system), addressing this problem does not require changes to the RNK. Instead, one can devise the function $\Phi_{\mathsf{sc}}$ so that it reasons over a unique contract address identifier $\mathsf{id}$, which is recorded within the *payload* of every record.[5] The function $\Phi_{\mathsf{sc}}$ can achieve contract state isolation by enforcing that each input and output state record considered by single execution of $\Phi_{\mathsf{sc}}$ shares the same contract address.

To realize "inter-contract calls" between two functions $\Phi_{\mathsf{sc}_1}$ and $\Phi_{\mathsf{sc}_2}$, one can use "ephemeral" records that communicate between the two functions. For example, if $\Phi_{\mathsf{sc}_1}$ wishes to call $\Phi_{\mathsf{sc}_2}$, the caller may construct a record $\mathbf{r}_{\mathsf{e}}$ that contains the "arguments" to the called function $\Phi_{\mathsf{sc}_2}$, as well as the result of the function call. A transaction would then show that both $\Phi_{\mathsf{sc}_1}$ and $\Phi_{\mathsf{sc}_2}$ are satisfied.

The above example outlines how to implement a general smart contract system atop the RNK. We leave to future work the task of developing this outline into a full-fledged smart contract framework, and instead focus on constructing a scheme that implements the RNK, and on illustrating how to directly program the RNK

[4] In concrete implementations such as Ethereum [Woo17], contract identification is accomplished through unique contract addresses, each of which can be bound to a possibly non-unique `codeHash` that identifies the code of the program implementing the contract.

[5] This identifier can be generated in a manner similar to the asset identifier in Example 2.1.

to construct specific applications such as **user-defined assets**, **private decentralized asset exchanges**, and **regulation-friendly private stablecoins**. We discuss these applications in more detail in Section 6.

**Remark 2.3** (working in the UTXO model)**.** In the records nano-kernel, applications update their state by consuming records containing the old state, and producing new records that contain the updated state. This programming model is popularly known as the "unspent transaction output" (UTXO) model. This is in contrast to the "account-based" model which is used by many other smart contract systems [Goo14, Woo17, EOS18]. At present, it is not known how to efficiently achieve strong privacy properties in this model even for the simple case of privacy-preserving payments among any number of users, as we explain below.

In the account-based model, application state is stored in a persistent location associated with the application's account, and updates to this state are applied in-place. A smart contract that implements a currency in this model would store user balances in a persistent table $T$ that maps user account identifiers to user balances. Transactions from a user $A$ to another user $B$ would then decrement $A$'s balance in $T$ and increment $B$'s balance by a corresponding amount. A straightforward way to make this contract data-private (i.e., to hide the transaction value and the identities of $A$ and $B$) would be to replace the user balances in $T$ with *hiding commitments* to these balances; transactions would then update these commitments instead of directly updating the balances. However, while this hides transaction values, it does not hide user identities; to further hide these, every transaction would have to update *all* commitments in $T$, which entails a cost that grows linearly with the number of users. This approach is taken by zkLedger [NVV18], which enables private payments between a small number of known users (among other things).

Even worse, achieving function privacy when running multiple applications in such a system would require each transaction to hide which application's data was being updated, which means that the transaction would have to update the data of *all* applications at once, again severely harming the efficiency of the system.

In sum, it is unclear how to efficiently achieve strong data and function privacy in the account-based model when users can freely join and leave the system without notifying other users. On the other hand, we show in this paper that these properties can be achieved in the UTXO model at a modest cost.

## 2.4 Decentralized private computation

**A new cryptographic primitive.** We introduce a new cryptographic primitive called *decentralized private computation* (DPC) schemes, which capture the notion of a ledger-based system where privacy-preserving transactions attest to offline computations that follow the records nano-kernel. See Section 3 for the definition of DPC schemes, including the ideal functionality that we use to express security.

We construct a DPC scheme in Section 4, and prove it secure in Appendix A. We take Zerocash [BCG$^+$14] as a starting point, and then extend the protocol to support the records nano-kernel and also to facilitate proving security in the simulation paradigm relative to an ideal functionality (rather than via a collection of separate game-based definitions as in [BCG$^+$14]). Below we sketch the construction.

**Construction sketch.** Each transaction in the ledger consumes some old records and creates new records in a manner that is consistent with the records nano-kernel. To ensure privacy, a transaction only contains serial numbers of the consumed records, commitments of the created records, and a zero knowledge proof attesting that there exist records consistent with this information (and with the records nano-kernel). All commitments on the ledger are collected in a Merkle tree, which facilitates efficiently proving that a commitment appears on the ledger (by proving in zero knowledge the knowledge of a suitable authentication path). All serial numbers on the ledger are collected in a list that cannot contain duplicates. This implies that a record cannot be consumed twice because the same serial number is revealed each time a record is consumed. See Fig. 1.

The record data structure is summarized in Fig. 2. Each record is associated to an *address public key*, which is a commitment to a seed for a pseudorandom function acting as the corresponding *address secret key*; addresses determine ownership of records, and in particular consuming a record requires knowing its secret key. A *record* consists of an address public key, a data payload, a birth predicate, a death predicate, and a serial number nonce; a *record commitment* is a commitment to all of these attributes. The *serial number* of a record is the evaluation of a pseudorandom function, whose seed is the secret key for the record's address public key, evaluated at the record's serial number nonce. A record's commitment and serial number, which appear on the ledger when the record is created and consumed, reveal *no* information about the record attributes. This follows from the hiding properties of the commitment, and the pseudorandom properties of the serial number. The derivation of a record's serial number ensures that a user can create a record for another in such a way that its serial number is fully determined and yet cannot be predicted without knowing the other user's secret key.

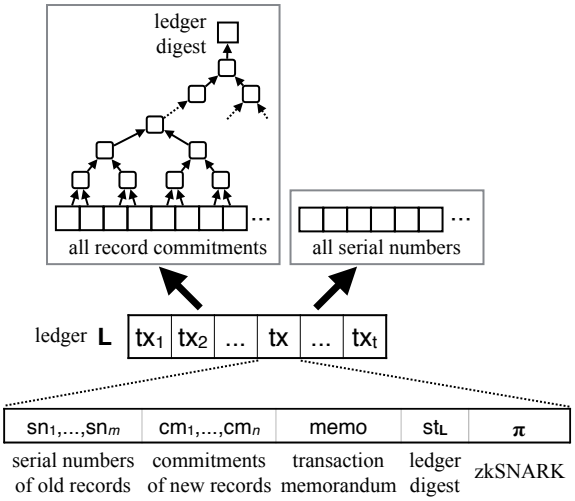


**Figure 1:** Construction of a transaction.

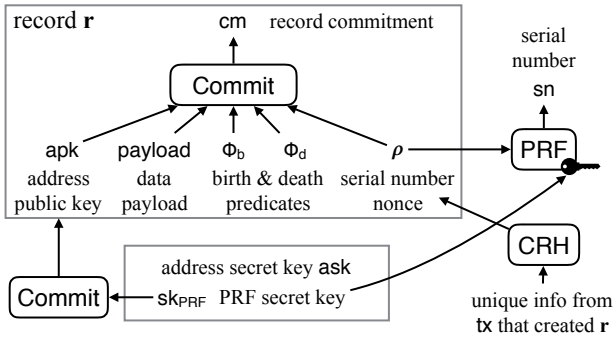


**Figure 2:** Construction of a record.

In order to produce a transaction, a user selects some previously-created records to consume, assembles some new records to create (including their payloads and predicates), and decides on other aspects of the local data such as the transaction memorandum (shared memory seen by all predicates and published on the ledger) and the auxiliary input (shared memory seen by all predicates but not published on the ledger); see Fig. 3. If the user knows the secret keys of the records to consume and if all relevant predicates are satisfied (death predicates of old records and birth predicates of new predicates), then the user can produce a zero knowledge proof to append to the transaction. See Fig. 4 for a summary of the NP statement being proved.

In sum, a transaction only reveals the number of consumed records and number of created records, as well as any data that was deliberately revealed in the transaction memorandum (possibly nothing).[6]

**Achieving succinctness.** Our discussions so far have focused on achieving (data and function) privacy. However, we also want to achieve succinctness, namely, that a transaction can be validated in "constant time". This follows from a straightforward modification: we take the protocol that we have designed so far and use a zero knowledge *succinct* argument rather than just any zero knowledge proof. Indeed, the NP statement being proved (summarized in Fig. 4) involves attesting the satisfiability of all (old) death and (new) birth predicates, and thus we need to ensure that verifying the corresponding proof can be done in time that does not depend

[6]By supporting the use of dummy records, we can in fact ensure that only *upper bounds* on the foregoing numbers are revealed.

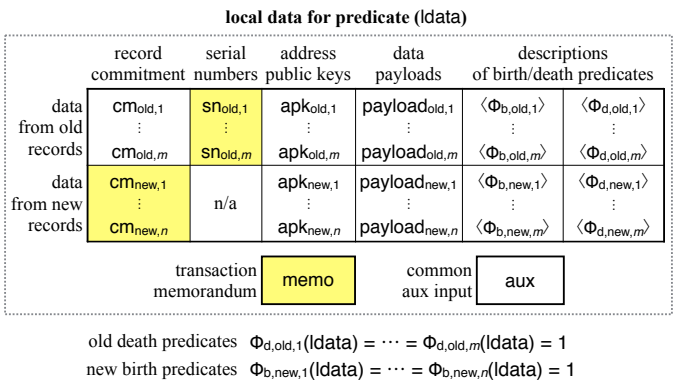


**Figure 3:** Predicates receive local data.

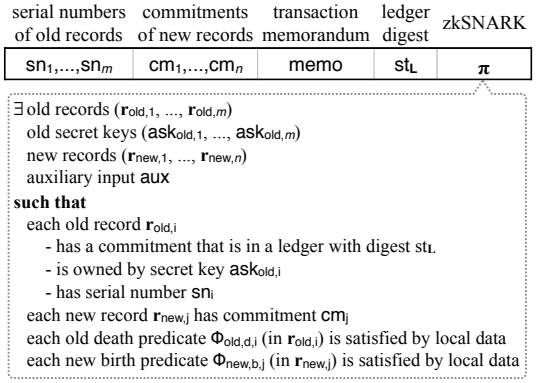


**Figure 4:** The execute statement.

on the complexity of these predicates. While turning this idea into an efficient implementation requires more ideas (as we discuss in Section 2.5), the foregoing modification suffices from a theoretical point of view.

**Delegation to an untrusted worker.** In our DPC scheme, a user must produce, and include in the transaction, a zero knowledge succinct argument that, among other things, attests that death predicates of consumed records are satisfied and, similarly, that birth predicates of created records are satisfied. This implies that the cost of creating a transaction grows with the complexity (and number of) predicates involved in the transaction. Such a cost can quickly become infeasible for weak devices such as mobile phones or hardware tokens.

We address this problem by enabling a user to *delegate* to an untrusted worker, such as a remote server, the computation that produces a transaction. This notion, which we call a *delegable DPC scheme*, empowers weak devices to produce transactions that they otherwise could not have produced on their own.

The basic idea is to augment address keys in such a way that the secret information needed to produce the cryptographic proof is separate from the secret information needed to authorize a transaction containing that proof. Thus, the user can communicate to the worker the secrets necessary to generate a cryptographic proof, while retaining the remaining secrets for authorizing this (and future) transactions. In particular, the worker has no way to produce valid transactions that have not been authorized by the user.

We use randomizable signatures to achieve the foregoing functionality, without violating either privacy or succinctness. Informally, we modify a record's serial number to be an unlinkable randomization of (part of) the record's address public key, and a user's authorization of a transaction consists of signing the instance and proof relative to every randomized key (i.e., serial number) in that transaction. See Section 5 for details.

## 2.5 Achieving an efficient implementation

Our system Zexe (*Zero knowledge EXEcution*) provides an implementation of two constructions: our "plain" DPC protocol, and its extension to a delegable DPC protocol. Achieving efficiency in our system required overcoming several challenges. Below we highlight some of these challenges, and explain how we addressed them; see Sections 7 and 8 for details. The discussions below equally apply to both types of DPC protocols.

**Avoiding the cost of universality.** The NP statement that we need to prove involves checking user-defined predicates, so it must support arbitrary computations that are not fixed in advance. However, state-of-the-art zkSNARKs for universal computations rely on expensive tools [BCG+13, BCTV14, WSR+15, BCTV17].

We address this problem by relying on one layer of *recursive proof composition* [Val08, BCCT13]. Instead of tasking the NP statement with directly checking user-defined predicates, we only task it with checking *succinct proofs* attesting to this. Checking these *inner* succinct proofs is a (relatively) inexpensive computation

that is fixed for *all* predicates, which can be "hardcoded" in the statement. Since the single *outer* succinct proof produced does not reveal information about the inner succinct proofs attesting to predicates' satisfiability (thanks to zero knowledge), the inner succinct proofs do *not* have to hide what predicate was checked, so they can be for NP statements *tailored* to the computations of particular user-defined predicates.

**A bespoke recursion.** Recursive proof composition has been empirically demonstrated for pairing-based SNARKs [BCTV17]. We thus focus our attention on these, and explain the challenges that arise in our setting. Recall that if we instantiate a SNARK's pairing via an elliptic curve $E$ defined over a prime field $\mathbb{F}_q$ and having a subgroup of prime order $r$, then (a) the SNARK supports NP statements expressed as arithmetic circuits over $\mathbb{F}_r$, while (b) proof verification involves arithmetic operations over $\mathbb{F}_q$. Being part of the NP statement, the SNARK verifier must also be expressed as an arithmetic circuit over $\mathbb{F}_r$, which is problematic because the verifier's "native" operations are over $\mathbb{F}_q$. Simulating $\mathbb{F}_q$ operations via $\mathbb{F}_r$ operations is expensive, and picking $E$ such that $q = r$ is impossible [BCTV17]. Prior work thus uses *multiple* curves [BCTV17]: a two-cycle of pairing-friendly elliptic curves, that is, two prime-order curves $E_1$ and $E_2$ such that the prime size of one's base field is the prime order of the other's group, and orchestrating SNARKs based on these so that fields "match up". However, known cycles are inefficient at 128 bits of security [BCTV17, CCW19].

We address this problem by noting that we merely need "a proof of a proof", and thus, instead of relying on a cycle, we can use the Cocks–Pinch method [FST10] to set up a bounded recursion [BCTV17]. First we pick a pairing-friendly elliptic curve that not only is suitable for 128 bits of security according to standard considerations but, moreover, is compatible with efficient SNARK provers in *both* levels of the recursion. Namely, letting $p$ be the prime order of the base field and $r$ the prime order of the group, we need that *both* $\mathbb{F}_r$ and $\mathbb{F}_p$ have multiplicative subgroups whose orders are large powers of 2. The condition on $\mathbb{F}_r$ ensures efficient proving for SNARKs over this curve, while the condition on $\mathbb{F}_p$ ensures efficient proving for SNARKs that verify proofs over this curve. In light of the above, we select a curve $E_{\mathsf{BLS}}$ from the Barreto–Lynn–Scott (BLS) family [BLS02, CLN11] with embedding degree 12. This family not only enables parameters that conservatively achieve 128 bits of security, but also enjoys properties that facilitate very efficient implementation [AFK$^+$12]. We ensure that both $\mathbb{F}_r$ and $\mathbb{F}_p$ have multiplicative subgroups of order $2^\alpha$ for $\alpha \geq 40$, by a suitable condition on the parameter of the BLS family.

Next we use the Cocks–Pinch method to pick a pairing-friendly elliptic curve $E_{\mathsf{CP}}$ over a field $\mathbb{F}_q$ such that the curve group $E_{\mathsf{CP}}(\mathbb{F}_q)$ contains a subgroup of prime order $p$ (the size of $E_{\mathsf{BLS}}$'s base field). Since the method outputs a prime $q$ that has about $2\times$ more bits than the desired $p$, and in turn $p$ has about $1.5\times$ more bits than $r$ (due to properties of the BLS family), we only need $E_{\mathsf{CP}}$ to have embedding degree of 6 in order to achieve 128 bits of security (as determined from the guidelines in [FST10]). We note that this approach is similar to that in Geppetto [CFH$^+$15], but adapted for BLS12 curves.

In sum, a SNARK over $E_{\mathsf{BLS}}$ is used to generate proofs of predicates' satisfiability; after that a zkSNARK over $E_{\mathsf{CP}}$ is used to generate proofs that these prior proofs are valid along with the remaining NP statement's checks. The matching fields between the two curves ensure that the former proofs can be efficiently verified.

**Minimizing operations over $E_{\mathsf{CP}}$.** While the curve $E_{\mathsf{CP}}$ facilitates efficient checking of SNARK proofs over $E_{\mathsf{BLS}}$, operations on it are at least $2\times$ more costly (in time and space) than operations over $E_{\mathsf{BLS}}$, simply because $E_{\mathsf{CP}}$'s base field is twice the size of $E_{\mathsf{BLS}}$'s base field. This makes checks in the NP relation $\mathcal{R}_{\mathsf{e}}$ that are not related to proof checking unnecessarily expensive.

To avoid this, we split $\mathcal{R}_{\mathsf{e}}$ into two NP relations, $\mathcal{R}_{\mathsf{BLS}}$ and $\mathcal{R}_{\mathsf{CP}}$. The latter is responsible only for verifying proofs of predicates' satisfaction, while the former is responsible for all other checks. We minimize the number of $E_{\mathsf{CP}}$ operations by proving satisfaction of $\mathcal{R}_{\mathsf{BLS}}$ and $\mathcal{R}_{\mathsf{CP}}$ with zkSNARKs over $E_{\mathsf{BLS}}$ and $E_{\mathsf{CP}}$ respectively. A transaction now includes both proofs.

**Optimizing the NP statement.** We note that the remaining NP statement's checks can themselves be

quite expensive, as they range from verifying authentication paths in a Merkle tree to verifying commitment openings, and from evaluating pseudorandom functions to evaluating collision resistant functions. Prior work realizing similar collections of checks required upwards of *four million gates* [BCG+14] to express such checks. This not only resulted in high latencies for producing transactions (several minutes) but also resulted in large public parameters for the system (hundreds of megabytes).

Commitments and collision-resistant hashing can be expressed as very efficient arithmetic circuits if one opts for Pedersen-type constructions over suitable Edwards elliptic curves (and techniques derived from these ideas are now part of deployed systems [HBHW20]). To achieve this, we pick two Edwards curves, $E_{\mathsf{Ed/BLS}}$ over the field $\mathbb{F}_r$ (thereby matching the group order of $E_{\mathsf{BLS}}$), and $E_{\mathsf{Ed/CP}}$ over the field $\mathbb{F}_p$ (thereby matching the group order of $E_{\mathsf{CP}}$). This allows to realise very efficient circuits for various primitives used in our NP relations, including commitments, collision-resistant hashing, and randomizable signatures. Overall, we obtain highly optimized realizations of all checks in Fig. 4.

### 2.6 Deployment considerations

DPC schemes include a setup algorithm that specifies how to sample public parameters, which are used to produce transactions and to verify transactions. The setup algorithm in our DPC construction (see Section 4) simply consists of running the setup algorithms for the various cryptographic building blocks that we rely on: commitment schemes, collision-resistant hash functions, and zero knowledge proofs.

In practice, deploying cryptography that relies on setup algorithms (such as DPC schemes) can be challenging because the entity running the setup algorithm may be able to break certain security properties of the scheme, by abusing knowledge of the randomness used to produce the public parameters. On the other hand, *some* setup algorithm is typically inevitable. For example, non-interactive zero knowledge proofs without any setup exist only for languages decidable in polynomial time [GO94]. Nevertheless, one could still aim for a *transparent setup*, one that consists of public randomness, because in practice it is cheaper to realize.

Our construction of a DPC scheme has a transparent setup algorithm whenever the setup algorithms for the underlying cryptographic building blocks also have transparent setups. For example, this would hold if we instantiated our construction via Pedersen commitments, Pedersen hash functions, and transparent zkSNARKs (as obtained from probabilistic checking tools in the random oracle model [Mic00, BCS16]).

However, due to efficiency considerations described in Section 2.5, our implemented system relies on pairing-based zkSNARKs whose setup is *not* transparent. (We use the simulation-extractable zkSNARK of Groth and Maller [GM17].) We should thus discuss how one may deploy our implemented system, and in particular the effects of compromise in the trusted setup phase of these SNARKs. (All other primitives in our system use a transparent setup.)

Recall that prior zkSNARK deployments have used secure multiparty computation [BCG+15, ZCa16, BGM17, BGG18], so that the sampled public parameters are guaranteed to be secure as long as even a single participating party is honest. One could leverage these same ideas to sample "master" parameters for proving/verifying the two NP relations $\mathcal{R}_{\mathsf{BLS}}$ and $\mathcal{R}_{\mathsf{CP}}$ (over the two elliptic curves $E_{\mathsf{BLS}}$ and $E_{\mathsf{CP}}$) mentioned in Section 2.5. Note that these public parameters do *not* depend on any user-defined functions (birth or death predicates), and can thus be sampled once and for all regardless of which applications will be run over the system. Note also that these public parameters must be trusted by *everyone*, because if they were compromised then the security (but not privacy) of *all* applications running over the system would be compromised as well.

The foregoing public parameters are not the only ones that need to be sampled in order to use our implemented system. Every (birth or death) predicate requires its own public parameters, because (the verification key contained in) these public parameters is part of the record that contains it, and is ultimately

used to recursively check a proof of the predicate's satisfiability. Since an application relies only on the public parameters of certain predicates, we call such parameters as "application" parameters.

Unlike "master" parameters, "application" parameters do not have to be sampled at the start of the system's lifetime, and also do not have to be trusted by every user in the system. Indeed, interactions across records are overseen by the NP relations $\mathcal{R}_{\mathsf{BLS}}$ and $\mathcal{R}_{\mathsf{CP}}$ (which rely on the "master" parameters) and thus compromised parameters for one application will not affect (the security and privacy of) an application that does not rely on them. This means that a user only needs to trust the parameters that are relied upon by the applications that the user cares about. In turn this means that the sampling of application parameters can be viewed as an organic process, which occurs as applications are developed and deployed, and each application can be in charge of deciding whichever method is most suitable for securely sampling its own parameters.

Very recent works [MBKM19, CFQ19, CHM$^+$20, GWC19] have proposed pairing-based SNARKs that have a universal setup that can be used for *any* circuit. Once such SNARK constructions mature into efficient implementations, our system can be easily modified to use these instead of [GM17] to mitigate the above concerns, as both our construction and implementation make use of the underlying SNARKs in a modular and black-box manner.

# 3 Definition of decentralized private computation schemes

We define *decentralized private computation* (DPC) schemes, a cryptographic primitive in which parties with access to an ideal append-only ledger execute computations offline and subsequently post privacy-preserving, publicly-verifiable transactions that attest to the correctness of these offline executions. This primitive generalizes prior notions [BCG+14] that were limited to proving correctness of simple financial invariants.

Below we introduce the data structures, interface, and security requirements for a DPC scheme: Section 3.1 describes the main data structures of a DPC scheme, Section 3.2 defines the syntax of the DPC algorithms, and finally in Section 3.3 we describe the security requirements for DPC schemes via an ideal functionality. We note that our definition of DPC schemes focuses on (correctness and) privacy, because we leave succinctness as a separate efficiency goal that easily follows from suitable building blocks (see Remark 4.1).

## 3.1 Data structures

In a DPC scheme there are three main data structures: *records*, *transactions*, and the *ledger*.

**Records.** A *record*, denoted by the symbol $\mathbf{r}$, is a data structure representing a unit of data. Records can be created or consumed, and these events denote state changes in the system. For example, in a currency application, records store units of the currency, and state changes represent the flow of units in that currency.

In more detail, a record $\mathbf{r}$ has the following attributes (see Fig. 5): (a) a *commitment* $\mathsf{cm}$, which binds together all other attributes of $\mathbf{r}$ while hiding all information about them; (b) an *address public key* $\mathsf{apk}$, which specifies the record's owner; (c) a *payload* $\mathsf{payload}$ containing arbitrary application-dependent information; (d) a *birth predicate* $\Phi_{\mathsf{b}}$ that must be satisfied when $\mathbf{r}$ is created; (e) a *death predicate* $\Phi_{\mathsf{d}}$ that must be satisfied when $\mathbf{r}$ is consumed; and (f) other construction-specific information. Both $\Phi_{\mathsf{b}}$ and $\Phi_{\mathsf{d}}$ are arbitrary non-deterministic boolean-valued functions. The payload $\mathsf{payload}$ contains a designated subfield $\mathsf{isDummy}$ which denotes whether $\mathbf{r}$ is dummy or not.

Informally, the "life" of a (non-dummy) record $\mathbf{r}$ is marked by two events: *birth* and *death*. The record $\mathbf{r}$ is *born* (or is *created*) when its commitment $\mathsf{cm}$ is posted to the ledger as part of a transaction. Then the record $\mathbf{r}$ *dies* (or is *consumed*) when its serial number $\mathsf{sn}$ appears on the ledger as part of a later transaction. At each of these times (birth or death) the corresponding predicate ($\Phi_{\mathsf{b}}$ or $\Phi_{\mathsf{d}}$) must be satisfied. Dummy records, on the other hand, can be created freely, but consuming them requires satisfaction of their death predicates. The purpose of dummy records is solely to enable the creation of new non-dummy records.

To consume $\mathbf{r}$, one must also know the address secret key $\mathsf{ask}$ corresponding to $\mathbf{r}$'s address public key $\mathsf{apk}$ because the serial number $\mathsf{sn}$ to be revealed can only be computed from $\mathbf{r}$ and $\mathsf{ask}$. The ledger forbids the same serial number to appear more than once, so that: (a) a record cannot be consumed twice because it is associated to exactly one serial number; (b) others cannot prevent one from consuming a record because it is computationally infeasible to create two distinct records that share the same serial number $\mathsf{sn}$ but have distinct commitments $\mathsf{cm}$ and $\mathsf{cm}'$.

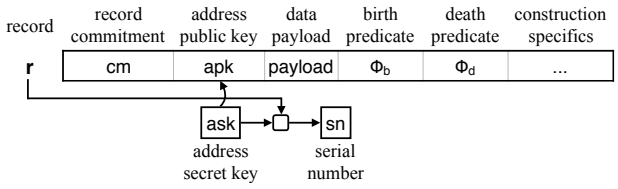


**Figure 5:** Diagram of a record.

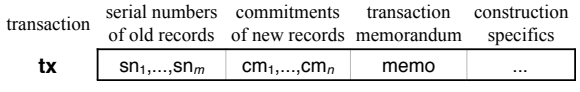


**Figure 6:** Diagram of a transaction.

**Transactions.** A transaction, denoted by the symbol $\mathsf{tx}$, is a data structure representing a state change that involves the consumption and creation of records (see Fig. 6). It is a tuple $([\mathsf{sn}_i]_1^m, [\mathsf{cm}_j]_1^n, \mathsf{memo}, \star)$ where (a) $[\mathsf{sn}_i]_1^m$ is the list of serial numbers of the $m$ old records, (b) $[\mathsf{cm}_j]_1^n$ is the list of commitments of the $n$ new records, (c) $\mathsf{memo}$ is an arbitrary string associated with the transaction, and (d) $\star$ is other construction-specific information. The transaction $\mathsf{tx}$ reveals only the following information about old and new records: (i) the old records' serial numbers; (ii) the new records' commitments; and (iii) the fact that the death predicates of all consumed records and birth predicates of all new records were satisfied.

Anyone can assemble a transaction and append it to the ledger, provided that it is "valid" in the sense that (all records are well-formed and) the death predicates of any consumed records and the birth predicates of any created records are satisfied. Note that all transactions reveal the number of old records ($m$) and the number of new records ($n$), but not how many of these were dummy or not.

**Ledger.** We consider a model where all parties have access to an append-only ledger, denoted $\mathbf{L}$, that stores all published transactions. Our definitions (and constructions) are *agnostic* to how this ledger is realized (e.g., the ledger may be centrally managed or a distributed protocol). When an algorithm needs to interact with the ledger, we specify $\mathbf{L}$ in the algorithm's superscript. The ledger exposes the following interface.

- $\mathbf{L}.\mathsf{Len}$: Return the number of transactions currently on the ledger.
- $\mathbf{L}.\mathsf{Push}(\mathsf{tx})$: Append a (valid) transaction $\mathsf{tx}$ to the ledger.
- $\mathbf{L}.\mathsf{Digest} \to \mathsf{st}_{\mathbf{L}}$: Return a (short) digest of the current state of the ledger.
- $\mathbf{L}.\mathsf{ValidateDigest}(\mathsf{st}_{\mathbf{L}}) \to b$: Check that $\mathsf{st}_{\mathbf{L}}$ is a valid digest for some (past) ledger state.
- $\mathbf{L}.\mathsf{Contains}(\mathsf{tx}) \to b$: Determine if $\mathsf{tx}$ (or a subcomponent thereof) appears on the ledger or not.
- $\mathbf{L}.\mathsf{Prove}(\mathsf{tx}) \to \mathbb{w}_{\mathbf{L}}$: If a transaction $\mathsf{tx}$ (or a subcomponent thereof) appears on the ledger, return a proof of membership $\mathbb{w}_{\mathbf{L}}$ for it. If there are duplicates, return a proof for the lexicographically first one.
- $\mathbf{L}.\mathsf{Verify}(\mathsf{st}_{\mathbf{L}}, \mathsf{tx}, \mathbb{w}_{\mathbf{L}}) \to b$: Check that $\mathbb{w}_{\mathbf{L}}$ certifies that $\mathsf{tx}$ (or a subcomponent thereof) is in a ledger with digest $\mathsf{st}_{\mathbf{L}}$.

We stress that only "valid" transactions can be appended to the ledger. While the full definition of a valid transaction is implementation dependent, in all cases it must be that the commitments and serial numbers in a transaction (including any appearing in the $\star$ field of a transaction) do not already appear on the ledger.

## 3.2 Algorithms

A DPC scheme is a tuple of algorithms (some of which may read information from $\mathbf{L}$):

$$\mathsf{DPC} = (\mathsf{Setup}, \mathsf{GenAddress}, \mathsf{Execute}^{\mathbf{L}}, \mathsf{Verify}^{\mathbf{L}}) \ .$$

The syntax and semantics of these algorithms are informally described below.

**Setup:** $\mathsf{DPC.Setup}(1^\lambda) \to \mathsf{pp}$.

On input a security parameter $1^\lambda$, $\mathsf{DPC.Setup}$ outputs public parameters $\mathsf{pp}$ for the system. A trusted party runs this algorithm once and then publishes its output; afterwards the trusted party is not needed anymore.

For some constructions, the trusted party can be replaced by an efficient multiparty computation that securely realizes the $\mathsf{DPC.Setup}$ algorithm (see [BCG+15, ZCa16, BGM17, BGG18] for how this has been done in some systems); in other constructions, the trusted party may not be needed, as the public parameters may simply consist of a random string of a certain length.

**Create address:** $\mathsf{DPC.GenAddress}(\mathsf{pp}) \to (\mathsf{apk}, \mathsf{ask})$.

On input public parameters pp, DPC.GenAddress outputs an address key pair $(\mathsf{apk}, \mathsf{ask})$. Any user may run this algorithm to create an address key pair. Each record is bound to an address public key, and the corresponding secret key is used to consume it.

**Execute:** Any user may invoke DPC.Execute to consume records and create new ones.

$$\mathsf{DPC.Execute}^{\mathbf{L}}\left(\begin{array}{ll} \text{public parameters} & \mathsf{pp} \\ \text{old records} & [\mathbf{r}_i]_1^m \\ \text{old address secret keys} & [\mathsf{ask}_i]_1^m \\ \text{new address public keys} & [\mathsf{apk}_j]_1^n \\ \text{new record payloads} & [\mathsf{payload}_j]_1^n \\ \text{new record birth predicates} & [\Phi_{\mathsf{b},j}]_1^n \\ \text{new record death predicates} & [\Phi_{\mathsf{d},j}]_1^n \\ \text{auxiliary predicate input} & \mathsf{aux} \\ \text{transaction memorandum} & \mathsf{memo} \end{array}\right) \to \left(\begin{array}{ll} \text{new records} & [\mathbf{r}_j]_1^n \\ \text{transaction} & \mathsf{tx} \end{array}\right).$$

Given as input a list of old records $[\mathbf{r}_i]_1^m$ with corresponding secret keys $[\mathsf{ask}_i]_1^m$, attributes for new records, private auxiliary input aux to birth and death predicates of new and old records respectively,[7] and an arbitrary transaction memorandum memo, DPC.Execute produces new records $[\mathbf{r}_j]_1^n$ and a transaction tx. The transaction attests that the input records' death predicates and the output records' birth predicates are all satisfied. The user subsequently pushes tx to the ledger by invoking $\mathbf{L}.\mathsf{Push}(\mathsf{tx})$.

**Verify:** $\mathsf{DPC.Verify}^{\mathbf{L}}(\mathsf{pp}, \mathsf{tx}) \to b$.

On input public parameters pp and a transaction tx, and given oracle access to the ledger $\mathbf{L}$, DPC.Verify outputs a bit $b$ denoting whether the transaction tx is valid relative to the ledger $\mathbf{L}$.

### 3.3 Security

Informally, a DPC scheme achieves the following security goals.

- *Execution correctness.* Malicious parties cannot create valid transactions if the death predicate of some consumed record or the birth predicate of some created record is not satisfied.
- *Execution privacy.* Transactions reveal only the information revealed in the memorandum field, a bound on the number of consumed records, and a bound on the number of created records.[8] All other information is hidden, including the payloads and predicates of all involved records. For example, putting aside the information revealed in the memorandum (which is arbitrary), one cannot link a transaction that consumes a record with the prior transaction that created it.
- *Consumability.* Every record can be consumed at least once and at most once by parties that know its secrets. Thus, a malicious party cannot create two valid records for another party such that only one of them can be consumed. (This captures security against "faerie-gold" attacks [HBHW20].)
- *Transaction non-malleability.* Malicious parties cannot modify a transaction "in flight" to the ledger.

Formally, we prove *standalone* security against *static corruptions*, in a model where every party has private anonymous channels to all other parties [IKOS06].[9] (In Appendix C we discuss how to prove security under

[7] In addition to the "global" auxiliary input aux, each predicate may also take as input a "local" auxiliary input that is not (necessarily) shared with other predicates. For simplicity, we make these local inputs implicit.

[8] And any information implied by knowing that the birth (resp., death) predicates of consumed (resp., created) records are satisfied.

[9] Parties can, e.g., use these channels to communicate the contents of newly created records to other parties.

composition and against adaptive corruptions.) In more detail, we capture security of a DPC scheme via a *simulation-based* security definition that is akin to UC security [Can01], but restricted to a single execution.

**Definition 3.1.** *A DPC scheme* $\mathsf{DPC}$ *is* **secure** *if for every efficient real-world adversary* $\mathcal{A}$ *there exists an efficient ideal-world simulator* $\mathcal{S}_{\mathcal{A}}$ *such that for every efficient environment* $\mathcal{E}$ *the following are computationally indistinguishable:*

- *the output of* $\mathcal{E}$ *when interacting with the adversary* $\mathcal{A}$ *in a real-world execution of* $\mathsf{DPC}$ *in a model where parties can communicate with other parties via private anonymous channels; and*
- *the output of* $\mathcal{E}$ *when interacting with the simulator* $\mathcal{S}_{\mathcal{A}}$ *in an ideal-world execution with the ideal functionality* $\mathcal{F}_{\mathrm{DPC}}$ *specified in Fig. 7 (and further described below).*

We describe the data structures used by the ideal functionality $\mathcal{F}_{\mathrm{DPC}}$, the internal state of $\mathcal{F}_{\mathrm{DPC}}$, and the interface offered by $\mathcal{F}_{\mathrm{DPC}}$ to parties in the ideal-world execution.

**Ideal data structures.** The ideal functionality $\mathcal{F}_{\mathrm{DPC}}$ uses ideal counterparts of a DPC scheme's data structures. An *address public key* $\mathsf{apk}$ denotes the owner of an *ideal record* $\mathtt{r}$, which is a tuple $(\mathsf{cm}, \mathsf{apk}, \mathsf{payload}, \Phi_{\mathsf{b}}, \Phi_{\mathsf{d}})$, where $\mathsf{cm}$ is its commitment, $\mathsf{apk}$ is its address public key, $\mathsf{payload}$ is its payload, and $\Phi_{\mathsf{b}}$ and $\Phi_{\mathsf{d}}$ are its birth and death predicates. The record is also associated with a unique identifier (or *serial number*) $\mathsf{sn}$. We require that $\mathsf{apk}$, $\mathsf{cm}$, and $\mathsf{sn}$ are "globally unique"; this means that there cannot be two different ideal records $\mathtt{r}$ and $\mathtt{r}'$ having the same commitments or serial numbers.

The distribution of these components is specified by the simulator $\mathcal{S}$ as follows. Before the ideal execution begins, $\mathcal{S}$ specifies three functions $(\mathsf{SampleAddrPk}, \mathsf{SampleCm}, \mathsf{SampleSn})$ that, on input a random string, sample $(\mathsf{apk}, \mathsf{cm}, \mathsf{sn})$ respectively. When $\mathcal{F}_{\mathrm{DPC}}$ needs to sample one of these, it invokes the respective functions. (Note that $\mathcal{F}_{\mathrm{DPC}}$ cannot directly ask $\mathcal{S}$ to sample these because that would reveal to $\mathcal{S}$ when an honest party was invoking $\mathcal{F}_{\mathrm{DPC}}.\mathsf{GenAddress}$ or $\mathcal{F}_{\mathrm{DPC}}.\mathsf{Execute}$, and we cannot afford this leakage.)

**Internal state.** The ideal functionality $\mathcal{F}_{\mathrm{DPC}}$ maintains several internal tables.

- $\mathsf{Addr}$, which stores address public keys.
- $\mathsf{AddrUsers}$, which maps an address public key to the set of parties that are authorized to use it.
- $\mathsf{Records}$, which maps a record's commitment to that record's information (address public key, payload, birth predicate, and death predicates).
- $\mathsf{RecUsers}$, which maps a record's commitment to the set of parties that are authorized to consume it. Note that, for a record $\mathtt{r}$, the set $\mathsf{RecUsers}[\mathtt{r}.\mathsf{cm}]$ can be different from the set in $\mathsf{AddrUsers}[\mathtt{r}.\mathsf{apk}]$, but a party $\mathcal{P}$ has to be in both sets to consume $\mathtt{r}$.
- $\mathsf{SerialNumbers}$, which maps a record's commitment to that record's (unique) serial number.
- $\mathsf{State}$, which maps a record's commitment to that record's state, either `alive` or `dead`.

**Ideal algorithms.** The ideal functionality $\mathcal{F}_{\mathrm{DPC}}$ provides the following interface to parties.

- *Address generation:* $\mathcal{F}_{\mathrm{DPC}}.\mathsf{GenAddress}$ outputs a new address public key $\mathsf{apk}$.
- *Execution:* $\mathcal{F}_{\mathrm{DPC}}.\mathsf{Execute}$ performs an execution that consumes old records and creates new records. All parties are notified that an execution has occurred, and learn the serial numbers of input records, commitments of output records, and the transaction memorandum $\mathsf{memo}$. Concurrent $\mathcal{F}_{\mathrm{DPC}}.\mathsf{Execute}$ calls are serialized arbitrarily.
- *Record consumption authorization:* $\mathcal{F}_{\mathrm{DPC}}.\mathsf{ShareRecord}$ allows a party $\mathcal{P}$ to authorize another party $\mathcal{P}'$ to consume a record $\mathtt{r}$ (provided that $\mathcal{P}'$ is also authorized to use $\mathtt{r}$'s address public key).

**Operation of honest parties.** In both the real and ideal executions, the environment $\mathcal{E}$ can send instructions to honest parties. These instructions can be one of $\mathsf{GenAddress}$, $\mathsf{Execute}$, or $\mathsf{ShareRecord}$. In the real world honest parties translate these instructions into corresponding invocations of $\mathsf{DPC}$ algorithms (or messages sent via private anonymous channels as in the case of $\mathsf{ShareRecord}$), while in the ideal world they translate

them into corresponding invocations of $\mathcal{F}_{\text{DPC}}$ algorithms. In both worlds, honest parties immediately invoke ShareRecord on records obtained from an Execute instruction. Finally, in the ideal world, when invoking $\mathcal{F}_{\text{DPC}}$, honest parties do not provide any inputs marked as optional; instead, they let $\mathcal{F}_{\text{DPC}}$ sample these.

**Intuition.** We explain how $\mathcal{F}_{\text{DPC}}$ enforces the informal security notions described at this section's beginning.

- *Execution correctness.* $\mathcal{F}_{\text{DPC}}$.Execute ensures that the death predicates of consumed records and birth predicates of created records are satisfied by the local data. Note that each predicate receives its own position as input so that it knows to which record in the local data it belongs.
- *Execution privacy.* Transactions contain serial numbers $[\mathsf{sn}_i]_1^m$ of consumed records, commitments $[\mathsf{cm}_j]_1^n$ of created records, and a memorandum memo. Serial numbers and commitments are sampled via SampleSn and SampleCm, so they are independent of the contents of any record, and thus reveal no information about them. Transactions thus reveal no information (beyond what is contained in memo).
- *Consumability.* From the point of view of $\mathcal{F}_{\text{DPC}}$, two records are different if and only if they have different commitments. In such a case, both records can be consumed as long as their death predicates are satisfied. If a DPC scheme realizes $\mathcal{F}_{\text{DPC}}$, then it must satisfy this same requirement: if two valid records have distinct commitments, then they must both be consumable.
- *Transaction non-malleability.* The adversary has no power to modify the inputs to, or output of, an honest party's invocation of $\mathcal{F}_{\text{DPC}}$.Execute.

$\mathcal{F}_{\text{DPC}}.\mathsf{GenAddress}[\mathcal{P}]$ ((optional) address public key $\mathsf{apk}$)

1. Sample randomness $r$ for generating address public key.
2. If $\mathsf{apk} = \perp$ then $\mathsf{apk} \leftarrow \mathsf{SampleAddrPk}(r)$.
3. Check that $\mathsf{apk}$ is **unique**: $\mathsf{Addr}[\mathsf{apk}] = \perp$.
4. Set $\mathsf{Addr}[\mathsf{apk}] := r$.
5. If $\mathcal{P}$ is corrupted: set $S$ to be the set of corrupted parties.
6. If $\mathcal{P}$ is honest: set $S := \{\mathcal{P}\}$.
7. Set $\mathsf{AddrUsers}[\mathsf{apk}] := \mathsf{AddrUsers}[\mathsf{apk}] \cup S$.
8. **Send to $\mathcal{P}$:** address public key $\mathsf{apk}$.

$\mathcal{F}_{\text{DPC}}.\mathsf{ShareRecord}[\mathcal{P}]$ $\begin{pmatrix} \text{record} & \mathbb{r} \\ \text{recipient party} & \mathcal{P}' \end{pmatrix}$

1. If $\mathsf{Records}[\mathbb{r}.\mathsf{cm}] \neq \perp$:
   (a) Check that $\mathcal{P} \in \mathsf{RecUsers}[\mathbb{r}.\mathsf{cm}]$.
   (b) Retrieve $((\mathsf{cm}, \mathsf{apk}, \mathsf{payload}, \Phi_{\mathsf{b}}, \Phi_{\mathsf{d}}), r) := \mathsf{Records}[\mathbb{r}.\mathsf{cm}]$.
2. If $\mathcal{P}'$ is corrupted: set $S$ to be the set of corrupted parties.
3. If $\mathcal{P}'$ is honest: set $S := \{\mathcal{P}'\}$.
4. Set $\mathsf{RecUsers}[\mathbb{r}.\mathsf{cm}] := \mathsf{RecUsers}[\mathbb{r}.\mathsf{cm}] \cup S$.
5. If $\mathcal{P}$ is honest and $\mathcal{P}'$ isn't, **Send to $\mathcal{P}'$:** $(\texttt{RecordAuth}, (\mathbb{r}, r))$.
6. Else, **Send to $\mathcal{P}'$:** $(\texttt{RecordAuth}, \mathbb{r})$.

$\mathcal{F}_{\text{DPC}}.\mathsf{Execute}[\mathcal{P}]$ $\begin{pmatrix} \text{old records} & [\mathbb{r}_i]_1^m \\ \text{(optional) old serial numbers} & [\mathsf{sn}_i]_1^m \\ \text{(optional) new record commitments} & [\mathsf{cm}_j]_1^n \\ \text{new address public keys} & [\mathsf{apk}_j]_1^n \\ \text{new record payloads} & [\mathsf{payload}_j]_1^n \\ \text{new record birth predicates} & [\Phi_{\mathsf{b},j}]_1^n \\ \text{new record death predicates} & [\Phi_{\mathsf{d},j}]_1^n \\ \text{auxiliary predicate input} & \mathsf{aux} \\ \text{transaction memorandum} & \mathsf{memo} \end{pmatrix}$

1. For each $i \in \{1, \ldots, m\}$:
   (a) Sample randomness $r_i$.
   (b) If $\mathsf{sn}_i = \perp$ then generate **serial number**: $\mathsf{sn}_i \leftarrow \mathsf{SampleSn}(r_i)$.
   (c) Check that $\mathsf{sn}_i$ is **unique**: $\mathsf{SerialNumbers}[\mathsf{sn}_i] = \perp$.
2. For each $j \in \{1, \ldots, n\}$:
   (a) Sample randomness $r_j$.
   (b) If $\mathsf{cm}_j = \perp$ then generate **commitment**: $\mathsf{cm}_j \leftarrow \mathsf{SampleCm}(r_j)$.
   (c) Check that $\mathsf{cm}_j$ is **unique**: $\mathsf{Records}[\mathsf{cm}_j] = \perp$.
   (d) **Construct record**: $\mathbb{r}_j := (\mathsf{cm}_j, \mathsf{apk}_j, \mathsf{payload}_j, \Phi_{\mathsf{b},j}, \Phi_{\mathsf{d},j})$.
3. Define the local data $\mathsf{ldata} := ([\mathbb{r}_i]_1^m, [\mathsf{sn}_i]_1^m, [\mathbb{r}_j]_1^n, \mathsf{aux}, \mathsf{memo})$.
4. For each $i \in \{1, \ldots, m\}$:
   (a) Parse $\mathbb{r}_i$ as $(\mathsf{cm}_i, \mathsf{apk}_i, \mathsf{payload}_i, \Phi_{\mathsf{b},i}, \Phi_{\mathsf{d},i})$.
   (b) Check that, for some randomness $r_i$, **old record $\mathbb{r}_i$ exists**: $((\mathsf{apk}_i, \mathsf{payload}_i, \Phi_{\mathsf{b},i}, \Phi_{\mathsf{d},i}), r_i) = \mathsf{Records}[\mathsf{cm}_i]$.
   (c) Check that $\mathcal{P}$ **is authorized to use** $\mathsf{apk}_i$: $\mathcal{P} \in \mathsf{AddrUsers}[\mathsf{apk}_i]$.
   (d) If $\mathsf{payload}_i.\mathsf{isDummy} = 0$:
      i. Check that **record is unconsumed**: $\mathsf{State}[\mathbb{r}_i] = \texttt{alive}$.
      ii. Check that $\mathcal{P}$ **is authorized to consume** $\mathbb{r}_i$: $\mathcal{P} \in \mathsf{RecUsers}[\mathsf{cm}_i]$.
      iii. Check that $\mathcal{P}$ **is authorized to use** $\mathsf{apk}_i$: $\mathcal{P} \in \mathsf{AddrUsers}[\mathsf{apk}_i]$.
   (e) Check that **death predicate is satisfied**: $\Phi_{\mathsf{d},i}(i \| \mathsf{ldata}) = 1$.
   (f) **Mark it as consumed**: $\mathsf{State}[\mathsf{cm}_i] := \texttt{dead}$.
5. For each $j \in \{1, \ldots, n\}$:
   (a) Check that **birth predicate is satisfied**: $\Phi_{\mathsf{b},j}(j \| \mathsf{ldata}) = 1$.
   (b) **Insert new record $\mathbf{r}_j$**: $\mathsf{Records}[\mathsf{cm}_j] := ((\mathsf{apk}_j, \mathsf{payload}_j, \Phi_{\mathsf{b},j}, \Phi_{\mathsf{d},j}), r_j)$.
   (c) **Mark new record as unconsumed**: $\mathsf{State}[\mathsf{cm}_j] := \texttt{alive}$.
6. **Send to $\mathcal{P}$:** $([\mathbb{r}_j]_1^n)$.
7. **Send to all parties:** $(\texttt{Execute}, [\mathsf{sn}_i]_1^m, [\mathsf{cm}_j]_1^n, \mathsf{memo})$.

**Figure 7:** Ideal functionality $\mathcal{F}_{\text{DPC}}$ of a DPC scheme.

# 4 Construction of decentralized private computation schemes

We describe our construction of a DPC scheme. In Section 4.1 we introduce the building blocks that we use, and in Section 4.2 we describe each algorithm in the scheme. The security proof is provided in Appendix A. We also describe some extensions of our construction, in functionality and in security, in Appendix C.

## 4.1 Building blocks

**CRHs.** A collision-resistant hash function $\mathsf{CRH} = (\mathsf{Setup}, \mathsf{Eval})$ works as follows.
- *Setup:* on input a security parameter, $\mathsf{CRH.Setup}$ samples public parameters $\mathsf{pp}_{\mathsf{CRH}}$.
- *Hashing:* on input public parameters $\mathsf{pp}_{\mathsf{CRH}}$ and message $m$, $\mathsf{CRH.Eval}$ outputs a short hash $h$ of $m$.

Given public parameters $\mathsf{pp}_{\mathsf{CRH}} \leftarrow \mathsf{CRH.Setup}(1^\lambda)$, it is computationally infeasible to find distinct inputs $x$ and $y$ such that $\mathsf{CRH.Eval}(\mathsf{pp}_{\mathsf{CRH}}, x) = \mathsf{CRH.Eval}(\mathsf{pp}_{\mathsf{CRH}}, y)$.

**PRFs.** A pseudorandom function family $\mathsf{PRF} = \{\mathsf{PRF}_x \colon \{0,1\}^* \to \{0,1\}^{O(|x|)}\}_x$, where $x$ denotes the seed, is computationally indistinguishable from a random function family.

**Commitments.** A commitment scheme $\mathsf{CM} = (\mathsf{Setup}, \mathsf{Commit})$ enables a party to generate a (perfectly) hiding and (computationally) binding commitment to a given message.
- *Setup:* on input a security parameter, $\mathsf{CM.Setup}$ samples public parameters $\mathsf{pp}_{\mathsf{CM}}$.
- *Commitment:* on input public parameters $\mathsf{pp}_{\mathsf{CM}}$, message $m$, and randomness $r_{\mathsf{cm}}$, $\mathsf{CM.Commit}$ outputs a commitment $\mathsf{cm}$ to $m$.

We also use a *trapdoor* commitment scheme $\mathsf{TCM} = (\mathsf{Setup}, \mathsf{Commit})$, with the same syntax as above. Auxiliary algorithms (beyond those in $\mathsf{CM}$) enable producing a trapdoor and using it to open a commitment, originally to an empty string, to an arbitrary message. These algorithms are used only in the proof of security, and so we introduce them there (see Appendix A).

**NIZKs.** Non-interactive zero knowledge arguments of knowledge enable a party, known as the *prover*, to convince another party, known as the *verifier*, about knowledge of the witness for an NP statement without revealing any information about the witness (besides what is already implied by the statement being true). This primitive is a tuple $\mathsf{NIZK} = (\mathsf{Setup}, \mathsf{Prove}, \mathsf{Verify})$ with the following syntax.
- *Setup:* on input a security parameter and the specification of an NP relation $\mathcal{R}$, $\mathsf{NIZK.Setup}$ outputs a set of public parameters $\mathsf{pp}_{\mathsf{NIZK}}$ (also known as a *common reference string*).
- *Proving:* on input $\mathsf{pp}_{\mathsf{NIZK}}$ and an instance-witness pair $(\mathbb{x}, \mathbb{w}) \in \mathcal{R}$, $\mathsf{NIZK.Prove}$ outputs a proof $\pi$.
- *Verifying:* on input $\mathsf{pp}_{\mathsf{NIZK}}$, instance $\mathbb{x}$, and proof $\pi$, $\mathsf{NIZK.Verify}$ outputs a decision bit.

*Completeness* states that honestly generated proofs make the verifier accept; *(computational) proof of knowledge* states that if the verifier accepts a proof for an instance then the prover "knows" a witness for it; and *perfect zero knowledge* states that honestly generated proofs can be perfectly simulated, when given a trapdoor to the public parameters. In fact, we require a strong form of (computational) proof of knowledge known as *simulation-extractability*, which states that proofs continue to be proofs of knowledge even when the adversary has seen prior simulated proofs. For more details, see [Sah99, DDO+01, Gro06].

**Remark 4.1.** If $\mathsf{NIZK}$ is additionally *succinct* (i.e., it is a simulation-extractable zkSNARK) then the DPC scheme constructed in this section is also *succinct*. This is the case in our implementation; see Section 8.

## 4.2 Algorithms

Pseudocode for our construction of a DPC scheme is in Fig. 8. The construction involves invoking zero knowledge proofs for the NP relation $\mathcal{R}_{\mathsf{e}}$ described in Fig. 9. The text below is a summary of the construction.

**System setup.** DPC.Setup is a wrapper around the setup algorithms of cryptographic building blocks. It invokes CM.Setup, TCM.Setup, CRH.Setup, and NIZK.Setup to obtain (plain and trapdoor) commitment public parameters $\mathsf{pp}_{\mathsf{CM}}$ and $\mathsf{pp}_{\mathsf{TCM}}$, CRH public parameters $\mathsf{pp}_{\mathsf{CRH}}$, and NIZK public parameters for the NP relation $\mathcal{R}_{\mathsf{e}}$ (see Fig. 9). It then outputs $\mathsf{pp} := (\mathsf{pp}_{\mathsf{CM}}, \mathsf{pp}_{\mathsf{TCM}}, \mathsf{pp}_{\mathsf{CRH}}, \mathsf{pp}_{\mathsf{e}})$.

**Address creation.** DPC.GenAddress constructs an address key pair as follows. The address secret key $\mathsf{ask} = (\mathsf{sk}_{\mathsf{PRF}}, r_{\mathsf{pk}})$ consists of a secret key $\mathsf{sk}_{\mathsf{PRF}}$ for the pseudorandom function PRF, and commitment randomness $r_{\mathsf{pk}}$. The address public key apk is a perfectly hiding commitment to $\mathsf{sk}_{\mathsf{PRF}}$ with randomness $r_{\mathsf{pk}}$.

**Execution.** DPC.Execute produces a transaction attesting that some old records $[\mathbf{r}_i]_1^m$ were consumed and some new records $[\mathbf{r}_j]_1^n$ were created, and that their death and birth predicates were satisfied. First, DPC.Execute computes a ledger membership witness and serial number for every old record. Then, DPC.Execute invokes the following auxiliary function to create record commitments for the new records.

$\mathsf{DPC.ConstructRecord}(\mathsf{pp}, \mathsf{apk}, \mathsf{payload}, \Phi_{\mathsf{b}}, \Phi_{\mathsf{d}}, \rho) \to (\mathbf{r}, \mathsf{cm})$

1. Sample new commitment **randomness** $r$.
2. Construct new record **commitment**: $\mathsf{cm} \leftarrow \mathsf{TCM.Commit}(\mathsf{pp}_{\mathsf{TCM}}, \mathsf{apk} \| \mathsf{payload} \| \Phi_{\mathsf{b}} \| \Phi_{\mathsf{d}} \| \rho;\ r)$.
3. Construct new **record** $\mathbf{r} := \begin{pmatrix} \text{address public key} & \mathsf{apk} & \text{payload} & \mathsf{payload} & \text{comm. rand.} & r \\ \text{serial number nonce} & \rho & \text{predicates} & (\Phi_{\mathsf{b}}, \Phi_{\mathsf{d}}) & \text{commitment} & \mathsf{cm} \end{pmatrix}$.
4. Output $(\mathbf{r}, \mathsf{cm})$.

Information about all records, secret addresses of old records, the desired transaction memorandum memo, and desired auxiliary predicate input aux are collected into the local data ldata (see Fig. 9).

Finally, DPC.Execute produces a proof that all records are well-formed and that several conditions hold.

- *Old records are properly consumed*, namely, for every old record $\mathbf{r}_i \in [\mathbf{r}_i]_1^m$:
  - (if $\mathbf{r}_i$ is not dummy) $\mathbf{r}_i$ *exists*, demonstrated by checking a ledger membership witness for $\mathbf{r}_i$'s commitment;
  - $\mathbf{r}_i$ *has not been consumed*, demonstrated by publishing $\mathbf{r}_i$'s serial number $\mathsf{sn}_i$;
  - $\mathbf{r}_i$*'s death predicate* $\Phi_{\mathsf{d},i}$ *is satisfied*, demonstrated by checking that $\Phi_{\mathsf{d},i}(i \| \mathsf{ldata}) = 1$.
- *New records are property created*, namely, for every new record $\mathbf{r}_j \in [\mathbf{r}_j]_1^n$:
  - $\mathbf{r}_j$*'s serial number is unique*, achieved by generating the nonce $\rho_j$ as $\mathsf{CRH.Eval}(\mathsf{pp}_{\mathsf{CRH}}, j \| \mathsf{sn}_1 \| \ldots \| \mathsf{sn}_m)$;
  - $\mathbf{r}_j$*'s birth predicate* $\Phi_{\mathsf{b},j}$ *is satisfied*, demonstrated by checking that $\Phi_{\mathsf{b},j}(j \| \mathsf{ldata}) = 1$.

The serial number sn of a record $\mathbf{r}$ relative to an address secret key $\mathsf{ask} = (\mathsf{sk}_{\mathsf{PRF}}, r_{\mathsf{pk}})$ is derived by evaluating PRF at $\mathbf{r}$'s serial number nonce $\rho$ with seed $\mathsf{sk}_{\mathsf{PRF}}$. This ensures that sn is pseudorandom even to a party that knows all of $\mathbf{r}$ but not ask (e.g., to a party that created the record for some other party). Note that each predicate receives its own position as input so that it knows to which record in the local data it belongs.

**$\mathsf{DPC.Setup}$**
*Input:* security parameter $1^\lambda$
*Output:* public parameters $\mathsf{pp}$

1. Generate **commitment parameters**:
   $\mathsf{pp}_{\mathsf{CM}} \leftarrow \mathsf{CM.Setup}(1^\lambda)$, $\mathsf{pp}_{\mathsf{TCM}} \leftarrow \mathsf{TCM.Setup}(1^\lambda)$.
2. Generate **CRH parameters**: $\mathsf{pp}_{\mathsf{CRH}} \leftarrow \mathsf{CRH.Setup}(1^\lambda)$.
3. Generate **NIZK parameters for** $\mathcal{R}_{\mathsf{e}}$ (see Figure 9):
   $\mathsf{pp}_{\mathsf{e}} \leftarrow \mathsf{NIZK.Setup}(1^\lambda, \mathcal{R}_{\mathsf{e}})$.
4. Output $\mathsf{pp} := (\mathsf{pp}_{\mathsf{CM}}, \mathsf{pp}_{\mathsf{TCM}}, \mathsf{pp}_{\mathsf{CRH}}, \mathsf{pp}_{\mathsf{e}})$.

**$\mathsf{DPC.GenAddress}$**
*Input:* public parameters $\mathsf{pp}$
*Output:* address key pair $(\mathsf{apk}, \mathsf{ask})$

1. Sample secret key $\mathsf{sk}_{\mathsf{PRF}}$ for pseudorandom function PRF.
2. Sample randomness $r_{\mathsf{pk}}$ for commitment scheme CM.
3. Set **address public key**
   $\mathsf{apk} := \mathsf{CM.Commit}(\mathsf{pp}_{\mathsf{CM}}, \mathsf{sk}_{\mathsf{PRF}};\ r_{\mathsf{pk}})$.
4. Set **address secret key** $\mathsf{ask} := (\mathsf{sk}_{\mathsf{PRF}}, r_{\mathsf{pk}})$.
5. Output $(\mathsf{apk}, \mathsf{ask})$.

**$\mathsf{DPC.Execute}^{\mathbf{L}}$**
*Input:*
- public parameters $\mathsf{pp}$
- old { records $[\mathbf{r}_i]_1^m$; address secret keys $[\mathsf{ask}_i]_1^m$ }
- new { address public keys $[\mathsf{apk}_j]_1^n$; record payloads $[\mathsf{payload}_j]_1^n$; record birth predicates $[\Phi_{\mathsf{b},j}]_1^n$; record death predicates $[\Phi_{\mathsf{d},j}]_1^n$ }
- auxiliary predicate input $\mathsf{aux}$
- transaction memorandum $\mathsf{memo}$

*Output:* new records $[\mathbf{r}_j]_1^n$ and transaction $\mathsf{tx}$

1. For each $i \in \{1, \ldots, m\}$, process the $i$-th old record as follows:
   (a) Parse old record $\mathbf{r}_i$ as $\begin{pmatrix} \text{address public key} & \mathsf{apk}_i & \text{payload} & \mathsf{payload}_i & \text{comm. rand.} & r_i \\ \text{serial number nonce} & \rho_i & \text{predicates} & (\Phi_{\mathsf{b},i}, \Phi_{\mathsf{d},i}) & \text{commitment} & \mathsf{cm}_i \end{pmatrix}$.
   (b) If $\mathsf{payload}_i.\mathsf{isDummy} = 1$, set **ledger membership witness** $\mathrm{w}_{\mathbf{L},i} := \bot$.
   If $\mathsf{payload}_i.\mathsf{isDummy} = 0$, compute **ledger membership witness** for commitment: $\mathrm{w}_{\mathbf{L},i} \leftarrow \mathbf{L}.\mathsf{Prove}(\mathsf{cm}_i)$.
   (c) Parse address secret key $\mathsf{ask}_i$ as $(\mathsf{sk}_{\mathsf{PRF},i}, r_{\mathsf{pk},i})$.
   (d) Compute **serial number**: $\mathsf{sn}_i \leftarrow \mathsf{PRF}_{\mathsf{sk}_{\mathsf{PRF},i}}(\rho_i)$.
2. For each $j \in \{1, \ldots, n\}$, construct the $j$-th new record as follows:
   (a) Compute **serial number nonce**: $\rho_j := \mathsf{CRH.Eval}(\mathsf{pp}_{\mathsf{CRH}}, j \| \mathsf{sn}_1 \| \ldots \| \mathsf{sn}_m)$.
   (b) Construct **new record**: $(\mathbf{r}_j, \mathsf{cm}_j) \leftarrow \mathsf{DPC.ConstructRecord}(\mathsf{pp}_{\mathsf{TCM}}, \mathsf{apk}_j, \mathsf{payload}_j, \Phi_{\mathsf{b},j}, \Phi_{\mathsf{d},j}, \rho_j)$.
3. Retrieve current **ledger digest**: $\mathsf{st}_{\mathbf{L}} \leftarrow \mathbf{L}.\mathsf{Digest}$.
4. Construct **instance** $\mathrm{x}_{\mathsf{e}}$ **for** $\mathcal{R}_{\mathsf{e}}$: $\mathrm{x}_{\mathsf{e}} := (\mathsf{st}_{\mathbf{L}}, [\mathsf{sn}_i]_1^m, [\mathsf{cm}_j]_1^n, \mathsf{memo})$.
5. Construct **witness** $\mathrm{w}_{\mathsf{e}}$ **for** $\mathcal{R}_{\mathsf{e}}$: $\mathrm{w}_{\mathsf{e}} := ([\mathbf{r}_i]_1^m, [\mathrm{w}_{\mathbf{L},i}]_1^m, [\mathsf{ask}_i]_1^m, [\mathbf{r}_j]_1^n, \mathsf{aux})$.
6. Generate **proof for** $\mathcal{R}_{\mathsf{e}}$: $\pi_{\mathsf{e}} \leftarrow \mathsf{NIZK.Prove}(\mathsf{pp}_{\mathsf{e}}, \mathrm{x}_{\mathsf{e}}, \mathrm{w}_{\mathsf{e}})$.
7. Construct **transaction**: $\mathsf{tx} := ([\mathsf{sn}_i]_1^m, [\mathsf{cm}_j]_1^n, \mathsf{memo}, \star)$, where $\star := (\mathsf{st}_{\mathbf{L}}, \pi_{\mathsf{e}})$.
8. Output $([\mathbf{r}_j]_1^n, \mathsf{tx})$.

**$\mathsf{DPC.Verify}^{\mathbf{L}}$**
*Input:* public parameters $\mathsf{pp}$ and transaction $\mathsf{tx}$
*Output:* decision bit $b$

1. Parse $\mathsf{tx}$ as $([\mathsf{sn}_i]_1^m, [\mathsf{cm}_j]_1^n, \mathsf{memo}, \star)$ and $\star$ as $(\mathsf{st}_{\mathbf{L}}, \pi_{\mathsf{e}})$.
2. Check that **there are no duplicate serial numbers**
   (a) within the transaction $\mathsf{tx}$: $\mathsf{sn}_i \neq \mathsf{sn}_j$ for every distinct $i, j \in \{1, \ldots, m\}$;
   (b) on the ledger: $\mathbf{L}.\mathsf{Contains}(\mathsf{sn}_i) = 0$ for every $i \in \{1, \ldots, m\}$.
3. Check that **the ledger state is valid**: $\mathbf{L}.\mathsf{ValidateDigest}(\mathsf{st}_{\mathbf{L}}) = 1$.
4. Construct **instance for the relation** $\mathcal{R}_{\mathsf{e}}$: $\mathrm{x}_{\mathsf{e}} := (\mathsf{st}_{\mathbf{L}}, [\mathsf{sn}_i]_1^m, [\mathsf{cm}_j]_1^n, \mathsf{memo})$.
5. Check **proof for the relation** $\mathcal{R}_{\mathsf{e}}$: $\mathsf{NIZK.Verify}(\mathsf{pp}_{\mathsf{e}}, \mathrm{x}_{\mathsf{e}}, \pi_{\mathsf{e}}) = 1$.

**Figure 8:** Construction of a DPC scheme.

$$\mathbb{x}_{\mathsf{e}} = \begin{pmatrix} \text{ledger digest} & \mathsf{st}_{\mathbf{L}} \\ \text{old record serial numbers} & [\mathsf{sn}_i]_1^m \\ \text{new record commitments} & [\mathsf{cm}_j]_1^n \\ \text{transaction memorandum} & \mathsf{memo} \end{pmatrix} \quad \text{and} \quad \mathbb{w}_{\mathsf{e}} = \begin{pmatrix} \text{old records} & [\mathbf{r}_i]_1^m \\ \text{old record membership witnesses} & [\mathbb{w}_{\mathbf{L},i}]_1^m \\ \text{old address secret keys} & [\mathsf{ask}_i]_1^m \\ \text{new records} & [\mathbf{r}_j]_1^n \\ \text{auxiliary predicate input} & \mathsf{aux} \end{pmatrix}$$

where

- for each $i \in \{1, \ldots, m\}$, $\mathbf{r}_i = (\mathsf{apk}_i, \mathsf{payload}_i, \Phi_{\mathsf{b},i}, \Phi_{\mathsf{d},i}, \rho_i, r_i, \mathsf{cm}_i)$;
- for each $j \in \{1, \ldots, n\}$, $\mathbf{r}_j = (\mathsf{apk}_j, \mathsf{payload}_j, \Phi_{\mathsf{b},j}, \Phi_{\mathsf{d},j}, \rho_j, r_j, \mathsf{cm}_j)$.

Define the local data $\mathsf{ldata} := \begin{pmatrix} [\mathsf{cm}_i]_1^m & [\mathsf{apk}_i]_1^m & [\mathsf{payload}_i]_1^m & [\Phi_{\mathsf{d},i}]_1^m & [\Phi_{\mathsf{b},i}]_1^m & [\mathsf{sn}_i]_1^m & \mathsf{memo} \\ [\mathsf{cm}_j]_1^n & [\mathsf{apk}_j]_1^n & [\mathsf{payload}_j]_1^n & [\Phi_{\mathsf{d},j}]_1^n & [\Phi_{\mathsf{b},j}]_1^n & \mathsf{aux} & \end{pmatrix}$.

Then, a witness $\mathbb{w}_{\mathsf{e}}$ is valid for an instance $\mathbb{x}_{\mathsf{e}}$ if the following conditions hold:

1. For each $i \in \{i, \ldots, m\}$:
   - If $\mathbf{r}_i$ is not dummy, $\mathbb{w}_{\mathbf{L},i}$ proves that the commitment $\mathsf{cm}_i$ is in a ledger with digest $\mathsf{st}_{\mathbf{L}}$: $\mathbf{L}.\mathsf{Verify}(\mathsf{st}_{\mathbf{L}}, \mathsf{cm}_i, \mathbb{w}_{\mathbf{L},i}) = 1$.
   - The address public key $\mathsf{apk}_i$ and secret key $\mathsf{ask}_i$ form a valid key pair:
     $\mathsf{apk}_i = \mathsf{CM}.\mathsf{Commit}(\mathsf{pp}_{\mathsf{CM}}, \mathsf{sk}_{\mathsf{PRF},i};\ r_{\mathsf{pk},i})$ and $\mathsf{ask}_i = (\mathsf{sk}_{\mathsf{PRF},i}, r_{\mathsf{pk},i})$.
   - The serial number $\mathsf{sn}_i$ is valid: $\mathsf{sn}_i = \mathsf{PRF}_{\mathsf{sk}_{\mathsf{PRF},i}}(\rho_i)$.
   - The old record commitment $\mathsf{cm}_i$ is valid: $\mathsf{cm}_i = \mathsf{TCM}.\mathsf{Commit}(\mathsf{pp}_{\mathsf{TCM}}, \mathsf{apk}_i \| \mathsf{payload}_i \| \Phi_{\mathsf{b},i} \| \Phi_{\mathsf{d},i} \| \rho_i;\ r_i)$.
   - The death predicate $\Phi_{\mathsf{d},i}$ is satisfied by local data: $\Phi_{\mathsf{d},i}(i \| \mathsf{ldata}) = 1$.
2. For each $j \in \{1, \ldots, n\}$:
   - The serial number nonce $\rho_j$ is computed correctly: $\rho_j = \mathsf{CRH}.\mathsf{Eval}(\mathsf{pp}_{\mathsf{CRH}}, j \| \mathsf{sn}_1 \| \ldots \| \mathsf{sn}_m)$.
   - The new record commitment $\mathsf{cm}_j$ is valid: $\mathsf{cm}_j = \mathsf{TCM}.\mathsf{Commit}(\mathsf{pp}_{\mathsf{TCM}}, \mathsf{apk}_j \| \mathsf{payload}_j \| \Phi_{\mathsf{b},j} \| \Phi_{\mathsf{d},j} \| \rho_j;\ r_j)$.
   - The birth predicate $\Phi_{\mathsf{b},j}$ is satisfied by local data: $\Phi_{\mathsf{b},j}(j \| \mathsf{ldata}) = 1$.

**Figure 9:** The execute NP relation $\mathcal{R}_{\mathsf{e}}$.

# 5 Delegating zero knowledge execution

The cost of creating a transaction in the DPC scheme from Section 4 grows with the complexity (and number of) predicates involved in the transaction. The user must produce, and include in the transaction, a cryptographic proof that, among other things, attests that death predicates of consumed records are satisfied and, similarly, that birth predicates of created records are satisfied. This implies that producing transactions on weak devices such as mobile phones or hardware tokens quickly becomes infeasible.

In Sections 5.1 to 5.3 we explain how to address this problem by enabling a user to *delegate* to an untrusted worker, such as a remote server, the computation that produces a transaction. This empowers weak devices to produce transactions that they otherwise could not have produced on their own. Then, in Section 5.4, we explain how the ideas that we use for delegating transactions also yield solutions for achieving *threshold transactions* and *blind transactions* in a DPC scheme, which are also valuable in applications. Techniques derived from these ideas are now part of deployed systems [HBHW20].

## 5.1 Approach

A naive approach is for the user to simply ask the worker to produce the cryptographic proof on its behalf, and then include this proof in the transaction. The intuition behind this idea is that the user can check that the proof received from the worker is valid, by simply running the proof verification procedure. Indeed, whenever the DPC scheme uses a succinct argument (see Remark 4.1), the verification procedure is *succinct*.

However, this approach is *insecure*, because the worker, in order to produce a proof, would have to learn not only the instance but also the secret witness for the NP statement being proved. Since the secret witness includes the user's address secret key, if the worker learns this information then the worker can impersonate the user, e.g., by producing further transactions that the user never *authorized*. This naive approach also fails in prior proof-based ledger protocols, including Zerocash [BCG$^+$14]. New ideas are needed.

Taking our construction of a DPC scheme from Section 4 as a starting point, we explain how to enable a user to delegate the expensive proof computation to a worker in such a way that the worker *cannot* produce valid transactions that have not been authorized by the user; see Fig. 11. (Additional security goals, such as ensuring that the worker learns no information about the user, are left to future work.)

The basic idea is to augment address keys in such a way that the secret information needed to produce the cryptographic proof is separate from the secret information needed to authorize a transaction containing that proof. Thus, the user can communicate to the worker the secrets necessary to generate a cryptographic proof, while retaining the remaining secrets for authorizing this (and future) transactions. In particular, the worker has no way to produce valid transactions that have not been authorized by the user.

We stress that the simplistic solution in which the user authorizes the proof produced by the worker by signing it via a secret key not shared with the worker *does not work* because it violates privacy. Indeed, others would have to use the same public key to verify signatures across multiple transactions containing signatures produced by the same secret key, thereby linking these transactions together.

The next two sub-sections explain how we achieve delegation: first, in Section 5.2, we describe a variant of randomizable signatures, which we use as a building block; then, in Section 5.3, we provide a high-level description of a *delegable* DPC scheme. The detailed construction is provided in Appendix B.

## 5.2 Additional building block: randomizable signatures

A *randomizable* signature scheme is a tuple of algorithms $\mathsf{SIG} = (\mathsf{Setup}, \mathsf{Keygen}, \mathsf{Sign}, \mathsf{Verify}, \mathsf{RandPk}, \mathsf{RandSig})$ that enables a party to sign messages, while also allowing randomization of public keys and

signatures to prevent linking across multiple signatures. We first discuss the syntax of the usual algorithms.

- *Setup:* on input a security parameter, $\mathsf{SIG.Setup}$ samples public parameters $\mathsf{pp}_{\mathsf{SIG}}$.
- *Key generation:* on input public parameters $\mathsf{pp}_{\mathsf{SIG}}$, $\mathsf{SIG.Keygen}$ samples a key pair $(\mathsf{pk}_{\mathsf{SIG}}, \mathsf{sk}_{\mathsf{SIG}})$.
- *Message signing:* on input public parameters $\mathsf{pp}_{\mathsf{SIG}}$, secret key $\mathsf{sk}_{\mathsf{SIG}}$, and message $m$, $\mathsf{SIG.Sign}$ produces a signature $\sigma$.
- *Signature verification:* on input public parameters $\mathsf{pp}_{\mathsf{SIG}}$, public key $\mathsf{pk}_{\mathsf{SIG}}$, message $m$, and signature $\sigma$, $\mathsf{SIG.Verify}$ outputs a bit $b$ denoting whether $\sigma$ is a valid signature for $m$ under public key $\mathsf{pk}_{\mathsf{SIG}}$.

In addition to the usual algorithms, $\mathsf{SIG}$ has two algorithms for randomizing public keys and signatures.

- *Public key randomization:* $\mathsf{SIG.RandPk}(\mathsf{pp}_{\mathsf{SIG}}, \mathsf{pk}_{\mathsf{SIG}}, r_{\mathsf{SIG}})$ samples a randomized public key $\hat{\mathsf{pk}}_{\mathsf{SIG}}$.
- *Signature randomization:* $\mathsf{SIG.RandSig}(\mathsf{pp}_{\mathsf{SIG}}, \sigma, r_{\mathsf{SIG}})$ samples a randomized signature $\hat{\sigma}$.

The signature scheme $\mathsf{SIG}$ must satisfy the following security properties.

- *Existential unforgeability.* Given a public key $\mathsf{pk}_{\mathsf{SIG}}$, it is infeasible to produce a forgery under $\mathsf{pk}_{\mathsf{SIG}}$ or *under under any randomization of* $\mathsf{pk}_{\mathsf{SIG}}$. This notion strengthens the standard unforgeability notion, and is similar to that of randomizable signatures in [FKM+16].
- *Unlinkability.* Given a public key $\mathsf{pk}_{\mathsf{SIG}}$ and a tuple $(\hat{\mathsf{pk}}_{\mathsf{SIG}}, m, \hat{\sigma})$ where $\hat{\sigma}$ is a valid signature for $m$ under $\hat{\mathsf{pk}}_{\mathsf{SIG}}$, no efficient adversary can determine if $\hat{\mathsf{pk}}_{\mathsf{SIG}}$ is a fresh public key and $\hat{\sigma}$ a fresh signature, or if instead $\hat{\mathsf{pk}}_{\mathsf{SIG}}$ is a randomization of $\mathsf{pk}_{\mathsf{SIG}}$ and $\hat{\sigma}$ a randomization of a signature for $\mathsf{pk}_{\mathsf{SIG}}$. This property is a computational relaxation of the perfect unlinkability property of randomizable signatures in [FKM+16].
- *Injective randomization.* Randomization of public keys is *(computationally) injective* with respect to randomness. Informally, given public parameters $\mathsf{pp}_{\mathsf{SIG}}$, it is infeasible to find a public key $\mathsf{pk}_{\mathsf{SIG}}$ and $r_1 \neq r_2$ such that $\mathsf{SIG.RandPk}(\mathsf{pp}_{\mathsf{SIG}}, \mathsf{pk}_{\mathsf{SIG}}, r_1) = \mathsf{SIG.RandPk}(\mathsf{pp}_{\mathsf{SIG}}, \mathsf{pk}_{\mathsf{SIG}}, r_2)$.

### 5.3 A delegable DPC scheme

We describe how to construct a *delegable DPC scheme*, namely, a DPC scheme in which a user can delegate to an untrusted worker the expensive computations associated with producing a transaction. The security goal is that the worker should not be able to produce valid transactions that have not been authorized by the user. Below we assume familiarity with our "plain" DPC construction (see Section 4).

The user will maintain (among other things) a key pair $(\mathsf{pk}_{\mathsf{SIG}}, \mathsf{sk}_{\mathsf{SIG}})$ for a randomizable signature scheme $\mathsf{SIG}$ (see Section 5.2). The public key $\mathsf{pk}_{\mathsf{SIG}}$ will be embedded in the user's public key $\mathsf{apk}$ and also be used to derive the serial numbers of records "owned" by $\mathsf{apk}$. In contrast, the secret key $\mathsf{sk}_{\mathsf{SIG}}$ will not be a part of any data structures, and will *only* be used to *authorize* transactions by signing the cryptographic proofs produced by untrusted workers.

In more detail, we first describe how addresses and records are generated (also see summary in Fig. 10).

- **Addresses.** In Section 4 an address public key $\mathsf{apk}$ was a commitment to a secret key $\mathsf{sk}_{\mathsf{PRF}}$ for a pseudorandom function $\mathsf{PRF}$. Now $\mathsf{apk}$ is a commitment to this same information *as well as* the public key of a key pair $(\mathsf{pk}_{\mathsf{SIG}}, \mathsf{sk}_{\mathsf{SIG}})$ for $\mathsf{SIG}$. The corresponding address secret key $\mathsf{ask}$ consists of all the committed information and the commitment randomness.

- **Records.** The structure of a record, including how a record commitment is computed, is as in Section 4. However, a record's serial number $\mathsf{sn}$ is now derived in a different way: while previously $\mathsf{sn} := \mathsf{PRF}_{\mathsf{sk}_{\mathsf{PRF}}}(\rho)$ now we set $\mathsf{sn} := \mathsf{SIG.RandPk}(\mathsf{pp}_{\mathsf{SIG}}, \mathsf{pk}_{\mathsf{SIG}}, \mathsf{PRF}_{\mathsf{sk}_{\mathsf{PRF}}}(\rho))$ where $\rho$ is the record's serial number nonce. Namely, while before serial numbers were outputs of a pseudorandom function keyed by $\mathsf{sk}_{\mathsf{PRF}}$, now they are randomizations of the authorization public key $\mathsf{pk}_{\mathsf{SIG}}$ when using suitable pseudorandomness.

Note that the foregoing new derivation of serial numbers does not break important security properties.

– *Unlinkability of serial numbers:* serial numbers of different records that share the same authorization public $\mathsf{pk}_{\mathsf{SIG}}$ are computationally indistinguishable. This follows rather directly from the fact $\mathsf{sn}$, being a randomization of $\mathsf{pk}_{\mathsf{SIG}}$, does not reveal information (to efficient distinguishers) about $\mathsf{pk}_{\mathsf{SIG}}$ itself.
– *No double spending:* a user cannot "spend" (i.e., consume) $\mathbf{r}$ in two different transactions by revealing different serial numbers because $r_{\mathsf{SIG}}$ (and thus $\mathsf{sn}$) is generated deterministically from $\mathbf{r}$. Since $\mathsf{SIG}$ is randomness-injective in $\mathsf{SIG.RandPk}$, $\mathsf{sn}$ is (computationally) unique to $\mathbf{r}$.

Having described the modified data structures of addresses and serial numbers, we now explain how a user can task a worker to produce the cryptographic proofs that need to be included in a transaction. For simplicity, in this high-level discussion we focus on the case where the transaction involves only one input (old) record $\mathbf{r}$ and one output (new) record $\mathbf{r}'$. In this case, the transaction contains a serial number $\mathsf{sn}$ (supposedly corresponding to $\mathbf{r}$), and a commitment $\mathsf{cm}'$ (supposedly corresponding to $\mathbf{r}'$).

Previously, the user had to generate a proof $\pi_{\mathsf{e}}$ that $\mathsf{sn}$ is consistent with $\mathbf{r}$, that $\mathsf{cm}'$ can be opened to $\mathbf{r}'$, and that the death and birth predicates of $\mathbf{r}$ and $\mathbf{r}'$ respectively are satisfied. Now the user can delegate to a worker the generation of the proof $\pi_{\mathsf{e}}$ because the modified derivation of $\mathsf{apk}$ and $\mathsf{sn}$ allows the user to communicate to the worker only $\mathbf{r}$, $\mathbf{r}'$ and a *part* of the address secret key of $\mathbf{r}$. Namely, the user sends to the worker only the pseudorandom function key $\mathsf{sk}_{\mathsf{PRF}}$ and the commitment randomness $r_{\mathsf{pk}}$. Crucially, the user does not have to communicate to the worker the authorization secret key $\mathsf{sk}_{\mathsf{SIG}}$.

After receiving the proof $\pi_{\mathsf{e}}$ from the worker, the user uses the authorization secret key $\mathsf{sk}_{\mathsf{SIG}}$ to sign $\pi_{\mathsf{e}}$ (along with the instance that $\pi_{\mathsf{e}}$ attests to), and then randomizes the resulting signature $\sigma$ to obtain $\hat{\sigma}$. The final transaction $\mathsf{tx}$ not only includes the serial number $\mathsf{sn}$ (consuming the old record), the commitment $\mathsf{cm}'$ (creating the new record), and $\pi_{\mathsf{e}}$ (attesting to the correct state transition) as before, but also includes $\hat{\sigma}$. Transaction verification involves checking the proof $\pi_{\mathsf{e}}$ and also checking that $\hat{\sigma}$ is valid with respect to the randomized public key $\mathsf{sn}$.

This completes our high-level description of our delegable DPC scheme; see Appendix B for details.

| | Plain DPC | Delegable DPC |
|---|---|---|
| Address secret key | $(\mathsf{sk}_{\mathsf{PRF}}, r_{\mathsf{pk}})$ | $(\mathsf{sk}_{\mathsf{SIG}}, \mathsf{sk}_{\mathsf{PRF}}, r_{\mathsf{pk}})$ |
| Address public key | $\mathsf{apk} := \mathsf{CM.Commit}\left(\begin{smallmatrix}\mathsf{pp}_{\mathsf{CM}},\\ \mathsf{sk}_{\mathsf{PRF}}\end{smallmatrix}; r_{\mathsf{pk}}\right)$ | $\mathsf{apk} := \mathsf{CM.Commit}\left(\begin{smallmatrix}\mathsf{pp}_{\mathsf{CM}},\\ \mathsf{pk}_{\mathsf{SIG}}\|\mathsf{sk}_{\mathsf{PRF}}\end{smallmatrix}; r_{\mathsf{pk}}\right)$ |
| Serial number derivation | $\mathsf{sn} \leftarrow \mathsf{PRF}_{\mathsf{sk}_{\mathsf{PRF}}}(\rho)$ | 1. $r_{\mathsf{SIG}} \leftarrow \mathsf{PRF}_{\mathsf{sk}_{\mathsf{PRF}}}(\rho)$<br>2. $\mathsf{sn} \leftarrow \mathsf{SIG.RandPk}(\mathsf{pp}_{\mathsf{SIG}}, \mathsf{pk}_{\mathsf{SIG}}, r_{\mathsf{SIG}})$ |
| Transaction construction | $\mathsf{tx} := ([\mathsf{sn}_i]_1^m, [\mathsf{cm}_j]_1^n, \mathsf{memo}, \star)$,<br>where $\star := (\mathsf{st}_{\mathbf{L}}, \pi_{\mathsf{e}})$. | 1. **Sign transaction contents:**<br>(a) $\sigma_i \leftarrow \mathsf{SIG.Sign}(\mathsf{pp}_{\mathsf{SIG}}, \mathsf{sk}_{\mathsf{SIG},i}, \mathbb{x}_{\mathsf{e}}\|\pi_{\mathsf{e}})$.<br>(b) $\hat{\sigma}_i \leftarrow \mathsf{SIG.RandSig}(\mathsf{pp}_{\mathsf{SIG}}, \sigma_i, r_{\mathsf{SIG},i})$.<br>2. $\mathsf{tx} := ([\mathsf{sn}_i]_1^m, [\mathsf{cm}_j]_1^n, \mathsf{memo}, \star)$,<br>where $\star := (\mathsf{st}_{\mathbf{L}}, \pi_{\mathsf{e}}, [\hat{\sigma}_i]_1^m)$. |
| Transaction verification | Check that serial numbers do not appear on ledger, that the ledger state digest is valid, and that the NIZK proof verifies. | As in plain DPC, but additionally check that each **signature verifies:**<br>$\mathsf{SIG.Verify}(\mathsf{pp}_{\mathsf{SIG}}, \mathsf{tx.sn}_i, \mathbb{x}_{\mathsf{e}}\|\pi_{\mathsf{e}}, \sigma_i) = 1$. |

**Figure 10:** Summary of differences between plain DPC and delegable DPC (highlighted).

## 5.4 Threshold transactions and blind transactions

We explain how the delegable DPC scheme described above can be modified, in a straightforward way, to achieve additional features: *threshold transactions* or *blind transactions*.

**Threshold transactions.** A DPC scheme has threshold transactions if the power to authorize transactions can be vested unto any $t$ out of $n$ parties, for any desired choice of $t$ and $n$ (as opposed to a single user as discussed thus far, which corresponds to the special case of $t = n = 1$); see Fig. 12. Threshold transactions are useful in many settings, e.g., to enhance operational security by realizing two-factor authentication.

We can achieve threshold transactions by simply using, in our delegable DPC scheme, a randomizable signature scheme $\mathsf{SIG}$ that also supports *threshold key generation* and *threshold signing algorithms* [DF91]. Such a *threshold signature scheme* distributes signing ability among $n$ parties such that at least $t$ of them are needed to authorize a signature. Threshold key generation would then be used to create an address, and threshold signing would be used to authorize a transaction by signing the corresponding cryptographic proof.

**Blind transactions.** A DPC scheme has blind transactions if there is a way for a user to authorize a transaction without learning of its contents; see Fig. 13. Blind transactions, in conjunction with prior techniques [CGL+17], can be used to construct efficient lottery tickets and thereby probabilistic micropayments.

We can achieve blind transactions by simply using, in our delegable DPC scheme, a randomizable signature scheme $\mathsf{SIG}$ that has a *blind signing algorithm*, which can then be used for signing the relevant cryptographic proof in order to authorize a transaction.

**Instantiating randomizable threshold and blind signatures.** As we explain in Appendix B.1, we construct randomizable signature schemes by modifying Schnorr signatures. To further construct threshold or blind randomizable signatures, it is enough to note that public key and signature randomization occurs *after* the public key or signature has been created. Thus one can use existing protocols for threshold key-generation and signing [SS01, NKDM03, Dod07], and blind signing [PS00, SJ99] to obtain public keys and signatures, and then use the algorithms from Appendix B.1 to randomize these. A nice feature of this approach is that all these types of delegated transactions (regular, threshold, blind) cannot be distinguished from one another.

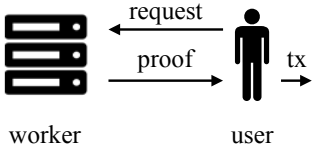


**Figure 11:** Delegable transactions.

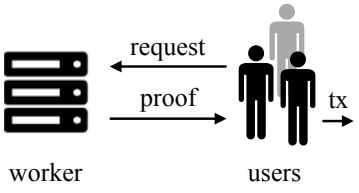


**Figure 12:** Threshold transactions.

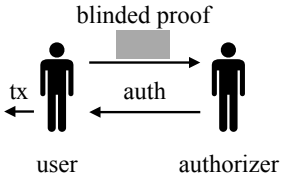


**Figure 13:** Blind transactions.

# 6 Applications

We describe example applications of DPC schemes, by showing how to "program" these within the records nano-kernel. We draw inspiration from current uses of smart contract systems (e.g., Ethereum), which largely focus on financial applications where privacy is an important goal. First, in Section 6.1 we describe how to enable users to privately create and transact with **custom assets** (expanding on Example 2.1). Second, in Section 6.2 we describe how to realize **private DEXs**, which enable users to privately trade assets while retaining custody of their assets. Finally, in Section 6.3, we describe how a central authority can issue assets with self-enforcing, *and updatable*, policies, and use these to realize regulation-friendly **private stablecoins**.

## 6.1 User-defined assets

One of the most basic applications of smart contract systems like Ethereum is the construction of *assets* (or *tokens*) that can be used for financial applications. For example, the Ethereum ERC20 specification [VB15] defines a general framework for such assets. These assets have two phases: asset minting (creation), and asset conservation (expenditure). We show below how to express such custom assets via the records nano-kernel.

We consider records whose payloads encode: an asset identifier $\mathsf{id}$, the initial asset supply $\mathbb{v}$, a value $v$, and application-dependent data $c$ (we will use this in Sections 6.2 and 6.3). We fix the birth predicate in all such records to be a *mint-or-conserve* function $\mathsf{MoC}$ that is responsible for asset minting and conservation. In more detail, the birth predicate $\mathsf{MoC}$ can be invoked in two modes, *mint mode* or *conserve mode*.

When invoked in mint mode, $\mathsf{MoC}$ creates the initial supply $\mathbb{v}$ of the asset in, say, a single output record, by deterministically deriving a fresh unique identifier $\mathsf{id}$ for the asset (see below for how), and storing the tuple $(\mathsf{id}, \mathbb{v}, \mathbb{v}, \bot)$ in the record's payload. The predicate $\mathsf{MoC}$ also ensures that the given transaction contains no input records or other output records (dummy records are allowed). If $\mathsf{MoC}$ is invoked in mint mode in other transactions, a *different* identifier $\mathsf{id}$ is created, ensuring that multiple assets can be distinguished even though anyone can use $\mathsf{MoC}$ as the birth predicate of a record.

When invoked in conserve mode, $\mathsf{MoC}$ inspects all records in a transaction whose birth predicates all equal $\mathsf{MoC}$ (i.e., all the transaction's user-defined assets) and whose asset identifiers all equal to the identifier of the current record. For these records it ensures that no new value is created: that is, the sum of the value across all output records is less than or equal to the sum of the value in all input records.

Below we provide pseudocode for $\mathsf{MoC}$, making the informal discussion above more precise.

Mint-or-conserve predicate $\mathsf{MoC}(k, \mathsf{ldata}; \mathsf{mode})$ (mode is the private input of the predicate)

1. Parse $\mathsf{ldata}$ as $\begin{pmatrix} [\mathsf{cm}_i^{\mathsf{in}}]_1^2 & [\mathsf{apk}_i^{\mathsf{in}}]_1^2 & [\mathsf{payload}_i^{\mathsf{in}}]_1^2 & [\Phi_{\mathsf{d},i}^{\mathsf{in}}]_1^2 & [\Phi_{\mathsf{b},i}^{\mathsf{in}}]_1^2 & [\mathsf{sn}_i^{\mathsf{in}}]_1^2 & \mathsf{memo} \\ [\mathsf{cm}_j^{\mathsf{out}}]_1^2 & [\mathsf{apk}_j^{\mathsf{out}}]_1^2 & [\mathsf{payload}_j^{\mathsf{out}}]_1^2 & [\Phi_{\mathsf{d},j}^{\mathsf{out}}]_1^2 & [\Phi_{\mathsf{b},j}^{\mathsf{out}}]_1^2 & & \mathsf{aux} \end{pmatrix}$.
2. If $\mathsf{mode} = (\mathtt{mint}, \mathbb{v}, r)$, ensure that the first output record contains the initial supply of the asset:
   (a) the index of the current output record is correct: $k = 1$.
   (b) all other records are dummy: $\mathsf{payload}_1^{\mathsf{in}}.\mathsf{isDummy} = \mathsf{payload}_2^{\mathsf{in}}.\mathsf{isDummy} = \mathsf{payload}_2^{\mathsf{out}}.\mathsf{isDummy} = 1$.
   (c) the asset identifier is derived correctly: $\mathsf{id} = \mathsf{CM.Commit}(\mathsf{pp}_{\mathsf{CM}}, \mathsf{sn}_1 \| \mathsf{sn}_2; r)$. (See explanation below.)
   (d) the current output record's payload is correct: $\mathsf{payload}_1^{\mathsf{out}}.\mathsf{isDummy} = 0$ and $\mathsf{payload}_1^{\mathsf{out}} = (\mathsf{id}, \mathbb{v}, \mathbb{v}, \bot)$.
3. If $\mathsf{mode} = \mathtt{conserve}$, check that the value of the current asset is conserved:
   (a) parse the current output record's payload $\mathsf{payload}_k^{\mathsf{out}}$ as $(\mathsf{id}^\star, \mathbb{v}^\star, v^\star, c^\star)$.
   (b) for $i \in \{1, 2\}$, parse the $i$-th input record's payload $\mathsf{payload}_i^{\mathsf{in}}$ as $(\mathsf{id}_i^{\mathsf{in}}, \mathbb{v}_i^{\mathsf{in}}, v_i^{\mathsf{in}}, c_i^{\mathsf{in}})$.
   (c) for $j \in \{1, 2\}$, parse the $i$-th output record's payload $\mathsf{payload}_j^{\mathsf{out}}$ as $(\mathsf{id}_j^{\mathsf{out}}, \mathbb{v}_j^{\mathsf{out}}, v_j^{\mathsf{out}}, c_j^{\mathsf{out}})$.
   (d) initialize $v^{\mathsf{in}} := 0$ and $v^{\mathsf{out}} := 0$, representing the value of asset $\mathsf{id}^\star$ consumed and created (respectively).
   (e) for $i \in \{1, 2\}$, if $\Phi_{\mathsf{b},i}^{\mathsf{in}} = \Phi_{\mathsf{b}}^\star$, $\mathsf{id}_i^{\mathsf{in}} = \mathsf{id}^\star$, $\mathsf{payload}_i^{\mathsf{in}}.\mathsf{isDummy} = 0$, set $v^{\mathsf{in}} := v^{\mathsf{in}} + v_i^{\mathsf{in}}$ and check that $\mathbb{v}_i^{\mathsf{in}} = \mathbb{v}^\star$.
   (f) for $j \in \{1, 2\}$, if $\Phi_{\mathsf{b},j}^{\mathsf{out}} = \Phi_{\mathsf{b}}^\star$, $\mathsf{id}_j^{\mathsf{out}} = \mathsf{id}^\star$, $\mathsf{payload}_j^{\mathsf{out}}.\mathsf{isDummy} = 0$, set $v^{\mathsf{out}} := v^{\mathsf{out}} + v_j^{\mathsf{out}}$ and check that $\mathbb{v}_j^{\mathsf{out}} = \mathbb{v}^\star$.
   (g) check that the value of asset $\mathsf{id}^\star$ is conserved: $v^{\mathsf{in}} = v^{\mathsf{out}}$.

Most of the lines above are self-explanatory, but for the line that derives a fresh unique identifier in the "mint" case (Step 2c), which deserves an explanation. Informally, the idea is to derive the identifier from the (globally unique) serial numbers of records consumed in the minting transaction. In more detail, we set the identifier to be a commitment to the serial numbers of consumed input records. To see why this works, first note that the commitment scheme's binding property guarantees that opening a commitment to two different messages is computationally difficult. Next, note that in our case these messages are the input records' serial numbers, and hence *are* different. Together, these facts imply that the identifier is globally unique (and hence non-repeating). A benefit of this method is that the commitment scheme's hiding property further guarantees that the identifier reveals no information about the underlying serial numbers, which in turn guarantees that the identifier hides all information about the initial minting transaction (given that $r$ is random).

## 6.2 Decentralized exchanges

We describe how to use death predicates that enforce custom-access policies to build *privacy-preserving decentralized exchanges*. These allow users to exchange custom assets with strong privacy guarantees without requiring users to give up custody of these assets. We proceed by first providing background on centralized and decentralized exchanges. Then, we formulate desirable privacy properties for decentralized exchanges. Finally, we describe constructions that achieve these properties.

**Motivation.** Exchanging digital assets is a compelling use case of ledger-based systems. A straightforward method to exchange digital assets is via a *centralized exchange*: users entrust the exchange with custody of their assets via an on-chain transaction, and the exchange can then credit or debit assets to users' accounts according to off-chain trades without any on-chain activity; users can "exit" by requesting to withdraw assets, which generates another on-chain transaction that transfers those assets from the exchange to the user. Examples of such exchanges include Coinbase [coi] and Binance [bin]. This architecture provides centralized exchanges with two attractive properties: (a) efficiency, namely, all trades occur in the exchange's off-chain database, resulting in low latency and high throughput for all users; and (b) privacy, namely, only the exchange knows the details of individual trades, and only asset deposits and withdrawals require on-chain activity; this activity can further be concealed by using private (Zerocash-style) transactions to realize deposits/withdrawals. However, this centralized architecture has a serious drawback: having given up custody of their assets, users are exposed to the risk of security breaches, fraud, or front-running at the exchange. These risks are not hypothetical: users have lost funds deposited at centralized exchanges [PA14, De18, Zha18, Cim18].

In light of the above, *decentralized exchanges* (DEXs) have been proposed as an alternative method for exchanging digital assets that enable users to retain custody of their assets. However, existing DEX constructions have poor efficiency and privacy guarantees. Below we describe how we can provide strong privacy for DEXs. (We leave improving the efficiency of DEXs to future work.)

**DEX architectures.** A DEX is a ledger-based application that enables users to trade digital assets without giving up custody of these assets to a third party. There are different DEX architectures with different trade-offs; see [Pro18] for a survey. In the following, we consider DEX architectures where the exchange has no state or maintains its state off-chain.[10] There are two main categories of such DEXs:

- *Intent-based DEX.* The DEX maintains an index, which is a table where makers publish their intention to trade (say, a particular asset pair) without committing any assets. A taker interested in a maker's intention to trade can directly communicate with the maker to agree on terms. They can jointly produce a transaction for the trade, to be broadcast for on-chain processing. An example of such a DEX is AirSwap [air].

[10] This is in contrast to DEX architectures that involve, say, a smart contract that stores on-chain the standing orders of all users.

An attractive feature of intent-based DEXs is that they reduce exposure to front-running because the information required for front-running (like prices or identities of the involved parties) has been finalized by the time the transaction representing the trade is broadcast for processing. Note that the aforementioned lack of information also makes it difficult for the market to discover appropriate exchange rates because listings in the index cannot directly be linked with completed transactions.

- *Order-based DEX.* The DEX maintains an order book, which is a table where makers can publish orders by committing the funds for those orders up front. A taker can then interact with the order book to fill orders. In an *open-book DEX*, the taker manually picks an order from the order book, while in a *closed-book DEX*, the taker is matched off-chain with a maker's offer by the order book operator. An example of an open-book DEX is Radar Relay [rad], and an example of a closed-book DEX is Paradex [par].

  Note that order books (which are typically public) give more information about market activity than indexes, and hence enable better price discovery. However, existing constructions of order-based DEXs also allow other parties to link a standing order with a transaction that fills the order before the transaction is finalized, enabling them to front-run the order. Which parties can front-run depends on the kind of order-based DEX: in the open-book variant, anyone can front-run, while in the closed-book variant, only the order book operator can front-run (as it is the sole entity that can invoke the trade smart-contract).

The architectures described above offer different trade-offs with respect to market price discovery and front-running exposure, and hence can be useful in different scenarios.

**Privacy shortcomings and goals.** While the foregoing DEX architectures offer attractive security and functionality, they do *not* provide strong privacy guarantees, as we now explain. First, each transaction reveals information about the corresponding trade, such as the assets and amounts that were exchanged. Prior work [BDJT17, BBD$^+$17, EMC19, DGK$^+$20] shows that such leakage enables front-running that harms user experience and market transparency, and proposes mitigations that, while potentially useful, do not provide strong privacy guarantees. Even if one manages to hide these trade details, transactions in existing DEXs *also* reveal the identities of transacting parties. Onlookers can use this information to extract trading patterns and frequencies of users. This reduces the privacy of users, violates the fungibility of assets, and increases exposure to front-running, because onlookers can use the aforementioned patterns to infer when particular assets are being traded.

These shortcomings motivate the following privacy goals for DEXs. Throughout, we assume that an order is defined by a pair of assets (that are to be exchanged), and their exchange rates.

1. *Trade confidentiality:* No efficient adversary $\mathcal{A}$ should be able to learn the trade details of completed or cancelled trades. That is, a transaction that completes or cancels a trade should not reveal to $\mathcal{A}$ the asset pairs or amounts involved in the trade.
2. *Trade anonymity:* No efficient adversary $\mathcal{A}$ should be able to learn the identities of parties involved in a trade. That is, a transaction that completes or cancels a trade should not reveal to $\mathcal{A}$ any information about the maker or taker of the trade.

A protocol that achieves trade confidentiality and trade anonymity against an adversary $\mathcal{A}$ is secure against front-running by $\mathcal{A}$. The flip-side of this is that $\mathcal{A}$ cannot easily discover the rates used in successful trades, leading to poorer visibility into the trading market. We now describe constructions of intent-based and order-based DEXs that achieve trade confidentiality and anonymity.[11]

[11]Throughout, we assume that users interact with index or-order book operators via anonymous channels. (If this is not the case, operators can use network information to link users across different interactions regardless of any cryptographic solutions used.).

**Record format.** Recall from Section 6.1 that records representing units of an asset have payloads of the form $(\mathsf{id}, \mathbb{v}, v, c)$, where $\mathsf{id}$ is the asset identifier, $\mathbb{v}$ is the initial asset supply, $v$ is the asset amount, and $c$ is arbitrary auxiliary information. In the following, we make use of records that, in addition to the mint-or-conserve birth predicate MoC, have an *exchange-or-cancel* death predicate EoC described next. Informally, EoC allows a record $\mathbf{r}$ to be consumed either by exchanging it for $v^\star$ units of an asset with birth predicate $\Phi^\star_\mathsf{b}$ and identifier $\mathsf{id}^\star$ ($\mathsf{id}^\star$, $\Phi^\star_\mathsf{b}$ and $v^\star$ are specified in $c$), or by "cancelling" the exchange and instead sending new records with $\mathbf{r}$'s asset identifier to an address $\mathsf{apk}^\star$ (also specified in $c$). The information required for the exchange includes the asset's birth predicate in addition to its identifier, as it enables users to interact with assets that have birth predicate different from MoC (such as the stablecoins in Section 6.3). The predicate is described below.

Exchange-or-cancel predicate $\mathsf{EoC}(k, \mathsf{ldata}; \mathsf{mode})$ ($\mathsf{mode}$ is the private input for the predicate.)

1. Parse $\mathsf{ldata}$ as $\begin{pmatrix} [\mathsf{cm}^{\mathsf{in}}_i]^2_1 & [\mathsf{apk}^{\mathsf{in}}_i]^2_1 & [\mathsf{payload}^{\mathsf{in}}_i]^2_1 & [\Phi^{\mathsf{in}}_{\mathsf{d},i}]^2_1 & [\Phi^{\mathsf{in}}_{\mathsf{b},i}]^2_1 & [\mathsf{sn}^{\mathsf{in}}_i]^2_1 & \mathsf{memo} \\ [\mathsf{cm}^{\mathsf{out}}_j]^2_1 & [\mathsf{apk}^{\mathsf{out}}_j]^2_1 & [\mathsf{payload}^{\mathsf{out}}_j]^2_1 & [\Phi^{\mathsf{out}}_{\mathsf{d},j}]^2_1 & [\Phi^{\mathsf{out}}_{\mathsf{b},j}]^2_1 & & \mathsf{aux} \end{pmatrix}$.
2. Recall that $k \in \{1, 2\}$ is the index of the current input record. Let $l \in \{1, 2\}$ denote the index of the other input record. (If $k = 1$ then set $l := 2$; if instead $k = 2$ then set $l := 1$.)
3. Parse the current input record's payload $\mathsf{payload}^{\mathsf{in}}_k$ as $(\mathsf{id}^{\mathsf{in}}_k, \mathbb{v}_k, v^{\mathsf{in}}_k, c^{\mathsf{in}}_k)$, and the application data $c^{\mathsf{in}}_k$ as $(\Phi^\star_{\mathsf{b},k}, \mathsf{id}^\star_k, v^\star_k, \mathsf{apk}^\star_k)$.
4. Parse the other input record's payload $\mathsf{payload}^{\mathsf{in}}_l$ as $(\mathsf{id}^{\mathsf{in}}_l, \mathbb{v}_l, v^{\mathsf{in}}_l, c^{\mathsf{in}}_l)$, and the application data $c^{\mathsf{in}}_l$ as $(\Phi^\star_{\mathsf{b},l}, \mathsf{id}^\star_l, v^\star_l, \mathsf{apk}^\star_l)$.
5. If $\mathsf{mode} = \mathtt{exch}$, ensure that the assets are correctly exchanged, by checking the following.
   (a) the input records are not dummy: $\mathsf{payload}^{\mathsf{in}}_1.\mathsf{isDummy} = \mathsf{payload}^{\mathsf{in}}_2.\mathsf{isDummy} = 0$.
   (b) the conditions of the trade are satisfied:
      i. the current input record has the expected identifier, birth predicate, and value: $\Phi^{\mathsf{in}}_{\mathsf{b},k} = \Phi^\star_{\mathsf{b},l}$, $\mathsf{id}^{\mathsf{in}}_k = \mathsf{id}^\star_l$, and $v^{\mathsf{in}}_k = v^\star_l$.
      ii. the other input record has the expected identifier, birth predicate, and value: $\Phi^{\mathsf{in}}_{\mathsf{b},l} = \Phi^\star_{\mathsf{b},k}$, $\mathsf{id}^{\mathsf{in}}_l = \mathsf{id}^\star_k$, and $v^{\mathsf{in}}_l = v^\star_k$.
      iii. the output records' birth predicates are correctly swapped: $\Phi^{\mathsf{in}}_{\mathsf{b},1} = \Phi^{\mathsf{out}}_{\mathsf{b},2}$ and $\Phi^{\mathsf{in}}_{\mathsf{b},2} = \Phi^{\mathsf{out}}_{\mathsf{b},1}$.
      iv. the output records have the correct asset identifier, initial supply, and value: $\mathsf{payload}^{\mathsf{out}}_1 = (\mathsf{id}^{\mathsf{in}}_2, \mathbb{v}_2, v^{\mathsf{in}}_2, \bot)$ and $\mathsf{payload}^{\mathsf{out}}_2 = (\mathsf{id}^{\mathsf{in}}_1, \mathbb{v}_1, v^{\mathsf{in}}_1, \bot)$.
      v. the output records are addressed correctly: $\mathsf{apk}^{\mathsf{out}}_k = \mathsf{apk}^\star_k$ and $\mathsf{apk}^{\mathsf{out}}_l = \mathsf{apk}^\star_l$.
6. Else if $\mathsf{mode} = \mathtt{cancel}$, ensure that the trade is cancelled by checking that the $\mathsf{id}_k$-value is transferred to the specified "redemption" address public key $\mathsf{apk}^\star_k$, by checking the following.
   (a) the current input record is non-dummy: $\mathsf{payload}_k.\mathsf{isDummy} = 0$.
   (b) the other input record is dummy: $\mathsf{payload}_l.\mathsf{isDummy} = 1$.
   (c) the output records are custom assets with identifier $\mathsf{id}^{\mathsf{in}}_k$:
      i. the output records have the correct birth predicate: $\Phi^{\mathsf{out}}_{\mathsf{b},1} = \Phi^{\mathsf{out}}_{\mathsf{b},2} = \Phi^{\mathsf{in}}_{\mathsf{b},k}$.
      ii. the output records have the correct asset identifier and initial supply: $\mathsf{payload}^{\mathsf{out}}_1 = (\mathsf{id}^{\mathsf{in}}_k, \mathbb{v}_k, v^{\mathsf{out}}_1, \bot)$ and $\mathsf{payload}^{\mathsf{out}}_2 = (\mathsf{id}^{\mathsf{in}}_k, \mathbb{v}_k, v^{\mathsf{out}}_2, \bot)$.
   (d) the output records preserve $\mathsf{id}^{\mathsf{in}}_k$-value: $v^{\mathsf{out}}_1 + v^{\mathsf{out}}_2 = v^{\mathsf{in}}_k$.
   (e) the address public key of the output records is correct: $\mathsf{apk}^{\mathsf{out}}_1 = \mathsf{apk}^{\mathsf{out}}_2 = \mathsf{apk}^\star_k$.

**The case of intent-based DEXs.** We describe an intent-based DEX that hides all information information about orders and transacting parties.

1. A maker $\mathsf{M}$ can publish to the index an intention to trade, which is a tuple $(\mathsf{id}_A, \mathsf{id}_B, \mathsf{pk}_\mathsf{M})$ to be interpreted as: "I want to buy assets with identifier $\mathsf{id}_B$ in exchange for assets with identifier $\mathsf{id}_A$. Please contact me using the encryption public key $\mathsf{pk}_\mathsf{M}$ if you would like to discuss the terms."
2. A taker $\mathsf{T}$ who is interested in this offer can use $\mathsf{pk}_\mathsf{M}$ to privately communicate with $\mathsf{M}$ and agree on the terms of the trade (the form of communication is irrelevant). If $\mathsf{T}$ and $\mathsf{M}$ do not reach an agreement, then $\mathsf{T}$ can always pursue other entries in the index. So suppose that $\mathsf{T}$ and $\mathsf{M}$ do reach an agreement. For the sake of example, $\mathsf{T}$ will give 10 units of asset $\mathsf{id}_B$ to $\mathsf{M}$ and will receive 5 units of asset $\mathsf{id}_A$ from $\mathsf{M}$.
3. The taker $\mathsf{T}$ creates a new record $\mathbf{r}$ with payload $(\mathsf{id}_B, \mathbb{v}_B, 10, c)$ for auxiliary data $c = (\mathsf{id}_A, 5, \mathsf{apk}_{\mathsf{new}})$, and with death predicate EoC. Then $\mathsf{T}$ sends $\mathbf{r}$ (along with the information necessary to redeem $\mathbf{r}$) to $\mathsf{M}$.
4. If $\mathsf{M}$ possesses a record worth 5 units of asset $\mathsf{id}_A$, he can use $\mathsf{T}$'s message to construct a DPC transaction that completes the exchange by consuming $\mathbf{r}$ and by producing appropriate new records for $\mathsf{M}$ and $\mathsf{T}$. (This

step deviates from existing intent-based DEXs in that it is the *maker* that broadcasts the trade transaction.)

The record $\mathbf{r}$ produced by the taker $\mathsf{T}$ can be redeemed by $\mathsf{M}$ only via an appropriate record in exchange. If $\mathsf{M}$ does not possess such a record, $\mathsf{T}$ can cancel the trade (at any time) and retrieve his funds by satisfying the "cancel" branch of the predicate $\mathsf{EoC}$ (which requires knowing the secret key corresponding to $\mathsf{apk}_{\mathsf{new}}$).

Note that regardless of whether the trade was successful or not, this protocol achieves trade anonymity and trade confidentiality against all parties (including the index operator). Indeed, the only information revealed in the final transaction is that some records were consumed and others created; no information is revealed about $\mathsf{M}$, $\mathsf{T}$, the assets involved in the trade ($\mathsf{id}_A$ and $\mathsf{id}_B$), or the amounts exchanged.

**The case of order-based DEXs.** We describe private order-based DEXs, with open or closed books.

- *Open-book DEX:* The variant below hides all information about $\mathsf{M}$ and $\mathsf{T}$, but reveals the assets and amounts involved. This implies achieving trade anonymity but not trade confidentiality.

  For the sake of example, assume again that the maker $\mathsf{M}$ will trade $5$ units of asset $\mathsf{id}_A$ for $10$ units of asset $\mathsf{id}_B$. $\mathsf{M}$ constructs a record $\mathbf{r}$ with payload $(\mathsf{id}_B, \mathbb{v}_B, v = 10, c = (\mathsf{id}_A, 5, \mathsf{apk}))$ and death predicate $\mathsf{EoC}$. He uses this to construct an order $o = (\mathbf{r}, \mathsf{info})$ consisting of the record and the information necessary to consume it, and publishes $o$ to the order book. An interested taker $\mathsf{T}$ can then construct and publish a transaction $\mathsf{tx}$ that consumes $\mathbf{r}$ and creates new records with the appropriate values and asset identifiers.

  The transaction $\mathsf{tx}$ hides information about the maker $\mathsf{M}$ and taker $\mathsf{T}$, but because it reveals $\mathbf{r}$'s serial number, it can be linked with its originating order $o$. This allows onlookers to learn the assets and amounts of $\mathsf{tx}$. Hence, this protocol achieves trade anonymity, but not trade confidentiality.

- *Closed-book DEX:* The variant below hides all order information from everyone but the order book operator. Hence it achieves trade anonymity and confidentiality against everyone but the order book operator.

  The maker $\mathsf{M}$ creates a record $\mathbf{r}$ as above, and sends the record and its consumption information $\mathsf{info}$ to the order book. The order book does not publish these; it publishes only the terms of the order. Takers can publish orders of their own, and if two orders match then the order book operator constructs a transaction $\mathsf{tx}$ that consumes both records and produces new records, completing the order. At no point does either party surrender custody of their funds, thus preserving the self-custodial nature of the exchange protocol.

  The foregoing achieves trade anonymity and confidentiality against everyone but the order book operator because only the order book operator learns the details of the records consumed by the transaction $\mathsf{tx}$, and $\mathsf{tx}$ itself (which once published anyone can see) does not reveal any information about these records. As a consequence, this protocol also protects against front-running by everyone but the order book.

  Note that in our protocol, the maker acts as the taker's counterparty (and vice versa), while in non-private closed-book DEXs, only the order book operator can act as the counterparty for both the maker and the taker. Our protocol can be modified to support such a flow by straightforward modifications to $\mathsf{EoC}$.

**Operator fees.** In the foregoing we have omitted a discussion of fees due to the operators of DEX infrastructure (such as index or order book operators). Support for such fees can be achieved, in a straightforward way, by the following small modifications to the exchange-or-cancel predicate $\mathsf{EoC}$. First, one would need to increase the number of output records of DPC transactions to $n = 3$; the third record would be used to pay fees to the operator. Second, one would have to decide how these fees are calculated. This can be done, e.g., by hardcoding a fee percentage into the predicate or by allowing users to specify fees that they are willing to pay.

**Remark 6.1** (preventing a denial-of-funds attack)**.** In the foregoing protocols, the maker $\mathsf{M}$ could refuse to provide the taker $\mathsf{T}$ with information about its output record, thus denying $\mathsf{T}$ the ability to consume

its exchanged record. A simple approach to prevent this is to modify the exchange-or-cancel predicate to additionally enforce that the memorandum field of the created transaction contains an encryption of the output record information under a public key specified by $\mathsf{T}$.

### 6.3 Stablecoins and centrally-managed assets

Recently there has been growing interest in custom assets that are managed by a central authority. These include stablecoins, which are assets whose value relative to an another is *fixed* (see [Har18] for an overview). Centrally-managed assets are more compatible with regulations like taxes or blacklists, because the central authority can enforce monetary policies that follow these regulations. Indeed, existing stablecoins like the Gemini dollar [gem] and the Paxos standard [pax] have mechanisms for reversing transactions or freezing funds in response to legal rulings. In this section, we show how to construct *private* centrally-managed assets that support arbitrary, and updatable, policies issued by the central authority; this in particular shows how to create and manage policies for *private stablecoins*. We stress that the ideas described below are compatible with applications that reason about other custom assets. For example, one can use DEXs from Section 6.2 to exchange units of a private stablecoin with units of any other user-defined asset (like one from Section 6.1).

We enforce policies by extending the mint-or-conserve predicate $\mathsf{MoC}$ from Section 6.1 into a *mint-or-enforce* predicate $\mathsf{MoE}_\Pi$ whose "enforce" mode enforces a desired policy $\Pi$. In more detail, say that a central authority $\mathsf{A}$ wishes to issue an asset satisfying policy $\Pi$ (initially). To do so, $\mathsf{A}$ generates a signature public key $\mathsf{pk}_\mathsf{A}$, and then invokes $\mathsf{MoE}_\Pi$ in *mint* mode. In this mode, $\mathsf{MoE}_\Pi$, like $\mathsf{MoC}$, generates the asset identifier $\mathsf{id}$ and creates the initial supply $\mathbb{v}$ of the asset in a single output record whose payload stores the tuple $(\mathsf{id}, \mathbb{v}, \mathbb{v}, \bot)$. Unlike $\mathsf{MoC}$, $\mathsf{MoE}_\Pi$ binds $\mathsf{id}$ not only to the serial numbers of input records (to achieve uniqueness), but also to the public key $\mathsf{pk}_\mathsf{A}$ that authorized $\Pi$. This means that, when receiving payments in such assets, the recipient can immediately deduce the asset's identifier and (authorized) policy.

In a transaction with multiple records, policies are applied and updated by the *enforce* mode of $\mathsf{MoE}_\Pi$. In this mode, $\mathsf{MoE}_\Pi$ ensures that the new record's payload stores $(\mathsf{id}, \mathbb{v}, v, c)$, and that the policy $\Pi$ is satisfied. To update a record $\mathbf{r}$ having policy $\Pi$ to a record $\mathbf{r}'$ having policy $\Pi'$, one can create a transaction that consumes $\mathbf{r}$ and creates $\mathbf{r}'$ such that $\mathbf{r}'$ has birth predicate $\mathsf{MoE}_{\Pi'}$. To ensure that this update is authorized by $\mathsf{A}$, $\mathsf{MoE}_{\Pi'}$ checks that a signature over $\Pi'$ with respect to $\mathsf{pk}_\mathsf{A}$ has been provided, and that $\mathsf{id}$ has been correctly derived from $\mathsf{pk}_\mathsf{A}$. These checks ensure that every record with identifier $\mathsf{id}$ only has authorized policies.

Below we provide pseudocode for $\mathsf{MoE}_\Pi$.

Mint-or-enforce predicate $\mathsf{MoE}_\Pi(k, \mathsf{ldata}; \mathsf{mode})$ ($\mathsf{mode}$ is the private input of the predicate)

1. Parse $\mathsf{ldata}$ as $\begin{pmatrix} [\mathsf{cm}_i^{\mathsf{in}}]_1^2 & [\mathsf{apk}_i^{\mathsf{in}}]_1^2 & [\mathsf{payload}_i^{\mathsf{in}}]_1^2 & [\Phi_{\mathsf{d},i}^{\mathsf{in}}]_1^2 & [\Phi_{\mathsf{b},i}^{\mathsf{in}}]_1^2 & [\mathsf{sn}_i^{\mathsf{in}}]_1^2 & \mathsf{memo} \\ [\mathsf{cm}_j^{\mathsf{out}}]_1^2 & [\mathsf{apk}_j^{\mathsf{out}}]_1^2 & [\mathsf{payload}_j^{\mathsf{out}}]_1^2 & [\Phi_{\mathsf{d},j}^{\mathsf{out}}]_1^2 & [\Phi_{\mathsf{b},j}^{\mathsf{out}}]_1^2 & & \mathsf{aux} \end{pmatrix}$.
2. If $\mathsf{mode} = (\mathtt{mint}, \mathbb{v}, r, \mathsf{pk}_{\mathsf{SIG}}, \sigma_\Pi)$, ensure that the first output record contains the initial supply of the asset:
   (a) the index of the current output record is correct: $k = 1$.
   (b) all other records are dummy: $\mathsf{payload}_1^{\mathsf{in}}.\mathsf{isDummy} = \mathsf{payload}_2^{\mathsf{in}}.\mathsf{isDummy} = \mathsf{payload}_2^{\mathsf{out}}.\mathsf{isDummy} = 1$.
   (c) the asset identifier is derived correctly: $\mathsf{id} = \mathsf{CRH}(\mathsf{pp}_{\mathsf{CRH}}, \mathsf{CM.Commit}(\mathsf{pp}_{\mathsf{CM}}, \mathsf{sn}_1 \| \mathsf{sn}_2; r) \| \mathsf{pk}_{\mathsf{SIG}})$.
   (d) the policy $\Pi$ is authorized by $\mathsf{pk}_{\mathsf{SIG}}$: $\mathsf{SIG.Verify}(\mathsf{pp}_{\mathsf{SIG}}, \mathsf{pk}_{\mathsf{SIG}}, \Pi, \sigma_\Pi) = 1$.
   (e) the current output record's payload is correct: $\mathsf{payload}_1^{\mathsf{out}}.\mathsf{isDummy} = 0$ and $\mathsf{payload}_1^{\mathsf{out}} = (\mathsf{id}, \mathbb{v}, \mathbb{v}, c = \bot)$.
3. If $\mathsf{mode} = (\mathtt{enforce}, \rho, \mathsf{pk}_{\mathsf{SIG}}, \sigma_\Pi)$, check that the policy $\Pi$ is enforced:
   (a) parse the current output record's payload $\mathsf{payload}_k^{\mathsf{out}}$ as $(\mathsf{id}^\star, \mathbb{v}^\star, v^\star, c)$.
   (b) check that $\mathsf{pk}_{\mathsf{SIG}}$ is valid for the asset: $\mathsf{id}^\star = \mathsf{CRH}(\mathsf{pp}_{\mathsf{CRH}}, \rho \| \mathsf{pk}_{\mathsf{SIG}})$.
   (c) check that the policy $\Pi$ is authorized under pk: $\mathsf{SIG.Verify}(\mathsf{pp}_{\mathsf{SIG}}, \mathsf{pk}_{\mathsf{SIG}}, \Pi, \sigma_\Pi) = 1$.
   (d) check that the policy $\Pi$ is satisfied: $\Pi(k, \mathsf{ldata}) = 1$.

By way of example, we now show how a central authority $\mathsf{A}$ can use the mint-or-enforce predicate to construct a stablecoin that enforces a *blacklisting* (in addition to the default value-conservation policy).

Namely, if an address is on a blacklist $B$ of addresses, the address is not allowed to participate in transactions. To do so, $\mathsf{A}$ follows the above procedure to construct and publish a mint-or-enforce predicate $\mathsf{MoE}_{\Pi_B}$ implementing a policy $\Pi_B$ that inspects the address public keys of consumed records, and ensures that none of them are in $B$. Now suppose that later on $\mathsf{A}$ wishes to update $B$ into a new blacklist $B'$ that includes a new address $\mathsf{apk}$. It does so by publishing a corresponding updated predicate $\mathsf{MoE}_{\Pi_{B'}}$ for this new blacklist, and users can use the above update mechanism to move their records from policy $\Pi_B$ to policy $\Pi_{B'}$. Now, any funds stored at the newly-blacklisted address $\mathsf{apk}$ cannot be moved to the new policy.

# 7 Implementation strategy

The straightforward approach to implement our construction of a DPC scheme (described in Section 4) is to instantiate the proof system via a simulation-extractable zkSNARK (e.g., [GM17]) and then select the other cryptographic building blocks so that the circuit (more precisely, constraint system) for deciding the NP relation $\mathcal{R}_{\mathsf{e}}$ has as small a size as possible. While the straightforward approach sounds promising, closer inspection reveals significant costs that we need to somehow reduce. In this section we discuss, in a "problem and solution" format, the challenges that we encountered and how we addressed them. (The implementation strategy for plain DPC schemes directly ports over to delegable DPC schemes so we do not discuss them.)

**Problem 1: universality is expensive.** The NP relation $\mathcal{R}_{\mathsf{e}}$ involves checking arbitrary predicates, which means that one must rely on proof systems for *universal* computations. However, checking universal computations via state-of-the-art zkSNARKs involves expensive tools for universal circuits/machines [BCG$^+$13, BCTV14, WSR$^+$15, BCTV17]. These tools would not only yield an expensive solution but would also penalize users who only produce transactions that attest to simple inexpensive predicates, because these users would have to incur the costs of using these "heavy duty" proof systems.

**Solution 1: recursive proof verification.** We address this problem by relying on one layer of *recursive proof composition* [Val08, BCCT13]. Instead of tasking $\mathcal{R}_{\mathsf{e}}$ with checking satisfiability of general predicates, we only task it with checking *succinct proofs* attesting to this. Checking succinct proofs is a (relatively) inexpensive computation that is universal for *all* predicates, which can be "hardcoded" in $\mathcal{R}_{\mathsf{e}}$. Crucially, since the "outer" succinct proofs produced for $\mathcal{R}_{\mathsf{e}}$ do not reveal information about the "inner" succinct proofs attesting to predicates' satisfiability (thanks to zero knowledge), the inner succinct proofs do *not* have to hide what predicate was checked, removing the need for expensive universal circuits; in fact, inner proofs do not even have to be zero knowledge. Rather, these inner succinct proofs can be for NP relations *tailored* to the computations needed by particular birth and death predicates. Furthermore, this approach ensures that a user only has to incur the cost of proving satisfiability of the specific predicates involved in his transactions, regardless of the complexity of predicates used by other users in their transactions.

In more detail, taking the case of one input and one output record as an example, we modify $\mathsf{DPC.Execute}$ to additionally take as input SNARK proofs $\pi_{\mathsf{d}}$ and $\pi_{\mathsf{b}}$, and also modify the NP relation $\mathcal{R}_{\mathsf{e}}$ so that, instead of directly checking that $\Phi_{\mathsf{d}}$ and $\Phi_{\mathsf{b}}$ are satisfied, it instead checks that $\pi_{\mathsf{d}}$ *and* $\pi_{\mathsf{b}}$ *attest to the satisfaction of* $\Phi_{\mathsf{d}}$ *and* $\Phi_{\mathsf{b}}$. That is, $\mathcal{R}_{\mathsf{e}}$ checks that $\mathsf{NIZK.Verify}(\mathsf{pp}_{\Phi_{\mathsf{d}}}, \mathbb{x}_{\mathsf{e}}, \pi_{\mathsf{d}}) = 1$ and $\mathsf{NIZK.Verify}(\mathsf{pp}_{\Phi_{\mathsf{b}}}, \mathbb{x}_{\mathsf{e}}, \pi_{\mathsf{b}}) = 1$, where $\mathsf{pp}_{\Phi_{\mathsf{d}}}$ are public parameters for the NP relation $\mathcal{R}_{\Phi_{\mathsf{d}}} := \{(\mathbb{x}_{\mathsf{e}}, \mathbb{w}_{\mathsf{e}}) \text{ s.t. } \Phi_{\mathsf{d}}(\mathbb{x}_{\mathsf{e}}, \mathbb{w}_{\mathsf{e}}) = 1\}$ and similarly for $\Phi_{\mathsf{b}}$. The public parameters $\mathsf{pp}_{\Phi_{\mathsf{d}}}$ and $\mathsf{pp}_{\Phi_{\mathsf{b}}}$ are stored in the record, in place of (a description of) the predicates.[12]

More generally, we modify $\mathsf{DPC.Execute}$ to additionally take as input SNARK proofs $[\pi_{\mathsf{d},i}]_1^m$ attesting that the old records' death predicates are satisfied and SNARK proofs $[\pi_{\mathsf{b},j}]_1^n$ attesting that the new records' birth predicates are satisfied. Moreover, we similarly modify the NP relation $\mathcal{R}_{\mathsf{e}}$ to check that these proofs are valid, instead of directly checking that the relevant predicates are satisfied.

In sum, $\mathcal{R}_{\mathsf{e}}$ is not tasked with checking general predicates. Instead, it merely has to check SNARK proofs, a fixed computation of size $O_\lambda(m+n)$. Separately, a user wishing to prove that a predicate $\Phi$ is satisfied will invoke a SNARK on an NP statement of size $|\Phi|$ (tailored for $\Phi$).[13] The approach described so far, however,

[12] More precisely, to verify a proof for a predicate $\Phi$, the proof verifier does not need to read all of $\mathsf{pp}_\Phi$, which has size $O_\lambda(|\Phi|)$ in some zkSNARKs (i.e., it is large). Rather, the proof verifier only needs to read $O_\lambda(|\mathbb{x}_{\mathsf{e}}|)$ bits of $\mathsf{pp}_\Phi$, which are collectively known as the *verification key*. The record would then store this verification key (or a hash thereof) rather than $\mathsf{pp}_\Phi$.

[13] An additional benefit of each predicate $\Phi$ having its own public parameters $\mathsf{pp}_\Phi$ is flexible trust: users are not obliged to trust parameters used in each others' transactions and, moreover, if some parameters are known to be compromised, predicates can safely refuse to interact with records associated with them. We view this *isolation mechanism* as a novel and valuable feature in practice.

hides additional costs that we need to overcome.

**Problem 2: recursion is expensive.** Recursive proof composition has so far been empirically demonstrated for pairing-based SNARKs [BCTV17], whose proofs are extremely short and cheap to verify. We thus focus our attention on these, and explain the efficiency challenges that we must overcome in our setting.

Recall that pairings are instantiated via elliptic curves of small embedding degree. If we instantiate a SNARK's pairing via an elliptic curve $E$ defined over a prime field $\mathbb{F}_q$ and having a subgroup of large prime order $r$, then (a) the SNARK supports NP relations $\mathcal{R}$ expressed as arithmetic circuits over $\mathbb{F}_r$, while (b) proof verification involves arithmetic operations over $\mathbb{F}_q$. This means that we need to express $\mathcal{R}_{\mathsf{e}}$ via arithmetic circuits over $\mathbb{F}_r$. In turn, since the SNARK verifier is part of $\mathcal{R}_{\mathsf{e}}$, this means that we need to also express the verifier via an arithmetic circuit over $\mathbb{F}_r$, which is problematic because the verifier's "native" operations are over $\mathbb{F}_q$. Simulating $\mathbb{F}_q$ operations via $\mathbb{F}_r$ operations introduces significant overheads, and picking $E$ such that $q = r$, in order to avoid simulation, is impossible [BCTV17].

Prior work thus suggests using *multiple* curves [BCTV17], such as a two-cycle of pairing-friendly elliptic curves, that is, two prime-order curves $E_1$ and $E_2$ such that the prime size of one's base field is the prime order of the other's group, and orchestrating SNARKs based on these so that fields always "match up". Unfortunately, known curves with these properties are inefficient at 128 bits of security [BCTV17, CCW19].

**Solution 2: tailored set of curves.** In our setting we merely need "a proof of a proof", with the latter proof not itself depending on further proofs.

This implies that we do not actually need a cycle of pairing-friendly elliptic curves (which enables recursion of arbitrary depth), but rather only a "two-chain" of two curves $E_1$ and $E_2$ such that the size of the base field of $E_1$ is the size of the prime order subgroup of $E_2$. We can use the Cocks–Pinch method [FST10] to set up such a bounded recursion [BCTV17]. We now elaborate on this.

First we pick a pairing-friendly elliptic curve $E_1$ that not only is suitable for 128 bits of security according to standard considerations (involving, e.g., its embedding degree and the ratio of the sizes of its base field and prime order group) but, moreover, is compatible with efficient SNARK provers in *both* levels of the recursion. Namely, letting $p$ be the prime order of the base field and $r$ the prime order of the group, we need that *both* $\mathbb{F}_r$ and $\mathbb{F}_p$ have multiplicative subgroups whose orders are large powers of 2. The condition on $\mathbb{F}_r$ ensures efficient proving for SNARKs over $E_1$, while the condition on $\mathbb{F}_p$ ensures efficient proving for SNARKs that verify proofs over $E_1$. In light of the above, we set $E_1$ to be $E_{\mathsf{BLS}}$, a curve from the Barreto–Lynn–Scott (BLS) family [BLS02, CLN11] with embedding degree 12. This family not only enables parameters that conservatively achieve 128 bits of security, but also enjoys properties that facilitate very efficient implementation [AFK+12]. We ensure that both $\mathbb{F}_r$ and $\mathbb{F}_p$ have multiplicative subgroups of order $2^\alpha$ for $\alpha \geq 40$, by choosing the parameter $x$ of the BLS family to satisfy $x \equiv 1 \bmod 3 \cdot 2^\alpha$; indeed, for such a choice of $x$ both $r(x) = x^4 - x^2 + 1$ and $p(x) = (x-1)^2 r(x)/3 + x$ are divisible by $2^\alpha$. This also ensures that $x \equiv 1 \bmod 3$, which ensures that there are efficient towering options for the relevant fields [Cos12].

Next we use the Cocks–Pinch method to pick a pairing-friendly elliptic curve $E_2 = E_{\mathsf{CP}}$ over a field $\mathbb{F}_q$ such that the curve group $E_{\mathsf{CP}}(\mathbb{F}_q)$ contains a subgroup of prime order $p$ (the size of $E_{\mathsf{BLS}}$'s base field). Since the method outputs a prime $q$ that has about $2\times$ more bits than the desired $p$, and in turn $p$ has about $1.5\times$ more bits than $r$ (due to properties of the BLS family), we only need $E_{\mathsf{CP}}$ to have embedding degree 6 in order to achieve 128 bits of security (as determined from the guidelines in [FST10]).

In sum, proofs of predicates' satisfiability are produced via a SNARK over $E_{\mathsf{BLS}}$, and proofs for the NP relation $\mathcal{R}_{\mathsf{e}}$ are produced via a zkSNARK over $E_{\mathsf{CP}}$. The matching fields between the two curves ensure that the former proofs can be efficiently verified.

**Problem 3: Cocks–Pinch curves are costly.** While the curve $E_{\mathsf{CP}}$ was chosen to facilitate efficient checking of proofs over $E_{\mathsf{BLS}}$, the curve $E_{\mathsf{CP}}$ is at least $2\times$ more expensive (in time and space) than $E_{\mathsf{BLS}}$ simply

because $E_{\mathsf{CP}}$'s base field has about twice as many bits as $E_{\mathsf{BLS}}$'s base field. Checks in the NP relation $\mathcal{R}_{\mathsf{e}}$ that are not directly related to proof checking are now unnecessarily carried over a less efficient curve.

**Solution 3: split relations across two curves.** We split $\mathcal{R}_{\mathsf{e}}$ into two NP relations $\mathcal{R}_{\mathsf{BLS}}$ and $\mathcal{R}_{\mathsf{CP}}$ (see Fig. 14), with the latter containing just the proof check and the former containing all other checks. We can then use a zkSNARK over the curve $E_{\mathsf{BLS}}$ (an efficient curve) to produce proofs for $\mathcal{R}_{\mathsf{BLS}}$, and a zkSNARK over $E_{\mathsf{CP}}$ (the less efficient curve) to produce proofs for $\mathcal{R}_{\mathsf{CP}}$. This approach significantly reduces the running time of DPC.Execute (producing proofs for the checks in $\mathcal{R}_{\mathsf{BLS}}$ is more efficient over $E_{\mathsf{BLS}}$ than over $E_{\mathsf{CP}}$), at the expense of a modest increase in transaction size (a transaction now includes a zkSNARK proof over $E_{\mathsf{BLS}}$ in addition to a proof over $E_{\mathsf{CP}}$). An important technicality that must be addressed is that the foregoing split relies on certain secret information to be shared across the NP relations, namely, the identities of relevant predicates and the local data. We can store this information in suitable commitments that are part of the NP instances for the two NP relations (doing this efficiently requires some care as we discuss below).

**Problem 4: the NP relations have many checks.** Even using $E_{\mathsf{CP}}$ only for SNARK verification and $E_{\mathsf{BLS}}$ for all other checks does not suffice: the NP relations $\mathcal{R}_{\mathsf{BLS}}$ and $\mathcal{R}_{\mathsf{CP}}$ still have to perform expensive checks like verifying Merkle tree authentication paths and commitment openings, and evaluating pseudorandom functions and collision resistant functions. Similar NP relations, like the one in Zerocash [BCG$^+$14], require upwards of *four million gates* to express such checks, resulting in high latencies for producing transactions (several minutes) and large public parameters for the system (hundreds of megabytes).

**Solution 4: efficient EC primitives.** Commitments and collision-resistant hashing can be expressed as very efficient arithmetic circuits if one opts for Pedersen-type constructions over suitable Edwards elliptic curves (and techniques derived from these ideas are now part of deployed systems [HBHW20]). To do this, we pick two Edwards curves, $E_{\mathsf{Ed/BLS}}$ over the field $\mathbb{F}_r$ (matching the group order of $E_{\mathsf{BLS}}$) and $E_{\mathsf{Ed/CP}}$ over the field $\mathbb{F}_p$ (matching the group order of $E_{\mathsf{CP}}$). This enables us to achieve very efficient circuits for primitives used in our NP relations, including commitments, collision-resistant hashing, and randomizable signatures. (Note that $E_{\mathsf{Ed/BLS}}$ and $E_{\mathsf{Ed/CP}}$ do not need to be pairing-friendly as the primitives only rely on their group structure.)

**Problem 5: sharing information between NP relations is costly.** We have said that splitting $\mathcal{R}_{\mathsf{e}}$ into two NP relations $\mathcal{R}_{\mathsf{BLS}}$ and $\mathcal{R}_{\mathsf{CP}}$ relies on sharing secret information via commitments across NP statements; namely, a commitment $\mathsf{cm}_{\Phi}$ to the identities of predicates and a commitment $\mathsf{cm}_{\mathsf{ldata}}$ to the local data. But if both relations open these commitments, we cannot make an efficient use of Pedersen commitments because the two NP relations are over different fields: $\mathcal{R}_{\mathsf{BLS}}$ is over $\mathbb{F}_r$, while $\mathcal{R}_{\mathsf{CP}}$ is over $\mathbb{F}_p$. For example, if we used a Pedersen commitment over the order-$r$ subgroup of the Edwards curve $E_{\mathsf{Ed/BLS}}$, then: (a) opening a commitment in $\mathcal{R}_{\mathsf{BLS}}$ would be cheap, but (b) opening a commitment in $\mathcal{R}_{\mathsf{CP}}$ would involve expensive simulation of $\mathbb{F}_r$-arithmetic via $\mathbb{F}_p$-arithmetic. (And similarly if we used a Pedersen commitment over the order-$p$ subgroup of the Edwards curve $E_{\mathsf{Ed/CP}}$.) To make matters worse, the predicate identities and the local data are large, so an inefficient solution for committing to these would add significant costs to $\mathcal{R}_{\mathsf{BLS}}$ and $\mathcal{R}_{\mathsf{CP}}$.

**Solution 5: hash predicate verification keys and commit to local data.** In a record, instead of storing predicate verification keys, we store collision-resistant hashes of these. This reduces the cost of producing the commitment $\mathsf{cm}_{\Phi}$ in $\mathcal{R}_{\mathsf{BLS}}$ and $\mathcal{R}_{\mathsf{CP}}$, as $\mathsf{cm}_{\Phi}$ contains hashes that are much smaller than verification keys. We realize $\mathsf{cm}_{\Phi}$ via Blake2s, a boolean primitive of modest cost in $\mathbb{F}_r$ and $\mathbb{F}_p$. Crucially, *only* $\mathcal{R}_{\mathsf{CP}}$ needs to access the verification keys themselves, so we can efficiently use a Pedersen hash over the Edwards curve $E_{\mathsf{Ed/CP}}$ to let $\mathcal{R}_{\mathsf{CP}}$ check the keys (supplied as non-deterministic advice) against the hashes inside $\mathsf{cm}_{\Phi}$.

We realize the local data commitment $\mathsf{cm}_{\mathsf{ldata}}$ via a Pedersen commitment over $E_{\mathsf{Ed/BLS}}$, and assume that predicates take $\mathsf{cm}_{\mathsf{ldata}}$ as input rather than local data in the clear. Since both $\mathcal{R}_{\mathsf{BLS}}$ and the predicate relations are defined over the field $\mathbb{F}_r$ (the prime-order subgroup of the curve $E_{\mathsf{BLS}}$), non-deterministically opening

$\mathsf{cm}_{\mathsf{ldata}}$ is efficient in both relations. This approach significantly reduces costs because $\mathcal{R}_{\mathsf{CP}}$ no longer needs to reason about the contents of $\mathsf{cm}_{\mathsf{ldata}}$, and can simply pass $\mathsf{cm}_{\mathsf{ldata}}$ as input to the SNARK verifier.

The NP relation $\mathcal{R}_{\mathsf{BLS}}$ has instances $\mathbb{x}_{\mathsf{BLS}}$ and witnesses $\mathbb{w}_{\mathsf{BLS}}$ of the form

$$\mathbb{x}_{\mathsf{BLS}} = \begin{pmatrix} \text{ledger digest} & \mathsf{st}_{\mathbf{L}} \\ \text{old record serial numbers} & [\mathsf{sn}_i]_1^m \\ \text{new record commitments} & [\mathsf{cm}_j]_1^n \\ \text{predicate commitment} & \mathsf{cm}_{\Phi} \\ \text{local data commitment} & \mathsf{cm}_{\mathsf{ldata}} \\ \text{transaction memorandum} & \mathsf{memo} \end{pmatrix} \quad \text{and} \quad \mathbb{w}_{\mathsf{BLS}} = \begin{pmatrix} \text{old records} & [\mathbf{r}_i]_1^m \\ \text{old record membership witnesses} & [\mathbb{w}_{\mathbf{L},i}]_1^m \\ \text{old address secret keys} & [\mathsf{ask}_i]_1^m \\ \text{new records} & [\mathbf{r}_j]_1^n \\ \text{predicate comm. randomness} & r_{\Phi} \\ \text{local data randomness} & r_{\mathsf{ldata}} \\ \text{auxiliary predicate input} & \mathsf{aux} \end{pmatrix}$$

where

- for each $i \in \{1, \ldots, m\}$, $\mathbf{r}_i = (\mathsf{apk}_i, \mathsf{payload}_i, h_{\mathsf{b},i}, h_{\mathsf{d},i}, \rho_i, r_i, \mathsf{cm}_i)$;
- for each $j \in \{1, \ldots, n\}$, $\mathbf{r}_j = (\mathsf{apk}_j, \mathsf{payload}_j, h_{\mathsf{b},j}, h_{\mathsf{d},j}, \rho_j, r_j, \mathsf{cm}_j)$.

Define the local data $\mathsf{ldata} := \begin{pmatrix} [\mathsf{cm}_i]_1^m & [\mathsf{apk}_i]_1^m & [\mathsf{payload}_i]_1^m & [h_{\mathsf{d},i}]_1^m & [h_{\mathsf{b},i}]_1^m & [\mathsf{sn}_i]_1^m & \mathsf{memo} \\ [\mathsf{cm}_j]_1^n & [\mathsf{apk}_j]_1^n & [\mathsf{payload}_j]_1^n & [h_{\mathsf{d},j}]_1^n & [h_{\mathsf{b},j}]_1^n & \mathsf{aux} & \end{pmatrix}$.

A witness $\mathbb{w}_{\mathsf{BLS}}$ is valid for an instance $\mathbb{x}_{\mathsf{BLS}}$ if the following conditions hold:

1. For each $i \in \{i, \ldots, m\}$:
   - If $\mathbf{r}_i$ is not dummy, $\mathbb{w}_{\mathbf{L},i}$ proves that the commitment $\mathsf{cm}_i$ is in a ledger with digest $\mathsf{st}_{\mathbf{L}}$: $\mathbf{L}.\mathsf{Verify}(\mathsf{st}_{\mathbf{L}}, \mathsf{cm}_i, \mathbb{w}_{\mathbf{L},i}) = 1$.
   - The address public key $\mathsf{apk}_i$ and secret key $\mathsf{ask}_i$ form a valid key pair:
     $\mathsf{apk}_i = \mathsf{CM.Commit}(\mathsf{pp}_{\mathsf{CM}}, \mathsf{sk}_{\mathsf{PRF},i};\ r_{\mathsf{pk},i})$ and $\mathsf{ask}_i = (\mathsf{sk}_{\mathsf{PRF},i}, r_{\mathsf{pk},i})$.
   - The serial number $\mathsf{sn}_i$ is valid: $\mathsf{sn}_i = \mathsf{PRF}_{\mathsf{sk}_{\mathsf{PRF},i}}(\rho_i)$.
   - The old record commitment $\mathsf{cm}_i$ is valid: $\mathsf{cm}_i = \mathsf{TCM.Commit}(\mathsf{pp}_{\mathsf{TCM}}, \mathsf{apk}_i \| \mathsf{payload}_i \| h_{\mathsf{b},i} \| h_{\mathsf{d},i} \| \rho_i;\ r_i)$.
2. For each $j \in \{1, \ldots, n\}$:
   - The serial number nonce $\rho_j$ is computed correctly: $\rho_j = \mathsf{CRH.Eval}(\mathsf{pp}_{\mathsf{CRH}}, j \| \mathsf{sn}_1 \| \ldots \| \mathsf{sn}_m)$.
   - The new record commitment $\mathsf{cm}_j$ is valid: $\mathsf{cm}_j = \mathsf{TCM.Commit}(\mathsf{pp}_{\mathsf{TCM}}, \mathsf{apk}_j \| \mathsf{payload}_j \| h_{\mathsf{b},j} \| h_{\mathsf{d},j} \| \rho_j;\ r_j)$.
3. The predicate commitment $\mathsf{cm}_{\Phi}$ is valid: $\mathsf{cm}_{\Phi} = \mathsf{b2s}([h_{\mathsf{d},i}]_1^m \parallel [h_{\mathsf{b},j}]_1^n \parallel r_{\Phi})$.
4. The local data commitment $\mathsf{cm}_{\mathsf{ldata}}$ is valid: $\mathsf{cm}_{\mathsf{ldata}} = \mathsf{CM.Commit}(\mathsf{pp}_{\mathsf{CM}}, \mathsf{ldata}; r_{\mathsf{ldata}})$

---

The NP relation $\mathcal{R}_{\mathsf{CP}}$ has instances $\mathbb{x}_{\mathsf{CP}}$ and witnesses $\mathbb{w}_{\mathsf{CP}}$ of the form

$$\mathbb{x}_{\mathsf{CP}} = \begin{pmatrix} \text{predicate commitment} & \mathsf{cm}_{\Phi} \\ \text{local data commitment} & \mathsf{cm}_{\mathsf{ldata}} \end{pmatrix} \quad \text{and} \quad \mathbb{w}_{\mathsf{CP}} = \begin{pmatrix} \text{old death pred. ver. keys} & [\mathsf{vk}_{\mathsf{d},i}]_1^m \\ \text{old death pred. proofs} & [\pi_{\mathsf{d},i}]_1^m \\ \text{new birth pred. ver. keys} & [\mathsf{vk}_{\mathsf{b},j}]_1^n \\ \text{new birth pred. proofs} & [\pi_{\mathsf{b},j}]_1^n \\ \text{predicate comm. randomness} & r_{\Phi} \end{pmatrix}$$

A witness $\mathbb{w}_{\mathsf{CP}}$ is valid for an instance $\mathbb{x}_{\mathsf{CP}}$ if the following conditions hold:

1. For each $i \in \{i, \ldots, m\}$:
   - The death predicate hash $h_{\mathsf{d},i}$ is computed correctly: $h_{\mathsf{d},i} = \mathsf{CRH.Eval}(\mathsf{pp}_{\mathsf{CRH}}, \mathsf{vk}_{\mathsf{d},i})$.
   - The death predicate proof $\pi_{\mathsf{d},i}$ is valid: $\mathsf{NIZK.Verify}(\mathsf{vk}_{\mathsf{d},i}, i \| \mathsf{cm}_{\mathsf{ldata}}, \pi_{\mathsf{d},i})$.
2. For each $j \in \{1, \ldots, n\}$:
   - The birth predicate hash $h_{\mathsf{b},j}$ is computed correctly: $h_{\mathsf{b},j} = \mathsf{CRH.Eval}(\mathsf{pp}_{\mathsf{CRH}}, \mathsf{vk}_{\mathsf{b},j})$.
   - The birth predicate proof $\pi_{\mathsf{b},j}$ is valid: $\mathsf{NIZK.Verify}(\mathsf{vk}_{\mathsf{b},j}, j \| \mathsf{cm}_{\mathsf{ldata}}, \pi_{\mathsf{b},j})$.
3. The predicate commitment $\mathsf{cm}_{\Phi}$ is valid: $\mathsf{cm}_{\Phi} = \mathsf{b2s}([h_{\mathsf{d},i}]_1^m \parallel [h_{\mathsf{b},j}]_1^n \parallel r_{\Phi})$.

**Figure 14:** Splitting the NP relation $\mathcal{R}_{\mathsf{e}}$ into two NP relations $\mathcal{R}_{\mathsf{BLS}}$ and $\mathcal{R}_{\mathsf{CP}}$, over $\mathbb{F}_r$ and $\mathbb{F}_p$ respectively.

# 8 System implementation

We implemented our "plain" DPC scheme (Section 4) and our delegable DPC scheme (Section 5), by following the strategy described in Section 7. The resulting system, named **Zexe** (*Zero knowledge EXEcution*), consists of several Rust libraries: (a) a library for finite field and elliptic curve arithmetic, adapted from [Bow17b]; (b) a library for cryptographic building blocks, including zkSNARKs for constraint systems (using components from [Bow17a]); (c) a library with constraints for many of these building blocks; and (d) a library that realizes our constructions of plain and delegable DPC. Our code base, like our construction, is written in terms of abstract building blocks, which allows to easily switch between different instantiations of the building blocks. In the rest of this section we describe the efficient instantiations used in the experiments reported in Section 9.

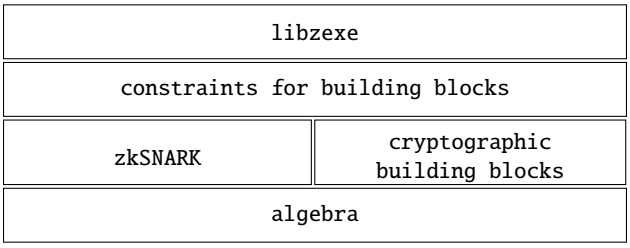


**Figure 15:** Stack of libraries comprising Zexe.

**Ledger.** The ledger $\mathbf{L}$ in our prototype is simply an ideal ledger, i.e., an append-only log of valid transactions that is stored in memory. Of course, in a real-world deployment, this ideal ledger would be replaced by a distributed protocol that realizes (a suitable approximation of) an ideal ledger. Recall from Section 3.1 that we require the ledger $\mathbf{L}$ to provide a method to efficiently prove and verify membership of a transaction, or one of its subcomponents, in $\mathbf{L}$. For this, we maintain a Merkle tree [Mer87] atop the list of transactions, using the collision-resistant hash function $\mathsf{CRH}$ described below. This results in the following algorithms for $\mathbf{L}$.

- $\mathbf{L}.\mathsf{Push}(\mathsf{tx})$: Append $\mathsf{tx}$ to the transaction list and update the Merkle tree.
- $\mathbf{L}.\mathsf{Digest} \rightarrow \mathsf{st}_{\mathbf{L}}$: Return the root of the Merkle tree.
- $\mathbf{L}.\mathsf{Prove}(\mathsf{tx}) \rightarrow \mathbb{w}_{\mathbf{L}}$: Return the authentication path for $\mathsf{tx}$ in the Merkle tree.
- $\mathbf{L}.\mathsf{Verify}(\mathsf{st}_{\mathbf{L}}, \mathsf{tx}, \mathbb{w}_{\mathbf{L}}) \rightarrow b$: Check that $\mathbb{w}_{\mathbf{L}}$ is a valid authentication path for $\mathsf{tx}$ in a tree with root $\mathsf{st}_{\mathbf{L}}$.

Our prototype maintains the Merkle tree in memory, but a real-world deployment would have to maintain it via a distributed protocol. (Such data structures atop distributed ledgers are used in existing systems [ZCa15].)

**Pseudorandom function.** Fixing key length and input length at 256 bits, we instantiate $\mathsf{PRF}$ using the Blake2s hash function [ANWW13]: $\mathsf{PRF}_k(x) := \mathsf{b2s}(k \| x)$ for $k, x \in \{0,1\}^{256}$.

**Elliptic curves.** Our implementation strategy (see Section 7) involves several elliptic curves: two pairing-friendly curves $E_{\mathsf{BLS}}$ and $E_{\mathsf{CP}}$, and two "plain" curves $E_{\mathsf{Ed/BLS}}$ and $E_{\mathsf{Ed/CP}}$ whose base field respectively matches the prime-order subgroup of $E_{\mathsf{BLS}}$ and $E_{\mathsf{CP}}$. Details about these curves are in Figure 16; the parameter used to generate the BLS curve $E_{\mathsf{BLS}}$ is $x = 3 \cdot 2^{46} \cdot (7 \cdot 13 \cdot 499) + 1$ (see Section 7 for why).

**NIZKs.** We instantiate the NIZKs used for the NP relation $\mathcal{R}_{\mathsf{e}}$ via zero-knowledge *succinct* non-interactive arguments of knowledge (zk-SNARKs), which makes our DPC schemes succinct (see Remark 4.1). Concretely, we rely on the simulation-extractable zkSNARK of Groth and Maller [GM17], used over the pairing-friendly elliptic curves $E_{\mathsf{BLS}}$ (for proving predicates' satisfiability) and $E_{\mathsf{CP}}$ (for proving validity of these latter proofs).

**DLP-hard group.** Several instantiations of cryptographic primitives introduced below rely on the hardness of extracting discrete logarithms in a prime order group. We generate these groups via a group generator $\mathsf{SampleGrp}$, which on input a security parameter $\lambda$ (represented in unary), outputs a tuple $(\mathbb{G}, q, g)$ that

| name | curve type | embedding degree | size of prime-order subgroup | size of base field | size of compressed group elements (rounded to multiples of 8 bytes) $\mathbb{G}_1$ | $\mathbb{G}_2$ |
|---|---|---|---|---|---|---|
| $E_{\mathsf{Ed/BLS}}$ | twisted Edwards | — | $s$ | $r$ | 32 | — |
| $E_{\mathsf{BLS}}$ | BLS | 12 | $r$ | $p$ | 48 | 96 |
| $E_{\mathsf{Ed/CP}}$ | twisted Edwards | — | $t$ | $p$ | 48 | — |
| $E_{\mathsf{CP}}$ | short Weierstrass | 6 | $p$ | $q$ | 104 | 312 |

| prime | value | size in bits | 2-adicity |
|---|---|---|---|
| $s$ | `0x4aad957a68b2955982d1347970dec005293a3afc43c8afeb95aee9ac33fd9ff` | 251 | 1 |
| $r$ | `0x12ab655e9a2ca55660b44d1e5c37b00159aa76fed00000010a11800000000001` | 253 | 47 |
| $t$ | `0x35c748c2f8a21d58c760b80d94292763445b3e601ea271e1d75fe7d6eeb84234066d10f5d893814103486497d95295` | 374 | 2 |
| $p$ | `0x1ae3a4617c510eac63b05c06ca1493b1a22d9f300f5138f1ef3622fba094800170b5d44300000008508c00000000001` | 377 | 46 |
| $q$ | `0x3848c4d2263babf8941fe959283d8f526663bc5d176b746af0266a7223ee72023d07830c728d80f9d78bab3596c8617c579252a3fb77c79c13201ad533049cfe6a399c2f764a12c4024bee135c065f4d26b7545d85c16dfd424adace79b57b942ae9` | 782 | 3 |

**Figure 16:** The elliptic curves $E_{\mathsf{BLS}}, E_{\mathsf{CP}}, E_{\mathsf{Ed/BLS}}, E_{\mathsf{Ed/CP}}$.

describes a group $\mathbb{G}$ of prime order $q$ generated by $g$. The discrete-log problem is hard in $\mathbb{G}$. In our prototype we fix $\mathbb{G}$ to be the largest prime-order subgroup of either $E_{\mathsf{Ed/BLS}}$ or $E_{\mathsf{Ed/CP}}$, depending on the context.

**Commitments.** We instantiate (plain and) trapdoor commitments via Pedersen commitments over $\mathbb{G}$, as defined in Figure 17; note that the setup algorithm takes as additional input the message length $n$. Pedersen commitments are perfectly hiding, and are computationally binding if the discrete-log problem is hard in $\mathbb{G}$.

**Collision-resistant hashing.** We instantiate $\mathsf{CRH}$ via a Pedersen hash function over $\mathbb{G}$, as specified in Figure 18; note that the setup algorithm takes as additional input the message length $n$. Collision resistance follows from hardness of the discrete-logarithm problem [MRK03].

**Remark 8.1.** Hopwood et al. [HBHW20] note that projecting a twisted Edwards curve point $(x, y)$ to its $x$-coordinate is injective when the point is in the curve's largest prime-order subgroup. Our implementation uses this fact to reduce the output size of $\mathsf{TCM}$ and $\mathsf{CRH}$ by projecting their output to its $x$-coordinate.

$\mathsf{TCM.Setup}(1^\lambda, n) \to \mathsf{pp}_{\mathsf{TCM}}$:
1. Sample a group: $(\mathbb{G}, q, g) \leftarrow \mathsf{SampleGrp}(1^\lambda)$.
2. For $i \in \{1, \ldots, n\}$, sample generator $h_i$: $r_i \leftarrow \mathbb{Z}_q; h_i := g^{r_i}$.
3. Output $\mathsf{pp}_{\mathsf{TCM}} := (\mathbb{G}, q, g, [h_i]_1^n)$.

$\mathsf{TCM.Commit}(\mathsf{pp}_{\mathsf{TCM}}, m \in \{0,1\}^n; r_{\mathsf{cm}}) \to \mathsf{cm}$:
1. Parse $\mathsf{pp}_{\mathsf{TCM}}$ as $(\mathbb{G}, q, g, [h_i]_1^n)$.
2. Output $\mathsf{cm} := g^{r_{\mathsf{cm}}} \prod_{i=1}^n h_i^{m_i}$.

**Figure 17:** Pedersen commitment scheme.

$\mathsf{CRH.Setup}(1^\lambda, n) \to \mathsf{pp}_{\mathsf{CRH}}$:
1. Sample a group: $(\mathbb{G}, q, g_1) \leftarrow \mathsf{SampleGrp}(1^\lambda)$.
2. For $i \in \{2, \ldots, n\}$, sample generator $g_i$: $r_i \leftarrow \mathbb{Z}_q; g_i := g^{r_i}$.
3. Output $\mathsf{pp}_{\mathsf{CRH}} := (\mathbb{G}, q, [g_i]_1^n)$.

$\mathsf{CRH.Eval}(\mathsf{pp}_{\mathsf{CRH}}, m \in \{0,1\}^n) \to h$:
1. Parse $\mathsf{pp}_{\mathsf{CRH}}$ as $(\mathbb{G}, q, [g_i]_1^n)$.
2. Output $h := \prod_{i=1}^n g_i^{m_i}$.

**Figure 18:** Pedersen collision-resistant hash.

# 9 System evaluation

In Section 9.1 we evaluate individual cryptographic building blocks. In Section 9.2 we evaluate the cost of NP relations expressed as constraints, as required by the underlying zkSNARK. In Section 9.3 we evaluate the running time of DPC algorithms. In Section 9.4 we evaluate the sizes of DPC data structures. All reported measurements were taken on a machine with an Intel Xeon 6136 CPU at 3.0 GHz with 252 GB of RAM.

## 9.1 Cryptographic building blocks

We are interested in two types of costs associated with a given cryptographic building block: the *native execution cost*, which are the running times of certain algorithms on a CPU; and the *constraint cost*, which are the numbers of constraints required to express certain invariants, to be used by the underlying zkSNARK.

**Native execution cost.** The zkSNARK dominates native execution cost, and the costs of all other building blocks are negligible in comparison. Therefore we separately report only the running times of the zkSNARK, which in our case is a protocol due to Groth and Maller [GM17], abbreviated as GM17. When instantiated over the elliptic curve $E_{\mathsf{BLS}}$, the GM17 prover takes $25\,\mu\mathrm{s}$ per constraint (with 12 threads), while the GM17 verifier takes $250n\,\mu\mathrm{s} + 9.5\,\mathrm{ms}$ on an input with $n$ field elements (with 1 thread). When instantiated over the elliptic curve $E_{\mathsf{CP}}$, the respective prover and verifier costs are $147\,\mu\mathrm{s}$ per constraint and $1.6n\,\mathrm{ms} + 34\,\mathrm{ms}$.

**Constraint cost.** There are three building blocks that together account for the majority of the cost of NP statements that we use. These are: (a) the Blake2s PRF, which requires 21792 constraints to map a 64-byte input to a 32-byte output; (b) the Pedersen collision-resistant hash, which requires $5n$ constraints for an input of $n$ bits; and (c) the GM17 verifier, which requires $14n + 52626$ constraints for an $n$-bit input.

## 9.2 The execute NP relation

In many zkSNARK constructions, including the one that we use, one must express all the relevant checks in the given NP relation as (rank-1) *quadratic constraints* over a certain large prime field. The goal is to minimize the number of such constraints because the prover's costs grow (quasi)linearly in this number.

In our DPC scheme we use a zkSNARK for the NP relation $\mathcal{R}_{\mathsf{e}}$ in Fig. 9 and, similarly, in our delegable DPC scheme we use it for the NP relation $\mathcal{R}_{\mathsf{e}}^{\mathsf{del}}$ in Fig. 23. More precisely, for efficiency reasons explained in Section 7, we split $\mathcal{R}_{\mathsf{e}}$ into the two NP relations $\mathcal{R}_{\mathsf{BLS}}$ and $\mathcal{R}_{\mathsf{CP}}$ in Fig. 14, which we prove via zkSNARKs over the pairing-friendly curves $E_{\mathsf{BLS}}$ and $E_{\mathsf{CP}}$, respectively. (We also similarly split $\mathcal{R}_{\mathsf{e}}^{\mathsf{del}}$.)

Table 3 reports the number of constraints that that we use to express $\mathcal{R}_{\mathsf{BLS}}$, as a function of the number of input ($m$) and output ($n$) records, and additionally reports its primary contributors. Table 4 does the same for $\mathcal{R}_{\mathsf{CP}}$. These tables show that for each input record costs are dominated by verification of a Merkle tree path and the verification of a (death predicate) proof; while for each output record costs are dominated by the verification of a (birth predicate) proof. We also report the cumulative number of constraints when setting $m := 2$ and $n := 2$ because this is a representative instantiation of $m$ and $n$ that enables useful applications.

## 9.3 DPC algorithms

In Table 1 we report the running times of algorithms in our plain DPC and delegable DPC implementations for two input and two output records (i.e., $m := 2$ and $n := 2$). Note that for Execute and Verify, we have excluded costs of ledger operations (such as retrieving an authentication path or scanning for duplicate serial numbers) because these depend on how a ledger is realized, which is orthogonal to our work. Also, we assume that Execute receives as inputs the application-specific SNARK proofs checked by the NP relation. Producing

each of these proofs requires invoking the GM17 prover, over the elliptic curve $E_{\mathsf{BLS}}$, for the relevant birth or death predicate; we describe the cost of doing so for representative applications in Section 9.5.

Observe that the overhead incurred by delegable DPC over plain DPC is negligible, and that, as expected, Setup and Execute are the most costly algorithms, as they invoke costly zkSNARK setup and proving algorithms. To mitigate these costs, Setup and Execute are executed on 12 threads; everything else is executed with 1 thread. Overall, we learn that Execute takes less than a minute, and Verify takes tens of milliseconds. Furthermore, both Setup and Execute consume less than 5 GB of RAM. These costs are comparable with those of similar systems such as Zerocash [BCG+14] and Hawk [KMS+16].

## 9.4 DPC data structures

**Addresses.** An address public key in a DPC scheme is a point on the elliptic curve $E_{\mathsf{Ed/BLS}}$, which is 32 bytes when compressed (see Fig. 16); the corresponding secret key is 64 bytes and consists of a PRF seed (32 bytes) and commitment randomness (32 bytes). In a delegable DPC scheme, address public keys do not change, but address secret keys are 96 bytes, because they additionally contain the 32-byte secret key of a randomizable signature scheme over the elliptic curve $E_{\mathsf{Ed/BLS}}$ (see Fig. 10).

**Transactions.** A transaction in a DPC scheme, with two input and two output records, is 968 bytes. It contains two zkSNARK proofs: $\pi_{\mathsf{BLS}}$, over the elliptic curve $E_{\mathsf{BLS}}$, and $\pi_{\mathsf{CP}}$, over the curve $E_{\mathsf{CP}}$. Each proof consists of two $\mathbb{G}_1$ and one $\mathbb{G}_2$ elements from its respective curve, amounting to 192 bytes for $\pi_{\mathsf{BLS}}$ and 520 for $\pi_{\mathsf{CP}}$ (both in compressed form). In general, for $m$ input records and $n$ output records, transactions are $32m + 32n + 840$ bytes. In a delegable DPC scheme, a transaction additionally contains a 64-byte signature for each input record. See Table 2 for a detailed break down of all of these costs.

**Record contents.** We set a record's payload to be 32 bytes long; if a predicate needs longer data then it can set the payload to be the hash of this data, and use non-determinism to access the data. The foregoing choice means that all contents of a record add up to 224 bytes, since a record consists of an address public key (32 bytes), the 32-byte payload, hashes of birth and death predicates (48 bytes each), a serial number nonce (32 bytes), and commitment randomness (32 bytes).

| | Plain DPC | Delegable DPC |
|---|---|---|
| Setup | 109.62 s | 109.3 s |
| GenAddress | 380 µs | 780 µs |
| Execute | 52.5 s | 53.4 s |
| Verify | 46 ms | 47 ms |

**Table 1:** Cost of DPC algorithms for 2 inputs and 2 outputs.

| | Plain DPC | Delegable DPC |
|---|---|---|
| 2 inputs and 2 outputs | 968 | 1096 |
| $m$ inputs and $n$ outputs | $32m + 32n + 840$ | $96m + 32n + 840$ |
| Per input record: | | |
| Serial number | 32 | 32 |
| Signature | — | 64 |
| Per output record: | | |
| Commitment | 32 | 32 |
| Memorandum | 32 | 32 |
| zkSNARK proof over $E_{\mathsf{CP}}$ | 520 | 520 |
| zkSNARK proof over $E_{\mathsf{BLS}}$ | 192 | 192 |
| Predicate commitment | 32 | 32 |
| Local data commitment | 32 | 32 |
| Ledger digest | 32 | 32 |

**Table 2:** Size of a DPC transaction (in bytes).

| | | Plain DPC | Delegable DPC |
|---|---|---|---|
| **Total with 2 inputs and 2 outputs** | | **387412** | **414339** |
| Below we provide a breakdown of the number of constraints with $m$ input and $n$ output records. | | | |
| Per input record | Total | 117699 | 125401 |
| | Enforce validity of: | | |
| | Merkle tree path | 81824 | 81824 |
| | Address key pair | 3822 | 8435 |
| | Serial number computation | 22301 | 25390 |
| | Record commitment | 9752 | 9752 |
| Per output record | Total | 15427 | 19523 |
| | Enforce validity of: | | |
| | Serial number nonce | 5417 | 9513 |
| | Record commitment | 10010 | 10010 |
| Other: | Enforce validity of: | | |
| | Predicate commitment | $21792 \cdot \lceil \frac{3}{4}(m+n) + \frac{1}{2} \rceil$ | $21792 \cdot \lceil \frac{3}{4}(m+n) + \frac{1}{2} \rceil$ |
| | Local data commitment | $7168 \cdot m + 6144 \cdot n$ | $8192 \cdot m + 6144 \cdot n$ |
| | Miscellaneous | 7368 | 8651 |

**Table 3:** Number of constraints for $\mathcal{R}_{\mathsf{BLS}}$.

| | | Plain DPC | Delegable DPC |
|---|---|---|---|
| **Total with 2 inputs and 2 outputs** | | **439224** | **439476** |
| Below we provide a breakdown of the number of constraints with $m$ input and $n$ output records. | | | |
| Per input record | Total | 87569 | 87569 |
| | Enforce validity of: | | |
| | Death predicate ver. key | 45827 | 45827 |
| | Death predicate proof | 41742 | 41742 |
| Per output record | Total | 87569 | 87569 |
| | Enforce validity of: | | |
| | Birth predicate ver. key | 45827 | 45827 |
| | Birth predicate proof | 41742 | 41742 |
| Other | Enforce validity of: | | |
| | Predicate commitment | $21792 \cdot \lceil \frac{3}{4}(m+n) + \frac{1}{2} \rceil$ | $21792 \cdot \lceil \frac{3}{4}(m+n) + \frac{1}{2} \rceil$ |
| | Miscellaneous | 1780 | 2032 |

**Table 4:** Number of constraints for $\mathcal{R}_{\mathsf{CP}}$.

### 9.5 Applications

We do not report total costs for producing transactions for the applications in Section 6 because the additional application-specific costs are *negligible* compared to the base cost reported in Table 1. This is because all application-specific proofs are produced over the efficient elliptic curve $E_{\mathsf{BLS}}$, and moreover, for each application we consider, the heaviest computation checked by these proofs is the relatively lightweight one of opening the local data commitment; the remaining costs consist of a few cheap range and equality checks. Indeed, with two input and two output records, these applications require fewer than $35,000$ constraints (compared to over $350,000$ for $\mathcal{R}_{\mathsf{BLS}}$ and $\mathcal{R}_{\mathsf{CP}}$), and producing the corresponding proofs takes tens of milliseconds (compared to tens of seconds for the base cost of $\mathsf{DPC.Execute}$).

# A Proof of security for our DPC scheme

We prove that our DPC construction (see Section 4) satisfies the security definition in Section 3.3. To do this, for every real-world (efficient) adversary $\mathcal{A}$, we construct an ideal-world (efficient) simulator $\mathcal{S}$ such that the ideal-world and real-world executions are computationally indistinguishable with respect to any (efficient) environment $\mathcal{E}$. We proceed in three parts: in Appendix A.1 we describe building blocks used to construct the simulator $\mathcal{S}$; in Appendix A.2 we describe the simulator $\mathcal{S}$; in Appendix A.3 we argue that the ideal-world and the real-world executions are computationally indistinguishable.

## A.1 Building blocks for the simulator

We describe various algorithms that are used as sub-routines in the simulator $\mathcal{S}$.

**Trapdoor commitments.** Recall from Section 4.1 that a *trapdoor* commitment scheme is a commitment scheme with auxiliary algorithms $(\mathsf{SimSetup}, \mathsf{Equivocate})$ that enable one to open a commitment $\mathsf{cm}$ to any chosen message. Below we restrict $\mathsf{cm}$ to be a commitment to the empty string $\varepsilon$ because this is sufficient for the proof of security of our DPC scheme.

- *Trapdoor setup:* on input a security parameter, $\mathsf{TCM.SimSetup}$ samples public parameters $\mathsf{pp}_{\mathsf{TCM}}$ and a trapdoor $\mathsf{td}_{\mathsf{TCM}}$ such that $\mathsf{pp}_{\mathsf{TCM}}$ is indistinguishable from public parameters sampled by $\mathsf{TCM.Setup}$.
- *Equivocation:* on input public parameters $\mathsf{pp}_{\mathsf{TCM}}$, trapdoor $\mathsf{td}_{\mathsf{TCM}}$, commitment $\mathsf{cm}$ to $\varepsilon$, corresponding commitment randomness $r_{\mathsf{cm}}$ (so that $\mathsf{TCM.Commit}(\mathsf{pp}_{\mathsf{TCM}}, \varepsilon; r_{\mathsf{cm}}) = \mathsf{cm}$), and target message $m'$, $\mathsf{TCM.Equivocate}$ outputs commitment randomness $r'_{\mathsf{cm}}$ such that $\mathsf{TCM.Commit}(\mathsf{pp}_{\mathsf{TCM}}, m'; r'_{\mathsf{cm}}) = \mathsf{cm}$. Moreover, if $r_{\mathsf{cm}}$ is uniformly random then $r'_{\mathsf{cm}}$ is statistically close to uniformly random.

In Figure 19 we instantiate these algorithms for the Pedersen commitment scheme. Note that the real and simulated public parameters are identical; moreover, the trapdoor randomness $r'_{\mathsf{cm}}$ is the real randomness $r_{\mathsf{cm}}$ shifted by uniformly random field elements, and is hence statistically close to $r_{\mathsf{cm}}$.

$\mathsf{TCM.SimSetup}(1^\lambda, n) \to (\mathsf{pp}_{\mathsf{TCM}}, \mathsf{td}_{\mathsf{TCM}})$
1. Sample a group: $(\mathbb{G}, q, g) \leftarrow \mathsf{SampleGrp}(1^\lambda)$.
2. For $i \in \{1, \ldots, n\}$:
   sample $r_i$ uniformly from $\mathbb{Z}_q$, and set $h_i := g^{r_i}$.
3. Output $(\mathsf{pp}_{\mathsf{TCM}} := (\mathbb{G}, q, g, [h_i]_1^n), \mathsf{td}_{\mathsf{TCM}} := [r_i]_1^n)$.

$\mathsf{TCM.Equivocate}(\mathsf{pp}_{\mathsf{TCM}}, \mathsf{td}_{\mathsf{TCM}}, \mathsf{cm}, r_{\mathsf{cm}}, m' \in \{0,1\}^n) \to r'_{\mathsf{cm}}$
1. Parse $\mathsf{pp}_{\mathsf{TCM}}$ as $(\mathbb{G}, q, g, [h_i]_1^n)$.
2. Parse $\mathsf{td}_{\mathsf{TCM}}$ as $[r_i]_1^n$.
3. Output $r'_{\mathsf{cm}} := r_{\mathsf{cm}} - \sum_{i=1}^n r_i m'_i \bmod q$.

**Figure 19:** Simulated setup and equivocation algorithms for the Pedersen commitment scheme.

**NIZKs.** The scheme $\mathsf{NIZK} = (\mathsf{Setup}, \mathsf{Prove}, \mathsf{Verify})$ is a *simulation-extractable* non-interactive zero knowledge argument. Formally stating the properties of this scheme involves several auxiliary algorithms.

– *Trapdoor setup:* on input a security parameter and a description of an NP relation $\mathcal{R}$, $\mathsf{NIZK.SimSetup}$ outputs a set of public parameters $\mathsf{pp}_{\mathsf{NIZK}}$ and a trapdoor $\mathsf{td}_{\mathsf{NIZK}}$.
– *Simulation:* on input public parameters $\mathsf{pp}_{\mathsf{NIZK}}$, trapdoor $\mathsf{td}_{\mathsf{NIZK}}$, NP instance $\mathbb{x}$, and (optionally) auxiliary information $\mathsf{aux}$, $\mathsf{NIZK.Simulate}$ outputs a simulated proof $\pi$.
– *Extraction:* on input public parameters $\mathsf{pp}_{\mathsf{NIZK}}$, trapdoor $\mathsf{td}_{\mathsf{NIZK}}$, NP instance $\mathbb{x}$, and proof $\pi$, $\mathsf{NIZK.Extract}$ outputs a witness $\mathbb{w}$ such that $(\mathbb{x}, \mathbb{w}) \in \mathcal{R}$ (allegedly).

We can now state the properties satisfied by $\mathsf{NIZK}$.

- *Completeness:* for every NP relation $\mathcal{R}$ and instance-witness pair $(\mathbb{x}, \mathbb{w}) \in \mathcal{R}$,

$$\Pr\left[\, \mathsf{NIZK.Verify}(\mathsf{pp}_{\mathsf{NIZK}}, \mathbb{x}, \pi) = 1 \;\middle|\; \begin{array}{r} \mathsf{pp}_{\mathsf{NIZK}} \leftarrow \mathsf{NIZK.Setup}(1^\lambda, \mathcal{R}) \\ (\mathbb{x}, \pi) \leftarrow \mathsf{NIZK.Prove}(\mathsf{pp}_{\mathsf{NIZK}}, \mathbb{x}, \mathbb{w}) \end{array} \right] = 1 \ .$$

- *Perfect zero knowledge:* for every relation $\mathcal{R}$ and efficient adversary $\mathcal{A}$,

$$\Pr\left[\begin{array}{l}\mathsf{pp}_{\mathsf{NIZK}} \leftarrow \mathsf{NIZK.Setup}(1^\lambda, \mathcal{R}) \\ \mathcal{A}^{S_1(\cdot,\cdot)}(\mathsf{pp}_{\mathsf{NIZK}}, \mathsf{aux}) = 1\end{array}\right] = \Pr\left[\begin{array}{l}(\mathsf{pp}_{\mathsf{NIZK}}, \mathsf{td}_{\mathsf{NIZK}}) \leftarrow \mathsf{NIZK.SimSetup}(1^\lambda, \mathcal{R}) \\ \mathcal{A}^{S_2(\cdot,\cdot)}(\mathsf{pp}_{\mathsf{NIZK}}, \mathsf{aux}) = 1\end{array}\right]$$

  where the two oracles are defined as follows
  - $S_1(\mathbb{x}, \mathbb{w}) :=$ "if $(\mathbb{x}, \mathbb{w}) \in \mathcal{R}$ then $\mathsf{NIZK.Prove}(\mathsf{pp}_{\mathsf{NIZK}}, \mathbb{x}, \mathbb{w})$, else abort";
  - $S_2(\mathbb{x}, \mathbb{w}) :=$ "if $(\mathbb{x}, \mathbb{w}) \in \mathcal{R}$ then $\mathsf{NIZK.Simulate}(\mathsf{pp}_{\mathsf{NIZK}}, \mathsf{td}_{\mathsf{NIZK}}, \mathbb{x})$, else abort".
- *Simulation extractability:* for every relation $\mathcal{R}$ and efficient adversary $\mathcal{A}$,

$$\Pr\left[\begin{array}{l}(\mathbb{x}, \pi) \notin Q \\ (\mathbb{x}, \mathbb{w}) \notin \mathcal{R} \\ \mathsf{NIZK.Verify}(\mathsf{pp}_{\mathsf{NIZK}}, \mathbb{x}, \pi) = 1\end{array}\middle|\begin{array}{r}(\mathsf{pp}_{\mathsf{NIZK}}, \mathsf{td}_{\mathsf{NIZK}}) \leftarrow \mathsf{NIZK.SimSetup}(1^\lambda, \mathcal{R}) \\ (\mathbb{x}, \pi) \leftarrow \mathcal{A}^{S(\cdot)}(\mathsf{pp}_{\mathsf{NIZK}}) \\ \mathbb{w} \leftarrow \mathsf{NIZK.Extract}(\mathsf{pp}_{\mathsf{NIZK}}, \mathsf{td}_{\mathsf{NIZK}}, \mathbb{x}, \pi)\end{array}\right] = \mathrm{negl}(\lambda) \quad ,$$

  where $S(\mathbb{x}) := \mathsf{NIZK.Simulate}(\mathsf{pp}_{\mathsf{NIZK}}, \mathsf{td}_{\mathsf{NIZK}}, \mathbb{x})$ and $Q$ is the set of query-answer pairs between the adversary $\mathcal{A}$ and the simulated-proof oracle $S$.

## A.2 The ideal-world simulator

The ideal-word simulator $\mathcal{S}$ will interact with the ideal functionality $\mathcal{F}_{\text{DPC}}$ and with the environment $\mathcal{E}$. Note that for UC security it suffices to show security against a dummy real-world adversary $\mathcal{A}$ that simply forwards all instructions from the environment $\mathcal{E}$ [Can01]. Since our security definition is a special case of UC security, we inherit this simplification, and thus only consider such an adversary $\mathcal{A}$. The pseudocode for $\mathcal{S}$ is provided below; auxiliary subroutines are provided in Figure 20.

**Setup.**

1. Initialize an empty table $\mathcal{S}.\mathsf{Records}$ that maps record commitments to their contents.
2. Initialize an empty table $\mathcal{S}.\mathsf{AddrPk}$ that maps address public keys to their secret keys.
3. Initialize an empty transaction ledger $\mathbf{L}$.
4. Sample simulated public parameters and trapdoor: $(\mathsf{pp}, \mathsf{td}) \leftarrow \mathsf{DPC.SimSetup}(1^\lambda)$. (See Fig. 20.)
5. Define

$$\begin{aligned}\mathsf{SampleAddrPk}(\cdot) &:= \mathsf{CM.Commit}(\mathsf{pp}_{\mathsf{CM}}, \varepsilon; \cdot) \quad , \\ \mathsf{SampleCm}(\cdot) &:= \mathsf{TCM.Commit}(\mathsf{pp}_{\mathsf{TCM}}, \varepsilon; \cdot) \quad , \\ \mathsf{SampleSn}(\cdot) &:= \text{ "sample uniformly random string of correct length" } .\end{aligned}$$

6. Start ideal-world execution with the above $(\mathsf{SampleAddrPk}, \mathsf{SampleCm}, \mathsf{SampleSn})$.

At this point, the simulator will receive messages notifying it of transactions and of messages sharing contents of newly-created records. The simulator handles each case separately.

**Transaction notifications.**

- **From environment.** When $\mathcal{E}$ instructs a corrupted party to invoke $\mathbf{L}.\mathsf{Push}(\mathsf{tx})$:
  1. If $\mathsf{DPC.Verify}^{\mathbf{L}}(\mathsf{pp}, \mathsf{tx}) \neq 1$, abort.
  2. Parse the real-world transaction $\mathsf{tx}$ as $([\mathsf{sn}_i]_1^m, [\mathsf{cm}_j]_1^n, \mathsf{memo}, \star)$.
  3. Compute $([\mathbf{r}_i]_1^m, [\mathsf{ask}_i]_1^m, [\mathbf{r}_j]_1^n, \mathsf{aux}) \leftarrow \mathsf{DPC.ExtractExecute}(\mathsf{pp}, \mathsf{td}, \mathsf{tx})$. [See Figure 20.]
  4. For every $i \in \{1, \ldots, m\}$:

(a) Parse the real-world record $\mathbf{r}_i$ as $(\mathsf{apk}_i, \mathsf{payload}_i, \Phi_{\mathsf{b},i}, \Phi_{\mathsf{d},i}, \rho_i, r_i, \mathsf{cm}_i)$.
(b) Parse the address secret key $\mathsf{ask}_i$ as $(\mathsf{sk}_{\mathsf{PRF},i}, r_{\mathsf{pk},i})$.
(c) If $\mathcal{S}.\mathsf{Records}[\mathsf{cm}_i] \neq \mathbf{r}_i$, abort. (***Note:*** *Captures binding property of the commitment.*)
(d) If $\mathbf{L}.\mathsf{Contains}(\mathsf{cm}_i) = 0$, abort. (***Note:*** *Captures existence of record.*)
(e) Create the ideal-world record $\mathbb{r}_i := (\mathsf{cm}_i, \mathsf{apk}_i, \mathsf{payload}_i, \Phi_{\mathsf{b},i}, \Phi_{\mathsf{d},i})$.
(f) If $\mathcal{S}.\mathsf{AddrPk}[\mathsf{apk}_i] = \perp$:
i. Invoke $\mathcal{F}_{\text{DPC}}.\mathsf{GenAddress}(\mathsf{apk}_i)$.
ii. Insert $\mathsf{apk}_i$ into $\mathcal{S}.\mathsf{AddrPk}$: $\mathcal{S}.\mathsf{AddrPk}[\mathsf{apk}_i] := \mathsf{ask}_i$.
(g) Else, if $\mathcal{S}.\mathsf{AddrPk}[\mathsf{apk}_i] \neq \mathsf{ask}_i$, abort. (***Note:*** *Captures uniqueness of secret key.*)
5. For every $j \in \{1, \ldots, n\}$:
(a) Parse the real-world record $\mathbf{r}_j$ as $(\mathsf{apk}_j, \mathsf{payload}_j, \Phi_{\mathsf{b},j}, \Phi_{\mathsf{d},j}, \rho_j, r_j, \mathsf{cm}_j)$.
(b) If the serial number nonce $\rho_j$ was seen in a prior extracted transaction, or if $\rho_j = \rho_k$ for $k \neq j$, abort. (***Note:*** *Captures uniqueness of nonce.*)
(c) Set $\mathcal{S}.\mathsf{Records}[\mathsf{cm}_j] := \mathbf{r}_j$.
6. Construct instance for $\mathcal{R}_{\mathsf{e}}$: $\mathbb{x}_{\mathsf{e}} := (\mathsf{st}_{\mathbf{L}}, [\mathsf{sn}_i]_1^m, [\mathsf{cm}_j]_1^n, \mathsf{memo})$.
7. Construct witness for $\mathcal{R}_{\mathsf{e}}$: $\mathbb{w}_{\mathsf{e}} := ([\mathbf{r}_i]_1^m, [\mathbb{w}_{\mathbf{L},i}]_1^m, [\mathsf{ask}_i]_1^m, [\mathbf{r}_j]_1^n, \mathsf{aux})$.
8. If $(\mathbb{x}_{\mathsf{e}}, \mathbb{w}_{\mathsf{e}}) \notin \mathcal{R}_{\mathsf{e}}$, abort.
9. Invoke $\mathcal{F}_{\text{DPC}}.\mathsf{Execute}([\mathbb{r}_i]_1^m, [\mathsf{meta}_i]_1^m, [\mathsf{sn}_i]_1^m, [\mathsf{cm}_j]_1^n, [\mathsf{apk}_j]_1^n, [\mathsf{payload}_j]_1^n, [\Phi_{\mathsf{b},j}]_1^n, [\Phi_{\mathsf{d},j}]_1^n, \mathsf{aux}, \mathsf{memo})$.
10. **Receive from $\mathcal{F}_{\text{DPC}}$:** $[\mathbb{r}_j]_1^n$.
11. **Receive from $\mathcal{F}_{\text{DPC}}$:** $(\texttt{Execute}, [\mathsf{sn}_i]_1^m, [\mathsf{cm}_j]_1^n, \mathsf{memo})$.
12. Append the real-world transaction $\mathsf{tx}$ to the ledger $\mathbf{L}$.

- **From ideal functionality.** When $\mathcal{F}_{\text{DPC}}$ broadcasts $(\texttt{Execute}, [\mathsf{sn}_i]_1^m, [\mathsf{cm}_j]_1^n, \mathsf{memo})$:
1. Compute $([\mathbf{r}_j]_1^n, \mathsf{tx}) \leftarrow \mathsf{DPC.SimExecute}^{\mathbf{L}}(\mathsf{pp}, \mathsf{td}, [\mathsf{sn}_i]_1^m, [\mathsf{cm}_j]_1^n, \mathsf{memo})$. (See Fig. 20.)
2. For each $j \in \{1, \ldots, n\}$, set $\mathcal{S}.\mathsf{Records}[\mathsf{cm}_j] := \mathbf{r}_j$.
3. Append the real-world transaction $\mathsf{tx}$ to the ledger $\mathbf{L}$.

**Record authorization notification.**

- **From environment.** When $\mathcal{E}$ instructs a corrupted party to send $(\texttt{RecordAuth}, \mathbf{r}, \mathcal{P})$ to $\mathcal{P}$:
1. Parse the real-world record $\mathbf{r}$ as $(\mathsf{apk}, \mathsf{payload}, \Phi_{\mathsf{b}}, \Phi_{\mathsf{d}}, \rho, r, \mathsf{cm})$.
2. Invoke $\mathcal{F}_{\text{DPC}}.\mathsf{ShareRecord}(\mathbb{r}, \mathcal{P})$ with $\mathbb{r} := (\mathsf{cm}, \mathsf{apk}, \mathsf{payload}, \Phi_{\mathsf{b}}, \Phi_{\mathsf{d}})$.
- **From ideal functionality.** When $\mathcal{F}_{\text{DPC}}$ sends $(\texttt{RecordAuth}, \mathbb{r}, r)$:
1. Parse the ideal record $\mathbb{r}$ as $(\mathsf{cm}, \mathsf{apk}, \mathsf{payload}, \Phi_{\mathsf{b}}, \Phi_{\mathsf{d}})$.
2. Retrieve the real-world record $\mathbf{r} = \mathcal{S}.\mathsf{Records}[\mathsf{cm}]$, and set the serial number nonce $\rho := \mathbf{r}.\rho$.
3. Define new record commitment message $m := (\mathsf{apk} \| \mathsf{payload} \| \Phi_{\mathsf{b}} \| \Phi_{\mathsf{d}} \| \rho)$.
4. Compute new commitment randomness $r' \leftarrow \mathsf{TCM.Equivocate}(\mathsf{pp}_{\mathsf{TCM}}, \mathsf{td}_{\mathsf{TCM}}, \mathsf{cm}, r, m)$.
5. Construct the new real-world record $\mathbf{r}' := (\mathsf{apk}, \mathsf{payload}, \Phi_{\mathsf{b}}, \Phi_{\mathsf{d}}, \rho, r', \mathsf{cm})$.
6. Set $\mathcal{S}.\mathsf{Records}[\mathsf{cm}] := \mathbf{r}'$.
7. **Send to $\mathcal{A}$:** $(\texttt{RecordAuth}, \mathbf{r}')$ .

DPC.SimSetup

*Input:* security parameter $1^\lambda$

*Output:* simulated public parameters $\mathsf{pp}$ and trapdoor $\mathsf{td}$

1. Sample parameters for commitment: $\mathsf{pp}_{\mathsf{CM}} \leftarrow \mathsf{CM.Setup}(1^\lambda)$.
2. Sample simulated parameters for trapdoor commitment: $(\mathsf{pp}_{\mathsf{TCM}}, \mathsf{td}_{\mathsf{TCM}}) \leftarrow \mathsf{TCM.SimSetup}(1^\lambda)$.
3. Sample parameters for CRH: $\mathsf{pp}_{\mathsf{CRH}} \leftarrow \mathsf{CRH.Setup}(1^\lambda)$.
4. Sample simulated parameters for NIZK for $\mathcal{R}_{\mathsf{e}}$: $(\mathsf{pp}_{\mathsf{e}}, \mathsf{td}_{\mathsf{e}}) \leftarrow \mathsf{NIZK.SimSetup}(1^\lambda, \mathcal{R}_{\mathsf{e}})$.
5. Set $\mathsf{pp} := (\mathsf{pp}_{\mathsf{CM}}, \mathsf{pp}_{\mathsf{TCM}}, \mathsf{pp}_{\mathsf{CRH}}, \mathsf{pp}_{\mathsf{e}})$.
6. Set $\mathsf{td} := (\mathsf{td}_{\mathsf{TCM}}, \mathsf{td}_{\mathsf{e}})$.
7. Output $(\mathsf{pp}, \mathsf{td})$.

$\mathsf{DPC.SimExecute}^{\mathbf{L}}$

*Input:*
- public parameters $\mathsf{pp}$ and trapdoor $\mathsf{td}$
- old serial numbers $[\mathsf{sn}_i]_1^m$
- new record commitments $[\mathsf{cm}_j]_1^n$
- transaction memorandum $\mathsf{memo}$

*Output:* new records $[\mathbf{r}_j]_1^n$ and transaction $\mathsf{tx}$

1. For $j \in \{1, \ldots, n\}$:
   (a) Set new serial number nonce $\rho_j := \mathsf{CRH.Eval}(\mathsf{pp}_{\mathsf{CRH}}, j \| \mathsf{sn}_1 \| \ldots \| \mathsf{sn}_m)$.
   (b) Set address public key, payload, predicates, and commitment randomness to be the empty string: $\mathsf{apk}_j, \mathsf{payload}_j, \Phi_{\mathsf{b},j}, \Phi_{\mathsf{d},j}, r_j := \varepsilon$.
   (c) Construct dummy record: $\mathbf{r}_j := (\mathsf{apk}_j, \mathsf{payload}_j, \Phi_{\mathsf{b},j}, \Phi_{\mathsf{d},j}, \rho_j, r_j, \mathsf{cm}_j)$.
2. Retrieve current ledger digest: $\mathsf{st}_{\mathbf{L}} \leftarrow \mathbf{L}.\mathsf{Digest}$.
3. Construct instance for relation $\mathcal{R}_{\mathsf{e}}$: $\mathbb{x}_{\mathsf{e}} := (\mathsf{st}_{\mathbf{L}}, [\mathsf{sn}_i]_1^m, [\mathsf{cm}_j]_1^n, \mathsf{memo})$.
4. Generate simulated proof for $\mathcal{R}_{\mathsf{e}}$: $\pi_{\mathsf{e}} \leftarrow \mathsf{NIZK.Simulate}(\mathsf{pp}_{\mathsf{e}}, \mathsf{td}_{\mathsf{e}}, \mathbb{x}_{\mathsf{e}})$.
5. Construct transaction: $\mathsf{tx} := ([\mathsf{sn}_i]_1^m, [\mathsf{cm}_j]_1^n, \mathsf{memo}, \star)$, where $\star := (\mathsf{st}_{\mathbf{L}}, \pi_{\mathsf{e}})$.
6. Output $([\mathbf{r}_j]_1^n, \mathsf{tx})$.

DPC.ExtractExecute

*Input:*
- public parameters $\mathsf{pp}$ and trapdoor $\mathsf{td}$
- transaction $\mathsf{tx}$

*Output:*
- old $\begin{cases} \text{records } [\mathbf{r}_i]_1^m \\ \text{address secret keys } [\mathsf{ask}_i]_1^m \end{cases}$
- new records $[\mathbf{r}_j]_1^n$
- auxiliary predicate input $\mathsf{aux}$

1. Parse $\mathsf{tx}$ as $([\mathsf{sn}_i]_1^m, [\mathsf{cm}_j]_1^n, \mathsf{memo}, \star)$ and $\star$ as $(\mathsf{st}_{\mathbf{L}}, \pi_{\mathsf{e}})$.
2. Construct instance for relation $\mathcal{R}_{\mathsf{e}}$: $\mathbb{x}_{\mathsf{e}} := (\mathsf{st}_{\mathbf{L}}, [\mathsf{sn}_i]_1^m, [\mathsf{cm}_j]_1^n, \mathsf{memo})$.
3. Obtain witness: $\mathbb{w}_{\mathsf{e}} \leftarrow \mathsf{NIZK.Extract}(\mathsf{pp}_{\mathsf{e}}, \mathsf{td}_{\mathsf{e}}, \mathbb{x}_{\mathsf{e}}, \pi_{\mathsf{e}})$.
4. Parse the witness $\mathbb{w}_{\mathsf{e}}$ as $([\mathbf{r}_i]_1^m, [\mathbb{w}_{\mathbf{L},i}]_1^m, [\mathsf{ask}_i]_1^m, [\mathbf{r}_j]_1^n, \mathsf{aux})$.
5. Output $([\mathbf{r}_i]_1^m, [\mathsf{ask}_i]_1^m, [\mathbf{r}_j]_1^n, \mathsf{aux})$.

**Figure 20:** Several subroutines used by the ideal-world simulator $\mathcal{S}$.

## A.3 Proof of security by hybrid argument

We use a sequence of hybrids, each identified by a game $\mathcal{G}_i$, to prove that the outputs of the environment $\mathcal{E}$ when interacting with the real-world (dummy) adversary $\mathcal{A}$ and the ideal-world simulator $\mathcal{S}$ are computationally indistinguishable. We denote by $\mathsf{Output}_i(\mathcal{E})$ the output of $\mathcal{E}$ in game $\mathcal{G}_i$, and by $\mathcal{G}_0$ the real-world execution.

- $\mathcal{G}_1$ (sample parameters):
  This game is the real-world execution modified as follows.
  - $\mathcal{E}$ interacts with $\mathcal{S}$ instead of $\mathcal{A}$.
  - $\mathcal{S}$ uses DPC.Setup to generate public parameters pp, and gives these to $\mathcal{E}$.
  - $\mathcal{S}$ maintains the ledger $\mathbf{L}$ for $\mathcal{E}$ (it appends to $\mathbf{L}$ any pushed transaction passing the checks in DPC.Verify).
  - $\mathcal{S}$ forwards messages from $\mathcal{E}$ to $\mathbf{L}$ and other parties.
  - $\mathcal{S}$ forwards messages from other honest parties to $\mathcal{E}$.

  $\mathsf{Output}_1(\mathcal{E})$ is perfectly indistinguishable from $\mathsf{Output}_0(\mathcal{E})$ since $\mathcal{S}$ samples the public parameters honestly, maintains the ledger identically to the ideal ledger, and otherwise behaves like the dummy adversary.

- $\mathcal{G}_2$ (simulate setup):
  $\mathcal{S}$ invokes DPC.SimSetup instead of DPC.Setup. $\mathsf{Output}_2(\mathcal{E})$ is perfectly indistinguishable from $\mathsf{Output}_1(\mathcal{E})$ since NIZK is perfect zero knowledge.

- $\mathcal{G}_3$ (simulate proofs):
  In all honest party transactions, $\mathcal{S}$ replaces NIZK proofs with simulated proofs produced via NIZK.Simulate. $\mathsf{Output}_3(\mathcal{E})$ is perfectly indistinguishable from $\mathsf{Output}_2(\mathcal{E})$ since NIZK is perfect zero knowledge.

- $\mathcal{G}_4$ (simulate serial numbers):
  In all honest party transactions, $\mathcal{S}$ replaces all serial numbers with uniformly random elements sampled from PRF's codomain. Since PRF is a pseudorandom function, and $\mathcal{E}$ does not know the secret key used to compute it, $\mathsf{Output}_4(\mathcal{E})$ is computationally indistinguishable from $\mathsf{Output}_3(\mathcal{E})$.

- $\mathcal{G}_5$ (simulate commitments and equivocate commitment openings):
  In all honest party transactions, $\mathcal{S}$ replaces record commitments with commitments to the empty string $\varepsilon$. In all messages from honest parties to corrupted parties containing record contents, $\mathcal{S}$ replaces the actual commitment randomness with randomness produced by TCM.Equivocate. $\mathsf{Output}_5(\mathcal{E})$ is perfectly indistinguishable from $\mathsf{Output}_4(\mathcal{E})$ since TCM is perfectly hiding and equivocation produces commitment randomness that is statistically close to uniform.

- $\mathcal{G}_6$ (handle adversarial transactions):
  For every corrupted party transaction, $\mathcal{S}$ extracts an NP instance $\mathtt{x}_\mathsf{e}$ and witness $\mathtt{w}_\mathsf{e}$ for $\mathcal{R}_\mathsf{e}$ from the included proof and then proceeds as follows.
  - If $(\mathtt{x}_\mathsf{e}, \mathtt{w}_\mathsf{e}) \notin \mathcal{R}$, $\mathcal{S}$ aborts. If NIZK is simulation-extractable, this occurs with negligible probability.
  - For all $i \in \{1, \ldots, m\}$, if the contents of any $\mathbf{r}_i$ are different from those seen in any `RecordAuth` from an honest party or in the output of a previously extracted transaction, $\mathcal{S}$ aborts. If TCM is a binding commitment scheme, then this occurs with negligible probability.
  - For all $i \in \{1, \ldots, m\}$, if the extracted secret key $\mathsf{ask}_i$ for $\mathsf{apk}_i$ differs from the secret key extracted for $\mathsf{apk}_i$ in a prior transaction, $\mathcal{S}$ aborts. If CM is a binding commitment scheme, then this occurs with negligible probability.

– For all $j \in \{1, \ldots, n\}$, if the serial number nonce $\rho_j$ matches one extracted in a prior transaction, $\mathcal{S}$ aborts. If $\mathsf{CRH}$ is a collision-resistant hash, then this occurs with negligible probability because the serial number nonce is the output of $\mathsf{CRH}$ evaluated (in part) over the serial numbers of the input records. If this input is distinct across two different invocations of $\mathsf{CRH}$, then collision resistance guarantees that a nonce collision happens with negligible probability. Now for the transaction to be valid, it must contain serial numbers not seen before on the ledger. Therefore, the inputs to $\mathsf{CRH}$ are never repeated.

$\mathsf{Output}_6(\mathcal{E})$ is therefore computationally indistinguishable from $\mathsf{Output}_5(\mathcal{E})$.

The final game is distributed identically to the operation of $\mathcal{S}$ from the point of view of $\mathcal{E}$. We have thus shown that $\mathcal{E}$'s advantage in distinguishing the interaction with $\mathcal{S}$ from the interaction with $\mathcal{A}$ is negligible.

# B Construction of a delegable DPC scheme

We provide more details on the delegable DPC scheme discussed in Section 5. First we give details on randomizable signatures (Appendix B.1), and then give pseudocode for the DPC construction (Appendix B.2).

## B.1 Definition and construction of a randomizable signature scheme

A *randomizable* signature scheme is a tuple of algorithms $\mathsf{SIG} = (\mathsf{Setup}, \mathsf{Keygen}, \mathsf{Sign}, \mathsf{Verify}, \mathsf{RandPk}, \mathsf{RandSig})$ that enables a party to sign messages, while also allowing randomization of public keys and signatures to prevent linking across multiple signatures.

We have already described the syntax of the scheme's algorithms, and summarized its security properties, in Section 5.2. Now we discuss in more detail the security properties, and the construction used in our code.

**Security properties.** The signature scheme $\mathsf{SIG}$ satisfies the following security properties.

- *Existential unforgeability under randomization (EUR).* For every efficient adversary $\mathcal{A}$, the following probability is negligible:

$$\Pr\left[\begin{array}{c} \left(m^* \notin Q \text{ and } \mathsf{SIG.Verify}(\mathsf{pp_{SIG}}, \mathsf{pk_{SIG}}, m^*, \sigma^*)\right) \\ \text{or} \\ \left(m^* \notin Q \text{ and } \mathsf{SIG.Verify}(\mathsf{pp_{SIG}}, \hat{\mathsf{pk}}_{\mathsf{SIG}}, m^*, \sigma^*)\right) \end{array} \middle| \begin{array}{r} \mathsf{pp_{SIG}} \leftarrow \mathsf{SIG.Setup}(1^\lambda) \\ (\mathsf{pk_{SIG}}, \mathsf{sk_{SIG}}) \leftarrow \mathsf{SIG.Keygen}(\mathsf{pp_{SIG}}) \\ (m^*, \sigma^*, r^*_{\mathsf{SIG}}) \leftarrow \mathcal{A}^{S(\cdot)}(\mathsf{pp_{SIG}}, \mathsf{pk_{SIG}}) \\ \hat{\mathsf{pk}}_{\mathsf{SIG}} \leftarrow \mathsf{SIG.RandPk}(\mathsf{pp_{SIG}}, \mathsf{pk_{SIG}}, r^*_{\mathsf{SIG}}) \end{array}\right]$$

  where $S(m) := \mathsf{SIG.Sign}(\mathsf{pp_{SIG}}, \mathsf{sk_{SIG}}, m)$ and $Q$ are the queries made by $\mathcal{A}$ to the signing oracle $S$.

- *Unlinkability.* Every efficient adversary $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ has at most negligible advantage in guessing the bit $b$ in the IND-RSIG game below.

$\mathsf{IND\text{-}RSIG}^{\mathsf{SIG}}_{\mathcal{A}}(1^\lambda)$:
1. Generate public parameters: $\mathsf{pp_{SIG}} \leftarrow \mathsf{SIG.Setup}(1^\lambda)$.
2. Generate key pair: $(\mathsf{pk_{SIG}}, \mathsf{sk_{SIG}}) \leftarrow \mathsf{SIG.Keygen}(\mathsf{pp_{SIG}})$.
3. Obtain message from adversary: $m \leftarrow \mathcal{A}_1^{\mathsf{SIG.Sign}(\mathsf{pp_{SIG}}, \mathsf{sk_{SIG}}, \cdot)}(\mathsf{pp_{SIG}}, \mathsf{pk_{SIG}})$.
4. Sample a bit $b$ uniformly at random.
5. If $b = 0$:
   (a) Sample new key pair: $(\mathsf{pk'_{SIG}}, \mathsf{sk'_{SIG}}) \leftarrow \mathsf{SIG.Keygen}(\mathsf{pp_{SIG}})$.
   (b) Sign message: $\sigma \leftarrow \mathsf{SIG.Sign}(\mathsf{pp_{SIG}}, \mathsf{sk'_{SIG}}, m)$.
   (c) Set $c := (\mathsf{pk'_{SIG}}, \sigma)$.
6. If $b = 1$:
   (a) Sign message: $\sigma \leftarrow \mathsf{SIG.Sign}(\mathsf{pp_{SIG}}, \mathsf{sk_{SIG}}, m)$.
   (b) Sample randomness $r_{\mathsf{SIG}}$.
   (c) Randomize public key: $\hat{\mathsf{pk}}_{\mathsf{SIG}} \leftarrow \mathsf{SIG.RandPk}(\mathsf{pp_{SIG}}, \mathsf{pk_{SIG}}, r_{\mathsf{SIG}})$.
   (d) Randomize signature: $\hat{\sigma} \leftarrow \mathsf{SIG.RandSig}(\mathsf{pp_{SIG}}, \sigma, r_{\mathsf{SIG}})$.
   (e) Set $c := (\hat{\mathsf{pk}}_{\mathsf{SIG}}, \hat{\sigma})$.
7. Output $\mathcal{A}_2^{\mathsf{SIG.Sign}(\mathsf{pp_{SIG}}, \mathsf{sk_{SIG}}, \cdot)}(\mathsf{pp_{SIG}}, \mathsf{pk_{SIG}}, c)$.

- *Injective randomization.* For every efficient adversary $\mathcal{A}$, the following probability is negligible:

$$\Pr\left[\begin{array}{c} r_1 \neq r_2 \\ \mathsf{SIG.RandPk}(\mathsf{pp_{SIG}}, \mathsf{pk_{SIG}}, r_1) = \mathsf{SIG.RandPk}(\mathsf{pp_{SIG}}, \mathsf{pk_{SIG}}, r_2) \end{array} \middle| \begin{array}{r} \mathsf{pp_{SIG}} \leftarrow \mathsf{SIG.Setup}(1^\lambda) \\ (\mathsf{pk_{SIG}}, r_1, r_2) \leftarrow \mathcal{A}(\mathsf{pp_{SIG}}) \end{array}\right].$$

**Construction.** In Fig. 21 we provide a modification of the Schnorr signature scheme [Sch91] that is randomizable. We briefly explain why this modification satisfies the security properties above.

- *Existential unforgeability under randomization (EUR).* Given an efficient adversary $\mathcal{A}$ that breaks EUR of randomizable Schnorr signatures, we construct an efficient adversary $\mathcal{A}'$ that breaks existential unforgeability of standard Schnorr signatures. In detail, $\mathcal{A}'$ forwards signature queries from $\mathcal{A}$ to its own signing oracle and returns the answers to $\mathcal{A}$ and then, when $\mathcal{A}$ outputs a tuple $(m^*, \sigma^*, r^*_{\mathsf{SIG}})$, $\mathcal{A}'$ outputs the tuple $(m^*, \sigma)$ where $\sigma$ is computed as follows. If $\sigma^*$ is a valid signature for $m^*$ under $\mathsf{pk}_{\mathsf{SIG}}$ then $\sigma := \sigma^*$. Otherwise, $\mathcal{A}'$ "undoes" the randomization of $\sigma^* = (s, e)$ by setting $\sigma := (s + e \cdot r^*_{\mathsf{SIG}}, e)$; thus if $\mathcal{A}$ outputs a forgery for a randomization of $\mathsf{pk}_{\mathsf{SIG}}$, $\mathcal{A}'$ translates it back into a forgery for $\mathsf{pk}_{\mathsf{SIG}}$. In sum, since standard Schnorr signatures are secure in the random oracle model assuming hardness of discrete logarithms [PS00], so is the randomizable variant under the same assumptions.
- *Unlinkability of public keys.* Public keys are unlinkable because SIG.RandPk multiplies the public key pk (which is a group element) by a random group element; the result is statistically independent of pk.
- *Unlinkability of signatures.* The only part of a Schnorr signature that depends on the public or secret key is the scalar $s$. Since SIG.RandSig adds a random shift to $s$, the result is statistically independent of the signature's original key pair.
- *Injective randomization.* Fixing all inputs but for $r_{\mathsf{SIG}}$, SIG.RandPk is a permutation over $\mathbb{G}$. Hence, finding collisions over the randomness is not possible.

$\mathsf{SIG.Setup}(1^\lambda) \to \mathsf{pp}_{\mathsf{SIG}}$
1. Sample a group: $(\mathbb{G}, q, g) \leftarrow \mathsf{SampleGrp}(1^\lambda)$.
2. Sample cryptographic hash function $H$.
3. Output $\mathsf{pp}_{\mathsf{SIG}} := (\mathbb{G}, q, g, H)$.

$\mathsf{SIG.Keygen}(\mathsf{pp}_{\mathsf{SIG}}) \to (\mathsf{pk}_{\mathsf{SIG}}, \mathsf{sk}_{\mathsf{SIG}})$
1. Parse $\mathsf{pp}_{\mathsf{SIG}}$ as $(\mathbb{G}, q, g, H)$.
2. Sample a scalar $x$ uniformly from $\mathbb{Z}_q$.
3. Output $(\mathsf{pk}_{\mathsf{SIG}}, \mathsf{sk}_{\mathsf{SIG}}) := (g^x, x)$.

$\mathsf{SIG.Verify}(\mathsf{pp}_{\mathsf{SIG}}, \mathsf{pk}_{\mathsf{SIG}}, m, \sigma) \to b$
1. Parse $\mathsf{pp}_{\mathsf{SIG}}$ as $(\mathbb{G}, q, g, H)$.
2. Parse $\sigma$ as $(s, e)$.
3. Set $r_v := g^s \mathsf{pk}_{\mathsf{SIG}}^e = g^{s+xe}$.
4. Set $e_v := H(r_v \| m)$.
5. Check if $e = e_v$.

$\mathsf{SIG.Sign}(\mathsf{pp}_{\mathsf{SIG}}, \mathsf{sk}_{\mathsf{SIG}}, m) \to \sigma$
1. Parse $\mathsf{pp}_{\mathsf{SIG}}$ as $(\mathbb{G}, q, g, H)$.
2. Sample a scalar $k$ uniformly from $\mathbb{Z}_q$.
3. Set $r := g^k$ and $e := H(r \| m)$.
4. Set $s := k - xe$.
5. Output $\sigma := (s, e)$.

$\mathsf{SIG.RandPk}(\mathsf{pp}_{\mathsf{SIG}}, \mathsf{pk}_{\mathsf{SIG}}, r_{\mathsf{SIG}}) \to \hat{\mathsf{pk}}_{\mathsf{SIG}}$
1. Parse $\mathsf{pp}_{\mathsf{SIG}}$ as $(\mathbb{G}, q, g, H)$.
2. Output $\hat{\mathsf{pk}}_{\mathsf{SIG}} := \mathsf{pk}_{\mathsf{SIG}} \cdot g^{r_{\mathsf{SIG}}}$.

$\mathsf{SIG.RandSig}(\mathsf{pp}_{\mathsf{SIG}}, \sigma, r_{\mathsf{SIG}}) \to \hat{\sigma}$
1. Parse $\mathsf{pp}_{\mathsf{SIG}}$ as $(\mathbb{G}, q, g, H)$.
2. Parse $\sigma$ as $(s, e)$.
3. Output $\hat{\sigma} := (s - e \cdot r_{\mathsf{SIG}}, e)$.

**Figure 21:** Construction of a randomizable signature scheme based on the Schnorr signature scheme [Sch91].

## B.2 Construction of a delegable DPC scheme

Fig. 22 provides pseudocode that, together with the modified NP relation $\mathcal{R}^{\mathsf{del}}_{\mathsf{e}}$ given in Fig. 23, formalizes the high-level description of a delegable DPC scheme from Section 5.3. In both figures, we highlighted changes from the "plain" DPC scheme in Section 4.2. *The only step in* DPC.Execute *that must be performed by the delegator is Step 7a; all other steps can be performed by the worker without knowing the signature secret key.*

**DPC.Setup**

*Input:* security parameter $1^\lambda$

*Output:* public parameters $\mathsf{pp}$

1. Generate **commitment parameters**: $\mathsf{pp}_{\mathsf{CM}} \leftarrow \mathsf{CM.Setup}(1^\lambda)$, $\mathsf{pp}_{\mathsf{TCM}} \leftarrow \mathsf{TCM.Setup}(1^\lambda)$.
2. Generate **CRH parameters**: $\mathsf{pp}_{\mathsf{CRH}} \leftarrow \mathsf{CRH.Setup}(1^\lambda)$.
3. Generate **signature parameters**: $\mathsf{pp}_{\mathsf{SIG}} \leftarrow \mathsf{SIG.Setup}(1^\lambda)$.
4. Generate **NIZK parameters for** $\mathcal{R}_{\mathsf{e}}^{\mathsf{del}}$ (Fig. 23): $\mathsf{pp}_{\mathsf{e}} \leftarrow \mathsf{NIZK.Setup}(1^\lambda, \mathcal{R}_{\mathsf{e}}^{\mathsf{del}})$.
5. Output $\mathsf{pp} := (\mathsf{pp}_{\mathsf{CM}}, \mathsf{pp}_{\mathsf{TCM}}, \mathsf{pp}_{\mathsf{CRH}}, \mathsf{pp}_{\mathsf{SIG}}, \mathsf{pp}_{\mathsf{e}})$.

**DPC.GenAddress**

*Input:* public parameters $\mathsf{pp}$

*Output:* address key pair $(\mathsf{apk}, \mathsf{ask})$

1. Generate **authorization key pair**: $(\mathsf{pk}_{\mathsf{SIG}}, \mathsf{sk}_{\mathsf{SIG}}) \leftarrow \mathsf{SIG.Keygen}(\mathsf{pp}_{\mathsf{SIG}})$ .
2. Sample secret key $\mathsf{sk}_{\mathsf{PRF}}$ for pseudorandom function PRF.
3. Sample randomness $r_{\mathsf{pk}}$ for commitment scheme TCM.
4. Set **address public key** $\mathsf{apk} := \mathsf{CM.Commit}(\mathsf{pp}_{\mathsf{CM}}, \mathsf{pk}_{\mathsf{SIG}} \| \mathsf{sk}_{\mathsf{PRF}};\ r_{\mathsf{pk}})$.
5. Set **address secret key** $\mathsf{ask} := (\mathsf{sk}_{\mathsf{SIG}}, \mathsf{sk}_{\mathsf{PRF}}, r_{\mathsf{pk}})$.
6. Output $(\mathsf{apk}, \mathsf{ask})$.

**DPC.Execute$^{\mathbf{L}}$**

*Input:*

- public parameters $\mathsf{pp}$
- old { records $[\mathbf{r}_i]_1^m$; address secret keys $[\mathsf{ask}_i]_1^m$ }
- auxiliary predicate input $\mathsf{aux}$
- transaction memorandum $\mathsf{memo}$
- new { address public keys $[\mathsf{apk}_j]_1^n$; record payloads $[\mathsf{payload}_j]_1^n$; record birth predicates $[\Phi_{\mathsf{b},j}]_1^n$; record death predicates $[\Phi_{\mathsf{d},j}]_1^n$ }

*Output:* new records $[\mathbf{r}_j]_1^n$ and transaction $\mathsf{tx}$

1. For each $i \in \{1, \ldots, m\}$, process the $i$-th old record as follows:
   (a) Parse old record $\mathbf{r}_i$ as $\mathbf{r}_i = \begin{pmatrix} \text{address public key} & \mathsf{apk}_i & \text{payload} & \mathsf{payload}_i & \text{comm. rand.} & r_i \\ \text{serial number nonce} & \rho_i & \text{predicates} & (\Phi_{\mathsf{b},i}, \Phi_{\mathsf{d},i}) & \text{commitment} & \mathsf{cm}_i \end{pmatrix}$.
   (b) If $\mathsf{payload}_i.\mathsf{isDummy} = 1$, set **ledger membership witness** $\mathrm{w}_{\mathbf{L},i} := \perp$.
   If $\mathsf{payload}_i.\mathsf{isDummy} = 0$, compute **ledger membership witness** for commitment: $\mathrm{w}_{\mathbf{L},i} \leftarrow \mathbf{L}.\mathsf{Prove}(\mathsf{cm}_i)$.
   (c) Parse address secret key $\mathsf{ask}_i$ as $(\mathsf{sk}_{\mathsf{SIG},i}, \mathsf{sk}_{\mathsf{PRF},i}, r_{\mathsf{pk},i})$ and derive $\mathsf{pk}_{\mathsf{SIG},i}$ from $\mathsf{sk}_{\mathsf{SIG},i}$.
   (d) Compute **signature randomness**: $r_{\mathsf{SIG},i} \leftarrow \mathsf{PRF}_{\mathsf{sk}_{\mathsf{PRF},i}}(\rho_i)$.
   (e) Compute **serial number**: $\mathsf{sn}_i \leftarrow \mathsf{SIG.RandPk}(\mathsf{pp}_{\mathsf{SIG}}, \mathsf{pk}_{\mathsf{SIG},i}, r_{\mathsf{SIG},i})$.
2. For each $j \in \{1, \ldots, n\}$, construct the $j$-th new record as follows:
   (a) Compute **serial number nonce**: $\rho_j := \mathsf{CRH.Eval}(\mathsf{pp}_{\mathsf{CRH}}, j \| \mathsf{sn}_1 \| \ldots \| \mathsf{sn}_m)$.
   (b) Construct **new record**: $\mathbf{r}_j \leftarrow \mathsf{DPC.ConstructRecord}(\mathsf{pp}, \mathsf{apk}_j, \mathsf{payload}_j, \Phi_{\mathsf{b},j}, \Phi_{\mathsf{d},j}, \rho_j)$.
3. Retrieve current **ledger digest**: $\mathsf{st}_{\mathbf{L}} \leftarrow \mathbf{L}.\mathsf{Digest}$.
4. Construct **instance for relation** $\mathcal{R}_{\mathsf{e}}^{\mathsf{del}}$: $\mathbb{x}_{\mathsf{e}} := (\mathsf{st}_{\mathbf{L}}, [\mathsf{sn}_i]_1^m, [\mathsf{cm}_j]_1^n, \mathsf{memo})$.
5. Construct **witness for relation** $\mathcal{R}_{\mathsf{e}}^{\mathsf{del}}$: $\mathrm{w}_{\mathsf{e}} := ([\mathbf{r}_i]_1^m, [\mathrm{w}_{\mathbf{L},i}]_1^m, [\mathsf{sk}_{\mathsf{PRF},i}]_1^m, [\mathsf{pk}_{\mathsf{SIG},i}]_1^m, [\mathsf{meta}_i]_1^m, [r_{\mathsf{pk},i}]_1^m, [\mathbf{r}_j]_1^n, \mathsf{aux})$.
6. Generate **proof for relation** $\mathcal{R}_{\mathsf{e}}^{\mathsf{del}}$: $\pi_{\mathsf{e}} \leftarrow \mathsf{NIZK.Prove}(\mathsf{pp}_{\mathsf{e}}, \mathbb{x}_{\mathsf{e}}, \mathrm{w}_{\mathsf{e}})$.
7. For each $i \in \{1, \ldots, m\}$:
   (a) **Sign message**: $\sigma_i \leftarrow \mathsf{SIG.Sign}(\mathsf{pp}_{\mathsf{SIG}}, \mathsf{sk}_{\mathsf{SIG},i}, \mathbb{x}_{\mathsf{e}} \| \pi_{\mathsf{e}})$.
   (b) **Randomize signature**: $\hat{\sigma}_i \leftarrow \mathsf{SIG.RandSig}(\mathsf{pp}_{\mathsf{SIG}}, \sigma_i, r_{\mathsf{SIG},i})$.
8. Construct **transaction**: $\mathsf{tx} := ([\mathsf{sn}_i]_1^m, [\mathsf{cm}_j]_1^n, \mathsf{memo}, \star)$, where $\star := (\mathsf{st}_{\mathbf{L}}, \pi_{\mathsf{e}}, [\hat{\sigma}_i]_1^m)$.
9. Output $([\mathbf{r}_j]_1^n, \mathsf{tx})$.

**DPC.Verify$^{\mathbf{L}}$**

*Input:* public parameters $\mathsf{pp}$ and transaction $\mathsf{tx}$

*Output:* decision bit $b$

1. Parse $\mathsf{tx}$ as $([\mathsf{sn}_i]_1^m, [\mathsf{cm}_j]_1^n, \mathsf{memo}, \star)$ and $\star$ as $(\mathsf{st}_{\mathbf{L}}, \pi_{\mathsf{e}}, [\hat{\sigma}_i]_1^m)$.
2. Check that **there are no duplicate serial numbers**
   (a) within the transaction $\mathsf{tx}$: $\mathsf{sn}_i \neq \mathsf{sn}_j$ for every distinct $i, j \in \{1, \ldots, m\}$;
   (b) on the ledger: $\mathbf{L}.\mathsf{Contains}(\mathsf{sn}_i) = 0$ for every $i \in \{1, \ldots, m\}$.
3. Check that **the ledger state is valid**: $\mathbf{L}.\mathsf{ValidateDigest}(\mathsf{st}_{\mathbf{L}}) = 1$.
4. Construct **instance for the relation** $\mathcal{R}_{\mathsf{e}}^{\mathsf{del}}$: $\mathbb{x}_{\mathsf{e}} := (\mathsf{st}_{\mathbf{L}}, [\mathsf{sn}_i]_1^m, [\mathsf{cm}_j]_1^n, \mathsf{memo})$.
5. Check **proof for the relation** $\mathcal{R}_{\mathsf{e}}^{\mathsf{del}}$: $\mathsf{NIZK.Verify}(\mathsf{pp}_{\mathsf{e}}, \mathbb{x}_{\mathsf{e}}, \pi_{\mathsf{e}}) = 1$.
6. For every $i \in \{1, \ldots, m\}$, check that **signature verifies**: $\mathsf{SIG.Verify}(\mathsf{pp}_{\mathsf{SIG}}, \mathsf{sn}_i, \mathbb{x}_{\mathsf{e}} \| \pi_{\mathsf{e}}, \hat{\sigma}_i) = 1$.

**Figure 22:** Construction of a delegable DPC scheme. Highlights denote differences from Figure 8.

The NP relation $\mathcal{R}_{\mathsf{e}}^{\mathsf{del}}$ has instances $\mathbb{x}_{\mathsf{e}}$ and witnesses $\mathbb{w}_{\mathsf{e}}$ of the following form.

$$\mathbb{x}_{\mathsf{e}} = \begin{pmatrix} \text{ledger digest} & \mathsf{st}_{\mathbf{L}} \\ \text{old record serial numbers} & [\mathsf{sn}_i]_1^m \\ \text{new record commitments} & [\mathsf{cm}_j]_1^n \\ \text{transaction memorandum} & \mathsf{memo} \end{pmatrix} \quad \text{and} \quad \mathbb{w}_{\mathsf{e}} = \begin{pmatrix} \text{old records} & [\mathbf{r}_i]_1^m \\ \text{old record membership witnesses} & [\mathbb{w}_{\mathbf{L},i}]_1^m \\ \text{old record authorization public keys} & [\mathsf{sk}_{\mathsf{PRF},i}]_1^m \\ \text{old record serial number secret keys} & [\mathsf{pk}_{\mathsf{SIG},i}]_1^m \\ \text{old record address randomness} & [r_{\mathsf{pk},i}]_1^m \\ \text{new records} & [\mathbf{r}_j]_1^n \\ \text{auxiliary predicate input} & \mathsf{aux} \end{pmatrix}$$

where

- for each $i \in \{1, \ldots, m\}$, $\mathbf{r}_i = (\mathsf{apk}_i, \mathsf{payload}_i, \Phi_{\mathsf{b},i}, \Phi_{\mathsf{d},i}, \rho_i, r_i, \mathsf{cm}_i)$;
- for each $j \in \{1, \ldots, n\}$, $\mathbf{r}_j = (\mathsf{apk}_j, \mathsf{payload}_j, \Phi_{\mathsf{b},j}, \Phi_{\mathsf{d},j}, \rho_j, r_j, \mathsf{cm}_j)$.

Define the local data $\mathsf{ldata} := \begin{pmatrix} [\mathsf{cm}_i]_1^m & [\mathsf{apk}_i]_1^m & [\mathsf{payload}_i]_1^m & [\Phi_{\mathsf{d},i}]_1^m & [\Phi_{\mathsf{b},i}]_1^m & [\mathsf{sn}_i]_1^m & \mathsf{memo} \\ [\mathsf{cm}_j]_1^n & [\mathsf{apk}_j]_1^n & [\mathsf{payload}_j]_1^n & [\Phi_{\mathsf{d},j}]_1^n & [\Phi_{\mathsf{b},j}]_1^n & \mathsf{aux} & \end{pmatrix}$.

A witness $\mathbb{w}_{\mathsf{e}}$ is valid for an instance $\mathbb{x}_{\mathsf{e}}$ if the following conditions hold:

1. For each $i \in \{i, \ldots, m\}$:
   - If $\mathbf{r}_i$ is not dummy, $\mathbb{w}_{\mathbf{L},i}$ proves that the commitment $\mathsf{cm}_i$ is in a ledger with digest $\mathsf{st}_{\mathbf{L}}$: $\mathbf{L}.\mathsf{Verify}(\mathsf{st}_{\mathbf{L}}, \mathsf{cm}_i, \mathbb{w}_{\mathbf{L},i}) = 1$.
   - The address public key $\mathsf{apk}_i$ matches the authorization public key $\mathsf{pk}_{\mathsf{SIG},i}$ and the serial number secret key $\mathsf{sk}_{\mathsf{PRF},i}$: $\mathsf{apk}_i = \mathsf{CM}.\mathsf{Commit}(\mathsf{pp}_{\mathsf{CM}}, \mathsf{pk}_{\mathsf{SIG},i} \| \mathsf{sk}_{\mathsf{PRF},i}; r_{\mathsf{pk},i})$ .
   - The serial number $\mathsf{sn}_i$ is valid: $r_{\mathsf{SIG},i} = \mathsf{PRF}_{\mathsf{sk}_{\mathsf{PRF},i}}(\rho_i)$ and $\mathsf{sn}_i = \mathsf{SIG}.\mathsf{RandPk}(\mathsf{pp}_{\mathsf{SIG}}, \mathsf{pk}_{\mathsf{SIG},i}, r_{\mathsf{SIG},i})$.
   - The old record commitment $\mathsf{cm}_i$ is valid: $\mathsf{cm}_i = \mathsf{TCM}.\mathsf{Commit}(\mathsf{pp}_{\mathsf{TCM}}, \mathsf{apk}_i \| \mathsf{payload}_i \| \Phi_{\mathsf{b},i} \| \Phi_{\mathsf{d},i} \| \rho_i;\ r_i)$.
   - The death predicate $\Phi_{\mathsf{d},i}$ is satisfied by the local data: $\Phi_{\mathsf{d},i}(i \| \mathsf{ldata}) = 1$.
2. For each $j \in \{1, \ldots, n\}$:
   - The serial number nonce $\rho_j$ is computed correctly: $\rho_j = \mathsf{CRH}.\mathsf{Eval}(\mathsf{pp}_{\mathsf{CRH}}, j \| \mathsf{sn}_1 \| \ldots \| \mathsf{sn}_m)$.
   - The new record commitment $\mathsf{cm}_j$ is valid: $\mathsf{cm}_j = \mathsf{TCM}.\mathsf{Commit}(\mathsf{pp}_{\mathsf{TCM}}, \mathsf{apk}_j \| \mathsf{payload}_j \| \Phi_{\mathsf{b},j} \| \Phi_{\mathsf{d},j} \| \rho_j;\ r_j)$.
   - The birth predicate $\Phi_{\mathsf{b},j}$ is satisfied by the local data: $\Phi_{\mathsf{b},j}(j \| \mathsf{ldata}) = 1$.

**Figure 23:** The NP relation $\mathcal{R}_{\mathsf{e}}^{\mathsf{del}}$. Highlights denote differences from Figure 9.

# C Extensions in functionality and in security

We summarize some natural extensions of our DPC construction that give richer functionality, as well as methods to prove security notions beyond standalone non-adaptive security.

**Storing data in addresses.** For some applications it can be useful to verifiably associate address public keys with additional metadata meta. One can easily modify our construction to achieve this by using the address public key commitment to additionally commit to meta. To prove that a given address public key is bound to the metadata string meta, one can use a standard non-interactive zero knowledge proof of knowledge.[14]

With such a mechanism in hand, we can realise various useful functionality like *on-ledger encryption*: a user stores an encryption public key in the metadata of one of her addresses, and others can later use this public key to encrypt information about records created for her, and store the resulting ciphertext in the transaction's memorandum. This method, used for example in Zerocash [BCG+14], gives users the option to not use other out-of-band secure communication channels.

**Selective disclosure.** For compliance purposes, it may be useful to selectively reveal information about about a transaction to certain parties. Our implementation can be extended to support this by changing how hashes of predicate verification keys are committed to in a transaction: instead of committing all the verification keys together, one can instead commit to them in separate commitments. To disclose the predicates that were invoked in a transaction, a user can then simply open the relevant commitments.

**Ledger position.** In some applications it may be useful to know the unique ledger position of a record, i.e., to have this information be part of the local data ldata given as input to predicates. For example, one can use a record's ledger position to implement a "time lock" that prevents the record's consumption until a pre-specified amount of time has passed since the record's creation. However, the ledger interface we described in Section 3.1 does not expose this functionality: $\mathbf{L}.\mathsf{Prove}$ only returns a proof that a transaction (or a subcomponent thereof) appears on the ledger, and *not its position*. One can augment $\mathbf{L}.\mathsf{Prove}$ to instead output the transaction's ledger position $\mathsf{pos}_{\mathbf{L}}$, and a proof that $\mathsf{pos}_{\mathbf{L}}$ is the transaction's position on the ledger. Our instantiation of the ledger with a Merkle tree supports this augmentation inherently: the path to the transaction in the Merkle tree is also its position the tree.

**Composable security.** The security definition in Section 3.3 is a restriction of UC security definitions to a single execution at any given time. We can avoid this restriction and prove our construction UC-secure by replacing our simulation-extractable NIZKs with UC-secure NIZKs. The remainder of the proof would go through unchanged, and this would achieve composition of multiple protocol instances.

**Adaptive security.** We can prove adaptive security, with a minor modification to our protocol in Section 4. The barrier to proving security against adaptive corruptions (even in a standalone setting) is a lack of forward-secure privacy. Namely, when the adversary corrupts a party $\mathcal{P}$, it gets access to $\mathcal{P}$'s state, which includes contents of records held by $\mathcal{P}$ and address secret keys belonging to $\mathcal{P}$. The adversary can then use this information to break unlinkability of $\mathcal{P}$'s transactions by deriving the serial numbers of consumed records and matching these against those present on the ledger.

In the proof, this problem is reflected in how the simulator $\mathcal{S}$ handles serial numbers in honest party transactions (see Appendix A.2). For honest party transactions, serial numbers are sampled uniformly at random via SampleSn. When the environment $\mathcal{E}$ corrupts an honest party, it can attempt to carry out the aforementioned linking attack by computing serial numbers via the PRF. Since serial numbers already

[14] We do not need these NIZK proofs to be simulation-extractable since we do not extract from them. In fact, in our implementation we can even use specific sigma-protocols designed to prove knowledge of openings of Pedersen commitments.

published in transactions were derived randomly, they would not match the output of the PRF, allowing $\mathcal{E}$ to distinguish the ideal world from the real world.

We address this issue as follows. First, we work in the secure-erasure model and ensure that honest parties delete (a) all records output from $\mathsf{Execute}$ (after sending their contents to the intended recipients), and (b) all records that have been consumed. Hence, at the time a party is corrupted, the state revealed to the adversary does not contain secrets of past records, so the adversary cannot derive those records' serial numbers. Next, we have to convincingly match the address public keys of unconsumed records with corresponding address secret keys. To do this, we modify $\mathsf{DPC.GenAddress}$ to use trapdoor commitments to construct address public keys. The trapdoor property then allows us to open public keys to the correct secret keys.

However, these measures by themselves are not enough. Consider the following scenario: the adversary corrupts an honest user and learns her secret key. For every transaction in the ledger, it computes the serial number nonces of the output records from the serial numbers of the input records. The adversary can then use these nonces along with the secret key to derive candidate serial numbers for the output records. If these candidate serial numbers appear on the ledger, then the adversary learns that the record has been consumed.

To prevent this, we randomize the serial number nonces of all records output by $\mathsf{Execute}$ by deriving them as $\rho_j := \mathsf{CRH}(j \| r_{\rho,j} \| \mathsf{sn}_1 \| \cdots \| \mathsf{sn}_m)$ for some randomness $r_{\rho,j}$ that is deleted after invoking $\mathsf{Execute}$. This randomization ensures that the serial number nonce of an output record cannot be derived deterministically from the (publicly visible) serial numbers of the input records.

The above measures, however, are still insufficient: the adversary still knows the secrets of records that a corrupted party sent to an honest party. After corrupting this honest party, the adversary can learn its address secret key and therefore derive the serial number of those records. To overcome this obstacle, one can replace the PRF with a *programmable PRF* [PS18], for which the owner of the secret key can "program" the PRF to output pre-determined values on specific inputs: for all polynomial-sized sets $S = \{(x_i, y_i)\}_i$, the owner of a PRF secret key $\mathsf{sk}$ can derive a second key $\mathsf{sk}_S$ such that $\mathsf{PRF}_{\mathsf{sk}_s}(x_i) = y_i$ for each $(x_i, y_i) \in S$, while $\mathsf{PRF}_{\mathsf{sk}_S}(x) = \mathsf{PRF}_{\mathsf{sk}}(x)$ for other inputs $x$. This fixes the foregoing issue because $\mathcal{S}$ can now give $\mathcal{E}$ a programmed PRF secret key for the set $S = \{(\rho_i, \mathsf{sn}_i)\}_i$, where $\rho_i$ is the serial number nonce of the $i$-th record received from a corrupted party.

## Acknowledgments

We thank Samuel Steffen for bringing our attention to the "denial-of-funds" attack mentioned in Remark 6.1, which can be addressed as explained there. This work was supported in part by: a Google Faculty Award; the National Science Foundation under awards CNS-1653110 and CNS-1801479; the UC Berkeley Center for Long-Term Cybersecurity; and donations from the Ethereum Foundation, the Interchain Foundation, and Qtum.

## References

[ADMM14a] Marcin Andrychowicz, Stefan Dziembowski, Daniel Malinowski, and Lukasz Mazurek. Fair two-party computations via Bitcoin deposits. In *Proceedings of the BITCOIN workshop at the 18th International Conference on Financial Cryptography and Data Security*, FC '14, pages 105–121, 2014.

[ADMM14b] Marcin Andrychowicz, Stefan Dziembowski, Daniel Malinowski, and Lukasz Mazurek. Secure multiparty computations on Bitcoin. In *Proceedings of the 35th IEEE Symposium on Security and Privacy*, SP '14, pages 443–458, 2014.

[AFK+12] Diego F. Aranha, Laura Fuentes-Castañeda, Edward Knapp, Alfred Menezes, and Francisco Rodríguez-Henríquez. Implementing pairings at the 192-bit security level. In *Proceedings of the 5th International Conference on Pairing-Based Cryptography*, Pairing '12, pages 177–195, 2012.

[air] AirSwap. `https://www.airswap.io/`. Accessed 2018-12-27.

[AKR+13] Elli Androulaki, Ghassan Karame, Marc Roeschlin, Tobias Scherer, and Srdjan Capkun. Evaluating user privacy in Bitcoin. In *Proceedings of the 17th International Conference on Financial Cryptography and Data Security*, FC '13, pages 34–51, 2013.

[ANWW13] Jean-Phillipe Aumasson, Samuel Neves, Zooko Wilcox-O'Hearn, and Christian Winnerlein. BLAKE2: simpler, smaller, fast as MD5. In *Proceedings of the 11th International Conference on Applied Cryptography and Network Security*, ACNS '13, pages 119–135, 2013.

[BAZB20] Benedikt Bünz, Shashank Agrawal, Mahdi Zamani, and Dan Boneh. Zether: Towards privacy in a smart contract world. In *Proceedings of the 20th International Conference on Financial Cryptography and Data Security*, FC '20, pages 423–443, 2020.

[BBB+18] Benedikt Bünz, Jonathan Bootle, Dan Boneh, Andrew Poelstra, Pieter Wuille, and Gregory Maxwell. Bulletproofs: Short proofs for Confidential Transactions and more. In *Proceedings of the 39th IEEE Symposium on Security and Privacy*, S&P '18, pages 315–334, 2018.

[BBD+17] Iddo Bentov, Lorenz Breidenbach, Phil Daian, Ari Juels, Yunqi Li, and Xueyuan Zhao. The cost of decentralization in 0x and EtherDelta. `http://hackingdistributed.com/2017/08/13/cost-of-decent/`, 2017. Accessed 2019-01-03.

[BCC+16] Jonathan Bootle, Andrea Cerulli, Pyrros Chaidos, Jens Groth, and Christophe Petit. Efficient zero-knowledge arguments for arithmetic circuits in the discrete log setting. In *Proceedings of the 35th Annual International Conference on the Theory and Applications of Cryptographic Techniques*, EUROCRYPT '16, pages 327–357, 2016.

[BCCT13] Nir Bitansky, Ran Canetti, Alessandro Chiesa, and Eran Tromer. Recursive composition and bootstrapping for SNARKs and proof-carrying data. In *Proceedings of the 45th ACM Symposium on the Theory of Computing*, STOC '13, pages 111–120, 2013.

[BCG+13] Eli Ben-Sasson, Alessandro Chiesa, Daniel Genkin, Eran Tromer, and Madars Virza. SNARKs for C: verifying program executions succinctly and in zero knowledge. In *Proceedings of the 33rd Annual International Cryptology Conference*, CRYPTO '13, pages 90–108, 2013.

[BCG+14] Eli Ben-Sasson, Alessandro Chiesa, Christina Garman, Matthew Green, Ian Miers, Eran Tromer, and Madars Virza. Zerocash: Decentralized anonymous payments from Bitcoin. In *Proceedings of the 2014 IEEE Symposium on Security and Privacy*, SP '14, pages 459–474, 2014.

[BCG+15] Eli Ben-Sasson, Alessandro Chiesa, Matthew Green, Eran Tromer, and Madars Virza. Secure sampling of public parameters for succinct zero knowledge proofs. In *Proceedings of the 36th IEEE Symposium on Security and Privacy*, SP '15, pages 287–304, 2015.

[BCS16] Eli Ben-Sasson, Alessandro Chiesa, and Nicholas Spooner. Interactive oracle proofs. In *Proceedings of the 14th Theory of Cryptography Conference*, TCC '16-B, pages 31–60, 2016.

[BCTV14] Eli Ben-Sasson, Alessandro Chiesa, Eran Tromer, and Madars Virza. Succinct non-interactive zero knowledge for a von neumann architecture. In *Proceedings of the 23rd USENIX Security Symposium*, USENIX '14, pages 781–796, 2014.

[BCTV17] Eli Ben-Sasson, Alessandro Chiesa, Eran Tromer, and Madars Virza. Scalable zero knowledge via cycles of elliptic curves. *Algorithmica*, 79(4):1102–1160, 2017.

[BDJT17] Lorenz Breidenbach, Phil Daian, Ari Juels, and Florian Tramèr. To sink frontrunners, send in the submarines. `http://hackingdistributed.com/2017/08/28/submarine-sends/`, 2017. Accessed 2019-01-03.

[BGG18] Sean Bowe, Ariel Gabizon, and Matthew D. Green. A multi-party protocol for constructing the public parameters of the pinocchio zk-SNARK, 2018.

[BGM17] Sean Bowe, Ariel Gabizon, and Ian Miers. Scalable multi-party computation for zk-SNARK parameters in the random beacon model. Cryptology ePrint Archive, Report 2017/1050, 2017.

[bin] Binance. `https://www.binance.com/`. Accessed 2019-01-03.

[Bit15] Bitcoin. Some miners generating invalid blocks. `https://bitcoin.org/en/alert/2015-07-04-spv-mining`, 2015.

[BKM17] Iddo Bentov, Ranjit Kumaresan, and Andrew Miller. Instantaneous decentralized poker. In *Proceedings of the 23rd Annual International Conference on the Theory and Application of Cryptology and Information Security*, ASIACRYPT '17, pages 410–440, 2017.

[BLS02] Paulo S. L. M. Barreto, Ben Lynn, and Michael Scott. Constructing elliptic curves with prescribed embedding degrees. In *Proceedings of the 3rd International Conference on Security in Communication Networks*, SCN '02, pages 257–267, 2002.

[Bow17a] Sean Bowe. Bellman, 2017. URL: `https://github.com/zkcrypto/bellman`.

[Bow17b] Sean Bowe. Pairing, 2017. URL: `https://github.com/zkcrypto/pairing`.

[Can01] Ran Canetti. Universally composable security: A new paradigm for cryptographic protocols. In *Proceedings of the 42nd Annual Symposium on Foundations of Computer Science*, FOCS '01, pages 136–145, 2001.

[CCW19] Alessandro Chiesa, Lynn Chua, and Matthew Weidner. On cycles of pairing-friendly elliptic curves. *SIAM Journal on Applied Algebra and Geometry*, 3(2):175–192, 2019.

[CFH+15] Craig Costello, Cédric Fournet, Jon Howell, Markulf Kohlweiss, Benjamin Kreuter, Michael Naehrig, Bryan Parno, and Samee Zahur. Geppetto: Versatile verifiable computation. In *Proceedings of the 2015 IEEE Symposium on Security and Privacy*, IEEE S&P 2015, pages 253–270, 2015.

[CFQ19] Matteo Campanelli, Dario Fiore, and Anaïs Querol. LegoSNARK: modular design and composition of succinct zero-knowledge proofs. In *Proceedings of the 26th ACM Conference on Computer and Communications Security*, CCS '19, pages 2075–2092, 2019.

[CGL+17] Alessandro Chiesa, Matthew Green, Jingcheng Liu, Peihan Miao, Ian Miers, and Pratyush Mishra. Decentralized anonymous micropayments. In *Proceedings of the 36th Annual International Conference on the Theory and Applications of Cryptographic Techniques*, EUROCRYPT '17, pages 609–642, 2017.

[Cha14] Chainalysis. Chainalysis inc. `https://chainalysis.com/`, 2014.

[CHM+20] Alessandro Chiesa, Yuncong Hu, Mary Maller, Pratyush Mishra, Noah Vesely, and Nicholas Ward. Marlin: Preprocessing zkSNARKs with universal and updatable SRS. In *Proceedings of the 39th Annual International Conference on the Theory and Applications of Cryptographic Techniques*, EUROCRYPT '20, pages 738–768, 2020.

[Cim18] Catalin Cimpanu. Zaif cryptocurrency exchange loses $60 million in recent hack. `https://www.zdnet.com/article/zaif-cryptocurrency-exchange-loses-60-million-in-july-hack/`, 2018. Accessed 2018-12-27.

[CLN11] Craig Costello, Kristin E. Lauter, and Michael Naehrig. Attractive subfamilies of BLS curves for implementing high-security pairings. In *Proceedings of the 12th International Conference on Cryptology in India*, INDOCRYPT '11, pages 320–342, 2011.

[coi] Coinbase. `https://www.coinbase.com/`. Accessed 2019-01-03.

[Cos12] Craig Costello. Particularly friendly members of family trees. Cryptology ePrint Archive, Report 2012/072, 2012.

[CRR11] Ran Canetti, Ben Riva, and Guy N. Rothblum. Practical delegation of computation using multiple servers. In *Proceedings of the 19th ACM Conference on Computer and Communications Security*, CCS '11, pages 445–454, 2011.

[CRR13] Ran Canetti, Ben Riva, and Guy N. Rothblum. Refereed delegation of computation. *Information and Computation*, 226:16–36, 2013.

[CZJ+17] Ethan Cecchetti, Fan Zhang, Yan Ji, Ahmed E. Kosba, Ari Juels, and Elaine Shi. Solidus: Confidential distributed ledger transactions via PVORM. In *Proceedings of the 24th ACM Conference on Computer and Communications Security*, CCS '17, pages 701–717, 2017.

[CZK+18] Raymond Cheng, Fan Zhang, Jernej Kos, Warren He, Nichola Hynes, Noah Johnson, Ari Juels, Andrew Miller, and Dawn Song. Ekiden: A platform for confidentiality-preserving, trustworthy, and performant smart contract execution. arXiv cs.CR/1804.05141, 2018.

[DDO+01] Alfredo De Santis, Giovanni Di Crescenzo, Rafail Ostrovsky, Giuseppe Persiano, and Amit Sahai. Robust non-interactive zero knowledge. In *Proceedings of the 21st Annual International Cryptology Conference*, CRYPTO '01, pages 566–598, 2001.

[De18] Nikhilesh De. Coincheck confirms crypto hack loss larger than Mt. Gox. `https://www.coindesk.com/coincheck-confirms-crypto-hack-loss-larger-than-mt-gox`, 2018. Accessed 2018-12-27.

[DF91] Yvo Desmedt and Yair Frankel. Shared generation of authenticators and signatures (extended abstract). In *Proceedings of the 11th Annual International Cryptology Conference*, CRYPTO '91, pages 457–469, 1991.

[DGK+20] Philip Daian, Steven Goldfeder, Tyler Kell, Yunqi Li, Xueyuan Zhao, Iddo Bentov, Lorenz Breidenbach, and Ari Juels. Flash Boys 2.0: Frontrunning, transaction reordering, and consensus instability in decentralized exchanges. In *Proceedings of the 41st IEEE Symposium on Security and Privacy*, S&P '20, pages 910–927, 2020.

[Dod07] Yevgeniy Dodis. Lecture notes: Exposure-resilient cryptography, 2007. `https://www.cs.nyu.edu/courses/spring07/G22.3033-013/`.

[DSC+15] Tien Tuan Anh Dinh, Prateek Saxena, Ee-Chien Chang, Beng Chin Ooi, and Chunwang Zhang. M2R: enabling stronger privacy in MapReduce computation. In *Proceedings of the 24th USENIX Security Symposium*, USENIX Security '15, pages 447–462, 2015.

[Ell13] Elliptic. Elliptic enterprises limited. https://www.elliptic.co/, 2013.

[EMC19] Shayan Eskandari, Seyedehmahsa Moosavi, and Jeremy Clark. SoK: Transparent Dishonesty: front-running attacks on blockchain. arXiv cs.CR/1902.05164, 2019.

[EOS18] EOS. EOS.IO technical white paper, 2018. https://github.com/EOSIO/Documentation/blob/master/TechnicalWhitePaper.md.

[Eth16] Ethereum. I thikn the attacker is this miner - today he made over $50k. https://www.reddit.com/r/ethereum/comments/55xh2w/i_thikn_the_attacker_is_this_miner_today_he_made/, 2016.

[Eth18] Etherscan. The ethereum block explorer. https://etherscan.io/tokens, 2018.

[FK97] Uriel Feige and Joe Kilian. Making games short. In *Proceedings of the 29th ACM Symposium on the Theory of Computing*, STOC '97, pages 506–516, 1997.

[FKM+16] Nils Fleischhacker, Johannes Krupp, Giulio Malavolta, Jonas Schneider, Dominique Schröder, and Mark Simkin. Efficient unlinkable sanitizable signatures from signatures with re-randomizable keys. In *Proceedings of the 19th International Conference on Practice and Theory in Public-Key Cryptography*, PKC '16, pages 301–330, 2016.

[FST10] David Freeman, Michael Scott, and Edlyn Teske. A taxonomy of pairing-friendly elliptic curves. *Journal of Cryptology*, 23(2):224–280, 2010.

[gem] Gemini. https://gemini.com/dollar. Accessed 2019-01-24.

[GGPR13] Rosario Gennaro, Craig Gentry, Bryan Parno, and Mariana Raykova. Quadratic span programs and succinct NIZKs without PCPs. In *Proceedings of the 32nd Annual International Conference on Theory and Application of Cryptographic Techniques*, EUROCRYPT '13, pages 626–645, 2013.

[GM17] Jens Groth and Mary Maller. Snarky signatures: Minimal signatures of knowledge from simulation-extractable snarks. In *Proceedings of the 37th Conference on the Theory and Applications of Cryptographic Techniques*, CRYPTO '17, pages 581–612, 2017.

[GO94] Oded Goldreich and Yair Oren. Definitions and properties of zero-knowledge proof systems. *Journal of Cryptology*, 7(1):1–32, 1994.

[Goo14] L.M. Goodman. Tezos — a self-amending crypto-ledger, 2014. https://tezos.com/static/white_paper-2dc8c02267a8fb86bd67a108199441bf.pdf.

[Gro06] Jens Groth. Simulation-sound NIZK proofs for a practical language and constant size group signatures. In *Proceedings of the 12th International Conference on the Theory and Application of Cryptology and Information Security*, ASIACRYPT '06, pages 444–459, 2006.

[GWC19] Ariel Gabizon, Zachary J. Williamson, and Oana Ciobotaru. Plonk: Permutations over lagrange-bases for oecumenical noninteractive arguments of knowledge. Cryptology ePrint Archive, Report 2019/953, 2019.

[Har18] Colin Harper. In search of stability: An overview of the budding stablecoin ecosystem. https://bitcoinmagazine.com/articles/search-stability-overview-budding-stablecoin-ecosystem/, 2018. Accessed 2019-1-24.

[HBHW20] Daira Hopwood, Sean Bowe, Taylor Hornby, and Nathan Wilcox. Zcash protocol specification, 2020. URL: https://github.com/zcash/zips/blob/master/protocol/protocol.pdf.

[IKOS06] Yuval Ishai, Eyal Kushilevitz, Rafail Ostrovsky, and Amit Sahai. Cryptography from anonymity. In *Proceedings of the 47th Annual Symposium on Foundations of Computer Science*, FOCS '06, pages 239–248, 2006.

[JSST16] Sanjay Jain, Prateek Saxena, Frank Stephan, and Jason Teutsch. How to verify computation with a rational network. arXiv cs.GT/1606.05917, 2016.

[KB16] Ranjit Kumaresan and Iddo Bentov. Amortizing secure computation with penalties. In *Proceedings of the 23rd ACM Conference on Computer and Communications Security*, CCS '16, pages 418–429, 2016.

[KGC+17] Harry Kalodner, Steven Goldfeder, Alishah Chator, Malte Möser, and Arvind Narayanan. BlockSci: design and applications of a blockchain analysis platform. arXiv cs.CR/1709.02489, 2017.

[KGC+18] Harry A. Kalodner, Steven Goldfeder, Xiaoqi Chen, S. Matthew Weinberg, and Edward W. Felten. Arbitrum: Scalable, private smart contracts. In *Proceedings of the 27th USENIX Security Symposium*, USENIX Security '18, pages 1353–1370, 2018.

[KGM19] Gabriel Kaptchuk, Matthew Green, and Ian Miers. Giving state to the stateless: Augmenting trustworthy computation with ledgers. In *Proceedings of the 26th Annual Network and Distributed System Security Symposium*, NDSS '19, 2019.

[KMB15] Ranjit Kumaresan, Tal Moran, and Iddo Bentov. How to use Bitcoin to play decentralized poker. In *Proceedings of the 22nd ACM Conference on Computer and Communications Security*, CCS '15, pages 195–206, 2015.

[KMS+16] Ahmed E. Kosba, Andrew Miller, Elaine Shi, Zikai Wen, and Charalampos Papamanthou. Hawk: The blockchain model of cryptography and privacy-preserving smart contracts. In *Proceedings of the 2016 IEEE Symposium on Security and Privacy*, SP '16, pages 839–858, 2016.

[LTKS15] Loi Luu, Jason Teutsch, Raghav Kulkarni, and Prateek Saxena. Demystifying incentives in the consensus computer. In *Proceedings of the 22nd ACM Conference on Computer and Communications Security*, CCS '15, pages 706–719, 2015.

[MAB+13] Frank McKeen, Ilya Alexandrovich, Alex Berenzon, Carlos V. Rozas, Hisham Shafi, Vedvyas Shanbhogue, and Uday R. Savagaonkar. Innovative instructions and software model for isolated execution. In *Proceedings of the Second Workshop on Hardware and Architectural Support for Security and Privacy*, HASP '13, page 10, 2013.

[MBKM19] Mary Maller, Sean Bowe, Markulf Kohlweiss, and Sarah Meiklejohn. Sonic: Zero-knowledge SNARKs from linear-size universal and updateable structured reference strings. In *Proceedings of the 26th ACM Conference on Computer and Communications Security*, CCS '19, 2019.

[Mer87] Ralph C. Merkle. A digital signature based on a conventional encryption function. In *Proceedings of the 7th Conference on the Theory and Applications of Cryptographic Techniques*, CRYPTO '87, pages 369–378, 1987.

[MGGR13] Ian Miers, Christina Garman, Matthew Green, and Aviel D. Rubin. Zerocoin: Anonymous distributed e-cash from Bitcoin. In *Proceedings of the 2013 IEEE Symposium on Security and Privacy*, SP '13, pages 397–411, 2013.

[Mic00] Silvio Micali. Computationally sound proofs. *SIAM Journal on Computing*, 30(4):1253–1298, 2000. Preliminary version appeared in FOCS '94.

[MPJ+13] Sarah Meiklejohn, Marjori Pomarole, Grant Jordan, Kirill Levchenko, Damon McCoy, Geoffrey M. Voelker, and Stefan Savage. A fistful of Bitcoins: characterizing payments among men with no names. In *Proceedings of the 2013 Internet Measurement Conference*, IMC '13, pages 127–140, 2013.

[MRK03] Silvio Micali, Michael O. Rabin, and Joe Kilian. Zero-knowledge sets. In *Proceedings of the 44th Annual Symposium on Foundations of Computer Science*, FOCS '03, pages 80–91, 2003.

[MS18] Izaak Meckler and Evan Shapiro. Coda: Decentralized cryptocurrency at scale, 2018. `https://cdn.codaprotocol.com/v2/static/coda-whitepaper-05-10-2018-0.pdf`.

[Nak09] Satoshi Nakamoto. Bitcoin: a peer-to-peer electronic cash system, 2009. URL: http://www.bitcoin.org/bitcoin.pdf.

[NKDM03] Antonio Nicolosi, Maxwell N. Krohn, Yevgeniy Dodis, and David Mazières. Proactive two-party signatures for user authentication. In *Proceedings of the Network and Distributed System Security Symposium*, NDSS '03, 2003.

[NVV18] Neha Narula, Willy Vasquez, and Madars Virza. zkLedger: Privacy-preserving auditing for distributed ledgers. In *Proceedings of the 15th USENIX Symposium on Networked Systems Design and Implementation*, NSDI '18, pages 65–80, 2018.

[PA14] Nathaniel Popper and Rachel Abrams. Apparent theft at Mt. Gox shakes Bitcoin world. https://www.nytimes.com/2014/02/25/business/apparent-theft-at-mt-gox-shakes-bitcoin-world.html, 2014. Accessed 2018-12-27.

[par] Paradex. https://paradex.io/. Accessed 2018-12-27.

[pax] Paxos Standard. https://paxos.com/standard. Accessed 2019-01-24.

[PB17] Joseph Poon and Vitalik Buterin. Plasma: Scalable autonomous smart contracts. https://plasma.io/, 2017.

[Pro18] The Brooklyn Project. An overview of decentralized trading of digital assets. https://collaborate.thebkp.com/project/TL/document/9/version/10/, 2018. Accessed 2018-12-27.

[PS00] David Pointcheval and Jacques Stern. Security arguments for digital signatures and blind signatures. *Journal of Cryptology*, 13(3):361–396, 2000.

[PS18] Chris Peikert and Sina Shiehian. Privately constraining and programming PRFs, the LWE way. In *Proceedings of the 21st International Conference on Practice and Theory in Public-Key Cryptography*, PKC '18, pages 675–701, 2018.

[rad] Radar Relay. https://radarrelay.com/. Accessed 2018-12-27.

[RCGJ+17] Arka Rai Choudhuri, Matthew Green, Abhishek Jain, Gabriel Kaptchuk, and Ian Miers. Fairness in an unfair world: Fair multiparty computation from public bulletin boards. In *Proceedings of the 24th ACM Conference on Computer and Communications Security*, CCS '17, pages 719–728, 2017.

[Rei16] Christian Reiwießner. From smart contracts to courts with not so smart judges. https://blog.ethereum.org/2016/02/17/smart-contracts-courts-not-smart-judges/, 2016.

[RH11] Fergal Reid and Martin Harrigan. An analysis of anonymity in the Bitcoin system. In *Proceedings of the 3rd IEEE International Conference on Privacy, Security, Risk and Trust (PASSAT), and the 3rd IEEE International Conference on Social Computing (SocialCom)*, SocialCom/PASSAT '11, pages 1318–1326, 2011.

[RS13] Dorit Ron and Adi Shamir. Quantitative analysis of the full Bitcoin transaction graph. In *Proceedings of the 17th International Conference on Financial Cryptography and Data Security*, FC '13, pages 6–24, 2013.

[Sah99] Amit Sahai. Non-malleable non-interactive zero knowledge and adaptive chosen-ciphertext security. In *Proceedings of the 40th Annual Symposium on Foundations of Computer Science*, FOCS '99, pages 543–553, 1999.

[SCF+15] Felix Schuster, Manuel Costa, Cédric Fournet, Christos Gkantsidis, Marcus Peinado, Gloria Mainar-Ruiz, and Mark Russinovich. VC3: Trustworthy data analytics in the cloud using SGX. In *Proceedings of the 36th IEEE Symposium on Security and Privacy*, SP '15, pages 38–54, 2015.

[Sch91] Claus-Peter Schnorr. Efficient signature generation by smart cards. *Journal of Cryptology*, 4(3):161–174, 1991.

[SJ99] Claus-Peter Schnorr and Markus Jakobsson. Security of discrete log cryptosystems in the random oracle and the generic model. In *The Mathematics of Public-Key Cryptography*, MPKC '99, 1999.

[SMZ14] Michele Spagnuolo, Federico Maggi, and Stefano Zanero. BitIodine: Extracting intelligence from the Bitcoin network. In *Proceedings of the 18th International Conference on Financial Cryptography and Data Security*, FC '14, pages 457–468, 2014.

[SS01] Douglas R. Stinson Stinson and Reto Strobl. Provably secure distributed Schnorr signatures and a $(t, n)$ threshold scheme for implicit certificates. In *Proceedings of the 5th Australasian Conference on Information Security and Privacy*, ACISP '01, pages 417–434, 2001.

[TR17] Jason Teutsch and Christian Reiwießner. A scalable verification solution for blockchains. `https://people.cs.uchicago.edu/~teutsch/papers/truebit.pdf`, 2017.

[Val08] Paul Valiant. Incrementally verifiable computation or proofs of knowledge imply time/space efficiency. In *Proceedings of the 5th Theory of Cryptography Conference*, TCC '08, pages 1–18, 2008.

[VB15] Fabian Vogelsteller and Vitalik Buterin. ERC-20 token standard. `https://github.com/ethereum/EIPs/blob/master/EIPS/eip-20.md`, 2015.

[VBMW+19] Jo Van Bulck, Marina Minkin, Ofir Weisse, Daniel Genkin, Baris Kasikci, Frank Piessens, Mark Silberstein, Thomas F. Wenisch, Yuval Yarom, and Raoul Strackx. Breaking virtual memory protection and the SGX ecosystem with foreshadow. *IEEE Micro*, 39(3):66–74, 2019.

[Woo17] Gavin Wood. Ethereum: A secure decentralised generalised transaction ledger, 2017. `http://yellowpaper.io`.

[WSR+15] Riad S. Wahby, Srinath T. V. Setty, Zuocheng Ren, Andrew J. Blumberg, and Michael Walfish. Efficient RAM and control flow in verifiable outsourced computation. In *Proceedings of the 22nd Annual Network and Distributed System Security Symposium*, NDSS '15, 2015.

[ZCa15] ZCash Company. https://z.cash/, 2015.

[ZCa16] ZCash parameter generation. `https://z.cash/technology/paramgen.html`, 2016. Accessed: 2017-09-28.

[ZDB+17] Wenting Zheng, Ankur Dave, Jethro G. Beekman, Raluca Ada Popa, Joseph E. Gonzalez, and Ion Stoica. Opaque: An oblivious and encrypted distributed analytics platform. In *Proceedings of the 14th USENIX Symposium on Networked Systems Design and Implementation*, NSDI '17, pages 283–298, 2017.

[Zha18] Wolfie Zhao. Bithumb $31 million crypto exchange hack: What we know (and don't). `https://www.coindesk.com/bithumb-exchanges-31-million-hack-know-dont-know`, 2018. Accessed 2018-12-27.