

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.jhu.edu
Cornell Tech

Pratyush Mishra
pratyush@berkeley.edu
UC Berkeley

Howard Wu
howardwu@berkeley.edu
UC Berkeley

October 8, 2018

Abstract

Ledger-based systems that enable 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. Unfortunately, expensive re-execution and lack of privacy rule out many use cases.

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 by anyone in constant time, regardless of the offline computation.

The core of ZEXE is a protocol for a new cryptographic primitive that we introduce, *decentralized private computation* (DPC). The security guarantees of DPC are concisely expressed via an ideal functionality, which our protocol provably achieves. In order to achieve an efficient implementation of our protocol, 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 2 minutes plus a time that grows with the offline computation.

To facilitate real-world deployments, ZEXE also provides support for delegating the process of producing a transaction to an untrusted worker, and support for threshold transactions and blind transactions.

Keywords: decentralized computation; zero knowledge proofs; succinct arguments

Contents

1	Introduction	3
1.1	Our contributions	4
1.2	Related work	5
2	Techniques	7
2.1	Achieving privacy for a single arbitrary function	7
2.2	Difficulties with achieving privacy for user-defined functions	8
2.3	The records nano-kernel: a minimalist shared execution environment	8
2.4	Decentralized private computation	10
2.5	Achieving an efficient implementation	12
3	Definition of decentralized private computation schemes	14
3.1	Data structures	14
3.2	Algorithms	15
3.3	Security	16
4	Construction of decentralized private computation schemes	20
4.1	Building blocks	20
4.2	Algorithms	20
5	Delegating zero knowledge execution	24
5.1	Approach	24
5.2	Additional building block: randomizable signatures	24
5.3	A delegable DPC scheme	25
5.4	Threshold transactions and blind transactions	27
6	A simple example: user-defined assets	28
7	Implementation strategy	29
8	System implementation	33
9	System evaluation	35
9.1	Cryptographic building blocks	35
9.2	The execute NP relation	35
9.3	DPC algorithms	35
9.4	DPC data structures	36
A	Proof of security for our DPC scheme	38
A.1	Building blocks for the simulator	38
A.2	The ideal-world simulator	39
A.3	Proof of security by hybrid argument	42
B	Construction of a delegable DPC scheme	44
B.1	Definition and construction of a randomizable signature scheme	44
B.2	Construction of a delegable DPC scheme	45
C	Extensions in functionality and in security	48
	Acknowledgments	50
	References	50

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*.

The 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 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 also *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.

The re-execution problem. Public auditability in the aforementioned systems (and many others) is achieved via direct verification of state transitions. This creates two problems. First, validating a transaction involves re-executing the associated computation and so, 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. Second, even with such mechanisms, validating an expensive transaction may simply not be economically profitable, 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.

¹Even 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 [Eli13, Cha14].

²For 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 equally cheap to validate (they are all indistinguishable), there is no “Verifier’s Dilemma” nor a need for mechanisms such as *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 that we introduce, *decentralized private computation* (DPC), a new approach to performing computations on a ledger. Informally, DPC supports a simple, yet expressive, programming model in which units of data, which we call *records*, contain within them scripts (arbitrary programs) that determine under what conditions they can be first created and then consumed. 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 concisely expressed via an ideal functionality, which our protocol provably achieves.

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).

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 2 minutes 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).

³One 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.

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 provides tolerable efficiency but is by no means a lightweight construction. We have, after all, set ambitious goals: *data/function privacy and succinctness, for a 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.

In light of the foregoing ambitious goals, we have, in our opinion, managed to achieve excellent transaction sizes (less than a kilobyte) and transaction verification times (tens of milliseconds).

The main undesirable cost in our system is, not surprisingly, the cost to generate the cryptographic proofs to include in transactions. We have managed to keep this cost to under 2 minutes plus a cost that grows with the offline computation, which is similar to what prior systems have achieved for *more limited* functionalities [BCG⁺14]. But we are optimistic that this cost can be significantly reduced via more careful engineering.

1.2 Related work

Avoiding naive re-execution. 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. Informally, users encrypt payment details and prove their validity, without disclosing what the payment details are. The Zerocash protocol, with some modifications, is now commercially deployed in several currencies, including Zcash [ZCa15]. In Zerocash, however, there is no support for scripting, that is, specifying small programs that dictate how funds can be spent. Even more so, in Zerocash there is no support for complex financial logic, and more generally for programming arbitrary state transitions like in smart contract systems such as Ethereum.

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. The privacy guarantee that Hawk achieves, known as transactional privacy, protects the private inputs used in a computation (directly so from the proofs’ zero knowledge

property) but does not protect the information of *which* computation was carried out. 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 hide *which* computation was carried out offline.

MPC with ledgers. Several works [ADMM14b, ADMM14a, KMB15, KB16, BKM17] 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. Ekiden [CZK⁺18] is a ledger-based system that uses 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).

2 Techniques

We 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 begin by noting that privacy is the “harder” of the two goals, since there is a straightforward folklore approach that provides succinctness alone: 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 then simply verify the cryptographic proof, and do not have to re-execute the computation. 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 Φ (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 Φ 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 Φ . Function privacy also holds but for trivial reasons: Φ is known in advance to everybody, and every transaction is about computations of Φ .

Note that Zerocash is indeed a special case of the above: it corresponds to fixing Φ to the particular (and publicly known) choice of a function $\Phi_{\mathfrak{S}}$ that governs value transfers within a single currency.

However the foregoing protocol supports only a single hard-coded function Φ , 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, while maintaining function privacy and without seeking prior permission from anyone. That is, when preparing a transaction, a user may pick *any* function Φ of his choice for creating new records by consuming some old records.

This alone *can* be achieved via the approach sketched 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. 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, he would not know what functions were used to create those records. For 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 Φ_{\S} (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 Φ_{\S} 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 Φ_{\S} . More generally, the lack of any enforced rules about how user-defined functions can interact precludes productive cooperation between users that do not trust one another. We stress that this challenge arises specifically due to function privacy, because 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 a “good” functions.

One way to address the foregoing problem would be to augment records with an attribute that *must* equal the identity of the function that created them, and then impose the restriction that in a valid transaction only records created by the same function may participate. This new attribute is never revealed on the ledger (just like a record's payload), and 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 Φ_{\S} . 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: limited functionality. The foregoing forbids any inter-process communication, and so one cannot realize even simple functionalities like transferring value between different currencies on the same ledger. 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 interactions cannot realize even basic applications such as a single currency, while complete process segregation is too limiting.

Balancing these extremes requires a shared execution environment, namely an *operating system*, that 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. This execution environment can be viewed as a “nano-kernel” that manages records, and can be summarized as follows.

The records nano-kernel (RNK): In addition to a data payload, a record r will now contain two user-defined functions, or more precisely two user-defined *predicates* (boolean functions). These are a *birth* predicate Φ_b , which is executed when r is created, and a *death* predicate Φ_d , which is executed when r is consumed. By suitably programming r ’s birth and death predicates, a user fully dictates the conditions under which r can be first created and later consumed.

As before, a transaction in the ledger attests that some old records were consumed in order to create new records. However, now it also attests that the old records’ death predicates and the new records’ birth predicates were all simultaneously satisfied when given a certain common input. This input is the transaction’s *local data*, 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. See Fig. 3.

In this way, *each predicate can individually decide if the local data is valid according to its own logic*. For example, a record can protect itself from other records that contain “bad” birth or death predicates, because the record’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.

In Example 2.1 below we illustrate this flexibility via a simple application. More generally, “applications” arise as emergent systems on top of carefully designed sets of predicates. Records appearing in the same transaction can use the shared memory to perform joint computation and communication. In fact, computation and communication may span across multiple transactions, by storing suitable intermediate state/message data in record payloads, or by publishing that data in transactions memoranda (as plaintext or ciphertext as needed). Thus, an arbitrary computation may be either conducted within the scope of a single transaction, or instead it could take place across many transactions, e.g., by thinking of transactions as realizing state transitions, with consumed records viewed as data loaded from memory, and created records viewed as data stored back to memory. These perspectives demonstrate that the records nano-kernel is a powerful model for computation.

Example 2.1 (user-defined tokens). We briefly explain how to use the records nano-kernel to realize an application wherein users can define new tokens and then exchange them according to desired exchange rates.

We consider records whose payloads encode two pieces of information: a token identifier id and a value v . We fix the birth predicate in all such records to be a function Φ_b^* that is responsible for creating the initial supply of a new token, and then subsequently conserving the value of the token across all transactions. In more detail, Φ_b^* can be invoked in one of two modes. In *mint mode*, given as input a desired initial supply v_0 , Φ_b^* deterministically derives a fresh unique identifier id for a new token (this can be done via information unique to a transaction) and stores the pair (id, v_0) in a newly created *genesis record*. In *conserve mode*, Φ_b^* inspects all records in a transaction whose birth predicates equal to Φ_b^* and whose token identifiers equal the identifier of the current record, and ensures that among these the token values are conserved.

Users can program death predicates of records to enforce conditions on how tokens can be consumed, e.g., by realizing *conditional exchanges* with other counter-parties. Suppose that w wishes to exchange 100 units of a token id_1 for 50 units of another token id_2 , but does not have a counter-party for the exchange. She creates a record r with 100 units of id_1 whose death predicate enforces that any transaction consuming r must also create another record, consumable by Alice, with 50 units of id_1 . She then publishes out of band information about r , and anyone can subsequently claim it by creating a transaction doing the exchange.

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.

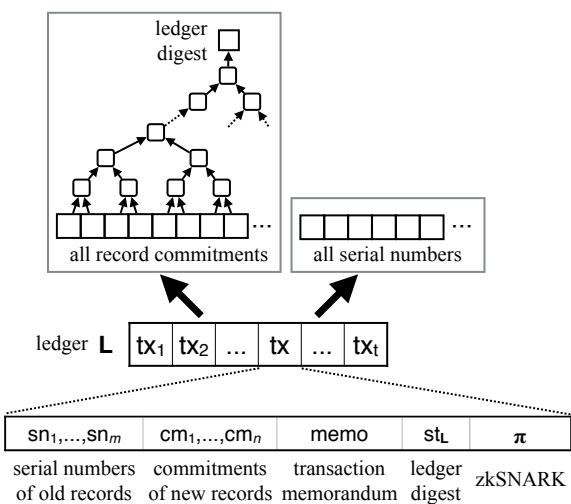


Figure 1: Construction of a transaction.

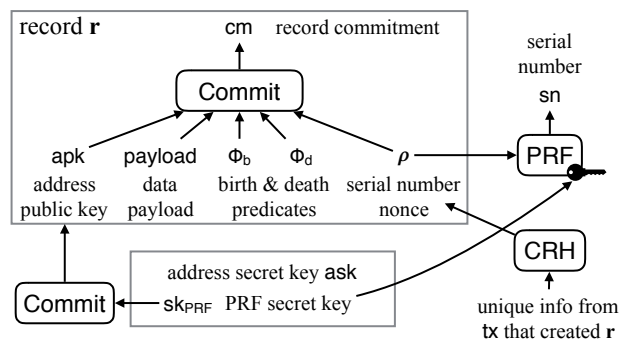


Figure 2: Construction of a record.

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).⁴

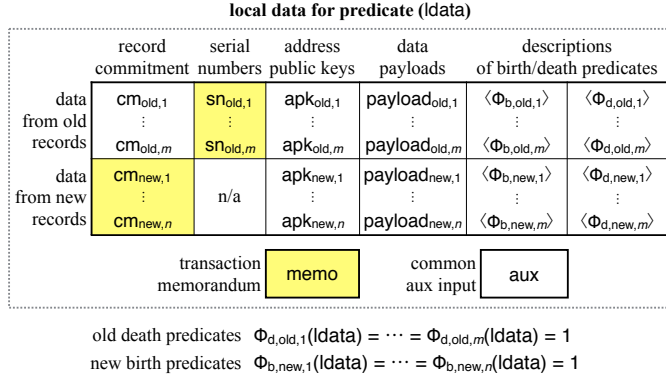


Figure 3: Predicates receive local data.

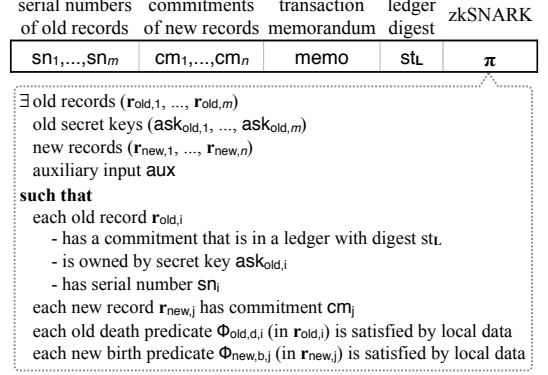


Figure 4: The execute statement.

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 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.

⁴By supporting the use of dummy records, we can in fact ensure that only *upper bounds* on the foregoing numbers are revealed.

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, CCW18].

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_{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^α 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_{CP} over a field \mathbb{F}_q such

that the curve group $E_{CP}(\mathbb{F}_q)$ contains a subgroup of prime order p (the size of E_{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_{CP} to have embedding degree of 6 in order to achieve 128 bits of security (as determined from the guidelines in [FST10]).

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

To avoid this, we split \mathcal{R}_e into two NP relations, \mathcal{R}_{BLS} and \mathcal{R}_{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_{CP} operations by proving satisfaction of \mathcal{R}_{BLS} and \mathcal{R}_{CP} with zkSNARKs over E_{BLS} and E_{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 [HBHW18]). To achieve this, we pick two Edwards curves, $E_{Ed/BLS}$ over the field \mathbb{F}_r (thereby matching the group order of E_{BLS}), and $E_{Ed/CP}$ over the field \mathbb{F}_p (thereby matching the group order of E_{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.

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

Informally, the “life” of a (non-dummy) record r is marked by two events: *birth* and *death*. The record r is *born* (or is *created*) when its commitment cm is posted to the ledger as part of a transaction. Then the record r *dies* (or is *consumed*) when its serial number sn appears on the ledger as part of a later transaction. At each of these times (birth or death) the corresponding predicate (Φ_b or Φ_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 r , one must also know the address secret key ask corresponding to r ’s address public key apk because the serial number sn to be revealed can only be computed from r and 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 sn but have distinct commitments cm and cm' .

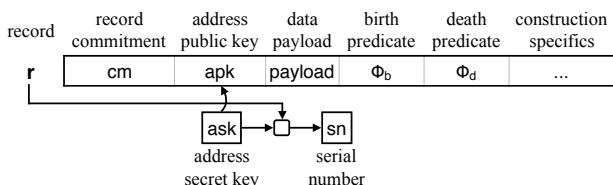


Figure 5: Diagram of a record.

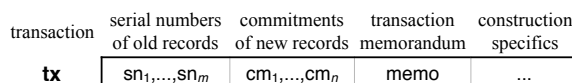


Figure 6: Diagram of a transaction.

Transactions. A transaction, denoted by the symbol tx , is a data structure representing a state change that involves the consumption and creation of records (see Fig. 6). It is a tuple $([\text{sn}_i]_1^m, [\text{cm}_j]_1^n, \text{memo}, \star)$ where (a) $[\text{sn}_i]_1^m$ is the list of serial numbers of the m old records, (b) $[\text{cm}_j]_1^n$ is the list of commitments of the n new records, (c) memo is an arbitrary string associated with the transaction, and (d) \star is other construction-specific information. The transaction 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.

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

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}):

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

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

Setup: $\text{DPC}.\text{Setup}(1^\lambda) \rightarrow \text{pp}$.

On input a security parameter 1^λ , $\text{DPC}.\text{Setup}$ outputs public parameters 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 $\text{DPC}.\text{Setup}$ algorithm (see [BCG⁺15, ZCa16, BGG17] 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: $\text{DPC.GenAddress}(\text{pp}, \text{meta}) \rightarrow (\text{apk}, \text{ask})$.

On input public parameters pp and address metadata meta , DPC.GenAddress outputs an address key pair (apk, ask) . The address metadata is bound to apk so that it is computationally infeasible to create two equal address public keys with different metadata. 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.

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

Given as input a list of old records $[\mathbf{r}_i]_1^m$ with corresponding secret keys $[\mathbf{ask}_i]_1^m$, attributes for new records, private auxiliary input aux to birth and death predicates of new and old records respectively,⁵ 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 $\mathbf{L.Push}(\text{tx})$.

Verify: $\text{DPC.Verify}^{\mathbf{L}}(\text{pp}, \text{tx}) \rightarrow b$.

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.⁶ 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 [HBHW18].)
- *Transaction non-malleability.* Malicious parties cannot modify a transaction “in flight” to the ledger.

⁵In 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.

⁶And any information implied by knowing that the birth (resp., death) predicates of consumed (resp., created) records are satisfied.

Formally, we prove *standalone* security against *static corruptions*, in a model where every party has private anonymous channels to all other parties [IKOS06].⁷ (In Appendix C we discuss how to prove security under 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 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 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}_{DPC} specified in Fig. 7 (and further described below).

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

Ideal data structures. The ideal functionality \mathcal{F}_{DPC} uses ideal counterparts of a DPC scheme’s data structures. An *address public key* apk denotes the owner of an *ideal record* \mathfrak{r} , which is a tuple $(cm, apk, \text{payload}, \Phi_b, \Phi_d)$, where cm is its commitment, apk is its address public key, payload is its payload, and Φ_b and Φ_d are its birth and death predicates. The record is also associated with a unique identifier (or *serial number*) sn . We require that apk , cm , and sn are “globally unique”; this means that there cannot be two different ideal records \mathfrak{r} and \mathfrak{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 (SampleAddrPk , SampleCm , SampleSn) that, on input a random string, sample (apk, cm, sn) respectively. When \mathcal{F}_{DPC} needs to sample one of these, it invokes the respective functions. (Note that \mathcal{F}_{DPC} cannot directly ask \mathcal{S} to sample these because that would reveal to \mathcal{S} when an honest party was invoking $\mathcal{F}_{DPC}.\text{GenAddress}$ or $\mathcal{F}_{DPC}.\text{Execute}$, and we cannot afford this leakage.)

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

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

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

- *Address generation:* $\mathcal{F}_{DPC}.\text{GenAddress}$ outputs a new address public key apk .
- *Execution:* $\mathcal{F}_{DPC}.\text{Execute}$ 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 `memo`. Concurrent $\mathcal{F}_{DPC}.\text{Execute}$ calls are serialized arbitrarily.
- *Record consumption authorization:* $\mathcal{F}_{DPC}.\text{ShareRecord}$ allows a party \mathcal{P} to authorize another party \mathcal{P}' to consume a record \mathfrak{r} (provided that \mathcal{P}' is also authorized to use \mathfrak{r} ’s address public key).

⁷Parties can, e.g., use these channels to communicate the contents of newly created records to other parties.

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 GenAddress, Execute, or ShareRecord. In the real world honest parties translate these instructions into corresponding invocations of DPC algorithms (or messages sent via private anonymous channels as in the case of ShareRecord), while in the ideal world they translate them into corresponding invocations of \mathcal{F}_{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}_{DPC} , honest parties do not provide any inputs marked as optional; instead, they let \mathcal{F}_{DPC} sample these.

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

- *Execution correctness.* \mathcal{F}_{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 $[\text{sn}_i]_1^m$ of consumed records, commitments $[\text{cm}_j]_1^n$ 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}_{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}_{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}_{DPC} .Execute.

$\mathcal{F}_{\text{DPC}}.\text{GenAddress}[\mathcal{P}]$ <p>(address metadata meta (optional) address public key apk)</p> <ol style="list-style-type: none"> 1. Sample randomness r for generating address public key. 2. If $\text{apk} = \perp$ then $\text{apk} \leftarrow \text{SampleAddrPk}(r)$. 3. Check that apk is unique: $\text{Addr}[\text{apk}] = \perp$. 4. Set $\text{Addr}[\text{apk}] := (\text{meta}, r)$. 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 $\text{AddrUsers}[\text{apk}] := \text{AddrUsers}[\text{apk}] \cup S$. 8. Send to \mathcal{P}: address public key apk. 	$\mathcal{F}_{\text{DPC}}.\text{ShareRecord}[\mathcal{P}]$ <p>(record \mathbb{r} recipient party \mathcal{P}')</p> <ol style="list-style-type: none"> 1. If $\text{Records}[\mathbb{r}.\text{cm}] \neq \perp$: <ol style="list-style-type: none"> (a) Check that $\mathcal{P} \in \text{RecUsers}[\mathbb{r}.\text{cm}]$. (b) Retrieve $((\text{cm}, \text{apk}, \text{payload}, \Phi_b, \Phi_d), r) := \text{Records}[\mathbb{r}.\text{cm}]$. 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 $\text{RecUsers}[\mathbb{r}.\text{cm}] := \text{RecUsers}[\mathbb{r}.\text{cm}] \cup S$. 5. If \mathcal{P} is honest and \mathcal{P}' isn't, Send to \mathcal{P}': $(\text{RecordAuth}, (\mathbb{r}, r))$. 6. Else, Send to \mathcal{P}': $(\text{RecordAuth}, \mathbb{r})$.
$\mathcal{F}_{\text{DPC}}.\text{Execute}[\mathcal{P}]$ <p>(old records $[\mathbb{r}_i]_1^m$ old address metadata $[\text{meta}_i]_1^m$ (optional) old serial numbers $[\text{sn}_i]_1^m$ (optional) new record commitments $[\text{cm}_j]_1^n$ new address public keys $[\text{apk}_j]_1^n$ new record payloads $[\text{payload}_j]_1^n$ new record birth predicates $[\Phi_{b,j}]_1^n$ new record death predicates $[\Phi_{d,j}]_1^n$ auxiliary predicate input aux transaction memorandum memo)</p> <ol style="list-style-type: none"> 1. For each $i \in \{1, \dots, m\}$: <ol style="list-style-type: none"> (a) Sample randomness r_i. (b) If $\text{sn}_i = \perp$ then generate serial number: $\text{sn}_i \leftarrow \text{SampleSn}(r_i)$. (c) Check that sn_i is unique: $\text{SerialNumbers}[\text{sn}_i] = \perp$. 2. For each $j \in \{1, \dots, n\}$: <ol style="list-style-type: none"> (a) Sample randomness r_j. (b) If $\text{cm}_j = \perp$ then generate commitment: $\text{cm}_j \leftarrow \text{SampleCm}(r_j)$. (c) Check that cm_j is unique: $\text{Records}[\text{cm}_j] = \perp$. (d) Construct record: $\mathbb{r}_j := (\text{cm}_j, \text{apk}_j, \text{payload}_j, \Phi_{b,j}, \Phi_{d,j})$. 3. Define the local data $\text{ldata} := ([\mathbb{r}_i]_1^m, [\text{sn}_i]_1^m, [\text{meta}_i]_1^m, [\mathbb{r}_j]_1^n, \text{aux}, \text{memo})$. 4. For each $i \in \{1, \dots, m\}$: <ol style="list-style-type: none"> (a) Parse \mathbb{r}_i as $(\text{cm}_i, \text{apk}_i, \text{payload}_i, \Phi_{b,i}, \Phi_{d,i})$. (b) Check that, for some randomness r_i, old record \mathbb{r}_i exists: $((\text{apk}_i, \text{payload}_i, \Phi_{b,i}, \Phi_{d,i}), r_i) = \text{Records}[\text{cm}_i]$. (c) Check that \mathcal{P} is authorized to use apk_i: $\mathcal{P} \in \text{AddrUsers}[\text{apk}_i]$. (d) If $\text{payload}_i.\text{isDummy} = 0$: <ol style="list-style-type: none"> i. Check that record is unconsumed: $\text{State}[\mathbb{r}_i] = \text{alive}$. ii. Check that \mathcal{P} is authorized to consume \mathbb{r}_i: $\mathcal{P} \in \text{RecUsers}[\text{cm}_i]$. iii. Check that \mathcal{P} is authorized to use apk_i: $\mathcal{P} \in \text{AddrUsers}[\text{apk}_i]$. (e) Check that death predicate is satisfied: $\Phi_{d,i}(i \text{ldata}) = 1$. (f) Mark it as consumed: $\text{State}[\text{cm}_i] := \text{dead}$. 5. For each $j \in \{1, \dots, n\}$: <ol style="list-style-type: none"> (a) Check that birth predicate is satisfied: $\Phi_{b,j}(j \text{ldata}) = 1$. (b) Insert new record \mathbb{r}_j: $\text{Records}[\text{cm}_j] := ((\text{apk}_j, \text{payload}_j, \Phi_{b,j}, \Phi_{d,j}), r_j)$. (c) Mark new record as unconsumed: $\text{State}[\text{cm}_j] := \text{alive}$. 6. Send to \mathcal{P}: $([\mathbb{r}_j]_1^n)$. 7. Send to all parties: $(\text{Execute}, [\text{sn}_i]_1^m, [\text{cm}_j]_1^n, \text{memo})$. 	

Figure 7: Ideal functionality \mathcal{F}_{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 $\text{CRH} = (\text{Setup}, \text{Eval})$ works as follows.

- *Setup*: on input a security parameter, CRH.Setup samples public parameters pp_{CRH} .
- *Hashing*: on input public parameters pp_{CRH} and message m , CRH.Eval outputs a short hash h of m .

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

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

Trapdoor commitments. A trapdoor commitment scheme $\text{TCM} = (\text{Setup}, \text{Commit})$ enables a party to generate a (perfectly) hiding and (computationally) binding commitment to a given message.

- *Setup*: on input a security parameter, TCM.Setup samples public parameters pp_{TCM} .
- *Commitment*: on input public parameters pp_{TCM} , message m , and randomness r_{cm} , TCM.Commit outputs a commitment cm to m .

Auxiliary algorithms 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 $\text{NIZK} = (\text{Setup}, \text{Prove}, \text{Verify})$ with the following syntax.

- *Setup*: on input a security parameter and the specification of an NP relation \mathcal{R} , NIZK.Setup outputs a set of public parameters pp_{NIZK} (also known as a *common reference string*).
- *Proving*: on input pp_{NIZK} and an instance-witness pair $(\mathbb{x}, \mathbb{w}) \in \mathcal{R}$, NIZK.Prove outputs a proof π .
- *Verifying*: on input pp_{NIZK} , instance \mathbb{x} , and proof π , 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 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}_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 TCM.Setup , CRH.Setup , and NIZK.Setup to obtain trapdoor commitment public parameters pp_{TCM} , CRH public parameters pp_{CRH} , and NIZK public parameters for the NP relation \mathcal{R}_e (see Fig. 9). It then outputs $\text{pp} := (\text{pp}_{\text{TCM}}, \text{pp}_{\text{CRH}}, \text{pp}_e)$.

Address creation. DPC.GenAddress , on input address metadata meta , constructs an address key pair as follows. The address secret key ask is a tuple $(\text{sk}_{\text{PRF}}, \text{meta}, r_{\text{pk}})$ consisting of a secret key sk_{PRF} for the pseudorandom function PRF, the address metadata meta , and commitment randomness r_{pk} . The address public key apk is a commitment to $\text{sk}_{\text{PRF}} \parallel \text{meta}$ with randomness r_{pk} . Thus, apk binds together sk_{PRF} and meta , while simultaneously hiding all information about them.

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.

$\text{DPC.ConstructRecord}(\text{pp}, \text{apk}, \text{payload}, \Phi_b, \Phi_d, \rho) \rightarrow (\mathbf{r}, \text{cm})$

1. Sample new commitment **randomness** r .
2. Construct new record **commitment**: $\text{cm} \leftarrow \text{TCM.Commit}(\text{pp}_{\text{TCM}}, \text{apk} \parallel \text{payload} \parallel \Phi_b, \parallel \Phi_d, \parallel \rho; r)$.
3. Construct new **record** $\mathbf{r} := \begin{pmatrix} \text{address public key} & \text{apk} & \text{payload} & \text{payload} & \text{comm. rand.} & r \\ \text{serial number nonce} & \rho & \text{predicates} & (\Phi_b, \Phi_d) & \text{commitment} & \text{cm} \end{pmatrix}$.
4. Output (\mathbf{r}, 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 sn_i ;
 - \mathbf{r}_i 's *death predicate* $\Phi_{d,i}$ *is satisfied*, demonstrated by checking that $\Phi_{d,i}(i \parallel \text{ldata}) = 1$.
- *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 ρ_j as $\text{CRH.Eval}(\text{pp}_{\text{CRH}}, j \parallel \text{sn}_1 \parallel \dots \parallel \text{sn}_m)$;
 - \mathbf{r}_j 's *birth predicate* $\Phi_{b,j}$ *is satisfied*, demonstrated by checking that $\Phi_{b,j}(j \parallel \text{ldata}) = 1$.

The serial number sn of a record \mathbf{r} relative to an address secret key $\text{ask} = (\text{sk}_{\text{PRF}}, \text{meta}, r_{\text{pk}})$ is derived by evaluating PRF at \mathbf{r} 's serial number nonce ρ with seed sk_{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.

<p>DPC.Setup <i>Input:</i> security parameter 1^λ <i>Output:</i> public parameters pp</p> <ol style="list-style-type: none"> 1. Generate trapdoor commitment parameters: $pp_{\text{TCM}} \leftarrow \text{TCM.Setup}(1^\lambda)$. 2. Generate CRH parameters: $pp_{\text{CRH}} \leftarrow \text{CRH.Setup}(1^\lambda)$. 3. Generate NIZK parameters for \mathcal{R}_e (see Figure 9): $pp_e \leftarrow \text{NIZK.Setup}(1^\lambda, \mathcal{R}_e)$. 4. Output $pp := (pp_{\text{TCM}}, pp_{\text{CRH}}, pp_e)$. 	<p>DPC.GenAddress <i>Input:</i> public parameters pp and address metadata $meta$ <i>Output:</i> address key pair (apk, ask)</p> <ol style="list-style-type: none"> 1. Sample secret key sk_{PRF} for pseudorandom function PRF. 2. Sample randomness r_{pk} for commitment scheme TCM. 3. Set address public key $apk := \text{TCM.Commit}(pp_{\text{TCM}}, sk_{\text{PRF}} \ meta; r_{\text{pk}})$. 4. Set address secret key $ask := (sk_{\text{PRF}}, meta, r_{\text{pk}})$. 5. Output (apk, ask).
<p>DPC.Execute^L <i>Input:</i></p> <ul style="list-style-type: none"> • public parameters pp • old $\left\{ \begin{array}{l} \text{records } [r_i]_1^m \\ \text{address secret keys } [ask_i]_1^m \end{array} \right.$ • new $\left\{ \begin{array}{l} \text{address public keys } [apk_j]_1^n \\ \text{record payloads } [payload_j]_1^n \\ \text{record birth predicates } [\Phi_{b,j}]_1^n \\ \text{record death predicates } [\Phi_{d,j}]_1^n \end{array} \right.$ • auxiliary predicate input aux • transaction memorandum memo <p><i>Output:</i> new records $[r_j]_1^n$ and transaction tx</p> <ol style="list-style-type: none"> 1. For each $i \in \{1, \dots, m\}$, process the i-th old record as follows: <ol style="list-style-type: none"> (a) Parse old record r_i as $\left(\begin{array}{cccc} \text{address public key} & apk_i & \text{payload} & payload_i \\ \text{serial number nonce} & \rho_i & \text{predicates} & (\Phi_{b,i}, \Phi_{d,i}) \\ \text{comm. rand.} & r_i & \text{commitment} & cm_i \end{array} \right)$. (b) If $payload_i.isDummy = 1$, set ledger membership witness $w_{L,i} := \perp$. If $payload_i.isDummy = 0$, compute ledger membership witness for commitment: $w_{L,i} \leftarrow \text{L.Prove}(cm_i)$. (c) Parse address secret key ask_i as $(sk_{\text{PRF},i}, meta_i, r_{\text{pk},i})$. (d) Compute serial number: $sn_i \leftarrow \text{PRF}_{sk_{\text{PRF},i}}(\rho_i)$. 2. For each $j \in \{1, \dots, n\}$, construct the j-th new record as follows: <ol style="list-style-type: none"> (a) Compute serial number nonce: $\rho_j := \text{CRH.Eval}(pp_{\text{CRH}}, j \ sn_1 \ \dots \ sn_m)$. (b) Construct new record: $(r_j, cm_j) \leftarrow \text{DPC.ConstructRecord}(pp_{\text{TCM}}, apk_j, payload_j, \Phi_{b,j}, \Phi_{d,j}, \rho_j)$. 3. Retrieve current ledger digest: $st_L \leftarrow \text{L.Digest}$. 4. Construct instance \mathbb{x}_e for \mathcal{R}_e: $\mathbb{x}_e := (st_L, [sn_i]_1^m, [cm_j]_1^n, memo)$. 5. Construct witness w_e for \mathcal{R}_e: $w_e := ([r_i]_1^m, [w_{L,i}]_1^m, [ask_i]_1^m, [r_j]_1^n, aux)$. 6. Generate proof for \mathcal{R}_e: $\pi_e \leftarrow \text{NIZK.Prove}(pp_e, \mathbb{x}_e, w_e)$. 7. Construct transaction: $tx := ([sn_i]_1^m, [cm_j]_1^n, memo, \star)$, where $\star := (st_L, \pi_e)$. 8. Output $([r_j]_1^n, tx)$. 	
<p>DPC.Verify^L <i>Input:</i> public parameters pp and transaction tx <i>Output:</i> decision bit b</p> <ol style="list-style-type: none"> 1. Parse tx as $([sn_i]_1^m, [cm_j]_1^n, memo, \star)$ and \star as (st_L, π_e). 2. Check that there are no duplicate serial numbers <ol style="list-style-type: none"> (a) within the transaction tx: $sn_i \neq sn_j$ for every distinct $i, j \in \{1, \dots, m\}$; (b) on the ledger: $\text{L.Contains}(sn_i) = 0$ for every $i \in \{1, \dots, m\}$. 3. Check that the ledger state is valid: $\text{L.ValidateDigest}(st_L) = 1$. 4. Construct instance for the relation \mathcal{R}_e: $\mathbb{x}_e := (st_L, [sn_i]_1^m, [cm_j]_1^n, memo)$. 5. Check proof for the relation \mathcal{R}_e: $\text{NIZK.Verify}(pp_e, \mathbb{x}_e, \pi_e) = 1$. 	

Figure 8: Construction of a DPC scheme.

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

where

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

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

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

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

Figure 9: The execute NP relation \mathcal{R}_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 [HBHW18].

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 $SIG = (\text{Setup}, \text{Keygen}, \text{Sign}, \text{Verify}, \text{RandPk}, \text{RandSig})$ 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, SIG.Setup samples public parameters pp_{SIG} .
- *Key generation*: on input public parameters pp_{SIG} , SIG.Keygen samples a key pair $(\text{pk}_{\text{SIG}}, \text{sk}_{\text{SIG}})$.
- *Message signing*: on input public parameters pp_{SIG} , secret key sk_{SIG} , and message m , SIG.Sign produces a signature σ .
- *Signature verification*: on input public parameters pp_{SIG} , public key pk_{SIG} , message m , and signature σ , SIG.Verify outputs a bit b denoting whether σ is a valid signature for m under public key pk_{SIG} .

In addition to the usual algorithms, SIG has two algorithms for randomizing public keys and signatures.

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

The signature scheme SIG must satisfy the following security properties.

- *Existential unforgeability*. Given a public key pk_{SIG} , it is infeasible to produce a forgery under pk_{SIG} or under any randomization of pk_{SIG} . This notion strengthens the standard unforgeability notion, and is similar to that of randomizable signatures in [FKM⁺16].
- *Unlinkability*. Given a public key pk_{SIG} and a tuple $(\hat{\text{pk}}_{\text{SIG}}, m, \hat{\sigma})$ where $\hat{\sigma}$ is a valid signature for m under $\hat{\text{pk}}_{\text{SIG}}$, no efficient adversary can determine if $\hat{\text{pk}}_{\text{SIG}}$ is a fresh public key and $\hat{\sigma}$ a fresh signature, or if instead $\hat{\text{pk}}_{\text{SIG}}$ is a randomization of pk_{SIG} and $\hat{\sigma}$ a randomization of a signature for pk_{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 pp_{SIG} , it is infeasible to find a public key pk_{SIG} and $r_1 \neq r_2$ such that $\text{SIG.RandPk}(\text{pp}_{\text{SIG}}, \text{pk}_{\text{SIG}}, r_1) = \text{SIG.RandPk}(\text{pp}_{\text{SIG}}, \text{pk}_{\text{SIG}}, r_2)$.

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 $(\text{pk}_{\text{SIG}}, \text{sk}_{\text{SIG}})$ for a randomizable signature scheme SIG (see Section 5.2). The public key pk_{SIG} will be embedded in the user’s public key apk and also be used to derive the serial numbers of records “owned” by apk . In contrast, the secret key sk_{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 apk was a trapdoor commitment to a secret key sk_{PRF} for a pseudorandom function PRF and the address metadata meta . Now apk is a trapdoor commitment to this same information *as well as* the public key of a key pair $(\text{pk}_{\text{SIG}}, \text{sk}_{\text{SIG}})$ for SIG . The corresponding address secret key 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 sn is now derived in a different way: while previously $\text{sn} := \text{PRF}_{\text{sk}_{\text{PRF}}}(\rho)$ now we set $\text{sn} := \text{SIG.RandPk}(\text{pp}_{\text{SIG}}, \text{pk}_{\text{SIG}}, \text{PRF}_{\text{sk}_{\text{PRF}}}(\rho))$ where ρ is the record’s serial number nonce. Namely, while before serial numbers were outputs of a pseudorandom function keyed by sk_{PRF} , now they are randomizations of the authorization public key pk_{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 pk_{SIG} are computationally indistinguishable. This follows rather directly from the fact sn , being a randomization of pk_{SIG} , does not reveal information (to efficient distinguishers) about pk_{SIG} itself.
- *No double spending*: a user cannot “spend” (i.e., consume) r in two different transactions by revealing different serial numbers because r_{SIG} (and thus sn) is generated deterministically from r . Since SIG is randomness-injective in $SIG.RandPk$, sn is (computationally) unique to 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 r and one output (new) record r' . In this case, the transaction contains a serial number sn (supposedly corresponding to r), and a commitment cm' (supposedly corresponding to r').

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

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

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

	Plain DPC	Delegable DPC
Address secret key	$(sk_{PRF}, meta, r_{pk})$	$(sk_{SIG}, sk_{PRF}, meta, r_{pk})$
Address public key	$apk := TCM.Commit \left(\begin{array}{c} PP_{TCM}, \\ sk_{PRF} meta \end{array}; r_{pk} \right)$	$apk := TCM.Commit \left(\begin{array}{c} PP_{TCM}, \\ pk_{SIG} sk_{PRF} meta \end{array}; r_{pk} \right)$
Serial number derivation	$sn \leftarrow PRF_{sk_{PRF}}(\rho)$	<ol style="list-style-type: none"> $r_{SIG} \leftarrow PRF_{sk_{PRF}}(\rho)$ $sn \leftarrow SIG.RandPk(pp_{SIG}, pk_{SIG}, r_{SIG})$
Transaction construction	$tx := ([sn_i]_1^m, [cm_j]_1^n, memo, \star)$, where $\star := (st_L, \pi_e)$.	<ol style="list-style-type: none"> Sign transaction contents: <ol style="list-style-type: none"> $\sigma_i \leftarrow SIG.Sign(pp_{SIG}, sk_{SIG,i}, \mathbb{X}_e \pi_e)$. $\hat{\sigma}_i \leftarrow SIG.RandSig(pp_{SIG}, \sigma_i, r_{SIG,i})$. $tx := ([sn_i]_1^m, [cm_j]_1^n, memo, \star)$, where $\star := (st_L, \pi_e, [\hat{\sigma}_i]_1^m)$.
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: $SIG.Verify(pp_{SIG}, tx.sn_i, \mathbb{X}_e \pi_e, \sigma_i) = 1$.

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 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 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.

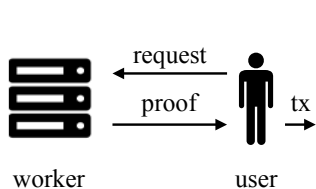


Figure 11: Delegable transactions.

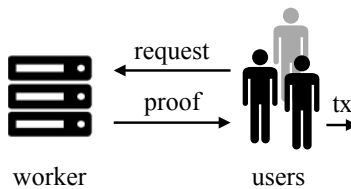


Figure 12: Threshold transactions.

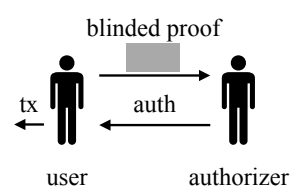


Figure 13: Blind transactions.

6 A simple example: user-defined assets

In this section we provide several example applications of ZEXE, and demonstrate how these can be “programmed” within the records nano-kernel. In this section we draw inspiration from current applications for smart contract systems such as Ethereum, which are largely focused on financial applications. We note that in these settings, privacy of contract inputs can be particularly valuable, given the monetary advantage that information asymmetry provides traders.

Payments and 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).

In more detail, the birth predicate Φ_b^* can be invoked in two modes, *mint mode* or *conserve mode*. When invoked in mint mode, Φ_b^* creates the initial supply v of the asset in, say, a single output record, by deterministically deriving a fresh unique identifier id for the asset,⁸ and storing the pair (id, v) in the record’s payload. The predicate Φ_b^* also ensures that in the given transaction there are no input records or other output records (dummy records are allowed). If Φ_b^* is invoked in mint mode in other transactions, a *different* identifier id is created, ensuring that multiple assets can be reliably distinguished even though anyone can run Φ_b^* .

When invoked in conserve mode, Φ_b^* inspects all records in a transaction whose birth predicates all equal Φ_b^* (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 the 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.

Custom access to assets. We can further refine the above application by noting that users can program death predicates to dictate access controls to their funds. This akin to how in Bitcoin scripts specify conditions on how funds can be spent (except that in Bitcoin scripts have limited expressivity). Access control mechanisms could vary from saying that the transaction must be authorized by two of three public keys, is valid only after a given amount of time, or must reveal the pre-image of a hash function. For example, this can be used to create so-called hash-timelock transactions that enable payment channels or enforce that a transaction must be authorized by multiple parties. We work out a concrete example next.

Conditional exchange of assets. In this application we want to enable users to exchange user-defined assets. For example, suppose that Alice owns 100 units of an asset with identifier id_1 , and wishes to exchange them for 50 units of an asset with identifier id_2 , but does not have a counter-party for the exchange. She could create a record r containing 100 units of id_1 that is consumable (i.e. payable to) anyone provided that it is done in a transaction that pays Alice 50 units of id_2 . She can do so by simply setting the death predicate to require that the transaction consuming r also creates a record r' belonging to Alice for 50 units of asset id_2 . Alice then publishes a transaction creating r and out-of-band announces the offer along with the details needed to claim it (including the contents of the record and the secret of of the address owning it). Anyone with this knowledge can take on Alice’s offer by creating a transaction that: (a) consumes the 100 units of id_1 in r from Alice; (b) consumes 50 units of id_2 in some other records; (c) creates 50 units of id_2 in r' for Alice; (d) creates 100 units of id_1 in some other records.

⁸One can securely generate this unique identifier by setting it to be the serial number of a dummy record that is consumed during the creation of the initial supply; since serial numbers are globally unique, so is the identifier.

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}_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}_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}_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}_e . Crucially, since the “outer” succinct proofs produced for \mathcal{R}_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 DPC.Execute to additionally take as input SNARK proofs π_d and π_b , and also modify the NP relation \mathcal{R}_e so that, instead of directly checking that Φ_d and Φ_b are satisfied, it instead checks that π_d and π_b *attest to the satisfaction of Φ_d and Φ_b* . That is, \mathcal{R}_e checks that $\text{NIZK.Verify}(\text{pp}_{\Phi_d}, \mathbb{x}_e, \pi_d) = 1$ and $\text{NIZK.Verify}(\text{pp}_{\Phi_b}, \mathbb{x}_e, \pi_b) = 1$, where pp_{Φ_d} are public parameters for the NP relation $\mathcal{R}_{\Phi_d} := \{(\mathbb{x}_e, \mathbb{w}_e) \text{ s.t. } \Phi_d(\mathbb{x}_e, \mathbb{w}_e) = 1\}$ and similarly for Φ_b . The public parameters pp_{Φ_d} and pp_{Φ_b} are stored in the record, in place of (a description of) the predicates.⁹

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

In sum, \mathcal{R}_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 Φ is satisfied will invoke a SNARK on an NP statement of size $|\Phi|$ (tailored for Φ).¹⁰ The approach described so far, however, hides additional costs that we need to overcome.

⁹More precisely, to verify a proof for a predicate Φ , the proof verifier does not need to read all of pp_Φ , which has size $O_\lambda(|\Phi|)$ in some zkSNARKs (i.e., it is large). Rather, the proof verifier only needs to read $O_\lambda(|\mathbb{x}_e|)$ bits of pp_Φ , which are collectively known as the *verification key*. The record would then store this verification key (or a hash thereof) rather than pp_Φ .

¹⁰An additional benefit of each predicate Φ having its own public parameters pp_Φ 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.

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}_e via arithmetic circuits over \mathbb{F}_r . In turn, since the SNARK verifier is part of \mathcal{R}_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, CCW18].

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) and, in particular, we can use the Cocks–Pinch method [FST10] to set up a bounded recursion [BCTV17]. We now elaborate on this.

First we pick a pairing-friendly elliptic curve 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 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_{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^α for $\alpha \geq 40$, by choosing the parameter x of the BLS family to satisfy $x \equiv 1 \pmod{3 \cdot 2^\alpha}$; 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^α . This also ensures that $x \equiv 1 \pmod{3}$, which ensures that there are efficient tower options for the relevant fields [Cos12].

Next we use the Cocks–Pinch method to pick a pairing-friendly elliptic curve E_{CP} over a field \mathbb{F}_q such that the curve group $E_{\text{CP}}(\mathbb{F}_q)$ contains a subgroup of prime order p (the size of E_{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_{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_{BLS} , and proofs for the NP relation \mathcal{R}_e are produced via a zkSNARK over E_{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_{CP} was chosen to facilitate efficient checking of proofs over E_{BLS} , the curve E_{CP} is at least $2\times$ more expensive (in time and space) than E_{BLS} simply because E_{CP} ’s base field has about twice as many bits as E_{BLS} ’s base field. Checks in the NP relation \mathcal{R}_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}_e into two NP relations \mathcal{R}_{BLS} and \mathcal{R}_{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_{BLS} (an efficient curve) to produce proofs for \mathcal{R}_{BLS} , and a zkSNARK over E_{CP} (the less efficient curve) to produce proofs for \mathcal{R}_{CP} . This approach significantly reduces the running time of DPC.Execute (producing proofs for the checks in \mathcal{R}_{BLS} is more efficient over E_{BLS} than over E_{CP}), at the expense of a modest increase in transaction size (a transaction now includes a zkSNARK proof over E_{BLS} in addition to a proof over E_{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 if we moved all checks not tied to the curve E_{CP} to the more efficient curve E_{BLS} , the NP relation \mathcal{R}_e (more precisely, its split into \mathcal{R}_{BLS} and \mathcal{R}_{CP}) collectively involves many checks that range from verifying authentication paths in a Merkle tree to verifying commitment openings, and from evaluating pseudorandom functions to evaluating collision resistant functions. NP relations with similar structure, such as those used in Zerocash [BCG⁺14], required upwards of *four million gates* to express all of these 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).

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 [HBHW18]). To do this, we pick two Edwards curves, $E_{\text{Ed}/\text{BLS}}$ over the field \mathbb{F}_r (matching the group order of E_{BLS}) and $E_{\text{Ed}/\text{CP}}$ over the field \mathbb{F}_p (matching the group order of E_{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 $E_{\text{Ed}/\text{BLS}}$ and $E_{\text{Ed}/\text{CP}}$ 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}_e into two NP relations \mathcal{R}_{BLS} and \mathcal{R}_{CP} relies on sharing secret information via commitments across NP statements; namely, a commitment cm_ϕ to the identities of predicates and a commitment cm_{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}_{BLS} is over \mathbb{F}_r , while \mathcal{R}_{CP} is over \mathbb{F}_p . For example, if we used a Pedersen commitment over the order- r subgroup of the Edwards curve $E_{\text{Ed}/\text{BLS}}$, then: (a) opening a commitment in \mathcal{R}_{BLS} would be cheap, but (b) opening a commitment in \mathcal{R}_{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 $E_{\text{Ed}/\text{CP}}$.) 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}_{BLS} and \mathcal{R}_{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 cm_ϕ in \mathcal{R}_{BLS} and \mathcal{R}_{CP} , as cm_ϕ contains hashes that are much smaller than verification keys. We realize cm_ϕ via Blake2s, a boolean primitive of modest cost in \mathbb{F}_r and \mathbb{F}_p . Crucially, *only* \mathcal{R}_{CP} needs to access the verification keys themselves, so we can efficiently use a Pedersen hash over the Edwards curve $E_{\text{Ed}/\text{CP}}$ to let \mathcal{R}_{CP} check the keys (supplied as non-deterministic advice) against the hashes inside cm_ϕ .

We realize the local data commitment cm_{ldata} via a Pedersen commitment over $E_{\text{Ed}/\text{BLS}}$, and assume that predicates take cm_{ldata} as input rather than local data in the clear. Since both \mathcal{R}_{BLS} and the predicate relations are defined over the field \mathbb{F}_r (the prime-order subgroup of the curve E_{BLS}), non-deterministically opening cm_{ldata} is efficient in both relations. This approach significantly reduces costs because \mathcal{R}_{CP} no longer needs to reason about the contents of cm_{ldata} , and can simply pass cm_{ldata} as input to the SNARK verifier.

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

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

where

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

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

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

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

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

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

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

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

Figure 14: Splitting the NP relation \mathcal{R}_e into two NP relations \mathcal{R}_{BLS} and \mathcal{R}_{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.

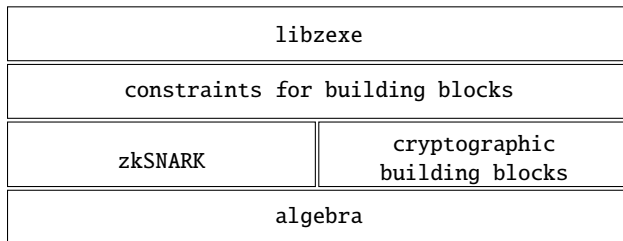


Figure 15: Stack of libraries comprising ZEXE.

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 CRH described below. This results in the following algorithms for \mathbf{L} .

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

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 PRF using the Blake2s hash function [ANWW13]: $\text{PRF}_k(x) := \text{b2s}(k||x)$ for $k, x \in \{0, 1\}^{256}$.

Elliptic curves. Our implementation strategy (see Section 7) involves several elliptic curves: two pairing-friendly curves E_{BLS} and E_{CP} , and two “plain” curves $E_{\text{Ed/BLS}}$ and $E_{\text{Ed/CP}}$ whose base field respectively matches the prime-order subgroup of E_{BLS} and E_{CP} . Details about these curves are in Figure 16; the parameter used to generate the BLS curve E_{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}_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_{BLS} (for proving predicates’ satisfiability) and E_{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 `SampleGrp`, which on input a security parameter λ (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
$E_{\text{Ed}/\text{BLS}}$	twisted Edwards	—	s	r	32	—
E_{BLS}	BLS	12	r	p	48	96
$E_{\text{Ed}/\text{CP}}$	twisted Edwards	—	t	p	48	—
E_{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	0x1ae3a4617c510eac63b05c06ca1493b1a22d9f300f5138f1ef3622fba094800170b5d44300000008508c0000000001	377	46
q	0x3848c4d2263babf8941fe959283d8f526663bc5d176b746af0266a7223ee72023d07830c728d80f9d78bab3596c8617c579252a3fb77c79c13201ad533049cfe6a399c2f764a12c4024bee135c065f4d26b7545d85c16dfd424adace79b57b942ae9	782	3

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

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 $E_{\text{Ed}/\text{BLS}}$ or $E_{\text{Ed}/\text{CP}}$, depending on the context.

Trapdoor commitments. We instantiate 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 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. [HBHW18] 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 TCM and CRH by projecting their output to its x -coordinate.

TCM.Setup($1^\lambda, n$) \rightarrow pp _{TCM} : 1. Sample a group: $(\mathbb{G}, q, g) \leftarrow \text{SampleGrp}(1^\lambda)$. 2. For $i \in \{1, \dots, n\}$, sample generator h_i : $r_i \leftarrow \mathbb{Z}_q; h_i := g^{r_i}$. 3. Output pp _{TCM} := $(\mathbb{G}, q, g, [h_i]_1^n)$.	CRH.Setup($1^\lambda, n$) \rightarrow pp _{CRH} : 1. Sample a group: $(\mathbb{G}, q, g_1) \leftarrow \text{SampleGrp}(1^\lambda)$. 2. For $i \in \{2, \dots, n\}$, sample generator g_i : $r_i \leftarrow \mathbb{Z}_q; g_i := g^{r_i}$. 3. Output pp _{CRH} := $(\mathbb{G}, q, [g_i]_1^n)$.
TCM.Commit(pp _{TCM} , $m \in \{0, 1\}^n; r_{\text{cm}}$) \rightarrow cm: 1. Parse pp _{TCM} as $(\mathbb{G}, q, g, [h_i]_1^n)$. 2. Output cm := $g^{r_{\text{cm}}} \prod_{i=1}^n h_i^{m_i}$.	CRH.Eval(pp _{CRH} , $m \in \{0, 1\}^n$) \rightarrow h : 1. Parse pp _{CRH} as $(\mathbb{G}, q, [g_i]_1^n)$. 2. Output $h := \prod_{i=1}^n g_i^{m_i}$.

Figure 17: Pedersen commitment scheme.

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_{BLS} , the GM17 prover takes $25 \mu\text{s}$ per constraint (with 12 threads), while the GM17 verifier takes $250 n \mu\text{s} + 9.5 \text{ms}$ on an input with n field elements (with 1 thread). When instantiated over the elliptic curve E_{CP} , the respective prover and verifier costs are $147 \mu\text{s}$ per constraint and $1.6 n \text{ms} + 34 \text{ms}$.

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

Table 3 reports the number of constraints that we use to express \mathcal{R}_{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}_{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.

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. 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 SNARK proofs checked by the NP relation. Producing each of these proofs requires invoking the

GM17 prover, over the elliptic curve E_{BLS} , for the relevant birth or death predicate; we provide the amortized time per constraint for this in Section 9.1.

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 2 minutes, and Verify takes tens of milliseconds.

9.4 DPC data structures

Addresses. An address public key in a DPC scheme is a point on the elliptic curve $E_{\text{Ed/BLS}}$, which is 32 bytes when compressed (see Fig. 16); the corresponding secret key is 96 bytes and consists of a PRF seed (32 bytes), address metadata (32 bytes), and commitment randomness (32 bytes). In a delegable DPC scheme, address public keys do not change, but address secret keys are 128 bytes, because they additionally contain the 32-byte secret key of a randomizable signature scheme over the elliptic curve $E_{\text{Ed/BLS}}$ (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: π_{BLS} , over the elliptic curve E_{BLS} , and π_{CP} , over the curve E_{CP} . Each proof consists of two \mathbb{G}_1 and one \mathbb{G}_2 elements from its respective curve, amounting to 192 bytes for π_{BLS} and 520 for π_{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		Plain DPC	Delegable DPC
Setup	188.63 s	187.76 s	2 inputs and 2 outputs	968	1096
GenAddress	14 ms	24 ms	m inputs and n outputs	$32m + 32n + 840$	$96m + 32n + 840$
Execute	86.9 s	87.02 s	Per input record:		
Verify	46 ms	69 ms	Serial number	32	32
			Signature	—	64
			Per output record:		
			Commitment	32	32
			Memorandum	32	32
			zkSNARK proof over E_{CP}	520	520
			zkSNARK proof over E_{BLS}	192	192
			Predicate commitment	32	32
			Local data commitment	32	32
			Ledger digest	32	32

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

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

		Plain DPC	Delegable DPC
Total with 2 inputs and 2 outputs		454122	467737
Below we provide a breakdown of the number of constraints with m input and n output records.			
Per input record	Total	141045	147211
	Enforce validity of:		
	Merkle tree path	98208	98208
	Address key pair	6663	9740
	Serial number computation	22301	25390
	Record commitment	13873	13873
Per output record	Total	20828	20828
	Enforce validity of:		
	Serial number nonce	6697	6697
	Record commitment	14131	14131
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	$10240 \cdot m + 7680 \cdot n$	$10240 \cdot m + 7680 \cdot n$
	Miscellaneous	7368	8651

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

		Plain DPC	Delegable DPC
Total with 2 inputs and 2 outputs		556678	556930
Below we provide a breakdown of the number of constraints with m input and n output records.			
Per input record	Total	116902	116902
	Enforce validity of:		
	Death predicate ver. key	56797	56797
	Death predicate proof	60105	60105
Per output record	Total	116902	116902
	Enforce validity of:		
	Birth predicate ver. key	56797	56797
	Birth predicate proof	60105	60105
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	1902	2154

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

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 (SimSetup, Equivocate) that enable one to open a commitment cm to any chosen message. Below we restrict cm to be a commitment to the empty string ε because this is sufficient for the proof of security of our DPC scheme.

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

$\text{TCM.SimSetup}(1^\lambda, n) \rightarrow (\text{pp}_{\text{TCM}}, \text{td}_{\text{TCM}})$ 1. Sample a group: $(\mathbb{G}, q, g) \leftarrow \text{SampleGrp}(1^\lambda)$. 2. For $i \in \{1, \dots, n\}$: sample r_i uniformly from \mathbb{Z}_q , and set $h_i := g^{r_i}$. 3. Output $(\text{pp}_{\text{TCM}} := (\mathbb{G}, q, g, [h_i]_1^n), \text{td}_{\text{TCM}} := [r_i]_1^n)$.	$\text{TCM.Equivocate}(\text{pp}_{\text{TCM}}, \text{td}_{\text{TCM}}, \text{cm}, r_{\text{cm}}, m' \in \{0, 1\}^n) \rightarrow r'_{\text{cm}}$ 1. Parse pp_{TCM} as $(\mathbb{G}, q, g, [h_i]_1^n)$. 2. Parse td_{TCM} as $[r_i]_1^n$. 3. Output $r'_{\text{cm}} := r_{\text{cm}} - \sum_{i=1}^n r_i m'_i \bmod q$.
---	--

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

NIZKs. The scheme $\text{NIZK} = (\text{Setup}, \text{Prove}, \text{Verify})$ 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} , NIZK.SimSetup outputs a set of public parameters pp_{NIZK} and a trapdoor td_{NIZK} .
- *Simulation:* on input public parameters pp_{NIZK} , trapdoor td_{NIZK} , NP instance \mathfrak{x} , and (optionally) auxiliary information aux , NIZK.Simulate outputs a simulated proof π .
- *Extraction:* on input public parameters pp_{NIZK} , trapdoor td_{NIZK} , NP instance \mathfrak{x} , and proof π , NIZK.Extract outputs a witness \mathfrak{w} such that $(\mathfrak{x}, \mathfrak{w}) \in \mathcal{R}$ (allegedly).

We can now state the properties satisfied by NIZK.

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

$$\Pr \left[\text{NIZK.Verify}(\text{pp}_{\text{NIZK}}, \mathfrak{x}, \pi) = 1 \mid \begin{array}{l} \text{pp}_{\text{NIZK}} \leftarrow \text{NIZK.Setup}(1^\lambda, \mathcal{R}) \\ (\mathfrak{x}, \pi) \leftarrow \text{NIZK.Prove}(\text{pp}_{\text{NIZK}}, \mathfrak{x}, \mathfrak{w}) \end{array} \right] = 1 .$$

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

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

where the two oracles are defined as follows

- $S_1(\mathbf{x}, \mathbf{w}) :=$ “if $(\mathbf{x}, \mathbf{w}) \in \mathcal{R}$ then $\text{NIZK.Prove}(\text{pp}_{\text{NIZK}}, \mathbf{x}, \mathbf{w})$, else abort”;
- $S_2(\mathbf{x}, \mathbf{w}) :=$ “if $(\mathbf{x}, \mathbf{w}) \in \mathcal{R}$ then $\text{NIZK.Simulate}(\text{pp}_{\text{NIZK}}, \text{td}_{\text{NIZK}}, \mathbf{x})$, else abort”.
- *Simulation extractability*: for every relation \mathcal{R} and efficient adversary \mathcal{A} ,

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

where $S(\mathbf{x}) := \text{NIZK.Simulate}(\text{pp}_{\text{NIZK}}, \text{td}_{\text{NIZK}}, \mathbf{x})$ 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-world simulator \mathcal{S} will interact with the ideal functionality \mathcal{F}_{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 $\mathcal{S}.\text{Records}$ that maps record commitments to their contents.
2. Initialize an empty table $\mathcal{S}.\text{AddrPk}$ that maps address public keys to their secret keys.
3. Initialize an empty transaction ledger \mathbf{L} .
4. Sample simulated public parameters and trapdoor: $(\text{pp}, \text{td}) \leftarrow \text{DPC.SimSetup}(1^\lambda)$. (See Fig. 20.)
5. Define

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

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

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 $\mathbf{L}.\text{Push}(\text{tx})$:
 1. If $\text{DPC.Verify}^{\mathbf{L}}(\text{pp}, \text{tx}) \neq 1$, abort.
 2. Parse the real-world transaction tx as $([\text{sn}_i]_1^m, [\text{cm}_j]_1^n, \text{memo}, \star)$.
 3. Compute $([\mathbf{r}_i]_1^m, [\text{ask}_i]_1^m, [\mathbf{r}_j]_1^n, \text{aux}) \leftarrow \text{DPC.ExtractExecute}(\text{pp}, \text{td}, \text{tx})$. [See Figure 20.]
 4. For every $i \in \{1, \dots, m\}$:

- (a) Parse the real-world record \mathbf{r}_i as $(\text{apk}_i, \text{payload}_i, \Phi_{\text{b},i}, \Phi_{\text{d},i}, \rho_i, r_i, \text{cm}_i)$.
 - (b) Parse the address secret key ask_i as $(\text{sk}_{\text{PRF},i}, \text{meta}_i, r_{\text{pk},i})$.
 - (c) If $\mathcal{S}.\text{Records}[\text{cm}_i] \neq \mathbf{r}_i$, abort. (**Note:** Captures binding property of the commitment.)
 - (d) If $\mathbf{L}.\text{Contains}(\text{cm}_i) = 0$, abort. (**Note:** Captures existence of record.)
 - (e) Create the ideal-world record $\mathbb{r}_i := (\text{cm}_i, \text{apk}_i, \text{payload}_i, \Phi_{\text{b},i}, \Phi_{\text{d},i})$.
 - (f) If $\mathcal{S}.\text{AddrPk}[\text{apk}_i] = \perp$:
 - i. Invoke $\mathcal{F}_{\text{DPC}}.\text{GenAddress}(\text{meta}_i, \text{apk}_i)$.
 - ii. Insert apk_i into $\mathcal{S}.\text{AddrPk}$: $\mathcal{S}.\text{AddrPk}[\text{apk}_i] := \text{ask}_i$.
 - (g) Else, if $\mathcal{S}.\text{AddrPk}[\text{apk}_i] \neq \text{ask}_i$, abort. (**Note:** Captures uniqueness of secret key.)
5. For every $j \in \{1, \dots, n\}$:
 - (a) Parse the real-world record \mathbf{r}_j as $(\text{apk}_j, \text{payload}_j, \Phi_{\text{b},j}, \Phi_{\text{d},j}, \rho_j, r_j, \text{cm}_j)$.
 - (b) If the serial number nonce ρ_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 $\mathcal{S}.\text{Records}[\text{cm}_j] := \mathbf{r}_j$.
 6. Construct instance for \mathcal{R}_e : $\mathbb{x}_e := (\text{st}_{\mathbf{L}}, [\text{sn}_i]_1^m, [\text{cm}_j]_1^n, \text{memo})$.
 7. Construct witness for \mathcal{R}_e : $\mathbb{w}_e := ([\mathbf{r}_i]_1^m, [\mathbb{w}_{\mathbf{L},i}]_1^m, [\text{ask}_i]_1^m, [\mathbf{r}_j]_1^n, \text{aux})$.
 8. If $(\mathbb{x}_e, \mathbb{w}_e) \notin \mathcal{R}_e$, abort.
 9. Invoke $\mathcal{F}_{\text{DPC}}.\text{Execute}([\mathbb{r}_i]_1^m, [\text{meta}_i]_1^m, [\text{sn}_i]_1^m, [\text{cm}_j]_1^n, [\text{apk}_j]_1^n, [\text{payload}_j]_1^n, [\Phi_{\text{b},j}]_1^n, [\Phi_{\text{d},j}]_1^n, \text{aux}, \text{memo})$.
 10. **Receive from** \mathcal{F}_{DPC} : $[\mathbb{r}_j]_1^n$.
 11. **Receive from** \mathcal{F}_{DPC} : $(\text{Execute}, [\text{sn}_i]_1^m, [\text{cm}_j]_1^n, \text{memo})$.
 12. Append the real-world transaction tx to the ledger \mathbf{L} .
- **From ideal functionality.** When \mathcal{F}_{DPC} broadcasts $(\text{Execute}, [\text{sn}_i]_1^m, [\text{cm}_j]_1^n, \text{memo})$:
 1. Compute $([\mathbf{r}_j]_1^n, \text{tx}) \leftarrow \text{DPC}.\text{SimExecute}^{\mathbf{L}}(\text{pp}, \text{td}, [\text{sn}_i]_1^m, [\text{cm}_j]_1^n, \text{memo})$. (See Fig. 20.)
 2. For each $j \in \{1, \dots, n\}$, set $\mathcal{S}.\text{Records}[\text{cm}_j] := \mathbf{r}_j$.
 3. Append the real-world transaction tx to the ledger \mathbf{L} .

Record authorization notification.

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

<p>DPC.SimSetup</p> <p><i>Input:</i> security parameter 1^λ</p> <p><i>Output:</i> simulated public parameters pp and trapdoor td</p> <ol style="list-style-type: none"> 1. Sample simulated parameters for trapdoor commitment: $(pp_{\text{TCM}}, td_{\text{TCM}}) \leftarrow \text{TCM.SimSetup}(1^\lambda)$. 2. Sample parameters for CRH: $pp_{\text{CRH}} \leftarrow \text{CRH.Setup}(1^\lambda)$. 3. Sample simulated parameters for NIZK for \mathcal{R}_e: $(pp_e, td_e) \leftarrow \text{NIZK.SimSetup}(1^\lambda, \mathcal{R}_e)$. 4. Set $pp := (pp_{\text{TCM}}, pp_{\text{CRH}}, pp_e)$. 5. Set $td := (td_{\text{TCM}}, td_e)$. 6. Output (pp, td).
<p>DPC.SimExecute^L</p> <p><i>Input:</i></p> <ul style="list-style-type: none"> • public parameters pp and trapdoor td • old serial numbers $[sn_i]_1^m$ • new record commitments $[cm_j]_1^n$ • transaction memorandum memo <p><i>Output:</i> new records $[r_j]_1^n$ and transaction tx</p> <ol style="list-style-type: none"> 1. For $j \in \{1, \dots, n\}$: <ol style="list-style-type: none"> (a) Set new serial number nonce $\rho_j := \text{CRH.Eval}(pp_{\text{CRH}}, j \ sn_1 \ \dots \ sn_m)$. (b) Set address public key, payload, predicates, and commitment randomness to be the empty string: $apk_j, payload_j, \Phi_{b,j}, \Phi_{d,j}, r_j := \varepsilon$. (c) Construct dummy record: $r_j := (apk_j, payload_j, \Phi_{b,j}, \Phi_{d,j}, \rho_j, r_j, cm_j)$. 2. Retrieve current ledger digest: $st_L \leftarrow \mathbf{L.Digest}$. 3. Construct instance for relation \mathcal{R}_e: $x_e := (st_L, [sn_i]_1^m, [cm_j]_1^n, memo)$. 4. Generate simulated proof for \mathcal{R}_e: $\pi_e \leftarrow \text{NIZK.Simulate}(pp_e, td_e, x_e)$. 5. Construct transaction: $tx := ([sn_i]_1^m, [cm_j]_1^n, memo, \star)$, where $\star := (st_L, \pi_e)$. 6. Output $([r_j]_1^n, tx)$.
<p>DPC.ExtractExecute</p> <p><i>Input:</i></p> <ul style="list-style-type: none"> • public parameters pp and trapdoor td • transaction tx <p><i>Output:</i></p> <ul style="list-style-type: none"> • old $\begin{cases} \text{records } [r_i]_1^m \\ \text{address secret keys } [ask_i]_1^m \end{cases}$ • new records $[r_j]_1^n$ • auxiliary predicate input aux <ol style="list-style-type: none"> 1. Parse tx as $([sn_i]_1^m, [cm_j]_1^n, memo, \star)$ and \star as (st_L, π_e). 2. Construct instance for relation \mathcal{R}_e: $x_e := (st_L, [sn_i]_1^m, [cm_j]_1^n, memo)$. 3. Obtain witness: $w_e \leftarrow \text{NIZK.Extract}(pp_e, td_e, x_e, \pi_e)$. 4. Parse the witness w_e as $([r_i]_1^m, [w_{L,i}]_1^m, [ask_i]_1^m, [r_j]_1^n, aux)$. 5. Output $([r_i]_1^m, [ask_i]_1^m, [r_j]_1^n, aux)$.

Figure 20: Several subroutines used by the ideal-world simulator 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 $\text{Output}_i(\mathcal{E})$ 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} .

$\text{Output}_1(\mathcal{E})$ is perfectly indistinguishable from $\text{Output}_0(\mathcal{E})$ 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 . $\text{Output}_2(\mathcal{E})$ is perfectly indistinguishable from $\text{Output}_1(\mathcal{E})$ 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 . $\text{Output}_3(\mathcal{E})$ is perfectly indistinguishable from $\text{Output}_2(\mathcal{E})$ 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, $\text{Output}_4(\mathcal{E})$ is computationally indistinguishable from $\text{Output}_3(\mathcal{E})$.

- \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 ε . 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 . $\text{Output}_5(\mathcal{E})$ is perfectly indistinguishable from $\text{Output}_4(\mathcal{E})$ 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 \mathbf{x}_e and witness \mathbf{w}_e for \mathcal{R}_e from the included proof and then proceeds as follows.

- If $(\mathbf{x}_e, \mathbf{w}_e) \notin \mathcal{R}$, \mathcal{S} aborts. If NIZK is simulation-extractable, this occurs with negligible probability.
- For all $i \in \{1, \dots, m\}$, 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 $i \in \{1, \dots, m\}$, if the extracted secret key ask_i for apk_i differs from the secret key extracted for apk_i in a prior transaction, \mathcal{S} aborts. If TCM is a binding commitment scheme, then this occurs with negligible probability.

- For all $j \in \{1, \dots, n\}$, if the serial number nonce ρ_j matches one extracted in a prior transaction, \mathcal{S} aborts. If CRH is a collision-resistant hash, then this occurs with negligible probability because the serial number nonce is the output of CRH evaluated (in part) over the serial numbers of the input records. If this input is distinct across two different invocations of 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 CRH are never repeated.

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

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 $\text{SIG} = (\text{Setup}, \text{Keygen}, \text{Sign}, \text{Verify}, \text{RandPk}, \text{RandSig})$ 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 SIG satisfies the following security properties.

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

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

where $S(m) := \text{SIG.Sign}(\text{pp}_{\text{SIG}}, \text{sk}_{\text{SIG}}, m)$ 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.

IND-RSIG $_{\mathcal{A}}^{\text{SIG}}(1^\lambda)$:

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

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

$$\Pr \left[\begin{array}{l} r_1 \neq r_2 \\ \text{SIG.RandPk}(\text{pp}_{\text{SIG}}, \text{pk}_{\text{SIG}}, r_1) = \text{SIG.RandPk}(\text{pp}_{\text{SIG}}, \text{pk}_{\text{SIG}}, r_2) \end{array} \middle| \begin{array}{l} \text{pp}_{\text{SIG}} \leftarrow \text{SIG.Setup}(1^\lambda) \\ (\text{pk}_{\text{SIG}}, r_1, r_2) \leftarrow \mathcal{A}(\text{pp}_{\text{SIG}}) \end{array} \right].$$

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 $(m^*, \sigma^*, r_{\text{SIG}}^*)$, \mathcal{A}' outputs the tuple (m^*, σ) where σ is computed as follows. If σ^* is a valid signature for m^* under pk_{SIG} then $\sigma := \sigma^*$. Otherwise, \mathcal{A}' “undoes” the randomization of $\sigma^* = (s, e)$ by setting $\sigma := (s + e \cdot r_{\text{SIG}}^*, e)$; thus if \mathcal{A} outputs a forgery for a randomization of pk_{SIG} , \mathcal{A}' translates it back into a forgery for pk_{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_{SIG} , SIG.RandPk is a permutation over \mathbb{G} . Hence, finding collisions over the randomness is not possible.

$\text{SIG.Setup}(1^\lambda) \rightarrow \text{pp}_{\text{SIG}}$ 1. Sample a group: $(\mathbb{G}, q, g) \leftarrow \text{SampleGrp}(1^\lambda)$. 2. Sample cryptographic hash function H . 3. Output $\text{pp}_{\text{SIG}} := (\mathbb{G}, q, g, H)$.	$\text{SIG.Sign}(\text{pp}_{\text{SIG}}, \text{sk}_{\text{SIG}}, m) \rightarrow \sigma$ 1. Parse pp_{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 $\sigma := (s, e)$.
$\text{SIG.Keygen}(\text{pp}_{\text{SIG}}) \rightarrow (\text{pk}_{\text{SIG}}, \text{sk}_{\text{SIG}})$ 1. Parse pp_{SIG} as (\mathbb{G}, q, g, H) . 2. Sample a scalar x uniformly from \mathbb{Z}_q . 3. Output $(\text{pk}_{\text{SIG}}, \text{sk}_{\text{SIG}}) := (g^x, x)$.	$\text{SIG.RandPk}(\text{pp}_{\text{SIG}}, \text{pk}_{\text{SIG}}, r_{\text{SIG}}) \rightarrow \hat{\text{pk}}_{\text{SIG}}$ 1. Parse pp_{SIG} as (\mathbb{G}, q, g, H) . 2. Output $\hat{\text{pk}}_{\text{SIG}} := \text{pk}_{\text{SIG}} \cdot g^{r_{\text{SIG}}}$.
$\text{SIG.Verify}(\text{pp}_{\text{SIG}}, \text{pk}_{\text{SIG}}, m, \sigma) \rightarrow b$ 1. Parse pp_{SIG} as (\mathbb{G}, q, g, H) . 2. Parse σ as (s, e) . 3. Set $r_v := g^s \text{pk}_{\text{SIG}}^e = g^{s+xe}$. 4. Set $e_v := H(r_v m)$. 5. Check if $e = e_v$.	$\text{SIG.RandSig}(\text{pp}_{\text{SIG}}, \sigma, r_{\text{SIG}}) \rightarrow \hat{\sigma}$ 1. Parse pp_{SIG} as (\mathbb{G}, q, g, H) . 2. Parse σ as (s, e) . 3. Output $\hat{\sigma} := (s - e \cdot r_{\text{SIG}}, e)$.

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 $\mathcal{R}_e^{\text{del}}$ 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.*

<p>DPC.Setup</p> <p><i>Input:</i> security parameter 1^λ</p> <p><i>Output:</i> public parameters pp</p> <ol style="list-style-type: none"> 1. Generate trapdoor commitment parameters: $\text{pp}_{\text{TCM}} \leftarrow \text{TCM.Setup}(1^\lambda)$. 2. Generate CRH parameters: $\text{pp}_{\text{CRH}} \leftarrow \text{CRH.Setup}(1^\lambda)$. 3. Generate signature parameters: $\text{pp}_{\text{SIG}} \leftarrow \text{SIG.Setup}(1^\lambda)$. 4. Generate NIZK parameters for $\mathcal{R}_e^{\text{del}}$ (Fig. 23): $\text{pp}_e \leftarrow \text{NIZK.Setup}(1^\lambda, \mathcal{R}_e^{\text{del}})$. 5. Output $\text{pp} := (\text{pp}_{\text{TCM}}, \text{pp}_{\text{CRH}}, \text{pp}_{\text{SIG}}, \text{pp}_e)$. 	<p>DPC.GenAddress</p> <p><i>Input:</i> public parameters pp and address metadata meta</p> <p><i>Output:</i> address key pair (apk, ask)</p> <ol style="list-style-type: none"> 1. Generate authorization key pair: $(\text{pk}_{\text{SIG}}, \text{sk}_{\text{SIG}}) \leftarrow \text{SIG.Keygen}(\text{pp}_{\text{SIG}})$. 2. Sample secret key sk_{PRF} for pseudorandom function PRF. 3. Sample randomness r_{pk} for commitment scheme TCM. 4. Set address public key $\text{apk} := \text{TCM.Commit}(\text{pp}_{\text{TCM}}, \text{pk}_{\text{SIG}} \parallel \text{sk}_{\text{PRF}} \parallel \text{meta}; r_{\text{pk}})$. 5. Set address secret key $\text{ask} := (\text{sk}_{\text{SIG}}, \text{sk}_{\text{PRF}}, \text{meta}, r_{\text{pk}})$. 6. Output (apk, ask).
<p>DPC.Execute^L</p> <p><i>Input:</i></p> <ul style="list-style-type: none"> • public parameters pp • old $\left\{ \begin{array}{l} \text{records } [\mathbf{r}_i]_1^m \\ \text{address secret keys } [\text{ask}_i]_1^m \end{array} \right.$ • auxiliary predicate input aux • transaction memorandandum memo <p style="margin-left: 200px;">• new $\left\{ \begin{array}{l} \text{address public keys } [\text{apk}_j]_1^n \\ \text{record payloads } [\text{payload}_j]_1^n \\ \text{record birth predicates } [\Phi_{\text{b},j}]_1^n \\ \text{record death predicates } [\Phi_{\text{d},j}]_1^n \end{array} \right.$</p> <p><i>Output:</i> new records $[\mathbf{r}_j]_1^n$ and transaction tx</p> <ol style="list-style-type: none"> 1. For each $i \in \{1, \dots, m\}$, process the i-th old record as follows: <ol style="list-style-type: none"> (a) Parse old record \mathbf{r}_i as $\mathbf{r}_i = \left(\begin{array}{cccccc} \text{address public key} & \text{apk}_i & \text{payload} & \text{payload}_i & \text{comm. rand.} & r_i \\ \text{serial number nonce} & \rho_i & \text{predicates} & (\Phi_{\text{b},i}, \Phi_{\text{d},i}) & \text{commitment} & \text{cm}_i \end{array} \right)$. (b) If $\text{payload}_i.\text{isDummy} = 1$, set ledger membership witness $\text{w}_{\text{L},i} := \perp$. If $\text{payload}_i.\text{isDummy} = 0$, compute ledger membership witness for commitment: $\text{w}_{\text{L},i} \leftarrow \text{L.Prove}(\text{cm}_i)$. (c) Parse address secret key ask_i as $(\text{sk}_{\text{SIG},i}, \text{sk}_{\text{PRF},i}, \text{meta}_i, r_{\text{pk},i})$ and derive $\text{pk}_{\text{SIG},i}$ from $\text{sk}_{\text{SIG},i}$. (d) Compute signature randomness: $r_{\text{SIG},i} \leftarrow \text{PRF}_{\text{sk}_{\text{PRF},i}}(\rho_i)$. (e) Compute serial number: $\text{sn}_i \leftarrow \text{SIG.RandPk}(\text{pp}_{\text{SIG}}, \text{pk}_{\text{SIG},i}, r_{\text{SIG},i})$. 2. For each $j \in \{1, \dots, n\}$, construct the j-th new record as follows: <ol style="list-style-type: none"> (a) Compute serial number nonce: $\rho_j := \text{CRH.Eval}(\text{pp}_{\text{CRH}}, j \parallel \text{sn}_1 \parallel \dots \parallel \text{sn}_m)$. (b) Construct new record: $\mathbf{r}_j \leftarrow \text{DPC.ConstructRecord}(\text{pp}, \text{apk}_j, \text{payload}_j, \Phi_{\text{b},j}, \Phi_{\text{d},j}, \rho_j)$. 3. Retrieve current ledger digest: $\text{st}_{\text{L}} \leftarrow \text{L.Digest}$. 4. Construct instance for relation $\mathcal{R}_e^{\text{del}}$: $\mathbb{x}_e := (\text{st}_{\text{L}}, [\text{sn}_i]_1^m, [\text{cm}_j]_1^n, \text{memo})$. 5. Construct witness for relation $\mathcal{R}_e^{\text{del}}$: $\mathbb{w}_e := ([\mathbf{r}_i]_1^m, [\text{w}_{\text{L},i}]_1^m, [\text{sk}_{\text{PRF},i}]_1^m, [\text{pk}_{\text{SIG},i}]_1^m, [\text{meta}_i]_1^m, [r_{\text{pk},i}]_1^m, [\mathbf{r}_j]_1^n, \text{aux})$. 6. Generate proof for relation $\mathcal{R}_e^{\text{del}}$: $\pi_e \leftarrow \text{NIZK.Prove}(\text{pp}_e, \mathbb{x}_e, \mathbb{w}_e)$. 7. For each $i \in \{1, \dots, m\}$: <ol style="list-style-type: none"> (a) Sign message: $\sigma_i \leftarrow \text{SIG.Sign}(\text{pp}_{\text{SIG}}, \text{sk}_{\text{SIG},i}, \mathbb{x}_e \parallel \pi_e)$. (b) Randomize signature: $\hat{\sigma}_i \leftarrow \text{SIG.RandSig}(\text{pp}_{\text{SIG}}, \sigma_i, r_{\text{SIG},i})$. 8. Construct transaction: $\text{tx} := ([\text{sn}_i]_1^m, [\text{cm}_j]_1^n, \text{memo}, \star)$, where $\star := (\text{st}_{\text{L}}, \pi_e, [\hat{\sigma}_i]_1^m)$. 9. Output $([\mathbf{r}_j]_1^n, \text{tx})$. 	
<p>DPC.Verify^L</p> <p><i>Input:</i> public parameters pp and transaction tx</p> <p><i>Output:</i> decision bit b</p> <ol style="list-style-type: none"> 1. Parse tx as $([\text{sn}_i]_1^m, [\text{cm}_j]_1^n, \text{memo}, \star)$ and \star as $(\text{st}_{\text{L}}, \pi_e, [\hat{\sigma}_i]_1^m)$. 2. Check that there are no duplicate serial numbers <ol style="list-style-type: none"> (a) within the transaction tx: $\text{sn}_i \neq \text{sn}_j$ for every distinct $i, j \in \{1, \dots, m\}$; (b) on the ledger: $\text{L.Contains}(\text{sn}_i) = 0$ for every $i \in \{1, \dots, m\}$. 3. Check that the ledger state is valid: $\text{L.ValidateDigest}(\text{st}_{\text{L}}) = 1$. 4. Construct instance for the relation $\mathcal{R}_e^{\text{del}}$: $\mathbb{x}_e := (\text{st}_{\text{L}}, [\text{sn}_i]_1^m, [\text{cm}_j]_1^n, \text{memo})$. 5. Check proof for the relation $\mathcal{R}_e^{\text{del}}$: $\text{NIZK.Verify}(\text{pp}_e, \mathbb{x}_e, \pi_e) = 1$. 6. For every $i \in \{1, \dots, m\}$, check that signature verifies: $\text{SIG.Verify}(\text{pp}_{\text{SIG}}, \text{sn}_i, \mathbb{x}_e \parallel \pi_e, \hat{\sigma}_i) = 1$. 	

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

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

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

where

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

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

A witness \mathbb{w}_e is valid for an instance \mathbb{x}_e if the following conditions hold:

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

Figure 23: The NP relation $\mathcal{R}_e^{\text{del}}$. 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.

On-ledger encryption. A user can store an encryption public key in the metadata of one of its addresses. Others can then use this public key to encrypt information about records created for the user, 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.

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: $L.Prove$ only returns a proof that a transaction (or a subcomponent thereof) appears on the ledger, and *not its position*. One can augment $L.Prove$ to instead output the transaction’s ledger position pos_L , and a proof that pos_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 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 $Execute$ (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. However this by itself is 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 $Execute$ by deriving them as $\rho_j := CRH(j \| r_{\rho,j} \| sn_1 \| \dots \| sn_m)$ for some randomness $r_{\rho,j}$ that is deleted after invoking $Execute$. 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 sk can derive a second key sk_S such that $\text{PRF}_{sk_S}(x_i) = y_i$ for each $(x_i, y_i) \in S$, while $\text{PRF}_{sk_S}(x) = \text{PRF}_{sk}(x)$ 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 $S = \{(\rho_i, sn_i)\}_i$, where ρ_i is the serial number nonce of the i -th record received from a corrupted party.

Acknowledgments

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 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.
- [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.
- [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. Full version available at <http://eprint.iacr.org/2014/349>.
- [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.
- [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.
- [BGG17] Sean Bowe, Ariel Gabizon, and Matthew D. Green. A multi-party protocol for constructing the public parameters of the pinocchio zk-snark. Cryptology ePrint Archive, Report 2017/602, 2017. <http://eprint.iacr.org/2017/602>.

- [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 Bove. Bellman, 2017. <https://github.com/zkcrypto/bellman>.
- [Bow17b] Sean Bove. Pairing, 2017. <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.
- [CCW18] Alessandro Chiesa, Lynn Chua, and Matthew Weidner. On cycles of pairing-friendly elliptic curves. arXiv math.NT/1803.02067, 2018. <https://arxiv.org/abs/1803.02067>.
- [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.
- [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.
- [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.
- [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. <https://arxiv.org/abs/1804.05141>.
- [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.
- [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.
- [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*, Security '15, pages 447–462, 2015.

- [Ell13] Elliptic. Elliptic enterprises limited. <https://www.elliptic.co/>, 2013.
- [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, 2018. <https://etherscan.io/tokens>.
- [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.
- [GM17] Jens Groth and Mary Maller. Snarky signatures: Minimal signatures of knowledge from simulation-extractable SNARKs. Cryptology ePrint Archive, Report 2017/540, 2017. <http://eprint.iacr.org/2017/540>.
- [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.
- [HBHW18] Daira Hopwood, Sean Bowe, Tailor Hornby, and Nathan Wilcox. Zcash protocol specification, 2018. 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. <https://arxiv.org/abs/1606.05917>.
- [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. <https://arxiv.org/abs/1709.02489>.
- [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*, Security '18, pages 1353–1370, 2018.
- [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.
- [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.
- [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.
- [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.
- [PB17] Joseph Poon and Vitalik Buterin. Plasma: Scalable autonomous smart contracts. <https://plasma.io/>, 2017.
- [PS00] David Pointcheval and Jacques Stern. Security arguments for digital signatures and blind signatures. *J. 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.
- [Rei16] Christian Reiwißner. 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 and Reto Strohli. 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 Reiwi ner. 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, 2015. <https://github.com/ethereum/EIPs/blob/master/EIPS/eip-20.md>.
- [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, 2015. <https://z.cash/>.
- [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.