



On Publicly-Accountable Zero-Knowledge and Small Shuffle Arguments

Nils Fleischhacker¹(✉) and Mark Simkin²

¹ Ruhr University Bochum, Bochum, Germany
mail@nilsfleischhacker.de

² Aarhus University, Aarhus, Denmark
simkin@cs.au.dk

Abstract. Constructing interactive zero-knowledge arguments from simple assumptions with small communication complexity and good computational efficiency is an important, but difficult problem. In this work, we study interactive arguments with noticeable soundness error in their full generality and for the specific purpose of constructing concretely efficient shuffle arguments.

To counterbalance the effects of a larger soundness error, we show how to transform such three-move arguments into publicly-accountable ones which allow the verifier to convince third parties of detected misbehavior by a cheating prover. This may be particularly interesting for applications where a malicious prover has to balance the profits it can make from cheating successfully and the losses it suffers from being caught.

We construct interactive, public-coin, zero-knowledge arguments with noticeable soundness error for proving that a target vector of commitments is a pseudorandom permutation of a source vector. Our arguments do not rely on any trusted setup and only require the existence of collision-resistant hash functions. The communication complexity of our arguments is *independent* of the length of the shuffled vector. For a soundness error of $2^{-5} = 1/32$, the communication cost is 153 bytes without and 992 bytes with public accountability, meaning that our arguments are shorter than shuffle arguments realized using Bulletproofs (IEEE S&P 2018) and even competitive in size with SNARKs, despite only relying on simple assumptions.

1 Introduction

Zero-knowledge arguments allow a prover to convince a verifier of the truth of a statement without leaking any additional information. Such arguments are a fundamental building block, ubiquitous in cryptography, with various applications in both theory and practice.

The quality of an argument system can be measured in several different ways. One of the most important quality measures is the size of the argument,

N. Fleischhacker—Funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany’s Excellence Strategy - EXC 2092 CASA - 390781972.

M. Simkin—Supported by a DFF Sapere Aude Grant 9064-00068B.

i.e. how many bits the prover needs to exchange with the verifier to convince them of the validity of the statement. Minimizing this measure is important for real-world applications, where the statements itself may be over several gigabytes large and where communicating large amounts of data over a wide area network can quickly turn into the main efficiency bottleneck. Another important measure is the computational efficiency of both prover and verifier. We would like our argument to incur as little computational overhead on both parties as possible. Finally, we would also like our arguments to rely on simple and well-studied assumptions. Arguments that rely on highly structured or even non-falsifiable assumptions may be prone to cryptanalysis, those that rely on more popular number-theoretic assumptions, like the discrete logarithm or the factoring assumption, can be broken by quantum computers, and arguments that require a common reference string need to rely on a trusted third party that has to generate this string.

One particularly popular class of zero-knowledge arguments are those that enable a prover to convince a verifier that two vectors of commitments or encryptions contain the same multiset of plaintext messages without revealing the messages themselves or the permutation between the two vectors. Shuffle arguments are used in applications like e-voting protocols [42], anonymous communication systems [18, 38], decentralized online poker [10], cryptocurrencies [20, 22], and others.

The idea of shuffle arguments originates in the work of Chaum on mix-nets [18] and the first constructions were presented by Sako and Kilian [45] and Abe [1–3]. These, as well as early subsequent works [25, 30, 34, 42], all had argument sizes, which were linear in the size ℓ of the permuted vector. The first sublinear shuffle argument, with an argument size of $\mathcal{O}(\ell^{2/3})$, was presented by Groth and Ishai [33]. Following this work, Groth and Bayer [9, 31] presented arguments with an argument size of $\mathcal{O}(\sqrt{\ell})$ and recently Bünz et al. [17] showed how to obtain arguments, based on the discrete logarithm assumption, of size $\mathcal{O}(\log \ell)$ via sorting circuits. A different line of works construct so called SNARKs [37, 40], which are constant-sized arguments for arbitrary statements. Unfortunately, SNARKs inherently rely on strong non-falsifiable assumptions [26], require a trusted setup, and are computationally expensive for the prover. Zero-knowledge arguments based on the MPC-in-the-head technique [5, 35] do not require any trusted setup, base their security solely on the existence of collision-resistant hash functions, but have a proof size of $\mathcal{O}(\sqrt{\ell})$. Interactive proofs based on probabilistically checkable proofs [37] only rely on collision-resistant hash functions, have proofs of size $\mathcal{O}(\log \ell)$, but are prohibitively expensive from a computational perspective.

Given this state of the art, it is evident that constructing small shuffle arguments, and more generally arbitrary arguments, from simple assumptions with good computational efficiency is a challenging, but important task. In this work, we ask:

Can we construct shuffle arguments of size $o(\log \ell)$ with good computational efficiency from simple assumptions that satisfy a relaxed, but still meaningful security notion?

Answering this question is of practical importance. In certain real-world scenarios, arguments that satisfy the strongest possible security notion can simply be too inefficient. In these cases more efficient arguments that satisfy a weaker security notion may provide an important trade-off between efficiency and security.

1.1 Our Contribution

In this work, we study interactive arguments with noticeable soundness errors, i.e. arguments that allow a cheating prover to convince a verifier of a false statement with some noticeable probability. Such arguments can still be useful in scenarios, where a malicious prover has to balance the profit that it can make from successfully cheating and the loss it has, when cheating is detected. Consider for instance a decentralized online poker game, where a malicious prover wins \$1 for every incorrect shuffle argument that is accepted by the verifier, but loses a \$100 security deposit if cheating is detected. In such a scenario, a soundness error as large as $1/2$ may be acceptable, since even then cheating is not profitable for a rationally behaving prover.

We study arguments with noticeable soundness error both in their full generality and for specific purpose of constructing concretely efficient shuffle arguments for pseudorandom shuffles. Concretely, we make the following contributions:

Publicly-Accountable Zero-Knowledge Arguments. To realize the idea of punishing cheating provers, we need to take care of two things. First, we need to ensure that a verifier, upon detecting a cheating attempt, obtains a publicly verifiable certificate that can be used to convince a third party auditor of the prover's malicious behavior. Secondly, we need to ensure that an honest prover cannot be falsely accused. We introduce the notion of *publicly-accountable zero-knowledge arguments* that formally model the two requirements above.

In the full version of this work, we show how to transform *any* three-move, honest-verifier zero-knowledge argument with a soundness error that is at least inversely polynomial in the security parameter into a publicly-accountable argument with only slightly larger communication complexity. This is achieved via two steps. We prove that those honest-verifier zero-knowledge arguments already satisfy full zero-knowledge and then show how such zero-knowledge arguments can be transformed into their publicly-accountable counterparts with the help of symmetric private information retrieval. In this version of the work we show how to make the shuffle arguments described below publicly accountable in a concretely efficient manner.

It is interesting to note that in contrast to sequential repetition, which is commonly used to make the soundness error negligible at the cost of a multiplicative factor that is linear in the security parameter, our transformation only incurs a small additive factor in terms of round and bandwidth complexity.

Shuffle Arguments for Pseudorandom Shuffles. Next, we focus on constructing efficient three-move (public-coin) honest-verifier zero-knowledge shuffle arguments with inversely polynomial soundness error. For this purpose, we make one additional observation that allows us to further simplify the problem we aim to solve. When looking at the majority of applications, where shuffle arguments are actually used, the concrete permutation between the input and the output vector is chosen at random. Most often, the goal is to simply hide the relation between entries in the input and the output vector, but the concrete permutation itself is irrelevant as long as it is sufficiently random. In e-voting, for example, users send their encrypted votes to an untrusted shuffling authority, which shuffles them to hide the voting preference of any specific user. In anonymous communication systems, users send messages through one or more shuffling authorities to some recipients and shuffling of the ciphertexts ensures that no outside observer can see which sender communicates with which recipient. Thus, we focus on shuffle arguments for pseudorandom shuffles instead of arbitrary shuffles.

We introduce the notion of *zero-knowledge arguments for partially fixed statements* and present conceptually simple interactive shuffling arguments satisfying this notion. The main idea behind our new notion is to consider statements that are only partially fixed, i.e. that consist of a fixed and a non-fixed part. The fixed part is known to both prover and verifier, whereas the non-fixed part is chosen by the prover. At the end of an interaction between prover and verifier, the verifier learns the full statement and is convinced of its correctness. For the specific case of shuffling, the fixed part is the initial vector of commitments and the non-fixed part is a permutation thereof. Our notion aims to capture the fact that we only care about the initial vector being permuted, but not about the concrete permutation that is used. Rather than requiring that zero-knowledge holds for all statements, we require that zero-knowledge holds for all partially fixed statements with a uniformly random non-fixed part. Our notion is, in spirit, similar to distributional zero-knowledge [19,27], but focuses on a particular distribution over the statements.

For this notion, we present the first computationally efficient shuffle arguments for pseudorandom permutations, whose argument size is *independent* of the length of the vector that is being permuted. More specifically, we present public-coin, three-move arguments in the standalone model based on simple assumptions, such as collision-resistant hash functions.¹ The soundness error in our constructions can be set arbitrarily small as long as it remains inversely

¹ Naturally, our arguments are only useful for vectors of rerandomizable commitments/encryptions, which may require specific number-theoretic assumptions. The arguments itself, however, only rely on simple assumptions.

polynomial in the security parameter. The computational overhead of our construction grows with smaller soundness errors. We show how an arbitrary number of shuffle arguments can be batched without any additional communication cost.

We evaluate the practical efficiency of our constructions by providing concrete argument sizes when instantiated in the standard model. Our evaluation shows that, for a soundness error of 2^{-5} , the instantiation of our shuffle argument has a communication cost of 153 bytes without and 992 bytes with public accountability,

This is on the same order of magnitude as SNARKs such as [32] at 144 bytes and smaller than Bulletproofs even when the permuted vector of commitments² is reasonably short. The size of our argument is independent of the specific number of commitments being shuffled. The computational cost of shuffling an ℓ -length vector with soundness error $1/t$ is dominated by computing $t \cdot \ell$ rerandomizations of the shuffled commitments for both prover and verifier. In practice for Pedersen and similar commitments, the cost for the verifier can be reduced to roughly 2ℓ rerandomization at the cost of roughly doubling the communication complexity. A detailed description of this modification can be found in the full version of this paper.

Since the non-fixed part of the partially fixed shuffle statement is chosen randomly in each execution, it follows that the fully fixed statement will be different in each execution with high probability. For this reason, we cannot reduce the soundness error via sequential repetition. Due to the non-negligible soundness error, we can also not use the Fiat-Shamir transform [23] for making them non-interactive.

1.2 Comparison to SNARKs

In terms of size, our shuffle arguments are similar to SNARKs. Our underlying assumptions, however, are significantly weaker. We do not rely on non-falsifiable assumptions or a trusted setup. In exchange, we have a noticeable soundness error. As such, for the specific case of shuffle arguments, our work shows that the need for a trusted setup and non-falsifiable assumptions can be overcome in a practically efficient manner in applications that can tolerate a small, but noticeable soundness error.

1.3 Relation to Secure Computation with Covert Security

Our concept of publicly-accountable zero-knowledge arguments is strongly related to general two- and multiparty computation protocols with covert security and public verifiability [6, 7]. A protocol is said to be secure against covert adversaries and publicly verifiable, if an actively misbehaving party in the protocol execution is caught with some constant probability and the honest parties are

² For the sake of concreteness, we focus on commitments in this work. However, our arguments are also applicable to other primitives such as rerandomizable encryption schemes.

guaranteed to obtain a certificate that can be shown to third parties to incriminate the misbehaving party. Applying a generic secure computation compiler, like the one of Damgård, Orlandi, and Simkin [21], to transform zero-knowledge arguments into publicly-accountable ones can potentially work, but would result in arguments whose size has an exponentially worse dependence on the soundness error.

1.4 Technical Overview

In the following, we present the main ideas behind our public-coin, zero-knowledge shuffle argument and outline how it can be made publicly accountable. In the full version, we show how arbitrary public-coin arguments with a polynomially large challenge space can be made publicly accountable.

The Shuffle Argument. Initially, both prover and verifier are given an initial vector of commitments V . The goal of the prover is to choose some vector V' and convince the verifier that there exists some permutation π , such that $V' = \pi(V)$. To be precise, the permutation π here does two things. It first rerandomizes and then permutes all commitments in V . The high-level idea behind our construction is to let the prover choose a permutation π , which can be represented as a sequence of t pseudorandom permutations π_1, \dots, π_t in a space-efficient manner. The prover first computes V' by sequentially applying each permutation π_i for $1 \leq i \leq t$ and then sends a hash of the intermediate vectors V_i and V' to the verifier, who picks $t - 1$ permutations that shall be opened. The prover sends descriptions of these $t - 1$ permutations to the verifier. Skipping over some details, the verifier now uses V to check every permutation π_j with $j < i$ and V' to check every permutation π_j with $j > i$ by recomputing the intermediate vectors. Since π_i remains hidden from the verifier, it cannot learn the overall permutation π . A malicious prover can only convince the verifier of a false statement if it chose all π_j with $i \neq j$ as correct and π_i as an incorrect permutation. Thus the probability of a prover cheating successfully is $1/t$. An interesting open question is whether our approach can be modified to get a better dependence between t and the resulting soundness error.

If done naively, then the size of the argument described above is $\mathcal{O}(t \cdot \log(\ell!))$, since the prover has to send each permutation π_j for $1 \leq j \leq t$ separately. This can easily be brought down to $\mathcal{O}(t)$ by sending only short random seeds which can be expanded into a pseudorandom permutation using a regular pseudorandom generator. To further reduce the size of our argument, we use a puncturable pseudorandom functions (PPRF), which behaves like a regular PRF but has an additional algorithm `Puncture` that allows the holder of a secret key k to compute a key $k\{x\}$, which can be used to evaluate the PPRF on every point $x' \neq x$. Importantly, for a PPRF it holds that the key $k\{x\}$ does not reveal anything about its evaluation at the point x . Assuming that we have a PPRF \mathcal{F} whose range is the set of all possible permutations, we can now succinctly represent π_j as the evaluation $\mathcal{F}(k, j)$ for $1 \leq j \leq t$. When the verifier asks to open all but




Fig. 1. The prover chooses the permutation and rerandomization by sampling a key k for a puncturable PRF. It then derives individual permutations and rerandomization factors for each stage by feeding the stage index through the PRF. The final result of performing the individual permutations forms the final shuffling and thus the prover-chosen part of the statement. The prover then hashes all of the intermediate permutations. The puncturable key k will allow the prover to partially open the computation of the intermediate permutations.

the i -th permutation, we return the punctured key $k\{i\}$ to the verifier. Using this approach, the size of the argument now mainly depends on the size of the punctured key and not directly on t . Using a PPRF based on the GGM construction [28], this brings down the size of the argument to $\mathcal{O}(\log t)$. Using a recent construction of PPRFs due to Aviram, Gellert, and Jager [8] we can make the size of our arguments even completely independent of t in the random oracle model. However, due to the large constants in Aviram et al.’s construction, this approach is only of theoretical interest. A visual illustration of our construction can be found in Fig. 1.

Making the Argument Publicly Accountable. To make the shuffle argument publicly-accountable, we need to ensure that a cheating prover produces some form of self-incriminating evidence whenever it fails to cheat. Since we want this evidence to be publicly verifiable, we need to assume that there exists a public signature key pk that is associated with the prover, who holds the corresponding signing key sk .

Ideally, we would like the prover to sign the transcript of all exchanged messages at the end of each execution and send this signature to the verifier. If the prover were to always do this, then we would be done, since the verifier would obtain a signature incriminating the prover, whenever it attempts to cheat, but is caught; assuming the prover always sends some last message, even if it can not respond correctly. Obviously this does not work, since the prover can simply abort the execution without signing anything, when it receives a challenge that it does not like. Our idea is to let the verifier receive the prover’s last message corresponding to the verifier’s challenge, without revealing the challenge to the

prover. On a high-level, we achieve this through the use of symmetric private information retrieval, which enables a receiver holding an index $i \in \{1, \dots, t\}$ to obtain a value x_i from the sender’s input vector $X = (x_1, \dots, x_t)$ in a manner that does not reveal i to the sender and does not reveal any x_j for $j \neq i$ to the receiver.

For the specific case of our shuffle arguments, the senders inputs will be a sequence of punctured keys and the receiver will retrieve one of them. We observe that in this instance the symmetric PIR can in fact be replaced by a very efficient oblivious key puncturing protocol implicit in the work of Boyle et al. [14].

2 Preliminaries

We denote by $\lambda \in \mathbb{N}$ the security parameter that is implicitly given as input to all algorithms in unary representation 1^λ . We denote by $\{0, 1\}^\ell$ the set of all bit-strings of length ℓ . For a finite set S , we denote the action of sampling x uniformly at random from S by $x \leftarrow S$, and we denote the cardinality of S by $|S|$. An algorithm is efficient or PPT if it runs in time polynomial in the security parameter. If \mathcal{A} is randomized then by $y := \mathcal{A}(x; r)$ we denote that \mathcal{A} is run on input x and with random coins r and produces output y . If no randomness is specified, then it is assumed that \mathcal{A} is run with freshly sampled uniform random coins. We write this as $y \leftarrow \mathcal{A}(x)$. A function $\text{negl}(\lambda)$ is negligible if for every positive polynomial $\text{poly}(\lambda)$ there exists an $N \in \mathbb{N}$ such that for all $\lambda > N$, $\text{negl}(\lambda) \leq \frac{1}{\text{poly}(\lambda)}$.

For two interactive Turing machines A and B we denote by $\langle A(a), B(b) \rangle$ the execution of the protocol between A and B an inputs a and b . We denote by $(t, s) \leftarrow \langle A(a), B(b) \rangle$ the outputs of A and B after the protocol execution respectively. In protocols where A does not receive an output, we write $s \leftarrow \langle A(a), B(b) \rangle$ to denote the output of B . We further denote by $\mathcal{T} := \langle A(a), B(b) \rangle$ the transcript resulting from the interaction.

2.1 Puncturable Pseudorandom Functions

Puncturable pseudorandom functions (PPRFs) can be constructed from one-way functions, where the key-length is $\mathcal{O}(\log |D|)$ and D is the input domain of the PRF [12, 15, 36]. Subsequent works have shown how to construct PPRFs with short keys from the strong RSA [8] and lattice-based assumptions [16].

Definition 1 (Puncturable PRFs). *The tuple $(\mathcal{F}, \text{Puncture})$ of PPT algorithms is a secure puncturable pseudorandom function with key length $\kappa(\lambda)$, input length $i(\lambda)$, and range $O(\lambda)$ if the following conditions hold:*

Functionality: *For every $\lambda \in \mathbb{N}$, $k \in \{0, 1\}^{\kappa(\lambda)}$, $x, x' \in \{0, 1\}^{i(\lambda)}$ with $x \neq x'$, and $k' \leftarrow \text{Puncture}(k, x)$ it holds that $\mathcal{F}(k, x') = \mathcal{F}(k', x')$.*

Pseudorandomness: For any PPT adversary \mathcal{A} it holds that

$$\left| \begin{array}{l} \Pr[x \leftarrow \mathcal{A}(1^\lambda) : \mathcal{A}(\text{Puncture}(k, x), \mathcal{F}(k, x)) = 1] \\ - \Pr[x \leftarrow \mathcal{A}(1^\lambda); y \leftarrow O(\lambda) : \mathcal{A}(\text{Puncture}(k, x), y) = 1] \end{array} \right| \leq \text{negl}(\lambda).$$

For our shuffle arguments (for vectors of length ℓ) we require a PPRF with range $\text{Perm}_\ell \times \mathcal{R}^\ell$ where \mathcal{R} is the randomness space of a perfectly and inversely rerandomizable commitment scheme and Perm_ℓ is the set of all $\ell!$ permutations over $\{0, \dots, \ell-1\}$. To obtain a PPRF over this range, one can simply use a PPRF that outputs bit strings, potentially stretching the output using a pseudorandom generator, and combine it with a shuffling algorithm like the Fisher-Yates shuffle [24] by using the stretched output of the PPRF as the random tape of the shuffling algorithm.

2.2 Oblivious Key Puncturing

We formalize the notion of an oblivious key puncturing protocol (OPP) between a receiver R , who has a secret index i , and a sender S , who has a secret PRF key k . At the end of the protocol execution, the receiver should learn the key punctured at i , while the sender should learn nothing. An oblivious key puncturing protocol is effectively a special case of a symmetric PIR, where the sender’s inputs are all possible punctured keys.

Definition 2 (Oblivious Key Puncturing). Let $(\mathcal{F}, \text{Puncture})$ be a secure puncturable PRF with key length $\kappa(\lambda)$, input length $i(\lambda)$, and range $O(\lambda)$. A pair of PPT algorithms (S, R) along with a setup algorithm Setup that outputs a crs is a secure receiver-extractable, oblivious key puncturing protocol for $(\mathcal{F}, \text{Puncture})$, if the following conditions hold:

Completeness: For any $k \in \{0, 1\}^{\kappa(\lambda)}$ and $i \in \{0, 1\}^{i(\lambda)}$, it holds that

$$\Pr \left[\begin{array}{l} \text{crs} \leftarrow \text{Setup}(1^\lambda); k' \leftarrow \text{Puncture}(k, i); \\ k'' \leftarrow \langle S(\text{crs}, k), R(\text{crs}, i) \rangle \end{array} : k' = k'' \right] = 1.$$

Receiver Privacy: For any $i, i' \in \{0, 1\}^{i(\lambda)}$ and any malicious PPT sender S^* , it holds that

$$\left| \begin{array}{l} \Pr[\text{crs} \leftarrow \text{Setup}(1^\lambda); b \leftarrow \langle S^*(\text{crs}), R(\text{crs}, i) \rangle : b = 1] \\ - \Pr[\text{crs} \leftarrow \text{Setup}(1^\lambda); b \leftarrow \langle S^*(\text{crs}), R(\text{crs}, i') \rangle : b = 1] \end{array} \right| \leq \text{negl}(\lambda),$$

where the probabilities are taken over the random coins of Setup , S^* and R .

Sender Simulation: There exists a PPT simulator $\text{Sim} = (\text{Sim}_0, \text{Sim}_1)$ such that for any key $k \in \{0, 1\}^{\kappa(\lambda)}$ and any malicious PPT receiver R^* it holds that

$$\left| \begin{array}{l} \Pr[\text{crs} \leftarrow \text{Setup}(1^\lambda); b \leftarrow \langle S(\text{crs}, k), R^*(\text{crs}, 1^\lambda) \rangle : b = 1] \\ - \Pr \left[\begin{array}{l} (\text{crs}, \text{td}) \leftarrow \text{Sim}_0(1^\lambda); \\ b \leftarrow \langle \text{Sim}_1^{\text{Puncture}(k, \cdot)}(\text{crs}, \text{td}), R^*(\text{crs}, 1^\lambda) \rangle \end{array} : b = 1 \right] \end{array} \right| \leq \text{negl}(\lambda),$$

where Sim_1 can query its oracle at most once and the probability is taken over the random coins of the involved parties.

Remark 1. Protocols that need a (non-empty) CRS, require a trusted setup. For those protocols, using standard techniques, the trusted setup can be avoided at the cost of a constant number of additional rounds of communication.

It turns out that for the PPRF based on one-way functions [12, 15, 36], highly efficient instantiations of such an oblivious key puncturing protocol already exist implicitly in [14]³. For a PPRF with domain D , the communication and computational complexity of their protocol is effectively that of $\log |D|$ invocations of an actively secure 1-out-of-2 oblivious transfer.

2.3 Commitments

Shuffle proofs are generally only of interest for rerandomizable commitment schemes. Our construction of shuffle proofs requires more than just perfect rerandomizability. Specifically we require that rerandomization can also be performed in reverse.

Definition 3 (Perfectly and Inversely Rerandomizable Commitments). Let $\mathcal{C} = (\text{Setup}, \text{Com})$ be a commitment scheme with message space \mathcal{M} and randomness space \mathcal{R} . \mathcal{C} is perfectly and inversely rerandomizable, if there exist PPT algorithms $\text{Rerand}, \text{Rerand}^{-1}$ such that the following conditions hold:

Perfect Rerandomization: For every $\text{ck} \leftarrow \text{Setup}(1^\lambda)$, $m \in \mathcal{M}$, and $c \leftarrow \text{Com}(\text{ck}, m)$ it holds that for a uniformly chosen $r \leftarrow \mathcal{R}$, $\text{Rerand}(\text{ck}, c, r)$ and $\text{Com}(\text{ck}, c; r)$ are distributed identically.

Inverse Rerandomization: For every $\text{ck} \leftarrow \text{Setup}(1^\lambda)$, $m \in \mathcal{M}$, $r \in \mathcal{R}$, and $c \leftarrow \text{Com}(\text{ck}, m)$ it holds that $\text{Rerand}^{-1}(\text{ck}, \text{Rerand}(\text{ck}, c, r), r) = c$.

One example of a popular commitment scheme satisfying the properties described above is the Pedersen commitment scheme [43]. Note that since we assume perfect rerandomizability, it is guaranteed that for any ck, c, r_1 , and r_2 , there exists an r_3 such that $\text{Rerand}(\text{ck}, \text{Rerand}(\text{ck}, c, r_1), r_2) = \text{Rerand}(\text{ck}, c, r_3)$.

3 Zero-Knowledge Argument for Partially Fixed Statements

In a regular proof or argument system, the full statement is fixed a priori and given to both the prover and verifier which then run an interactive protocol between them. In contrast in an arguments for partially fixed statements only a part x of the statement is fixed and the *prover gets to sample the rest of the*

³ The authors prove (in Theorem 7 in [14]) that their construction satisfies a weaker notion than the one we defined here, but it can be easily seen that their construction satisfies our notion as well, when instantiated with an actively secure oblivious transfer.

statement y together with auxiliary information \mathbf{aux} that will allow the prover to efficiently prove that $(x, y) \in \mathcal{L}$. Note, that \mathbf{aux} is not necessarily just a regular witness for (x, y) . In fact, in our shuffle proof, a regular witness for the statement would merely be the permutation and rerandomization factors. However, the auxiliary information used by our prover is a highly compact representation of a decomposition of both the permutation and rerandomization. This also implies that a prover in such a system is not necessarily capable of proving all $(x, y) \in \mathcal{L}$, but merely a, potentially small, subset. However, the definition of zero-knowledge will imply that the full statement (x, y) chosen by the prover is indistinguishable from a uniform choice of $(x, y) \in \mathcal{L}$ conditioned on x . To formally define such argument systems, we first define partially fixable languages, as those languages where y can be efficiently uniformly sampled conditioned on x .

Definition 4 (Partially Fixable Languages). *Let X, Y be sets. Let $\mathcal{L} \subseteq X \times Y$ be an NP language consisting of pairs $(x, y) \in X \times Y$ with the corresponding NP-relation \mathcal{R} . For any $x \in X$ we denote by \mathcal{L}_x the language $\mathcal{L}_x = \{(x, y') \mid y' \in Y \wedge (x, y') \in \mathcal{L}\}$. \mathcal{L} is called partially fixable if for all $x \in X$ such that $\mathcal{L}_x \neq \emptyset$ it is possible to uniformly sample from \mathcal{L}_x in expected polynomial time.*

We can now define argument systems for such languages.

Definition 5 (Arguments for Partially Fixed Statements). *Let $\mathcal{L} \subseteq X \times Y$ be a partially fixable language with the corresponding NP-relation \mathcal{R} . A probabilistic polynomial time two-stage prover $\mathbf{P} = (\mathbf{P}_0, \mathbf{P}_1)$ and a probabilistic polynomial time verifier \mathbf{V} are said to be an interactive argument for partially fixed statements of \mathcal{L} with soundness error ϵ if the following conditions hold:*

Completeness: *For any $x \in X$ with $\mathcal{L}_x \neq \emptyset$ it holds that*

$$\Pr[(y, \mathbf{aux}) \leftarrow \mathbf{P}_0(x); b \leftarrow \langle \mathbf{P}_1(x, y, \mathbf{aux}), \mathbf{V}(x, y) \rangle : (x, y) \in \mathcal{L} \wedge b = 1] = 1.$$

Soundness: *For any malicious probabilistic polynomial time prover \mathbf{P}^* and any $(x, y) \notin \mathcal{L}$ it holds that*

$$\Pr[b \leftarrow \langle \mathbf{P}^*(1^\lambda), \mathbf{V}(x, y) \rangle : b = 1] \leq \epsilon + \text{negl}(\lambda).$$

We are only interested in arguments that are zero knowledge. We define two flavors of zero-knowledge.

Definition 6 (Zero-Knowledge Arguments for Partially Fixed Statements). *Let (\mathbf{P}, \mathbf{V}) be an interactive argument for partially fixed statements of \mathcal{L} . The argument is said to be zero knowledge if there exists an expected polynomial time simulator \mathbf{Sim} , such that for any (potentially malicious) polynomial time verifier \mathbf{V}^* , all probabilistic polynomial time distinguishers \mathcal{D} , and all $x \in X$ with $\mathcal{L}_x \neq \emptyset$ it holds that*

$$\left| \begin{array}{l} \Pr[(y, \mathbf{aux}) \leftarrow \mathbf{P}_0(x); s \leftarrow \langle \mathbf{P}_1(x, y, \mathbf{aux}), \mathbf{V}^*(x, y) \rangle : \mathcal{D}(x, y, s) = 1] \\ - \Pr[(x, y) \leftarrow \mathcal{L}_x; s \leftarrow \mathbf{Sim}^{\mathbf{V}^*(x, y)}(x, y) : \mathcal{D}(x, y, s) = 1] \end{array} \right| \leq \text{negl}(\lambda)$$

where \mathbf{Sim} has the power of rewinding \mathbf{V}^* .

Definition 7 (Honest Verifier Zero-Knowledge). *Let (P, V) be an interactive argument for partially fixed statements of \mathcal{L} . The argument is said to be honest verifier zero knowledge if there exists an expected polynomial time simulator Sim , such that for all probabilistic polynomial time distinguishers \mathcal{D} , and all $x \in X$ with $\mathcal{L}_x \neq \emptyset$ it holds that*

$$\left| \begin{array}{l} \Pr[(y, T) \leftarrow \text{SIMU}(x) : \mathcal{D}(x, y, T) = 1] \\ - \Pr[(y, T) \leftarrow \text{REAL}(x) : \mathcal{D}(x, y, T) = 1] \end{array} \right| \leq \text{negl}(\lambda),$$

where REAL and SIMU are defined as follows

$\text{REAL}(x)$	$\text{SIMU}(x)$
$(y, \text{aux}) \leftarrow P_0(x)$	$(x, y) \leftarrow \mathcal{L}_x$
$T := \langle P_1(x, y, \text{aux}), V(x, y) \rangle$	$T \leftarrow \text{Sim}(x, y)$
return (y, T)	return (y, T)

We note several important differences compared to regular argument systems. When defining an argument systems where the prover can choose part of the statement completeness can no longer be defined by simply quantifying over all valid statements. Instead, completeness explicitly specifies that the honest prover will always choose valid statements. Further, in the definition of zero-knowledge, it is not necessarily clear how y should be chosen in the simulated case. The definition above requires that y is chosen uniformly at random from \mathcal{L}_x in this case as opposed to also being chosen by the prover. This has an important implication. Namely it implicitly requires the honest prover to choose y in a way that is computationally indistinguishable from uniform, since otherwise there exists a trivial distinguisher. Lastly we note that these definitions coincide with the standard zero-knowledge argument definitions, when $|\mathcal{L}_x| = 1$.

4 On Three-Move Public-Coin HVZK Arguments and Zero-Knowledge

In the following, we show that any three-move public-coin⁴ argument with a polynomially large challenge space that satisfies (computational) HVZK is also (computationally) zero-knowledge against malicious verifiers. A corollary of this result is that our shuffle argument from Sect. 5, which we will prove to be HVZK, is automatically fully zero-knowledge.

Theorem 1. *Let (P, V) be some three-move public-coin honest verifier zero-knowledge argument for language $\mathcal{L} \subseteq X \times Y$ and let \mathcal{C} be the associated challenge space. If $|\mathcal{C}| \leq \text{poly}(\lambda)$ then (P, V) is also zero-knowledge against malicious verifiers.*

⁴ We call a three-move argument system public-coin if the second message is a uniformly random bit-string.

Proof. Let V^* be an arbitrary malicious polynomial time verifier. Let Sim' be the *honest verifier* zero-knowledge simulator for the 3-move public-coin argument as specified in Definition 7. To prove the theorem, we specify a zero-knowledge simulator Sim that takes as input a statement (x, y) , has blackbox access to V^* , and produces an output that is computationally indistinguishable from the output of V^* in a real protocol execution.

At first sight, the proof of the theorem statement may seem trivial. Intuitively, Sim picks a random challenge d , runs the simulator Sim' to obtain a transcript (e, d, z) , feeds the first message e to V^* and if the verifier outputs a challenge d^* with $d^* = d$, then we are done and otherwise we simply restart this whole process until we guess the verifier’s challenge correctly. Unfortunately, this approach only works if we have an argument that satisfies *perfect HVZK* and it turns out that this naive simulator is not guaranteed to run in expected polynomial time if our argument system is only computationally HVZK.

To make sure that our simulator V^* does indeed run in expected polynomial time, we closely follow a proof strategy due to Goldreich and Kahan [29].⁵ We specify the zero-knowledge simulator Sim in Fig. 2.

$\text{Sim}^{V^*(x,y)}(x, y)$	$\text{EstimateDelta}(x, y)$
1 : $(e, d, z) \leftarrow \text{Sim}'(x, y)$	$k := 0, m := 0$
2 : $d^* \leftarrow V^*(e)$	while $k < 12\lambda$
3 : if $d^* = \text{abort}$	$m := m + 1$
4 : return abort	rewind V^*
5 : $\tilde{d} \leftarrow \text{EstimateDelta}(x, y)$	$(e, d, \tilde{z}) \leftarrow \text{Sim}'(x, y)$
6 : for $0 \leq i < \min(4 C \lambda\tilde{d}^{-1}, 2^\lambda)$	if $V^*(e) \neq \text{abort}$
7 : rewind V^*	$k = k + 1$
8 : $(e, d, z) \leftarrow \text{Sim}'(x, y)$	return k/m
9 : $d^* \leftarrow V^*(e)$	
10 : if $d = d^*$	
11 : return $V^*(z)$	
12 : return fail	

Fig. 2. Zero-knowledge simulator for any three-move public coin honest verifier zero knowledge argument with a polynomially large challenge space.

We first observe in Lemma 2 that for any verifier V^* the probability of aborting after seeing a simulated first message output by Sim' does not differ significantly from the probability of aborting after seeing a real first protocol message.

⁵ See [39] for a very nice and detailed discussion of this proof strategy.

Lemma 2. *For any polynomial time algorithm V^* and any $x \in X$, such that $\mathcal{L}_x \neq \emptyset$ it holds that*

$$\left| \Pr[(y, e, d, z) \leftarrow \text{SIMU}(x) : V^*(x, y, e) = \text{abort}] - \Pr[(y, e, d, z) \leftarrow \text{REAL}(x) : V^*(x, y, e) = \text{abort}] \right| \leq \text{negl}(\lambda)$$

Proof. Let V^* be an arbitrary malicious polynomial time verifier. Consider the following distinguisher \mathcal{D} against the *honest verifier* zero-knowledge property of the argument: Upon receiving (x, y) and (e, d, z) as input, the distinguisher \mathcal{D} invokes V^* with fresh random coins and input (x, y, e) . If V^* aborts then \mathcal{D} outputs 1. Otherwise it outputs 0. We observe that \mathcal{D} 's distinguishing advantage against the honest verifier zero-knowledge property of the argument corresponds exactly to the difference in the abort probabilities of V^* . Since the argument is honest verifier zero-knowledge, \mathcal{D} 's distinguishing advantage must be negligible and Lemma 2 thus follows. \square

Furthermore, we use an observation made previously by Goldreich and Kahan [29].

Lemma 3 ([29]). *For any algorithm V^* , let $\delta = \delta(\lambda)$ be the probability that it does not abort upon seeing a simulated first message. With probability at least $1 - 2^{-\lambda}$, the estimate $\tilde{\delta}$ in line 5 of Sim in Fig. 2 is within a constant factor of δ .*

Using these two observations we will now analyze the probability that one iteration of Sim's main loop is successful. I.e. we will show that the probability that for a precomputed transcript (e, d, z) , V^* upon receiving input e will return d^* with $d = d^*$ with probability at least $\delta/|\mathcal{C}| - \text{negl}(\lambda)$.

Lemma 4. *Let (P, V) be a three-move public-coin honest verifier zero-knowledge argument for language $\mathcal{L} \subseteq X \times Y$ and let \mathcal{C} be the associated challenge space with $|\mathcal{C}| \leq \text{poly}(\lambda)$. Let further V^* be any polynomial time verifier. For any $x \in X$, such that $\mathcal{L}_x \neq \emptyset$ it then holds that*

$$\Pr[(y, e, d, z) \leftarrow \text{SIMU}(x) : V^*(x, y, e) = d] \geq \frac{\delta}{|\mathcal{C}|} - \text{negl}(\lambda).$$

Proof. Let V^* be an arbitrary polynomial time verifier. Now consider the following distinguisher \mathcal{D} against the honest verifier zero-knowledge property of the argument. The distinguisher \mathcal{D} receives as input (x, y) and (e, d, z) . It initializes V^* , provides it with (x, y, e) , and receives back d^* . If $d^* = \text{abort}$, then \mathcal{D} flips a random coin b and return that as its guess. Otherwise, \mathcal{D} outputs 1 if $d = d^*$ and 0 if $d \neq d^*$. Let $\delta + \gamma$ be the probability that V^* aborts after seeing a first real message, where $|\gamma| = \text{negl}(\lambda)$ by Lemma 2. By the honest verifier zero-knowledge property of the argument it must then hold that

$$\text{negl}(\lambda) \geq \left| \Pr[(y, e, d, z) \leftarrow \text{SIMU}(x) : \mathcal{D}(x, y, e, d, z) = 1] - \Pr[(y, e, d, z) \leftarrow \text{REAL}(x) : \mathcal{D}(x, y, e, d, z) = 1] \right|$$

$$\begin{aligned}
 & \left| \begin{aligned}
 & \Pr [(y, e, d, z) \leftarrow \text{SIMU}(x) : \mathcal{D}(x, y, e, d, z) = 1 \mid \mathbf{V}^*(x, y, e) \neq \text{abort}] \\
 & \cdot \Pr [(y, e, d, z) \leftarrow \text{SIMU}(x) : \mathbf{V}^*(x, y, e) \neq \text{abort}] \\
 + & \Pr [(y, e, d, z) \leftarrow \text{SIMU}(x) : \mathcal{D}(x, y, e, d, z) = 1 \mid \mathbf{V}^*(x, y, e) = \text{abort}] \\
 & \cdot \underbrace{\Pr [(y, e, d, z) \leftarrow \text{SIMU}(x) : \mathbf{V}^*(x, y, e) = \text{abort}]}_{=1-\delta} \\
 = & -\Pr [(y, e, d, z) \leftarrow \text{REAL}(x) : \mathcal{D}(x, y, e, d, z) = 1 \mid \mathbf{V}^*(x, y, e) = \text{abort}] \\
 & \cdot \underbrace{\Pr [(y, e, d, z) \leftarrow \text{REAL}(x) : \mathbf{V}^*(x, y, e) = \text{abort}]}_{=1-\delta-\gamma} \\
 - & \Pr [(y, e, d, z) \leftarrow \text{REAL}(x) : \mathcal{D}(x, y, e, d, z) = 1 \mid \mathbf{V}^*(x, y, e) \neq \text{abort}] \\
 & \cdot \underbrace{\Pr [(y, e, d, z) \leftarrow \text{REAL}(x) : \mathbf{V}^*(x, y, e) \neq \text{abort}]}_{=\delta+\gamma}
 \end{aligned} \right| \\
 = & \left| \begin{aligned}
 & \Pr [(y, e, d, z) \leftarrow \text{SIMU}(x) : \mathbf{V}^*(x, y, e) = d \mid \mathbf{V}^*(x, y, e) \neq \text{abort}] \\
 & \cdot \Pr [(y, e, d, z) \leftarrow \text{SIMU}(x) : \mathbf{V}^*(x, y, e) \neq \text{abort}] \\
 + & \underbrace{\Pr [b \leftarrow \{0, 1\} : b = 1]}_{=1/2} \cdot (1 - \delta) - \underbrace{\Pr [b \leftarrow \{0, 1\} : b = 1]}_{=1/2} \cdot (1 - \delta - \gamma) \\
 - & \underbrace{\Pr [(y, e, d, z) \leftarrow \text{REAL}(x) : \mathbf{V}^*(x, y, e) = d \mid \mathbf{V}^*(x, y, e) \neq \text{abort}]}_{=1/|\mathcal{C}|} \\
 & \cdot (\delta + \gamma)
 \end{aligned} \right| \\
 = & \left| \begin{aligned}
 & \Pr [(y, e, d, z) \leftarrow \text{SIMU}(x) : \mathbf{V}^*(x, y, e) = d \mid \mathbf{V}^*(x, y, e) \neq \text{abort}] \\
 & \cdot \Pr [(y, e, d, z) \leftarrow \text{SIMU}(x) : \mathbf{V}^*(x, y, e) \neq \text{abort}] + \frac{\gamma}{2} - \frac{\delta + \gamma}{|\mathcal{C}|}
 \end{aligned} \right| \\
 = & \left| \begin{aligned}
 & \Pr [(y, e, d, z) \leftarrow \text{SIMU}(x) : \mathbf{V}^*(x, y, e) = d \mid \mathbf{V}^*(x, y, e) \neq \text{abort}] \\
 & \cdot \Pr [(y, e, d, z) \leftarrow \text{SIMU}(x) : \mathbf{V}^*(x, y, e) \neq \text{abort}] - \frac{\delta}{|\mathcal{C}|} + \frac{(|\mathcal{C}| - 2)\gamma}{2|\mathcal{C}|}
 \end{aligned} \right| \\
 = & \left| \begin{aligned}
 & \Pr [(y, e, d, z) \leftarrow \text{SIMU}(x) : \mathbf{V}^*(x, y, e) = d \mid \mathbf{V}^*(x, y, e) \neq \text{abort}] \\
 & \cdot \Pr [(y, e, d, z) \leftarrow \text{SIMU}(x) : \mathbf{V}^*(x, y, e) \neq \text{abort}] - \frac{\delta}{|\mathcal{C}|} + \text{negl}(\lambda)
 \end{aligned} \right|. \quad (1)
 \end{aligned}$$

We can now consider the two cases of the value between the absolute value bars in Eq. 1 being positive, or negative. If it's positive, then it holds that

$$\begin{aligned}
 \frac{\delta}{|\mathcal{C}|} - \text{negl}(\lambda) & \leq \Pr [(y, e, d, z) \leftarrow \text{SIMU}(x) : \mathbf{V}^*(x, y, e) = d \mid \mathbf{V}^*(x, y, e) \neq \text{abort}] \\
 & \cdot \Pr [(y, e, d, z) \leftarrow \text{SIMU}(x) : \mathbf{V}^*(x, y, e) \neq \text{abort}]. \quad (2)
 \end{aligned}$$

If it's negative, then it must hold that

$$\begin{aligned}
 \text{negl}(\lambda) & \geq -\Pr [(y, e, d, z) \leftarrow \text{SIMU}(x) : \mathbf{V}^*(x, y, e) = d \mid \mathbf{V}^*(x, y, e) \neq \text{abort}] \\
 & \cdot \Pr [(y, e, d, z) \leftarrow \text{SIMU}(x) : \mathbf{V}^*(x, y, e) \neq \text{abort}] + \frac{\delta}{|\mathcal{C}|} - \text{negl}(\lambda)
 \end{aligned}$$

and thereby that

$$\frac{\delta}{|\mathcal{C}|} - \text{negl}(\lambda) \leq \Pr [(y, e, d, z) \leftarrow \text{SIMU}(x) : \mathbf{V}^*(x, y, e) = d \mid \mathbf{V}^*(x, y, e) \neq \text{abort}] \cdot \Pr [(y, e, d, z) \leftarrow \text{SIMU}(x) : \mathbf{V}^*(x, y, e) \neq \text{abort}]. \tag{3}$$

Combining the two cases, i.e., Eqs. 2 and 3 we can use the law of total probability to conclude that

$$\Pr [(y, e, d, z) \leftarrow \text{SIMU}(x) : \mathbf{V}^*(x, y, e) = d] \geq \frac{\delta}{|\mathcal{C}|} - \text{negl}(\lambda).$$

as claimed. □

We now want to use Lemma 4 to argue that the output of the simulator is indistinguishable from the output of \mathbf{V}^* in a real execution. For this, consider the following. By Lemma 4, there exists a negligible function ϵ , such that

$$\Pr [(y, e, d, z) \leftarrow \text{SIMU}(x) : \mathbf{V}^*(x, y, e) = d] \geq \frac{\delta}{|\mathcal{C}|} - \epsilon(\lambda).$$

For each security parameter $\lambda \in \mathbb{N}$ we can consider two cases:

Case i. If it holds that $\delta(\lambda) > 2|\mathcal{C}|\epsilon(\lambda)$, then we have $\epsilon(\lambda) < \delta/(2|\mathcal{C}|)$ and it therefore holds that

$$\Pr [(y, e, d, z) \leftarrow \text{SIMU}(x) : \mathbf{V}^*(x, y, e) = d] \geq \frac{\delta}{|\mathcal{C}|} - \epsilon(\lambda) > \frac{\delta}{2|\mathcal{C}|}.$$

It follows that in expectation the simulator needs at most $2|\mathcal{C}|/\delta$ rewinding attempts to obtain one non-aborting and correctly guessed execution. Via markov-inequality it follows that the probability of not having seen a single non-aborting correctly guessed execution after $4\lambda|\mathcal{C}|/\delta$ rewindings is negligible.

Lastly observe that by Lemma 3 the estimate $\tilde{\delta}$ is within a constant factor of δ with probability $1 - 2^{-\lambda}$. Therefore, the simulator will output a valid transcript with a probability of $1 - \text{negl}(\lambda)$, ensuring that the output of the simulator is indistinguishable from the output of \mathbf{V}^* in a real execution with overwhelming probability.

Case ii. If it holds that $\delta(\lambda) \leq 2|\mathcal{C}|\epsilon(\lambda)$, then δ is smaller than a negligible function for this λ . Assume that in this case the rewinding strategy *always fails*. Then a real execution of \mathbf{V}^* results in **abort** with probability at least $1 - (\delta - \text{negl}(\lambda))$ by Lemma 2, while **Sim** outputs **abort** with probability $1 - \delta$. That means the statistical distance between the two output distributions is at most $\delta + \text{negl}(\lambda)$ which is an overall negligible function. Combining the two cases, we can conclude that the distinguishing advantage against **Sim** is upper bounded by a negligible function for each $\lambda \in \mathbb{N}$ and thus is overall computationally indistinguishable from the output of \mathbf{V}^* in a real execution.

It remains to bound the expected runtime of `Sim`. Again, by Lemma 3, the estimate $\tilde{\delta}$ is within a constant factor of δ with probability $1 - 2^{-\lambda}$. But whenever the estimate is *wrong*, the runtime of the main loop is still bounded by the worst case running time of 2^λ with the simulator outputting `fail`. We thus have an upper bound on the expected runtime of

$$\begin{aligned} & \overbrace{\text{poly}(\lambda)}^{\text{lines 1-4}} + \delta \left(\overbrace{\text{poly}(\lambda) + \frac{12\lambda}{\delta} \text{poly}(\lambda)}^{\text{EstimateDelta}} + \overbrace{\left((1 - 2^{-\lambda}) \frac{4\lambda |\mathcal{C}|}{\tilde{\delta}} + 2^{-\lambda} 2^\lambda \right) \cdot \text{poly}(\lambda)}^{\text{lines 6-12}} \right) \\ & \leq \text{poly}(\lambda) + (12\lambda + \mathcal{O}(1)) \cdot 4\lambda |\mathcal{C}| + \delta \text{poly}(\lambda) \cdot \text{poly}(\lambda) \leq \text{poly}(\lambda). \end{aligned}$$

□

5 An Efficient Shuffle Argument

Let $\mathcal{C} = (\text{Setup}, \text{Com}, \text{Rerand})$ be a perfectly and inversely rerandomizable commitment scheme with message space \mathcal{M} and randomness space \mathcal{R} . To define shuffle arguments for \mathcal{C} , we first need to define partially fixable language of valid shuffles relative to a rerandomizable commitment scheme. To this end, we first define π as an algorithm that takes a vector of \mathcal{C} commitments V , a permutation $p \in \text{Perm}_\ell$, and randomnesses $r_0, \dots, r_{\ell-1} \in \mathcal{R}^\ell$ as input, permutes the elements of V and randomizes each commitment. The algorithm π as well as its inverse is described in Fig. 3. We can now define the partially fixable language of valid shuffles relative to π as follows.

Definition 8 (Valid Shuffle). *Let \mathcal{C} be a perfectly rerandomizable commitment scheme with commitment space \mathcal{C} . The language $\text{Shuffle}_\ell \subseteq \mathcal{C}^\ell \times \mathcal{C}^\ell$ of valid shuffles of vectors of length ℓ is defined as*

$$\text{Shuffle}_\ell = \{ (V, V') \in \mathcal{C}^\ell \times \mathcal{C}^\ell \mid \exists (p, \vec{r}) \in \text{Perm}_\ell \times \mathcal{R}^\ell. V' = \pi(V, p, \vec{r}) \}$$

Shuffles are transitive as stated by the following lemma.

Lemma 5. *If $(V, V') \in \text{Shuffle}_\ell$ and $(V', V'') \in \text{Shuffle}_\ell$, then $(V, V'') \in \text{Shuffle}_\ell$.*

Proof. Since permutations are closed under composition and since, by assumption on the commitment scheme, it holds that for any $r_1, r_2 \in \mathcal{R}$, there exists an $r_3 \in \mathcal{R}$ such that $\text{Rerand}(\text{ck}, \text{Rerand}(\text{ck}, c, r_1), r_2) = \text{Rerand}(\text{ck}, c, r_3)$, the lemma immediately follows. □

```

π(V, p, r₀, ..., r_{ℓ-1})
-----
for 0 ≤ i < ℓ
    V'[p(i)] := Rerand(ck, V[i], r_{p(i)})
return V'
```

```

π⁻¹(V', p, r₀, ..., r_{ℓ-1})
-----
for 0 ≤ i < ℓ
    V[i] := Rerand⁻¹(ck, V'[p(i)], r_i)
return V'
```

Fig. 3. The algorithms for rerandomizing and permuting a vector of ciphertexts, as well as its inverse.

Definition 9 (Shuffle Argument). *An interactive arguments for partially fixed statements of Shuffle_ℓ relative to any perfectly rerandomizable commitment scheme \mathcal{C} is called a shuffle argument for \mathcal{C} .*

Now, let $(\mathcal{F}, \text{Puncture})$ be a puncturable pseudorandom function with key length $k(\lambda)$, input length $i(\lambda)$, and range $\text{Perm}_\ell \times \mathcal{R}^\ell$. Let H be a collision resistant hash function. In Fig. 4 we then describe a simple three-move shuffle argument. We will first prove that this protocol is a zero knowledge shuffle argument as stated in the following theorem.

$P_0(V_0)$	$V_1(V'_0, V'_t, c)$	$V_2(k')$
$k \leftarrow \{0, 1\}^{k(\lambda)}$	$d \leftarrow \{1, \dots, t\}$	for $1 \leq i < d$
for $1 \leq i \leq t$	return d	$V'_i := \pi(V'_{i-1}, \mathcal{F}(k', i))$
$V_i := \pi(V_{i-1}, \mathcal{F}(k, i))$	$P_2(d)$	for $t > i \geq d$
$c := H(V_1, \dots, V_{t-1})$	if $d \in \{1, \dots, t\}$	$V'_i := \pi^{-1}(V'_{i+1}, \mathcal{F}(k', i + 1))$
return $(V_t, (c, k))$	$k' \leftarrow \text{Puncture}(k, d)$	if $H(V'_1, \dots, V'_{t-1}) = c$
$P_1(V_0, (c, k))$	return k'	return 1
return c	else	else return 0
	return \perp	

Fig. 4. An algorithmic description of the shuffle argument.

Theorem 6. *Let $\mathcal{C} = (\text{Setup}, \text{Com}, \text{Rerand})$ be a perfectly and inversely rerandomizable commitment scheme with message space \mathcal{M} and randomness space \mathcal{R} . Let $(\mathcal{F}, \text{Puncture})$ be a puncturable pseudorandom function with key length $k(\lambda)$, input length $i(\lambda)$, and range $\text{Perm}_\ell \times \mathcal{R}^\ell$. Let H be a collision resistant hash function. Then the argument system $\langle P = (P_0, P_1, P_2), V = (V_1, V_2) \rangle$ described in Fig. 4 is a zero-knowledge shuffle argument with soundness error $1/t$ for \mathcal{C} .*

Theorem 6 follows from Lemmas 7 and 10 as well as Corollary 12, which we prove in the following.

Lemma 7. *Let \mathcal{C} , $(\mathcal{F}, \text{Puncture})$, and H be as in Theorem 6. Then the argument system described in Fig. 4 is complete.*

Proof. We need to show that it always holds that $(V_0, V'_t) \in \text{Shuffle}_\ell$ and that, in an interaction with the honest prover, the verifier always accepts and outputs 1.

Claim 8. *For any $V_0 \in \mathcal{C}^\ell$ and any $(V_t, (c, k)) \leftarrow P_0(V_0)$ it holds that $(V_0, V'_t) \in \text{Shuffle}_\ell$.*

The prover computes each V_i as $V_i := \pi(V_{i-1}, \mathcal{F}(k, i))$, where $\mathcal{F}(k, i)$ outputs the description of a permutation and ℓ random values in \mathcal{R} . It then follows from the definition of π in Fig. 3 and the perfect rerandomizability of \mathcal{C} , that for $0 < i \leq t$, $(V_{i-1}, V_i) \in \text{Shuffle}_\ell$. Using Lemma 5 we can then conclude by induction, that $(V_0, V_t) \in \text{Shuffle}_\ell$.

Claim 9. *For any $V_0 \in \mathcal{C}^\ell$ and any $(V_t, (c, k)) \leftarrow P_0(V_0)$ any honest execution of $\langle P(V_0, V_t, (c, k)), V(V_0, V_t) \rangle$ will always output 1.*

We note that V_2 uses k' to recompute all V'_i for $1 \leq i < t$. For $1 \leq i < d$, this happens by computing $V'_i := \pi(V'_{i-1}, \mathcal{F}(k', i)) = \pi(V'_{i-1}, \mathcal{F}(k, i))$, where the last equality holds by the functionality requirement of the puncturable PRF, since $i \neq d$. As it always holds that $V'_0 = V_0$, it follows by induction over i , that $V'_i = V_i$ for $1 \leq i < d$.

For $d \leq i < t$, the verifier computes $V'_i := \pi^{-1}(V'_{i+1}, \mathcal{F}(k', i + 1))$. It always holds that $V_t = V'_t$ and the prover computes $V_t := \pi(V_{t-1}, \mathcal{F}(k, t))$. This gives us $V'_{t-1} = \pi^{-1}(\pi(V_{t-1}, \mathcal{F}(k, t)), \mathcal{F}(k, t))$, since $t \neq d$. By the definition of π and π^{-1} we thus have that for $0 \leq j < \ell$ it holds that

$$V'_{t-1}[j] = \text{Rerand}^{-1}(\text{Rerand}(V'[j], r_j), r_j) = V_{t-1}[j]$$

for some value of r_j . The last equality follows from the inverse rerandomizability of \mathcal{C} . Therefore, it follows that $V'_{t-1} = V_{t-1}$ and by induction that $V'_i = V_i$ for $d \leq i < t$.

We thus have that with probability 1, $(V'_1, \dots, V'_{t-1}) = (V_1, \dots, V_{t-1})$ and therefore $H(V'_1, \dots, V'_{t-1}) = c$. It thus follows that V_2 outputs 1.

Combining the two claims Lemma 7 immediately follows. □

Lemma 10. *Let \mathcal{C} , $(\mathcal{F}, \text{Puncture})$, and H be as in Theorem 6. Then the argument system described in Fig. 4 is sound with soundness error $1/t$.*

Proof. Let $(V_0^*, V_t^*) \notin \text{Shuffle}_\ell$ and let P^* be an arbitrary probabilistic polynomial time prover. We will show, that

$$\Pr[b \leftarrow \langle P^*(1^\lambda), V(x, y) \rangle : b = 1] \leq 1/t + \text{negl}(\lambda).$$

We will assume without loss of generality, that P^* actually sends a first message c and that c is fixed.⁶

Let $d_0, d_1 \in \{1, \dots, t\}$ be two arbitrary distinct challenges and let $k'_0 \leftarrow P^*(d_0)$ and $k'_1 \leftarrow P^*(d_1)$ be the corresponding responses. Consider, that the verifier works by recomputing $V_b = (V_1^b, \dots, V_{t-1}^b)$ and checking that it hashes to c . The verifier computes V_i^b as $V_i^b = \pi(V_{i-1}^b, \mathcal{F}(k'_b, i))$ for $i < d_b$ and as $\pi^{-1}(V_{i+1}^b, \mathcal{F}(k'_b, i + 1))$ for $i \geq d_b$.

By definition of π and π^{-1} , this implies for $i < d_0$ that $(V_{i-1}^0, V_i^0) \in \text{Shuffle}_\ell$ and for $i \geq d_0$ that $(V_i^0, V_{i+1}^0) \in \text{Shuffle}_\ell$. By Lemma 5 we can thus conclude that

$$(V_0, V_{d_0-1}^0) \in \text{Shuffle}_\ell \quad \text{and} \quad (V_{d_0}^0, V_t) \in \text{Shuffle}_\ell. \tag{4}$$

⁶ This is without loss of generality, since we can fix the prover's random coins to those that lead to the highest success probability.

Since $d_1 \neq d_0$, we have by the same reasoning, that

$$(V_{d_0-1}^1, V_{d_0}^1) \in \text{Shuffle}_\ell. \tag{5}$$

If it were true, that $(V_{d_0-1}^0, V_{d_0}^0) = (V_{d_0-1}^1, V_{d_0}^1)$ then it would follow from Eqs. 4 and 5 by Lemma 5 that $(V_0, V_t) \in \text{Shuffle}_\ell$ which would contradict the initial assumption. Therefore, it must hold that $(V_{d_0-1}^0, V_{d_0}^0) \neq (V_{d_0-1}^1, V_{d_0}^1)$ and $\mathbf{V}_0 \neq \mathbf{V}_1$. This would mean, however, that we could break collision resistance of \mathbf{H} by presenting $\mathbf{V}_0, \mathbf{V}_1$ with probability

$$\Pr[(k'_0, r_0) \leftarrow \mathbf{P}^*(d_0) : \mathbf{V}_2(k'_0) = 1] \cdot \Pr[(k'_1, r_1) \leftarrow \mathbf{P}^*(d_1) : \mathbf{V}_2(k'_1) = 1].$$

Since the hash function is collision resistant, it follows that the above probability can be bounded by a negligible function. Thus, at least one of the two probabilities must be itself negligible. Since we have shown the above for all pairs of distinct challenges, this means that there can exist at most one challenge $d \in \{1, \dots, t\}$ such that $\Pr[(k', r) \leftarrow \mathbf{P}^*(d) : \mathbf{V}_2(k') = 1]$ is non-negligible. It thus ultimately follows that

$$\begin{aligned} & \Pr[b \leftarrow \langle \mathbf{P}^*(1^\lambda), \mathbf{V}(x, y) \rangle : b = 1] \\ &= \sum_{d=1}^t \Pr[\mathbf{V}_1(c) = d] \cdot \Pr[k' \leftarrow \mathbf{P}^*(d) : \mathbf{V}_2(k') = 1] \\ &= \frac{1}{t} \cdot \sum_{d=1}^t \Pr[k' \leftarrow \mathbf{P}^*(d) : \mathbf{V}_2(k') = 1] \\ &\leq \frac{1}{t} \cdot (1 + (t - 1) \cdot \text{negl}(\lambda)) \leq \frac{1}{t} + \text{negl}(\lambda) \end{aligned}$$

as claimed.

Lemma 11. *Let \mathcal{C} , $(\mathcal{F}, \text{Puncture})$, and \mathbf{H} be as in Theorem 6. Then the argument system described in Fig. 4 is honest verifier zero-knowledge.*

Before we prove Lemma 11, we first state the following simple corollary which follows immediately from by combining Lemma 11 with Theorem 1.

Corollary 12. *Let \mathcal{C} , $(\mathcal{F}, \text{Puncture})$, and \mathbf{H} be as in Theorem 6. Then the argument system described in Fig. 4 is zero-knowledge.*

Proof (Lemma 11). By Definition 7 we need to show that there exists a simulator Sim , such that

$$\left| \begin{array}{l} \Pr[(V_t, \mathcal{T}) \leftarrow \text{SIMU}(V_0) : \mathcal{D}(V_0, V_t, \mathcal{T}) = 1] \\ - \Pr[(V_t, \mathcal{T}) \leftarrow \text{REAL}(V_0) : \mathcal{D}(V_0, V_t, \mathcal{T}) = 1] \end{array} \right| \leq \text{negl}(\lambda). \tag{6}$$

We specify the honest-verifier zero-knowledge simulator in Fig. 5 and use a series of game hops specified in Figs. 6 and 7 to prove that the above equation holds.

```

Sim( $V_0, V_t$ )
-----
 $d \leftarrow \{1, \dots, t\}$ 
 $k \leftarrow \{0, 1\}^{k(\lambda)}$ 
for  $1 \leq i < d$ 
     $V_i := \pi(V_{i-1}, \mathcal{F}(k, i))$ 
for  $t > i \geq d$ 
     $V_i := \pi^{-1}(V_{i+1}, \mathcal{F}(k, i + 1))$ 
 $c := \mathbf{H}(V_1, \dots, V_{t-1})$ 
 $k' \leftarrow \mathbf{Puncture}(k, d)$ 
return ( $c; d; k'$ )
    
```

Fig. 5. Honest-verifier zero knowledge simulator for the three-move protocol specified in Fig. 4

```

Game1( $V_0$ )
-----
 $d \leftarrow \{1, \dots, t\}$ 
 $k \leftarrow \{0, 1\}^{k(\lambda)}$ 
for  $1 \leq i < d$ 
     $V_i := \pi(V_{i-1}, \mathcal{F}(k, i))$ 
 $p \leftarrow \text{Perm}_\ell, \vec{r} \leftarrow \mathcal{R}^\ell$ 
 $V_t := \pi(V_0, p, \vec{r})$ 
for  $t > i \geq d$ 
     $V_i := \pi^{-1}(V_{i+1}, \mathcal{F}(k, i + 1))$ 
 $c := \mathbf{H}(V_1, \dots, V_{t-1})$ 
 $k' \leftarrow \mathbf{Puncture}(k, d)$ 
 $b \leftarrow \mathcal{D}(V_0, V_t, (c; d; k'))$ 
return  $b$ 
    
```

Fig. 6. Game 1 for the proof of honest verifier zero-knowledge.

We first observe that

$$\Pr[(V_t, \mathcal{T}) \leftarrow \text{SIMU}(V_0) : \mathcal{D}(V_0, V_t, \mathcal{T}) = 1] = \Pr[\text{Game}_1(V_0) = 1] \tag{7}$$

This is easily verified. V_t is chosen uniformly at random from all valid shufflings in both cases. Further, c, d and k' are all computed in exactly the same way. Similarly, we observe that

$$\Pr[(V_t, \mathcal{T}) \leftarrow \text{REAL}(V_0) : \mathcal{D}(V_0, V_t, \mathcal{T}) = 1] = \Pr[\text{Game}_4(V_0)] \tag{8}$$

This is also easily verified. Game_4 computes V_t and (c, k) in exactly the same way as P_0 , lets the honest verifier $\mathbf{V}_1(V_0, V_t, c)$ choose d and finally computes k' in exactly the same manner as P_2 . What remains is to bound the differences between each pair of consecutive games.

Hop from Game₁ to Game₂. The changes between the two games are purely syntactic. In Game_1 the final shuffling V_t is sampled uniformly at random from all valid shuffles of V_0 . In Game_2 the final shuffling V_t is computed as the composition of several intermediate *valid* shuffles. The shuffling at position d is chosen uniformly at random and *independently* from all other shuffles. Since all previous shuffles are valid, this makes V_d a uniformly random shuffling of V_0 . Further, since all following shuffles are valid and the shuffling of V_d was independent, this makes V_t a uniformly random shuffling of V_0 . Therefore V_t is distributed identically in both games. By the perfect and inverse rerandomizability of \mathcal{C} , it

Game ₂ (V ₀)	Game ₃ (V ₀)	Game ₄ (V ₀)
$d \leftarrow \{1, \dots, t\}$	$d \leftarrow \{1, \dots, t\}$	
$k \leftarrow \{0, 1\}^{k(\lambda)}$	$k \leftarrow \{0, 1\}^{k(\lambda)}$	$k \leftarrow \{0, 1\}^{k(\lambda)}$
for $1 \leq i < d$	for $1 \leq i \leq t$	for $1 \leq i \leq t$
$V_i := \pi(V_{i-1}, \mathcal{F}(k, i))$	$V_i := \pi(V_{i-1}, \mathcal{F}(k, i))$	$V_i := \pi(V_{i-1}, \mathcal{F}(k, i))$
$p \leftarrow \text{Perm}_\ell, \vec{r} \leftarrow \mathcal{R}^\ell$		
$V_d := \pi(V_{d-1}, p, \vec{r})$		
for $d < i \leq t$		
$V_i := \pi(V_{i-1}, \mathcal{F}(k, i))$		
$c := \text{H}(V_1, \dots, V_{t-1})$	$c := \text{H}(V_1, \dots, V_{t-1})$	$c := \text{H}(V_1, \dots, V_{t-1})$
		$d \leftarrow \mathbf{V}_1(V_0, V_t, c)$
$k' \leftarrow \text{Puncture}(k, d)$	$k' \leftarrow \text{Puncture}(k, d)$	$k' \leftarrow \text{Puncture}(k, d)$
$b \leftarrow \mathcal{D}(V_0, V_t, (c; d; k'))$	$b \leftarrow \mathcal{D}(V_0, V_t, (c; d; k'))$	$b \leftarrow \mathcal{D}(V_0, V_t, (c; d; k'))$
return b	return b	return b

Fig. 7. Game 2 through 4 for the proof of honest verifier zero-knowledge.

makes no difference, whether the V_i for $d < i < t$ are computed in the “forward” direction from V_{i-1} or in the “backwards” direction from V_{i+1} . Therefore the two games are perfectly equivalent, and it holds that

$$\Pr[\text{Game}_1(V_0) = 1] = \Pr[\text{Game}_2(V_0) = 1]. \tag{9}$$

Hop from Game₂ to Game₃. Note that the only difference between the two games is in the computation of V_d , which is computed as a uniformly random shuffle in *Game₂* and as a *pseudorandom* shuffle in *Game₃*. This means we can bound the difference between the two games using a reduction to the pseudorandomness of the puncturable pseudorandom function. Specifically we use \mathcal{D} as a distinguisher against \mathcal{F} by requesting a key punctured on d and after receiving k' and $y = (p, \vec{r})$, computing V_i with key k' as in *Game₂* except that we compute V_d using $y = (p, \vec{r})$. It is easy to see, that if y is uniformly random, then we perfectly simulate *Game₂*, whereas if $y = \mathcal{F}(k, d)$ we perfectly simulate *Game₃*. By the security of \mathcal{F} it must therefore hold that

$$|\Pr[\text{Game}_2(V_0) = 1] - \Pr[\text{Game}_3(V_0) = 1]| \leq \text{negl}(\lambda) \tag{10}$$

Hop from Game₃ to Game₄. The changes between the two games are again merely syntactic. In particular, the games behave identically, except that *Game₃* chooses d uniformly at random from $\{1, \dots, t\}$ whereas *Game₄* lets the verifier choose $d \leftarrow \mathbf{V}_0(V_0, V_t, c)$. However, by definition of \mathbf{V}_0 these two sampling strategies are identical and it holds that

$$\Pr[\text{Game}_3(V_0) = 1] = \Pr[\text{Game}_4(V_0) = 1]. \tag{11}$$

Finally, combining Eqs. 9 through 11, we get

$$|\Pr[\text{Game}_1(V_0) = 1] - \Pr[\text{Game}_4(V_0) = 1]| \leq \text{negl}(\lambda), \quad (12)$$

which combined with Eq. 7 and Eq. 8 gives us Eq. 6 thus concluding the proof. \square

In the full version of this paper, we additionally show how to modify the construction to achieve straightline simulation in strict polynomial time in the CRS model. We also show how the amortized efficiency of the construction can be improved through batching and how the verifier's computational overhead can be optimized at the cost of a slightly worse communication complexity.

6 Public Accountability

If a public-key infrastructure (PKI), which associates the prover with a public key pk , is available, then we can discourage malicious provers from attempting to cheat by ensuring that the verifier obtains a publicly verifiable certificate that attests any failed cheating attempt by the prover. In the context of blockchains, such a certificate could for example be used to punish the prover through financial penalties.

In the following definition, we define this property by requiring the existence of a judge algorithm that can verify valid certificates and cannot be fooled by invalid certificates that falsely accuse an honest prover of misbehavior. For the sake of readability, we implicitly assume that the verifier has access to the prover's public key and the prover has access to their own public and secret keys.

Definition 10 (Publicly-Accountable Zero-Knowledge Arguments). *Let $\mathcal{L} \subseteq X \times Y$ be a partially fixable language. Let (P, V) be an interactive zero-knowledge argument for \mathcal{L} with soundness error ϵ in the PKI model. We say that the argument system is publicly accountable, if there exists a PPT algorithm Setup and a deterministic polynomial time judge algorithm J , such that the following conditions hold:*

Accountability: *Fix any $(x, y) \notin \mathcal{L}$ and let P^* be a malicious probabilistic polynomial time prover with*

$$\Pr[\text{crs} \leftarrow \text{Setup}(1^\lambda); b \leftarrow \langle P^*(\text{crs}), V(\text{crs}, x, y) \rangle : b = 1] \geq \delta\epsilon,$$

where the probability is taken over the random coins of the prover and the verifier and $0 < \delta \leq 1$. Then it holds that

$$\begin{aligned} & \Pr[\text{crs} \leftarrow \text{Setup}(1^\lambda); \text{cert} \leftarrow \langle P^*(\text{crs}), V(\text{crs}, x, y) \rangle : J(\text{crs}, \text{pk}, \text{cert}) = 1] \\ & \geq \delta(1 - \epsilon) - \text{negl}(\lambda), \end{aligned}$$

where the probability is taken over the random coins of the prover and the verifier.

Defamation-Freeness: For any $x \in X$ with $\mathcal{L}_x \neq \emptyset$, for any honest prover P and malicious probabilistic polynomial time verifier V^* , it holds that

$$\Pr \left[\begin{array}{l} \text{crs} \leftarrow \text{Setup}(1^\lambda); \\ (y, \text{aux}) \leftarrow P_0(\text{crs}, x); \\ \text{cert} \leftarrow \langle P_1(x, y, \text{aux}), V^*(\text{crs}, x, y) \rangle \end{array} : J(\text{pk}, \text{cert}) = 1 \right] \leq \text{negl}(\lambda).$$

We show that the three move zero-knowledge shuffle argument (P, V) described in Fig. 4 in Sect. 5 can be transformed into a publicly-accountable zero-knowledge argument.

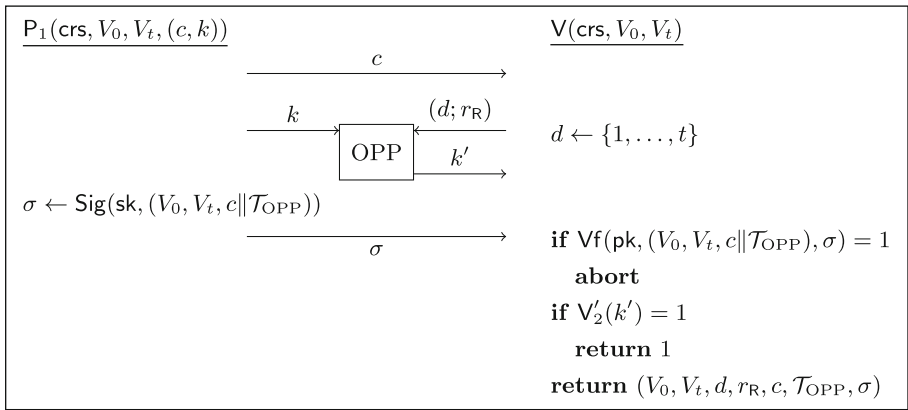


Fig. 8. The publicly-accountable transformation of the three-move shuffle argument. Here r_R refers to the random tape V uses to execute R in the OPP. \mathcal{T}_{OPP} refers to the full transcript resulting from the OPP execution.

Theorem 13. Let $\langle P', V' \rangle$ be the three-move zero-knowledge shuffle argument described in Fig. 4 in Sect. 5. Let (S, R) be a secure receiver-extractable, oblivious key puncturing protocol for the puncturable PRF used in $\langle P', V' \rangle$. Let $(\text{Gen}, \text{Sig}, \text{Vf})$ be an existentially unforgeable signature scheme. Then the protocol $\langle P, V \rangle$ with $P = (P'_0, P_1)$, with P_1 and V as specified in Fig. 8 is a publicly-accountable zero-knowledge argument with soundness error $1/t$.

Proof. We now show that this construction is indeed a publicly-accountable zero-knowledge argument.

Lemma 14. The argument system (P, V) is complete.

Proof. This directly follows from the completeness of the underlying argument system (P', V') , the completeness of the oblivious key puncturing protocol, and the correctness of the signature scheme. □

Lemma 15. The argument system (P, V) is sound with soundness error $1/t$.

Proof. Let $(V_0, V_t) \notin \text{Shuffle}_\ell$ and let P^* be an arbitrary malicious probabilistic polynomial time prover against (P, V) . We use P^* to construct a probabilistic polynomial time prover \tilde{P} against the original three move argument system (P', V') as follows. \tilde{P} initializes P^* with uniform random coins. Without loss of generality assume that P^* outputs a first message c , which \tilde{P} forwards to the verifier. The verifier outputs a challenge d . \tilde{P} , acting on behalf of a simulated verifier towards P^* , engages in an execution of the oblivious puncturing protocol as the receiver, where d is the choice index. Let k' be the output that \tilde{P} receives from P^* in this execution. \tilde{P} forwards k' to the verifier. It is straightforward to see that \tilde{P} 's success probability is at least as large as the success probability of P^* . Since by Theorem 6 \tilde{P} 's success probability is bounded by $1/t + \text{negl}(\lambda)$ the lemma follows. \square

Lemma 16. *The argument system (P, V) is zero-knowledge.*

Proof. Let V^* be an arbitrary verifier to which we have blackbox access. Let $(\text{Sim}_0, \text{Sim}_1)$ be the simulator from the sender simulation property of the OPP. We construct a verifier \tilde{V}' for the underlying argument system, which makes blackbox use of V^* and works as follows. First we let \tilde{V}' sample a signature key pair (sk, pk) and whenever V^* asks the (simulated) PKI for the verification key of the prover, our verifier will return pk . Then \tilde{V}' generates a simulated common reference string for the OPP as $(\text{crs}, \text{td}) \leftarrow \text{Sim}_0(1^\lambda)$ and initializes $V^*(\text{crs}, 1^\lambda)$. Upon receiving a first prover message c , verifier \tilde{V}' forwards the message to V^* . Then, \tilde{V}' uses $\text{Sim}_1(\text{crs}, \text{td})$ to execute the oblivious key puncturing protocol with V^* . If and when Sim_1 makes its single query d to its puncturing oracle, \tilde{V}' simulates the oracle by outputting d as its challenge message and returning the response k' as the answer. Finally, \tilde{V}' outputs whatever V^* outputs. Due to the sender simulation property of the OPP, the view of V^* when simulated by \tilde{V}' is computationally indistinguishable from a real execution. Therefore the output of \tilde{V}' is computationally indistinguishable from the output of V^* in a real execution of the publically accountable protocol. Since, \tilde{V}' is a verifier for the underlying three-move argument and by Corollary 12 that argument is zero-knowledge, there exists a zero-knowledge simulator Sim' that can simulate this output given only blackbox access to \tilde{V}' . We can, thus, define the zero-knowledge simulator Sim with blackbox access to V^* simply as Sim' with blackbox access to \tilde{V}' . \square

Lemma 17. *The argument system (P, V) satisfies accountability.*

Proof. The judge algorithm J is specified in Fig. 9. Let $(V_0, V_t) \notin \text{Shuffle}_\ell$, let P^* be an arbitrary malicious PPT prover. Let $\text{noAbort}(\vec{r})$ be the event that P^* does not abort and sends a valid signature when the parties run on random tapes $\vec{r} = (r_P, r_V)$. Let $d(\vec{r})$ be the function that returns the verifier's challenge. Without loss of generality, we assume that for any fixed first message of the prover, there exists a set of verifier challenges $D(\vec{r})$, which the prover successfully answers in such a way that the verifier accepts.

```

J(crs, pk, cert = (V0, Vt, d, rR, c, T, σ))
-----
if ∀f(pk, (V0, Vt, c||T), σ) = 0
    return 0
k' := TCheck(crs, T, d, rR)
if k' = ⊥
    return 0
b := V'2(k')
return b ⊕ 1
    
```

Fig. 9. The judge algorithm J for the publicly-accountable transformation of a three-move zero-knowledge argument. Here $TCheck$ refers to an algorithm that given a CRS, an OPP transcript \mathcal{T} , an input d and a random tape r_R outputs R 's output, if the transcript is consistent with d and r_R and \perp otherwise.

From the receiver privacy of the OPP, it follows that

$$|\Pr[\text{noAbort}(\vec{r}) \mid d(\vec{r}) \in D(\vec{r})] - \Pr[\text{noAbort}(\vec{r}) \mid d(\vec{r}) \notin D(\vec{r})]| \leq \text{negl}(\lambda),$$

where the probability is taken over the random coins \vec{r} of the parties. We observe that

$$\begin{aligned}
 & \Pr[\text{cert} \leftarrow \langle P^*(1^\lambda; r_P), V(x, y; r_V) \rangle : J(\text{pk}, \text{cert}) = 1 \mid (V_0, V_t) \notin \text{Shuffle}_\ell] \\
 & \geq \Pr[\text{noAbort}(\vec{r}) \mid d(\vec{r}) \notin D(\vec{r})] \cdot \Pr[d(\vec{r}) \notin D(\vec{r})] \\
 & \geq (\Pr[\text{noAbort}(\vec{r}) \mid d(\vec{r}) \in D(\vec{r})] - \text{negl}(\lambda)) \cdot \Pr[d(\vec{r}) \notin D(\vec{r})] \\
 & \geq (\delta - \text{negl}(\lambda)) \cdot (1 - \epsilon - \text{negl}(\lambda)) \\
 & \geq \delta(1 - \epsilon) - \text{negl}(\lambda)
 \end{aligned}$$

□

Lemma 18. *The argument system (P, V) is defamation-free.*

Proof. Let $(V_0, V_t, d, r_R, c, T, \sigma)$ be the certificate presented by V^* . We observe, that the existential unforgeability of the signature scheme implies that except with negligible probability, c and T must have originated from an interaction with the honest prover on input (V_0, V_t) . The completeness of the OPP implies that $TCheck$ will either output \perp , in which case J will output 0, or the correct response k' to the challenge d . It follows that (c, d, k') is the view of V' in an honest execution of the three-move shuffle argument. If $(V_0, V_t) \in \text{Shuffle}_\ell$, the completeness of the argument implies that have $V'_2(k') = 1$ and J will output 0. The lemma directly follows from the above observations. □

The theorem statement follows from Lemmas 15, 16, 17, and 18. □

7 Instantiation and Comparison

To evaluate the practical usefulness of the shuffle argument from Sect. 5 and its publicly accountable counterpart from Sect. 6, we explore how to instantiate

Table 1. Transcript size of shuffle arguments for vectors of length 100,000. The reported numbers for our constructions correspond to the instantiations described in Sect. 7 and are independent of the vector length and the type of commitment being shuffled. The numbers for Bayer-Groth [9] are taken from their paper for an instantiation in a q order subgroup of \mathbb{Z}_p^* with $|q| = 160$ and $|p| = 1024$ shuffling ElGamal ciphertexts. For Bulletproofs [17], we consider an instantiation in ristretto255 [48] for shuffling Pedersen commitments. For the Groth SNARK [32] we consider an instantiation with curve BLS12-381 [13] and observe that the numbers are independent of the kind of commitments being shuffled and the vector length.

Scheme	Assumptions	Trusted setup	Soundness	Communication cost (byte)
This Work	CRHF, PRG	None	2^{-2}	81
			2^{-5}	153
			$2^{-\gamma}$	$32 + \lceil \gamma \cdot 24 \frac{1}{8} \rceil$
This Work (with accountability)	CRHF, PRG, DDH	CRS	2^{-2}	416
			2^{-5}	992
			$2^{-\gamma}$	$32 + \gamma \cdot 192$
This Work (with accountability)	CRHF, PRG, RLWE	CRS	2^{-2}	$0.44 \cdot 2^{20}$
			2^{-5}	$1.1 \cdot 2^{20}$
			$2^{-\gamma}$	$32 + \gamma \cdot 0.215 \cdot 2^{20}$
Bayer-Groth [9]	Discrete logarithm	CRS	$\text{negl}(\lambda)$	700,000
Bulletproofs [17]	Discrete logarithm	CRS	$\text{negl}(\lambda)$	1,600
SNARKs [32]	Generic bilinear group	SRS	$\text{negl}(\lambda)$	144

them in practice. We are particularly interested in the concrete communication complexities of such instantiations, since this is where our constructions shine. Towards this goal, we need to pick specific instantiations of the underlying building blocks, such as the collision-resistant hash function and the PPRF with its oblivious puncturing protocol. We aim for a security level of roughly 128 bits.

The Hash Function. The collision resistant hash function can in practice be instantiated using any of SHA-256 [46], SHA3-256, or SHAKE256 [47] with 256 bit output length. Any of these instantiations leads to hashes of size $|c| = 256$ bits.

The Puncturable PRF. The PPRF can be instantiated using the construction of Goldreich, Goldwasser and Micali (GGM) [28]. This construction relies on an internal length-doubling pseudorandom generator which can be instantiated using a secure stream cipher, such as AES [4] in CTR mode or ChaCha20 [11]. Taking the losses introduced by both the security proof of GGM as well as the proof of zero-knowledge into account, we require a PRG with $128 + \lceil 2 \log t + \log \log t \rceil$ bits of security to achieve our security goal of 128 bits. For reasonable values of t , using AES-192 in CTR mode would thus suffice. Simply instantiating GGM like this yields a PPRF with range $\{0, 1\}^{192}$. As explained in Sect. 2, to

construct a PPRF with range $\text{Perm}_\ell \times \mathcal{R}^\ell$, the output can be stretched using a PRG and combined with the Fisher-Yates shuffle [24] by using the stretched output as the random tape of the shuffling algorithm. The necessary PRG can again be instantiated using, e.g., AES-192 in CTR mode or ChaCha20. Overall, the size of a punctured key with the AES-192 instantiation is $|k'| = \lceil \log t \rceil \cdot 192$ bits.

The Oblivious Puncturing Protocol. Using the PPRF construction of GGM mentioned above, we can use the oblivious puncturing protocol described in [14], which itself relies on $\lceil \log t \rceil$ many oblivious transfers with active security. These can be instantiated with the 2-round UC secure protocol of Peikert, Vaikunathan, and Waters [44] over ristretto255 [48] by relying on the DDH assumption. From a computational point of view, the parties need to perform $10 \lceil \log t \rceil$ exponentiations, $6 \lceil \log t \rceil$ multiplications, and $2t$ evaluations of a PRG for one oblivious puncturing. Using our instantiation, the OPP runs in two rounds and thus the overall protocol runs in three, since the signature σ can be sent in parallel with the second message of the OPP from the prover to the verifier.

To obtain post-quantum security, we can instantiate the oblivious transfer with the protocol of Micciancio and Sorrell [41], which relies on the ring learning with errors assumption.

Acknowledgments. The authors would like to thank the anonymous PKC reviewer for pointing out the efficient OPP instantiation implicit in [14] as a practical replacement for general purpose PIR. The authors would also like to thank Ivan Damgård for the insightful discussions about the Goldreich and Kahan proof technique as well as Diego F. Aranha and Cathie Yun for information about practical instantiations of SNARKs and Bulletproofs respectively.

References

1. Abe, M.: Universally verifiable mix-net with verification work independent of the number of mix-servers. In: Nyberg, K. (ed.) EUROCRYPT 1998. LNCS, vol. 1403, pp. 437–447. Springer, Heidelberg (1998). <https://doi.org/10.1007/BFb0054144>
2. Abe, M.: Mix-networks on permutation networks. In: Lam, K.-Y., Okamoto, E., Xing, C. (eds.) ASIACRYPT 1999. LNCS, vol. 1716, pp. 258–273. Springer, Heidelberg (1999). https://doi.org/10.1007/978-3-540-48000-6_21
3. Abe, M., Hoshino, F.: Remarks on mix-network based on permutation networks. In: Kim, K. (ed.) PKC 2001. LNCS, vol. 1992, pp. 317–324. Springer, Heidelberg (2001). https://doi.org/10.1007/3-540-44586-2_23
4. Advanced Encryption Standard (AES). National Institute of Standards and Technology (NIST), FIPS PUB 197, U.S. Department of Commerce, November 2001
5. Ames, S., Hazay, C., Ishai, Y., Venkatasubramanian, M.: Liger: lightweight sub-linear arguments without a trusted setup. In: ACM CCS 2017, pp. 2087–2104 (2017)
6. Asharov, G., Orlandi, C.: Calling out cheaters: covert security with public verifiability. In: Wang, X., Sako, K. (eds.) ASIACRYPT 2012. LNCS, vol. 7658, pp. 681–698. Springer, Heidelberg (2012). https://doi.org/10.1007/978-3-642-34961-4_41

7. Aumann, Y., Lindell, Y.: Security against covert adversaries: efficient protocols for realistic adversaries. In: Vadhan, S.P. (ed.) TCC 2007. LNCS, vol. 4392, pp. 137–156. Springer, Heidelberg (2007). https://doi.org/10.1007/978-3-540-70936-7_8
8. Aviram, N., Gellert, K., Jager, T.: Session resumption protocols and efficient forward security for TLS 1.3 0-RTT. In: Ishai, Y., Rijmen, V. (eds.) EUROCRYPT 2019. LNCS, vol. 11477, pp. 117–150. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-17656-3_5
9. Bayer, S., Groth, J.: Efficient zero-knowledge argument for correctness of a shuffle. In: Pointcheval, D., Johansson, T. (eds.) EUROCRYPT 2012. LNCS, vol. 7237, pp. 263–280. Springer, Heidelberg (2012). https://doi.org/10.1007/978-3-642-29011-4_17
10. Bentov, I., Kumaresan, R., Miller, A.: Instantaneous decentralized poker. In: Takagi, T., Peyrin, T. (eds.) ASIACRYPT 2017. LNCS, vol. 10625, pp. 410–440. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-70697-9_15
11. Bernstein, D.J.: Chacha, a variant of salsa20. In: Workshop Record of SASC, vol. 8, pp. 3–5 (2008)
12. Boneh, D., Waters, B.: Constrained pseudorandom functions and their applications. In: Sako, K., Sarkar, P. (eds.) ASIACRYPT 2013. LNCS, vol. 8270, pp. 280–300. Springer, Heidelberg (2013). https://doi.org/10.1007/978-3-642-42045-0_15
13. Bowe, S.: BLS12-381: New zk-SNARK Elliptic Curve Construction, March 2017. <https://electriccoin.co/blog/new-snark-curve/>
14. Boyle, E., et al.: Efficient two-round OT extension and silent non-interactive secure computation. In: ACM CCS 2019, pp. 291–308 (2019)
15. Boyle, E., Goldwasser, S., Ivan, I.: Functional signatures and pseudorandom functions. In: Krawczyk, H. (ed.) PKC 2014. LNCS, vol. 8383, pp. 501–519. Springer, Heidelberg (2014). https://doi.org/10.1007/978-3-642-54631-0_29
16. Brakerski, Z., Vaikuntanathan, V.: Constrained key-homomorphic PRFS from standard lattice assumptions. In: Dodis, Y., Nielsen, J.B. (eds.) TCC 2015. LNCS, vol. 9015, pp. 1–30. Springer, Heidelberg (2015). https://doi.org/10.1007/978-3-662-46497-7_1
17. Büinz, B., Bootle, J., Boneh, D., Poelstra, A., Wuille, P., Maxwell, G.: Bulletproofs: short proofs for confidential transactions and more. In: 2018 IEEE Symposium on Security and Privacy, pp. 315–334 (2018)
18. Chaum, D.: Untraceable electronic mail, return addresses, and digital pseudonyms. *Commun. ACM* **24**(2), 84–88 (1981)
19. Chung, K.-M., Lui, E., Pass, R.: From weak to strong zero-knowledge and applications. In: Dodis, Y., Nielsen, J.B. (eds.) TCC 2015. LNCS, vol. 9014, pp. 66–92. Springer, Heidelberg (2015). https://doi.org/10.1007/978-3-662-46494-6_4
20. Dagher, G.G., Büinz, B., Bonneau, J., Clark, J., Boneh, D.: Provisions: Privacy-preserving proofs of solvency for bitcoin exchanges. In: ACM CCS 2015, pp. 720–731 (2015)
21. Damgård, I., Orlandi, C., Simkin, M.: Black-box transformations from passive to covert security with public verifiability. In: Micciancio, D., Ristenpart, T. (eds.) CRYPTO 2020. LNCS, vol. 12171, pp. 647–676. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-56880-1_23
22. Fauzi, P., Meiklejohn, S., Mercer, R., Orlandi, C.: Quisquis: a new design for anonymous cryptocurrencies. In: Galbraith, S.D., Moriai, S. (eds.) ASIACRYPT 2019. LNCS, vol. 11921, pp. 649–678. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-34578-5_23

23. Fiat, A., Shamir, A.: How to prove yourself: practical solutions to identification and signature problems. In: Odlyzko, A.M. (ed.) CRYPTO 1986. LNCS, vol. 263, pp. 186–194. Springer, Heidelberg (1987). https://doi.org/10.1007/3-540-47721-7_12
24. Fisher, R.A., Yates, F.: Statistical Tables for Biological, Agricultural and Medical Research, 3rd edn. Oliver and Boyd, Edinburgh (1949)
25. Furukawa, J., Sako, K.: An efficient scheme for proving a shuffle. In: Kilian, J. (ed.) CRYPTO 2001. LNCS, vol. 2139, pp. 368–387. Springer, Heidelberg (2001). https://doi.org/10.1007/3-540-44647-8_22
26. Gentry, C., Wichs, D.: Separating succinct non-interactive arguments from all falsifiable assumptions. In: 43rd ACM STOC, pp. 99–108 (2011)
27. Goldreich, O.: A uniform-complexity treatment of encryption and zero-knowledge. *J. Cryptol.* **6**(1), 21–53 (1993). <https://doi.org/10.1007/BF02620230>
28. Goldreich, O., Goldwasser, S., Micali, S.: How to construct random functions (extended abstract). In: 25th FOCS, pp. 464–479 (1984)
29. Goldreich, O., Kahan, A.: How to construct constant-round zero-knowledge proof systems for NP. *J. Cryptol.* **9**(3), 167–189 (1996). <https://doi.org/10.1007/BF00208001>
30. Groth, J.: A verifiable secret shuffle of homomorphic encryptions. In: Desmedt, Y.G. (ed.) PKC 2003. LNCS, vol. 2567, pp. 145–160. Springer, Heidelberg (2003). https://doi.org/10.1007/3-540-36288-6_11
31. Groth, J.: Linear algebra with sub-linear zero-knowledge arguments. In: Halevi, S. (ed.) CRYPTO 2009. LNCS, vol. 5677, pp. 192–208. Springer, Heidelberg (2009). https://doi.org/10.1007/978-3-642-03356-8_12
32. Groth, J.: On the size of pairing-based non-interactive arguments. In: Fischlin, M., Coron, J.-S. (eds.) EUROCRYPT 2016. LNCS, vol. 9666, pp. 305–326. Springer, Heidelberg (2016). https://doi.org/10.1007/978-3-662-49896-5_11
33. Groth, J., Ishai, Y.: Sub-linear zero-knowledge argument for correctness of a shuffle. In: Smart, N. (ed.) EUROCRYPT 2008. LNCS, vol. 4965, pp. 379–396. Springer, Heidelberg (2008). https://doi.org/10.1007/978-3-540-78967-3_22
34. Groth, J., Lu, S.: Verifiable shuffle of large size ciphertexts. In: Okamoto, T., Wang, X. (eds.) PKC 2007. LNCS, vol. 4450, pp. 377–392. Springer, Heidelberg (2007). https://doi.org/10.1007/978-3-540-71677-8_25
35. Ishai, Y., Kushilevitz, E., Ostrovsky, R., Sahai, A.: Zero-knowledge from secure multiparty computation. In: 39th ACM STOC, pp. 21–30 (2007)
36. Kiayias, A., Papadopoulos, S., Triandopoulos, N., Zacharias, T.: Delegatable pseudorandom functions and applications. In: ACM CCS 2013, pp. 669–684 (2013)
37. Kilian, J.: A note on efficient zero-knowledge proofs and arguments (extended abstract). In: 24th ACM STOC, pp. 723–732 (1992)
38. Kwon, A., Lazar, D., Devadas, S., Ford, B.: Riffle: an efficient communication system with strong anonymity. *PoPETs* **2016**(2), 115–134 (2016)
39. Lindell, Y.: How to simulate it - A tutorial on the simulation proof technique. *Cryptology ePrint Archive*, Report 2016/046 (2016). <http://eprint.iacr.org/2016/046>
40. Micali, S.: CS proofs (extended abstracts). In: 35th FOCS, pp. 436–453 (1994)
41. Micciancio, D., Sorrell, J.: Simpler statistically sender private oblivious transfer from ideals of cyclotomic integers. In: Moriai, S., Wang, H. (eds.) ASIACRYPT 2020. LNCS, vol. 12492, pp. 381–407. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-64834-3_13
42. Neff, C.A.: A verifiable secret shuffle and its application to e-voting. In: ACM CCS 2001, pp. 116–125 (2001)

43. Pedersen, T.P.: Non-interactive and information-theoretic secure verifiable secret sharing. In: Feigenbaum, J. (ed.) CRYPTO 1991. LNCS, vol. 576, pp. 129–140. Springer, Heidelberg (1992). https://doi.org/10.1007/3-540-46766-1_9
44. Peikert, C., Vaikuntanathan, V., Waters, B.: A framework for efficient and composable oblivious transfer. In: Wagner, D. (ed.) CRYPTO 2008. LNCS, vol. 5157, pp. 554–571. Springer, Heidelberg (2008). https://doi.org/10.1007/978-3-540-85174-5_31
45. Sako, K., Kilian, J.: Receipt-free mix-type voting scheme. In: Guillou, L.C., Quisquater, J.-J. (eds.) EUROCRYPT 1995. LNCS, vol. 921, pp. 393–403. Springer, Heidelberg (1995). https://doi.org/10.1007/3-540-49264-X_32
46. Secure Hash Standard (SHS). National Institute of Standards and Technology, NIST FIPS PUB 180–4, U.S. Department of Commerce, August 2015
47. SHA-3 Standard: Permutation-Based Hash and Extendable-Output Functions. National Institute of Standards and Technology, NIST FIPS PUB 202, U.S. Department of Commerce, Aug 2015
48. de Valence, H., Grigg, J., Tankersley, G., Valsorda, F., Isis Lovecruft, M.H.: The ristretto255 and decaf448 groups. IETF CFRG Internet Draft (2020)