

Private Access Control for Function Secret Sharing^{*}

Sacha Servan-Schreiber¹, Simon Beyzerov^{2,3}, Eli Yablon^{2,4}, and Hyojae Park^{2,4}

¹ MIT CSAIL

² MIT PRIMES

³ Massachusetts Academy of Math and Science at WPI

⁴ Sharon High School

November 2, 2023

Abstract. Function Secret Sharing (FSS; Eurocrypt 2015) allows a dealer to share a function f with two or more evaluators. Given secret shares of a function f , the evaluators can locally compute secret shares of $f(x)$ for any input x , without learning information about f .

In this paper, we initiate the study of access control for FSS. Given the shares of f , the evaluators can ensure that the dealer is authorized to share the provided function. For a function family \mathcal{F} and an access control list defined over the family, the evaluators receiving the shares of $f \in \mathcal{F}$ can efficiently check that the dealer knows the access key for f .

This model enables new applications of FSS, such as: (1) anonymous authentication in a multi-party setting, (2) access control in private databases, and (3) authentication and spam prevention in anonymous communication systems.

Our definitions and constructions abstract and improve the concrete efficiency of several recent systems that implement ad-hoc mechanisms for access control over FSS. The main building block behind our efficiency improvement is a discrete-logarithm zero-knowledge proof-of-knowledge over secret-shared elements, which may be of independent interest.

We evaluate our constructions and show a 50–70 \times reduction in computational overhead compared to existing access control techniques used in anonymous communication. In other applications, such as private databases, the processing cost of introducing access control is only 1.5–3 \times , when amortized over databases with 500,000 or more items.

Keywords: Function secret sharing, verifiable FSS, access control, authentication, anonymous communication, private databases, zero-knowledge multi-verifier proofs

^{*} This is the extended and corrected version of [48].

Table of Contents

1	Introduction	3
1.1	Background on FSS	4
1.2	Prior work	4
1.3	Followup work	5
1.4	Contributions	5
2	Background and Preliminaries	6
2.1	Access Control Lists (ACLs)	6
2.2	Notation and cryptographic preliminaries	7
3	Private Access Control Lists	8
3.1	Defining PACLs	8
3.2	Symmetric-key PACL	10
3.3	Security against malicious verifiers	10
3.4	Key distribution	11
4	PACLs for Distributed Point Functions	11
4.1	DPF-PACL constructions	11
4.1.1	DPF-PACL for match predicate	12
4.1.2	DPF-PACL for inclusion predicates	12
4.2	Optimizations and extensions	14
4.3	Aggregating PACLs	15
4.4	Security analysis	15
5	Faster PACLs for DPFs from Verifiable DPFs	17
5.1	Schnorr Proof over Secret Shares (SPoSS)	17
5.1.1	Protocol overview	18
5.2	VDPF-PACL using SPoSS	20
5.3	Security analysis	20
6	PACLs for FSS from DPF-PACLs	22
7	Generic PACLs from distributed zero-knowledge proofs	23
7.1	Preliminaries	23
7.2	Construction	25
8	Implementation and evaluation	27
8.1	Prover costs	27
8.2	Communication costs	27
8.3	Verification costs	29
8.4	Applications of PACLs	29
9	Conclusion	31
A	Verifiable DPFs	36
B	Deferred proofs	36
B.1	Proof of Theorem 4 (security of Algorithm 2)	36
B.2	Security of SPoSS	38
C	Beaver’s Protocol	39
D	Naive SPoSS using SNIPs	39

1 Introduction

Function secret sharing (FSS) [9, 11] is at the core of many privacy-preserving systems, including private databases [21, 22, 54], private telemetry [8], privacy-preserving machine learning [40, 45], distributed oblivious RAM (ORAM) [26], anonymous communication [19, 27, 41, 53], and efficient multi-party computation [12]. Since these applications involve the processing of private user data, often in settings where users may be behaving maliciously, access control becomes an important problem [27, 41, 53]. For example, in applications of FSS that involve privately reading from—or writing to—a database [8, 26, 27, 41, 53], access control is necessary to prevent malicious users from accessing or overwriting data belonging to other users.

FSS lets any user (called the *dealer*) distribute succinct secret shares of a function to a set of function evaluators. These evaluators can efficiently—and non-interactively—evaluate the function on a common input x to obtain secret shares of $f(x)$. FSS guarantees that the function remains private to any strict subset of the evaluators, which means that the evaluators do not learn anything about f (beyond the function family that f belongs to).

In this paper, we investigate the problem of privately enforcing access control in the context of FSS. We identify several existing applications of FSS that construct different ad-hoc solutions for access control [27, 41, 53], highlighting the utility of formally studying this paradigm.

For example, FSS is often used for private information retrieval (PIR) [9, 16, 30, 54]. In PIR, a dealer secret shares a selection function f_i with the evaluators. The evaluators use the shares to evaluate f_i on a database \mathcal{DB} and send back secret shares $[f_i(\mathcal{DB})]$, which encodes the i -th item in the database. The dealer then recovers the i -th item by combining the returned shares. Importantly, the evaluators who are given shares of f_i learn nothing about f_i (beyond the fact that f_i is from the “selection function” family) and therefore do not learn which item the dealer retrieved from \mathcal{DB} .

In the PIR setting, access control could require that only users with an *access key* for the i -th item in the database can successfully share f_i with the evaluators. More specifically, in applications involving e-commerce [34, 35], web queries [54], and media consumption [33], where users are only entitled to retrieve some (but not all) items in the database, such access control is imperative. Likewise, in private information *writing* applications, such as anonymous communication systems [27, 41, 53] and private telemetry [8, 18], access control is crucial to prevent attacks by malicious users sending invalid writes (e.g., by overwriting mailboxes of honest users [27, 41, 53]). Only users with a valid access key for the i -th database row should be able to write to it.

Defining the problem of private access control. We model access control as a one-to-one mapping between functions and keys. Each function (in a family of functions) is mapped to a verification key and an access key. The evaluators hold the verification keys, and a dealer has one or more access keys (we discuss key distribution in Section 3.4). A dealer secret shares the function f_i through FSS and, using the corresponding access key, provides a proof π proving access rights to f_i under some subset of verification keys. The evaluators (whom we also call the *verifiers*) can check the proof π using the verification keys (without learning which keys were used) and decide whether or not the dealer is entitled to an evaluation of the function f_i . For example, in the PIR setting, knowledge of the access key for the selection function f_i , allows a user to distribute secret shares of f_i to the evaluators along with a proof π . The evaluators check π before evaluating the function to ensure that the user is entitled to retrieve the i -th item (without learning i).

Challenges. Privately enforcing access control over a secret-shared function is challenging. As mentioned above, the evaluators are oblivious to the function they are evaluating, which bars obvious approaches to access control. That is, access control must maintain the privacy of the function (see Section 1.1), even in the presence of *malicious* evaluators that may try to exploit the access control mechanism to violate privacy (this is a problem when designing *any* form of verification over FSS [8, 11, 23]). Additionally, FSS is concerned with *efficiency* (computation and communication overheads for the dealer and the evaluators). As such, the access control mechanism must preserve the efficiency guarantees of the FSS scheme. Preventing malicious evaluators from violating privacy, without relying on strong assumptions or inefficient solutions, can be difficult [7, 11, 23].

Goals. We identify *efficiency* and *malicious security* as the primary goals when modeling and designing access control for FSS. More specifically—and following the requirements of FSS—we will require *communication efficiency* and *minimal interaction* between verifiers. Our definition (described in Section 3) captures these efficiency requirements by demanding (1) succinct proofs, (2) no interaction

with the prover, and (3) at most *one* message exchanged between verifiers to check access rights (note that this is in fact optimal as it is necessary to exchange one message to verify any proof over secret-shares [7]). Importantly, by minimizing interaction between verifiers, we also obtain security against any subset of malicious verifiers. More concretely, (3) ensures that any construction satisfying our model never provides “feedback” to any subset of malicious verifiers which, in turn, ensures that malicious verifiers obtain no information through the access control mechanism.

1.1 Background on FSS

FSS [9] takes a different approach to “traditional” secret sharing of data. With traditional secret sharing, a dealer shares a value v with a set of s parties such that (1) knowledge of up to some threshold number of shares does not reveal any information on v and (2) shares can be efficiently recombined to recover v . FSS applies the same idea to *functions* where the dealer instead secret shares a description of a function $f : \{0, 1\}^n \rightarrow \{0, 1\}^*$ with a set of s evaluators. The advantage of FSS (over traditional secret sharing) is *succinctness* of the secret shares, requiring the shares of f to be much smaller than the truth table of f . Denote the shares of f as $[f] := ([f]_1, \dots, [f]_s)$. The evaluators can then locally compute shares $[y] := [f(x)]$ on any input x . Informally, FSS must satisfy three properties:

- **Correctness.** The j -th evaluator can evaluate their secret share of f on a public input x to obtain a secret share of $f(x)$.
- **Privacy.** A secret share $[f]_i$ reveals no information on f .⁵
- **Efficiency.** The secret shares $[f]$ are succinct (sublinear in the size of the truth table for f).

Boyle et al. [9, 11] provide constructions for several function families. Specifically, they describe efficient FSS constructions for NC^0 functions, constant conjunction search queries, and interval functions. Their constructions are based only on the assumption that one-way functions exist [9]. The main primitive behind their results is an FSS family for distributed point functions (DPFs) [9, 11, 30]. Subsequent work of Boyle et al. [11] extends DPFs to FSS for *decision trees* and products of distinct secret-shared functions. FSS schemes from stronger cryptographic assumptions yield constructions for *all* efficient function families [9, 10, 25].

1.2 Prior work

Recent work on anonymous communication provides ad-hoc solutions to access control in the context of FSS. Express [27] (USENIX 2021), Spectrum [41] (NSDI 2022), and Sabre [53] (S&P 2022) use FSS for anonymous communication. In these systems, users privately write messages into mailboxes using a DPF. To prevent malicious users from corrupting mailboxes belonging to honest users, all three systems require a form of access control applied over DPFs, which they enforce through a lightweight multi-party computation protocol.

The access control mechanism in Express associates each mailbox with a secret “address.” The evaluators keep the addresses secret. Only a dealer with knowledge of the secret mailbox address can successfully write to that mailbox. Unfortunately, the mechanism used in Express has several drawbacks: (1) it does not generalize to larger families of DPFs and currently remains specific to the two-party DPF construction for point functions [9, 11], (2) the use of “addresses” for access control requires a large output range and leads to a $5\times$ computational overhead for the evaluators (and, more importantly, prevents optimization techniques for DPFs [11]), and (3) the multi-party computation requires extra communication between evaluators and the dealer. In contrast, our model and constructions (Sections 4 and 5) require only one message exchanged between verifiers and no interaction with the dealer. Additionally, our constructions make minimal assumptions on the underlying DPF scheme, making them compatible with DPF optimizations [11, 23].

Sabre extends Express by developing a different access control mechanism using zero-knowledge proofs over secret shares. Their techniques reduce the computational overhead on the evaluators (especially in the context of anonymous communication where many users are assumed to be malicious) at the cost of significantly increasing communication between the dealer and the evaluators. Like

⁵ Function secret sharing assumes computationally bounded adversaries compared to the information-theoretic security guarantees of traditional secret sharing such as Shamir [49].

Express, Sabre is designed around two-party DPF constructions [9, 11] and requires a round of interaction between the evaluators to enforce access control. The access control mechanism of Sabre is actually a special case of a generic approach to access control realized via zero-knowledge proofs over secret-shares, which we describe in Section 7. However, their techniques are tailored to the anonymous communication setting where (1) a large number of users are assumed to be malicious and (2) where access control is verified in batches.

Spectrum uses yet a different technique for access control via secret-shared hashing. The idea behind Spectrum is to have the evaluators “hash” the value being written to each mailbox using a unique hash key associated with the mailbox. The evaluators only process writes from a dealer that can prove, in zero-knowledge, that it knows the resulting hash value (which, in turn, proves that the dealer knows the hash key of the mailbox). The technique is communication efficient and only requires one message exchanged between verifiers to enforce access control (which aligns with our modeling of access control). In Section 4, we start by abstracting the access control construction of Spectrum and fitting it into our definition. We then generalize it further and develop new techniques to realize more efficient private access control schemes.

In Section 8, we show that our improved constructions reduce the computational overhead by 50–70× in both Express and Spectrum and have 1,000× smaller proof sizes compared to Sabre.

1.3 Followup work

Ji, Zhang, and Ren [36] improve the performance of public-key PACLs by a factor of 2.5–3× and provide several improvements to our schemes to support fine-grained access policies. Importantly, they also identify an attack on our SPoSS-based public-key PACL scheme (Section 5.2) if it is instantiated with a (verifiable) DPF that has an arbitrary output message at the special index (our work implicitly assumed that the verifiable DPF has an output $m = 1$ and rejects otherwise, as done in related work with extractable DPFs [8]). However, to address this oversight, we updated the construction (Algorithm 4) to make this requirement on the output message explicit. We stress that the concrete performance and evaluation of our construction is *not* impacted by this change, since our implementation uses a verifiable DPF with an output in \mathbb{Z}_2 as opposed to \mathbb{Z}_p . We note that Ji et al. [36] provide an improved and extended SPoSS-based public-key PACL construction for the general setting where $m \in \mathbb{Z}_p$ is arbitrary, at very minimal performance cost.

1.4 Contributions

We make the following five contributions:

Contribution I. A model for Private Access Control Lists (PACLs) for FSS. Our definitions capture the functionality requirements of several existing applications of access control for FSS [27, 41, 53] and demand a stringent set of efficiency requirements that align closely with the goals of FSS.

Contribution II. PACL constructions for black-box DPFs and lightweight FSS classes derived from DPFs. Our constructions are secure against malicious provers, guarantee privacy in the face of malicious verifiers, have a constant-factor overhead (relative to sharing and evaluating the function), and can be used as drop-in replacements in existing applications for significant efficiency improvements.

Contribution III. For the special case of *verifiable* FSS [23], we construct an optimized public-key PACL for black-box verifiable DPFs (which gives rise to PACLs for a large class of verifiable FSS). For this construction, we develop a zero-knowledge proof of discrete-logarithm knowledge over *secret-shared* group elements. Our construction has 2,400× smaller proof sizes compared to a naive approach and is possibly of independent interest.

Contribution IV. PACLs for FSS for functions in P/poly (not just classes of FSS derived from DPFs). Our generic construction is based on non-interactive zero-knowledge proofs over secret shares, instantiated in the random oracle model.

Contribution V. An open-source implementation which we evaluate on several canonical applications, such as anonymous user authentication in a distributed setting, access control in private databases, and anonymous communication.

2 Background and Preliminaries

Here, we define **non-private** access control for functions. We define **private** access control for FSS in [Section 3](#).

2.1 Access Control Lists (ACLs)

We define ACLs from a cryptographic lens in order to facilitate the definitions of *private* ACLs, which we introduce in [Section 3](#). Specifically, we define ACLs as a set of access and verification keys associated with a collection of objects (in our case, functions). An access key can be used to prove access rights for the function under the corresponding verification key.

Definition 1 (Access Control Lists). *Let $\lambda \in \mathbb{N}$ be a security parameter, $\mathcal{F}: \{0, 1\}^n \rightarrow \{0, 1\}^*$ be a function family, and $f_i \in \mathcal{F}$. An ACL scheme consists of an access control list A_λ (parameterized by λ) containing verification keys and an efficiently computable predicate $\text{CheckAccess}(A_\lambda, f_i, \text{sk})$ that outputs yes if and only if the access key sk satisfies a relation R with respect to the verification key associated with f_i in A_λ . For notational convenience, we let $N := |\mathcal{F}|$ and omit the λ subscript when it is clear from context.*

We note that [Definition 1](#) is general and not specific to FSS (indeed, [Definition 1](#) does not even capture the notion of secret shares or distributed verifiers). We will define these notions in [Section 3](#) when formalizing private ACLs for FSS. It is also natural to equip [Definition 1](#) with *completeness* and *soundness* properties (with respect to an adversary). These are likewise deferred to the formalization of private ACLs in [Definitions 3](#) and [4](#).

We now describe several instantiations of CheckAccess from [Definition 1](#).

Match predicate. In most cases, the goal of CheckAccess is to check if the provided access key matches with some unique verification key associated with the function f_i . In this setting, only dealers that have knowledge of the unique access key corresponding to the function f_i should be able to secret-share it with the evaluators. The *match* predicate consists of an efficiently computable relation $R(\cdot, \cdot)$, N verification keys $(\text{vk}_1, \dots, \text{vk}_N)$ and access keys $(\text{sk}_1, \dots, \text{sk}_N)$, such that $R(\text{vk}_i, \text{sk}_j) = 1$ if and only if $i = j$ and each tuple is uniquely associated with a canonical instance of $f_i \in \mathcal{F}$. CheckAccess is defined as:

$$\begin{array}{l} \text{CheckAccess}(A_\lambda, f_i, \text{sk}) \\ \hline \text{parse } A_\lambda = (\text{vk}_1, \dots, \text{vk}_N) \\ \text{if } R(\text{vk}_i, \text{sk}) = 1 \text{ return yes,} \\ \text{else return no} \end{array}$$

That is, CheckAccess outputs yes if and only if the provided sk is related to the verification key associated with f_i .

Inclusion predicate. A generalization of the match predicate satisfying [Definition 1](#) is the *inclusion* predicate that associates each function with $\ell \geq 1$ verification keys,

$$A_\lambda := \begin{pmatrix} (\text{vk}_{1,1} \dots \text{vk}_{1,\ell}) \\ \vdots \\ (\text{vk}_{N,1} \dots \text{vk}_{N,\ell}) \end{pmatrix}.$$

Any key $\text{vk}_{i,j}$ in the row $(\text{vk}_{i,1} \dots \text{vk}_{i,\ell})$ can be used to satisfy the relation for f_i . CheckAccess is defined as:

$$\begin{array}{l} \text{CheckAccess}(A_\lambda, f_i, \text{sk}) \\ \hline \text{parse } A_\lambda = ((\text{vk}_{1,1}, \dots, \text{vk}_{1,\ell}), \dots, (\text{vk}_{N,1}, \dots, \text{vk}_{N,\ell})) \\ \text{if } \exists \text{vk}_{i,j} \text{ such that } R(\text{vk}_{i,j}, \text{sk}) = 1 \text{ return yes,} \\ \text{else return no} \end{array}$$

That is, `CheckAccess` outputs `yes` if and only if `sk` matches with *any* of the verification keys associated with f_i . Having multiple access keys per function enables efficient access key revocation in a setting with many dealers: an access key can be removed for a given function *without* impacting the validity of the remaining access keys.

Boolean predicate. A generalization of the inclusion predicate satisfying [Definition 1](#) is a monotone boolean predicate (ANDs and ORs) over a list of `CheckAccess` outputs. Here, Λ_λ consists of ℓ sublists $\Lambda_\lambda^{(1)}, \dots, \Lambda_\lambda^{(\ell)}$ and a predicate P defined over the bits $b_i \leftarrow \text{CheckAccess}_i(\Lambda_\lambda^{(i)}, f_i, \text{sk})$ for $i \in \{1, \dots, \ell\}$. `CheckAccess`($\Lambda_\lambda, f_i, \text{sk}$) outputs `yes` if and only if $P(\text{sk}, b_1, \dots, b_\ell) = 1$.

Remark 1 (Access Control Theory). Constructing effective access control policies for different real-world scenarios has been studied from a theoretical lens [[5](#), [14](#), [17](#)]. Our access predicates are fully compatible with different high-level access policies that could be defined at an organization level.

2.2 Notation and cryptographic preliminaries

Notation. We use $x \leftarrow \text{Alg}$ to denote assignment from a possibly randomized algorithm `Alg` and $x \leftarrow_R D$ to denote a random sample from a distribution D . We denote a set of linear secret shares of x (resp. function secret shares of f) as $[x]$ (resp. $[f]$) and $[x]_i$ (resp. $[f]_i$) as the i -th secret share in the set of shares encoding x (resp. f). Specifically, we let $[x] := ([x]_1, \dots, [x]_s)$ when x is secret-shared between s parties. We say an algorithm is *efficient* if it runs in probabilistic polynomial time.

Linear Secret Sharing. A linear secret-sharing (LSS) scheme [[49](#)] is parameterized by integers $2 \leq t \leq s$ and a field \mathbb{F} , and consists of two (possibly randomized) algorithms `LSS.Share`_(\mathbb{F}, t, s) and `LSS.Recover`. `LSS.Share`_(\mathbb{F}, t, s) generates s shares of a secret value in the field \mathbb{F} such that (1) any subset of t or more shares can be combined using the linear function `LSS.Recover` to reveal the encoded value in the field \mathbb{F} , (2) no subset of fewer than $t - 1$ shares provides any information on the secret, and (3) shares can be added together to obtain a new share encoding the sum of the secrets. We will occasionally parameterize `LSS.Share` with a ring \mathcal{R} instead of a field \mathbb{F} .

Remark 2. A consequence of the linearity of `LSS.Recover` is that it can be evaluated “in the exponent” of a group. That is, given $g^{[v]_1}, \dots, g^{[v]_t}$, it is possible to efficiently compute $g^v := g^{\text{LSS.Recover}([v]_1, \dots, [v]_t)}$. For simplicity of notation, we define `ExpRecover`: $\mathbb{G}^t \rightarrow \mathbb{G}$ to be the efficiently computable function which takes as input $(g^{[v]_1}, \dots, g^{[v]_t})$ and outputs g^v with $v := \text{LSS.Recover}([v]_1, \dots, [v]_t)$.

Discrete logarithm problem and assumption. Let $\lambda \in \mathbb{N}$ be a security parameter. For a cyclic group \mathbb{G} of prime order $p = p(\lambda)$ with generator g , the Discrete Logarithm (DL) assumption states that no efficient algorithm \mathcal{A} can find $x \in \mathbb{F}_p$ satisfying $y = g^x$ for a uniformly random $y \in \mathbb{G}$ [[38](#)].

Function Secret Sharing. FSS is a generalization of LSS; rather than secret-sharing a *value*, FSS captures the notion of secret sharing *functions*.

Definition 2 (FSS; adapted from [[9](#), [19](#)]). Let $2 \leq t \leq s$ be integers, \mathbb{G} be an Abelian group, and $\mathcal{F} = \{f_i: \{0, 1\}^n \rightarrow \mathbb{G}\}_{1 \leq i \leq N}$ be a family of N functions. A (t, s) -FSS scheme for \mathcal{F} consists of efficiently computable (possibly randomized) algorithms `(FSS.Gen, FSS.Eval)` with the following syntax:

- `FSS.Gen`($1^\lambda, f$) $\rightarrow (\kappa_1, \dots, \kappa_s)$. Takes as input a security parameter λ and description of a function $f \in \mathcal{F}$. Outputs s evaluation keys $\kappa_1, \dots, \kappa_s$.
- `FSS.Eval`(κ_i, x) $\rightarrow [y]_i$. Takes as input an evaluation key κ_i and an input value $x \in \{0, 1\}^n$. Outputs secret share $[y]_i$.

The above functionality must satisfy the following properties:

- *Correctness.* A (t, s) -FSS scheme is correct if there exists an efficient output decoder `Decode` such that for all $f \in \mathcal{F}$, and all subsets $I \subseteq \{1, \dots, s\}$ where $|I| \geq t$,

$$\Pr \left[\begin{array}{l} (\kappa_1, \dots, \kappa_s) \leftarrow \text{FSS.Gen}(1^\lambda, f) : \\ \text{Decode}(\{[y]_i \leftarrow \text{FSS.Eval}(\kappa_i, x) \mid i \in I\}) = f(x) \end{array} \right] = 1.$$




Fig. 1: Overview of FSS and our access control model. In FSS *without* access control, a dealer distributes shares of a function f_i with the evaluators. ① A setting *with* access control where the evaluators have Λ and ② the dealer has an access key pk_2 for f_2 . ③ the dealer (acting as the prover) with knowledge of the access key for f_2 distributes secret shares of the function f_2 and proof shares π to the evaluators. ④ The evaluators collectively check access rights to f_2 using the proof shares $[\pi]$ and the access control list Λ .

- *Privacy.* For all $I \subset \{1, \dots, s\}$ subset of indices such that $|I| < t$, let D_I be the distribution over $\{\kappa_i \mid i \in I\}$ where κ_i is sampled according to $\text{FSS.Gen}(1^\lambda, f)$. A (t, s) -FSS scheme $(\text{FSS.Gen}, \text{FSS.Eval})$ is private if there exists an efficient simulator \mathcal{S} such that $D_I \approx_c \mathcal{S}(1^\lambda, I)$. That is, the distribution of any subset of fewer than t FSS keys reveals no information on the function f to the subset of computationally bounded evaluators.
- *Efficiency.* A (t, s) -FSS scheme is efficient if (1) each key κ_i is at most $O(\lambda N^\epsilon)$ in size, for any $\epsilon < 1$ (possibly dependent on n) and (2) Decode runs in time $O(\lambda s)$.

Following Boyle et al. [9], we assume Decode is a linear function of the inputs (e.g., a summation over the outputs of Eval). As such, we will write $[f]_i$ to denote the i -th FSS key κ_i and $[f(x)]_i$ to denote the i -th share of the evaluation $\text{Eval}(\kappa_i, x)$. We will also assume that $\text{Decode} := \text{LSS.Recover}$ so as to allow for output reconstruction “in the exponent” of a group (see Remark 2).

3 Private Access Control Lists

In this section, we formalize the notion of private ACLs applied to FSS (Definition 2). A Private ACL (PACL) is instantiated between a prover and a set of s verifiers. The prover holds an access key sk and the function $f_i \in \mathcal{F}$. The verifiers hold secret-shares $[f_i]$ and have the access control list Λ (see Definition 1) for the function family \mathcal{F} . The verifiers determine whether or not CheckAccess outputs yes, without learning f_i . See Figure 1 for an overview.

Efficiency constraints. As highlighted in Section 1, a PACL scheme is efficient if it has a small communication overhead for the prover (relative to sharing f) and at most one message exchanged between verifiers. By requiring that only *one*, short message is exchanged between verifiers, we achieve optimal communication (up to constant factors) and ensure function privacy against malicious verifiers deviating from protocol. (Our definition will also implicitly eliminate all solutions that involve the prover in the verification process.)

3.1 Defining PACLs

A PACL scheme consists of four algorithms: PACL.KeyGen , PACL.Prove , PACL.Audit , and PACL.Verify , parameterized by a function family \mathcal{F} , and integers $2 \leq t \leq s$. PACL.Prove is used by the prover to generate an access control proof for a function f_i . The other algorithms are used by the verifiers to enforce access control. The PACL.Audit and PACL.Verify algorithms, combined, instantiate CheckAccess for the family of functions in the distributed setting. PACL.Audit and PACL.Verify only reveal the output of CheckAccess (yes or no), without revealing any other information to the verifiers. We leave any public parameters as an implicit input to all algorithms.

Definition 3 (PACL: Syntax, Completeness, & Efficiency). Let $\lambda \in \mathbb{N}$ be a security parameter, \mathbb{G} be an Abelian group, and $\mathcal{F} := \{f_i: \{0, 1\}^n \rightarrow \mathbb{G}\}_{1 \leq i \leq N}$ be a family of N functions. Fix integers $2 \leq t \leq s$ and let $(\text{FSS.Gen}, \text{FSS.Eval})$ instantiate a (t, s) -FSS scheme for \mathcal{F} . A (t, s) -PACL scheme consists of efficient algorithms $(\text{PACL.KeyGen}, \text{PACL.Prove}, \text{PACL.Audit}, \text{PACL.Verify})$ with the following syntax:

- $\text{PACL.KeyGen}(1^\lambda, f) \rightarrow (\text{vk}, \text{sk})$. Takes as input a security parameter λ and a function $f \in \mathcal{F}$. Outputs a new pair of verification and access keys (vk, sk) .
- $\text{PACL.Prove}(f, \text{sk}) \rightarrow ([\pi]_1, \dots, [\pi]_s)$. Takes as input a function description $f \in \mathcal{F}$ and access key sk . Outputs a vector of s proof secret shares $([\pi]_1, \dots, [\pi]_s)$.
- $\text{PACL.Audit}(\Lambda, [f]_i, [\pi]_i) \rightarrow \tau_i$. Takes as input access control list $\Lambda \subseteq \{\text{vk}_i\}_{1 \leq i \leq N}$, function secret share $[f]_i$ of f sampled according to Gen , and proof share $[\pi]_i$. Outputs audit token τ_i .
- $\text{PACL.Verify}(\mathcal{T} := \{\tau_i \mid i \in I\}) \rightarrow \text{yes/no}$. Takes as input a set of t or more audit tokens indexed by the set $I \subseteq \{1, \dots, s\}$. Outputs yes or no.

The above functionality must satisfy:

- *Completeness.* Let CheckAccess be as defined in [Definition 1](#). A (t, s) -PACL scheme is complete if for all security parameters λ , for all subsets $I \subseteq \{1, \dots, s\}$ with $|I| \geq t$, for all $\Lambda \subseteq \{\text{vk}_i\}_{1 \leq i \leq N}$ where $\forall i, \text{vk}_i$ is sampled according to PACL.KeyGen , and for all secret shares $([f]_1, \dots, [f]_s)$ of $f \in \mathcal{F}$ sampled according to $\text{FSS.Gen}(1^\lambda, f)$, it holds that:

$$\Pr \left[\begin{array}{l} ([\pi]_1, \dots, [\pi]_s) \leftarrow \text{PACL.Prove}(f, \text{sk}); \\ \{\tau_i \leftarrow \text{PACL.Audit}(\Lambda, [f]_i, [\pi]_i) \mid i \in I\} : \\ \text{PACL.Verify}(\{\tau_i \mid i \in I\}) = \text{CheckAccess}(\Lambda, f, \text{sk}) \end{array} \right] = 1,$$

where the probability is taken over the randomness of PACL.KeyGen and PACL.Prove .

- *Efficiency.* The size of each proof share $[\pi]_i$ is most $O(\lambda N^\epsilon)$ for any $\epsilon < 1$ (possibly dependent on n). The size of each audit token τ_i is $O(\lambda)$.

Remark 3. We will primarily be interested in PACLs where ϵ , as defined in the efficiency property of [Definition 3](#) (PACL), matches the ϵ of [Definition 2](#) (FSS). This translates to a constant overhead in communication over sharing f itself via FSS (i.e., *without* any access control).

Remark 4. The access control list Λ is defined over a subset of all N functions in the family f , since the family \mathcal{F} can be of exponential size. One can think of Λ as being an access control list for a polynomial-sized subset of the family \mathcal{F} or, alternatively, as an access control list for a polynomially-sized family of functions (e.g., when $n \in O(\log \lambda)$ and $|\mathbb{G}| \in O(1)$).

Definition 4 (PACL, Soundness & Privacy). A PACL scheme (as defined in [Definition 3](#)) must satisfy the soundness and privacy properties, which are defined as follows.

- *Soundness.* For all efficient adversaries \mathcal{A} , there exists a negligible function negl such that for all subsets $I \subseteq \{1, \dots, s\}$ where $|I| \geq t$, and subsets $S \subseteq \{1, \dots, N\}$ where $|S| \in \text{poly}(\lambda)$,

$$\Pr[\text{PKSOUNDNESS}_{\text{PACL}, \mathcal{A}, I, S}(\lambda) = \text{yes}] \leq \text{negl}(\lambda),$$

where $\text{PKSOUNDNESS}_{\text{PACL}, \mathcal{A}, I, S}(\lambda)$ is defined in [Figure 2](#).

Game $\text{PKSOUNDNESS}_{\text{PACL}, \mathcal{A}, I, S}(\lambda)$	Oracle $\text{GETKEY}(j)$
1 : for $i \in S$:	1 : $T := T \cup \{j\}$
2 : $(\text{vk}_i, \text{sk}_i) \leftarrow \text{PACL.KeyGen}(1^\lambda, f_i)$	2 : return sk_j
3 : $\Lambda := \{\text{vk}_i\}_{i \in S}$, $T = \{\}$	
4 : $([f_\gamma], [\pi_\gamma]) \leftarrow \mathcal{A}^{\text{GETKEY}}(1^\lambda, \Lambda, S)$	
5 : $f_\gamma \leftarrow \text{LSS.Recover}([f_\gamma])$	
6 : for $i \in I$:	
7 : $\tau_i \leftarrow \text{PACL.Audit}(\Lambda, [f_\gamma]_i, [\pi_\gamma]_i)$	
8 : return $\text{PACL.Verify}(\{\tau_i \mid i \in I\}) = \text{yes}$ and $f_\gamma \in \mathcal{F}$ and $\gamma \in S \setminus T$	

Fig. 2: PACL soundness game.

In words, no efficient algorithm \mathcal{A} can forge a proof π that verifies with non-negligible probability without knowledge of an access key for f_γ .

- *Privacy.* For all subsets $I \subset \{1, \dots, s\}$ such that $|I| < t$, define D_I to be the distribution over the set $\{[\pi]_i, \tau_1, \dots, \tau_s \mid i \in I\}$, where each $[\pi]_i$ is sampled according to $\text{PACL.Prove}(f, \text{sk})$, and τ_j is computed according to $\text{PACL.Audit}(\Lambda, [f]_j, [\pi]_j)$ for all $j \in \{1, \dots, s\}$. A (t, s) -PACL is private if there exists an efficient simulator \mathcal{S} such that: $D_I \approx_c \mathcal{S}(1^\lambda, I)$. That is, the distribution of proof shares and audit shares reveal nothing about f or the access key sk to a subset of at most $t - 1$ computationally bounded verifiers.

3.2 Symmetric-key PACL

In [Definition 4](#), we define soundness where the verification keys are assumed to be public and accessible to the adversary. This models a setting where an adversary may obtain a snapshot of one verifier's state (consisting of access control list Λ). For some applications [[27](#), [53](#)], it is useful to relax the definition of soundness of [Definition 4](#) and let the access control list Λ consist of the *secret* keys rather than *public* keys (see prior approaches in [Section 1.2](#)). In this regime, the soundness definition must exclude Λ from the inputs to the adversary \mathcal{A} . In practical terms, symmetric-key PACLs do not protect against snapshot attacks where an adversary might momentarily compromise a verifier and learn Λ (allowing it to subvert the access control at a later point in time) [[3](#), [29](#)].

Definition 5 (PACL: Symmetric-key soundness). For all efficient adversaries \mathcal{A} , there exists a negligible function negl such that for all subsets $I \subseteq \{1, \dots, s\}$ where $|I| \geq t$, and subsets $S \subseteq \{1, \dots, N\}$ where $|S| \in \text{poly}(\lambda)$,

$$\Pr[\text{SKSOUNDNESS}_{\text{PACL}, \mathcal{A}, I, S}(\lambda) = \text{yes}] \leq \text{negl}(\lambda),$$

where $\text{SKSOUNDNESS}_{\text{PACL}, \mathcal{A}, I, S}(\lambda)$ is defined in [Figure 3](#).

Game $\text{SKSOUNDNESS}_{\text{PACL}, \mathcal{A}, I, S}(\lambda)$	Oracle $\text{GETKEY}(j)$
1 : for $i \in S$:	1 : $T := T \cup \{j\}$
2 : $(-, \text{sk}_i) \leftarrow \text{PACL.KeyGen}(1^\lambda, f_i)$	2 : return sk_j
3 : $\Lambda := \{\text{sk}_i\}_{i \in S}$, $T = \{\}$	
4 : $([f_\gamma], [\pi_\gamma]) \leftarrow \mathcal{A}^{\text{GETKEY}}(1^\lambda, S)$	
5 : $f_\gamma \leftarrow \text{LSS.Recover}([f_\gamma])$	
6 : for $i \in I$:	
7 : $\tau_i \leftarrow \text{PACL.Audit}(\Lambda, [f_\gamma]_i, [\pi_\gamma]_i)$	
8 : return $\text{PACL.Verify}(\{\tau_i \mid i \in I\}) = \text{yes}$ and $f_\gamma \in \mathcal{F}$ and $\gamma \in S \setminus T$	

Fig. 3: Symmetric-key PACL access soundness game.

In words, no efficient algorithm \mathcal{A} , without knowledge of the access key, can forge a proof π that verifies with non-negligible probability. Unlike [Definition 4](#), here Λ is private to the verifiers and is not given to \mathcal{A} .

3.3 Security against malicious verifiers

Privacy. [Definition 4](#) guarantees privacy against any subset of fewer than t malicious verifiers. Only one message (the audit token) is exchanged by the verifiers to check the proof. Thus, the audit token of each honest verifier is guaranteed to be computed *independently* of audit tokens output by malicious verifiers (i.e., malicious verifiers have no influence over the output of the honest verifiers).

Soundness. In contrast, note that the *soundness* property of PACLs is only guaranteed if all verifiers follow the protocol. This is a natural consequence of the fact that FSS itself only guarantees integrity of the output if all evaluators adhere to the protocol (any malicious evaluator can incorrectly compute $[f_i(x)]$ to corrupt the final output). As such, access control is only well-defined when verifiers have a vetted interest in ensuring correctness of the function evaluation.

3.4 Key distribution

Key distribution is a challenging problem in many real-world systems. Systems using FSS and PACLs must handle distributing the verification and access keys to the users (dealers) and function evaluators (verifiers). This can be done through a variety of techniques. For example, a trusted setup can take place to generate and distribute the keys. Alternatively, anonymous communication channels can be used to register with the evaluators by providing a verification key for a particular function. Ultimately, the key distribution mechanism itself is orthogonal to the goals of PACLs as it depends significantly on the deployment setting (e.g., see Express [27] and Spectrum [41]).

4 PACLs for Distributed Point Functions

In this section, we describe our PACL constructions for the class of distributed point functions (DPFs). DPFs are the main primitive behind more complex FSS classes constructible from minimal assumptions [9, 11]. By focusing on DPFs, our PACL constructions become applicable to larger classes of FSS, which we explain further in Section 6.

Distributed Point Functions (DPFs). A *point function* P_i is a function that evaluates to 1 on input i and evaluates to 0 on all other inputs $j \neq i$. A *distributed* point function is an instance of FSS for the family of point functions. More generally, DPFs can be defined to output *any* value m at index i [30]. We focus on $m = 1$ for simplicity and note that our constructions generalize to any m .

4.1 DPF-PACL constructions

Parameters. Let \mathbb{G} be a group of prime order $p = p(\lambda) \in \text{poly}(\lambda)$ with generator g in which the discrete logarithm problem is assumed to be computationally intractable. We assume that the family of (distributed) point functions has range \mathbb{F}_p . In the special case of two-party DPF constructions, which output in a binary field \mathbb{F}_2 [9, 11], our constructions can be adapted by simply “interpreting” the binary secret share as an element of \mathbb{F}_p , resulting in subtractive secret shares of either -1 or 1 at the special index, which the prover knows.

Overview. In Section 4.1.1, we construct a DPF-PACL for the match predicate of Section 2.1. Our construction can be seen as a generalization of the technique used by Newman et al. [41]. In Section 4.1.2, we extend this technique to a DPF-PACL for the inclusion predicate of Section 2.1. Before delving into the technical details, we make the following two remarks about our constructions:

Remark 5. Our constructions are defined with $\Lambda := (\text{vk}_1, \dots, \text{vk}_N)$, which can potentially be exponential, as noted in Remark 4. As such, we can think of our constructions as either being suitable for a polynomial-sized function family \mathcal{F} or as being defined for an *implicit* list of N access keys. In Section 4.2, we explain how to enforce access control for $\Lambda \subseteq \{\text{vk}_i\}_{1 \leq i \leq N}$, where only a subset of the N (implicit) verification keys are used by the verifiers.

Remark 6. For DPFs, our constructions only enforce access control for the point function’s “special index,” letting the output message m be an arbitrary value (for simplicity, our constructions assume $m = 1$). Enforcing that the output be from a specific set of allowed elements of the group \mathbb{G} can be done using zero-knowledge proofs on secret-shared data [7].

4.1.1 DPF-PACL for match predicate In Algorithm 1, we present the construction for a DPF-PACL with CheckAccess instantiated for the match predicate described in Section 2.1. Loosely speaking, the idea behind the construction is to use the DPF to locally select shares of the i -th verification key in Λ . Two facts make this possible: (1) all the verifiers have $\Lambda = (g^{\alpha_1}, \dots, g^{\alpha_N})$ and (2) the FSS key κ_i encoding a DPF can be used to privately retrieve the k -th entry in any vector by first evaluating the DPF $[y_j]_i \leftarrow \text{DPF.Eval}(\kappa_k, j)$ for all $j \in \{1, \dots, N\}$, and then computing the inner-product “in the exponent” as $g^{[\alpha_k]_i} := g^{\langle (\alpha_1, \dots, \alpha_N), ([y_1]_i, \dots, [y_N]_i) \rangle}$, giving a multiplicative share of the k -th verification key $\text{vk}_k := g^{\alpha_k}$.

To verify knowledge of α_k , the prover distributes to the verifiers *additive* secret shares of $\pi := -\alpha_k = \text{sk}_k$ (described in PACL.Prove). Each verifier computes $\tau_i := (g^{[\alpha_k]_i})g^{[\pi]_i}$ using PACL.Audit and reveals τ_i to all other verifiers. All verifiers proceed to check that $\tau = g^0$ (described in PACL.Verify).

Theorem 1 (Informal). *There exists a DPF-PACL for the FSS family $\text{DPF}: \{0, 1\}^n \rightarrow \mathbb{F}_p$ with proof size $O(\lambda)$ and audit size $O(\lambda)$, where CheckAccess is instantiated as the match predicate of Section 2.1. The verifiers run in time $O(\lambda \cdot |\Lambda|)$.*

Algorithm 1: DPF-PACL FOR MATCH PREDICATES

Public parameters: integers $2 \leq t \leq s$, function family $\mathcal{F} = \{f_i : \{0, 1\}^n \rightarrow \mathbb{F}_p\}_{1 \leq i \leq N}$, and cyclic group \mathbb{G} of order p with generator g .

– PACL.KeyGen($1^\lambda, f_i$):

- 1: $\alpha_i \leftarrow_R \mathbb{F}_p$
- 2: $\text{vk}_i := g^{\alpha_i}, \text{sk}_i := -\alpha_i$
- 3: **return** $(\text{vk}_i, \text{sk}_i)$

– PACL.Prove(f, sk):

- 1: $([\pi]_1, \dots, [\pi]_s) \leftarrow \text{LSS.Share}_{(\mathbb{F}_p, t, s)}(\text{sk})$
- 2: **return** $([\pi]_1, \dots, [\pi]_s)$

– PACL.Audit($\Lambda, [f]_i, [\pi]_i$):

- 1: **parse** $\Lambda = (g^{\alpha_1}, \dots, g^{\alpha_N})$ and $[f]_i = \kappa_i$
- 2: **for** $j \in \{1, \dots, N\}$:
 - 2.1: $[y_j]_i \leftarrow \text{DPF.Eval}(\kappa_i, j)$
- 3: $A := \prod_{j=1}^N (g^{\alpha_j})^{[y_j]_i}$ // Inner product in \mathbb{G} .
- 4: $\tau_i := A \cdot g^{[\pi]_i}$
- 5: **return** τ_i

– PACL.Verify(\mathcal{T}):

- 1: **parse** $\mathcal{T} := \{\tau_1, \dots, \tau_t\}$.
- 2: $C \leftarrow \text{ExpRecover}(\tau_1, \dots, \tau_t)$ // See Remark 2 for definition.
- 3: **return** $C \stackrel{?}{=} 1_{\mathbb{G}}$

4.1.2 DPF-PACL for inclusion predicates We now describe how to instantiate a DPF-PACL with an inclusion predicate (Section 2.1). Each function is associated with ℓ access keys. As such, Λ consists of N verification keys, where each verification key consists of ℓ subkeys. For each $\text{vk}_i \in \Lambda$, any of the ℓ subkeys can be used to prove access rights for the function f_i .

Theorem 2 (Informal). *Let s_ℓ be the size of a DPF key for a point function with a domain of size ℓ . There exists a DPF-PACL for the FSS family $\text{DPF}: \{0, 1\}^n \rightarrow \mathbb{F}_p$ with proof size $O(\lambda + s_\ell)$*

Algorithm 2: DPF-PACL FOR INCLUSION PREDICATES

Public parameters: integers $2 \leq t \leq s$, function family $\mathcal{F} := \{f_i : \{0,1\}^n \rightarrow \mathbb{F}_p\}_{1 \leq i \leq N}$, and cyclic group \mathbb{G} of order p with generator g .

Let $(\text{Prove}', \text{Audit}', \text{Verify}')$ be as in [Algorithm 1](#).

// Parties can locally pre-compute correction terms ahead of time to avoid recomputing them in **Audit**.

// Only required for efficiency, not correctness.

– **Precomputation:** // Compute correction terms.

1: **parse** $\Lambda = (\text{vk}_1, \dots, \text{vk}_N)$, $\text{vk}_j = (\text{vk}_{j,1}, \dots, \text{vk}_{j,\ell})$

2: **for** $j \in \{1, \dots, N\}$, $k \in \{1, \dots, \ell\}$:

2.1: **parse** $\text{vk}_{j,k} = (g^{\alpha_{j,k}}, g^{\beta_{j,k}})$

2.2: $g^{w_{(j-1)\ell+k}} := \prod_{l=1, l \neq k}^{\ell} g^{\alpha_{j,l}}$

– **PACL.KeyGen** $(1^\lambda, f_i)$:

1: $(\alpha_{i,1}, \dots, \alpha_{i,\ell}) \leftarrow_R \mathbb{F}_p^\ell$, $(\beta_{i,1}, \dots, \beta_{i,\ell}) \leftarrow_R \mathbb{F}_p^\ell$

2: **for** $j \in \{1, \dots, \ell\}$:

2.1: $\text{vk}_{i,j} := (g^{\alpha_{i,j}}, g^{\beta_{i,j}})$, $\text{sk}_{i,j} := (-\alpha_{i,j}, -\beta_{i,j}, j)$

3: **return** $(\text{vk}_{i,1}, \dots, \text{vk}_{i,\ell}), (\text{sk}_{i,1}, \dots, \text{sk}_{i,\ell})$

– **PACL.Prove** (f, sk) :

1: **parse** $f = P_i$ and $\text{sk} = (\alpha, \beta, \gamma)$

2: $([\alpha]_1, \dots, [\alpha]_s) \leftarrow \text{LSS.Share}_{(\mathbb{F}_p, t, s)}(\alpha)$

3: $\omega := (i-1)\ell + \gamma$ // γ -th key in row i .

4: $(\kappa'_1, \dots, \kappa'_s) \leftarrow \text{DPF.Gen}(1^\lambda, P_\omega)$

5: $([\beta]_1, \dots, [\beta]_s) \leftarrow \text{Prove}'(P_\omega, \beta)$

6: $[\pi]_j := ([\alpha]_j, [\beta]_j, \kappa'_j)$ **for** $j \in \{1, \dots, s\}$

7: **return** $([\pi]_1, \dots, [\pi]_s)$

– **PACL.Audit** $(\Lambda, [f]_i, [\pi]_i)$:

1: **parse** $\Lambda = (\text{vk}_1, \text{vk}_2, \dots, \text{vk}_N)$,
 $\text{vk}_j = (\text{vk}_{j,1}, \dots, \text{vk}_{j,\ell})$, $\text{vk}_{j,k} = (g^{\alpha_{j,k}}, g^{\beta_{j,k}})$,
 $[f]_i = \kappa_i$, and $[\pi]_i = ([\alpha]_i, [\beta]_i, \kappa'_i)$

2: **for** $j \in \{1, \dots, N\}$:

2.1: $[y_j]_i \leftarrow \text{DPF.Eval}(\kappa_i, j)$

3: $(A_1, \dots, A_\ell) := \prod_{j=1}^N (g^{\alpha_{j,1}}, \dots, g^{\alpha_{j,\ell}})^{[y_j]_i}$

4: $A' := (g^{w_1}, \dots, g^{w_{N\ell}})$, $\tau_i^{(0)} \leftarrow \text{Audit}'(A', \kappa'_i, [\beta]_i)$

5: **for** $j \in \{1, \dots, N\ell\}$:

5.1: $[c_j]_i \leftarrow \text{DPF.Eval}(\kappa'_i, j)$

6: $W := \prod_{j=1}^{N\ell} (g^{w_j})^{[c_j]_i}$ // Use pre-computed correction term.

7: $A := \left(\prod_{j=1}^{\ell} A_j \right) \cdot (W)^{-1}$, $\tau_i^{(1)} := A \cdot g^{[\alpha]_i}$

8: **return** $\tau_i := (\tau_i^{(0)}, \tau_i^{(1)})$

– **PACL.Verify** (\mathcal{T}) :

1: **parse** $\mathcal{T} = \{\tau_1, \dots, \tau_t\}$ and $\tau_i = (\tau_i^{(0)}, \tau_i^{(1)})$

2: **if** $\text{PACL.Verify}'(\{\tau_1^{(0)}, \dots, \tau_t^{(0)}\}) = \text{no}$ **then return no**

3: $C \leftarrow \text{ExpRecover}(\tau_1^{(1)}, \dots, \tau_t^{(1)})$ // See [Remark 2](#).

4: **return** $C \stackrel{?}{=} 1_{\mathbb{G}}$

and audit size $O(\lambda)$, where `CheckAccess` is instantiated as the inclusion predicate of [Section 2.1](#). The verifiers run in time $O(\lambda \cdot \ell |A|)$.

[Algorithm 2](#) presents our construction of DPF-PACL for inclusion predicates. At a high level, the verifiers “select” the i -th row in the matrix A using f_i (similarly to [Algorithm 1](#)) by computing the inner product between the evaluation of f_i on its domain and the access control matrix. However, the challenge is then to have the verifiers obliviously select the j -th column in the selected row. Because the resulting row is secret-shared, the verifiers cannot recursively select the column using another DPF, as it would require the vector to be known by all verifiers. Revealing the column does not work either, as it could violate the privacy requirement. For example, revealing the column means that the verifiers would learn that the prover is using an access key contained in the column, which may only include a subset of functions (e.g., following revocations), thus leaking information about the secret-shared function. One option is to use zero-knowledge proofs over secret shares [\[7, 18\]](#). However, we opt for a simpler and more efficient approach. First, the verifiers generate ℓ sums of verification keys for each row in the access control list (resulting in a total of $N\ell$ terms). One of these sums can then be used as a “correction term” by the prover to select only the j -th column in the row. To see how, consider a row $R_i = (g^{\alpha_{i,1}}, \dots, g^{\alpha_{i,\ell}})$. Each of the ℓ correction terms $g^{w_{i,1}}, \dots, g^{w_{i,\ell}}$ associated with the i -th row is defined as: $g^{w_{i,j}} := \prod_{k=1, k \neq j}^{\ell} g^{\alpha_{i,k}}$. The j -th entry in the multiplicatively secret-shared row $[R_i] := (g^{[\alpha_{i,1}]}, \dots, g^{[\alpha_{i,\ell}]})$ can be recovered as: $g^{[\alpha_{i,j}]} := \left(\prod_{k=1}^{\ell} g^{[\alpha_{i,k}]} \right) / g^{[w_{i,j}]}$.

The prover can easily select the correction term $g^{w_{ij}}$ by generating a separate DPF for the point function P_ω and sending it to the verifiers. The verifiers use the DPF to select the ω -th term in the list of correction terms. To see how, notice that we can “flatten” the correction terms into a list of size $N\ell$ elements and take the inner product to get a secret share of the ω -th correction term, as in [Section 4.1.1](#).

Unfortunately, while the prover can now select the correct key in the list, this idea also creates an avenue for an attack. A malicious prover can subvert access control entirely by selecting multiple correction terms to “annihilate” a row. The prover can send a distributed *multi-point* function (a point function that evaluates to 1 on multiple inputs) to select all ℓ correction terms of a row in A . If, additionally, the prover is trying to authenticate the DPF for point function $P_{i,m}$ that evaluates to $m = (\ell - 1)$ on input i (recall [Section 4](#)), then the verifiers would compute, $\prod_{k=1}^{\ell} g^{(\ell-1) \cdot [\alpha_{i,k}]}$ from the evaluation of $P_{i,m}$ over the verification keys and $\prod_{k=1}^{\ell} g^{[w_{i,k}]}$ over the correction terms. Then, it holds that:

$$\prod_{k=1}^{\ell} g^{(\ell-1)[\alpha_{i,k}]} \Big/ \prod_{k=1}^{\ell} g^{[w_{i,k}]} = \prod_{k=1}^{\ell} g^{(\ell-1)[\alpha_{i,k}]} \Big/ \prod_{k=1}^{\ell} g^{(\ell-1)[\alpha_{i,k}]} = g^{[0]}.$$

Hence, the verifiers would recover shares of g^0 for which the discrete logarithm is simply zero. To prevent this attack, we leverage the following insight: each correction term is associated with a *unique* access key in A . As a consequence, we can instantiate a separate DPF-PACL to enforce access control over the vector of correction terms. Specifically, we generate an access key $\beta_{i,j}$ for $w_{i,j}$ and apply [Algorithm 1](#) to enforce the access control over the set of correction terms. The access key is now a *tuple* $(\alpha_{i,j}, \beta_{i,j})$, and verification consists of checking access control for *two* DPFs: the implicit DPF (P_i) and the DPF selecting the correction term.

4.2 Optimizations and extensions

We briefly highlight some optimizations and extensions that can be applied to [Algorithms 1](#) and [2](#).

Reducing communication and computation. We present [Algorithm 2](#) with a separate DPF for the selection of the correction term. This would result in an additive overhead of $O(\lambda(N\ell)^\epsilon)$ in communication (ϵ as defined in [Definition 2](#)). However, we observe that we can use the “FSS tensoring” transformation described by Boyle et al. [\[11\]](#) to capitalize on the common “backbone” of the underlying DPF being authenticated and reduce the communication overhead from $O(\lambda(N\ell)^\epsilon)$ down to $O(\lambda^\epsilon)$. Specifically, the prover can use $P_i(i)$ (the non-zero output of the DPF) as a mask for κ' (the key for the DPF selecting the correction term). In this way, κ' only needs a range of ℓ (rather than $N\ell$) leading to the reduced proof size. In the interest of space, we point the reader to Boyle et al. [\[11\]](#) for a full description of the FSS tensoring technique.

Sparse domain auditing. The constructions presented in [Algorithms 1](#) and [2](#) require $O(\lambda N)$ work per verifier to compute `PACL.Audit`. However, in practice, the evaluators (i.e., verifiers) might only

evaluate f on a *sparse* subset of the domain rather than the entire domain of the function. In this case, the verifiers only need to compute `PACL.Audit` on the matching subset of the domain on which they evaluate f . Taking this to its extreme, if the verifiers only evaluate f on a constant number of inputs, then this optimization leads to `PACL.Audit` running in $O(1)$ time. Furthermore, the ACL Λ need only contain $O(1)$ keys. More generally, for a subset $S \subseteq \{1, \dots, N\}$ of the DPF domain, we need $|S|$ keys in Λ and evaluate `PACL.Audit` on the $|S|$ inputs, making the verifier work $O(|S|)$. This optimization addresses [Remark 5](#), and makes our constructions satisfy [Definitions 4](#) and [5](#) where $\Lambda \subseteq \{vk_i\}_{1 \leq i \leq N}$. Given this optimization, the overhead of PACLs is essentially constant relative to the evaluation of the function itself. A downside, however, is that the prover may need to know S (or a subset thereof) when computing `PACL.Prove`. More specifically, we can view this optimization as enforcing access control on a smaller function f' that *coincides* with f on all inputs in the subset S . That is, $f'(x) = f(x)$ for all $x \in S$ but it may be the case that $f'(x') \neq f(x')$ for all $x' \notin S$, which naturally requires the prover to know f' .

Public-key vs. symmetric-key DPF-PACL. When \mathbb{G} is chosen to be a group in which the discrete logarithm problem is assumed to be computationally intractable [\[6\]](#) (e.g., when $\mathbb{G} = \mathbb{Z}_p^*$ is the multiplicative group of order p), then the construction satisfies the soundness property of PACLs as defined in [Definition 4](#), because the entire access policy can be made public. However, when \mathbb{G} is instead defined as an additive group \mathbb{Z}_p , then we get a symmetric-key PACL satisfying the relaxed soundness property in [Section 3.2](#) (i.e., no security against adversaries obtaining a snapshot of a server state, which contains the verification keys). The advantage of using the additive group $\mathbb{G} = \mathbb{Z}_p$ is that we can get far better concrete performance (see [Section 8](#)). However, this comes at the cost of requiring the evaluation servers to keep the access policy secret. In [Section 5](#), we describe how to get concrete performance of symmetric-key PACLs while still making it possible to have a public access policy consisting of verification keys in a group \mathbb{G} where the discrete logarithm problem is hard.

4.3 Aggregating PACLs

A nice property of our DPF-PACL constructions ([Sections 4.1.1](#) and [4.1.2](#)) is the ability to *aggregate* proofs across different DPFs and access control lists. Concretely, our constructions satisfy the following two aggregation properties. At a high level, for any integer q that is polynomial in the security parameter λ and family of point functions \mathcal{F} :

1. Let Λ be an ACL for the family \mathcal{F} and let $f_1, \dots, f_q \in \mathcal{F}$ have associated access keys $\alpha_1, \dots, \alpha_q \in \Lambda$, then $\alpha' := \sum_{i=1}^q sk_i$ is an access key for $f'(x) := \sum_{i=1}^q f_i(x)$. This aggregation property allows the verifiers to simultaneously enforce access control on q distinct functions in the family for the computational and bandwidth overhead of verifying a single function in the family.
2. Our constructions permit aggregating proofs from multiple *separate* ACLs to simultaneously enforce access control on a vector of functions $(f_1, \dots, f_q) \in \mathcal{F}^q$. For a vector of ACLs $(\Lambda_1, \dots, \Lambda_q)$, the ACL $\Lambda' := \odot_{i=1}^q \Lambda_i$ (where \odot denotes the group operation applied component-wise) is an ACL for the vector (f_1, \dots, f_q) such that `CheckAccess` $(\Lambda_i, f_i, sk_i) = \text{yes}$ for all i . This aggregation property allows for batched verification of q functions, each associated with a *separate* ACL: the verifiers first compute $g^{[\alpha^{(1)}]}, \dots, g^{[\alpha^{(q)}]}$ individually for each function using the corresponding ACL. Then $g^{[\alpha]} := \odot_{i=1}^q g^{[\alpha^{(i)}]}$ can be verified using Λ' . While the computational overhead of this aggregation property remains proportional to verifying each function individually, it permits more compact proofs and audits.

4.4 Security analysis

In this section, we prove security of [Algorithms 1](#) and [2](#) with respect to [Definitions 3](#) and [4](#). We first prove a useful lemma which says that any adversary that wins the $\text{PKSOUNDNESS}_{\text{PACL}, \mathcal{A}, I, S}(\lambda)$ game in our DPF-PACL constructions with some function (not necessarily a DPF), must also implicitly output a valid DPF and access key.

Lemma 1. *If there exists an efficient adversary \mathcal{A} that wins the $\text{PKSOUNDNESS}_{\text{PACL}, \mathcal{A}, I, S}(\lambda)$ game for our DPF-PACL constructions ([Algorithms 1](#) and [2](#)) with non-negligible probability $\delta(\lambda)$ for some function \hat{f}_γ and proof $\hat{\pi}$ where \hat{f}_γ is not sampled from \mathcal{F}_{DPF} , then there exists an efficient adversary \mathcal{A}' that wins the $\text{PKSOUNDNESS}_{\text{PACL}, \mathcal{A}, I, S}(\lambda)$ game with probability $\delta(\lambda)$, where \mathcal{A}' outputs f_γ and π such that $f_\gamma \in \mathcal{F}_{\text{DPF}}$.*

Proof. The proof hinges on the aggregation property of our construction (Section 4.3). Consider an efficient \mathcal{A} that outputs \hat{f}_γ and $\hat{\pi}$ where \hat{f} is *not* a point function (and also not the trivial identity function $f(x) = 0$ for all x). Then, it holds that $\hat{f}_\gamma = \sum_{j \in S} a_j f_j$, where $S \subseteq \{0, 1, \dots, N\}$, each f_j is a point function, and a_j are arbitrary scalars in $\mathbb{F}_p \setminus \{0\}$. Construct an adversary \mathcal{A}' that breaks the $\text{PKSOUNDNESS}_{\text{PACL}, \mathcal{A}, I, S}(\lambda)$ game with $f_\gamma = P_\gamma$ (a point function) as follows. First, run \mathcal{A} to get function \hat{f}_γ . Then compute $f_\gamma := \hat{f}_\gamma - \sum_{j \in S, j \neq \gamma} a_j f_j$ and $\pi := \hat{\pi} - \sum_{j \in S, j \neq \gamma} a_j \alpha_j$ (recall that \mathcal{A}' is allowed to query for all α_j provided $j \neq \gamma$). Finally, output f_γ and π . It must hold that $\gamma \in S$ (if this were not the case then \mathcal{A} does not succeed as it queried all the necessary access keys sk_j for $j \in S$). By the aggregation properties of our construction (described in Section 4.3), it follows that f_γ is a point function and π is a valid access proof for f_γ . Thus, \mathcal{A}' succeeds with the same probability. \square

Theorem 3. *Let λ be a security parameter, $p = p(\lambda)$ be a prime, \mathbb{G} be a group of order p in which the discrete logarithm problem is assumed to be computationally intractable. Algorithm 1 (DPF-PACL for match predicates) satisfies the completeness, efficiency, soundness, and privacy properties of Definitions 3 and 4 with CheckAccess as defined in Section 2.1 (match predicate).*

Proof. We prove each property in turn.

Completeness. Let i be the special index of the encoded point function. Consider the exponent of the recovered audit: $\log_g(C) = \log_g(A \cdot g^{-\alpha}) = (\sum_{j=1}^N \alpha_j y_j) - \alpha$. We have: $(\sum_{j=1}^N \alpha_j y_j) = \alpha_i$ if $i \neq 0$ and 0 otherwise. As such, $\log_g(C) = \alpha_i - \alpha$. By construction, $\alpha := \alpha_i$, so it follows that $\log_g(C) = 0$. Therefore, $C = g^0 = 1_{\mathbb{G}}$, as required.

Soundness (proof of knowledge). Assume, towards contradiction, that there exists an efficient adversary \mathcal{A} and non-negligible function δ such that for all $I \subseteq \{1, \dots, s\}$ where $|I| \geq t$:

$$\Pr[\text{PKSOUNDNESS}_{\text{PACL}, \mathcal{A}, I, S}(\lambda) = \text{yes}] \geq \delta(\lambda).$$

By Lemma 1, we can assume that f_γ (output by \mathcal{A} in Figure 2) is a point function with special index γ . Construct an efficient algorithm \mathcal{B} that solves the discrete logarithm problem in \mathbb{G} as follows. On input $y := g^x$, sample random $\gamma' \leftarrow_R \{1, \dots, N\}$ and $(\alpha_1, \dots, \alpha_N) \leftarrow_R \mathbb{F}_p^N$. Set $A := (g^{\alpha_1}, \dots, g^{\alpha_N})$ but replace $g^{\alpha_{\gamma'}}$ with y . Let $T := \{\}$. Run $\mathcal{A}^{\text{GETKEY}}(1^\lambda, A)$. Respond to each $\text{GETKEY}(j)$ query with α_j (unless $j = \gamma'$, in which case abort) adding j to T . Obtain f_γ and π from \mathcal{A} . If $\gamma \neq \gamma'$ output fail. Else, output $-\pi$. The list A constructed by \mathcal{B} matches the distribution of PACL.KeyGen because $y := g^x$ is a random element of \mathbb{G} . If \mathcal{A} succeeds, then PACL.Verify outputs yes, which means that $C = 1_{\mathbb{G}}$ and so it holds that $x = -\pi$. The probability that $\gamma = \gamma'$ is $\frac{1}{N}$ and so \mathcal{B} succeeds with probability at least $\frac{1}{N} \delta(\lambda)$, which remains non-negligible. As such, \mathcal{B} successfully recovers the discrete logarithm in \mathbb{G} , contradicting the assumption that the discrete logarithm is computationally intractable in \mathbb{G} . We note that \mathcal{B} extracts the access key which makes our construction a proof of knowledge.

Privacy. We construct an efficient simulator \mathcal{S} for the view of any subset of $t-1$ (possibly malicious) verifiers. On input $(1^\lambda, I)$, \mathcal{S} proceeds as follows:

- 1: $([0]_1, \dots, [0]_s) \leftarrow \text{LSS.Share}_{(\mathbb{F}_p, t, s)}(0)$.
- 2: $([\pi]_1, \dots, [\pi]_s) \leftarrow \text{LSS.Share}_{(\mathbb{F}_p, t, s)}(0)$.
- 3: $\tau_k := g^{[0]_k}$ for all $k \in \{1, \dots, s\}$.
- 4: Output $\{[\pi]_i, \tau_1, \dots, \tau_s \mid i \in I\}$.

The distribution output by \mathcal{S} matches the distribution of any subset $I \subset \{1, \dots, s\}$, where $|I| < t$, because in the real view, (1) the proof shares $[\pi]$ are output by LSS.Share which guarantees that any subset of fewer than t shares is information-theoretically hiding and (2) the audit tokens are (computationally-hiding) multiplicative secret shares of $g^0 = 1_{\mathbb{G}}$ (in contrast, \mathcal{S} outputs *information-theoretically* hiding multiplicative shares). The audit tokens in the real view are not information-theoretically hiding because they are computed using the output of the DPF, which consists of computationally-hiding secret shares. The output of \mathcal{S} thus only differs on (2). However, if there is an efficient distinguisher for (2), then the FSS scheme is not private, a contradiction.

Efficiency. Each proof share $[\pi]_i$ is an element of \mathbb{G} and thus is of size $O(\lambda)$. Each audit token is also of size $O(\lambda)$. \square

Theorem 4. *Let p be a prime chosen with respect to a security parameter λ and let \mathbb{G} be any group of order p in which the discrete logarithm problem is assumed to be computationally intractable. Algorithm 2 (DPF-PACL for Inclusion Predicates) satisfies the completeness, efficiency, soundness, and privacy properties of Definitions 3 and 4 for the inclusion predicate of Section 2.1.*

Proof. The proof follows a similar structure to the proof of Theorem 3 but involves more tedious calculations. We defer the proof to Appendix B.1. \square

5 Faster PACLs for DPFs from Verifiable DPFs

In this section, we introduce a concretely more efficient construction of DPF-PACL for the class of *verifiable* DPFs (VDPFs) [23] (also known as *extractable* DPFs [8]). A VDPF allows the evaluators to efficiently check if the DPF is well-formed (see Appendix A), which we will capitalize on to construct more efficient PACLs.

The primary source of inefficiency in Algorithms 1 and 2 is due to the group operations required in computing the PACL audit (when \mathbb{G} is DL-hard group, as required for public-key soundness). If, instead, the verifiers could “select” the public key over a field (e.g., \mathbb{F}_p) rather than in \mathbb{G} , then computing the audit token would be bottlenecked by operations over the field instead of (possibly expensive) group operations in \mathbb{G} . We note that this only applies to PACLs with public verification keys (the soundness definition of Definition 4), where \mathbb{G} a group in which the discrete logarithm is hard. In contrast, in the case of that the symmetric-key soundness definition (Definition 5) is used, then the DPF-PACL construction can be instantiated with $\mathbb{G} = \mathbb{Z}_p$, which does not require concretely expensive exponentiations (see Section 4.2).

There are two technical challenges with this approach. First, if the audit is not computed in \mathbb{G} , the verifiers end up with additive shares $[g^{\alpha_i}]$ (rather than multiplicative shares $g^{[\alpha_i]}$) which does not lend itself to the efficient verification procedure of Algorithms 1 and 2. To overcome this problem, we introduce a building block we call a *Schnorr Proof over Secret Shares* (SPoSS; Section 5.1), which allows a prover to efficiently prove to a set of verifiers that it knows the discrete logarithm of an additively secret-shared group element.

The second challenge is that, in the proof of security, the knowledge extractor (see Section 4.4) would *not* have the guarantee that the resulting additive secret shares encode a verification key from Λ (it could be any linear combination of group elements). This rather subtle problem is a barrier to proving soundness when taking this approach with (non-verifiable) DPFs. To overcome this, we restrict our focus to VDPFs, which ensures that the verifiers always obtain a valid group element from Λ . We then prove security similarly to the proof of Theorem 3.

5.1 Schnorr Proof over Secret Shares (SPoSS)

SPoSS is a non-interactive proof system instantiated in the random oracle model between a prover and a set of two or more verifiers. The verifiers hold additive secret shares of a group element $y := g^x$. The prover provides a zero-knowledge proof-of-knowledge of x (i.e., the discrete logarithm of y base g). SPoSS is a concrete instantiation of a general zero-knowledge proof system over secret shares [7, 18] and can be thought of as a secret-shared analog of a Schnorr proof [46]. We define the formal requirements of SPoSS in Definition 6 and prove security of our construction in Appendix B.2. The proof size of our SPoSS construction is significantly smaller compared to generic approaches based on zero-knowledge proofs (see Section 8).

Definition 6 (SPoSS). *Let $\lambda \in \mathbb{N}$ be a security parameter and let \mathbb{G} be a cyclic group of order $q = q(\lambda)$ with generator g . A non-interactive zero-knowledge proof of discrete-logarithm knowledge over a (t, s) -secret-shared element y , consists of efficient (possibly randomized) algorithms (SPoSS.Prove, SPoSS.Audit, SPoSS.Verify) with the following functionality. We leave \mathbb{G} and g as implicit inputs.*

- SPoSS.Prove(x) $\rightarrow ([\pi]_1, \dots, [\pi]_s)$. Takes as input integer $x \in \mathbb{Z}_q$. Outputs proof shares $([\pi]_1, \dots, [\pi]_s)$.
- SPoSS.Audit($[y]_i, [\pi]_i$) $\rightarrow \tau_i$. Takes as input a secret shares $[y]_i$ and $[\pi]_i$. Outputs audit token τ_i .
- SPoSS.Verify($\mathcal{T} := \{\tau_i \mid i \in I\}$) $\rightarrow \text{yes/no}$. Takes as input any subset of t or more audit tokens indexed by the set $I \subseteq \{1, \dots, s\}$. Outputs yes if and only if π is a valid proof of discrete logarithm knowledge with respect to $y \in \mathbb{G}$.

The above functionality must satisfy the following properties.

- *Completeness.* For all $x \in \mathbb{Z}_q$ and $y := g^x$, and all subsets $I \subseteq \{1, \dots, s\}$ such that $|I| \geq t$,

$$\Pr \left[\begin{array}{l} ([y]_1, \dots, [y]_s) \leftarrow \text{LSS.Share}_{(\mathbb{Z}_q, t, s)}(y); \\ ([\pi]_1, \dots, [\pi]_s) \leftarrow \text{SPoSS.Prove}(x); \\ \{\tau_i \leftarrow \text{SPoSS.Audit}([y]_i, [\pi]_i) \mid i \in I\}; \\ \text{SPoSS.Verify}(\{\tau_i \mid i \in I\}) = \text{yes} \end{array} \right] = 1,$$

where the probability is over the randomness of SPoSS.Prove .

- *Argument-of-knowledge.* If there exists an efficient (possibly malicious) prover \mathcal{P}^* such that for any group element y , \mathcal{P}^* produces shares of π^* such that $\text{SPoSS.Verify}(\{\tau_i \mid i \in I\}) = \text{yes}$ with probability $\delta(\lambda)$, when each τ_i is computed according to $\text{SPoSS.Audit}([y]_i, [\pi^*]_i)$, then there exists an efficient knowledge extractor \mathcal{E} and negligible function negl such that,

$$\Pr[x \leftarrow \mathcal{E}^{\mathcal{P}^*}(y) \mid y = g^x] \geq \delta(\lambda) - \text{negl}(\lambda),$$

where the probability is over the randomness of \mathcal{P}^* . In words, \mathcal{E} recovers the discrete logarithm x from valid proofs output by \mathcal{P}^* . SPoSS is an argument (rather than a proof) of knowledge because the prover is computationally bounded in the random oracle model.

- *Zero-knowledge.* For all subsets $I \subset \{1, \dots, s\}$ such that $|I| < t$, define \mathcal{D}_I to be the distribution over $\{[\pi]_i, \tau_1, \dots, \tau_s \mid i \in I\}$, where $[\pi]_i$ is sampled according to $\text{SPoSS.Prove}(x)$, and τ_i is computed according to $\text{SPoSS.Audit}([y]_i, [\pi]_i)$ for all $j \in \{1, \dots, s\}$. SPoSS is said to be zero-knowledge if there exists an efficient simulator \mathcal{S} such that $\mathcal{D}_I \approx \mathcal{S}(1^\lambda, I)$. That is, the view induced by the proof shares and audit tokens reveals no information about x or y to any subset of fewer than $t - 1$ computationally bounded (possibly malicious) verifiers.

SPoSS: Main idea. The main idea behind SPoSS is to leverage the additive and multiplicative homomorphism of secret shares over \mathbb{Z}_{p-1} and \mathbb{Z}_p^* , respectively. Our construction assumes $\mathbb{G} = \mathbb{Z}_p^*$. However, our approach can be generalized to any group where the group operation can be described as an arithmetic circuit over a ring (however, the construction becomes more complicated). Notice that, given an additive share $[x]_i$ in \mathbb{Z}_{p-1} , each verifier can obtain a *multiplicative* share of x by computing $g^{[x]_i}$. At a high level, the SPoSS verification procedure goes as follows: Each verifier holds additive secret shares $[y]_i$ and $[x]_i$ (secret shared over \mathbb{Z}_p and \mathbb{Z}_{p-1} , respectively). First, each verifier computes $g^{[x]_i}$ to obtain a multiplicative secret share of x . Notice that $g^{[x]_i}$ is defined over the field \mathbb{F}_p and that the group operation of \mathbb{Z}_p^* is multiplication modulo p . The verifiers then compute the group operation (multiplication in \mathbb{F}_p) over the additive shares using a prover-assisted computation (explained further in [Section 5.1.1](#)). Notice that as a result of this computation, the verifiers hold additive secret shares $[g^x]$. Third, the verifiers compute $[w] := [y] - [g^x]$ and swap their shares of $[w]$ to check if $w = 0$.

Proposition 1. *The SPoSS construction in [Algorithm 3](#) satisfies the correctness, argument-of-knowledge, and zero-knowledge properties ([Definition 6](#)) required of a secret-shared non-interactive proof system, in the random oracle model [7].*

Proof. Deferred to [Appendix B.2](#). □

5.1.1 Protocol overview We describe SPoSS in [Algorithm 3](#). For clarity, we describe the protocol with two verifiers but note that all our techniques extend to a many-verifier setting in the natural way. In [Algorithm 3](#), the verifiers first derive additive shares of $g^{[x]_i}$ (for $i \in \{A, B\}$), which we denote by $[g^{[x]_i}]$. With two verifiers, this is done by simply letting Verifier A set $[g^{[x]_A}]_A := g^{[x]_A}$ and verifier B set $[g^{[x]_A}]_B := 0$ (observe that $[g^{[x]_A}]_A + [g^{[x]_A}]_B = g^{[x]_A}$, as required of additive secret sharing). Verifier B proceeds to do the same with $g^{[x]_B}$. If it were possible to compute the product $[g^{[x]_A} \cdot g^{[x]_B}]$ non-interactively over the additive secret shares, then the verifiers could locally obtain $[g^x]$. Unfortunately, doing so requires *interaction* between the verifiers. Instead, in [Algorithm 3](#), we use a standard approach from zero-knowledge proofs over secret-shares [7, 18] and have the prover

Algorithm 3: SCHNORR PROOF OVER SECRET SHARES

Public parameters: Group $\mathbb{Z}_p^* = (g, p)$ and random oracle $H : \{0, 1\}^* \rightarrow \mathbb{Z}_p \setminus \{0\}$.

– SPoSS.Prove(x):

- 1: $([x]_A, [x]_B) \leftarrow \text{LSS.Share}_{(\mathbb{Z}_{p-1}, 2, 2)}(x)$
- 2: $y_A := g^{[x]_A}, y_B := g^{[x]_B}$ // Multiplicative shares of y .
- 3: $([a]_A, [b]_A, [c]_A, [a]_B, [b]_B, [c]_B) \leftarrow \text{Beaver}_{(2,2)}(\mathbb{F}_p)$
- 4: $a \leftarrow [a]_A + [a]_B, b \leftarrow [b]_A + [b]_B$ // See optimization in Section 5.1.1.
- 5: $z_A, z_B \leftarrow_R \{0, 1\}^\lambda$ // Random nonces.
- 6: $r_A \leftarrow H(z_A, [x]_A, a, [c]_A), r_B \leftarrow H(z_B, [x]_B, b, [c]_B)$
- 7: $r \leftarrow r_A + r_B$
- 8: $d \leftarrow rg^{[x]_A} - a, e \leftarrow g^{[x]_B} - b$
- 9: $[\pi]_A := (A, [x]_A, a, [c]_A, r, d, e, z_A)$
- 10: $[\pi]_B := (B, [x]_B, b, [c]_B, r, d, e, z_B)$
- 11: **return** $([\pi]_A, [\pi]_B)$

– SPoSS.Audit($[y]_i, [\pi]_i$):

- 1: **parse** $[\pi]_i := (i, [x]_i, u, [c]_i, r, d, e, z_i)$
- 2: $\hat{r}_i \leftarrow H(z_i, [x]_i, u, [c]_i)$
- 3: $\hat{y} \leftarrow g^{[x]_i}$
- 4: **if** $i = A$:
 - 4.1: $\hat{f} \leftarrow r\hat{y} - u$
 - 4.2: $[v]_i \leftarrow (de/2) + eu$
 - 4.3: $[w]_i \leftarrow [v]_i + [c]_i - r[y]_i$
- 4: **if** $i = B$:
 - 4.1: $\hat{f} \leftarrow \hat{y} - u$
 - 4.2: $[v]_i \leftarrow (de/2) + du$
 - 4.3: $[w]_i \leftarrow [v]_i + [c]_i - r[y]_i$
- 5: $\tau_i := ([w]_i, \hat{r}_i, r, \hat{f}, d, e)$
- 6: **return** τ_i

– SPoSS.Verify($\{\tau_A, \tau_B\}$):

- 1: **parse** $\tau_A = ([w]_A, \hat{r}_A, r, \hat{d}, d, e)$
- 2: **parse** $\tau_B = ([w]_B, \hat{r}_B, r, \hat{e}, d, e)$
- 3: $\hat{r} \leftarrow \hat{r}_A + \hat{r}_B, w \leftarrow [w]_A + [w]_B$
- 4: **return** $w = 0$ and $\hat{r} = r$ and $\hat{d} = d$ and $\hat{e} = e$

“assist” the verifiers in the computation. Specifically, the prover provides a Beaver multiplication triple [4], enabling the verifiers to compute the multiplication. (We provide an overview of Beaver multiplication in Appendix C for completeness.)

Preventing malicious provers. As observed in prior work [18], prover-assisted multiplication can allow the prover to cheat by introducing a linear term in the output of the multiplication, which would result in the verifiers computing $[\hat{y}] := [g^{[x]_A} \cdot g^{[x]_B} + \Delta]$, for some Δ . To defend against this attack, in Algorithm 3, the verifiers instead check that $[r(g^{[x]_A} \cdot g^{[x]_B})] - [r(y)] = [0]$ where r is a random scalar chosen by the verifiers. As long as the prover does not choose r , the proof is guaranteed to fail for any $\Delta \neq 0$ with probability $1 - \frac{1}{p-1}$, when instantiated over \mathbb{F}_p [18].

Removing interaction. Finally, in Algorithm 3, to avoid interaction between verifiers, we apply the Fiat-Shamir transform [28] and let the *prover* (instead of the verifiers) choose r using a random oracle H . This makes SPoSS mesh with our PACL definition (which only allows for one message exchanged between verifiers). Concretely, we use the distributed analog of Fiat-Shamir described in the full version of Boneh et al. [7]. Given a random oracle H , the prover generates a proof using H to simulate the choice of r by the verifiers. As noted in [7], in the distributed setting, the resulting r can leak information about the shares. To prevent this, in Algorithm 3, we follow the blueprint of [7] and generate random nonces z_A and z_B , that are independent of the proof shares and serve to “mask” the inputs to H .

Reducing proof size. We observe that because each verifier sets all but their own additive share of $\hat{y}_i := g^{[x]_i}$ to zero, the i th verifier knows the value of all other verifiers’ “secret” share of $[\hat{y}_i]$ (they are always zero). As a consequence, only the verifier holding the non-zero share needs to mask it when computing the Beaver multiplication (see Appendix C). This corresponds to revealing a (from the Beaver triple) to Verifier A and b to Verifier B, where the Beaver triple is of the form $([a], [b], [ab])$. Because a and b are random, they still serve as a mask when computing the Beaver multiplication. We apply this optimization in Algorithm 3.

5.2 VDPF-PACL using SPoSS

In this section, we describe how SPoSS can be used to construct a VDPF-PACL. We focus on constructing a VDPF-PACL for the match predicate since extending the construction to inclusion predicates can be achieved by following the blueprint of Algorithm 2. We describe our construction in Algorithm 4. The main idea is that, following private selection of the verification keys (as in Algorithm 1 but now over \mathbb{F}_p) with a VDPF (Definition 8), each verifier holds an (additive) secret-share of $y := g^{\alpha_i}$ (in contrast to Algorithm 1, where the verifiers hold multiplicative secret shares of y). To prove knowledge of α_i , the prover provides a SPoSS proof to the verifiers for the secret-shared group element y . The verifiers then proceed to verify the SPoSS proof and accept if it passes.

5.3 Security analysis

Theorem 5. *Let p be a prime chosen with respect to a security parameter $\lambda \in \mathbb{N}$ and let \mathbb{G} be any group of order p in which the discrete logarithm problem is assumed to be computationally intractable. Algorithm 4 (VDPF-PACL for match predicates) satisfies the completeness, efficiency, soundness, and privacy properties of Definitions 3 and 4, with CheckAccess as defined in Section 2.1 (match predicate).*

Proof. We prove each property in turn.

Completeness. The inner-product computed over the keys results in parties holding secret shares of g^{α_i} (where i is the special index of the encoded point function). It then follows from the completeness of the SPoSS protocol (see Appendix B.2) that SPoSS.Verify must pass following an honest audit procedure. The other half of the PACL.Verify conjunction follows from the completeness of VDPFs (see Appendix A).

Soundness. Assume, towards contradiction, that there exists an efficient adversary \mathcal{A} and non-negligible function δ such that for all $I \subset \{1, \dots, s\}$:

$$\Pr[\text{PKSOUNDNESS}_{\text{PACL}, \mathcal{A}, I, S}(\lambda) = \text{yes}] \geq \delta(\lambda).$$

Algorithm 4: VDPF-PACL FOR MATCH PREDICATES

Public parameters: integers $2 \leq t \leq s$, function family $\mathcal{F} = \{f_i : \{0, 1\}^n \rightarrow \mathbb{Z}_{p-1}\}_{1 \leq i \leq N}$, and group $\mathbb{Z}_p^* = (g, p)$.

– **PACL.KeyGen**($1^\lambda, f_i$): as in [Algorithm 1](#).

– **PACL.Prove**(f, sk):

- 1: **parse** $f = P_i$
- 2: $([\pi]_1, \dots, [\pi]_s) \leftarrow \text{SPoSS.Prove}(\text{sk})$
- 3: **return** $([\pi]_1, \dots, [\pi]_s)$

– **PACL.Audit**($\Lambda, [f]_i, [\pi]_i$):

- 1: **parse** $\Lambda = (g^{\alpha_1}, \dots, g^{\alpha_N})$ and $[f]_i = \kappa_i$
- 2: **for** $j \in \{1, \dots, N\}$:
 - 2.1: $([y_j]_i, \rho_i) \leftarrow \text{VDPF.Eval}(\kappa_i, j)$
 - 2.2: $[b]_i := \sum_{i=1}^N [y]_i$
- 3: $A := \sum_{j=1}^N (g^{\alpha_j}) \cdot [y_j]_i$
- 4: $\tilde{\tau}_i \leftarrow \text{SPoSS.Audit}(A, [\pi]_i)$
- 5: **return** $\tau_i := (\tilde{\tau}_i, \rho_i, [b]_i)$

– **PACL.Verify**(\mathcal{T}):

- 1: **parse** $\mathcal{T} = \{(\tau_1, \rho_1, [b]_1), \dots, (\tau_t, \rho_t, [b]_t)\}$
- 2: **return** $\text{SPoSS.Verify}(\{\tau_1, \dots, \tau_t\})$
and $\text{VDPF.Verify}(\{\rho_i, \dots, \rho_t\})$
and $\text{LSS.Recover}([b]_1, \dots, [b]_t) = 1$

By the soundness property of VDPFs (see [Definition 8](#) in [Appendix A](#)), we can assume that f_γ (output by \mathcal{A} in the $\text{PKSOUNDNESS}_{\text{PACL}, \mathcal{A}, I, S}(\lambda)$ game defined in [Figure 2](#)) is a point function with special index γ . Moreover, because the verifiers check that $b = 1$, the output message encoded in the VDPF must be 1. This restricts \mathcal{A} to outputting a well-formed VDPF, which implies that both verifiers obtain secret shares of $g^{\alpha\gamma}$ from Λ when computing the inner product in PACL.Audit of [Algorithm 4](#). We construct an efficient extractor \mathcal{B} that solves the discrete logarithm problem as follows.

- 1: On input $y := g^x$, generate a planted instance of Λ with y in a random index (as in the proof of [Theorem 3](#)).
- 2: Run \mathcal{A} on inputs $(1^\lambda, y, \Lambda)$ and obtain as output $([f_i], [\pi])$.
- 3: Simulate the SPoSS knowledge extractor on input $[\pi]$ (see [Appendix B.2](#)) and obtain x .
- 4: Output x .

If \mathcal{A} succeeds, then \mathcal{B} recovers the discrete logarithm of g^x via the SPoSS knowledge extractor, contradicting the assumption that the discrete logarithm is computationally intractable in \mathbb{G} . By contrapositive, soundness of [Algorithm 4](#) follows.

Privacy. As with the proof of the DPF-PACL privacy, we construct an efficient simulator \mathcal{S} for the view of any subset of $t - 1$ (possibly malicious) verifiers. We use the efficient simulators $\mathcal{S}_{\text{VDPF}}$ and $\mathcal{S}_{\text{SPoSS}}$ to generate the view of the VDPF ([Definition 8](#) in [Appendix A](#)) output and SPoSS proof ([Definition 6](#)), respectively. On input $(1^\lambda, I)$, \mathcal{S} proceeds as follows:

- 1: $\{[\pi]_i, \tau_1, \dots, \tau_s \mid i \in I\} \leftarrow \mathcal{S}_{\text{SPoSS}}(1^\lambda, I)$.
- 2: $([b]_1, \dots, [b]_s) \leftarrow \text{LSS.Share}_{(\mathbb{Z}_p, t, s)}(1)$
- 3: $\{\rho_1, \dots, \rho_s\} \leftarrow \mathcal{S}_{\text{VDPF}}(1^\lambda)$.
- 4: Output $\{[\pi]_i, (\tau_1, \rho_1, [b]_1), \dots, (\tau_s, \rho_s, [b]_s) \mid i \in I\}$.

The distribution output by \mathcal{S} matches the distribution of any subset of $t - 1$ verifiers because: (1) the proof share $[\pi]_i$ and audit tokens τ_1, \dots, τ_s are output by the SPoSS simulator $\mathcal{S}_{\text{SPoSS}}$, which guarantees statistical indistinguishability of the view, (2) the VDPF simulator $\mathcal{S}_{\text{VDPF}}$ guarantees computational indistinguishability of the VDPF verification tokens ρ_1, \dots, ρ_s , and (3) the $[b]_i$'s are computational secret shares of 1 in the real view and information-theoretic secret shares of 1 in the simulated view.

Efficiency. Each proof share $[\pi]_i$ and audit token τ_i consists of a constant number of elements in \mathbb{F}_p and thus is of size $O(\lambda)$ (see [Definition 6](#) and [Appendix B.2](#)). By definition of VDPFs, the size of ρ_i satisfies our efficiency constraint. (We note that under the VDPF formulation presented by de Castro and Polychroniadou [23], the efficiency of the verification procedure is not explicitly defined but can be seen as an implicit requirement.) \square

Remark 7 (Embedding the message in \mathbb{Z}_p). If we wish to interpret a VDPF output in \mathbb{Z}_2 as an output in \mathbb{Z}_p , we end up with a message -1 or 1 (in \mathbb{Z}_p). In this case, when applying the VDPF-PACL scheme ([Algorithm 4](#)), the verification of the message (by checking that $b = 1$ in PACL.Verify) would leak information and should *not* be performed. However, this check is not required because the output message is in \mathbb{Z}_2 , which is sufficient to guaranteed that the converted output has magnitude $1 \in \mathbb{Z}_p$.

6 PACLs for FSS from DPF-PACLs

We now describe a set of PACL constructions for classes of FSS derived from DPFs (and verifiable DPFs). These transformations are taken from Boyle et al. [9, Section 3.2] and form the class of functions that can be efficiently secret-shared using lightweight cryptographic assumptions. More expressive classes of FSS are believed to require heavier tools [9], for instance, fully-homomorphic encryption [25].

PACLs for range functions and decision trees. Boyle et al. [9, 11] describe how to apply linear combinations of DPFs to derive FSS for range functions (and more generally decision trees [11]). Range functions and decision trees can be viewed as special cases of distributed multi-point functions (DMPF), which evaluate to a non-zero value on multiple inputs. In turn, DMPFs can be viewed as an *aggregation* of DPFs (this also follows from the linear-composition of FSS [9, Section 3.2]). By the aggregation property of DPF-PACLs ([Section 4.3](#)), we immediately obtain PACLs for DMPFs and, as a result, PACLs for range functions and decision trees.

PACLs for small function classes. FSS for all functions with a small domain $|\mathcal{F}|$ can be obtained via a DPF that “selects” the function $f_i \in \mathcal{F}$ in the canonically ordered function family \mathcal{F} [9]. Our DPF-PACL construction applies to this class of FSS directly as a result. Following similar transformations, Boyle et al. [9] obtain FSS for data matching and NC^0 functions, which we briefly describe next.

PACLs for data matching functions. Data-matching functions are parameterized by a set $S \subseteq \{1, \dots, N\}$ of $\ell \in O(1)$ elements and a value $v \in \{0, 1\}^n$ such that $f_{S,v}(x) = 1$ if $x_i = v_i, \forall i \in S$. FSS for this class of functions can be realized using a DPF with a range large enough to describe all $\binom{n}{\ell} 2^\ell$ possible values of $f_{S,v}$ (hence the requirement that ℓ is constant). As a consequence, our DPF-PACL can be applied directly to this family of FSS by associating each access key with the corresponding canonically ordered function.

PACLs for NC^0 functions. The class NC^0 captures all functions that can be represented by constant-depth boolean circuits $C : \{0, 1\}^u \rightarrow \{0, 1\}^v$ with fan-in 2 (two inputs per gate). We can trivially consider a DPF-PACL where Λ corresponds to all possible such circuits, of which there are $v^{O(u^{2^d})}$ in total. However, this is a naive approach. As observed by Boyle et al. [9], it is possible to leverage the bit-wise *parallel* structure of NC^0 circuits and DPFs to realize efficient FSS for NC^0 functions. Specifically, any circuit $C \in \text{NC}^0$ can be decomposed into v 1-bit-output, depth- d sub-circuits. For u -bit inputs, each such sub-circuit has only $O(u^{2^d})$ possibilities. For each sub-circuit, we can generate a DPF for the i th canonical ordering of all $O(u^{2^d})$ possible circuits. Repeating this for all v sub-circuits yields an FSS scheme consisting of v DPF keys (one for each sub-circuit). Using the aggregation property of our DPF-PACL construction described in Section 4.3, it is possible to enforce access control over the v DPFs simultaneously. However, it becomes necessary to enforce access control over the unique *combination* of sub-circuits since each DPF-PACL operates independently of the *global* circuit C . To achieve this, we can apply a “generic PACL” (Section 7) over the combination of sub-circuits (in conjunction with DPF-PACLs for each sub-circuits) using a zero-knowledge proof over secret-shared data.

7 Generic PACLs from distributed zero-knowledge proofs

In this section, we describe how to construct PACLs for any FSS class (formally, FSS for all functions in P/poly [9]). Our approach relies on secret-shared non-interactive proofs (SNIPs) [18] and Fiat-Shamir over SNIPs [7]. We describe these preliminaries in Section 7.1.

We keep the descriptions in this section to a high level, since the ideas and construction follow generically from zero-knowledge proofs on secret-shared data. Our construction generalizes the approach of Vadapalli et al. [53], which uses zero-knowledge proofs on secret-shared data to realize access control. In our construction (described in Section 7), we provide an optimization for shrinking the proof size with the help of a vector commitment scheme.

7.1 Preliminaries

SNIPs [18] (and their generalizations [7]) can be used to prove that any (public) arithmetic circuit C evaluates to 1 on a secret-shared input x provided that the following two conditions are met: (1) the circuit C is known to the verifiers and (2) the prover knows the input x and C . SNIPs guarantee that the verifiers (who hold secret shares of x) do not learn any information except that $C(x) = 1$. The efficiency of SNIPs is measured by the size of a SNIP proof and the interaction between verifiers (note that SNIPs are non-interactive for the prover). The size of a SNIP is proportional to the number of multiplication gates in the circuit and can be verified in one round of interaction. We provide a formal definition of SNIPs in Definition 7. We frame the definition to follow the syntax of Definition 3 by abstracting the verifier interaction using algorithms `SNIP.Audit` and `SNIP.Verify`.

Definition 7 (Secret-shared Non-interactive Proof (SNIP)). *Let $\lambda \in \mathbb{N}$ be a security parameter, \mathbb{F} be a finite field, and t and s be integers such that $2 \leq t \leq s$. Let C be any arithmetic circuit (defined over \mathbb{F}) where for an $x \in \mathbb{F}$, it holds that $C(x) = 1$. A SNIP is a zero-knowledge proof system instantiated between a prover and s verifiers holding t -out-of- s secret shares of x , where the prover convinces the verifiers that $C(x) = 1$, in zero knowledge. In the random oracle model, a SNIP proof system consists of three (possibly randomized) algorithms (`SNIP.Prove`, `SNIP.Audit`, `SNIP.Verify`):*

- $\text{SNIP.Prove}(x, C) \rightarrow ([\pi]_1, \dots, [\pi]_s)$. Takes as input the input x and the arithmetic circuit C . Outputs a vector of t -out-of- s proof secret shares.
- $\text{SNIP.Audit}([x]_i, [\pi]_i) \rightarrow \tau_i$. Takes as input a t -out-of- s secret share of x and the corresponding proof share. Outputs a verification string τ_i .
- $\text{SNIP.Verify}(\mathcal{T} := \{\tau_i \mid i \in I\}) \rightarrow \text{yes/no}$. Takes as input any subset of t or more verification strings indexed by the set $I \subseteq \{1, \dots, s\}$. Outputs **yes** if and only if it holds that $C(x) = 1$.

A SNIP must satisfy the completeness, soundness, zero-knowledge, and efficiency properties of a multi-verifier zero knowledge proof system [7, 55].

- *Completeness.* For all arithmetic circuits C and for all x such that $C(x) = 1$, then for all subsets $I \subseteq \{1, \dots, s\}$ such that $|I| \geq t$,

$$\Pr \left[\begin{array}{l} ([x]_1, \dots, [x]_s) \leftarrow \text{SNIP.LSS.Share}_{(t,s)}(x); \\ ([\pi]_1, \dots, [\pi]_s) \leftarrow \text{SNIP.Prove}(x, C); \\ \{\tau_i \leftarrow \text{SNIP.Audit}([x]_i, [\pi]_i) \mid i \in I\} \end{array} \middle| \text{SNIP.Verify}(\{\tau_i \mid i \in I\}) = \text{yes} \right] = 1,$$

where the probability is taken over the randomness of LSS.Share and SNIP.Prove .

- *Soundness.* For all efficient adversaries \mathcal{A} , subsets $I \subseteq \{1, \dots, s\}$ where $|I| \geq t$, and x such that $C(x) \neq 1$, there exists a negligible function negl such that,

$$\Pr \left[\begin{array}{l} ([x]_1, \dots, [x]_s, [\pi^*]_1, \dots, [\pi^*]_s) \leftarrow \mathcal{A}(1^\lambda, C), \\ \{\tau_i \leftarrow \text{SNIP.Audit}([x]_i, [\pi^*]_i) \mid i \in I\} \end{array} \middle| \begin{array}{l} \text{LSS.Recover}(\{[x]_i \mid i \in I\}) = x \\ \wedge \\ \text{SNIP.Verify}(\{\tau_i \mid i \in I\}) = 1 \end{array} \right] \leq \text{negl}(\lambda),$$

where the probability is taken over the randomness of \mathcal{A}

- *Zero-knowledge.* For all subsets $I \subseteq \{1, \dots, s\}$ such that $|I| < t$, define \mathcal{D}_I to be the distribution over $\{[\pi]_i, \tau_1, \dots, \tau_s \mid i \in I\}$, where each π_i is sampled according to Prove , and each $\text{SNIP}.\tau_j$ is sampled according to SNIP.Audit . A SNIP is zero-knowledge if there exists an efficient simulator \mathcal{S} such that $\mathcal{D}_I \approx \mathcal{S}(1^\lambda, I)$. That is, no subset of fewer than t (possibly malicious) verifiers gain any information about x (in the information theoretic sense).
- *Efficiency.* The size of each proof share $[\pi]_i$ is bounded by the number of multiplication gates in C , and the size of τ_i is constant (in the security parameter).

Fiat-Shamir for SNIPs. The Fiat-Shamir transform [28] is a standard technique used to eliminate interaction in zero-knowledge proofs. At a high level, Fiat-Shamir allows the prover to generate its own challenges with the help of a random oracle by “simulating” the randomness chosen by the verifiers in the interactive proof system. With SNIPs, however, the situation is slightly different because SNIPs are already non-interactive for the prover. Instead, Fiat-Shamir can be applied to SNIPs to reduce the interaction required when *verifying* the proof [7, 55]. Specifically, the verification of the SNIP with Fiat-Shamir requires only one message exchanged between verifiers (instead of one round of interaction consisting of two sequential messages).

As observed in [7], the main challenge in the distributed setting is that the shares of each verifier must be kept secret from the other verifiers. If Fiat-Shamir were to be applied directly over the proof share given to each verifier, then it could leak information about the proof shares of other verifiers (and possibly leak information about x). The high-level idea to get around this problem is to have the prover *randomize* each share with a random nonce, which it distributes to the verifiers along with the proof shares. Using the nonce, the verifiers are able to verify consistency of the challenge used to check the proof without learning any information on the secret shares held by the other verifiers. (See [7, 6.2.3] for details.)

How Fiat-Shamir is applied to SNIPs. When verifying a SNIP proof, the verifiers first jointly sample a random value r and proceed to perform a randomized polynomial identity test using the DeMillo-Lipton-Schwartz-Zippel lemma [24, 47]. The Fiat-Shamir transform can be used to let the prover sample the “challenge” randomness r on behalf of the verifiers using a random oracle. The verifiers then simply check the consistency of the proof (i.e., verify the polynomial identity test), which only requires exchanging one message.

Vector commitment schemes. A vector commitment scheme $VC = (VC.Commit, VC.Open, VC.Verify)$ is a commitment scheme [31] to a vector with an efficient local opening to any particular coordinate of the vector [15]. Concretely, we will use a Merkle tree [39] to instantiate VC , where $Merkle.Commit$ computes a hash tree over the vector coordinates, and $Merkle.Open$ outputs the sibling hashes from the leaf to the root, and $Merkle.Verify$ checks that the path from the coordinate to the root is correct using the sibling paths. For simplicity, we do not provide a formal definition or construction of a vector commitment schemes here, and instead point to Catalano and Fiore [15] for formal definitions. Informally, the two important properties required of any commitment scheme are *hiding* and *binding*, which guarantee that (1) the commitment reveals no information on the committed vector, and (2) a commitment cannot be opened to a different coordinate than the one committed to.

7.2 Construction

At a high level, we construct PACLs for P/poly-FSS as follows. As in our group-based constructions, the dealer first distributes shares of the function f using the FSS scheme and shares of the access key sk to all verifiers. In addition, the prover distributes a SNIP proof showing that (1) $[f]$ corresponds to *some* valid output of $FSS.Gen$ and (2) the access key sk is such that $CheckAccess(\Lambda, f, sk) = 1$. (Without loss of generality, we assume that both $FSS.Gen$ and $CheckAccess$ are described as arithmetic circuits over a field \mathbb{F} .) The verifiers then check the SNIP using their shares of f and sk .

While this approach to generic PACLs is conceptually simple, there is an efficiency challenge that needs to be addressed. By definition, $CheckAccess$ takes the entire access control list Λ , which would result in a large (linear in $|\Lambda|$) proof size, violating the efficiency property of PACLs. Specifically, the naive approach requires the prover to incorporate the entire access control list into the SNIP when proving that $CheckAccess(\Lambda, f, sk) = 1$, even if sk only depends on *one* verification key in Λ . We overcome this efficiency problem by using any vector commitment scheme [15, 39] (described by algorithms $VC.Commit$ and $VC.Verify$). We instead define $vk_i := H(f_i || sk_i)$, where H is a collision-resistant hash function (CRHF) sampled by the verifiers and sk_i is the access key for function f_i . The verifiers compute a commitment to all verification keys using $VC.Commit$, publishing the resulting commitment c and all the openings e_1, \dots, e_N such that $VC.Verify(c, vk_i, e_i) = 1$. See Figure 4 for an example where we instantiate the vector commitment using a Merkle tree [39]. The prover then sends $[vk]$, $[sk]$, $[\tilde{e}]$, and three SNIP proofs: (1) a proof that $FSS.Gen(1^\lambda, f; r) = [f]$, where r describes the random coins of $FSS.Gen$, (2) a proof that $vk = H(f, sk)$, and (3) a proof that $Verify(c, vk, \tilde{e}) = 1$. The verifiers check the three SNIP proofs using their shares $[f]$ and $[sk]$. If the proofs are accepting, then the verifiers are convinced that the dealer knows the access key sk associated with f and that their shares of f were output according to $FSS.Gen$. The size of the SNIP proofs is bounded by $O(s + \log |\Lambda|)$ rather than $O(|\Lambda|)$, where s is the number of multiplication gates in $FSS.Gen$.

Proposition 2. *Algorithm 5 satisfies the completeness, soundness, privacy, and efficiency guarantees of Definition 4.*

Proof (sketch). We give a high-level proof sketch for each property in turn.

Completeness. Completeness follows from the correctness of vector commitment scheme and completeness of SNIPs (Definition 7). If the FSS shares of f are generated correctly (output by $FSS.Gen$), then π_0 is accepted by the completeness of SNIPs. Similarly, for π_1 and π_2 , where completeness is easy to see. As a result, $PACL.Verify$ outputs *yes*.

Soundness. Soundness follows from the binding property of the vector commitment scheme and the soundness guarantee of SNIPs. Specifically, π_0 proves that $[f]$ is the output of $FSS.Gen$, π_1 proves that the prover knows a pre-image of the form $f || sk$ for the verification key vk , and π_2 proves that vk is in the key list Λ (i.e., a valid verification key). Note that the same f (which the verifiers hold a secret share of) is used by the verifiers when checking the validity of π_0 , π_1 , and π_2 , which binds the proofs together.

Privacy. Privacy follows directly from the privacy guarantee of SNIPs since the only information revealed is (1) the proof shares $[\pi_0]$, $[\pi_1]$, and $[\pi_2]$ and (2) the SNIP verification strings.

Efficiency. The size of $[\pi_0]$ is proportional to the number of multiplication gates in $FSS.Gen$. The size of $[\pi_1]$ is proportional to the number of multiplication gates in $VC.Verify$ (bounded by $O(\log |\Lambda|)$ when instantiated using a Merkle tree [39]). The audit token consists of the SNIP verification strings, which are of size $O(\lambda)$. \square

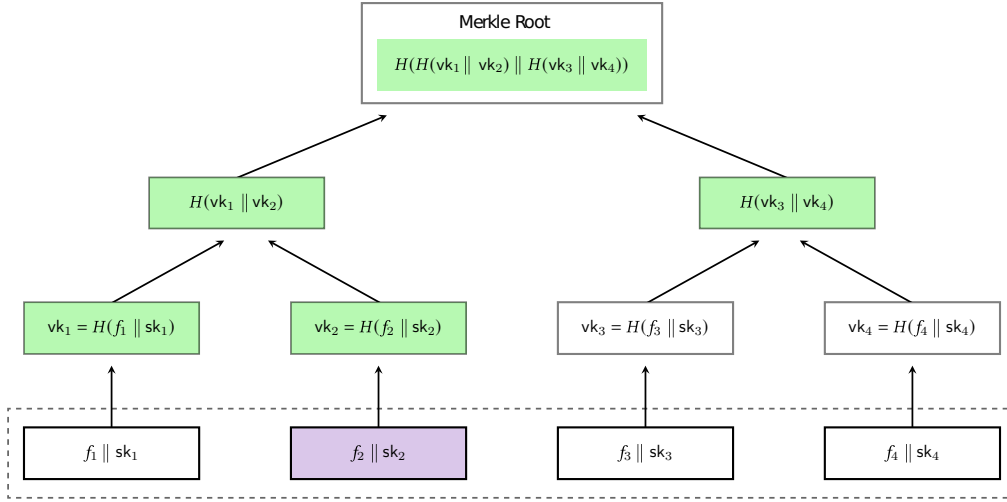


Fig. 4: Example of the Merkle tree description when proving $\text{CheckAccess}(\mathcal{A}, f_2, \text{sk}_2)$. The SNIP consists of secret-shared inputs to Merkle.Verify which takes as input the path from the root to the leaf (nodes highlighted in green), where the leaf node is the access key for the function (highlighted in purple). Nodes at the leaf layer consist of access keys and are not known to the verifiers. Verification keys are of the form $\text{vk}_i := H(f_i || \text{sk}_i)$ such that \mathcal{A} consists of all nodes in the second-to-last layer.

Algorithm 5: GENERIC PACL WITH SNIPS

Public parameters: CRHF H , vector commitment \mathbf{c} to the list of verification keys \mathcal{A} , and commitment openings e_1, \dots, e_N to each verification key in \mathcal{A} .

- $\text{PACL.KeyGen}(1^\lambda, f_i)$
 - 1: Sample $\text{sk}_i \leftarrow_R \{0, 1\}^\lambda$.
 - 2: Compute $\text{vk}_i := H(f_i || \text{sk}_i)$.
 - 3: Output $(\text{sk}_i, \text{vk}_i)$.
- $\text{PACL.Prove}(f, \text{sk})$:
 - 1: Compute $\tilde{\text{vk}} := H(f || \text{sk})$ and $\tilde{e} := e_j$ where $j \in \{1, \dots, N\}$ such that $f = f_j$.
 - 2: Compute SNIP proofs for the following statements:
 - 2.1: $[\pi_0]$ for the statement “ $\text{FSS.Gen}(1^\lambda, f; r) = [f]$ ”, where $[f]$ are the FSS shares of f .
 - 2.2: $[\pi_1]$ for the statement “ $\tilde{\text{vk}} = H(f || \text{sk})$ ”.
 - 2.3: $[\pi_2]$ for the statement “ $\text{VC.Verify}(\mathbf{c}, \tilde{\text{vk}}, \tilde{e}) = 1$ ”.
 - 3: Output proof shares $[\pi] := ([\pi_0], [\pi_1], [\pi_2], [\tilde{\text{vk}}], [\tilde{e}])$.
- $\text{PACL.Audit}(\mathcal{A}, [f]_i, [\pi]_i)$:
 - 1: Parse $[\pi]_i := ([\pi_0]_i, [\pi_1]_i, [\pi_2]_i, [\tilde{\text{vk}}]_i, [\tilde{e}]_i)$.
 - 2: Compute the verification strings $\tau^{(0)}$, $\tau^{(1)}$, and $\tau^{(2)}$ for $[\pi_0]_i$, $[\pi_1]_i$, and $[\pi_2]_i$, respectively.
 - 3: Output audit token $\tau_i := (\tau_i^{(0)}, \tau_i^{(1)}, \tau_i^{(2)})$.
- $\text{PACL.Verify}(\mathcal{T})$:
 - 1: Parse each $\tau_i \in \mathcal{T}$ as $\tau_i = (\tau_i^{(0)}, \tau_i^{(1)}, \tau_i^{(2)})$.
 - 2: $\mathcal{T}^{(j)} := \{\tau_i^{(j)} \mid i \in \{1, \dots, |\mathcal{T}|\}\}$, for all $j \in \{0, 1, 2\}$.
 - 3: Output yes if and only if all three sets $\mathcal{T}^{(0)}$, $\mathcal{T}^{(1)}$, $\mathcal{T}^{(2)}$ of SNIP verification strings are accepting under SNIP.Verify .

Remark 8 (Efficiency). The proof size for our generic FSS PACL construction is proportional to the size of the `Gen` circuit. However, FSS is only concerned with the *share size* and not the running time of `Gen`. As such, it is possible that `Gen` runs in time proportional to the truth table of the function being secret shared (e.g., see *programmable* DPFs [13]). In this case, the size of the SNIP proof—which has a size proportional to the circuit size of `Gen`—can negate the efficiency of the FSS scheme. An interesting direction for future work would be to make the proof size only depend on the size of the function secret shares and be independent of `Gen`. However, we note that in all known constructions of FSS for P/poly [9, 25], `Gen` runs in sublinear time relative to the truth table.

8 Implementation and evaluation

In this section, we describe our implementation and evaluation of the (V)DPF-PACL constructions from Sections 4 and 5. Our evaluation focuses on the state-of-the-art two-party FSS schemes [11, 23]. Multi-party FSS constructions are less efficient [9] or require heavier cryptographic assumptions, making them concretely slower [19, 41]. Because we are interested in evaluating the *overhead* of PACLs, evaluating our constructions in a two-party setting results in *worst-case overheads* relative to baseline FSS evaluations. We evaluate our implementation for applications of FSS, including private information retrieval, distributed anonymous authentication, and anonymous communication protocols.

Implementation. We implement PACLs in Go v1.16. and C in approximately 4,500 lines of code. Our implementation is open-source [1]. We instantiate \mathbb{G} as the P-256 elliptic curve group (part of the `crypto/elliptic` package in Go) in our DPF-PACL construction for public-key soundness (Definition 4) and as \mathbb{Z}_{2^λ} for symmetric-key soundness (Definition 5). For our public-key VDPF-PACL construction, we instantiate \mathbb{G} as \mathbb{F}_p^* with a 3072-bit prime p as specified in RFC3526.

Environment. We use Amazon Elastic Cloud Compute (EC2) for our experiments. We run experiments on a `c4.4xlarge` (16 vCPUs; 32 GB RAM) Amazon Linux general-purpose virtual machine. We use AES-NI-enabled CPUs for fast PRG evaluations, as well as other (V)DPF-specific optimizations [11, 23]. **All experiments are performed on a single core.**

Methodology. We run each experiment between 10 and 1,000 times (depending on the experiment) and report the average over the trials. 95% confidence interval is occasionally invisible.

Optimizations. In a setting with two verifiers, we observe that `Verify` simply has each verifier checking that each secret share corresponds to a share of zero. If the parties convert their shares to *subtractive* (rather than additive) secret shares, then this check becomes an equality check (both parties have the *same* subtractive share if it’s a share of zero). Therefore, to avoid sending all secret shares, the verifiers can send succinct hashes of their audit shares to reduce communication.

8.1 Prover costs

The proving costs for our PACL constructions are minimal. The prover only needs to generate (V)DPF keys (for inclusion predicates) in our (V)DPF-PACL constructions. The complexity of (V)DPF.Gen is linear in the (V)DPF domain parameter n [9, 11] and remains below 20 ms for practical values of n (i.e., $n \leq 128$). In our VDPF-PACL construction, the prover also has to compute the SPoSS proof. We benchmark SPoSS.Prove (with $s = 2$) at 30 ms of CPU time (bottlenecked by the two exponentiation in \mathbb{Z}_p^*). Proving costs grow linearly with s .

8.2 Communication costs

Tables 1 and 2 report the concrete communication overheads on the prover and the verifiers. To remain independent of the underlying DPF construction, we let s_ℓ denote the size of a DPF with range ℓ (e.g., $s_\ell = \log(\ell) \cdot (\lambda + 2)$ bits [11]).

In Table 1, we compare communication costs to other approaches for DPF access control. Express [27] and Sabre [53] operate in the symmetric-key setting, satisfying the relaxed PACL soundness definition (Section 3.2). Newman et al. [41] construct an access control mechanism similar to Algorithm 1 satisfying the soundness definition of Definition 4. Using SPoSS for verifying discrete-logarithm knowledge results in over $2,400\times$ less communication compared to a naive approach (described in Appendix D) and a $1,000\times$ smaller proof size compared to Sabre [53].

Table 1: Proof size (Prover \rightarrow Verifier) and audit token size (Verifier \leftrightarrow Verifier) for (V)DPF-PACL with the match predicate for access control (Algorithm 1). *Estimated (see Appendix D).

Match predicate	Prover \rightarrow Verifier	Verifier \leftrightarrow Verifier
DPF-PACL	32 B	64 B
VDPF-PACL (SPoSS)	1952 B	816 B
(V)DPF-PACL (symmetric-key)	16 B	16 B
VDPF-PACL (naive SPoSS*)	4.7 MB	816 B
Express [27] (symmetric-key)	2 kB	184 B
Spectrum [41]	32 B	64 B
Sabre [53] (symmetric-key)	(40 + 120 <i>n</i>) kB	16 B

Table 2: Proof size (Prover \rightarrow Verifier) and audit token size (Verifier \leftrightarrow Verifier) for (V)DPF-PACL with inclusion predicate for access control (Algorithm 2) and ℓ access keys per function in the FSS family. We denote by s_ℓ the size (in B) of a (V)DPF key with a range of $\{1, \dots, \ell\}$. Prior techniques [27, 41, 53] do not support inclusion predicates. *Estimated (see Appendix D).

(inclusion predicate)	Prover \rightarrow Verifier	Verifier \leftrightarrow Verifier
DPF-PACL	(16 + s_ℓ) B	32 B
VDPF-PACL (SPoSS)	(1952 + s_ℓ) B	816 B
(V)DPF-PACL (symmetric-key)	(16 + s_ℓ) B	16 B
VDPF-PACL (naive SPoSS*)	(4.7 $\times 10^6$ + s_ℓ) B	816 B

Table 3: Overhead of introducing public-key (Section 3) and symmetric key (Section 3.2) PACLs to DPF and DMPF classes of FSS. As the FSS class becomes more complex (e.g., FSS for DMPFs such as inequality and range functions), the overhead of enforcing access control diminishes. We set the (V)DPF domain to $\{0, 1\}^{32}$. All benchmarks are amortized over 100,000 evaluations of the FSS.

	FSS.Eval Baseline	DPF-PACL (match predicate)		DPF-PACL (inclusion predicate w/ $\ell = 20$)	
		symmetric-key	public-key	symmetric-key	public-key
DPF	1.41 μ s	1.42 μ s	5.66 μ s	33.10 μ s	84.11 μ s
DMPF	90.44 μ s	90.46 μ s	93.26 μ s	122.22 μ s	213.47 μ s
	VFSS.Eval Baseline	VDPF-PACL (match predicate)		VDPF-PACL (inclusion predicate w/ $\ell = 20$)	
		symmetric-key	public-key	symmetric-key	public-key
DPF	1.46 μ s	1.47 μ s	1.69 μ s	33.66 μ s	36.11 μ s
DMPF	93.14 μ s	93.19 μ s	93.50 μ s	125.52 μ s	128.71 μ s

8.3 Verification costs

We report the processing time in Figures 5 and 6. Introducing PACLs results in a concrete processing overhead relative to evaluating f_i (here, f_i is either a DPF or DMPF), especially when the number of evaluations of the function is small (e.g., less than 64). However, as the number of evaluations increases, the *amortized* cost of access control decreases (the overhead of the group exponentiation in *Verify* is amortized over the evaluations of f_i). FSS itself is typically only of interest in settings where the function is evaluated on a large number of inputs (otherwise, it is more efficient to just secret share $[f(x)]$ rather than $[f]$). As such, it is more reasonable to consider the amortized overhead that access control introduces. For our DPF-PACL construction (reported in Figure 5), the amortized overhead plateaus at approximately $5\times$ the baseline cost of evaluating f_i with around 2^8 evaluations. This is primarily due to the linear number of group (elliptic curve) exponentiations required in the *Audit* procedure. In contrast, for our VDPF-PACL construction (reported in Figure 6), which requires only one group exponentiation in \mathbb{F}_p^* , we observe a larger initial overhead but far better amortized overhead. The larger initial overhead is entirely due to the single exponentiation in \mathbb{F}_p^* (which we benchmark at approximately 13 ms). All our constructions have a lower overhead as the complexity of the FSS increases (e.g., when applying PACLs to DMPFs) thanks to the aggregation properties described in Section 4.3. To better understand the asymptotic amortization of our constructions, we report the tail values of Figures 5 and 6 in Table 3, where we amortize over 100,000 evaluations of the DPF and DMPF.




Fig. 5: DPF-PACL with match and inclusion predicate for access control, where ℓ is the number of access keys per function in the FSS family ($\ell = 1$ corresponds to match predicate). Our DPF-PACL construction amortizes with more evaluations of the DPF but plateaus at approximately 32 evaluations. For DMPFs, which have a higher baseline processing time, the overhead is less significant thanks to aggregation, which amortizes the overhead of access control (Section 4.3).

8.4 Applications of PACLs

Private databases with access control. Systems that use multi-server PIR (e.g., [19, 21, 22, 33, 37, 54]) can take advantage of (V)DPF-PACLs to restrict access to database items [34]. (Gupta et al. [33] specifically leave open the problem of supporting *authenticated* media consumption through Popcorn.) Other systems such as Dory [21] use DPFs for private keyword queries in a remote database. For example, Wang et al. [54] use PIR (realized using DPFs) to build privacy-preserving restaurant, geolocation, and flight searches. Gupta et al. [33] use PIR for privacy-preserving movie streaming. In Figure 7, we report the overhead of introducing access control in PIR via VDPF-PACLs. As the items in the database become larger, the overhead of introducing access control diminishes. Part of the overhead from introducing PACLs to PIR is due to switching from operations in a binary field (xors) to operations in \mathbb{F}_p , which are concretely slower. Ostrovsky and Shoup [44] describe a read-and-write private database, which can be realized using DPFs [11]. Applying VDPF-PACLs to this setting would result in similar overheads to the PIR setting.

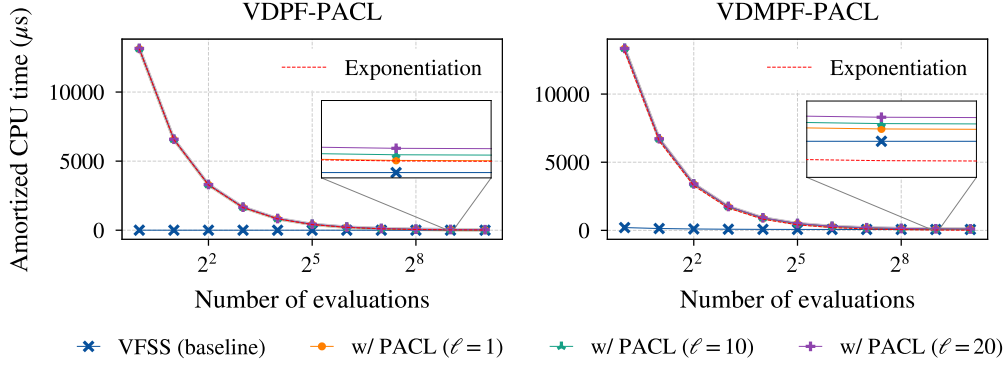


Fig. 6: VDPF-PACL with match and inclusion predicate for access control, where ℓ is the number of access keys per function in the FSS family ($\ell = 1$ corresponds to match predicate). Our VDPF-PACL construction is dominated by the exponentiation in \mathbb{F}_p^* (dashed lined) and has a higher *initial* overhead compared to the DPF-PACL construction. In contrast to DPF-PACLs, our VDPF-PACL construction amortizes almost entirely after 64 evaluations and has an asymptotically smaller overhead.




Fig. 7: Server-side CPU time for privately retrieving an item from a database through two-server PIR with and without VDPF-PACLs. Adding access control has a 1.5–3 \times impact on processing time. The access control overhead is automatically amortized over the number of items in the database. The relative overhead also diminishes as the items in the database become larger.

Anonymous authentication. We identify a potentially interesting application of PACLs for the purpose of anonymous user authentication (for instance, password-based authentication [32]) in a distributed setting. Denote each user pseudonym by i and let sk_i be the corresponding key (or password). We can set Λ to be the set of valid account keys and let f_i be the point function with special index i . Applying our VDPF-PACL construction over $[f_i]$, the verifiers can learn if user i has a valid account (and knows the key sk_i associated with vk_i) *without* learning which account was used to authenticate. In Table 4, we benchmark the processing time required to authenticate a user as a function of the number of accounts in the database.

Table 4: Evaluation of VDPF-PACLs applied to anonymous user authentication with varying number of accounts (evaluation points).

Anonymous authentication with VDPF-PACLs				
Number of accounts:	250K	500K	1M	2M
Authentication time:	103 ms	192 ms	381 ms	757 ms

Faster anonymous communication. Our VDPF-PACL construction can be applied out-of-the-box to the anonymous communication systems Express [27], Sabre [53], and Spectrum [41] to improve their concrete performance. In Figure 8, we show that swapping their implicit access control mechanism with our VDPF-PACL construction improves performance by a factor of 50–70×. Sabre [53]’s computational overhead is on-par with baseline FSS (and DPF-PACLs satisfying symmetric-key soundness Section 3.2 and Table 1) but requires significantly larger proofs for access control purposes (approximately 3.5 MB; see Table 1). However, we note that Sabre [53] achieves other nice properties that are tailored to anonymous communication (see Section 1.2).

Remark 9. In the anonymous authentication setting, we can use a VDPF that outputs an arbitrarily large message in the field extension \mathbb{Z}_{2^k} but only need to perform access control over the first bit of the output message in \mathbb{Z}_2 (which we set to 1, WLOG), letting the remaining $k - 1$ bits be arbitrary. See also Remark 7 pertaining to embedding \mathbb{Z}_2 outputs in \mathbb{Z}_p with VDPF-PACLs.

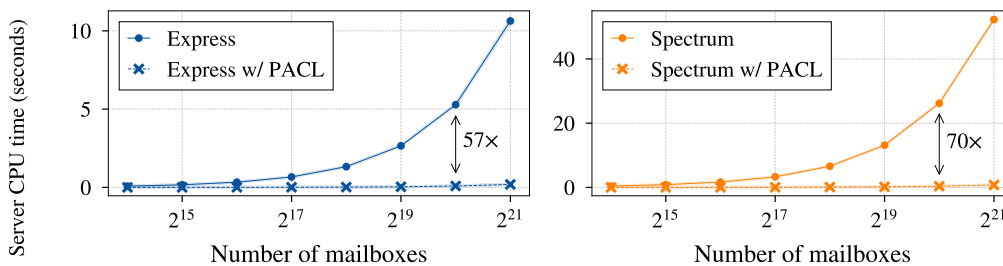


Fig. 8: Our VDPF-PACLs significantly improve performance of anonymous communication systems that require access control. These performance improvements come from two sources: (1) the ability to use optimized DPF constructions with our DPF-PACLs (e.g., the access control in Express requires concretely slower DPF constructions) and (2) the better amortization of VDPF-PACLs.

9 Conclusion

We modeled and formalized the notion of private access control for FSS. Our constructions can be applied to a variety of FSS applications and improve the performance of ad-hoc methods found in prior work. We also present a generic theoretical construction that has exciting potential for future work. Finally, we evaluate our constructions and showcase their performance on several concrete use cases, ranging from anonymous authentication and communication to access control in private databases. Our evaluation shows that introducing access control results in minimal overheads relative to baseline FSS and amortizes well asymptotically when the function is evaluated on many inputs.

Acknowledgements

We would like to thank Sridhar Devadas for his generous support and for reading several drafts of this paper. We thank Anish Athalye, Leo de Castro, Pawan Goyal, Kyle Hogan, Simon Langowski, Manon Revel, and Mayuri Sridhar for helpful discussions, comments, and suggestions. We thank Alexander Koch and Radhika Garg for pointing out a few typos and inconsistencies.

We would also like to thank the USENIX 2022 reviewers for providing us with exceptionally thorough feedback on an earlier version of this work and the S&P 2023 reviewers and our shepherd for their incredibly detailed comments that helped us to significantly improve the presentation of the paper. Finally, we would like to thank the organizers of the MIT PRIMES program for making it a such a fun and enriching experience.

Bibliography

- [1] PACL source code. <https://github.com/sachaservan/pacl>, 2022.
- [2] Abdelrahman Aly, Aysajan Abidin, and Svetla Nikova. Practically efficient secure distributed exponentiation without bit-decomposition. In *International Conference on Financial Cryptography and Data Security*. Springer, 2018.
- [3] Ghous Amjad, Seny Kamara, and Tarik Moataz. Breach-resistant structured encryption. *Proceedings on Privacy Enhancing Technologies*, 1, 2019.
- [4] Donald Beaver. Efficient multiparty protocols using circuit randomization. In *Annual International Cryptology Conference*. Springer, 1991.
- [5] D Elliott Bell and Leonard J La Padula. Secure computer system: Unified exposition and multics interpretation. Technical report, MITRE CORP BEDFORD MA, 1976.
- [6] Dan Boneh and Victor Shoup. A graduate course in applied cryptography. *Draft 0.5*, 2020.
- [7] Dan Boneh, Elette Boyle, Henry Corrigan-Gibbs, Niv Gilboa, and Yuval Ishai. Zero-knowledge proofs on secret-shared data via fully linear PCPs. In *Annual International Cryptology Conference*. Springer, 2019.
- [8] Dan Boneh, Elette Boyle, Henry Corrigan-Gibbs, Niv Gilboa, and Yuval Ishai. Lightweight techniques for private heavy hitters. In *2021 IEEE Symposium on Security and Privacy (SP)*. IEEE, 2021.
- [9] Elette Boyle, Niv Gilboa, and Yuval Ishai. Function secret sharing. In *Annual international conference on the theory and applications of cryptographic techniques*. Springer, 2015.
- [10] Elette Boyle, Niv Gilboa, and Yuval Ishai. Breaking the circuit size barrier for secure computation under DDH. In *Annual International Cryptology Conference*. Springer, 2016.
- [11] Elette Boyle, Niv Gilboa, and Yuval Ishai. Function secret sharing: Improvements and extensions. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, 2016.
- [12] Elette Boyle, Nishanth Chandran, Niv Gilboa, Divya Gupta, Yuval Ishai, Nishant Kumar, and Mayank Rathee. Function secret sharing for mixed-mode and fixed-point secure computation. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*. Springer, 2021.
- [13] Elette Boyle, Niv Gilboa, Yuval Ishai, and Victor I Kolobov. Programmable distributed point functions. In *Annual International Cryptology Conference*. Springer, 2022.
- [14] David FC Brewer and Michael J Nash. The chinese wall security policy. In *IEEE symposium on security and privacy*, volume 1989. Oakland, 1989.
- [15] Dario Catalano and Dario Fiore. Vector commitments and their applications. In *International Workshop on Public Key Cryptography*. Springer, 2013.
- [16] Benny Chor, Oded Goldreich, Eyal Kushilevitz, and Madhu Sudan. Private information retrieval. In *Proceedings of IEEE 36th Annual Foundations of Computer Science*. IEEE, 1995.
- [17] David D Clark and David R Wilson. A comparison of commercial and military computer security policies. In *1987 IEEE Symposium on Security and Privacy*. IEEE, 1987.
- [18] Henry Corrigan-Gibbs and Dan Boneh. Prio: Private, robust, and scalable computation of aggregate statistics. In *14th USENIX Symposium on Networked Systems Design and Implementation (NSDI 17)*, 2017.

- [19] Henry Corrigan-Gibbs, Dan Boneh, and David Mazières. Riposte: An anonymous messaging system handling millions of users. In *2015 IEEE Symposium on Security and Privacy*. IEEE, 2015.
- [20] Ivan Damgård, Matthias Fitz, Eike Kiltz, Jesper Buus Nielsen, and Tomas Toft. Unconditionally secure constant-rounds multi-party computation for equality, comparison, bits and exponentiation. In *Theory of Cryptography Conference*. Springer, 2006.
- [21] Emma Dauterman, Eric Feng, Ellen Luo, Raluca Ada Popa, and Ion Stoica. DORY: An encrypted search system with distributed trust. In *14th USENIX Symposium on Operating Systems Design and Implementation (OSDI 20)*, 2020.
- [22] Emma Dauterman, Mayank Rathee, Raluca Ada Popa, and Ion Stoica. Waldo: A private time-series database from function secret sharing. In *2022 IEEE Symposium on Security and Privacy (SP)*. IEEE, 2022.
- [23] Leo de Castro and Anitgoni Polychroniadou. Lightweight, maliciously secure verifiable function secret sharing. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*. Springer, 2022.
- [24] Richard A DeMillo and Richard J Lipton. A probabilistic remark on algebraic program testing. Technical report, Georgia Inst Tech Atlanta School of Information and Computer Science, 1977.
- [25] Yevgeniy Dodis, Shai Halevi, Ron D Rothblum, and Daniel Wichs. Spooky encryption and its applications. In *Annual International Cryptology Conference*. Springer, 2016.
- [26] Jack Doerner and abhi shelat. Scaling ORAM for secure computation. In *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security*, 2017.
- [27] Saba Eskandarian, Henry Corrigan-Gibbs, Matei Zaharia, and Dan Boneh. Express: Lowering the cost of metadata-hiding communication with cryptographic privacy. In *30th USENIX Security Symposium (USENIX Security 21)*, 2021.
- [28] Amos Fiat and Adi Shamir. How to prove yourself: Practical solutions to identification and signature problems. In *Conference on the theory and application of cryptographic techniques*. Springer, 1986.
- [29] Craig Gentry, Philip MacKenzie, and Zulfikar Ramzan. A method for making password-based key exchange resilient to server compromise. In *Annual International Cryptology Conference*. Springer, 2006.
- [30] Niv Gilboa and Yuval Ishai. Distributed point functions and their applications. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*. Springer, 2014.
- [31] Oded Goldreich. *Foundations of Cryptography, Volume 1*. Cambridge university press, 2000.
- [32] Matthew Green. Let’s talk about PAKE, october 2018. <https://blog.cryptographyengineering.com/2018/10/19/lets-talk-about-pake/>, 2018. Accessed December 2022.
- [33] Trinabh Gupta, Natacha Crooks, Whitney Mulhern, Srinath Setty, Lorenzo Alvisi, and Michael Walfish. Scalable and private media consumption with popcorn. In *13th USENIX Symposium on Networked Systems Design and Implementation (NSDI 16)*, 2016.
- [34] Ryan Henry, Femi Olumofin, and Ian Goldberg. Practical PIR for electronic commerce. In *Proceedings of the 18th ACM conference on Computer and communications security*, 2011.
- [35] Yizhou Huang and Ian Goldberg. Outsourced private information retrieval. In *Proceedings of the 12th ACM Workshop on Workshop on Privacy in the Electronic Society*, 2013.
- [36] Keyu Ji, Bingsheng Zhang, and Kui Ren. Fine-grained policy constraints for distributed point function. *Cryptology ePrint Archive*, 2023.

- [37] Ari Juels. Targeted advertising... and privacy too. In *Cryptographers' Track at the RSA Conference*. Springer, 2001.
- [38] Jonathan Katz and Yehuda Lindell. *Introduction to modern cryptography*. CRC press, 2020.
- [39] Ralph C Merkle. A digital signature based on a conventional encryption function. In *Conference on the theory and application of cryptographic techniques*. Springer, 1987.
- [40] Pratyush Mishra, Ryan Lehmkuhl, Akshayaram Srinivasan, Wenting Zheng, and Raluca Ada Popa. Delphi: A cryptographic inference service for neural networks. In *29th USENIX Security Symposium (USENIX Security 20)*, 2020.
- [41] Zachary Newman, Sacha Servan-Schreiber, and Srinivas Devadas. Spectrum: High-bandwidth anonymous broadcast. In *19th USENIX Symposium on Networked Systems Design and Implementation (NSDI 22)*. USENIX Association, April 2022.
- [42] Chao Ning and Qiuliang Xu. Constant-rounds, linear multi-party computation for exponentiation and modulo reduction with perfect security. In *International Conference on the Theory and Application of Cryptology and Information Security*. Springer, 2011.
- [43] Takashi Nishide and Kazuo Ohta. Multiparty computation for interval, equality, and comparison without bit-decomposition protocol. In *International Workshop on Public Key Cryptography*. Springer, 2007.
- [44] Rafail Ostrovsky and Victor Shoup. Private information storage. In *Proceedings of the twenty-ninth annual ACM symposium on Theory of computing*, 1997.
- [45] Théo Ryffel, Pierre Tholoniati, David Pointcheval, and Francis R. Bach. Ariann: Low-interaction privacy-preserving deep learning via function secret sharing. *Proc. Priv. Enhancing Technol.*, 2022(1):291–316, 2022.
- [46] Claus-Peter Schnorr. Efficient signature generation by smart cards. *Journal of cryptology*, 4(3), 1991.
- [47] Jacob T Schwartz. Fast probabilistic algorithms for verification of polynomial identities. *Journal of the ACM (JACM)*, 27(4), 1980.
- [48] Sacha Servan-Schreiber, Simon Beyzerov, Eli Yablon, and Hyojae Park. Private access control for function secret sharing. In *2023 IEEE Symposium on Security and Privacy (SP)*. IEEE Computer Society, 2022.
- [49] Adi Shamir. How to share a secret. *Communications of the ACM*, 22(11), 1979.
- [50] Joseph H Silverman. *The arithmetic of elliptic curves*, volume 106. Springer, 2009.
- [51] John T Tate. The arithmetic of elliptic curves. *Inventiones mathematicae*, 23(3), 1974.
- [52] Tomas Toft. Constant-rounds, almost-linear bit-decomposition of secret shared values. In *Cryptographers' Track at the RSA Conference*. Springer, 2009.
- [53] A. Vadapalli, K. Storrier, and R. Henry. Sabre: Sender-anonymous messaging with fast audits. In *2022 IEEE Symposium on Security and Privacy (SP) (SP)*, Los Alamitos, CA, USA, may 2022. IEEE Computer Society.
- [54] Frank Wang, Catherine Yun, Shafi Goldwasser, Vinod Vaikuntanathan, and Matei Zaharia. Splinter: Practical private queries on public data. In *14th USENIX Symposium on Networked Systems Design and Implementation (NSDI 17)*, 2017.
- [55] Kang Yang and Xiao Wang. Non-interactive zero-knowledge proofs to multiple verifiers. *Cryptology ePrint Archive*, 2022.

A Verifiable DPFs

de Castro and Polychroniadou [23] present definitions for (2, 2)-VDPF constructions. We generalize their definitions to any (t, s) -VDPF scheme.

Definition 8 (VDPF [23]). Let $\lambda \in \mathbb{N}$ be a security parameter and \mathbb{F} be any finite field. Fix a domain $\{0, 1\}^n$. A VDPF consists of three (possibly randomized) algorithms (VDPF.Gen, VDPF.Eval, VDPF.Verify):

- VDPF.Gen($1^\lambda, i \in \{0, 1\}^n, m \in \mathbb{F}$) $\rightarrow (\kappa_1, \dots, \kappa_s)$. Takes as input a security parameter, an index $i \in \{0, 1\}^n$, and message m . Outputs a set of evaluation keys encoding point function $P_{i,m}$ such that $P_{i,m}(i) = m$.
- VDPF.Eval($\kappa_j, X \subseteq \{0, 1\}^n$) $\rightarrow ([\mathbf{v}]_j, \rho_j)$. Takes as input an evaluation key κ_j and a subset of values in the domain. Outputs a secret share of a vector \mathbf{v} , where the k -th coordinate of \mathbf{v} corresponds to a share of $P_{i,m}(k)$ for $k \in X$, and a verification string ρ_j .
- VDPF.Verify($\{\rho_j \mid j \in J\}$) $\rightarrow \text{yes/no}$. Takes as input any subset of t or more verification strings indexed by the set $J \subseteq \{1, \dots, s\}$. Outputs yes if and only if $\{[\mathbf{v}]_j \mid j \in J\}$, as output by Eval using $\{\kappa_j \mid j \in J\}$, encodes a point function on the evaluated set of inputs X .

A (t, s) -VDPF must satisfy the correctness, privacy, and efficiency properties of FSS (Definition 2). Additionally, a VDPF must guarantee soundness. Informally, soundness requires that Verify outputs no for any set of evaluation keys that do not encode a point function on the evaluated points.

- *Correctness.* A (t, s) -VDPF is correct if there exists an efficient algorithm Decode such that for all $k \in \{0, 1\}^n$, for all $m \in \mathbb{F}$, and all subsets $J \subseteq \{1, \dots, s\}$, such that $|J| \geq t$,

$$\Pr \left[(\kappa_1, \dots, \kappa_s) \leftarrow \text{VDPF.Gen}(1^\lambda, i, m) : \text{Decode}(\{\text{VDPF.Eval}(\kappa_j, \{k\}) \mid j \in J\}) = P_{i,m}(k) \right] = 1,$$

where the probability is over VDPF.Gen.

- *Soundness.* A (t, s) -VDPF is sound if for all (possibly maliciously generated) keys $(\kappa_1^*, \dots, \kappa_s^*)$, adversarially chosen inputs $X^* \subseteq \{0, 1\}^n$, and ρ_j sampled according to $(\cdot, \rho_j) \leftarrow \text{VDPF.Eval}(\kappa_j^*, X^*)$ for $j \in J$, $|J| \geq t$, it holds that if the correctness property is not satisfied then:

$$\Pr[\text{VDPF.Verify}(\{\rho_j \mid j \in J\}) = \text{yes}] \leq \text{negl}(\lambda),$$

where the probability is taken over the adversary's randomness. Otherwise, VDPF.Verify($\{\rho_j \mid j \in J\}$) always outputs yes.

- *Privacy.* For all subsets $I \subset \{1, \dots, s\}$ such that $|I| < t$, define \mathcal{D}_I to be the distribution over $\{\kappa_i, \rho_1, \dots, \rho_s \mid i \in I\}$, where κ_i is sampled according to VDPF.Gen, each ρ_i^* is sampled arbitrarily, and each ρ_i is sampled according to Eval. A (t, s) -VDPF is private if there exists an efficient simulator \mathcal{S} such that $\mathcal{D}_I \approx_c \mathcal{S}(1^\lambda, I)$. That is, all subsets of fewer than s evaluation keys and the entire set of verification strings (of which $t - 1$ might be maliciously generated), reveal no information on the point function encoded in the set of keys $(\kappa_1, \dots, \kappa_s)$.

B Deferred proofs

B.1 Proof of Theorem 4 (security of Algorithm 2)

Completeness. Consider $C \in \mathbb{G}$ as computed in PAEL.Verify:

$$C := \left(\prod_{j=1}^{\ell} A_j \right) \cdot \left(\prod_{j=1}^{N\ell} g^{-c_j \cdot w_j} \right) \cdot g^\alpha.$$

Examining “the exponent,” we get that:

$$\log_g(C) = \sum_{j=1}^N \sum_{k=1}^{\ell} \alpha_{j,k} \cdot y_j - \left(\sum_{j=1}^{N\ell} w_j \cdot c_j \right) + \alpha.$$

If f_i is a DPF instance, then $y_j = 1$ only for $j = i$. Thus,

$$\log_g(C) = \sum_{k=1}^{\ell} \alpha_{i,k} - \left(\sum_{j=1}^{N\ell} w_j \cdot c_j \right) + \alpha.$$

Further, if $(\kappa'_1, \dots, \kappa'_s)$ encode a DPF for the $(i-1)\ell + \gamma = \omega$ -th index, then all $c_j = 0$ for $j \neq \omega$. Therefore,

$$\log_g(C) = \sum_{k=1}^{\ell} \alpha_{i,k} - w_{\omega} + \alpha.$$

However, by construction, $w_{\omega} = \sum_{k=1, k \neq \gamma}^{\ell} \alpha_{i,k}$, and so we get that: $\log_g(C) = \alpha_{i,\gamma} + \alpha = 0$, by construction since $\alpha := -\alpha_{i,\gamma}$. Therefore, it holds that $C = g^0 = 1_{\mathbb{G}}$ and `PACL.Verify` outputs yes, making the construction complete.

Soundness (proof of knowledge). Assume, towards contradiction, that there exists an efficient \mathcal{A} that wins the $\text{PKSOUNDNESS}_{\text{PACL}, \mathcal{A}, I, S}(\lambda)$ game with non-negligible probability $\delta(\lambda)$. Then, \mathcal{A} outputs secret shares of \hat{f}_{γ} (corresponding to point function P_{γ}) encoded as keys $(\hat{\kappa}_1, \dots, \hat{\kappa}_s)$ and proof shares $([\hat{\pi}]_1, \dots, [\hat{\pi}]_s)$ such that for all $I \subseteq \{1, \dots, s\}$ where $|I| \geq t$,

$$\Pr \left[\begin{array}{l} \tau_i \leftarrow \text{PACL.Audit}(\Lambda, [\hat{f}_{\gamma}]_i, [\hat{\pi}]_i), \forall i \in I : \\ \text{PACL.Verify}(\{\tau_i \mid i \in I\}) = \text{yes} \end{array} \right] \geq \delta(\lambda).$$

Without loss of generality (by [Lemma 1](#)), we can assume that \mathcal{A} outputs \hat{f}_{γ} sampled from the family of point functions when considering [Algorithm 2](#). For notational simplicity, we “expand” the N verification keys in Λ (each containing ℓ subkeys) so as to make Λ consist of $N\ell$ verification keys. We now construct an efficient algorithm \mathcal{B} that solves the discrete logarithm problem as follows. On input $y := g^x$,

- 1: $(\alpha_{1,1}, \dots, \alpha_{N,(\ell-1)}) \leftarrow_R \mathbb{F}_p^{N \times (\ell-1)}$.
- 2: $\Lambda := (g^{\alpha_{1,1}}, \dots, g^{\alpha_{1,\ell}}, \dots, g^{\alpha_{N,1}}, \dots, g^{\alpha_{N,(\ell-1)}}), y$.
- 3: Run $\mathcal{A}^{\text{GetKey}}(1^{\lambda}, \Lambda)$ and answer each $\text{GetKey}(i)$ query with $(\alpha_{i,1}, \dots, \alpha_{i,\ell})$ for all $i \neq N$. If \mathcal{A} queries GetKey on input N , then **abort**.
- 4: Obtain output $([\hat{f}_{\gamma}], [\hat{\pi}])$ from \mathcal{A} .
- 5: Recover \hat{f}_{γ} and $(\hat{f}_{\omega}, \hat{\alpha}, \hat{\beta})$ from $[\hat{f}_{\gamma}]$ and $[\hat{\pi}]$, respectively.
- 6: Output $\hat{\alpha}$.

By the aggregation properties of our DPF-PACL constructions and [Lemma 1](#) and soundness of [Algorithm 1](#) ([Section 4.4](#)), it holds that $\hat{\alpha}$ is the discrete logarithm of y whenever \mathcal{A} succeeds in $\text{PKSOUNDNESS}_{\text{PACL}, \mathcal{A}, I, S}(\lambda)$ with a DPF f_N and using the ℓ -th key in Λ for f_N . The probability that \mathcal{A} outputs f_N is $\frac{1}{N}$ and the probability \mathcal{A} uses the ℓ -th key for row N is $\frac{1}{\ell}$. Thus, \mathcal{B} succeeds with probability at least $\frac{1}{N\ell} \delta(\lambda)$ which remains non-negligible, contradicting the hardness of the discrete logarithm problem. Again, we note that \mathcal{B} extracts the access key which makes our construction a proof of knowledge.

Privacy. We construct an efficient simulator \mathcal{S} for the view of any subset of at most $t-1$ (possibly malicious) verifiers. Let \mathcal{S}' be the simulator in the proof of [Theorem 3](#). On input $(1^{\lambda}, I)$, \mathcal{S} proceeds as follows:

- 1: $([0]_1, \dots, [0]_s) \leftarrow \text{LSS.Share}_{(\mathbb{F}_p, t, s)}(0)$.
- 2: $([z']_1, \dots, [z']_s) \leftarrow \text{LSS.Share}_{(\mathbb{F}_p, t, s)}(0)$.
- 3: $(\kappa'_1, \dots, \kappa'_s) \leftarrow \text{DPF.Gen}(1^{\lambda}, P_1)$ for $P_1 \in \mathcal{F}_{\text{DPF}}$.
- 4: $\{[\pi']_i, \tau_1^{(0)}, \dots, \tau_s^{(0)} \mid i \in I\} \leftarrow \mathcal{S}'(1^{\lambda}, I)$.
- 5: $[\pi]_i := ([z']_i, [\pi']_i, \kappa'_i)$.
- 6: $\tau_k^{(1)} := g^{[0]_k}$ for all $k \in \{1, \dots, s\}$.
- 7: $\tau_k := (\tau_k^{(0)}, \tau_k^{(1)})$ for all $k \in \{1, \dots, s\}$.
- 8: Output $\{[\pi]_i, \tau_1, \dots, \tau_s \mid i \in I\}$.

The distribution output by \mathcal{S} matches the distribution of any subset $I \subset \{1, \dots, s\}$, where $|I| < t$ because: (1) any subset of $\{[0]_1, \dots, [0]_s\}$ of size $< t$ is uniformly distributed and therefore matches the distribution of any subset of $\{[\alpha]_1, \dots, [\alpha]_s\}$ of size $< t$ in the real view, (2) $[\pi']_i$ is guaranteed to

be computationally indistinguishable by the proof of [Theorem 3](#), (3) the DPF key for point function P_ω (in the real view) is computationally indistinguishable to κ'_i corresponding to point function P_1 by the privacy of FSS ([Definition 2](#)), and (4) the audit tokens are (computationally-hiding) multiplicative secret shares of $(1_{\mathbb{G}}, 1_{\mathbb{G}})$ in the real view and uniformly random multiplicative secret shares in the output of \mathcal{S} . An efficient distinguisher for (4) would also contradict the privacy property of FSS.

Efficiency. Using the ‘‘FSS tensoring’’ [[11](#)] optimization (see [Section 4.1.2](#)), the size of each proof share is $O(\lambda + s_\ell)$ where s_ℓ is the size of a DPF key encoding a point function with range $\{1, \dots, \ell\}$.

B.2 Security of SPoSS

Proof of [Proposition 1](#). We prove each property in turn: *completeness*, *argument-of-knowledge*, and *zero-knowledge*.

Completeness. We show that if the prover is honest, then `PACL.Verify` outputs yes. In [Algorithm 3](#), `PACL.Verify` outputs yes if and only if $w_A + w_B = 0$ and $\hat{r} = r$, $\hat{d} = d$, and $\hat{e} = e$. The equality of $\hat{r} = r$, $\hat{d} = d$, and $\hat{e} = e$ follows by inspection. To see why it holds that $w_A + w_B = 0$, observe that

$$\begin{aligned} w_A + w_B &= v_A - ry_A + v_B - ry_B \\ &= v - ry \\ &= 2\left(\frac{de}{2}\right) + ea + db + (c_A + c_B) - ry \\ &= \underbrace{(r\hat{y}_A - a)}_d \underbrace{(\hat{y}_B - b)}_e + ea + db + c - ry \\ &= r\hat{y} - a(\hat{y}_B) + a(\hat{y}_B) - ry \\ &= r(\hat{y} - y) = 0 \text{ by assumption that } \hat{y} = g^x = y. \end{aligned}$$

Argument-of-knowledge. We construct an efficient extractor \mathcal{E} that recovers the discrete logarithm of y from a proof $[\pi]$ output by a possible malicious prover \mathcal{P}^* . \mathcal{E} proceeds as follows:

- 1: Run $\mathcal{P}^*(y)$ to obtain as output (π_A, π_B) where

$$\begin{aligned} \pi_A &= (A, [x]_A, a, [c]_A, r, d, e, z_A) \\ \pi_B &= (B, [x]_B, b, [c]_B, r, d, e, z_B). \end{aligned}$$

- 2: Output $x = [x]_A + [x]_B$.

If (π_A, π_B) is a valid SPoSS proof valid, then $w_A + w_B = 0$. In turn, we have that $r\hat{y} - (r\hat{y}_A)b - a(\hat{y}_B) + ab + ea + db - ry + c = 0$ for some randomness r (see completeness proof). The malicious prover \mathcal{P}^* can choose arbitrary a, b, c . As such, we have that $c = ab + \Delta$ for some Δ [[18](#)], which yields $r\hat{y} - (r\hat{y}_A)b - a(\hat{y}_B) + 2ab + ea + db - ry + \Delta$ which reduces to $r(\hat{y} - y) + \Delta = 0$. Thus, either (1) the malicious prover obtained r from H such that $r(\hat{y} - y) = -\Delta$ (which happens with negligible probability given H is a random oracle) or (2) $\hat{y} = y$ and $\Delta = 0$ which implies that $(\hat{y} - y) = 0$ and therefore $g^{[x]_A + [x]_B} = \hat{y} = y$ and so $x = [x]_A + [x]_B$ is the discrete logarithm, as required.

Zero-knowledge. To prove that SPoSS is zero-knowledge, we construct an efficient simulator \mathcal{S} that given $i \in \{A, B\}$ and τ_i^* , outputs a statistically indistinguishable view to that of verifier i (the simulator generalizes to the many-verifier case). On input $(1^\lambda, \{i\})$, \mathcal{S} proceeds as follows:

- 1: $([w]_A, [w]_B) \leftarrow \text{LSS.Share}_{(\mathbb{F}_p, 2, 2)}(0)$.
- 2: $[x]_i \leftarrow_R \mathbb{Z}_{p-1}$.
- 3: $[c]_i, u, d, e \leftarrow_R \mathbb{F}_p$.
- 4: $\hat{r}_i, r \leftarrow_R \mathbb{Z}_p \setminus \{0\}$.
- 5: $z_i \leftarrow_R \{0, 1\}^\lambda$.
- 6: $[\pi]_i := ([x]_i, u, [c]_i, r, d, e, z_i)$.
- 7: $\tau_A \leftarrow ([w]_A, \hat{r}_i, r, d, d, e)$, $\tau_B \leftarrow ([w]_B, r - \hat{r}_i, r, e, d, e)$.
- 8: Output $\{[\pi]_i, \tau_A, \tau_B\}$.

Analysis. Consider the distribution of $[\pi]_i$ in the real view of the i -th verifier. Observe that $[\pi]_i$ consists of (1) secret shares $[x]_i$ and $[c]_i$, (2) masks z_i and u (see optimization described in [Section 5.1.1](#)), (3) Beaver multiplication openings d, e , and (4) the distributed Fiat-Shamir randomness r . All these values are generated by the prover. (1), (2) and (3) are uniformly distributed. (4) is uniformly distributed due to the mask z_i [7]. As such, the distribution of $[\pi]_i$, as output by \mathcal{S} , matches that of the real view of the i -th verifier.

Now, consider τ_j , which consists of (1) a secret share $[w]_i$, (2) random oracle outputs \hat{r} and r (which are identical when the prover is honest and distributed uniformly due to the mask z_j) and (3) Beaver multiplication openings f (computed by verifier j and d, e given by the prover). (1) and (3) are uniformly distributed due to $[c]_j$ being uniform share and the mask u_j . Moreover, f is either d or e and thus provides no new information. (2) reveals no new information because $\hat{r} = r$ and r is given to all verifiers.

We conclude that the output distribution of \mathcal{S} is distributed identically to the view of verifier i in [Algorithm 3](#), which concludes the proof of zero-knowledge.

C Beaver’s Protocol

We provide a brief overview of Beaver’s [4] approach to multiplication of secret shares as adapted by Corrigan-Gibbs and Boneh [18]. We focus on the setting with two parties; see [18] for a more general exposition. Given two parties holding additive shares $[x]$ and $[y]$, encoding field elements x and y , the parties must securely compute shares of $[xy]$, encoding the value $xy \in \mathbb{F}$. A Beaver triple consists of additive shares of $([a], [b], [c])$ such that a and b are random field elements and $c := ab \in \mathbb{F}$. If the parties are given shares of a Beaver triple, then the parties can compute a secret share encoding the product two secret shares x and y as follows. Each party locally computes:

$$[d] \leftarrow [x] - [a] \quad \text{and} \quad [e] \leftarrow [y] - [b],$$

and broadcasts its shares of d and e . The parties recover d and e and locally compute: $[xy] := d[b] + e[a] + [c] + \frac{de}{2}$. This works because

$$\begin{aligned} d[b] &= (x - a)[b] = [xb - ab], \\ e[a] &= (y - b)[a] = [ya - ab], \\ \frac{de}{2} &= \frac{xy - xb - ay + ab}{2} \end{aligned}$$

($\frac{de}{2}$ is a 2-out-of-2 “share” of $de = xy - xb - ay + ab$) and so we get that:

$$\begin{aligned} d[b] + e[a] + \frac{de}{2} + [c] &= \\ &= [xb - ab] + [ya - ab] + [xy - xb - ay + ab] + [c] \\ &= [xb - ab + ya - ab + xy - xb - ay + ab + c] \\ &= [xy - ab + c] \\ &= [xy]. \end{aligned}$$

As such, Beaver’s technique reduces the rounds of communication required to compute a multiplication over secret shares down to one round [4].

D Naive SPoSS using SNIPs

Here, we estimate the overhead of naively applying a SNIP for verifying a Schnorr proof over secret shares. Verification in Schnorr requires computing an exponentiation in \mathbb{G} . This translates to an exponentiation with a secret-shared exponent in our case. Using the textbook approach to modular exponentiation (e.g., repeated squaring) or even tailored techniques (e.g., [2, 20, 42, 43, 52]), requires computing a bit decomposition to convert the secret-shared exponent into binary secret-shares. A bit decomposition circuit requires $O(\lambda \log \lambda)$ multiplication gates, where $\lambda := \lceil \log p \rceil$ [20, 43, 52]. Using

the explicit multiplication complexity listed by Toft [52, Table 2] for their bit decomposition protocol, we get that the total number of multiplications to compute the circuit is $M = 31 \cdot \lambda \log(\lambda) + 71 \cdot \lambda + 30 \cdot \sqrt{\lambda}$. When $\mathbb{G} = \mathbb{Z}_p^*$ and $\log p \approx 3072$, then we have:

$$M = 32 \cdot 3072 \cdot \log(3072) + 71 \cdot 3072 + 30 \cdot \sqrt{3072} \approx 1.4 \times 10^6.$$

The SNIP proof requires sending one Beaver triple (3λ bits per verifier) and two elements of \mathbb{F}_p per multiplication gate in the circuit (the proof consists of a degree- $2M$ polynomial interpolating the multiplication gates). This results in $(3 \cdot 3072) + 2 \cdot (3072) \cdot M$ bits = $(9216) + (6144) \cdot (1.4 \times 10^6)$ bits ≈ 1 GB of communication per verifier.

However, we can be more clever and have the prover secret-share the *bit decomposed* exponent as part of the proof. In this case, the verifiers only need to check that the secret shares encode a binary number and then apply the group operation (multiplication in \mathbb{F}_p) λ times. We now describe the arithmetic circuit computing an exponentiation with a bit-decomposed exponent and repeated squaring. To the best of our knowledge, this is the most optimal generic approach to verifying a Schnorr proof over secret shares.

First, the verifiers check that each secret-shared bit a_i for $i \in \{1, \dots, \lambda\}$ is either 0 or 1. The arithmetic circuit computing this check is defined as $C(a_i) = 1 + a_i^2 - a_i$. Therefore, the arithmetic circuit for checking the validity of the binary decomposition requires λ multiplication gates and makes the SNIP proof consist of 2λ elements of \mathbb{F}_p . The total size, in bits, is therefore $2\lambda^2$.

Second, the verifiers much check the repeated squaring circuit, which requires computing the group operation (one multiplication in \mathbb{F}_p) λ times. This makes the SNIP proof consist of 2λ elements of \mathbb{F}_p . The total size of the repeated squaring proof, in bits, is therefore $2\lambda^2$ as well.

Combined, the total proof size is:

$$\begin{aligned} & \underbrace{(3 \cdot 3072)}_{\text{Beaver triple}} + \underbrace{2 \cdot (3072)^2}_{\text{binary check}} + \underbrace{2 \cdot (3072)^2}_{\text{repeated squaring}} \text{ bits} \\ & = (9216) + 2 \cdot (3072)^2 + 2 \cdot (3072)^2 \text{ bits} \\ & \approx 4.7 \text{ MB.} \end{aligned}$$

Note that using an elliptic curve instead of \mathbb{Z}_p^* (which would allow us to work over a field of roughly order $p \approx 2^{256}$ instead of $p \approx 2^{3072}$) does not improve the situation. While the proof size for the validity of the binary decomposition (checking that $C(a_i) = 1$, for all i) would be roughly $12\times$ smaller, these savings are negated by the complexity of the elliptic curve group operation, which requires multiple field multiplications to compute the group operation [50, 51]. The advantage of working with \mathbb{Z}_p^* is that we only require *one* multiplication in \mathbb{F}_p to apply the group operation.